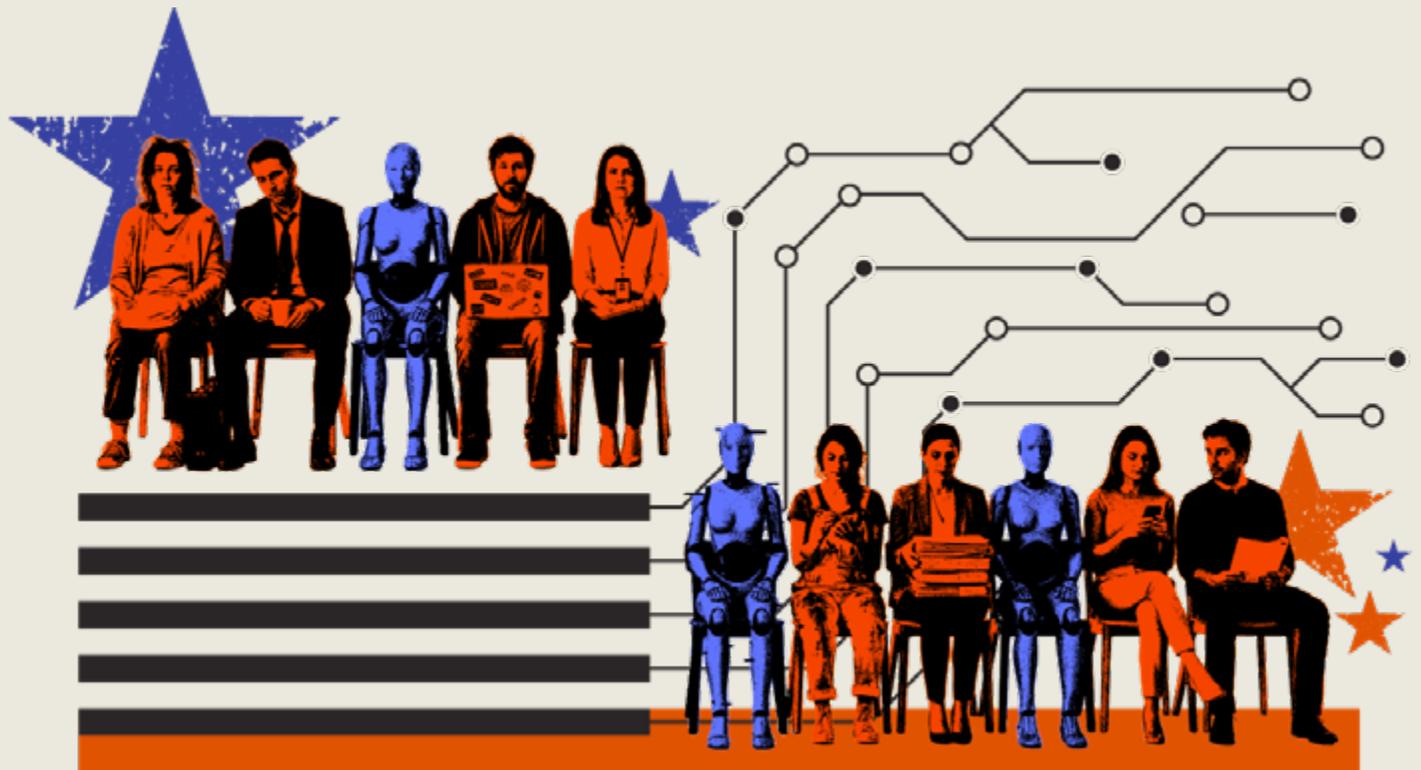


Will Wired Belts Become the New Rust Belts?

AI and the Emerging Geography of American Job Risk

Methodology



AI's Effects on the Workforce

AI Exposure Score Methodology

Occupational Level Exposure Scores

We gather occupational employment data at the metropolitan area, nonmetropolitan area, state, and territory levels using O*NET's database (specifically the May 2024 Occupational Employment and Wage Statistics data).

1. We gather various research that measure how exposed occupations are to AI in the workplace. These first scores rate the exposure of occupation to artificial intelligence rather than replacement, automatability, or vulnerability, focusing on how well AI can perform the tasks and activities of that occupation. We reach our score by aggregating the data from three other research papers, each viewing AI through a difference lens and providing a rating for most occupations:
 - **Task Based Score, or TB Score (Eloundou et. al):** This paper measures the exposure of occupation to AI by whether LLMs can reduce the time a human performs the task by at least half, with tasks in this paper referring to O*NET's Detailed Work Activities (DWA). They use a human and GPT rating of each DWA and aggregate at the occupational level. We use an average of their human and GPT ratings for each occupation in our score.
 - **Suitability for Machine Learning Score, or SML Score (Brynjolfsson et. al):** This paper focuses on machine learning and it uses a rubric to determine how applicable ML is to the 2,069 Detailed Work Activities (DWA) from O*NET. Its scores are then aggregated at the occupational level.
 - **Advances in AI Score, or AAI Score (Felten et. al):** This paper links categories of AI to the 52 O*NET abilities and measures the impact AI has had on that ability. It is then aggregated at the occupational level.
 - The TB and SML Scores cover AI and ML exposure respectively by viewing it at the DWA level, while the AAI Score focuses on past advancements at a broader level, which helps control for changes that may not be captured by the DWA level scores.
2. To determine the weight of each score, we use the principal component analysis (PCA) method. It objectively identifies the most significant patterns in the data while eliminating redundancy. PCA transforms the indices into uncorrelated principal components, with the first principal component (PC1) capturing the maximum variance across all scores. By using PC1 loadings as weights, we have created a composite index that optimally represents overall AI exposure, ensuring that each metric is weighted based on its true impact rather than arbitrary assumptions. The weights also reflect our belief that the TB Score offers the best insight on exposure, followed by the SML Score.
 - **PCA weights:**
 - i Task Based Score: 0.751384
 - ii Suitability for Machine Learning Score: 0.589772
 - iii Advances in AI Score: 0.295958

3. We matched the scores from the TB, SML, and AAI Scores with the 1,016 occupations provided by OEWS, as all three scores use the standardized occupation list from O*NET.
 4. To obtain the exposure score at the occupation level, we multiply the PCA weights by their respective scores and add the three outputs. This weighs each score according to PCA.
- The TB Score calculates data for 91% of occupations, SML Score for 70% of occupations, and AAI Score for 67% of occupations. Occupations may have all, some, or none of the three scores. To account for this, we normalize the weights of each occupation based on the scores that are available. Normalization allows us to view the numbers as a score out of 1 rather than an arbitrary number and handles the uneven distribution of the three scores effectively.
 - i For example, if all three scores are available, it would be divided by 1.6371141, the sum of all three PCA weights, to normalize it.
 - ii If only two were available, such as the TB and AAI Scores, it would be divided by the sum of their only PCA weights, in this case 1.047342.
 - iii If there is only one score available, it is divided by that score's PCA weight only.
 - iv However, if there are no scores available for that occupation, then that occupation will not be considered in the exposure score and is removed from the total employment numbers at every geographic level.

Geographic Level Exposure Scores

For each metropolitan and nonmetropolitan area as well as for each state and territory, we create a table with the area on the rows, the 1,016 occupations in the OEWS dataset at the columns, and the May 2024 OEWS employment numbers populating the table.

- We multiply each occupation by its respective exposure score and we take the sum of these numbers at the geographic level. We divide this value by the total number of employees considered in that area to receive the exposure score at the geographic level.
 - For example, Chief Executives have an exposure score of 0.63 and Abilene, TX has 50 people working in this position. We multiply these numbers to get 31.5, interpreted as the AI exposure all Chief Executives of Abilene have in total. This is repeated for every occupation and summed, equaling 31890.5. This is then divided by the number of valid workers in Abilene, 63940, which gives us the metropolitan level exposure score of 0.50.

LLM Score Methodology

Pre-processing: We gather data on LLM usage (Appel, McCrory, Tamkin, et al. 2025, Tomlinson et al. 2025) and compile the different sources into a combined statistic of usage by IWA.

We obtain data from Claude (Appel, McCrory, Tamkin, et al.) and convert it into an assessment of automation by IWA. During pre-processing of the Claude data, we use O*NET's 20.1 data dictionary, matching the structure of the data. (For all other computations, we use the more current 30.0 data dictionary.)

- 1 The Claude data is organized by task (from the O*NET 20.1 data dictionary) and interaction type, which allows us to obtain a score for each of the tasks they list, specifically for automation-type interactions.
 - i For each task, Anthropic lists a number of variables, including:
 - The number of Claude interactions in their dataset that they have assessed as pertaining to that task and originating in the United States.
 - Among all Claude interactions pertaining to that task (not limited to the United States), the proportion that they have assessed as belonging to each of five interaction-type categories. Three of these categories are considered to fall under augmentation, while two fall under automation.
 - ii For each task, we obtain an automation score by multiplying its number of total US interactions by the sum of its automated interaction proportions.
 - iii For tasks not appearing in Anthropic's data, we assume an automation score of 0.
- 2 For each IWA, we take the sum of automation scores for all tasks falling under DWAs that fall under that IWA: Where a task is assigned to multiple DWAs, we divide its score evenly among them.
- 3 This sum is taken to be the automation score by IWA for the Claude dataset: An IWA's score can also be interpreted as an estimate of the number of conversations in their dataset originating from the United States and dedicated to automating that IWA.

We obtain data from Microsoft Copilot (Tomlinson et al.) and convert it into an assessment of automation by IWA.

- 1 The Copilot data does list proportions by IWA rather than by task, but it does not divide the interactions into automation vs. augmentation. Instead, it assigns each conversation two (sometimes distinct) IWAs: one for the user's goal, and one for the AI's goal. For each IWA, they thus have two sets of data, each consisting of the following measures:
 - i Among all Copilot interactions, the proportion where the [user/AI] has that IWA as a goal.
 - ii Among all Copilot interactions where the [user/AI] has that IWA as a goal, what proportion they assess as having been completed successfully.

- iii Among all Copilot interactions where the [user/AI] has that IWA as a goal, the average assessed scope of impact of the interaction to the goal overall, given in the data as a score from 0 to 1.
“For each matching IWA in a conversation, we also perform an LLM classification of the fraction of work in the IWA that Copilot demonstrates the ability to assist or perform, which we call the impact scope (or simply scope), measured on a six-point Likert scale: none, minimal, limited, moderate, significant, complete. ... [W]e validate the scope classifiers with human judges blind to the classifier outputs” (Tomlinson et al.).
 - iv (Not used by us) Among all Copilot interactions where the [user/AI] has that IWA as a goal, and where the user has provided thumbs-up or thumbs-down feedback on Copilot’s performance, the proportion where the feedback has been thumbs-up.
 - v (Not used by us) A custom AI applicability score taking into account completion, scope, and coverage (coverage being a binary indicator of whether or not a significant number of interactions involve that IWA).
2. For each IWA, we multiply Microsoft’s score for its share of all IWAs, its success rate, and its average scope. We do this separately for the user-goal scores and the AI-goal scores, giving two scores for each IWA.
 3. We take the minimum of these two scores for each IWA as its automation score, based on the reasoning that interactions where the AI automates a task would involve both the user and the AI successfully completing the same goal.

We combine the data from the different sources into one metric.

1. For each of the above scores, we normalize by dividing through by the total sum for that score. Each constituent score now sums to 1.
2. We take the sum of the scores for each IWA and use this as our combined score.

Main computation: We predict task automation and consequent job loss, based on the findings of Autor and Thompson (2025). All O*NET data is from the O*NET 30.0 data dictionary; all employment data is from BLS Occupation and Wage Statistics.

For each occupation in the NEM (National Employment Matrix), we look at its associated O*NET occupation(s) and gather all tasks listed in O*NET under the associated occupation(s).

1. We use the crosswalk provided by BLS to find associated O*NET occupations for each NEM occupation. The crosswalk associates each O*NET occupation with one NEM occupation, while each NEM occupation is associated with one or more O*NET occupations.

2. Fishing and hunting workers (SOC code 45-3031, corresponding to O*NET occupation Fishing and Hunting Workers, code 45-3031.00) are excluded; their employment data does not appear in the NEM.

For each task, three independent figures are computed:

1. Weight, based on O*NET's measures of relevance and frequency.
 - i O*NET provides the share of respondents who rate a task as falling into each of seven frequency buckets: "Yearly or less," "More than yearly," "More than monthly," "More than weekly," "Daily," "Several times daily," and "Hourly or more." We use Tomlinson et al.'s associated annual counts of 1, 4, 24, 104, 260, 780, and 2080 respectively.
 - ii The weight for each task is given by its average annual count. This is computed by taking the average of the seven annual counts above, weighted by the number of responses falling into each bucket, and then multiplying by the Relevance of Task score, to account for those whose jobs do not involve the task at all.
 - iii Thus, the weight for each task corresponds to the number of times an average worker in the relevant occupation performs that task over the course of a year.
 - iv Tasks with no associated scores are assigned a uniform weight of 1, so their impact will be negligible for occupations that also have tasks with score data, but they can still be used to compute averages for occupations whose tasks are all missing data.
2. Degree of automation, based on the data compiled in pre-processing.
 - i Each task corresponds to one or more DWAs (Detailed Work Activities), each of which in turn corresponds to exactly one IWA (Intermediate Work Activity).
 - ii Where a task corresponds to more than one DWA, we split it among its DWAs, dividing the weight equally. From this point forward we treat these as separate tasks, so that we can say each task falls under only one DWA.
 - This split does not affect the results of any aggregations to the occupation level before downweighting is applied, since all have the same values and the same cumulative weights as the original tasks.
 - iii Since the data on LLM usage is expressed as a proportion of overall queries, we reweight the numbers to take into account the frequency of IWAs in the economy, following the method of Tomlinson et al.
 - Each task's weight is multiplied by the number of people employed in its associated occupation to obtain an estimate of the number of times that task is performed per year in total.
 - Each IWA is assigned a score, computed as the sum of the average annual frequencies over all tasks assigned to DWAs under that IWA.
 - These scores are divided by their sum, yielding the proportion of work dedicated to each IWA.
 - The LLM usage data from pre-processing is divided by the newly obtained scores, giving a set of figures describing LLM usage relative to overall IWA frequency.

- iv To each task, we assign the readjusted LLM usage figure of its associated IWA (that is, the IWA associated with the DWA associated with the task).
 - v Since these figures have no inherent unit, we scale them so that the maximum corresponds to a parameter we set: the automation factor, which we take to be the extent to which the most vulnerable tasks are already being automated by LLMs.
 - Our choice of automation factor is 0.25, which indicates an assumption that the most vulnerable tasks are already 25% automated.
 - Other tasks' scores are scaled by the same amount to produce an across-the-board assessment of automation levels by task.
 - vi To predict future levels of automation, we assume that a given task's capacity to be automated follows a logistic trajectory, with all tasks lying at different points along the same curve. We advance each task by a fixed amount, which is another parameter we set: time progression, measured in units along the x-axis of a standard logistic curve.
 - Our baseline choice of time progression is 5. For minuscule starting values, this corresponds to a multiplication by about 150; looking at larger starting values, it takes 1% to 60%, 10% to 94%, and 25% to 98%.
 - This progression represents the passage of time on a conceptual level, but there's no pre-defined conversion rate between units of time progression and real-world years.
 - vii We assign the results as the automation scores for our tasks.
3. Expertise level, based on the task statement given by O*NET.
- i Each task statement is converted to a 384-dimensional vector using Python's `sentence_transformers` library. This is a semantic embedding, meaning that geometric features of the embedding space correspond to semantic features of language.
 - The specific embedding model used is `sentence-transformers/all-MiniLM-L6-v2` (the library's most popular model as of 2025-11-12).
 - ii We perform a weighted least-squares linear regression (using the weights previously computed) from the task representations to O*NET's listed Job Zone score for each task's associated occupation, which we take to be an approximate measure of expertise. The logic is that this regression, by using tens of thousands of data points, will cut through the noise (e.g., some relatively inexpert tasks belonging to high-level jobs and vice versa) to get at a "true" mapping from semantic space to expertise level.
 - iii This function is then applied to each task to assign it an expertise score.

New weights are computed by scaling down the original weights, based on the degree of automation. (Expertise is not taken into account at this stage.)

1. For each task, the new weight is computed by subtracting the automation score from 100% and multiplying by the original weight.

For each occupation, we compare the new picture of the job with the old. Specifically, we measure two changes that are effected by the shift in weights:

1. The change in the occupation's expertise level, taken as the weighted average of the expertise levels among its tasks.
 - i This weighted average is taken once using the original weights, and once using the new weights. We compute the difference between the original-weight expertise level and the new-weight expertise level, and normalize by dividing by the standard deviation of original-weight expertise levels across all occupations.
2. The change in total task quantity (sum of all task weights): The sum of the new weights is subtracted from the sum of the old weights, and the result is divided by the sum of the old weights to obtain percentage task loss.

Using these numbers, we predict the percentage job loss for each occupation.

1. Conceptually, the change in expertise level acts as a shift in labor supply, with the number of qualified people changing in inverse proportion to the level of expertise required. The change in task quantity acts as a shift in demand.
2. Mathematically, Autor and Thompson found a function of best fit from these changes to changes in employment over recent decades, which we use to make our own projections.
 - i Autor and Thompson ran several linear regressions to express employment change as a function of shifts in expertise and quantity of tasks. Specifically, their predictors were (1) change in expertise, as a multiple of the standard deviation of expertise among all occupations at the beginning of the analysis period, and (2) change in number of tasks, as a percentage of the number of tasks for the given occupation at the beginning of the analysis period.¹ Their dependent variable was the change in the logarithm of the number of people employed.
 - ii We can choose which of these functions we use to predict employment change. Our choice is their regression in which they exclude four outliers, but do not treat task addition and removal separately, as this is the only one where both predictor variables have $p \leq 0.01$.

¹One of their regressions considered four predictor variables: change in expertise as a result of task loss, change in expertise as a result of task gain, percentage of tasks lost, and percentage of tasks gained.

- iii In the previous step, we performed the same normalizations that Autor and Thompson used to generate their predictor variables. Therefore, we just take the linear combination of these figures, using Autor and Thompson's coefficients, to predict log change in employment.
- iv We then take e to the power of the predicted log change to obtain predictions of the ratio between future employment and current employment, and subtract this from 100% to obtain percentage job loss.

Adjustments: We reduce the predicted job loss for some occupations based on the degree of physical activity required. Because legislators are not susceptible to market forces in the same way most occupations are, their predicted job loss is set to 0. All O*NET data from version 30.0.

1. Some physical tasks share IWAs with less physical tasks. For instance, the IWA "Provide information or assistance to the public" includes both "Write original content, such as headlines, cutlines, captions, and cover copy" (proofreaders and copy markers) and "Announce entrance of judge" (bailiffs). Thus, jobs involving physical tasks can be unfairly penalized if those tasks fall under the same category as a highly vulnerable task. We apply protections to jobs that require physical activity so that their predicted loss will not be unreasonably high.
2. We use five O*NET measures as proxies for physical activity: four Work Activities and one Work Context.
 - i Performing General Physical Activities: "Performing general physical activities includes doing activities that require considerable use of your arms and legs and moving your whole body, such as climbing, lifting, balancing, walking, stooping, and handling materials." This Work Activity is important for occupations including models, lifeguards, physical therapist assistants, and dancers.
 - ii Handling and Moving Objects: "Using hands and arms in handling, installing, positioning, and moving materials, and manipulating things." This Work Activity is important for occupations including library assistants, team assemblers, embalmers, and surgical technologists.
 - iii Controlling Machines and Processes: "Using either control mechanisms or direct physical activity to operate machines or processes (not including computers or vehicles)." This Work Activity is important for occupations including nuclear power reactor operators, bridge and lock tenders, anesthesiologists, and ophthalmologists.
 - iv Operating Vehicles, Mechanized Devices, or Equipment: "Running, maneuvering, navigating, or driving vehicles or mechanized equipment, such as forklifts, passenger vehicles, aircraft, or watercraft." This Work Activity is important for occupations including bus drivers, chauffeurs, airfield operations specialists, and truck drivers.
 - v Spend Time Sitting: "How much does this job require sitting?" The inverse of this Work Context indicator is used as a general proxy for physical activity, as time spent not sitting is assumed to be time spent on tasks that require a physical presence. Occupations deemed physical by this metric include short-order cooks, cafeteria attendants, bakers, and crossing guards.

3. All five are rated on a scale from 1 to 5 by O*NET: we use the Importance score for the Work Activities, and the measure of Spend Time Sitting is given by O*NET as an average of responses on a scale from 1 (Never) to 5 (Continually or almost continually). We reverse this last scale by subtracting the values from 6 (so that 1 now implies the most sitting and 5 the least), and then take the maximum of the five indicators for each O*NET occupation.
4. We subtract 1 and divide by 5 in order to obtain the scores on a scale from 0 to 1 rather than 1 to 5, and then reverse the scale (i.e. subtract from 1) to generate physicality coefficients: the lower the coefficient, the lower the degree of job loss.
5. Thus far, these numbers have been computed for O*NET occupations. We compute them for NEM occupations in two steps:
 - i First, for each NEM occupation that corresponds to at least one O*NET occupation with available data, the NEM occupation's physicality coefficient is taken as the average of the physicality coefficients of its associated O*NET occupations. (In most cases, there is a one-to-one correspondence between O*NET and NEM occupations, and the latter simply inherit the former's coefficients.)
 - ii Where an NEM occupation has no data available for any of its corresponding O*NET occupations, its coefficient is taken as the average of the coefficients of the most similar NEM occupations, taken to be those whose occupational codes share the greatest number of leading digits.
 - For instance, if one or more NEM occupations share all but the last digit of their occupational codes with the occupation in question (and have data available), the average of these is taken; otherwise, the average is taken of all those where only the last two digits differ, if available, and so forth.
6. Coefficients are grouped into 4 buckets: less than .25; at least .25 but less than .5; at least .5 but less than .75; and at least .75.
7. These buckets are used to apply modifiers to the job loss predictions: occupations in the first bucket have their predicted loss multiplied by 0 (i.e., they are predicted not to lose any employees); the second bucket by $\frac{1}{3}$; the third bucket by $\frac{2}{3}$; and the fourth bucket by 1 (i.e., they are not predicted to be protected by the physical components of their jobs).
8. Unrelated to the physicality scores, legislators' AI-related job loss is also set to 0, in recognition of the fact that their employment levels are not governed by the same market forces that drive employment in most occupations.

Job and Income Loss Methodology

Pre-processing

- We gather occupational employment data at the metropolitan area, nonmetropolitan area, state, and territory levels using O*NET's database (specifically the May 2024 Occupational Employment and Wage Statistics data).
- We create a new column called IS_VALID with the equation $=IF(ISNUMBER(TOT_EMP * A_MEDIAN), 1, 0)$, where TOT_EMP and A_MEDIAN are the corresponding cells in each row. This column will return 1 if both the TOT_EMP and A_MEDIAN values of that row are valid numbers, and 0 if either cell is not a number.
- We create a new column called JOB_LOSS. Using the OCC_CODE column for searching in the XLOOKUP function, we populate this column with the job loss numbers described in the LLM Score Methodology, then copy-paste the JOB_LOSS column to make it editable.

Preparing the Net Job Loss Data

- We create a new column called NET_JOB_LOSS and populate it column by multiplying the JOB_LOSS, TOT_EMP, and IS_VALID columns to get the net job loss for each occupation.
- If an occupation does not have employment or job loss data, this column will return an error (either #VALUE! or #N/A"). To delete these values, copy-paste the NET_JOB_LOSS column to make it editable and highlight the column, then use Control + F to replace all "#N/A" and all "#VALUE!" in this column with a blank cell.
- We create a new column, TOT_EMP_JOB_LOSS and populate it by using this function: $=IF(ISBLANK(NET_JOB_LOSS), 0, TOT_EMP) * IS_VALID$. This column returns the TOT_EMP if there is a valid NET_JOB_LOSS number available and zero if it does not. We then multiple it by IS_VALID to only count valid employment numbers. This column can now be used to consider the total employment numbers for just the occupations with valid job loss data.

Preparing the Income Loss Data

- We create a new column called INCOME_LOSS. We have used the median income to calculate the income loss because the median income is generally considered a more accurate representation of the average worker than the mean and is not impacted by high income outliers like the mean is. In most occupations, AI has the highest impact on entry level workers rather than experienced workers, however not exclusively. Therefore, the median income is more representative of the workers that are losing their jobs than mean income.

- We populate this column by using this equation: $=\text{NET_JOB_LOSS} * \text{A_MEDIAN} * \text{IS_VALID}$. The first part of the equation multiplies the job loss number by the median income, and by multiplying by IS_VALID, it will cause occupations without both median income and employment numbers to return 0. Copy and paste this column and highlight them, then use Control + F to replace all "#N/A" and all "#VALUE!" in these columns with a blank cell.
- We create a new column called TOT_EMP_INCOME_LOSS. Populate this column by multiplying the TOT_EMP_JOB_LOSS, the A_MEAN, and the IS_VALID columns, with the equation: $=\text{TOT_EMP_JOB_LOSS} * \text{IF}(\text{ISNUMBER}(\text{A_MEAN}), \text{A_MEAN}, \text{A_MEDIAN}) * \text{IS_VALID}$. This will multiply the valid occupations by the mean income if it exists, and if not, it will be substituted by the median income. This column returns the estimated total income earned by every worker for that occupation. Then copy and paste these new columns and highlight them, then use Control + F to replace all "#N/A" and all "#VALUE!" in these columns with a blank cell.

Metropolitan/Non-metropolitan Area Level and State/Territory Level Tables

- We create a pivot table with AREA_TITLE in the Rows section and the following in the Values section as Sums (not Counts): NET_JOB_LOSS, TOT_EMP_JOB_LOSS, INCOME_LOSS, and TOT_EMP_INCOME_LOSS. This aggregates all job and income loss data by area and state.
- We add O_GROUP as a filter and select "detailed" to only use the detailed occupations. We also add AREA_TYPE as a filter: When calculating metropolitan and non-metropolitan areas, check both the 4 and 6 boxes. When evaluating at the state and territory level, check the 2 and 3 boxes.
- We copy this table to a new sheet. Use the =ROUND() function to round all columns to the nearest whole number. Remove "Sum of " from all column names.
- We create two new columns, titled PERCENT_JOB_LOSS and PERCENT_INCOME_LOSS. We divide the NET_JOB_LOSS by the TOT_EMP_JOB_LOSS for PERCENT_JOB_LOSS; We also divide the INCOME_LOSS by the TOT_EMP_INCOME_LOSS for PERCENT_INCOME_LOSS.
- We create a new column titled INCOME_TO_LOSS_RATIO. We populate this column by dividing the PERCENT_INCOME_LOSS by the PERCENT_JOB_LOSS. This column generally ranges between 1 to 2, with values close to 1 representing regions where the average job losses are in median income jobs, while regions with a value closer to 2 have job losses in higher paying occupations.

Industry Level Tables

- Using the same pivot table, we remove all filters except the O_GROUP, to continue only using the detailed occupations. We add the I_GROUP column and select “sector” for the filter, which aggregates by the 20 main industries. “3-digit” can be used to break down those industries into the 85 subsectors and “4-digit” for the 247 industry groupings in place of using “sector.”
- We copy this table to a new sheet and use the =ROUND() function to round all columns to the nearest whole number. We also remove “Sum of “ from all column names.
- We create two new columns, titled PERCENT_JOB_LOSS and PERCENT_INCOME_LOSS. We divide the NET_JOB_LOSS by the TOT_EMP_JOB_LOSS for PERCENT_JOB_LOSS; We also divide the INCOME_LOSS by the TOT_EMP_INCOME_LOSS for PERCENT_INCOME_LOSS.
- We create a new column titled INCOME_TO_LOSS_RATIO and populate this column by dividing the PERCENT_INCOME_LOSS by the PERCENT_JOB_LOSS. This column generally ranges between 1 to 2, with values close to 1 representing regions where the average job losses are in median income jobs, while regions with a value closer to 2 have job losses in higher paying occupations.

Sources

- Appel, Ruth, Peter McCrory, Alex Tamkin, et al., Sep 2025. "Anthropic Economic Index September 2025 Report." <https://www.anthropic.com/research/anthropic-economic-index-september-2025-report>. Data at https://huggingface.co/datasets/Anthropic/EconomicIndex/tree/main/release_2025_09_15.
- Autor, David, and Neil Thompson, Jun 2025. "Expertise." <https://doi.org/10.1093/jeea/jvaf023>.
- Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock, May 2018. "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" AEA Papers and Proceedings 108 (2018): 43–47. <https://www.aeaweb.org/articles?id=10.1257/pandp.20181019>.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, Aug 2023. "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models." <https://arxiv.org/abs/2303.10130>.
- Felten, E., Raj, M., & Seamans, R. Apr 2021. "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses." Strategic Management Journal, 42(12), 2195–2217. <https://doi.org/10.1002/smj.3286>.
- Handa, Kunal, et al., Feb 2025. "Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations." <https://doi.org/10.48550/arXiv.2503.04761>. Data at https://huggingface.co/datasets/Anthropic/EconomicIndex/tree/main/release_2025_02_10.
- National Center for O*NET Development, Oct 2015. "O*NET 20.1 Data Dictionary." <https://www.onetcenter.org/dictionary/20.1/excel/>.
- National Center for O*NET Development, Aug 2025. "O*NET 30.0 Data Dictionary." <https://www.onetcenter.org/dictionary/30.0/excel/>.
- Tomlinson, Kiran, et al., Jul 2025. "Working with AI: Measuring the Applicability of Generative AI to Occupations." <https://doi.org/10.48550/arXiv.2507.07935>. Data at <https://github.com/microsoft/working-with-ai>.
- U.S. Bureau of Labor Statistics, Apr 2024. "Occupational Employment and Wages, May 2023." <https://www.bls.gov/oes/tables.htm>. Direct link to ZIP file (large): <https://www.bls.gov/oes/special-requests/oesm23all.zip>.
- U.S. Bureau of Labor Statistics, May 2024. "Occupational Employment and Wage Statistics (OEWS) Tables." <https://www.bls.gov/oes/tables.htm>. Direct link to ZIP file (large): <https://www.bls.gov/oes/special-requests/oesm24all.zip>.
- U.S. Bureau of Labor Statistics, Aug 2025. "National Employment Matrix." <https://www.bls.gov/emp/data/occupational-data.htm>. Direct link to spreadsheet: <https://www.bls.gov/emp/ind-occ-matrix/matrix.xlsx>.
- U.S. Bureau of Labor Statistics, Aug 2025. "O*NET-SOC to Occupational Outlook Handbook Crosswalk." <https://www.bls.gov/emp/documentation/crosswalks.htm>. Direct link to spreadsheet: <https://www.bls.gov/emp/classifications-crosswalks/nem-onet-to-soc-crosswalk.xlsx>.



*The graduate school
of global affairs
at Tufts University*