FEDERAL RESERVE BANK OF MINNEAPOLIS
Community Development and Engagement

WORKING PAPER

The Evolution of Employment Dispersion

April 2024

Ayushi Narayan
Federal Reserve Bank of Minneapolis

Ryan Nunn
Federal Reserve Bank of Minneapolis

Keywords: Employment, Inequality, Labor Market, Measurement

The Evolution of Employment Dispersion

Ayushi Narayan and Ryan Nunn[1]

April 1, 2024

Abstract

Labor market inequality encompasses both dispersion in earnings and employment. While earnings dispersion is relatively straightforward to assess, employment dispersion presents difficulties. We offer a new approach that uses observable worker-level information to predict employment propensities, then tracks changes in employment at different points in the predicted distribution. This approach is useful for a few purposes: first, to provide a more comprehensive picture of labor market inequality; second, to identify (in a fine-grained way) groups with relatively low employment propensities; and third, to facilitate explorations of how macroeconomic conditions (e.g., tight labor markets) may affect labor market inequality.

[1] The authors are grateful to Erik Hembre and Jennifer Hunt for insightful feedback on an earlier draft.

Introduction

Inequality in labor income stems from both disparities in the wages earned among those employed and disparities in employment itself. Individuals or groups at the bottom of the labor income distribution could have low wages, low levels of employment, or both.

Examining dispersion in wages is relatively straightforward because wages are a continuous variable. Analysts can examine a wage distribution and report a variety of conventional statistics like the ratio of the 90th and 10th wage percentiles.[2] When that ratio rises, wage inequality is understood to rise, and vice versa. The chief methodological difficulty in using wage dispersion alone to study inequality in labor income is that wages are only recorded for employed workers. It is thus important to study imbalances in employment, as well as wages.

However, employment—a binary variable—cannot be assessed in exactly the same way as wages.[3] Labor economists interested in employment inequality will typically calculate employment rates conditional on a single variable like education, sex, or race. For some questions, this approach is adequate: for example, if the analyst cares specifically about how individuals of differing education levels or demographic groups are faring in relation to each other. But for other questions, such as those about dispersion in employment probabilities more generally, this approach is likely to be unsatisfactory. Consider that in a regression of employment on bins of educational attainment in 2022, education can only account for 3 percent of the individual-level variation among prime-aged (ages 25–54) civilians. Education is therefore unlikely to fully capture the underlying variation in employment propensity that constitutes employment inequality.

To understand how labor market conditions are evolving more generally for individuals with differing propensities to work—i.e., to assess employment dispersion—education is an imperfect proxy for what the analyst cares about. In this case, it would be more useful to calculate employment rates across the distribution of the underlying propensity to work. For instance, how is actual employment evolving for individuals at the 90th percentile of employment propensity relative to the 10th percentile? Unfortunately, this is not immediately observable in the same way that the wage distribution is observable. For this type of question, we would like to approximate an individual-level distribution of employment propensity. These individuals can have varying employment propensities that are not perfectly correlated with their educational attainment or with any other single demographic variable. Of course, we are not able to fully recover the latent employment propensity distribution, but the use of a rich set of demographic characteristics allows for a better approximation of that distribution than is possible with education (or race or sex) alone.

We therefore propose a new way to unpack binary outcomes like employment and labor force participation.[4] Our approach uses observable worker-level information to predict employment propensities, then tracks changes in employment at different points in the predicted distribution. This approach is useful for a few purposes: first, to provide a more comprehensive picture of labor market inequality; second, to identify (in

[2] See Autor, Dube, and McGrew (2023) for a careful investigation of recent wage inequality trends.
[3] Another difference, which is not the focus of this paper, is that wages have a clearer mapping to worker welfare than does employment. With some limited exceptions (e.g., compensating differentials), lower wages indicate lower utility. By contrast, lower employment can reflect limited access to jobs, but it can also reflect a choice not to take paid employment.
[4] For some analytical questions, it may be appropriate to limit consideration to labor force participants, thereby excluding those who have not demonstrated a desire for market employment. In this article, we focus on employment.

a fine-grained way) groups with relatively low employment propensities; and third, to facilitate explorations of how macroeconomic conditions (e.g., tight labor markets) may affect labor market inequality.

Our proposed method aims to approximate an underlying individual-level distribution of employment propensity, then track actual employment rates over time by quantiles of that underlying distribution. We approximate the individual-level distribution with a basket of observable factors in the Current Population Survey microdata, including education but also demographic characteristics like age, sex, race, nativity, veteran status, marital status, and presence of own children in the household, as well as geographic characteristics. In our primary application of this approach, we estimate employment for individuals from 2017–18 and predict quintiles of employment propensities for 2015–23 using that estimation. Our baseline methodology uses a random forest research design that flexibly allows for interactions across variables to estimate employment in 2017–18, but the patterns we uncover are similar if we use a linear probability model or LASSO.

Results obtained from the new approach, in any of its variants, are not strictly comparable to those obtained from disaggregations of employment rate by education, race, or sex. Take education, for example. The lowest level of attainment does not generally correspond to a quintile of individuals (leaving aside the fact that it imperfectly corresponds to the *lowest* employment individuals). We see this as a virtue of the new approach: it provides a more consistent comparison over time between two groups within the underlying distribution of employment propensity.[5]

We find that the likelihood of employment increased most for the lowest predicted quintile of employment in the late 2010s as the labor market tightened. This corresponded with a falling dispersion in the propensity for employment across quintiles. Individuals with the lowest predicted likelihoods of working also experienced the sharpest drops in employment during 2020, which saw an increase in employment dispersion. However, by mid-2023 the lowest predicted quintile of individuals—as well as those across the entire distribution—were back at employment levels similar to those achieved at the beginning of 2020. The tight labor markets of 2023 and January 2020 were thus associated with relatively compressed employment propensities.

These results complement research examining recent trends in wage dispersion adjusted for the composition of the workforce (e.g., Autor, Dube, and McGrew 2023). We find that employment dispersion declined at the same time that adjusted wage dispersion fell. Together, these insights on wage and employment dispersion paint a clearer picture of recent trends in labor income inequality. Our approach is also useful for identifying observable characteristics that correspond to labor market disadvantage, which in turn can be used to investigate the relationship between labor market tightness and inequality. At the end of this article, we show how this investigation could proceed in the context of our approach.

Trends in employment by education, race, and sex

Though not always characterized as reflecting employment inequality per se, graphs of employment rates by education (Figure 1), sex (Figure 2), and race (Figure 3) are commonly generated to show how the extensive margin of the labor market is evolving. Each provides a window into how changing labor demand and supply are shaping the labor market, and each has been used as a launching point for studies of secular trends in employment and participation.

[5] A caveat to this statement applies to the extent that—in periods outside the estimation period (e.g., 2017–18 in our baseline case)—the share of individuals within each group may deviate somewhat from 20 percent. Later in the article, we discuss a variant of our approach that re-estimates the model in each period and thereby avoids this issue.

In the figures below, we observe declining gaps in employment, disaggregated in these three ways since 2015. Less-educated individuals, who are less likely to work overall than more highly educated individuals, increased their propensity to work between 2015 and January 2020 more than more highly educated individuals. In a temporary reversal of this trend, less-educated people experienced larger employment declines during the 2020 recession and aftermath but have returned to their 2020 employment levels by the end of 2023. Similar patterns are apparent by sex, with the likelihood of work increasing markedly for women between 2015 and January 2020, dropping slightly more during 2020, and reaching January 2020 levels by 2023. By race, there is likewise convergence in the Black-White employment rate gap through the beginning of 2020 that then sharply reverses during the 2020 recession, finally rebounding by mid-2023. Latino individuals follow a similar pattern as White individuals from 2015 through the beginning of 2020, though they experience sharper declines during 2020. After that, sharper growth for Latinos during the recovery period returns the Latino-White employment gap back to pre-2020 recession levels. Patterns are similar for the Asian American and Pacific Islander population with particularly strong growth since 2020 leading to convergence over the entire period. For American Indians and Alaska Natives, our estimates are noisier but remain suggestive of compression from 2015 to 2023.

Overall, these patterns suggest recent convergence in employment propensities based on various population characteristics. They inspire an examination of how employment rates have varied over time for individuals who ex-ante would be considered more or less likely to work. The following sections develop a more formalized way to assess how employment has evolved for groups with lower likelihoods to work compared to groups with higher likelihoods of working.

Figure 1. Employment by education, 2015–23

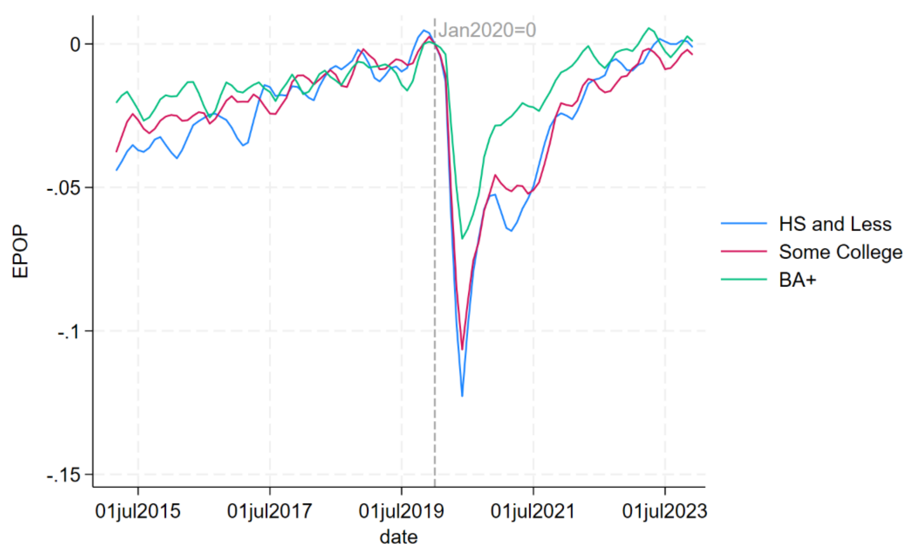

Notes: Figure presents data from the January 2015–December 2023 Current Population Surveys. Sample is restricted to civilians ages 25–54. Figure presents a trailing 3-month moving average of the employment-to-population ratio additively indexed at January 2020=0.

Figure 2. Employment by sex, 2015–23

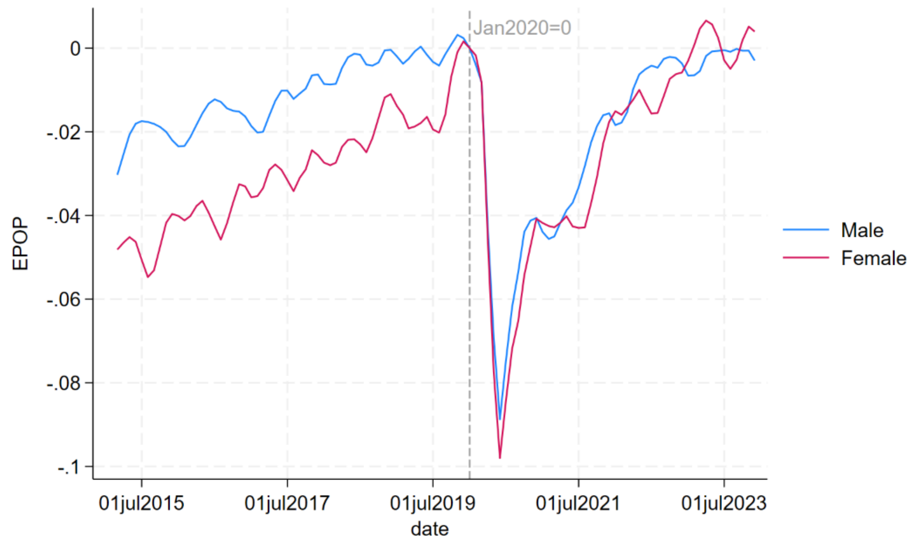

Notes: Figure presents data from the January 2015–December 2023 Current Population Surveys. Sample is restricted to civilians ages 25–54. Figure presents a trailing 3-month moving average of the employment-to-population ratio additively indexed at January 2020=0.

Figure 3. Employment by race/ethnicity, 2015–23

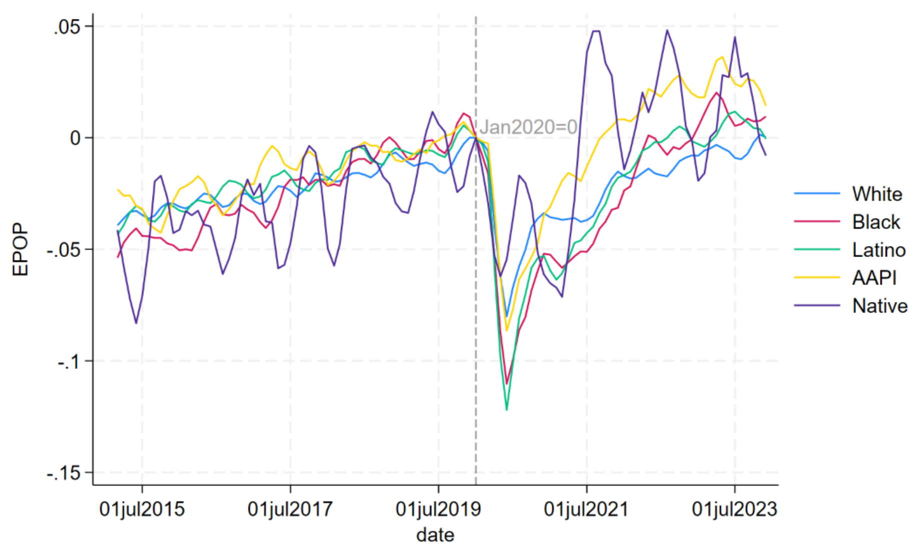

Notes: Figure presents data from the January 2015–December 2023 Current Population Surveys. Sample is restricted to civilians ages 25–54. Figure presents a trailing 3-month moving average of the employment-to-population ratio additively indexed at January 2020=0.

A new approach to measuring employment inequality

Disaggregations by education, race, or sex each have utility for answering specific questions. But to gain a more comprehensive understanding of recent trends in employment dispersion, we take the following approach. We restrict our population to prime-aged (25–54) civilians, i.e., those not employed in the armed forces. We use as our dependent variable an indicator for employment and as our independent variables five-year age categories, an indicator for sex, six race and ethnicity categories, indicators for whether an individual was born in the U.S. or is a U.S. citizen, an indicator for veteran status, an indicator for marital status, variables describing the number of own children in the household and the number of children under the age of 5, and finally geographic indicators including metropolitan status and state of residence. We exclude potential independent variables that could be affected by employment status, such as those related to disability or occupational licensing.

Our preferred methodology uses a random forest algorithm. A random forest algorithm is a machine learning method that uses numerous decision trees, which are essentially flowcharts for making decisions. For example, at the top of the tree might be a question of "Is the sex male?" Depending on the answer, there is another set of branches, such as, "Are there two children" and "Is the race Black?", and so forth, until the final outcome of the decision tree. In the random forest algorithm, each of the numerous decision trees is trained on a different subset of data with slightly different rules and makes its own prediction. The random forest algorithm combines all of these predictions to come up with a final decision. This helps reduce overfitting and makes the model more accurate in mapping worker characteristics to employment.

We implement a random forest algorithm with a depth of 5 and a regression decision tree type to predict employment based on the individual characteristics listed above. We implement the random forest algorithm for the years 2017–18 to avoid picking up on idiosyncratic employment patterns during the COVID-19 pandemic, as well as to make it visibly clear in the figures we present that over-fitting is not driving our results. We opt for a random forest algorithm in our preferred approach to allow for interactions among the observable variables while limiting the amount of structure we impose on the mapping between characteristics and employment. The patterns we uncover are similar when using a LASSO or linear probability model.

After implementing the algorithm for the 2017–2018 period, we use the algorithm results to predict employment for the entire analysis period of 2015–23. Once we have predicted individual-level employment for our analysis period, we construct quintiles of the predicted employment variable in the base period of 2017–18. For each month throughout the entire sample period, we calculate each group's average employment rate. That is, we use the predicted employment cutoffs implied by the 2017–18 quintiles as a way of consistently grouping individuals throughout the sample period, such that the number of individuals in any given group may rise or fall over time, but an individual with a given *predicted* employment probability will always fall into the same group. Among members of each group, the *actual* employment rate will vary over time.

Because our monthly data are somewhat noisy, our figures present trailing three-month moving averages of the observed (indexed) employment rates, which helps to more clearly see trends. We additively index the employment rate for each group to 0 in the base period of January 2020. As a result, the figures present a calculation that show *percentage point* changes in the group's employment rate relative to the base period. We later show that our findings are qualitatively similar when using multiplicative calculations that show percent changes; we also discuss the advantages and disadvantages of the two approaches.

Trends in employment inequality

Primary findings

Figure 4 presents the findings from our main approach, showing convergence in employment likelihoods from 2015 through the very beginning of 2020 as the labor market tightened. Lower quintiles of predicted employment experienced the largest increases in employment likelihoods during this time frame. During the labor market downturn caused by the COVID-19 pandemic, we observe a sharp increase in employment dispersion as the bottom quintiles experience the greatest declines in employment. As the labor market subsequently tightens, we again observe compression in employment across quintiles, and by mid-2023, the employment dispersion is similar to that observed in early 2020.

Figure 5 shows the difference between the top and bottom quintile from 2015 to 2023. This presentation of the top quintile-bottom quintile employment gap is similar in spirit to other inequality statistics like the 90-10 wage gap. The figure shows that the difference between the top and bottom quintile employment rates fell by 3.5 percentage points from December 2015–23.

Figure 4. Employment by predicted quintiles, 2015–23

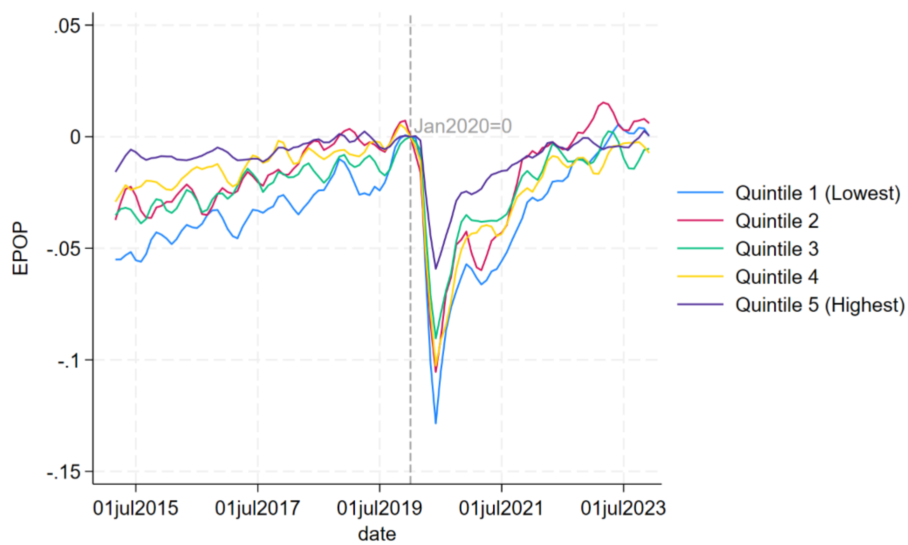

Notes: Figure presents data from the January 2015–December 2023 Current Population Surveys. Sample is restricted to civilians ages 25–54. Figure presents a trailing 3-month moving average of the employment-to-population ratio additively indexed at January 2020=0. Quintiles are constructed using a random forest algorithm for the base period of 2017–18 predicting employment using education, age, sex, race, nativity, veteran status, marital status, presence of own children in the household, and geography variables.

Figure 5. Predicted employment Q5 Less Q1, 2015–23

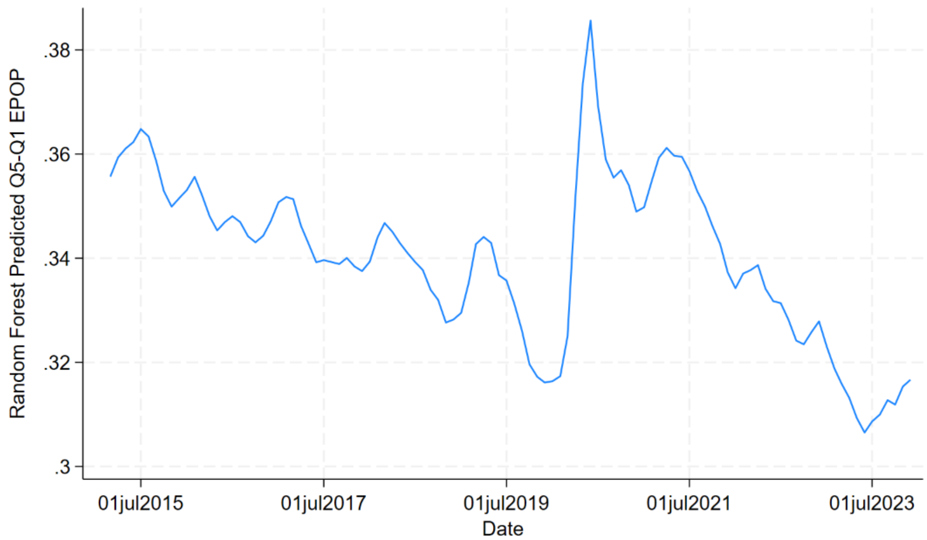

Notes: Figure presents data from the January 2015–December 2023 Current Population Surveys. Sample is restricted to civilians ages 25–54. Figure presents a trailing 3-month moving average of the employment-to-population ratio for the highest quintile less the trailing 3-month moving average for the bottom quintile. Quintiles are constructed using a random forest algorithm for the base period of 2017–18 predicting employment using education, age, sex, race, nativity, veteran status, marital status, presence of own children in the household, and geography variables.

Notably, the cumulative decline shown in Figure 5 occurs at the same time that adjusted wage compression is observed for employed workers (Autor, Dube, and McGrew 2023). The effect of *employment* compression strongly contributes to overall reduction in labor market inequality since the mid-2010s.

Robustness to alternate methodologies

Figures 6 and 7 repeat the analysis presented in Figure 4 above but are based on estimations of a LASSO regression (on all pairwise interactions between the variables noted above, excluding state of residence for computational feasibility) and a linear probability model regression, respectively. The patterns are similar to those shown in Figure 4. In part, this is because all three methods classify individuals into quintiles in relatively similar ways. Table 1 show the make-up of the quintiles yielded by the random forest, LASSO, and linear probability methodologies. For each of the methods, individual characteristics (including but not limited to education) are quite differently distributed across predicted quintiles. For example, the quintiles vary considerably by the share of each who are women, which is much larger in the bottom quintile than in the top.

Figure 6. Employment by predicted quintiles (LASSO variant), 2015–23

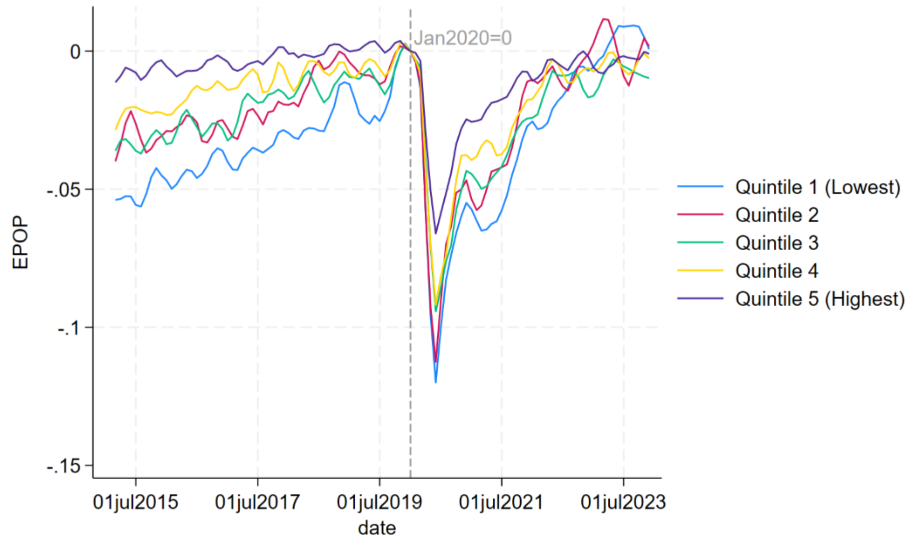

Notes: Figure presents data from the January 2015–December 2023 Current Population Surveys. Sample is restricted to civilians ages 25–54. Figure presents a trailing 3-month moving average of the employment-to-population ratio additively indexed at January 2020=0. Quintiles are constructed using a LASSO regression for the base period of 2017–18 predicting employment using all pairwise interactions of education, age, sex, race, nativity, veteran status, marital status, presence of own children in the household, and geography variables (excluding state for computational feasibility).

Figure 7. Employment by predicted quintile (LPM variant), 2015–23

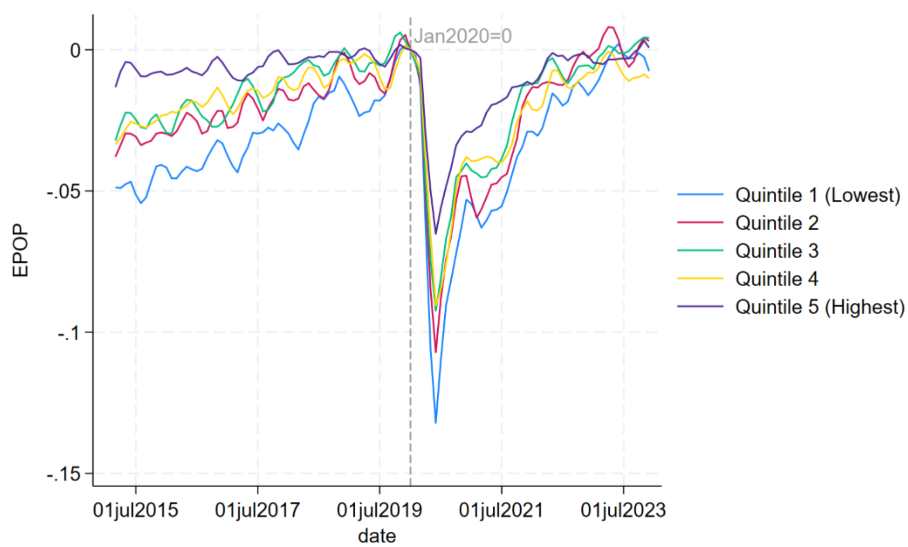

Notes: Figure presents data from the January 2015–December 2023 Current Population Surveys. Sample is restricted to civilians ages 25–54. Figure presents a trailing 3-month moving average of the employment-to-population ratio additively indexed at January 2020=0. Quintiles are constructed using a linear probability model for the base period of 2017–18 predicting employment using education, age, sex, race, nativity, veteran status, marital status, presence of own children in the household, and geography variables.

Table 1. Characteristics of quintiles by prediction methodology

A. Characteristics of random forest quintiles

Quintile	Age	College	Male	White NH	Black NH	Hispanic	Married	Veteran	NChild	NChld <5	Born US	Not Citizen	City
1 (Lowest)	40.10	8%	7%	43%	16%	29%	55%	1%	1.33	0.27	65%	25%	31%
2	38.32	20%	36%	46%	21%	21%	40%	4%	1.10	0.28	81%	4%	31%
3	40.24	51%	30%	65%	14%	14%	54%	4%	0.74	0.08	88%	5%	27%
4	37.49	42%	75%	63%	7%	22%	51%	6%	0.95	0.17	74%	16%	35%
5 (Highest)	39.94	68%	100%	75%	5%	11%	86%	8%	1.37	0.32	81%	8%	29%

B. Characteristics of LASSO quintiles

Quintile	Age	College	Male	White NH	Black NH	Hispanic	Married	Veteran	NChild	NChld <5	Born US	Not Citizen	City
1 (Lowest)	39.70	10%	15%	42%	19%	27%	54%	2%	1.32	0.34	67%	24%	33%
2	39.56	15%	35%	60%	15%	17%	45%	5%	0.93	0.18	84%	4%	28%
3	39.68	44%	43%	61%	14%	14%	52%	6%	0.87	0.12	85%	5%	29%
4	37.90	57%	66%	61%	10%	21%	58%	6%	0.89	0.15	75%	15%	36%
5 (Highest)	39.30	64%	89%	68%	6%	19%	79%	3%	1.47	0.32	78%	11%	26%

C. Characteristics of LPM quintiles

Quintile	Age	College	Male	White NH	Black NH	Hispanic	Married	Veteran	NChild	NChld <5	Born US	Not Citizen	City
1 (Lowest)	39.98	2%	6%	42%	19%	26%	50%	2%	1.37	0.32	69%	22%	32%
2	39.35	21%	36%	49%	18%	20%	50%	5%	0.99	0.25	74%	17%	32%
3	38.32	44%	43%	64%	9%	20%	55%	4%	0.95	0.18	82%	8%	30%
4	39.33	45%	63%	64%	10%	17%	64%	6%	1.12	0.18	81%	8%	28%
5 (Highest)	39.19	78%	98%	73%	7%	13%	68%	5%	1.06	0.18	84%	4%	30%

Notes: Table presents data from 2017–18 Current Population Surveys. Sample is restricted to civilians ages 25–54. Quintiles are constructed using various methodologies for this base period of 2017–18 predicting employment using education, age, sex, race, nativity, veteran status, marital status, presence of own children in the household, and geography variables. Table presents a subset of relevant summary statistics based on these variables.

Figure 8 replicates the main analysis, but rather than scaling the employment shares additively, it uses a multiplicative calculation that is derived by setting each group's employment rate to 100 in the January 2020 index period, then showing *percent* changes in the group's employment rate relative to that period. The patterns from this calculation need not be the same as those from the additive approach. For specificity, suppose that the bottom and top quintile employment rates started in the index period at 40 and 80 percent, respectively, subsequently falling to 36 and 72 percent, respectively. In the multiplicative approach, both groups would have a value of 90 in the subsequent period. In the additive approach, the bottom and top quintile values in the same subsequent period would be -0.04 and -0.08, respectively. In practice, the patterns are qualitatively similar using the multiplicative approach in Figure 8 as the additive approach in Figure 4.

The choice between additive and multiplicative approaches hinges on one's preferred measure of inequality. The multiplicative calculation assumes that inequality has not changed if the ratio of employment rates remains constant, e.g., if rates were to fall from 40 and 80 to 36 and 72 percent. By contrast, the additive calculation assumes that inequality has not changed if the *percentage point* group changes are identical, e.g., if rates were to fall from 40 and 80 to 36 and 76. We use the additive interpretation in our baseline approach given its ease of interpretation of percentage point changes and invariance to how a metric is expressed (e.g., employment rate vs. non-employment rate or participation rate vs. non-participation rate). But a reasonable argument can also be made for the multiplicative interpretation, which offers a similar interpretation to other labor market inequality measures that are conventionally expressed as ratios of group statistics, such as the Black/White unemployment ratio or the 90th–10th percentile wage ratio.

Figure 8. Employment by predicted quintile (multiplicative variant), 2015–23

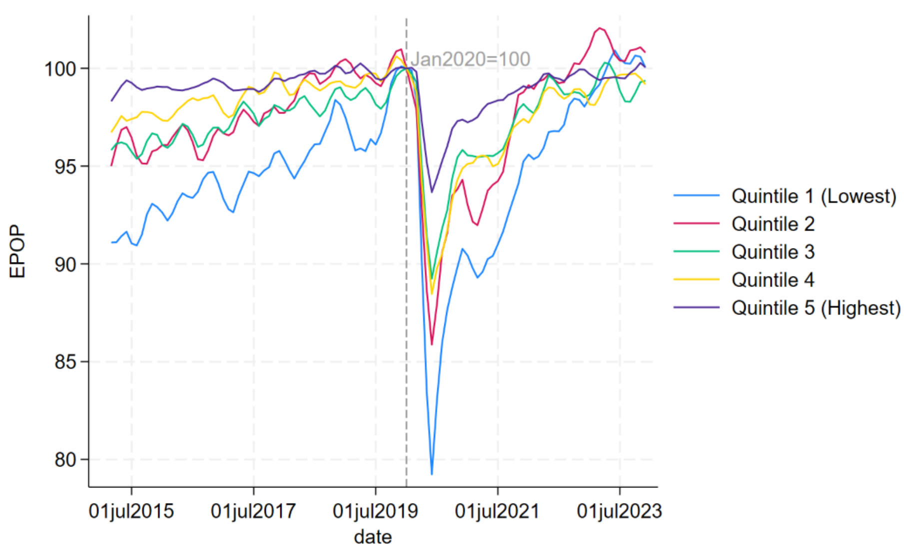

Notes: Figure presents data from the January 2015–December2023 Current Population Surveys. Sample is restricted to civilians ages 25–54. Figure presents a trailing 3-month moving average of the employment-to-population ratio multiplicatively indexed at January 2020=100. Quintiles are constructed using a random forest algorithm for the base period of 2017–18 predicting employment using education, age, sex, race, nativity, veteran status, marital status, presence of own children in the household, and geography variables.

Extension to other time periods

To determine whether the patterns we observe are new or are similar to those seen in other business cycles, we turn to examining employment dispersion leading up to and in the aftermath of the Great Recession (Figure 9). We show employment shares by predicted group for 2003–14, using a random forest algorithm implemented on 2005–06. That is, the mapping from observable factors to predicted employment rates was based on the period 2005–06 and quintile thresholds were based on that period. With the exception of differences in timing, the exercise is analogous to the one carried out for Figure 4 above. We observe a persistent, large increase in dispersion following the Great Recession. Unlike in our analysis of the recession associated with the COVID-19 pandemic, in our analysis of the Great Recession, we do not observe substantial compression in employment propensities as the labor market slowly recovers in the early to mid-2010s. We explore in more detail the relationship between labor market tightness and employment compression later in the paper after adapting our methodology below to track employment dispersion over longer time horizons.

Figure 9. Employment by predicted quintile, 2003–15

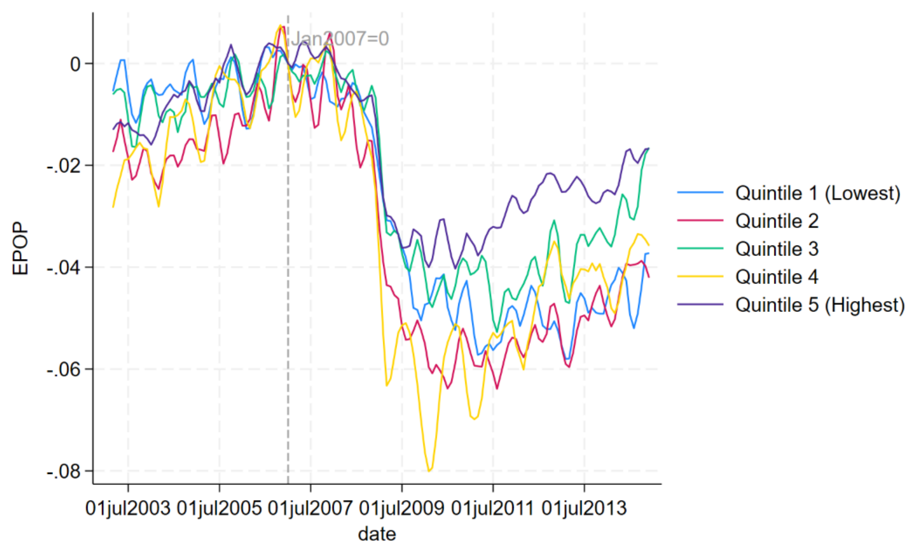

Notes: Figure presents data from the January 2003–December 2014 Current Population Surveys. Sample is restricted to civilians ages 25–54. Figure presents a trailing 3-month moving average of the employment-to-population ratio additively indexed at January 2007=0. Quintiles are constructed using a random forest algorithm for the base period of 2005–06 predicting employment using education, age, sex, race, nativity, veteran status, marital status, presence of own children in the household, and geography variables.

In our analysis thus far, we have chosen to fix the base period in each case to a short period prior to the business cycle peak. One reason for doing so is that it makes transparent any concerns about overfitting, i.e., it is apparent by inspection of the figures that there is no discontinuous shift in employment rates around the base periods. However, the choice of base periods is somewhat arbitrary. It also poses an obstacle to simpler application of the approach to a variety of time periods.

In order to expand our approach to an analysis across a longer time horizon, we modify the approach to calculate employment quintiles contemporaneously, rather than in a fixed base period. One advantage of this modification is that the model does not become systematically less predictive for time periods further away from the estimation sample, as could be the case in our baseline methodology. We also shift to a quarterly estimation to reduce computational burden. Figure 10 presents the findings from this methodology for 1982 onward, indexing at January 2000=0. The random forest algorithm is the same as previously described but excludes variables related to veteran status and nativity (which were not available in some early periods). We continue to observe the recent convergence in employment likelihoods with this variation in approach, along with the divergence that came after the Great Recession.

Figure 10. Employment by predicted quintile, 1982–2023

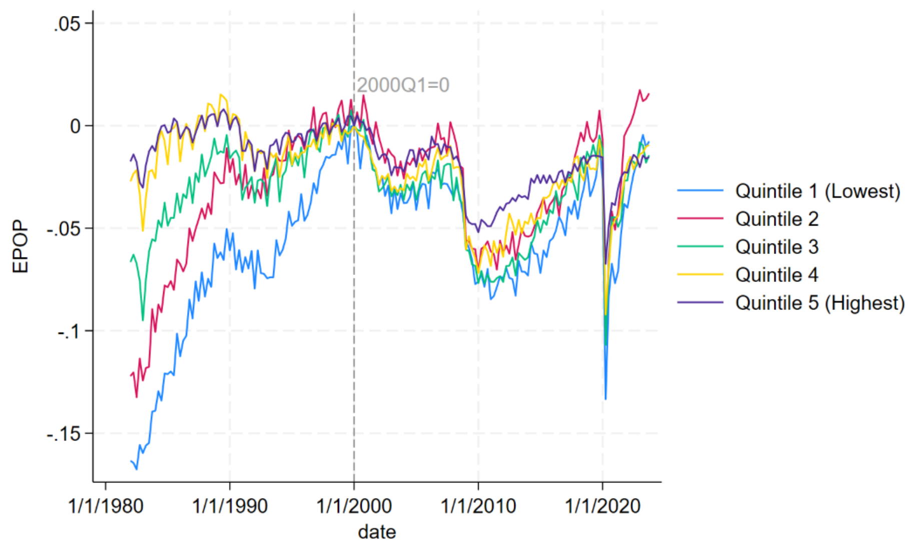

Notes: Figure presents data from the January 1982–September 2023 Current Population Surveys aggregated to the quarterly level. Sample is restricted to civilians ages 25–54. Figure presents the quarterly employment-to-population ratio additively indexed at 2000 Q1=0. Quintiles are constructed using a random forest algorithm for each quarter separately predicting employment using education, age, sex, race, nativity, veteran status, marital status, presence of own children in the household, and geography variables.

From the 1980s to 2000, we observe a secular trend of compression, i.e., a narrowing of employment rate gaps across our predicted quintiles. To the extent our predictor variables capture all facets of inequality we care about, this indicates genuine compression. However, if there are other variables of interest that do not perfectly correlate with those included in our prediction methodology, the graph instead might partially reflect a decline in how well our approach can predict employment, as employment likelihoods converge across the various characteristics we examine. Changes in the dispersion of employment likelihoods across the variables we examine are mechanically associated with changes in the efficacy of the prediction algorithm over time.

To demonstrate this point, Figure 11 shows how the R-squared from an analogous linear probability model evolves over time. We see a sharp decline in how predictive the regression can be from the 1980s to 2000, coinciding with a drop in the explanatory power of sex during this time. In other words, the narrowing of employment gaps between men and women, among other developments, made it more difficult to predict individual-level employment. If sex and its correlates represent an important facet of inequality we care about, the decline in explanatory power overall is a byproduct of convergence in employment by sex and not necessarily a matter of concern. Still, Figure 11 helps depict a potential limitation in using our methodology to assess employment inequality over time—a limitation that becomes more pronounced to the extent that our data exclude some relevant variables for predicting employment.

Figure 11. Explanatory power of LPM model, 1982–2023

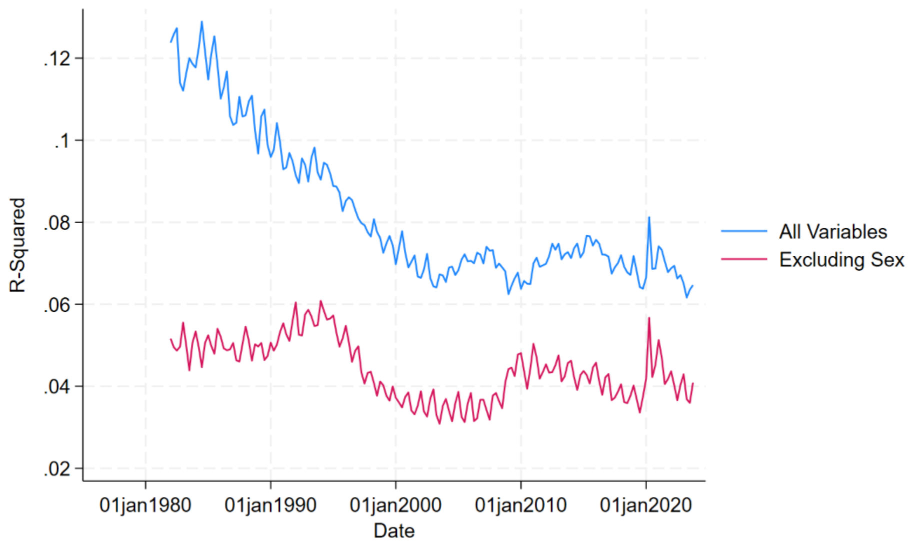

Notes: Figure presents data from the January 1982–September 2023 Current Population Surveys aggregated to the quarterly level. Sample is restricted to civilians ages 25–54. Figure presents the R-squared from a linear probability model that seeks to predict employment using education, age, sex, race, nativity, veteran status, marital status, presence of own children in the household, and geography variables.

Application to labor market tightness

Prior research has examined how labor market tightness affects workers in different demographic groups (e.g., Aaronson et al. 2019; Cajner et al. 2017; Hotchkiss and Moore 2022). We build upon that body of work to examine how labor market tightness matters differentially to those with high and low employment propensities.

In Figure 12, we plot quarterly employment gaps estimated (i.e., the difference in top and bottom quintile employment rates) against the national unemployment rate over 1982–2023.[6] We use a separate color for each business cycle, which helps visually adjust for secular trends in the explanatory power of observable factors concerns (depicted in Figure 11). Because we re-calculate the mapping between worker characteristics and employment rates in each period, we do not need to account for the model becoming less predictive for time periods further away from a fixed estimation sample, as would be the case with our baseline methodology.

[6] Thanks to Joshua Montes for this suggestion.

For most of the business cycles, we see that tighter labor markets are associated with less dispersion in our measure, in line with existing research on the relationship between labor market tightness and outcomes for different demographic groups (Aaronson et al. 2019; Cajner et al. 2017; Hotchkiss and Moore 2022). When labor market conditions are poor, the lowest predicted quintile experiences relatively worse employment outcomes, and when labor market conditions are tight, it experiences relatively better employment outcomes. The relationship between labor market tightness and employment dispersion appears weaker when labor markets are very slack. For example, we observe a limited relationship between the unemployment rate and employment compression when the unemployment rate has exceeded 6 percent since 1990 (or 8 percent in our overall time series). This limited relationship when labor market conditions are fairly slack is in line with the limited compression observed in the early recovery period of the Great Recession in Figure 8.

Figure 12. Employment dispersion and national unemployment rate, 1982–2023

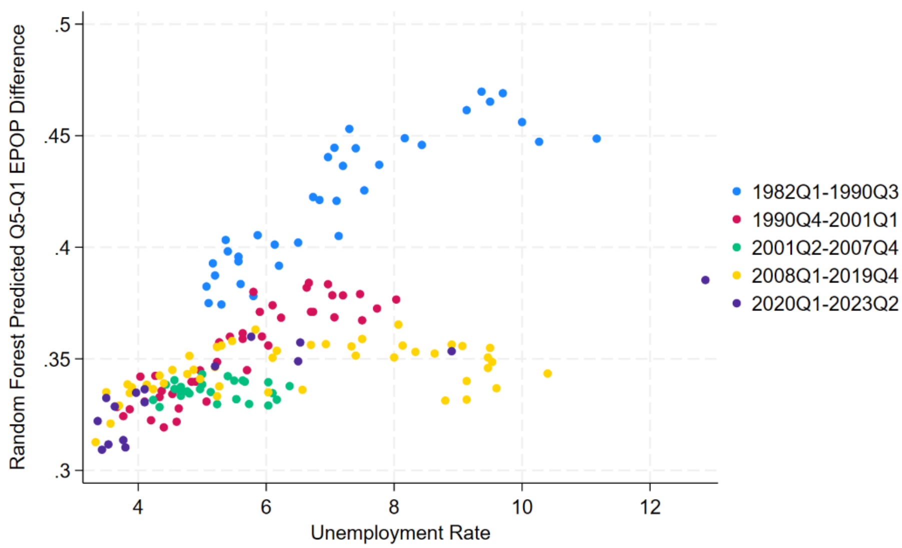

Notes: Figure presents data from the January 1982–December 2023 Current Population Surveys aggregated to the quarterly level. Sample is restricted to civilians ages 25–54. Figure presents the gap between the fifth and first quintile of quarterly employment-to-population ratio. Quintiles are constructed using a random forest algorithm for each quarter separately predicting employment using education, age, sex, race, nativity, veteran status, marital status, presence of own children in the household, and geography variables.

Conclusion

We present a new approach to understand the evolution of employment dispersion in this paper. By predicting employment using a comprehensive set of demographic and geographic variables, constructing quintiles of employment likelihoods based on those predictions, and calculating actual employment propensities for each of those quintiles, we are able to trace out how employment probabilities converge or diverge over time across a distribution. Our approach allows us to examine how dispersion in the likelihood of employment shifted in response to the COVID-19 pandemic and its aftermath, and we similarly apply our methodology to examine trends in employment dispersion during earlier business cycles. We find that individuals with the lowest predicted likelihoods of working experience the largest gains in employment during tight labor markets. The methodology introduced in our paper helps build a framework for future work that seeks to understand the determinants of employment inequality and how that relates to labor income inequality. Future research that could build upon our work could examine the role of labor force participation, study the evolution of voluntary vs. involuntary nonemployment, or examine the role of specific mechanisms in reducing employment inequality.

References

Aaronson, Stephanie, Mary C. Daly, William Wascher, and David W. Wilcox. 2019. "Okun Revisited: Who Benefits Most from a Strong Economy?" *Brookings Papers on Economic Activity* 2019 Spring: 333–404.

Autor, David, Arindrajit Dube, and Annie McGrew. 2023. "The Unexpected Compression: Competition at Work in the Low Wage Labor Market." NBER Working Paper 31010.

Cajner, Tomaz, Tyler Radler, David Ratner, and Ivan Vidangos. 2017. "Racial Gaps in Labor Market Outcomes in the Last Four Decades and over the Business Cycle." Finance and Economics Discussion Series 2017-071.

Hotchkiss, Julie L. and Robert E. Moore. 2022. "Some Like it Hot: Assessing Longer-Term Labor Market Benefits from a High-Pressure Economy." *International Journal of Central Banking* 18(2): 193–243.