The Chicago Fed Labor Market Indicators:

Bridging the Gap with Alternative Labor Data

Scott A. Brave, Benjamin Henken,
Ezra Karger, and Aryan Safi

April 2026

WP 2026-09

https://doi.org/10.21033/wp-2026-09

Federal Reserve Bank *of* Chicago

The Chicago Fed Labor Market Indicators Bridging the Gap with Alternative Labor Data

Scott A. Brave, Benjamin Henken,
Ezra Karger, Aryan Safi

Federal Reserve Bank of Chicago

April 2026

Abstract

We present the Chicago Fed Labor Market Indicators (LMI): a twice-monthly release that includes the job-finding rate, the job-separation rate, and a forecast for the U.S. Bureau of Labor Statistics (BLS) unemployment rate. To overcome limitations in data availability, the LMI uses partial least squares to combine CPS-based finding and separation rates with higher-frequency alternative and traditional labor market data series—such as unemployment insurance claims, Google Trends searches, online job postings, and survey-based indicators. Our resulting flow-consistent unemployment rate (FCR) correlates strongly with the BLS unemployment rate, and can be used to characterize the current "low-hire, low-fire" nature of the U.S. labor market. We use a Bayesian linear regression centered on a no-change prior to translate changes in our FCR into a real-time forecast for the next BLS unemployment rate reading. In backtesting spanning 2018–2026, our unemployment rate forecast improves on both a random-walk benchmark and the Bloomberg consensus forecast, with the largest accuracy gains during the Covid-19 pandemic when the unemployment rate rose rapidly in a way that was well-captured by high-frequency labor market data.

Keywords: unemployment, job separations, job finding, nowcast, big data, alternative data
JEL: C1, C53, E24, E37, J64

Corresponding author: Benjamin Henken, Federal Reserve Bank of Chicago, Economic Research, 230 S. LaSalle St., Chicago, IL 60604, benjamin.henken@chi.frb.org

1 Introduction

The monthly unemployment rate reported by the U.S. Bureau of Labor Statistics (BLS) is a benchmark indicator for assessing labor market slack and, by extension, the effects of macroeconomic policy. Deriving from the *Current Population Survey* (CPS) (or household survey as it is often referred to by economists and other labor market commentators), it is defined over a specific reference week of the month and released with a lag by the BLS in its monthly *Employment Situation* report. When labor market conditions shift quickly, or when short-lived disruptions happen to coincide with the reference week, policymakers and market participants benefit from incorporating higher-frequency signals into their assessment of labor market conditions. Against this backdrop, we present the Chicago Fed Labor Market Indicators (LMI): a new twice-monthly data release from the Federal Reserve Bank of Chicago that provides an early read on the BLS unemployment rate and the inflows and outflows of workers in the labor market that underlie its calculation.[1]

Attaching prominence to accurately measuring these inflows and outflows of workers is a central insight from modern labor economics. Movements in the unemployment rate reflect the net effect of gross flows into and out of unemployment. The "bathtub model" – as popularized by Şahin and Patterson [2012] – uses a flows-based decomposition to illustrate why unemployment can keep changing even after job flows begin to normalize. This framework emphasizes the importance of the flow-consistent (or steady-state) unemployment rate implied by current inflow and outflow rates to unemployment. Importantly for forecasting (and nowcasting), such a flows-based approach can be operationalized in real time. For example, Barnichon and Nekarda [2012], Şahin et al. [2021], and Brave et al. [2025b] show that statistical models built around these flows can yield substantial gains in real-time forecasting performance relative to widely used benchmarks. These contributions motivate our focus on job-finding and job-separation rates – as well as the flow-consistent unemployment rate implied by them – for understanding movements in the BLS unemployment rate in real time.

[1]Visit https://chicagofed.org/lmi for the latest Chicago Fed LMI release and further details.

The post-2020 labor market environment in the U.S. highlighted the need for timely measures of labor market slack. Nowhere is this clearer than in today's greatly expanded menu of high-frequency labor market indicators from the private sector. Researchers now benefit from a wealth of data that previously did not exist, much of which has matured to the point of having an established track record. For example, internet search intensity has been shown to improve near-term forecasts for initial unemployment insurance claims (Aaronson et al. [2022]). Online job postings have also proven valuable for tracking labor demand at a high frequency (Forsythe et al. [2020]). However, many of these newly developed *alternative* data are drawn from convenience samples, raising concerns about representativeness that can dominate statistical precision even in very large datasets. The practical challenge we face in creating a high-quality set of real-time labor market indicators is, therefore, to combine the speed of alternative data with careful benchmarking to related official statistics of interest.

In this paper, we summarize how we navigate this trade-off with the Chicago Fed LMI by combining real-time private and public sector labor market data with official statistics in order to estimate (i) a layoffs and other separations rate, (ii) a hiring rate for unemployed workers, and (iii) a real-time forecast of the monthly BLS unemployment rate. The design begins by targeting job flows: Using the Shimer [2005] and Shimer [2012] methodology, we construct job-finding and job-separation rates. We then use partial least squares to project these CPS-derived rates onto a set of higher-frequency traditional and alternative indicators (including, among others, unemployment insurance claims, Google Trends search data, job-posting-based opening rates, and survey-based measures of job search and labor market perceptions) and translate the projections into a flow-consistent unemployment rate. Finally, we relate changes in our flow-consistent unemployment rate to changes in the official unemployment rate to produce a real-time forecast ahead of each release from the BLS.

Our framework yields two complementary results. First, by providing real-time estimates of job-finding and job-separation rates aligned to the BLS reference week, our indicators support a structural decomposition of movements in the unemployment rate: an important output during times when similarly-sized changes in unemployment could reflect very different underlying labor market conditions (e.g., falling hiring versus rising layoffs). Second, in a pseudo-real-time evaluation over 2018–2026, our real-time unemployment rate forecast outperforms common benchmarks, including a random-walk (no-change) forecast and the Bloomberg consensus forecast, with particularly large gains when unemployment is changing rapidly. Our results underscore a broader point for the alternative data literature: even when individual high-frequency series are imperfect, their systematic combination when anchored to official statistics can deliver usable, timely signals of key policy-relevant outcomes.

The rest of this paper is organized as follows. Section 2 describes the traditional and alternative data series used in the Chicago Fed LMI and formalizes the trade-off between speed and coverage that motivates our approach. Section 3 then summarizes the output of our partial least squares model used to estimate real-time job-finding, job-separation, and flow-consistent unemployment rates and the Bayesian linear regression that is used to translate the latter into a forecast for the BLS unemployment rate. Section 4 presents two case studies: a decomposition of recent movements in our estimate of the flow-consistent unemployment rate into various sources and an assessment of the LMI's out-of-sample forecasting performance for the BLS unemployment rate. Finally, section 5 concludes.

2 Motivation

In this section, we provide some further background and motivation for the mix of data series used in the Chicago Fed Labor Market Indicators (LMI). Accounting for the differential strengths and weaknesses of alternative labor market data sources and official statistics is not an easy task. To highlight some of their key differences, it is useful to consider the construction of the unemployment rate reported by the U.S. Bureau of Labor Statistics (BLS) monthly in its *Employment Situation* report. The unemployment rate captures the number of people without a job who are searching for one as a share of the civilian labor force during a particular (reference) week each month. As such, it provides a timely snapshot of the labor market used as a reference point by both policymakers and market participants.

Like many official statistics, the unemployment rate is released with a lag – typically three to four weeks after the survey reference week. Although this reporting lag is considerably less than many other official statistics, when paired with the focus on a particular reference week, it can at times pose a concern. For instance, the focus on a single week per month had trouble keeping up with fast-moving public health developments during the Covid-19 pandemic that distorted unemployment statistics. Alternative data – often available at higher frequencies with little to no lag – proved invaluable during the pandemic period for this reason (Aaronson et al. [2020]). However, unlike official statistics, alternative data are rarely sourced from nationally representative samples. Instead, they typically derive from convenience samples, so-called because they are collected in the course of a firm's regular business. This fact introduces a trade-off between speed and coverage when using alternative data to capture labor market developments in the aggregate.

Despite these limitations, at the Chicago Fed we have found that alternative data can indeed be helpful in forecasting the unemployment rate in real time. The reason is simple: The unemployment rate summarizes the net effect of job finding and job separations, and alternative data are well-suited to capture these job flows. That said, while plentiful at this point, some alternative data correlate more strongly with aggregate job flows than others. We have identified the following measures as being particularly useful:

- Google: A seasonally adjusted Google Trends weekly index for the unemployment topic based on a daily sample of internet searches for the U.S.

- Indeed, Lightcast, and ADP: Seasonally adjusted job openings rates calculated from Indeed and Lightcast's daily measures of job postings divided by ADP's monthly measure of private payroll employment

- Morning Consult: Seasonally adjusted weekly indexes tracking self-reported unemployment, lost pay, and job search activity from online surveys of U.S. adults

Part of the verification process we used to arrive at this list of alternative data series was to confirm that they tend to be highly correlated with other, more traditional labor-market measures used to track job flows in real time. These traditional measures include:

- U.S. Department of Labor: Seasonally adjusted initial unemployment insurance (UI) claims (expressed as a share of covered employment) and the insured unemployment rate for the U.S.
- Conference Board: The labor market differential index from the Consumer Confidence Survey (jobs plentiful versus jobs hard to get)
- Federal Reserve Bank of New York: A labor market differential calculated as respondents' mean probability of finding a job in the next three months, less their mean probability of losing a job over the next year (both values from the previous month)
- Bloomberg: The consensus forecast for the yet-to-be-released unemployment rate as shown on Bloomberg's economic calendar (sourced from Investing.com)
- U.S. Bureau of Labor Statistics: The previous month's hiring and layoff rates from the Job Openings and Labor Turnover Survey (JOLTS)

To take advantage of differences in release timing as well as the differential strengths and weaknesses of these 13 measures, the Chicago Fed LMI combines them into three summary measures – each benchmarked to similarly constructed official statistics from the BLS. Published twice a month, these measures include: a job-finding rate, or outflow rate (the Chicago Fed Hiring Rate for Unemployed Workers); a job-separations rate, or inflow rate (the Chicago Fed Layoffs and Other Separations Rate); and an unemployment rate forecast (the Chicago Fed Real-Time Unemployment Rate Forecast). This particular decomposition of inflows and outflows to the unemployment rate is motivated by the "bathtub model" of unemployment (Şahin and Patterson [2012]). As flows into unemployment (job separations) rise more than those out of it (job finding), the unemployment rate rises – analogous to the water level rising in a bathtub with a barely open drain and a fully open faucet.

Figure 1 plots these 13 data series since July 2008, with each scaled using a log transformation. The first thing to note about the data in this figure is that not every measure is reported at the same frequency: Some are reported monthly (e.g., the Conference Board index, the Bloomberg consensus forecast, and the ADP employment data), while others are weekly (e.g., UI claims, the Google Trends index, Lightcast job postings data, Indeed job postings data, and the Morning Consult indexes). To relate them to each other, we need to be mindful of timing differences. We do so by focusing on only the values of these series that are observed (for weekly data) or known (for monthly data) by the reference week for the unemployment rate (typically the week containing the 12th of each month).

The Chicago Fed LMI are published twice monthly – first in an Advance release about 7 to 10 days ahead of the BLS employment report, and then in a Final release about 1 to 2 days ahead of the same report. Higher frequency data are typically reflected in the Advance release, with the traditional data more often reflected in the Final release. JOLTS is reported with a longer lag than the others. When missing, we estimate it from separate Bayesian linear regressions. We estimate the JOLTS hiring rate using the Revelio Labs hiring rate; the National Federation for Independent Business jobs hard-to-fill index, average change in employment per firm, and net percent of businesses with an increase in employment over the last three months; the Institute for Supply Management manufacturing and services employment indexes; and one lagged value of the JOLTS hiring rate. We estimate the JOLTS layoffs rate using the Revelio Labs attrition rate; announced job cuts from Challenger, Gray, & Christmas; continuing UI claims from the U.S. Department of Labor; the National Federation for Independent Business average change in employment per firm and net percent of businesses with an increase in employment over the last three months; the Morning Consult pay loss rate; and three lagged values of the JOLTS layoff rate.

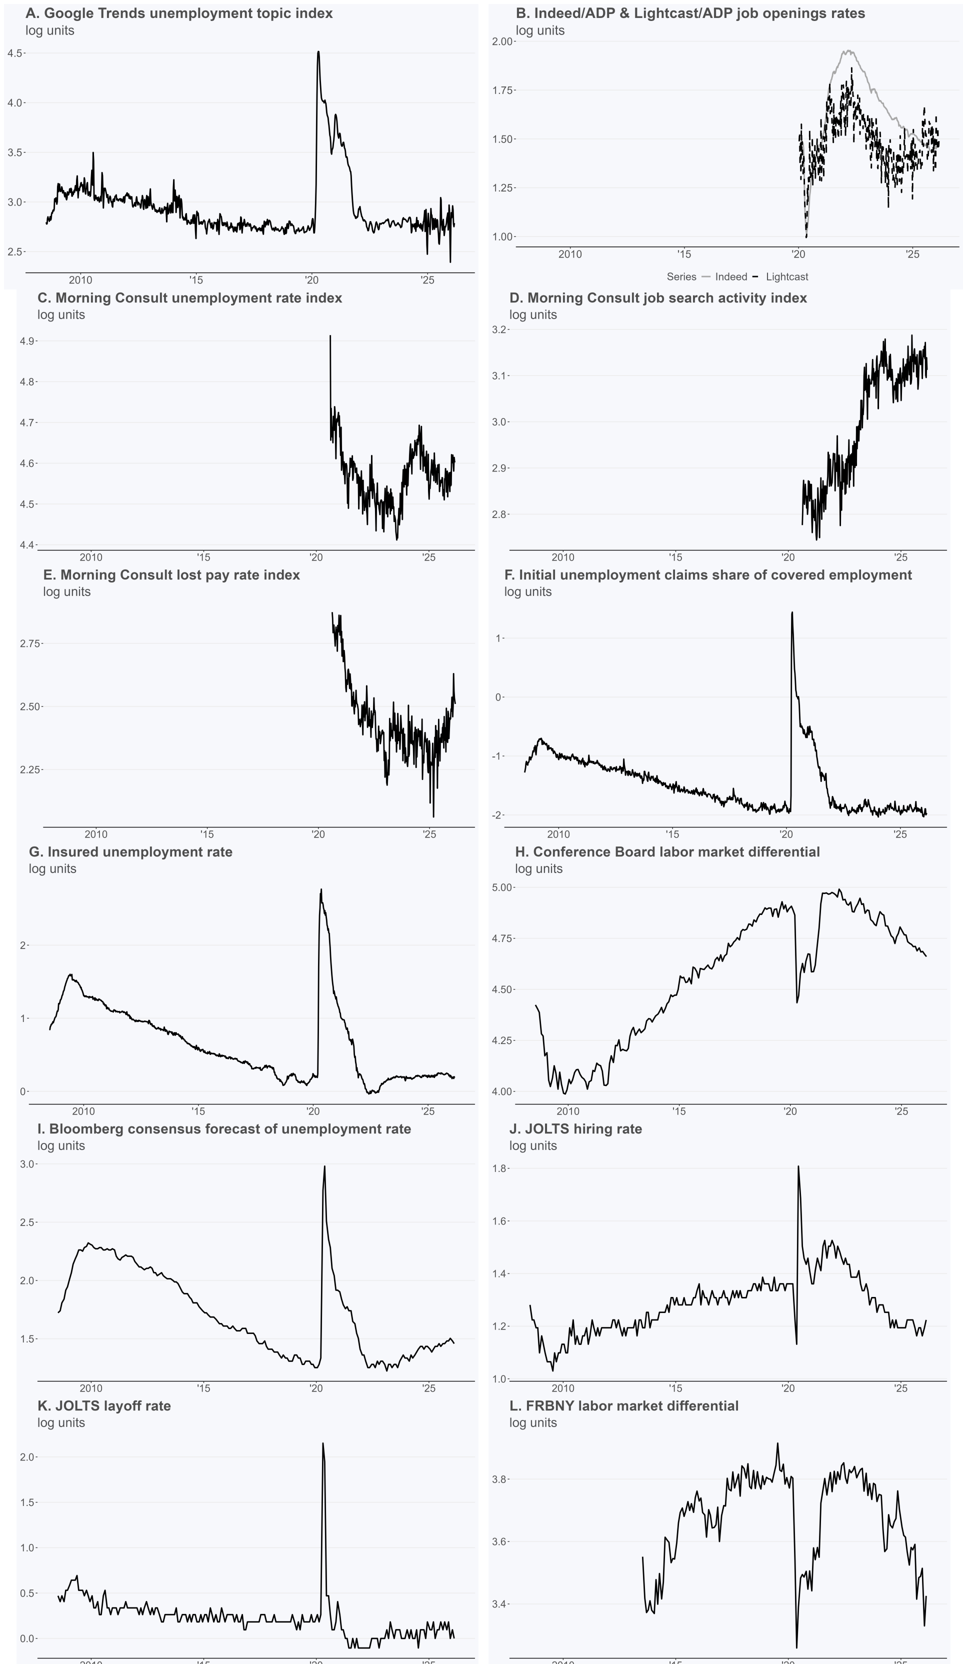

Figure 1: Alternative and Traditional Labor Market Indicators Used to Predict Job Flows

The use of the bathtub model as an organizing framework to relate our 13 alternative and traditional labor market data series back to official statistics for aggregate job flows – and in turn to the unemployment rate published by the BLS – is pivotal to our methodology.[2] It squarely puts the Chicago Fed LMI in the camp of viewing alternative data as useful complements to (and not substitutes for) official statistics. The next section describes how this process works, and the later sections describe what we learn from it. In doing so, we are following in the footsteps of many others who have tried to relate alternative data to official statistics, some with more success than others. We refer the reader to the examples found in Brave et al. [2022] for similar applications to ours, as well as for a more detailed discussion of the strengths and weaknesses of alternative data sources.

3 Constructing the Chicago Fed LMI

The Chicago Fed LMI are benchmarked to flows from the *Current Population Survey* (CPS) – a job-finding rate (unemployment rate outflows) and a job-separation rate (unemployment rate inflows). The job-finding rate measures the prevalence in the labor market of unemployed individuals transitioning into employment or out of the labor force, while the job-separations rate measures the prevalence of layoffs, discharges, and quits leading to nonemployment relative to the number of all employed persons. We calculate both CPS-derived measures using the methodologies of Shimer [2005] and Shimer [2012].[3] An advantage of this approach is that the two job flow rates can then be combined to create a flow-consistent unemployment rate (FCR) representing a steady-state unemployment rate if both the job-finding and job-separation rates stabilized at their current pace.

Others have already demonstrated the value of using the FCR to predict the unemployment rate (e.g., Barnichon and Nekarda [2012], Şahin et al. [2021], and Brave et al. [2025b]). Our contribution with the LMI is to instead provide a novel way to estimate both job flow

[2] See part I of the Technical Appendix for a complete data dictionary.

[3] See part II of the Technical Appendix for details on the calculation of the CPS-derived measures.

rates, the FCR, *and* the unemployment rate in real time. To blend the information in our alternative data series with the job-finding and job-separation rates derived from the CPS, we rely on a statistical method called partial least squares (PLS). PLS is designed to produce targeted linear combinations of a set of predictors (call them X) that maximize their covariance with a set of outcomes (call them Y). With the LMI, we include in Y the job-separation and job-finding rates derived from the CPS, and in X the set of traditional and alternative labor-market data shown in Figure 1. We then construct the model's implied job-separation, job-finding, and flow-consistent unemployment rates, taking advantage of the fact that many of the series in X are made available by their primary sources before those included in Y.[4]

Figure 2 reports the cumulative variance in X and Y explained by the two-factor PLS model underlying the Chicago Fed LMI. The blend of alternative and traditional data included in X both share a very strong common component and explain a considerable portion of the total variance in our sample of alternative and traditional labor market series. With two factors, we can explain about 90 percent of the total variation in X and about 70 percent of the total variation in Y. These results confirm the usefulness of combining alternative and traditional data series and using the bathtub model to relate them to aggregate job flows. In principle, one can further increase the variation explained in X by adding additional factors (common components) to the model. However, we have found that doing so adds very little value in explaining Y. With two job flow rates as our targets for PLS, a two-factor model provides just enough flexibility to fit each of them quite well without overfitting.

[4]See part II of the Technical Appendix for further details on the PLS model. PLS used in this context acts as a form of shrinkage estimator through targeted data-dimension reduction. It is also possible to incorporate sparsity in the PLS estimation procedure, similar to how estimators like the commonly used LASSO procedure go about doing so. We leave this option to future research.

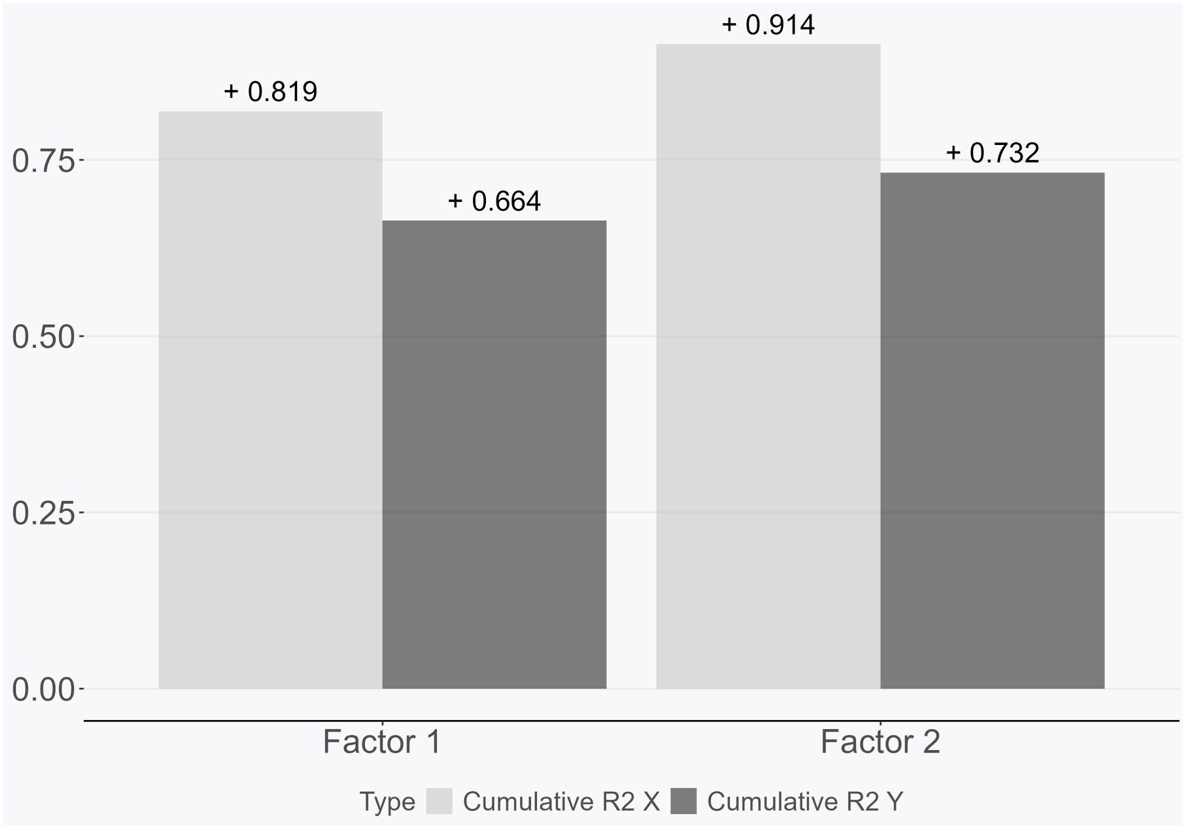

Figure 2: PLS Model Explained Variance of X and Y

In Figure 3, we go one step further and examine the factor loadings for each data series in X. The sign and magnitude of the loadings can tell us which information in X is helping to *drive* our estimates of the job flows in Y. The factor loadings represent the relative strength of the relationship between an observed variable in X and the corresponding latent factor used by PLS to explain Y. The greater the absolute value of the loading, the stronger its relationship to the underlying component. In Figure 3, we plot loadings for the first factor along the x-axis and loadings for the second factor along the y-axis. Overall, the diffuse spread of the indicators across the plot area suggests that each provides unique information about the two job flow rates. However, looking at the relative size of the loadings, some indicators are clearly more dominant than others.

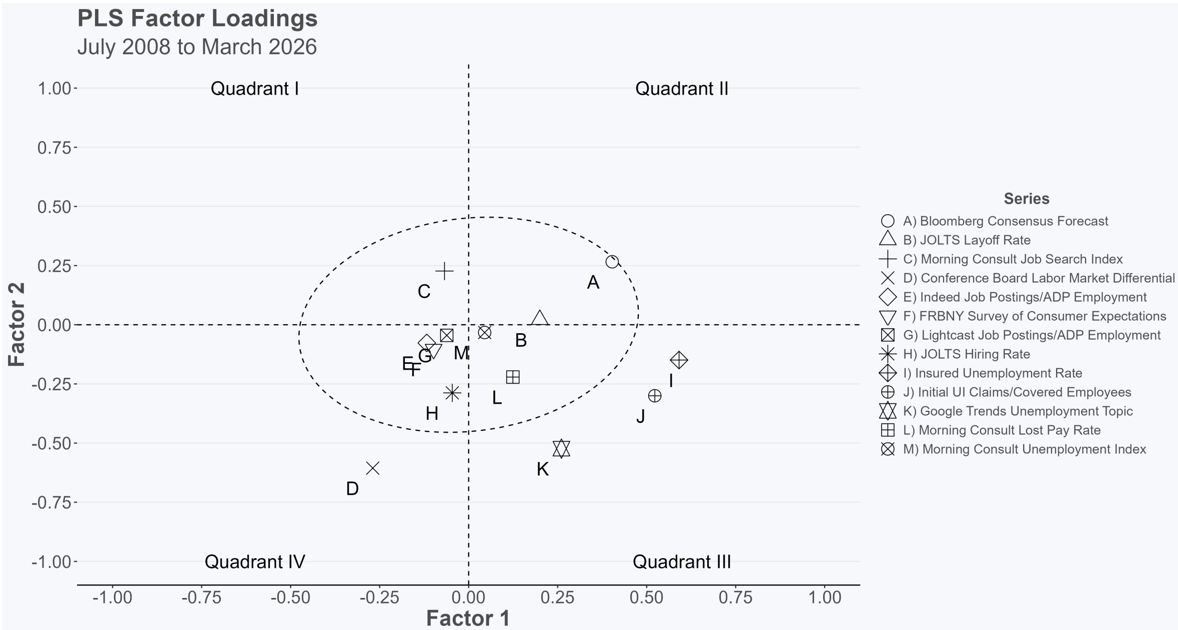

Figure 3: PLS First and Second Factor Loadings

The dotted black oval drawn around the figure's origin distinguishes the four series with the largest loadings across both factors (those falling outside this region): weekly initial unemployment insurance claims, the insured unemployment rate, the Google Trends unemployment topic, and the Conference Board labor market differential. Together, these four indicators account for roughly 80 percent of the model-explained variance in X and are among the longest-running time series in our sample. As shown in Figure 1, most of the remaining nine series in X have histories that do not begin until late 2020 or early 2021.

The next thing to notice about the figure is how the variables are distributed about the two axes. The indicators sort fairly cleanly (if imperfectly) by how well they capture patterns in job finding and job separations. For example, Quadrant III includes many of the series most closely associated with job separations (e.g., weekly initial claims, the insured unemployment rate, the Google Trends unemployment topic, and the Morning Consult lost pay rate), while Quadrant IV includes many of the series most closely associated with job finding (e.g., Indeed and Lightcast job postings divided by ADP (a pseudo job vacancy rate), the Conference Board and FRBNY labor market differentials, and the JOLTS hiring rate).

Finally, in Figure 4, we show the PLS model fits reported in the Chicago Fed LMI for the job-finding rate, the job-separations rate, and the FCR alongside their CPS-derived counterparts. Our PLS estimates track closely with the CPS-derived rates (see panels A, B, and C of Figure 4). The Chicago Fed Hiring Rate for Unemployed Workers is the model fit for the job-finding rate, and the Chicago Fed Layoffs and Other Separations Rate is the model fit for the job-separation rate. The FCR is constructed from these two job flow rates as the ratio of the job-separation rate to the sum of the job-finding and job-separation rates. In section 4.1, we take a closer look at what these two rates and their underlying indicators tell us about the current "low-hire, low-fire" labor market by decomposing changes in the FCR over various sample periods into contributions from its constituent parts.

Notice also that our model estimates in Figure 4 are all considerably smoother than their CPS-derived variants. This is a byproduct of blending the CPS data with alternative and traditional measures of job flows and a feature of using PLS to do so. This smoothing feature motivates our use of the FCR to forecast the BLS unemployment rate. To demonstrate, Figure 4, panel D, plots the time series of the BLS unemployment rate against our FCR estimate. The BLS unemployment rate exhibits a high degree of correlation with our FCR, with correlation coefficients of 0.98 in levels and 0.95 in first differences. The high correlation in levels motivates our use of the FCR to understand the drivers of changes in the unemployment rate, while the high correlation in first differences underpins our use of the FCR to forecast monthly changes in the BLS unemployment rate.

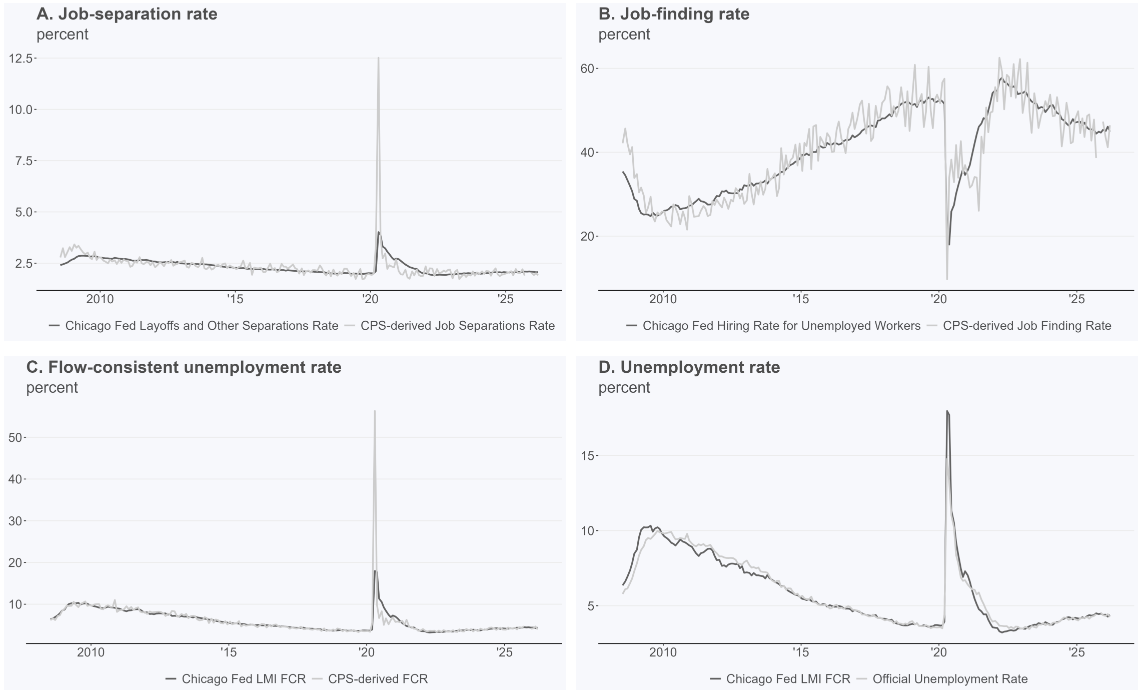

Figure 4: Job Finding, Separations, and Unemployment Rates

The Chicago Fed Real-Time Unemployment Rate Forecast is a prediction for the yet-to-be-released value of the unemployment rate that builds on this close relationship with our FCR. We use Bayesian estimation methods to form this prediction and bring in external (prior, or a priori) information that is not contained in the short sample history of BLS reference weeks since July 2008 used for PLS estimation. In particular, the average monthly change in the BLS unemployment rate over long periods is close to zero, despite the volatility around recessions. We account for this by centering the monthly change in the unemployment rate in our Bayesian linear regression model around a no-change prior.

Our Bayesian linear regression relates the reference-week-to-reference-week changes in both the unemployment rate and our model FCR, with a prior that shrinks the regression coefficients toward zero. In the limit, when our prior beliefs exactly hold, this model would predict no change in the unemployment rate. However, instead of imposing this exact relationship, we estimate the strength of this *a priori* belief balanced against the strength of the observed relationship between these two measures to form an *a posteriori* prediction.[5] In section 4.2, we summarize the out-of-sample forecasting ability of this statistical model.

4 Case Studies

In this section, we take a closer look at the Chicago Fed Labor Market Indicators (LMI) through the lens of two case studies. The first highlights the ability to decompose changes in the unemployment rate into their underlying flows and sources, while the second examines the out-of-sample forecast accuracy of our real-time unemployment rate estimates.

4.1 The "Low-Hire, Low-Fire" Labor Market

Since 2023, the U.S. labor market has experienced movements in the rate of job separations and job finding such that both inflows and outflows of workers have put upward pressure on the unemployment rate. Consistent with this development, the unemployment rate during this period has risen by roughly a percentage point from an all-time low of 3.4 percent in April 2023 to its current low level of 4.3 percent in January 2026. To identify the drivers of the recent increase in the unemployment rate, Figure 5 decomposes the Chicago Fed's model-implied FCR into separate contributions from job finding and job separations over the Jan 2023 – Mar 2026 and Jan 2017 – Mar 2026 periods.

Panel A of the figure demonstrates that almost 80 percent of the increase in the FCR since early 2023 can be traced back to lower job finding (the "low-hire" part). Panel B of the figure is helpful to then see why the remaining 20 percent due to an increase in job separations (the "low-fire" part) remains quite low historically. Although layoffs and other job separations have risen steadily since 2024, the increases have been small and gradual, with only minimal upward pressure on the unemployment rate through March 2026. This pattern is quite distinct from a typical labor market contraction, where job separations are

[5] See part III of the Technical Appendix for further details on the Bayesian linear regression.

often seen as an early signal of sharply rising unemployment (e.g., the 2020 recession).

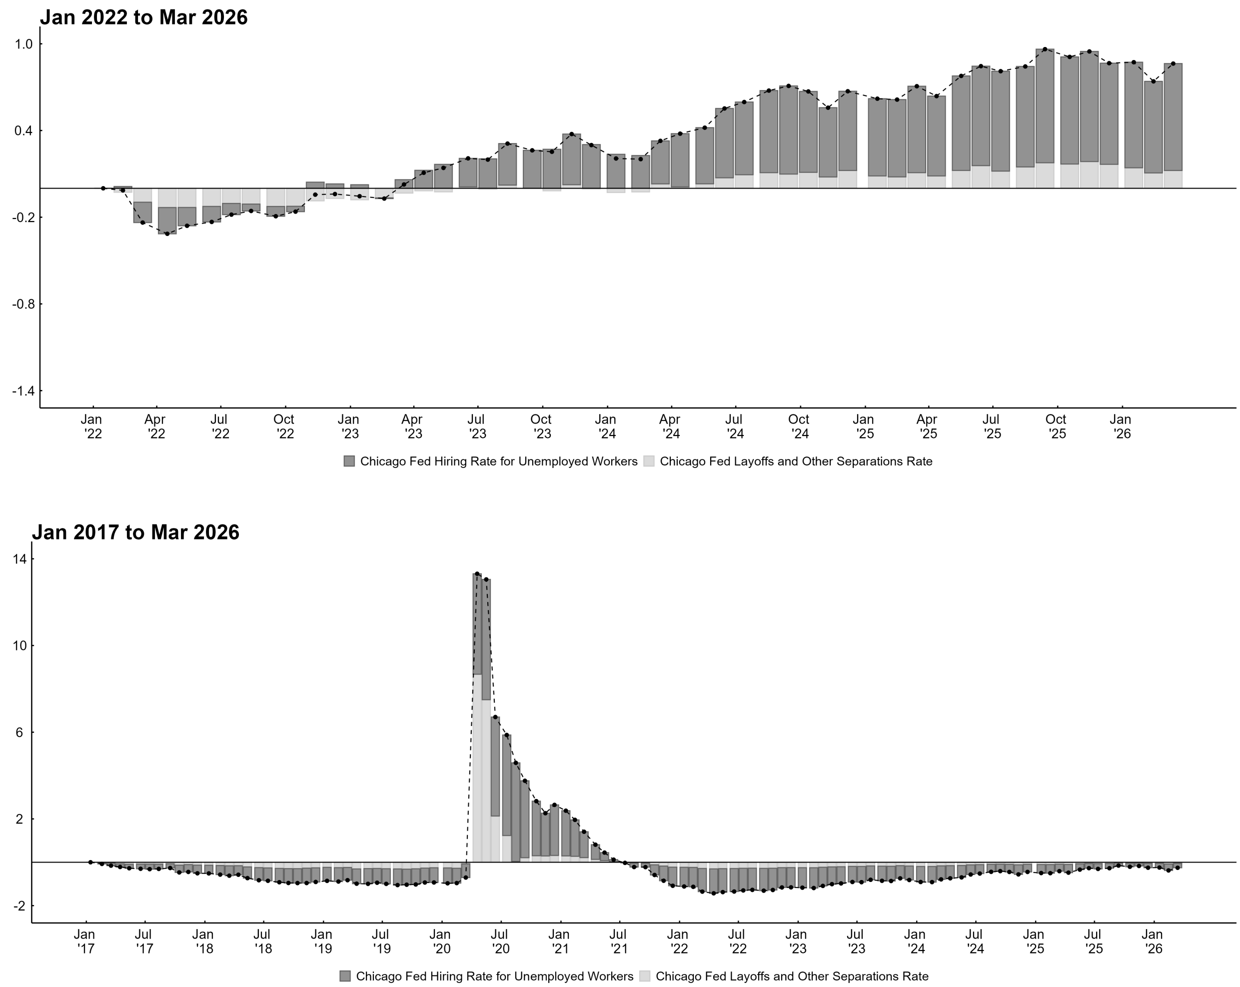

Figure 5: Cumulative Change in the Flow-Consistent Unemployment Rate

To better understand the "low-fire" dynamic, Figure 6 decomposes deviations in the Chicago Fed Layoffs and Other Job Separations Rate from its 2015-2019 mean into contributions from public- and private-sector data sources. This way of looking at the time series makes clear that the low-fire aspect of the labor market in the LMI is largely a function of traditional public sector data sources (particularly the low level of initial UI claims), while much of the increase since the second half of 2024 in job separations is explained instead by alternative private sector data sources. Many of these alternative data series correspond with measures of job openings and availability, suggesting both a reluctance on the part of employers to hire new workers or create new open roles and a recognition of this fact on the part of households in reporting tougher job-finding conditions.

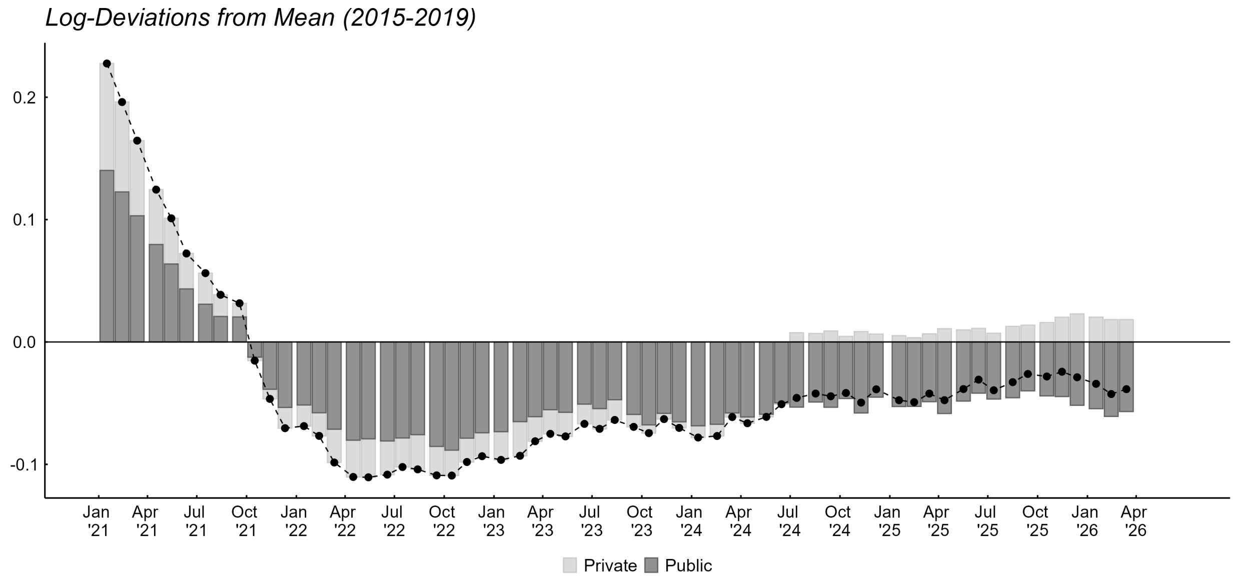

Figure 6: Chicago Fed Layoffs and Other Separations Rate Decomposition

Digging deeper into this "low-hire" dynamic, Figure 7 decomposes deviations in the Chicago Fed Hiring Rate for Unemployed Workers from its 2015-2019 mean into contributions from what we classify as job-availability and job-search oriented data sources based on their estimated impact on the CFLMI model. This way of looking at the time series for job finding makes clear that the low-hire aspect of the labor market is reflected in the CFLMI primarily through fluctuations in the job-availability oriented metrics (e.g., the Bloomberg Consensus Forecast, Morning Consult Job Search Activity Index, Conference Board's Labor Market Differential, Indeed & Lightcast Job Postings, and insured unemployment rate).

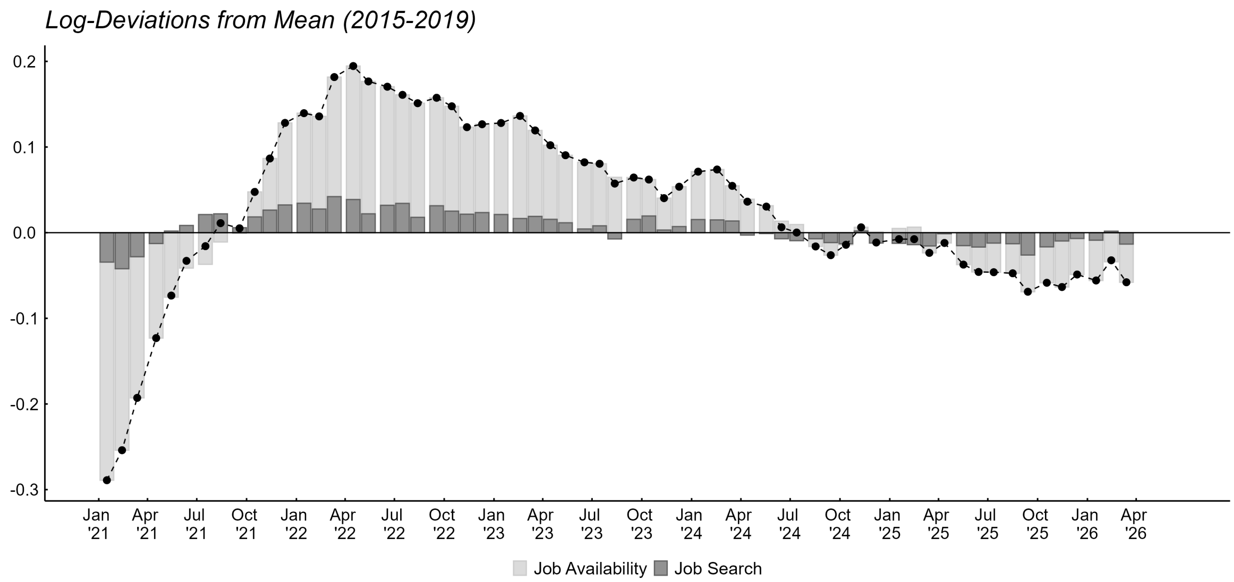

Figure 7: Chicago Fed Hiring Rate for Unemployed Workers Decomposition

While the limited historical availability of many of the private data sources used in the CFLMI prevents us from making a stronger judgment on their merits as leading indicators (hence why the figures start in early 2021), the divergence in signal between real-time public and private sector data sources is still a development worth monitoring. Some private-sector indicators in the CFLMI (e.g., the Conference Board Labor Market Differential) have long track records of capturing turning points in the labor market, and many others performed quite well during the recovery from the 2020 recession. Next, we take a closer look at the historical predictive ability of the CFLMI for the unemployment rate.

4.2 Forecasting the Unemployment Rate

To produce a real-time forecast for the BLS unemployment rate, the Chicago Fed LMI relies not only on the bathtub model as an organizing principle but also on the real-time data flow for the higher-frequency alternative and traditional data to generate a forecast *ahead* of each BLS report. These forecasts are published twice monthly: (1) an Advance release, 7 to 10 days out from the BLS report, which incorporates any data already available for the reference week covered in the upcoming BLS report; and (2) a Final release, 1-2 days out from the BLS report, incorporating any additional data updates from the intervening week.

While the LMI have only existed for a short time, we can still evaluate how its real-time unemployment rate forecasts might have performed historically through a pseudo-real-time forecasting exercise where we iteratively predict out-of-sample unemployment rates based on the information available at each point in time. For this exercise, we put ourselves in the typical position of a Final release of the LMI, using as much data as would likely have been available in real time. We do so here for the period January 2018 through March 2026, using both truncated data as currently exist in the LMI and vintage data releases for any data series found in the St. Louis Fed's Archival FRED (ALFRED) database.[6] To preserve a consistent data flow across months, we exclude from this analysis the fourth quarter of 2025 given the impact of the federal government shutdown on public data sources.

To judge the LMI's forecasting performance, we compute a series of forecast accuracy diagnostics (mean absolute error, MAE, and root mean squared error, RMSE) and compare them with those obtained using Bloomberg's consensus forecast and a random walk forecast embodying our no-change prior. Lower MAE and RMSE correspond to higher forecast accuracy, measured as percentage points of the unemployment rate. To then draw comparisons across models, we run Diebold-Mariano tests and report the statistical significance for percentage improvements in MAE or RMSE for the Chicago Fed LMI forecasts relative to both the Bloomberg consensus and random walk forecasts.

[6]For further details on the pseudo-real-time forecasting exercise, see part IV of the Technical Appendix.

Table 1: Pseudo-Real-Time Out-of-Sample Forecast Accuracy

A. Revised Government Data (Unemployment Rate, UI Claims, & JOLTS)

	Loss Function		*Improvement Consensus (%)*		*Improvement Random Walk (%)*	
Model	**MAE**	**RMSE**	**MAE**	**RMSE**	**MAE**	**RMSE**
Chicago Fed LMI	0.14	0.24	36*	67*	52*	79*
Bloomberg Consensus	0.23	0.72	0	0	24*	36*
Random Walk	0.30	1.13	-32	-57	0	0

B. Unrevised Government Data (CPS, UI Claims, & JOLTS)

	Loss Function		*Improvement Consensus (%)*		*Improvement Random Walk (%)*	
Model	**MAE**	**RMSE**	**MAE**	**RMSE**	**MAE**	**RMSE**
Chicago Fed LMI	0.15	0.31	35*	57*	51*	72*
Bloomberg Consensus	0.23	0.72	0	0	24*	36*
Random Walk	0.30	1.13	-32	-57	0	0

C. Revised Data with subsets of Chicago Fed LMI variables

	Loss Function		*Improvement Consensus (%)*		*Improvement Random Walk (%)*	
Model	**MAE**	**RMSE**	**MAE**	**RMSE**	**MAE**	**RMSE**
Full Model	0.14	0.24	36*	67*	52*	79*
Less Consensus	0.18	0.44	21*	38*	40*	61*
Less JOLTS & UI too	0.27	0.98	-18	-36	11*	13*

*The * markers denote statistical significance at the 95% confidence level.*

Notes: The table reports mean absolute error (MAE) and root mean-squared error (RMSE) values for out-of-sample nowcasts of the unemployment rate over a sample that extends from January 2018 through March 2026 (excluding the fourth quarter of 2025 due to the federal government shutdown). Panel A constructs nowcasts using data with revisions while Panel B uses unrevised (vintage) data available on ALFRED. Panel C repeats the results in Panel A for the full Chicago Fed LMI model and then reports results after dropping certain variables. Forecast accuracy is measured relative to the unemployment rate values as of April 2026. Forecast model comparisons are also reported relative to two benchmark forecasts: the Bloomberg consensus forecast for the same month and a random walk forecast that assumes an unchanged unemployment rate from the previous month. Negative percent improvements correspond to lower forecast accuracy than the benchmark forecast. Results for tests of equal forecast accuracy (Diebold–Mariano one-sided t-tests with HLN size correction and HAC standard errors) are shown as markers.

Table 1, part A, summarizes the out-of-sample performance of the LMI using truncated (and revised as noted in the table) data. Over the full sample period, the LMI statistically significantly outperforms both the random walk and Bloomberg consensus forecasts on average for both forecast accuracy metrics. To help put the magnitudes in Table 1, panel A in perspective, consider the inherent sampling uncertainty in the BLS unemployment rate. A 90% confidence interval for the one-month change in the U-3 civilian unemployment rate in January 2026 was roughly +/- 0.22 percentage points according to the BLS, or about equal in size to the MAE of the Bloomberg consensus forecasts since 2018. In contrast, the MAE of the LMI is slightly lower at around 0.15 percentage points, and the random walk forecast is slightly higher at 0.30 percentage points.

Table 1, part B, then repeats the same out-of-sample forecast performance evaluation using vintage (or unrevised, as noted in the table) data from the CPS, for UI claims, and JOLTS data. Several of these data series get revised once or more after they are first published. By using unrevised data whenever possible, we seek to more faithfully reflect the information set ahead of each historical BLS *Employment Situation* report, minimizing any advantage from data revisions that only became known in hindsight. With the unrevised data as inputs, the LMI's forecast performance only marginally worsens, and the model maintains its edge over the Bloomberg consensus and random walk forecasts.

Finally, in Table 1, part C, we repeat the same exercise with revised data, but now where we intentionally omit certain data series from the LMI to examine model performance in the absence of these variables. Above all, we find from this exercise that the LMI continue to outperform the Bloomberg consensus forecast on average even when we exclude this variable from the model. This suggests that our superior forecast performance is not due to the incorporation of Bloomberg consensus forecasts, but instead reflects additional information captured in the LMI. Forecast performance, however, worsens considerably when we also exclude the two UI claims series (and to a lesser extent the two JOLTS rates) from the model, reinforcing the importance of many of the traditional data series in the LMI. Even so, a model stripped of these five series (Bloomberg consensus forecast, weekly initial claims, insured unemployment rate, JOLTS hiring rate, JOLTS layoffs rate) continues to outperform a random walk forecast on average by 11-13%.

To show how the performances of the LMI and the other two benchmark forecasts have varied over time, Figure 7 plots MAE and RMSE values for all three based on 24-month moving averages of the forecast errors. The LMI matches or outperforms both the Bloomberg consensus and random walk forecasts for all but a few months of the period covered in the figure, with greater accuracy during periods when the unemployment rate was changing rapidly (e.g., 2020). This is true regardless of the data inputs' revision status. The Giacomini-Rossi fluctuation tests shown in Figure 8 using either revised or vintage data to contruct the LMI confirm this finding: the LMI forecast outperforms the Bloomberg consensus and random walk on a statistically significant basis for just under 2 years of the pandemic period, while performing statistically equally as well during more normal economic conditions.

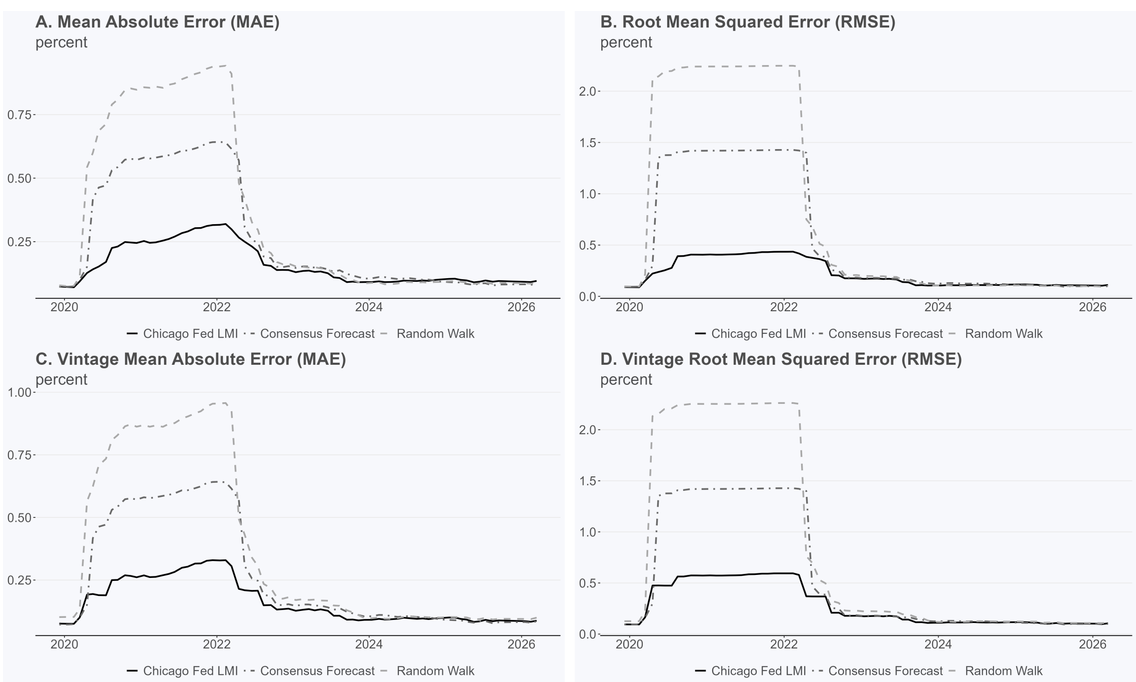

Figure 8: 24-Month Rolling MAE and RMSE

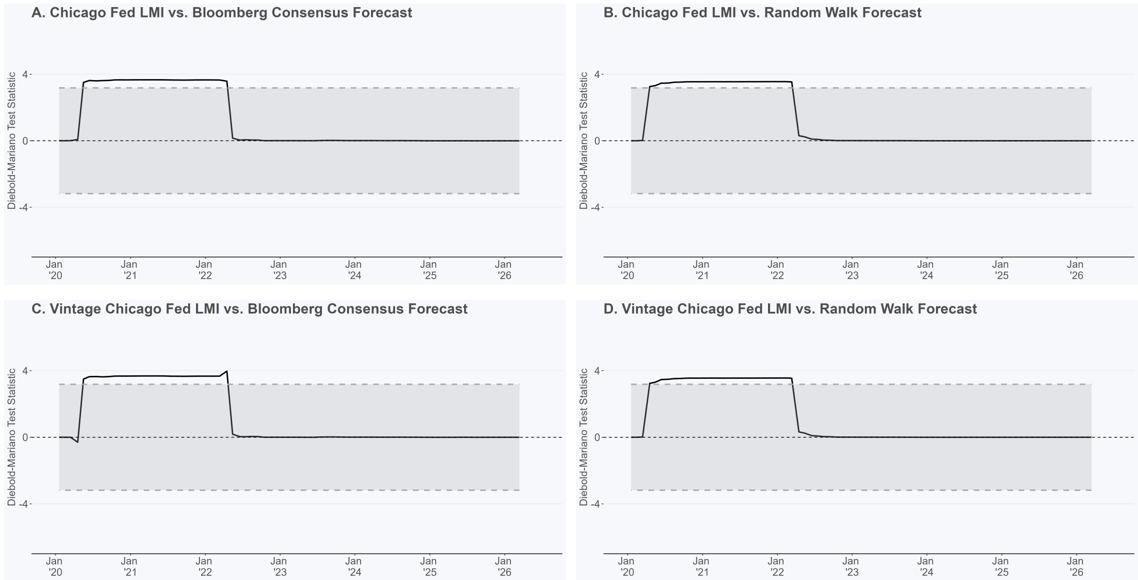

Figure 9: Giacomini-Rossi Fluctuation Tests

We can also assess how well the forecast model captures uncertainty. Calibration measures the frequency at which actual values of the unemployment rate fall within the forecast model's prediction intervals. In a well-calibrated model, actual values of the unemployment rate should fall within a given prediction interval at roughly the same frequency as the uncertainty level they capture (e.g., actual values of the unemployment rate should fall within a 90 percent prediction interval 90 percent of the time). To assess the forecast model's calibration, we use our out-of-sample estimates to generate finite-sample prediction intervals with conformal inference methods designed with calibration in mind.[7]

Figure 9 assesses out-of-sample calibration visually. The dark black line represents coverage values (y-axis) at various prediction intervals (x-axis), while the gray 45-degree line drawn from the origin represents what we would expect to see in a perfectly calibrated model. The forecast model closely tracks the 45-degree line, suggesting it captures forecast uncertainty quite well. The twice-monthly releases of the Chicago Fed LMI exploit this fact and include a probability distribution for various possible values of the next unemployment rate to be reported by the BLS (Brave et al. [2025a]) as well as a probabilistic assessment

[7]See part III of the Technical Appendix for further details on the construction of the prediction intervals.

of its likely direction of change summarizing the balance of risks to the forecast.

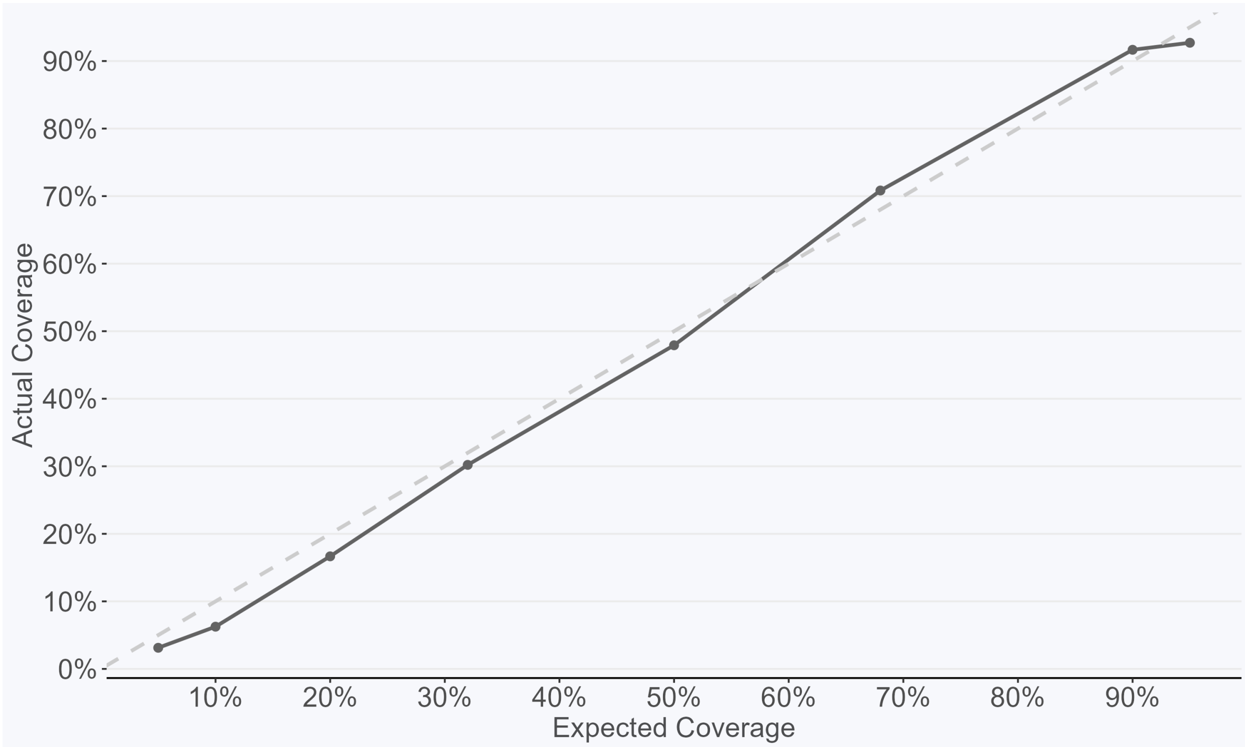

Figure 10: Out-of-Sample Model Calibration of Prediction Intervals

To assess the density forecast accuracy of the CFLMI, we benchmark its performance against a Random Walk forecast using the Continuous Ranked Probability Score (CRPS). The benchmark is generated by simulating forecasts from a normal distribution centered at the previous month's value of the unemployment rate with a standard deviation equal to the previous month's sample standard deviation for the monthly change in the unemployment rate. The resulting comparison in Figure 11 confirms for density forecast accuracy what we found in Table 1 for point forecast accuracy. The CFLMI's forecasts are both well-calibrated and sharp, with a mean CRPS of about 0.11 percentage points and a statistically significant percentage improvement in mean CRPS over the Random Walk model of about 65 percent. The percentage improvement over the Random Walk model in terms of density forecast accuracy is highly robust, with the mean improvement since 2020 still near 60 percent.

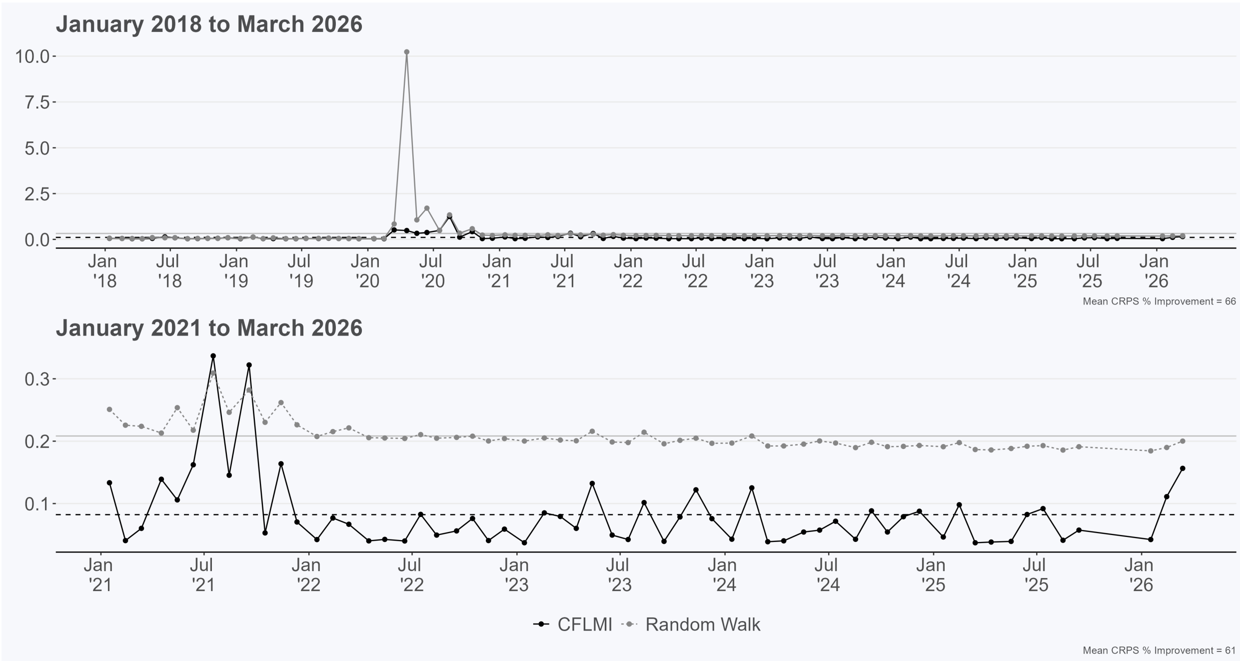

Figure 11: Continuous Ranked Probability Score Comparison

5 Conclusion

In this paper, we presented the Chicago Fed Labor Market Indicators (LMI), a new twice-monthly release from the Federal Reserve Bank of Chicago that combines real-time labor market data with official statistics. Our analysis yielded two complementary results relevant for the growing literature bridging the gap between alternative and traditional data sources. First, the real-time job flow rates in the LMI (i.e., the Chicago Fed Hiring Rate for Unemployed Workers and the Chicago Fed Layoffs and Other Separations Rate) support a structural decomposition of movements in the unemployment rate: a potentially important output for policymakers and market participants seeking to better understand labor market conditions in real time. To demonstrate their value in this regard, we used the Chicago Fed LMI to characterize the origins of the current "low-hire, low-fire" U.S. labor market.

Next, we showed that the Chicago Fed LMI can be highly useful for forecasters during periods of labor market transition. For instance, in a pseudo-real-time evaluation over 2018–

2026, the Chicago Fed Real-Time Unemployment Rate Forecast outperforms on average common benchmarks for nowcasting the unemployment rate, with particularly large gains when the unemployment rate was changing rapidly during the Covid-19 pandemic and the value to policymakers and market participants was exceptionally high. The Chicago Fed unemployment rate forecasts are also well-calibrated and can be used to generate sharp probabilistic assessments for the level and direction of change in the unemployment rate.

Our results underscore a broader point for the alternative data literature: even when individual high-frequency alternative data series are imperfect, their systematic combination when anchored to official statistics can deliver usable, timely signals of key policy-relevant outcomes. While both the quantity and quality of alternative data sources have progressed rapidly in the last few years, it is important to recognize that they remain at best complements (not substitutes) for the gold standard, nationally representative official statistics that policymakers and market participants rely heavily upon to form judgments on the state of labor market conditions. With the Chicago Fed LMI, we seek to have the best of both worlds, combining the speed of alternative data with the coverage of official statistics to reduce the information gap in monthly unemployment rate releases.

References

Daniel Aaronson, Scott A. Brave, R. Andrew Butters, and Michael Fogarty. The stay-at-home labor market: Google searches, unemployment insurance, and public health orders. *Chicago Fed Letter*, (436), 2020. URL https://www.chicagofed.org/-/media/publications/chicago-fed-letter/2020/cfl436-pdf.pdf?sc_lang=en.

Daniel Aaronson, Scott A. Brave, R. Andrew Butters, Michael Fogarty, Daniel W. Sacks, and Boyoung Seo. Forecasting unemployment insurance claims in realtime with google trends. *International Journal of Forecasting*, 38(2):567–581, 2022. ISSN

0169-2070. doi: https://doi.org/10.1016/j.ijforecast.2021.04.001. URL https://www.sciencedirect.com/science/article/pii/S0169207021000649.

Rina Foygel Barber, Emmanuel J. Candès, Aaditya Ramdas, and Ryan J. Tibshirani. Predictive inference with the jackknife+. *The Annals of Statistics*, 49(1):486 – 507, 2021. doi: 10.1214/20-AOS1965. URL https://doi.org/10.1214/20-AOS1965.

Regis Barnichon and Christopher J. Nekarda. The ins and outs of forecasting unemployment: Using labor force flows to forecast the labor market. *Brookings Papers on Economic Activity*, pages 83–131, 2012. URL https://www.brookings.edu/wp-content/uploads/2012/09/2012b_barnichon.pdf.

Scott A. Brave, R. Andrew Butters, and Michael Fogarty. The perils of working with big data, and a small checklist you can use to recognize them. *Business Horizons*, 65(4): 481–492, 2022. ISSN 0007-6813. doi: https://doi.org/10.1016/j.bushor.2021.06.004. URL https://www.sciencedirect.com/science/article/pii/S0007681321001178.

Scott A. Brave, Ben Henken, Katherine Jolley, Ezra Karger, and Aryan Safi. Introducing the chicago fed labor market indicators. *Chicago Fed Insights*, 2025a. URL https://www.chicagofed.org/publications/chicago-fed-insights/2025/introducing-chicago-fed-labor-market-indicators.

Scott A. Brave, Ben Henken, and Ezra Karger. Blending traditional and alternative labor market data with churn. *Chicago Fed Letter*, (506), 2025b. URL https://www.chicagofed.org/publications/chicago-fed-letter/2025/506.

Francis X. Diebold and Roberto S. Mariano. Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263, 1995. URL http://www.jstor.org/stable/1392185.

Vera Z. Eichenauer, Ronald Indergand, Isabel Z. Martínez, and Christoph Sax. Obtaining consistent time series from google trends. *Economic Inquiry*, 60(2):694–705,

2022. doi: https://doi.org/10.1111/ecin.13049. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/ecin.13049.

Eliza Forsythe, Lisa B. Kahn, Fabian Lange, and David Wiczer. Labor demand in the time of covid-19: Evidence from vacancy postings and ui claims. *Journal of Public Economics*, 189:104238, 2020. ISSN 0047-2727. doi: https://doi.org/10.1016/j.jpubeco.2020.104238. URL https://www.sciencedirect.com/science/article/pii/S004727272030102X.

Raffaella Giacomini and Barbara Rossi. Forecast comparisons in unstable environments. *Journal of Applied Econometrics*, 25(4):595–620, 2010. URL http://www.jstor.org/stable/40607048.

David Harvey, Stephen Leybourne, and Paul Newbold. Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13(2):281–291, 1997. doi: https://doi.org/10.1016/S0169-2070(96)00719-4. URL https://www.sciencedirect.com/science/article/pii/S0169207096007194.

Nicholas M. Kiefer and Timothy J. Vogelsang. A new asymptotic theory for heteroskedasticity-autocorrelation robust tests. *Econometric Theory*, 21(6): 1130–1164, December 2005. URL https://ideas.repec.org/a/cup/etheor/v21y2005i06p1130-1164_05.html.

Robert Shimer. The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1):25–49, 2005. URL https://www.aeaweb.org/articles?id=10.1257/0002828053828572.

Robert Shimer. Reassessing the ins and outs of unemployment. *Review of Economic Dynamics*, 12(2):127–148, 2012. URL https://www.sciencedirect.com/science/article/pii/S1094202512000063.

Ayşegül Şahin and Christina Patterson. The bathtub model of unemployment: The importance of labor market flow dynamics. *Liberty Street Economics*,

2012. URL https://libertystreeteconomics.newyorkfed.org/2012/03/the-bathtub-model-of-unemployment-the-importance-of-labor-market-flow-dynamics/.

Ayşegül Şahin, Murat Tasci, and Jin Yan. Unemployment in the time of covid-19: A flow-based approach to real-time unemployment projections. Working Paper 28445, National Bureau of Economic Research, February 2021. URL http://www.nber.org/papers/w28445.

6 Technical Appendix

This technical appendix is divided into four parts. Part I contains a more detailed description of the data used in the Chicago Fed Labor Market Indicators. Part II describes how we use partial least squares to combine official statistics with alternative and more traditional labor market indicators to track job flows – and, by connection, the unemployment rate –in real time. Part III shows how we translate our estimated real-time job flow rates into a real-time forecast for the unemployment rate. Finally, part IV provides additional details on the pseudo-real-time out-of-sample forecast evaluation exercise described in the main text.

6.1 Data Dictionary

The Chicago Fed Labor Market Indicators (LMI) combine monthly data on job flows from the *Current Population Survey* (CPS) with higher-frequency alternative and traditional labor market indicators to construct real-time estimates for the unemployment rate and the flows of workers out of unemployment (job finding) and into unemployment (job separations). The table below lists all the underlying data series in the Chicago Fed LMI, along with their primary sources, observation frequency, and any data transformations applied.

Table 2: Chicago Fed LMI Data Dictionary

Data Series	**Frequency**	**Source**	**Construction**
Civilian (U-3) unemployment rate	Monthly	U.S. Bureau of Labor Statistics *Current Population Survey*	Log units
Job-finding rate	Monthly	U.S. Bureau of Labor Statistics *Current Population Survey*	Log units
Job-separation rate	Monthly	U.S. Bureau of Labor Statistics *Current Population Survey*	Log units
Hiring rate	Monthly	U.S. Bureau of Labor Statistics *Job Openings and Labor Turnover Survey*	Log units; 1-month lag
Layoffs and discharges rate	Monthly	U.S. Bureau of Labor Statistics *Job Openings and Labor Turnover Survey*	Log units; 1-month lag
CB Labor market differential	Monthly	The Conference Board	Jobs plentiful vs. hard-to-get; log units
SCE Labor market differential	Monthly	Federal Reserve Bank of New York *Survey of Consumer Expectations*	Respondents' mean probability of finding a job in less than 3 months less their mean probability of losing a job over the next year; log units; 1-month lag
Initial claims for unemployment insurance	Weekly	U.S. Department of Labor	Initial claims as a share of covered employment; log units of observed values during BLS reference weeks
Insured unemployment rate	Weekly	U.S. Department of Labor	Continuing claims as a share of covered employment; log units of observed values during BLS reference weeks

Job search activity index	Weekly	Morning Consult	Seasonally adjusted using the Prophet tool developed by Facebook's data scientists for the seasonal adjustment of high-frequency data; log units of observed values during BLS reference weeks
Pay loss rate	Weekly	Morning Consult	Seasonally adjusted using Prophet; log units of observed values during BLS reference weeks
Lightcast/ADP job openings rate	Daily	Lightcast via Opportunity Insights, ADP	Lightcast average weekly job openings percentage difference from baseline indexed to January 4-31, 2020 * 5500 / last available monthly ADP nonfarm private payroll employment; log units of observed values during BLS reference weeks; seasonally adjusted using Prophet
Unemployment index	Weekly	Morning Consult	Seasonally adjusted using Prophet; log units of observed values during BLS reference weeks

Unemployment topic index	Daily	Google Trends	Seasonally adjusted using the Prophet tool developed by Facebook's data scientists for the seasonal adjustment of high-frequency data; log units of average daily values during BLS reference weeks constructed using the Eichenauer et al. [2022] method and the trendecon R package
Indeed/ADP job openings rate	Daily	Indeed, ADP	Indeed total job openings index * 5500 / last available monthly ADP nonfarm private payroll employment; log units of weekly average of daily values during BLS reference week

Note: Indeed and Lightcast job opening rates are multiplied by a constant (5500) in order to scale them to similar units as the JOLTS job openings rate. When the JOLTS data are unavailable by the time of one of our scheduled releases, we estimate them. We estimate the JOLTS hiring rate by regressing it on the Revelio Labs hiring rate; the National Federation for Independent Business jobs hard-to-fill index, average change in employment per firm, and net percent of businesses with an increase in employment over the last three months; the Institute for Supply Management manufacturing and services employment indexes; and one lagged value of the JOLTS hiring rate. We estimate the JOLTS layoffs rate by regressing it on the Revelio Labs attrition rate; announced job cuts from Challenger, Gray, & Christmas; continued claims for unemployment insurance from the U.S. Department of Labor; the National Federation for Independent Business average change in employment per firm and net percent of businesses with an increase in employment over the last three months; the Morning Consult pay loss rate; and three lagged values of the JOLTS layoffs rate.

6.2 Tracking Job Flows

The Chicago Fed LMI use the method described in Shimer [2005] and Shimer [2012] to compute job-finding, job-separation, and flow-consisent unemployment rates with CPS data. The *job-finding rate*, f_t, is given by

$$f_t = -log\frac{U_t - U_t^s}{U_{t-1}}, \tag{1}$$

where U is the unemployment level and U^s is the number of civilians unemployed for less than 5 weeks during the period. The *job-separation rate* is then computed by solving for s_t in the following equation:

$$0 = u_t - e^{-s_t - f_t} \times u_{t-1} - \frac{s_t}{s_t + f_t} \times (1 - e^{-s_t - f_t}). \tag{2}$$

The *flow-consistent unemployment rate* (FCR) is calculated as the job-separation rate divided by the sum of the job-separation and job-finding rates:

$$FCR = \frac{s}{f + s}. \tag{3}$$

Because the two job flow rates above are scaled in order to keep the FCR in BLS-unemployment-rate-like units, the FCR can be understood as the unemployment rate we would expect to see from the BLS if current job flow rates remained unchanged over time.

To blend the information from our higher-frequency indicators with the job-finding and job-separation rates derived from the CPS, we use a statistical method called partial least squares (PLS). PLS is designed to produce targeted linear combinations of a set of predictors (call them X) that maximize their covariance with a set of outcomes (call them Y). We include in Y log-transformed CPS-derived job-separation and job-finding rates, and in X the set of traditional and alternative labor-market indicators shown in Figure 1 of the main text. The use of PLS has two main advantages in this context: the first is to remove potential measurement error in monthly CPS flows and in higher-frequency alternative and traditional labor market data that could cloud our later forecasts for the BLS unemployment rate. The second is to leverage higher-frequency alternative and traditional labor market data to estimate the CPS job flows between monthly BLS releases.

The underlying static factor model that we use for this purpose can be summarized as

$$Y = \Lambda F + \varepsilon \tag{4}$$

$$X = \Gamma F + \eta \tag{5}$$

where F is a set of two common components, or factors, and ε and η represent classical measurement errors that are uncorrelated both with each other and separately within the set of indicators in X or Y. We estimate this two-factor model via PLS regression using the Bioconductor package ropls in the R programming language by finding the $\{F, \Gamma, \Lambda\}$ that maximizes $Cov(Y, X)$. We then construct model-implied job-separation and job-finding rates using the estimated factors (F), factor loadings (Γ), and prediction weights (Λ) and transforming our predictions for Y back into their natural units.[8] The model-implied job-separation rate is the Chicago Fed Layoffs and Other Separations Rate, while the model-implied job-finding rate is the Chicago Fed Hiring Rate for Unemployed Workers. Both rates can be decomposed additively into contributions from each data series in X.

With estimates of these two job flow rates, we can also calculate a model-implied FCR. Furthermore, as seen in the main text, we can decompose the change in the FCR over any two time periods into separate contributions from the model-implied job-separation and job-finding rates. To demonstrate, let the FCR be represented by u_t^* where

$$u_t^* = \frac{s_t}{s_t + f_t}$$

[8] Both X and Y are centered prior to estimation with PLS, such that transforming our predictions for Y back to natural units also entails adding back its mean.

in keeping with its definition above. Taking a first-order Taylor expansion around the previous period's values yields the following approximation for the change in FCR:

$$\Delta u_t^* \approx \frac{\partial u^*}{\partial s}|_{t-1}\Delta s_t + \frac{\partial u^*}{\partial f}|_{t-1}\Delta f_t$$

where the partial derivatives are:

$$\begin{aligned}\frac{\partial u^*}{\partial s} &= \frac{f}{(s+f)^2} \\ \frac{\partial u^*}{\partial f} &= \frac{-s}{(s+f)^2}\end{aligned}$$

We can, therefore, define the period-t contributions of job separations and job finding as

$$\begin{aligned}C_t^s &= \frac{f_{t-1}}{(s_{t-1}+f_{t-1})^2}\Delta s_t \\ C_t^f &= \frac{-s_{t-1}}{(s_{t-1}+f_{t-1})^2}\Delta f_t.\end{aligned}$$

Because of the approximation, the sum $C_t^s + C_t^f$ will generally differ slightly from the observed Δu_t^*. We distribute the approximation error proportionally to each component's absolute size. Letting $G_t = \Delta u_t^* - C_t^s - C_t^f$ denote the cumulative gap and

$$\begin{aligned}\omega_t^s &= \frac{|C_t^s|}{|C_t^s| + |C_t^f|} \\ \omega_t^f &= \frac{|C_t^f|}{|C_t^s| + |C_t^f|}\end{aligned}$$

denote the shares, the adjusted cumulative contributions are

$$\tilde{C_t^s} = C_t^s + \omega_t^s G_t$$
$$\tilde{C_t^f} = C_t^f + \omega_t^f G_t$$

such that $\tilde{C_t^f} + \tilde{C_t^s} = \Delta u_t^*$ and all contributions are expressed in percentage points.

6.3 Forecasting the Unemployment Rate

To predict the yet-to-be-released value for the BLS unemployment rate with our model-implied FCR, we focus on the *change* in FCR from the prior month to predict the *change* in the upcoming BLS unemployment rate. The intuition behind this is simple. By focusing on changes in the FCR rather than its level, we capture the upward or downward pressure on unemployment from job flows. This method is also convenient because it allows us to ignore any temporary level differences between the FCR and the BLS unemployment rate and to instead focus on their direction of change between reference weeks.

The average monthly change in the unemployment rate over extended periods is close to zero, despite the volatility around recessions. We account for this fact in our Bayesian linear regression for the reference-week-to-reference-week changes in the BLS unemployment rate and our FCR with a prior that shrinks the regression coefficients toward zero. To illustrate, consider the following linear regression equation relating the change in the BLS unemployment rate (ΔU) to the change in the FCR (ΔFCR):

$$\Delta U_t = \alpha + \beta \Delta FCR_t + \epsilon. \tag{6}$$

In the limit as $\{\alpha, \beta\} \to 0$ (where our prior beliefs exactly hold), this model would predict no change in the unemployment rate given $\epsilon \sim N(0, \sigma^2)$.

Instead of imposing this exact relationship when we estimate the linear regression model above, the strength of this a priori belief is balanced against the strength of the observed relationship between these two measures in our sample of observations that begins in July 2008 to form an a posteriori prediction. This is achieved by specifying priors which shrink α and the overall fit of the regression (i.e., R^2) toward zero at known rates that are themselves a function of σ^2 using the rstanarm package in the R programming language with an informative Lewandowski–Kurowicka–Joe (LKJ) prior.[9] To then obtain a forecast for the level of the unemployment rate, we add back to the predicted change from the Bayesian linear regression ($\Delta\hat{U}$) the previous month's unemployment rate level, i.e.,

$$\hat{U}_t = U_{t-1} + \Delta\hat{U}_t. \tag{7}$$

To assess forecast uncertainty, we next use the conformal inference methods described in Barber et al. [2021] and found in the conformalbayes package for the R programming language to construct finite-sample prediction intervals. These intervals adjust the Bayesian linear regression's credible sets by a correction factor based on historical out-of-sample performance. From these calibrated prediction intervals, it is then possible to construct a full probability distribution for various possible values of $\hat{U}_t$ by making use of the direct relationship between the quantiles of the forecast distribution and the coverage rates of the prediction intervals. We use the distfromq package for the R programming language to take draws from the approximate forecast distribution using a mixture of Normal and Cauchy distributions to match these quantiles and report the sample probabilities in each Chicago Fed LMI release.

6.4 Pseudo Real-Time Forecast Evaluation

The Chicago Fed Real-Time Unemployment Rate Forecast is published twice monthly in: (1) an Advance release, 7 to 10 days out from the U.S. Bureau of Labor Statistics (BLS)

[9] Further details on the prior settings can be found in the interactive dashboard available for download from the Download Center at https://chicagofed.org/lmi.

Employment Situation report, which incorporates any data already available for the reference week covered in the upcoming BLS report; and (2) a Final release, 1-2 days out from the BLS report, incorporating any additional data updates from the intervening week.

With the Advance release, the available higher-frequency alternative and traditional data series used in the LMI typically include the two weekly unemployment insurance (UI) claims series (updated 5-12 days after the weeks they cover) on the traditional side and Google Trends, Morning Consult, Lightcast, and Indeed (all updated on a daily or weekly basis) on the alternative side. When a particular data series is not yet available for the applicable reference week, we leverage its higher reporting frequency and use the latest available data point (usually from the week just before the reference week).

The Final release then accounts for any later data updates occurring closer to the BLS report. The Conference Board labor market differential, for example, typically gets updated on the last Tuesday of every month, while the two JOLTS rates often get published just a few days before the *Employment Situation* report. In the case of JOLTS and the FRBNY labor market differential, we use the previous months' data points to help predict the current month's job flow rates. Even though these two series are reported with a lag similar to that of the unemployment rate (preventing us from using their contemporaneous values as forecast inputs), they have proven useful in the LMI as leading indicators for unemployment.

The pseudo-real-time forecast evaluation exercise described in the main text is built around a real-time data flow resembling a typical Final release for the LMI. Using this simulated data flow, we estimate our Bayesian linear regression for each pseudo-month from January 2018 through March 2026 and take the median of the posterior distribution of the regression's fitted values as our point forecast for the change in the unemployment rate. With this median, we then compare our unemployment rate forecasts for each month against its current values as of April 2026 to construct a time series of historical forecast errors from which to calculate mean absolute and root mean-squared forecast errors. To maintain a consistent data flow across months, we exclude the fourth quarter of 2025 from our analysis, as public data sources were temporarily unavailable during the federal government shutdown that occurred during this time.[10]

To compare the performance of the Chicago Fed LMI with the benchmark forecasts in the main text, we use Diebold-Mariano (DM) tests (Diebold and Mariano [1995]). These one-sided hypothesis tests are for the null hypothesis that there is no difference between the LMI and the benchmark forecast and the alternative that the LMI are more accurate than the benchmark forecast. Following Harvey et al. [1997] and Kiefer and Vogelsang [2005], we apply a size correction to the DM test statistic and use Bartlett kernel weights with lag length equal to 33% of the sample size in order to ensure a valid heteroskedastic and autocorrelation consistent (HAC) variance estimate. We then preserve these choices for the Giacomini-Rossi fluctuation tests (Giacomini and Rossi [2010]) shown in the main text.

[10]Both the point and density forecast accuracy results from the main text are regularly updated in the interactive dashboard available for download from the Download Center at https://chicagofed.org/lmi.