A New Labor Market Stress Indicator*

Rohit Garimella† Óscar Jordá‡ Sanjay R. Singh§

December 2025

Abstract

Recessions are periods where the labor market deteriorates rapidly. Supporting business conditions to prevent such deterioration is a core objective of policymakers. In this paper we construct a labor market stress indicator (LMSI) primarily based on state-level unemployment insurance claims data that are observable as often as at weekly frequency. By examining both the geographical spread and the depth of labor market stress buildup, we provide an early indicator whose main function is to alert policymakers of potential economic slowdowns. Because the majority (but not all) of these slowdowns coincide with NBER recessions, the LMSI is also a useful signal of whether the economy is in recession. The paper then evaluates this feature of the LMSI compared with other recent indicators and highlights the strengths and weaknesses of each.

Keywords: labor market indicator, recession, receiver operating characteristic curve.
JEL Codes: E32, E65, E66, J11, J60

*We thank Manfred Keil, Sylvain Leduc, Thomas Mertens, Nicolas Petrosky-Nadeau, and Dan Wilson for useful comments and suggestions. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System. Remaining errors are our own.

†Federal Reserve Bank of San Francisco. Email: rohit.garimella@sf.frb.org.
‡Federal Reserve Bank of San Francisco, UC Davis and CEPR. Email: Oscar.Jorda@sf.frb.org.
§Federal Reserve Bank of San Francisco and UC Davis. Email: sanjay.singh@sf.frb.org.

1. Introduction

Timely assessments of a slowdown in the economy are central to macroeconomic stabilization policy. Monetary and fiscal authorities must decide, for example, whether a softening labor market signals a localized adjustment to idiosyncratic conditions, or the onset of a broad-based downturn. Because slowdowns can turn into recessions, early warning signals of rapidly deteriorating economic conditions are a valuable policymaker's tool. Traditionally, recession monitoring has relied on a small set of national indicators: the slope of the Treasury yield curve, which has a long track record as a predictor of U.S. recessions (see, e.g. Estrella and Mishkin, 1998; Rudebusch and Williams, 2009; Bauer and Mertens, 2018), and labor-market based rules, such as the Sahm (2019) rule, for example. The Sahm-rule calls the start of a recession when the three-month moving average of the national unemployment rate increases by at least 0.5 percentage point above its 12-month lowest value. Tools such as these are simple and powerful, but they refer to the aggregate economy as a whole. They are silent about geographic, heterogeneous labor-market stress. National signals provide little insight into regional conditions.

Recent episodes highlight the importance of this distinction. In mid-2024, the national Sahm-rule briefly crossed its 0.5 threshold. In addition, the yield curve had been inverted for some time. Together, these indicators prompted concerns that a recession had begun or was imminent. Yet labor-market stress was not evenly distributed across the country: many large states remained relatively resilient, and subsequent data revisions and outcomes suggest that this episode did not mark the start of a broad-based downturn. A framework that explicitly tracks the geographic dispersion and depth of labor-market stress can therefore sharpen both real-time diagnosis and medium-term recession risk assessments.

In this paper, we develop a new *Labor Market Stress Indicator* (LMSI) that places the geography of unemployment dynamics at the center of monitoring economic conditions. Building on recently digitized, historical, state-level unemployment insurance claims data (Fieldhouse, Howard, Koch, and Munro, 2022, 2024), we construct a monthly panel of state unemployment rates back to the late 1940s. The basic idea is to proxy state-level unemployment rates with unemployment-claims data. Thus, we first fit such a regression where the samples of state-level unemployment claims and rates overlap and then use the coefficient estimates to backfill the data back to the 1940s. Based on this expanded time series, we then extend the logic of the Sahm-rule to these fitted state level unemployment rates. That is, a state is classified as experiencing "accelerating unemployment" when its three-month moving average unemployment rate lies at least 0.5 percentage point above its minimum over the previous year. Our baseline LMSI counts the number of states in this accelerating bin at any point in time. In addition, we also calculate a labor-force–weighted version of the index, which measures the share of the national labor force residing in these states to get a sense of how many workers are affected.

As a result of this data effort, we create a long historical panel of geographically disaggregated labor-market stress series. In addition to generating real time signals on labor market conditions, we show that whenever 30 or more states first experience accelerating unemployment, the national economy has always been in an NBER recession. Moreover, during such episodes, about three-quarters of the U.S. labor force turn out to live in states with accelerating unemployment. Conversely, peaks in the LMSI below this range—such as the July 2024 episode, when only 25 states and about 47 percent of the labor force met the acceleration criterion—have not historically coincided with the onset of recessions. This dual perspective on *breadth* (number of states) and *depth* (labor-force share) allows us to distinguish genuine national downturns from more localized or sectoral imbalances.

We then evaluate the LMSI as a recession indicator in the spirit of the recent literature on unemployment-based rules. Sahm (2019) documents that her national rule has a near-perfect track record for identifying the start of recessions since 1960. Michaillat and Saez (2025) propose a "minimum indicator" that combines unemployment and vacancy data and introduces a two-sided threshold rule to detect recessions earlier than the Sahm-rule. Related work has explored alternative transformations or filters of national labor-market series (e.g. O'Trakoun and Scavette, 2025). We contribute to this literature by bringing high-frequency, claims-based state unemployment data to bear on labor market monitoring and by emphasizing geography as a key organizing dimension.

We assess the ability of the LMSI to classify each month into a recession/expansion period based on the National Bureau of Economic Research's (NBER) dating of recessions using the receiver operating characteristic (ROC) curve. The ROC curve, and in particular, the area under the ROC curve (AUROC), provide a formal statistical approach to compare alternative recession indicators, as Berge and Jordà (2011) show. As a reference, a coin-toss classifier has an AUROC of 0.5 whereas a perfect classifier has an AUROC of 1. Thus, the AUROC has a similar flavor to the R-squared in a regression. We find that, in real time, the LMSI attains an AUROC of around 0.87, similar to the national Sahm-rule and close to the vacancy–unemployment minimum indicator of Michaillat and Saez (2025). In real time, these indicators outperform the standard 10-year T-bond minus 3-month T-Bill term spread, though not, as we will see, at horizons beyond one year, approximately. At forecast horizons of one to two years, the LMSI's predictive power is comparable to that of the Michaillat-Saez indicator and somewhat weaker than the term spread. The latter is well-known to embed information about economic activity at longer horizons (see, e.g. Bauer and Mertens, 2018; Engstrom and Sharpe, 2019). However, why not combine the information across all indicators? A simple logit model that combines the LMSI, the Sahm-rule, the Michaillat-Saez indicator, and the term spread yields an AUROC of about 0.96, suggesting that there are substantial gains from exploiting their complementary strengths.

A second contribution of the paper is to use the LMSI to study the geography of labor-market stress across states and over the business cycle. Our approach is closely related to the literature on regional labor-market comovement and cyclicality (Blanchard, Katz, Hall, and Eichengreen, 1992; Fieldhouse et al., 2024; Russ, Shambaugh, and Singh, 2024), which documents persistent heterogeneity in exposure to aggregate shocks. We show that some "bellwether" states almost always appear among the first 30 accelerating states when the LMSI crosses its lower recession threshold. Others—often resource-intensive states in the Plains and Mountain West—are far less systematically aligned with national downturns. This cross-sectional heterogeneity helps explain why the July 2024 spike in the national Sahm-rule did not coincide with a recession. Many of the historically representative bellwether states remained outside the accelerating bin. Instead, labor market stress was concentrated in states that are typically less synchronized with the national business cycle.

Our paper is also directly inspired by the work of Edward E. Leamer on recession monitoring and early warning indicators. Leamer (2007) shows that residential investment offers by far the best early warning signal of postwar U.S. recessions, with virtually every downturn preceded by substantial weakness in housing and consumer durables. Leamer (2022) revisits recession forecasts based on monthly data, and Keil, Leamer, and Li (2023) develop binary-variable models that reinterpret traditional yield-curve probit models, both highlighting the enduring predictive power of the yield curve while advocating for models that combine simple, transparent indicators that can be communicated clearly to policymakers. Our Labor Market Stress Indicator is very much in this spirit: rather than relying on a single national aggregate, it tracks the diffusion of stress across states and over time, offering a simple, visual narrative about where recessions come from and how they spread.

Finally, we construct a weekly version of the LMSI using state-level unemployment insurance claims, available since 1987. Calibrating a weekly acceleration threshold by scaling the Sahm-style 0.5 percentage point rule to the level of the claims-based unemployment rate. Moreover, we show that a weekly LMSI closely tracks its monthly counterpart in correlating with past recessions, while providing a more timely read on evolving labor-market conditions. Unlike indicators that rely on official unemployment rates, the weekly LMSI can continue to be updated even when statistical releases are delayed or disrupted—for example, during federal government shutdowns—and it remains tightly focused on labor-market stress, in contrast to broader composite indexes such as the Weekly Economic Index (Lewis, Mertens, Stock, and Trivedi, 2022).

Taken together, our results suggest that counting states with accelerating unemployment is a simple, transparent, and empirically powerful way to summarize the geographic breadth and depth of labor-market stress. In addition to serving as a competitive recession indicator, the LMSI offers a flexible platform for analyzing regional heterogeneity, comparing the informational content of different signals (unemployment, vacancies, yields), and monitoring labor-market conditions at both monthly and weekly frequencies.

The rest of the paper is organized as follows. Section 2 describes the data and our methodology for constructing monthly and weekly claims-based state unemployment rates and the corresponding LMSI measures. Section 3 briefly describe the statistical metric that we use to formally assess the ability of different indicators to sort the sample into periods of expansion/recession. Section 4 studies the empirical properties of the LMSI, evaluates its in-sample and out-of-sample performance as a recession indicator, and compares it to the Sahm rule, the Michaillat–Saez minimum indicator, and the yield curve, including a detailed discussion of the July 2024 episode. Section 5 investigates the geography of the business cycle. Section 6 shows how best to combine indicators to generate assessments about the business cycle in real time. Section 7 concludes by discussing policy implications and avenues for future work, including the use of the LMSI for regional monitoring and real-time surveillance.

2. Constructing the LMSI: Data and methodology

In this section we discuss the sources of data that we use to construct the LMSI as well as the methods to expand the historical sample. The discussion is divided into the construction of the monthly indicator first, and then the construction of the weekly indicator.

2.1. The monthly LMSI

We rely on Fieldhouse et al. (2024) for their digitized historical data on monthly state-level unemployment claims and payroll data. We first convert the state unemployment insurance claims data to average weekly data by dividing by the number of weeks in each month. We then combine these data with averaged weekly unemployment claims data from the

Department of Labor to obtain a full claims dataset spanning December 1946 to September 2025. This way we expand the sample with 30 more years of data than what is available for state-level unemployment rates. To adjust for seasonality, we run an X-13 ARIMA-SEATS seasonal adjustment procedure on the initial and continuing unemployment claims as well as the nonfarm payroll figures. Hence, we construct a claims-based unemployment rate, $CBUR_{i,t}$, with the following formula:

$$CBUR_{i,t} = \frac{IC_{i,t} + CC_{i,t}}{IC_{i,t} + CC_{i,t} + NFP_{i,t}}, \tag{1}$$

where $IC_{i,t}$ and $CC_{i,t}$ are the three-month moving average of the initial and continued claims ending in month t for state i and $NFP_{i,t}$ is the nonfarm payroll employment.

Figure 1: *Raw claims-based unemployment rate, fitted and actual unemployment rate for California*

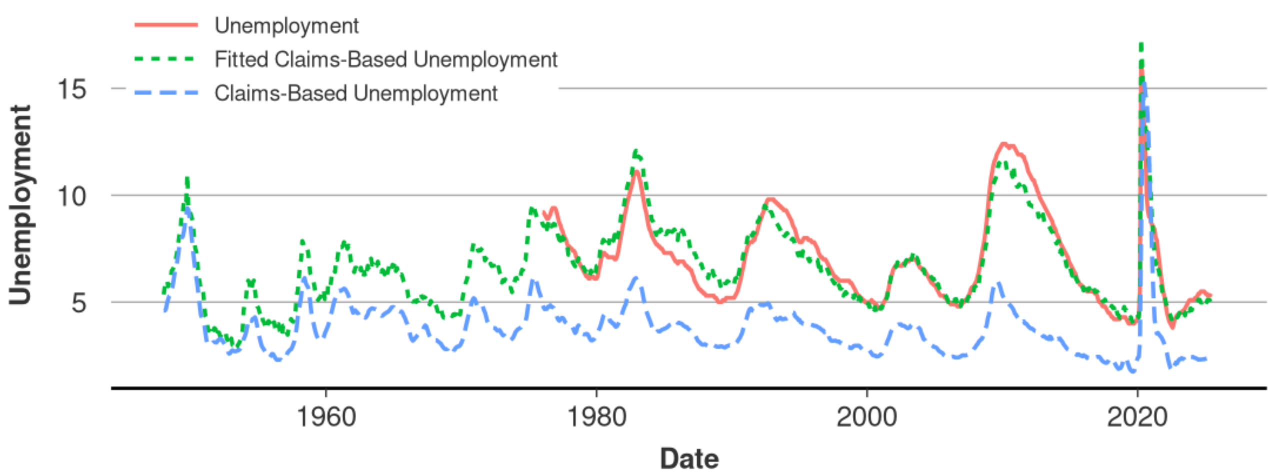

Notes: The blue long-dashed line shows the series that results from Equation 1. The green short-dashed line is the fitted unemployment rate based on regressing the actual unemployment rate on Equation 1 as shown in Equation 2. The solid red line is the raw unemployment rate. All data shown are for California. See text.

As an illustration, Figure 1 shows the claims-based unemployment rate for California compared to the actual state unemployment rate. The raw measure based on Equation 1 attains values consistently lower than the actual unemployment rate. This is likely due to the fact that there is large share of unemployed workers who do not apply for unemployment insurance. In fact, in 2022, only 26 percent of unemployed workers who had worked in the previous 12 months applied for unemployment insurance benefits since separating from their last job Bureau of Labor Statistics (2023). Even in this raw format, the correlation with the unemployment rate is 0.72. The green short-dash line in the figure shows the fitted unemployment rate over the sample where $CBUR_{i,t}$ and the unemployment rate overlap to show that indeed, $CBUR_{i,t}$ does a very good job at characterizing the unemployment rate and thus can be used to backcast the unemployment rate data back to 1949.

In particular, we do this with the following simple regression:

$$UR_{i,t} = \beta_{0,i} + \beta_{1,i}(CBUR_{i,t} - CBUR_{Natl,t}) + \beta_{2,i}UR_{Natl,t} + \epsilon_{i,t}, \tag{2}$$

where $CBUR$ represents the claims-based unemployment rate in Equation 1 and UR is the official national unemployment rate from the Bureau of Labor Statistics (BLS). We fit this regression individually for each of the 50 states and the District of Columbia (DC). A summary of the range of R^2 measures of fit is reported in Table 1. Individual results are reported in Table A.1 in the appendix. Using the fitted values of this regression, we then construct the monthly indicator using the official state-level unemployment rates in the available sample and backfill the data to 1949 with each state's fitted claims-based unemployment rate.

Table 1: *Claims-Based Unemployment Regression R^2 Range*

	Min	Max	Mean
R^2	0.64	0.95	0.84

Notes: range of R^2 measures of fit based on estimates of Equation 2 for each of the 50 states and DC. See text.

Based on the series of state-level unemployment rates extended using Equation 2, next we apply the Sahm-rule. Recall that this is defined as the three-month average of a state's unemployment rate increasing by at least 0.5 percentage point above its previous 12-month lowest value. Using this rule, we then count the number of states where the Sahm-rule is triggered and simply declare that these states are experiencing "accelerating unemployment." Since official state unemployment data begins in 1976, we use this approach on the fitted claims-based rates for March 1949 - February 1977 and on the official unemployment rates for March 1977 through the rest of the sample.

Figure 2: *Monthly Labor Market Stress Indicator*

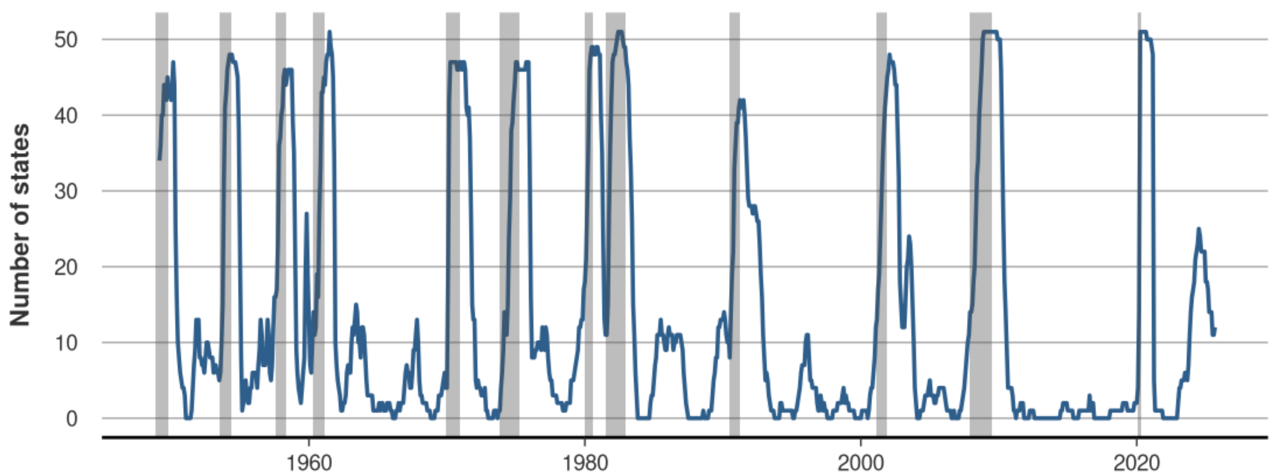

Notes: Number of states with accelerating unemployment based on the Sahm-rule and the state-level unemployment rate series constructed by backfilling the sample based on Equation 2. Shaded regions correspond to NBER recession periods. See text.

Figure 2 shows the monthly labor market stress indicator with shaded areas representing NBER-defined recession periods. This baseline LMSI counts the number of states in the accelerating bin at any point in time. The indicator correlates well with recessions throughout modern U.S. economic history. The national economy has invariably been in recession each time 30 or more states simultaneously experienced accelerating unemployment. Notably, the indicator hit a value of 25 in July 2024; this is also when the Sahm-rule first crossed its 0.5 threshold and reached 0.53. This episode shows that, while on aggregate the labor market had shown signs of stress in mid-2024, it was not a broad-based phenomenon.

One way to visualize where labor market stresses were and were not present on July 2024 is with Figure 3, which displays the map of the United States showing states experiencing accelerating unemployment as shaded. The map illustrates that states in the Midwest, Mountain West, and West Coast were experiencing an acceleration in their unemployment rates relative to states in the Mid-Atlantic or the South. Importantly, Texas, a large state in terms of population and economic relevance showed no evidence of labor market stress based on our indicator.

Figure 3: *Map of the monthly LMSI for July 2024*

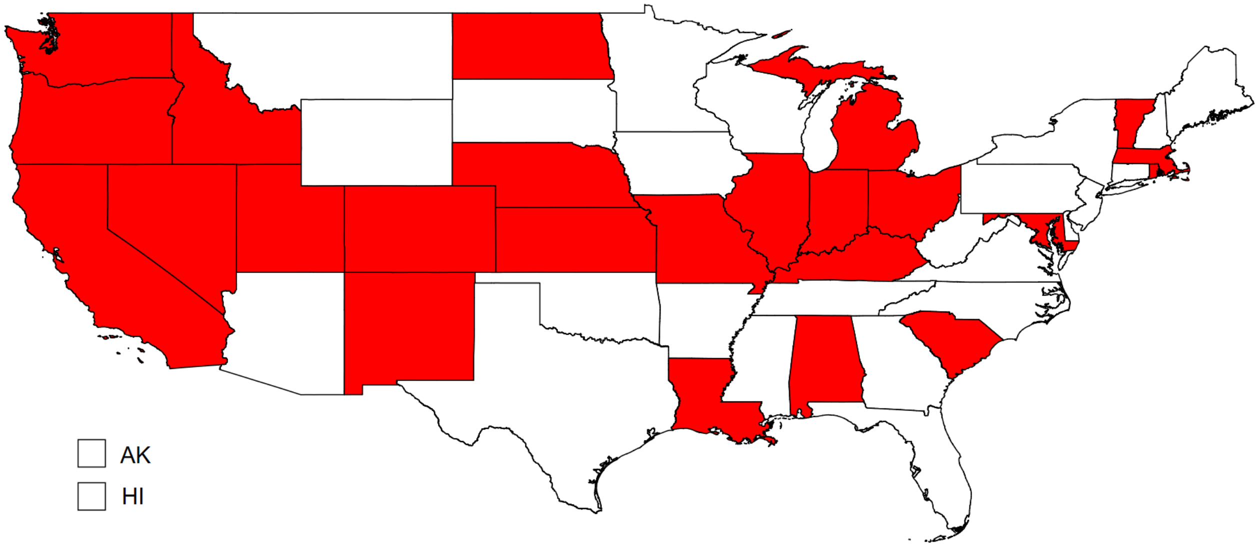

Notes: States where the LMSI showed evidence of accelerating unemployment based on our application of the Sahm-rule shaded in red. See text.

As Figure 3 clearly shows, states in the union vary considerably in size. Thus a simple count of how many states are distressed may over-represent the amount of stress in the economy if the majority of these states happen to be relatively small. California (the largest state), has a population that is nearly 36 times larger than that of the smallest state (Rhode Island). In order to adjust our indicator for the variation in size among states, we calculate what share of the national labor force lives in states with accelerating unemployment. This is shown in Figure 4. However, note that state-level labor force data are only available since 1976.

The figure displays a similar pattern to the LMSI. Again focusing on July 2024, we see that the share of the labor force in distressed states amounted to somewhat less than one-half (47% to be precise), again suggesting that this episode was far from being typical of a recessionary period. In previous recessions, roughly 75% of the U.S. labor force lived

Figure 4: *Monthly Labor Market Stress Indicator as Share of Labor Force*

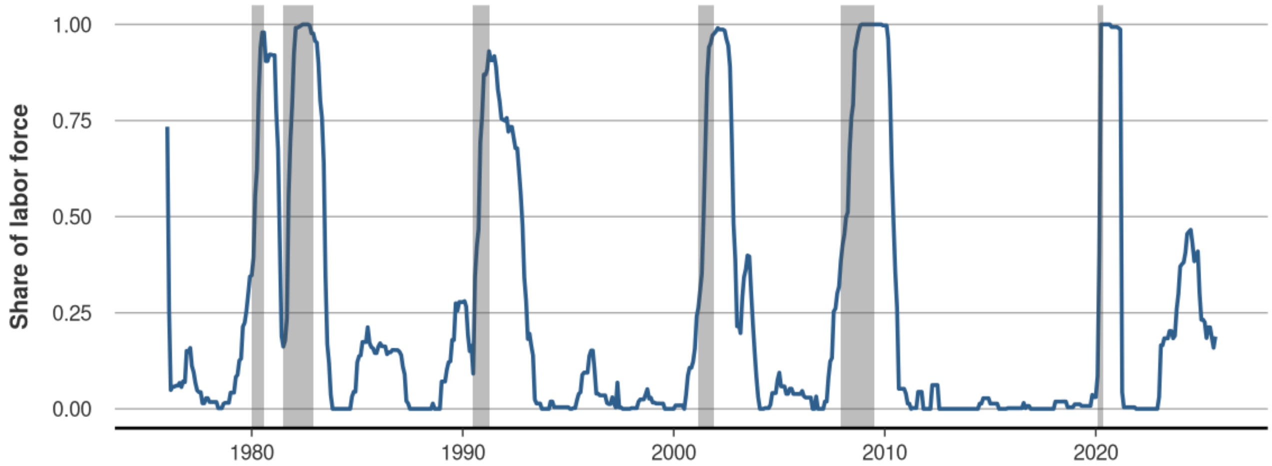

Notes: Share of the labor force based on states where the LMSI showed evidence of accelerating unemployment. Shaded regions correspond to NBER recession periods. See text.

in states experiencing accelerating unemployment. Comparing this metric to the 30 state threshold for the LMSI demonstrates the importance of both being aware of the geographic dispersion of labor market stress as well as measuring the size of the labor force affected by weak labor market conditions.

Another relevant feature of the LMSI and its labor force weighted counterpart is the timing of the index with respect to the duration of recessions. As is well documented (see, e.g. Shimer, 2012; Graetz and Michaels, 2017; Jaimovich and Siu, 2020), labor markets are much slower to recover than the return of GDP growth or even the return of the level of GDP back to its level before the recession started. This is quite clear in Figure 2 and Figure 4. For example, in the 1990-1991 recession, which was rather brief by historical standards, our LMSI indicators remained elevated for another 2-3 years.

The observation that the value of the our indicators remain elevated well after the recession has ended, by the dating provided by the NBER, has important ramifications for policy and for the evaluation of the indicators for the purposes of classifying periods into recession/expansion. From a policy point of view, as the 1990-1991 recession example shows, the relevant information is whether labor markets are in distress or not and hence whether policy needs to be supportive or not, well after the recession has ended. In contrast, if we evaluate the indicators solely on their ability to sort the data into expansions/recessions, the LMSI will be clearly handicapped. We see this as a feature, not a bug.

We conclude this section with a brief robustness check. One concern with our methodology is that it pastes fitted claims-based state unemployment rates with the available official data to create a longer time series than what is available. In order to assess whether this procedure may be affecting the results, Figure A.1 shows the results of creating the indicator using only the fitted CBUR for each state. This series is nearly identical to the monthly LMSI indicator reported in Figure 2 with a correlation of 0.98.

2.2. The weekly LMSI

The previous section showed that we can use unemployment insurance claims data to extend the series of state-level unemployment rates with a considerable degree of accuracy. Another benefit is that unemployment claims data are available at a weekly frequency, which allows us to create a weekly indicator of labor market conditions similar to the LMSI. State-level unemployment claims have been available at weekly frequency since 1987.

We approach the problem of constructing a weekly labor conditions indicator as follows. We rely directly on the CBUR constructed as shown in Equation 1. Hence we construct the number of states experiencing accelerating unemployment using the same process as the monthly indicator but where we adjust the 0.5 threshold to account for the differences shown in Figure 1. In particular, we multiply the 0.5 percentage point threshold by the average ratio of the national CBUR to the official national unemployment rate. This gives us a new acceleration threshold of approximately 0.2 percentage point. Using this new threshold value, we then construct the weekly version of the LMSI. Moreover, just as we did in the previous section, we can also calculate the share of the labor force in states with accelerating unemployment at weekly frequency. In addition to the increased timeliness of the weekly indicator, another benefit of this series is that it is mostly immune to national data availability. For example, during the US government shutdown of 2025, state unemployment data was not updated. However, unemployment claims data are produced by state governments and was therefore, readily available.

Figure 5: *Weekly Labor Market Stress Indicator*

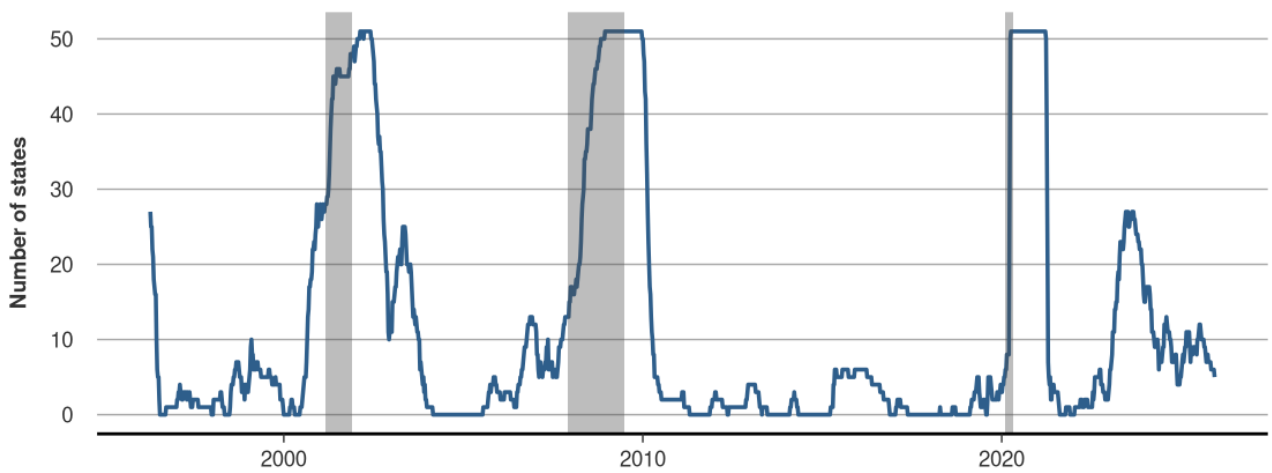

Notes: Number of states with accelerating unemployment based on the adjusted Sahm-rule with a 0.2 threshold and based on the CBUR constructed as in Equation 1 with weekly unemployment claims data. Shaded regions correspond to NBER recession periods. See text.

Figure 5 shows the past 30 years of the weekly version of the LMSI, that is, the number of states experiencing accelerating unemployment since 1995. Compared with Figure 2, we get a strikingly similar picture. In fact, focusing on the July 2024 episode, we clearly see that labor market distress was limited to about half the country. As a way to further inspect this episode, we can also construct a labor share weighted indicator like that of Figure 4. This is shown in Figure 6. The story is similar though the finer granularity of the weekly data suggests that the July 2024 episode briefly involved more than 50% share of the national labor force, a feature that gets lost when using the monthly indicator.

It is useful to illustrate how the weekly indicator can be used in real time. As of data from December 6, 2025, our weekly LMSI shows that five states are currently experiencing accelerating unemployment, representing only about 7% of the national labor force. Even though these readings appear slightly elevated relative to pre-COVID levels, the measure

Figure 6: *Weekly Labor Market Stress Indicator as Share of Labor Force*

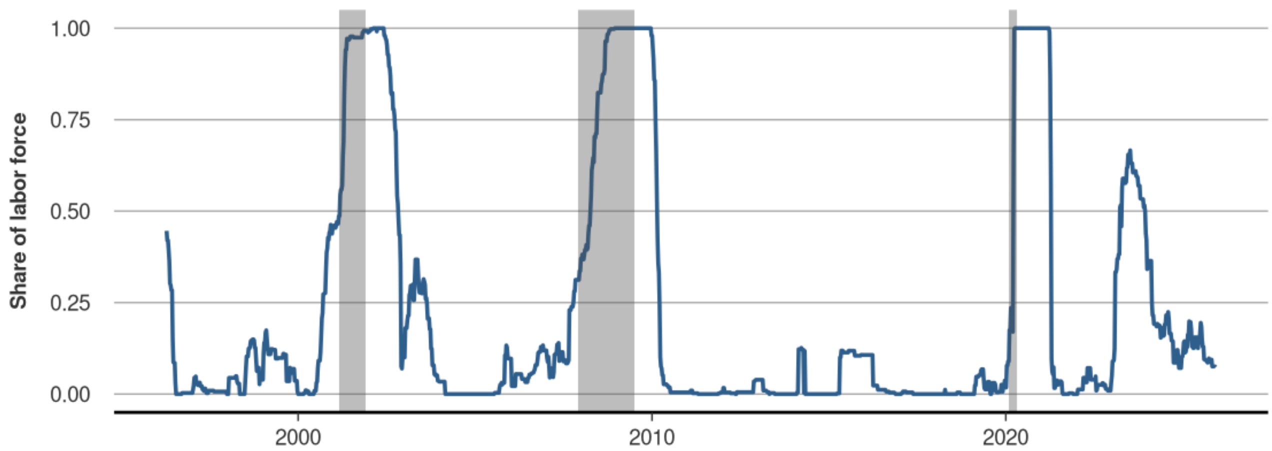

Notes: Share of the labor force based on states where the weekly LMSI showed evidence of accelerating unemployment. Shaded regions correspond to NBER recession periods. See text.

generally indicates that the labor market remained relatively stable at the end of 2025. While the weekly LMSI complements other weekly indicators, like the Weekly Economic Index from the Federal Reserve Bank of Dallas, the LMSI has the advantage of also capturing the geographic heterogeneity of economic conditions. For example Figure 7 shows the map of labor market stress as of December 6, 2025. Only five states, including DC, are experiencing accelerating unemployment as measured with the claims data.

Although the weekly LMSI provides a timely picture of labor market conditions, in the remainder of the paper we focus on the monthly LMSI version that starts in 1949. The longer sample provides with enough recession episodes to be able to conduct an analysis of its business cycle properties with some degree of confidence. First, we discuss the statistical methods that we use to do this type of analysis.

3. Statistical design

As we discussed in the introduction, there are several indicators whose aim is to signal, in real time and in the future, whether the economy will be in recession. In the U.S., the NBER dates periods of expansion and recession using a variety of indicators meant to capture a

Figure 7: *Map of Weekly LMSI for December 6, 2025*

Notes: States where the weekly LMSI showed evidence of accelerating unemployment based on our application of the adjusted Sahm-rule shaded in red. See text.

"[...] significant decline in economic activity that is spread across the economy and lasts more than a few months." "[...] These include real personal income less transfers, nonfarm payroll employment, employment as measured by the household survey, real personal consumption expenditures, manufacturing and trade sales adjusted for price changes, and industrial production." "[...] Two measures that are important in the determination of quarterly peaks and troughs, but that are not available monthly, are the expenditure-side and income-side estimates of real gross domestic product (GDP and GDI). The committee also considers quarterly averages of the monthly indicators described above, particularly payroll employment." (see `https://www.nber.org/research/business-cycle-dating`).

A common strategy to generate business cycle predictions consists of proposing a statistical model (such as a binary dependent model) where the state of the economy is determined as a function of a set of covariates. Instead, in this section we rely on the methods proposed in Berge and Jordà (2011), which we briefly review here.

Let $S_t \in \{0, 1\}$ denote the true state of the economy, with 0 denoting that period t is an expansion period, and 1 a recession period instead. Let Y_t denote a real-valued random variable, perhaps an index constructed to summarize a vector of variables, a probability prediction from a model, or perhaps simply an observable variable. Hence consider the binary prediction consisting of calling an expansion whenever $Y_t \geq c$, and a recession otherwise, where c is some threshold value. Associated with this prediction, we can define the following probabilities:

$$\begin{cases} TP(c) & = \mathbb{P}[Y_t \geq c | S_t = 1] \\ FP(c) & = \mathbb{P}[Y_t \geq c | S_t = 0] \\ & \\ TN(c) & = \mathbb{P}[Y_t < c | S_t = 0] \\ FN(c) & = \mathbb{P}[Y_t < c | S_t = 1] \end{cases} \quad ; \quad \text{with} \quad TP(c) + FN(c) = 1 \quad \text{and} \quad TN(c) + FP(c) = 1 \tag{3}$$

These probabilities highlight the issues surrounding the evaluation of any binary decision problem. Since recession periods represent about 1/10 of the sample, a prediction dating all periods as being in expansion would get a hit rate of 90%, an apparent success. Measures of fit in traditional binary dependent models suffer from the same issue. However, as is clear from Equation 3, calling all periods as being in expansion would have a true positive rate, $TP(c)$, of 1, but also a false positive rate, $FP(c) = 1$. Thus, we seek a measure that balances these two extremes. The earliest exposition of the decision theory behind problems such as this perhaps goes as far back as Peirce (1884).

One way to think about this classification problem is similar to a production possibilities (PP) frontier. Given a predictive model, we want to maximize the "production" of $TP(c)$ and $TN(c)$. A plot of this PP frontier in $TP - FP$ space is called the receiver operating characteristic (ROC) curve. Since $FP(c) = 1 - TN(c)$ it is clear that a representation in $TP - TN$ space contains the same information as the more traditional plot of the ROC curve in $TP - FP$ space. The ROC curve has a long tradition in statistics that started with the theory of radar detection (see, e.g. Peterson and Birdsall, 1953), and is used routinely in the evaluation of medical tests (the earliest citation is perhaps Youden, 1950).

The analogy with the PP frontier serves to illustrate the two extreme values typically associated with the area under the ROC curve of AUROC. A classifier consisting of flipping a coin will generate a ROC cuve that bisects the unit-square and thus the AUROC will be 0.5. A perfect classifier achieves perfect sorting of the data, that is $TP = TN = 1$ and hence the AUROC is 1. Thus, the closer the AUROC is to 1, the better the classification ability. This is the statistic on which we will base our comparisons in the next section.

4. Evaluating business cycles with the LMSI

How good a job does the monthly LMSI do in detecting periods of recession? We evaluate the LMSI by first constructing recession probabilities for each value of the LMSI by estimating a logit model where the dependent variable is a binary NBER recession dummy variable. This model serves two purposes: (1) it translates the LMSI into recession probabilities; and (2) it provides a first assessment of its AUROC. Figure 8 shows the time-series of recession probabilities based on the monthly LMSI, along with a projection of the recession probability going two years out. We estimate this probability by running 24 different logit regressions where the left hand side variable is shifted forward by one month at a time, that is, we use a direct forecasting approach.

One might have expected that the LMSI would have generated probabilities close to 1 for periods of recession and close to 0 otherwise. Alas, as is clear from the figure, the highest fitted probability values rarely exceed 0.6. This is a common finding in the literature. The reason is that, whereas the marginal contribution of another unit of the covariate at either end in a logit regression has a marginal effect on the predicted probability that approaches zero, covariates rarely asymptote toward $[-\infty, \infty]$ at the extremes. Thus, a direct readout of these probabilities is likely misleading. However, this simple model has an AUROC of 0.87, suggesting that the LMSI is a good classifier. That said, as of September 2025, the LMSI

Figure 8: *Recession probabilities based on the LMSI*

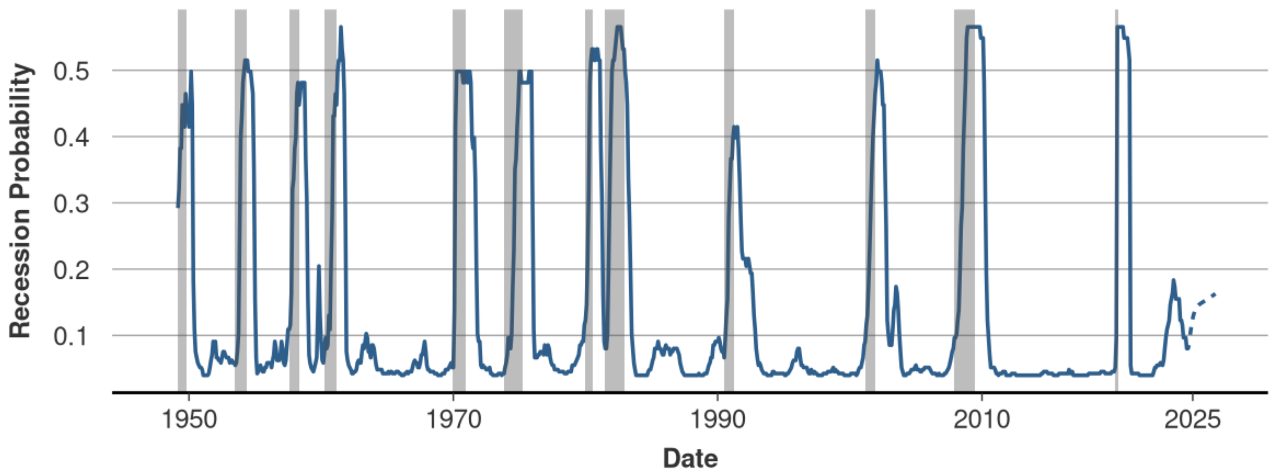

Notes: Predicted probabilities of recession based on a logit model with the NBER recession dummy as the dependent variable and the LMSI as the explanatory variable. The dashed portion is a prediction of the probability of recession over the next two years based on a direct forecasting approach. See text.

predicts a 9 percent chance of a recession, and a peak of 16 percent chance over the next 24 months.

Going back to the original LMSI, we find that the indicator never crosses 28 during non-recessionary periods. In other words $TN(28) = 1$ in the notation of Equation 3. The highest non-recession peak occurred in November 1959 with 27 states. This corresponds to the predicted logit probability never crossing 0.25. Similarly, we could ask, What is the smallest number of states associated with a recession state? By a similar calculation we find that threshold to be 41, which corresponds to the 1991 recession where the value of the LMSI was 42. That is, from Equation 3 then $TP(41) = 1$. With these two values we can then construct a dual, rule-of-thumb, threshold for the LMSI along the lines of what Michaillat and Saez (2025) do for their indicator and hence choose 30 and 40 as reasonable round numbers. This is shown in Figure 9

The indicator came close to the lower threshold $TN(28) = 1$ in July 2024 when there were 25 states experiencing accelerating unemployment. This was also when the national Sahm-rule was triggered at 0.53 before peaking at 0.57 the following month (recall that the Sahm-rule threshold is 0.5). This period thus represents a way in which the LMSI can

Figure 9: *Monthly LMSI with Upper and Lower Thresholds à là Michaillat and Saez (2025)*

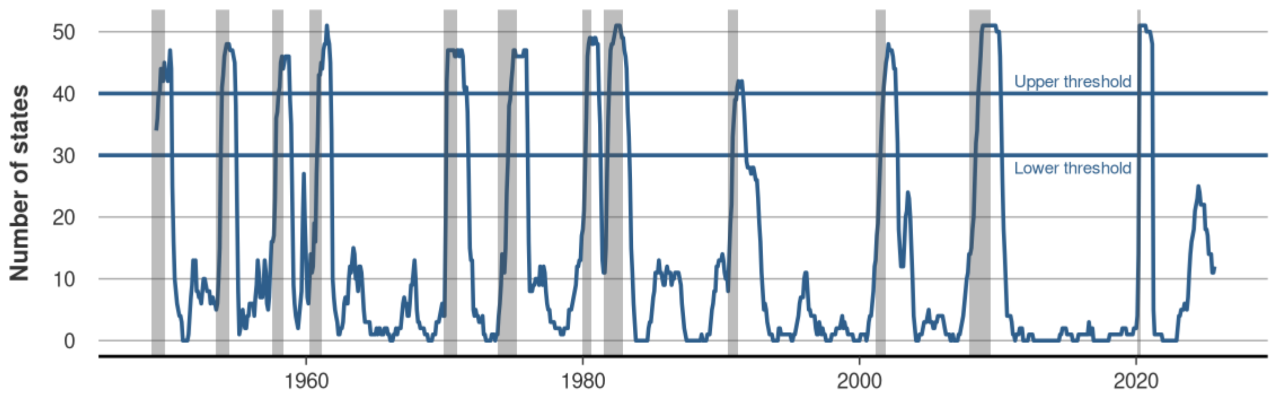

Notes: Monthly LMSI shown with the 30 and 40 thresholds that ensure $TN(28) = 1$ and $TP(41) = 1$. See text.

better describe the breadth of economic stress relative to nationally aggregated indicators. By looking at the state level composition, we find that only four of the ten largest states by population were experiencing labor market stress: California, Illinois, Ohio, and Michigan. This may have limited the spread of labor market stress to other states, thus precipitating a more serious downturn. In contrast to previous recessionary periods, the buildup of the LMSI to this peak was slower, as the indicator has generally quickly jumped to above 30 states in the previous recessions. Furthermore, in the July 2024 episode, while roughly half the number of states were experiencing accelerating unemployment, this translated to 47 percent of the labor force being affected, as we showed in Figure 4. This figure further captures the more localized nature of labor market disruptions as states with larger labor forces did not display signs of a worsening labor market.

5. The geography of business cycles

Our discussion of the July 2024 episode highlights the importance of understanding the dispersion of distress in the labor market across states. In that vein we next ask whether there are states that are more closely linked to the national business cycle than others. That is, are there states that are always implicated whenever there is distress in the labor market

Figure 10: *Frequency of States with Accelerating Unemployment when Threshold Crossed*

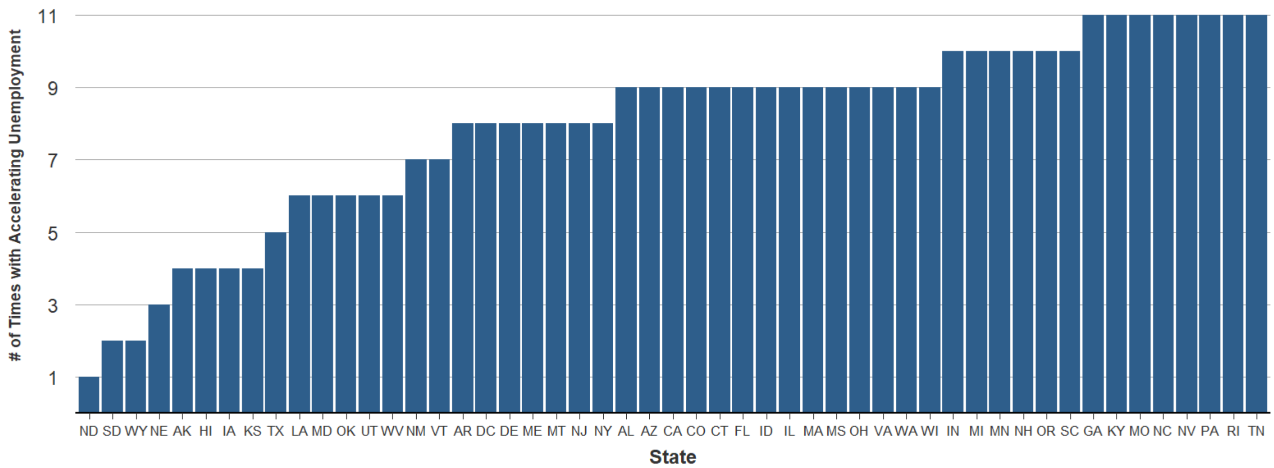

Notes: Frequency of times each state was implicated each time the LMSI crossed 30. See text.

(and hence a good chance that the economy is in recession)? And similarly, are there states that seldom follow national trends? To investigate these questions we do the following. For each episode in which 30 states or more experienced accelerating unemployment, we count the number of times that each state crossed the that threshold.

Figure 10 provides the rank of states by the frequency when they are one of the 30 states to have accelerating unemployment when the lower rule-of-thumb threshold of 30 is first crossed. On one side of the spectrum, Tennessee, Rhode Island, Pennsylvania, Nevada, North Carolina, Missouri, Kentucky, and Georgia have experienced accelerating unemployment in each of the previous 11 instances when the threshold is crossed, all of which coincided with a national recession. Thus, we could call these the "bellwether" states. They seem to be always associated with deteriorating conditions at the national level. In line with this observation, note that during the July 2024 episode, large bellwether states like Pennsylvania, Georgia, and North Carolina, were not in the accelerating unemployment bin. This might have been another clue indicating that the episode could not be counted as a national recession period.

On the other side of the spectrum, states like North Dakota, South Dakota, Wyoming, Nebraska, Alaska, Hawaii, Iowa, Kansas, and Texas, all have experienced labor market stress less than half the time when the lower threshold of 30 is crossed in the national economy. Their labor market is not as coincident with national labor market disruptions. It is notable that these states encompass the same region in the eastern plains. Together with Alaska, these states rely heavily on industries that extract natural resources, where swings in the unemployment rate are less likely to be tied to the national business cycle. In July 2024, we find that Nebraska, North Dakota, and Kansas all experienced accelerating unemployment, but this may have been driven by idiosyncratic regional rather than national drivers.

6. How well does the LMSI fare against other indicators?

To test the effectiveness of the LMSI relative to other recession indicators, we calculate the Area Under Receiver Opearting Characteristic (AUROC) for our monthly LMSI, the Sahm-rule, the minimum indicator described in Michaillat and Saez (2025) denoted simply as the MS indicator, and the inversion of the Treasury term spread between the yield of a 10-year T-Bond and the 3-month T-Bill (Bauer and Mertens, 2018). We evaluate both the real-time and the predictive performance of our indicator.

For real-time recession detection, the LMSI fares well against most indicators with an AUROC of 0.87 out of a maximum value of 1, compared to an AUROC of 0.51 and 0.86 for the yield curve inversion and national Sahm-rule, respectively. However, the MS indicator is the most successful for real-time detection with an AUROC of 0.93.

In order to evaluate the recession predictive power of the LMSI with respect to these alternative indicators, we calculate the AUROC from 1-month up to 24-months in the future. The result are displayed in Figure 11.

Panel (a) of Figure 11 shows that our indicator performs very similarly to the MS indicator in predicting recessions, with the MS indicator slightly edging the LMSI at shorter horizons and the LMSI narrowly beating the MS indicator at longer horizons, but these

Figure 11: *The classification ability of alternative indicators against the LMSI*

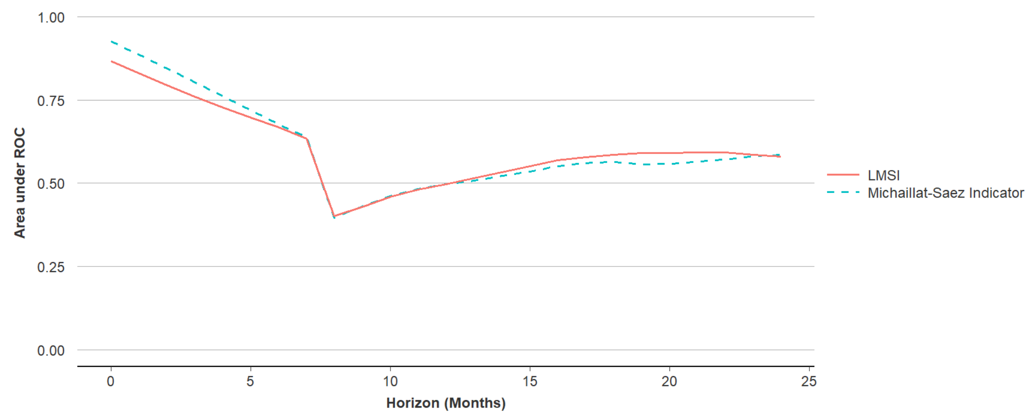

(a) *MS Indicator*

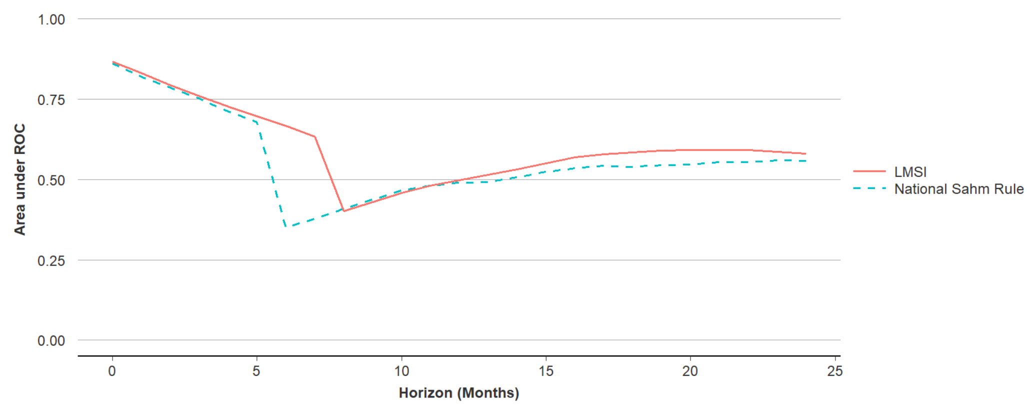

(b) *National Sahm Rule*

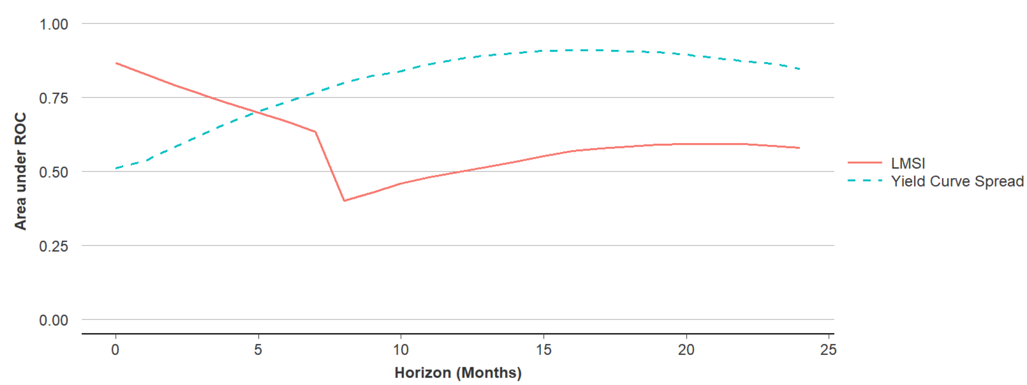

(c) *Yield Curve Inversion*

Notes: Each figure calculates the AUROC when the corresponding indices are used to predict 1 to 24 periods into the future. See text.

differences are de minimis. We should note that when the AUROC is below 0.5, then there is a negative relationship between the indicator and the probability of recession. The comparison with the Sahm-rule shown in panel (b) of Figure 11, follows a similar story, with some more noticeable differences around the mid-year mark, but again the differences are very small.

The pattern is more interesting when comparing the inversion of the yield curve with the LMSI as shown in panel (c) of Figure 11. At horizons below 6-months, the LMSI performs considerably better, but that pattern shifts after the mid-year mark. In fact, as has been previously documented (see, e.g. Berge and Jordà, 2011; Bauer and Mertens, 2018), the inversion of the yield curve is the best predictor at horizons of about 16-18 months out. Going back to the July 2024 episode, we note that the yield curve last inverted in October 2022, 19 months before the Sahm-rule was triggered, and when the LMSI peaked at 25 states. This episode neatly showcases the value of each of these indicators, with the yield curve being preferable at longer horizons, and the Sahm-rule and LMSI preferable at shorter horizons.

Based on this observation we then ask what is the best way to combine these indicators to obtain the best prediction at each of the following 1-24 months ahead. We do this by including all three indicators into a logit model where the left-hand side variable (the NBER recession indicator) is shifted forward in time from 1 to 24 periods, thus resulting in the estimation of 24 models. At the short end, we find that this model has an AUROC of 0.96 for the first month, higher than using any single indicator alone. Figure 12 shows the AUROC of the combined indicator over the 24-month horizon, revealing that combining these indicators results in an AUROC that never falls below 0.88 across two years of recession prediction.

Recall that the LMSI is meant to capture stress in the labor market and that it tends to lag the end of recessions precisely because labor markets tend to recover more slowly than output. We thus present our indicator not just as a recession indicator, but instead as a tool to more accurately describe the geographic heterogeneity of labor market conditions

Figure 12: *AUROC of All Indicators*

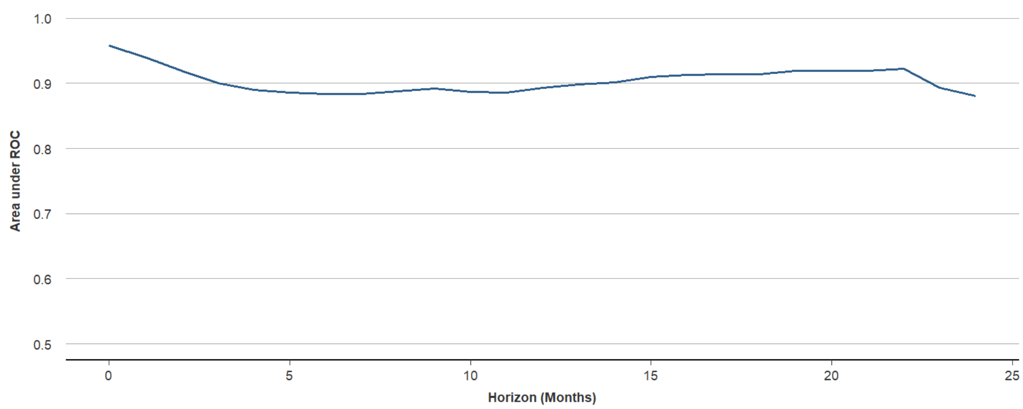

Notes: The line plots the AUROC for each of the 24 logit models that use the MS indicator, the Sahm-rule, and the LMSI as regressors. See text.

across the country. Although it has attractive properties for recession detection, we find it is more useful as a geographical lens on the depth and breadth of labor market disruptions. In conjunction with comparable indicators such as the Sahm-rule and MS Indicator, the LMSI can provide greater context behind indicator readings. This is most clearly shown in the context of the July 2024 episode. By examining the geography of labor market stress that produced the Sahm-rule trigger, we were able to point out the limited breadth of labor market disruption.

7. Conclusion

This paper introduces a new Labor Market Stress Indicator that brings geographic information to the center of recession monitoring. Using a long historical panel of state-level unemployment insurance claims and payroll data, we construct claims-based unemployment rates, map them into fitted state unemployment series, and extend the Sahm (2019) rule to the state level. Our baseline monthly LMSI counts how many states experience "accelerating" unemployment at any point in time, and a complementary labor-force–weighted version measures the share of the national labor force residing in those states. Together, these measures provide a parsimonious summary of the breadth and depth of labor-market stress across the United States.

Several facts emerge from our analysis. First, the monthly LMSI aligns closely with NBER business cycle dates: whenever roughly 30 or more states experience accelerating unemployment for the first time in an episode, the national economy has been in recession, and in those periods roughly three-quarters of the labor force resides in such states. By contrast, peaks below this range—including the July 2024 episode, when 25 states and about 47 percent of the labor force were in the accelerating bin—have not historically coincided with the onset of recessions. These patterns suggest a dual-threshold characterization of recession risk in terms of both geographic dispersion and labor-force coverage, and they highlight the value of looking beyond national aggregates when interpreting signals such as the Sahm-rule.

Second, we show that the LMSI performs competitively as a recession indicator relative to established benchmarks. In real time, the LMSI attains an AUROC that is comparable to the national Sahm-rule and the Michaillat–Saez minimum indicator, and higher than that of a standard term-spread-based rule. However, at forecast horizons of one to two years its performance is similar to the minimum indicator and weaker than the yield curve inversion metric. A simple logit specification that combines the LMSI with the Sahm-rule, the MS indicator, and the yield curve inversion delivers further gains in classification accuracy. This evidence underscores that geographically disaggregated labor-market information contains independent predictive content that is not fully captured by national unemployment or financial variables, and that recession monitoring can benefit from combining indicators with different horizons and information sets.

Third, the LMSI offers a useful lens on regional heterogeneity and exposure to aggregate downturns. We document systematic differences across states in how often they appear among the first 30 accelerating states when the lower threshold is crossed, identifying a set of "bellwether" states that consistently lead national recessions and a group of states—often in resource-intensive regions—that are more insulated from national labor-market fluctuations. This cross-sectional structure helps interpret borderline episodes such as

July-2024, when the national Sahm-rule briefly triggered while many historically bellwether states remained outside the accelerating bin and stress was concentrated in states that are typically less synchronized with the national business cycle. More generally, the LMSI can be used to study the spatial diffusion of labor-market stress and to evaluate how sectoral composition, industrial structure, and policy differences shape states' cyclical sensitivities.

Finally, we construct a weekly version of the LMSI based on raw state-level claims data. Although its shorter sample limits its usefulness for historical business-cycle analysis, the weekly LMSI tracks past recessions closely and provides a more timely read on evolving labor-market conditions. Because it relies on administrative claims data rather than on official unemployment rates, the weekly indicator can continue to be updated when statistical releases are delayed or disrupted, and it remains narrowly focused on labor-market stress in contrast to broader composite indexes. This high-frequency extension illustrates how the LMSI framework can be adapted to different data environments while preserving its core geographic interpretation.

References

Bauer, Michael D, and Thomas M Mertens. 2018. Economic Forecasts with the Yield Curve. *FRBSF Economic Letter* 2018(07).

Berge, Travis J, and Òscar Jordà. 2011. Evaluating the classification of economic activity into recessions and expansions. *American Economic Journal: Macroeconomics* 3(2): 246–277.

Blanchard, Olivier Jean, Lawrence F Katz, Robert E Hall, and Barry Eichengreen. 1992. Regional evolutions. *Brookings papers on economic activity* 1992(1): 1–75.

Bureau of Labor Statistics, U.S. Department of Labor. 2023. Most unemployed people in 2022 did not apply for unemployment insurance benefits. Technical Report, The Economics Daily.

Engstrom, Eric C, and Steven A Sharpe. 2019. The near-term forward yield spread as a leading indicator: A less distorted mirror. *Financial Analysts Journal* 75(4): 37–49.

Estrella, Arturo, and Frederic S Mishkin. 1998. Predicting US recessions: Financial variables as leading indicators. *Review of Economics and Statistics* 80(1): 45–61.

Fieldhouse, Andrew, Sean Howard, Christoffer Koch, and David Munro. 2022. A New Claims-Based Unemployment Dataset: Application to Postwar Recoveries Across US States. Technical Report, International Monetary Fund.

Fieldhouse, Andrew, Sean Howard, Christoffer Koch, and David Munro. 2024. The Emergence of a Uniform Business Cycle in the United States: Evidence from New Claims-Based Unemployment Data. *Brookings Papers on Economic Activity* 2004(1): 265–319.

Graetz, Georg, and Guy Michaels. 2017. Is modern technology responsible for jobless recoveries? *American Economic Review* 107(5): 168–173.

Jaimovich, Nir, and Henry E Siu. 2020. Job polarization and jobless recoveries. *Review of Economics and Statistics* 102(1): 129–147.

Keil, Manfred, Edward E. Leamer, and Yao Li. 2023. An investigation into the probability that this is the last year of the economic expansion. *Journal of Forecasting* 42(5): 1228–1244.

Leamer, Edward E. 2007. Housing is the business cycle. *Proceedings - Economic Policy Symposium - Jackson Hole* 149–233.

Leamer, Edward E. 2022. A New Way of Forecasting Recessions. NBER Working Papers 30247, National Bureau of Economic Research, Inc.

Lewis, Daniel J, Karel Mertens, James H Stock, and Mihir Trivedi. 2022. Measuring real activity using a weekly economic index. *Journal of Applied Econometrics* 37(4): 667–687.

Michaillat, Pascal, and Emmanuel Saez. 2025. Has the recession started? *Oxford Bulletin of Economics and Statistics* .

O'Trakoun, John, and Adam Scavette. 2025. A better Sahm rule? Introducing the SOS recession indicator. *Economics Letters* 247: 112175.

Peirce, Charles S. 1884. The numerical measure of the success of predictions. *Science* (93): 453–454.

Peterson, Wesley W, and Theodore G Birdsall. 1953. The theory of signal detectability. Part 1. the general theory. Technical Report 13, University of Michigan, Department of Electrical Engineering, Electronic Defense Group.

Rudebusch, Glenn D, and John C Williams. 2009. Forecasting recessions: The puzzle of the enduring power of the yield curve. *Journal of Business & Economic Statistics* 27(4): 492–503.

Russ, Katheryn N, Jay C Shambaugh, and Sanjay R Singh. 2024. Currency Areas, Labor Markets, and Regional Cyclical Sensitivity. *IMF Economic Review* 72(1): 152–195.

Sahm, Claudia. 2019. Direct stimulus payments to individuals. *Recession ready: Fiscal policies to stabilize the American economy* 67–92.

Shimer, Robert. 2012. Wage rigidities and jobless recoveries. *Journal of Monetary Economics* 59: S65–S77.

Youden, William J. 1950. Index for rating diagnostic tests. *Cancer* 3(1): 32–35.

A. Appendix

Figure A.1: *Monthly LMSI based on fitted state unemployment rates*

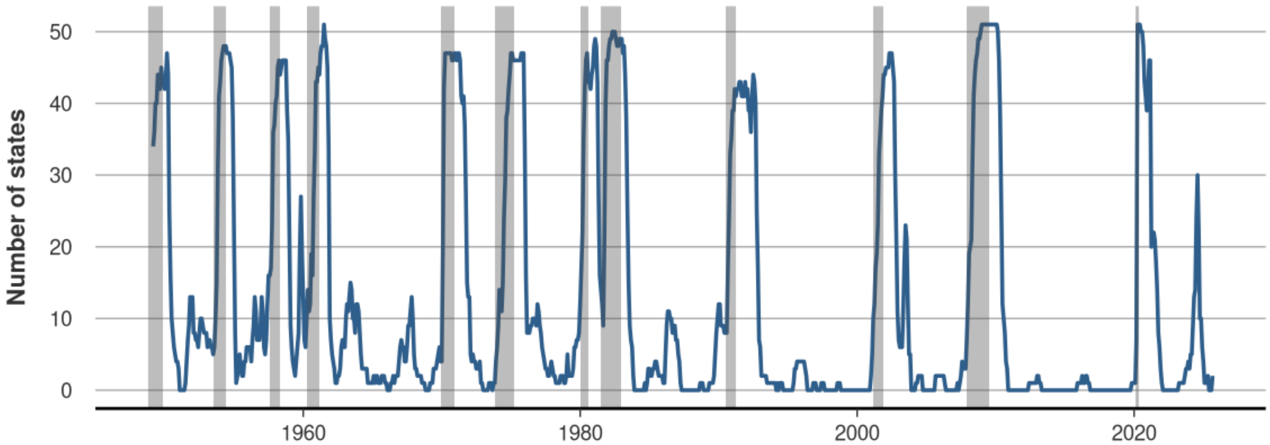

Notes: Number of states with accelerating unemployment based on the Sahm-rule and the state-level unemployment rate series constructed by fitting Equation 2 on the 1976-2025 sample and then using the fitted values for the entire 1949-2025 sample. Shaded regions correspond to NBER recession periods. See text.

Table A.1: *Claims-Based Unemployment Fitting Regression Results*

State	Coefficient	R^2	Residuals SD
AL	1.44	.91	.30
AK	.82	.87	.37
AZ	.90	.87	.37
AR	.85	.84	.40
CA	.88	.89	.33
CO	.88	.81	.43
CT	1.25	.87	.36
DE	1.49	.81	.43
DC	.52	.73	.52
FL	.92	.93	.27
GA	.21	.80	.44
HI	1.31	.81	.44
ID	.75	.92	.27
IL	.47	.90	.31
IN	1.46	.94	.24
IA	1.30	.84	.40
KS	.65	.75	.50
KY	.90	.80	.45
LA	1.39	.78	.47
ME	.85	.87	.36
MD	.65	.86	.38
MA	1.68	.87	.36
MI	.87	.92	.28
MN	.23	.89	.33
MS	1.69	.87	.36
MO	.57	.92	.29
MT	1.35	.81	.44
NE	.48	.67	.58
NV	1.15	.85	.39
NH	1.10	.78	.47
NJ	1.77	.90	.32
NM	.97	.64	.60
NY	.97	.89	.34
NC	.57	.83	.42
ND	.87	.72	.53
OH	1.30	.95	.23
OK	1.46	.79	.46
OR	1.09	.92	.28
PA	1.04	.91	.31
RI	.35	.75	.50
SC	1.09	.90	.32
SD	.31	.76	.49
TN	.58	.91	.29
TX	1.20	.78	.46
UT	.76	.84	.40
VT	1.19	.87	.36
VA	.10	.92	.28
WA	.61	.92	.27
WV	2.15	.82	.43
WI	.99	.93	.27
WY	.99	.73	.52