EconomicPerspectives

Economic Perspectives, No. 1, April 2026

Crossref

The U.S. Labor Market: A Long-Run Perspective

By Scott A. Brave, R. Andrew Butters, Ross Cole, Max Gillet

Mathematical and Quantitative Methods | Macroeconomics and Monetary Economics | Labor and Demographic Economics

As the economy recovered from the pandemic recession of 2020, the labor market stood center stage. Increased interest around job openings, separations, and quits data added kindling to the fire of discussion about the Great Resignation (Gittleman, 2022), and labor market tightness provided the foundation for a plausible supply-side cost-push explanation of the coinciding sharp increase in inflation (Benigno and Eggertsson, 2023).

When trying to understand the latest business cycle and judging what that means for the labor market going forward, it would be easy to say this time is different (Romer and Romer, 2018). Despite being easy (and potentially justified *this time*), being too dismissive of the lessons of past business cycles is often counterproductive (Stock and Watson, 2012). With that in mind, we lean into those lessons in this *Economic Perspectives* article, paying specific attention to what they might mean for the labor market at this moment *and* in the future.

To accomplish this, we rely on a mix of old and new tools. The former consists of what has quickly become a go-to empirical macro model—the vector autoregression, or VAR (Sims, 1980)—while the latter consists of recently developed priors that are well suited for capturing theoretical long-run relationships in macroeconomic time series. By combining the two into a Bayesian VAR (BVAR) framework, we are able to leverage both economic theory and past business cycle experiences to analyze both movements along and shifts in the Beveridge curve, which refers to the inverse relationship between the unemployment rate and the job vacancy rate (Lubik, 2021).

Our motivation in doing so is to provide important context for the labor market recovery following the pandemic recession. Here, we are following the old maxim attributed to Mark Twain that while history doesn't repeat itself, it often rhymes. There is no close precedent in the last 70 years of U.S. economic history for the magnitude of the recession that unfolded in the first half of 2020; and even six years out from those events, the continued economic expansion has several times since surprised many that expected it would lose steam much earlier. But that does not mean that past business cycle experiences are not relevant for the current labor market situation.

Figure 1 provides some context for this view, highlighting recent Beveridge curve movements in comparison with a select subset of economic expansions (see appendix 1 for data sources). While no single episode matches perfectly with the recent experience, all of them bear a resemblance. The early recovery from the pandemic recession occurred within a context of high unemployment and low job vacancies similar to that of the recoveries that followed the 1973–75 and 1981–82 recessions. More recently, the sharp decline in job vacancies from an elevated level with only a limited increase in unemployment has instead more closely resembled the latter stages of the expansions that followed the 1960–61 and 2007–09 recessions.

1. Select economic expansions and the Beveridge curve

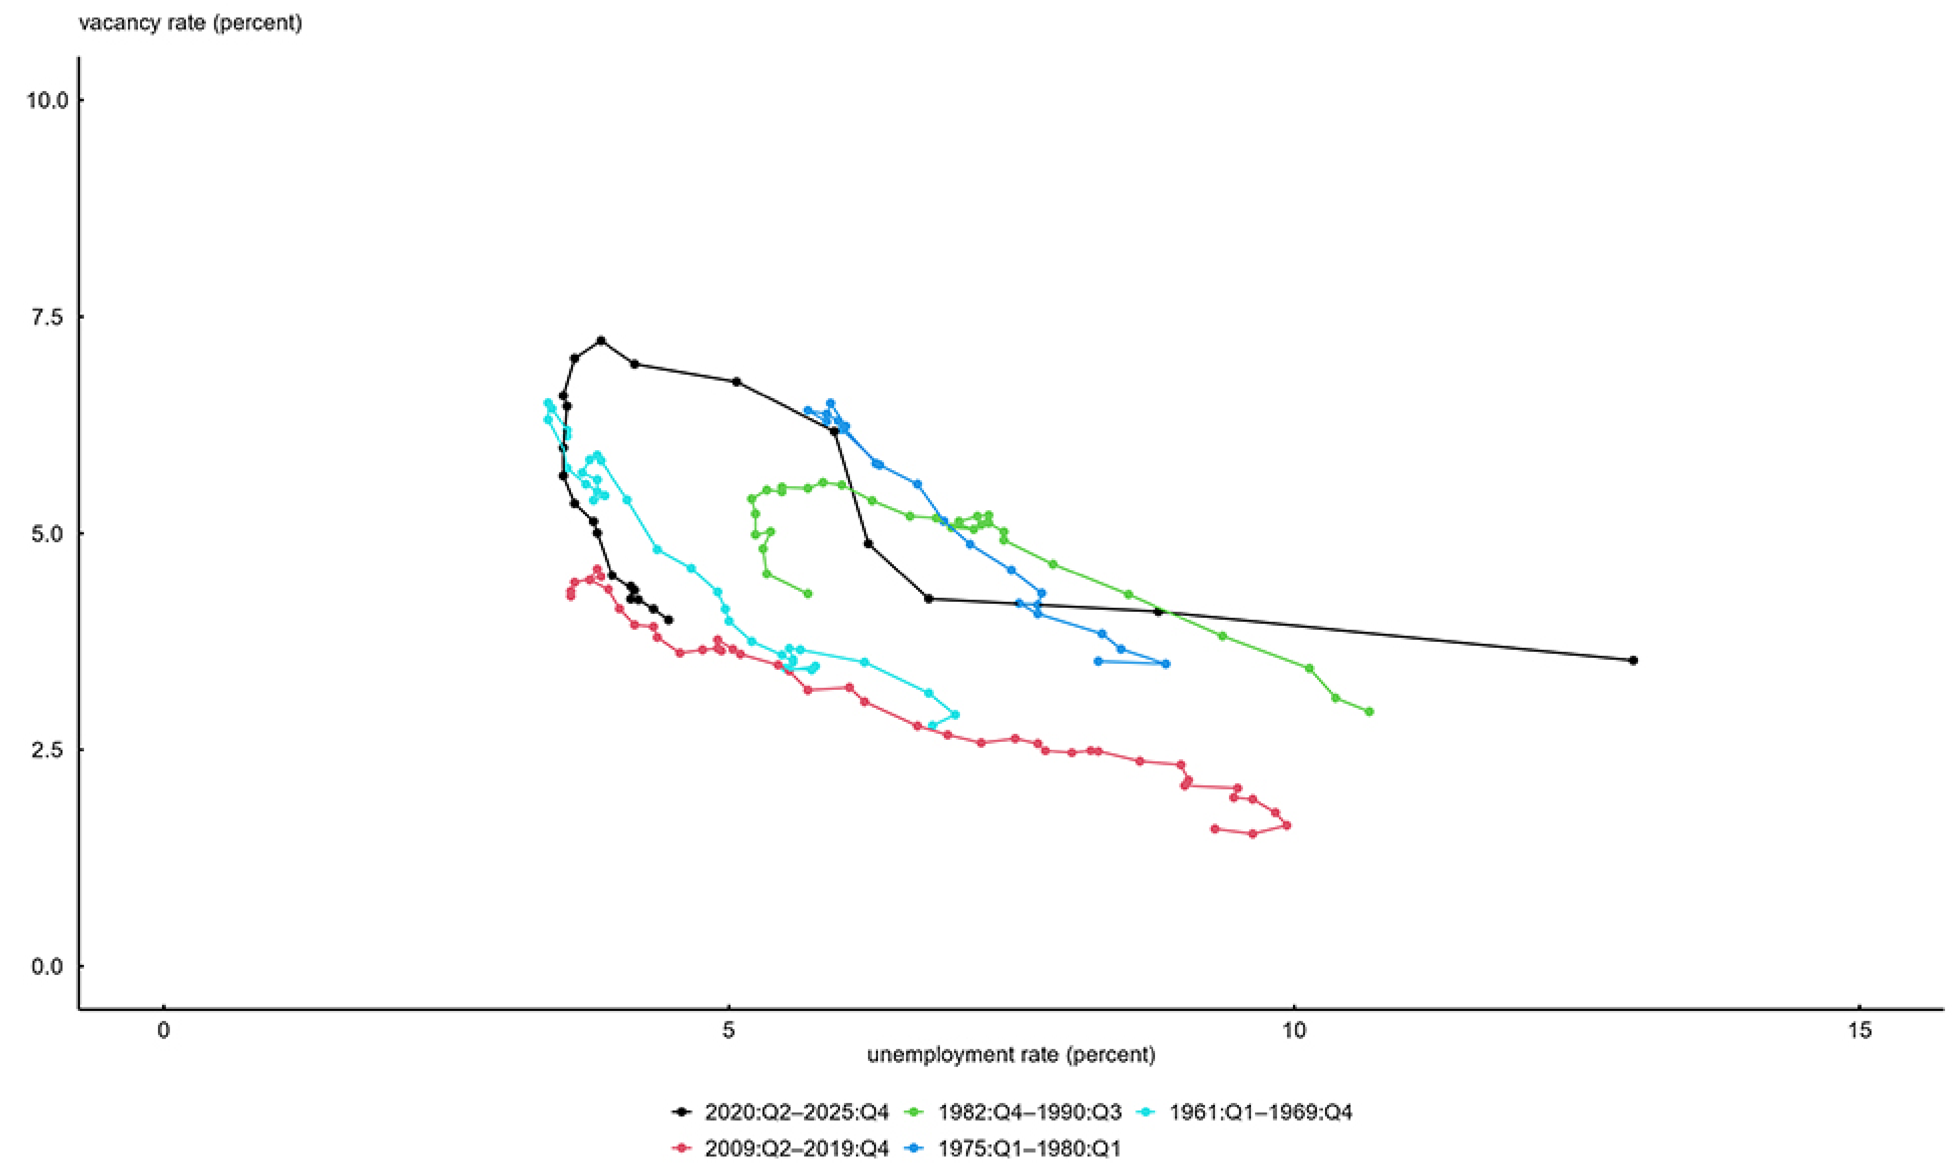

Note: The figure shows historical values of the unemployment (*U*) and job vacancy (*V*) rates during a subset of economic expansions as defined by the National Bureau of Economic Research.
Sources: Authors' calculations based on data from Haver Analytics.

The question of the hour is, of course, where are we likely to go from here? To answer this question, we use our BVAR to forecast the near-term trajectory of the U.S. labor market and conduct scenario analyses centered around that forecast. Lending some credence to this exercise in our view is the fact that our model's forecasts are disciplined both by priors reflecting the long-run relationships between labor market variables posited by economic theory and roughly 70 years of data capturing a multitude of various historical labor market experiences and movements in the Beveridge curve (see appendix 2 for technical details). Before we proceed, we want to emphasize that the views expressed in this article are our own and do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.

Our BVAR's baseline forecasts suggest that the vacancy rate is likely to stabilize at just under 4%, with the unemployment rate edging higher to around 4.5%. This represents a relatively flat forecast for the unemployment rate, which historically has tended to rise sharply once it turns upward in contrast with its more recent slow upward drift.

To provide greater context for the recent behavior of the unemployment rate, we use the method of sign restrictions (Uhlig, 2017) to identify structural shocks in our BVAR (as described in appendix 3) capturing shifts in the Beveridge curve versus movements along the curve. Shifts in the Beveridge curve result from structural changes in match efficiency (or the process of matching available workers with open jobs) and labor supply (or the size of the available labor force), whereas movements along the curve reflect cyclical changes to the wage bargaining process or aggregate demand and supply over the business cycle.

The recent cooling of the labor market according to our BVAR primarily reflects the receding tailwinds from past cyclical demand-side shocks as well as structural supply-side shocks—in particular, aggregate demand and match efficiency shocks. Their reduced impact on the unemployment and vacancy rates, however, has been partially offset at various times by increasing tailwinds coming from other cyclical and structural shocks—in particular, wage bargaining and labor supply shocks.

Like any model, however, our BVAR is only a *guide* (and not a *crystal ball*). It only allows us to evaluate the likelihood of potential outcomes based on past experience and our prior knowledge. Therefore, to highlight specific risks to the labor market outlook, we also conduct scenario analyses. These scenarios condition the BVAR's forecasts on an alternative future path for a particular variable of interest; therefore, these scenarios capture an array of potential alternative shocks that could be relevant in the medium term.

For example, in one such scenario we condition on a future path for initial unemployment insurance (UI) claims that rises faster and remains more persistently elevated than in the BVAR's baseline forecast. Unexpected spikes in initial UI claims are often seen as harbingers of further labor market weakening. This scenario is, thus, consistent with an unexpected decline in the demand for labor relative to the baseline forecast.

Similarly, we also consider scenarios where labor force participation or household wealth declines more than in our baseline forecast, capturing adverse shocks to labor supply (e.g., a lower-than-expected labor force participation rate) or aggregate demand (e.g., a potential decline in home and/or equity market prices from current levels). While on the positive side, we consider a scenario where gross private domestic investment growth significantly accelerates (e.g., from broader adoption of artificial intelligence (AI) and/or recent expansionary fiscal policy measures to promote domestic investment).

Of the four scenarios that we consider, the unemployment rate is higher than in our baseline forecast in three of them. However, only two of those cases result in broader forecasts that resemble recessionary dynamics (i.e., the scenarios capturing a decline in labor demand and household wealth). Across all four scenarios, the unemployment rate ranges from a low of about 4% to a high of about 4.75% over the 12-quarter projection period that we consider.

There are, of course, many more scenarios that one *could* consider than we *do* consider here. For this reason, we embed our model into an open-source application that can be tailored as needed to produce other alternative scenarios in the future (as described in appendix 4) and provide it in a downloadable form for public use.

A BVAR with long-run priors

The starting point for our analysis is a vector autoregression. Long a workhorse in empirical macroeconomics, VARs are statistical models that use historical relationships to make informed inferences about future events. Their flexibility in this regard is considerable and has been shown under certain conditions to approximate the complex dynamic relationships between households, firms, and policymakers in modern dynamic stochastic general equilibrium (DSGE) models (Giacomini, 2013).

Our VAR relies on 18 quarterly time series to summarize historical labor market relationships in the U.S. over a 70-year history. Moreover, most of the metrics commonly used to judge labor market activity can be constructed from combinations of these measures. All original data sources used in this article are available in figure A1 of appendix 1. Figure 2 shows the time series of each indicator with shaded periods corresponding to recessions for the U.S. as defined by the National Bureau of Economic Research (NBER).

2. Labor market indicators

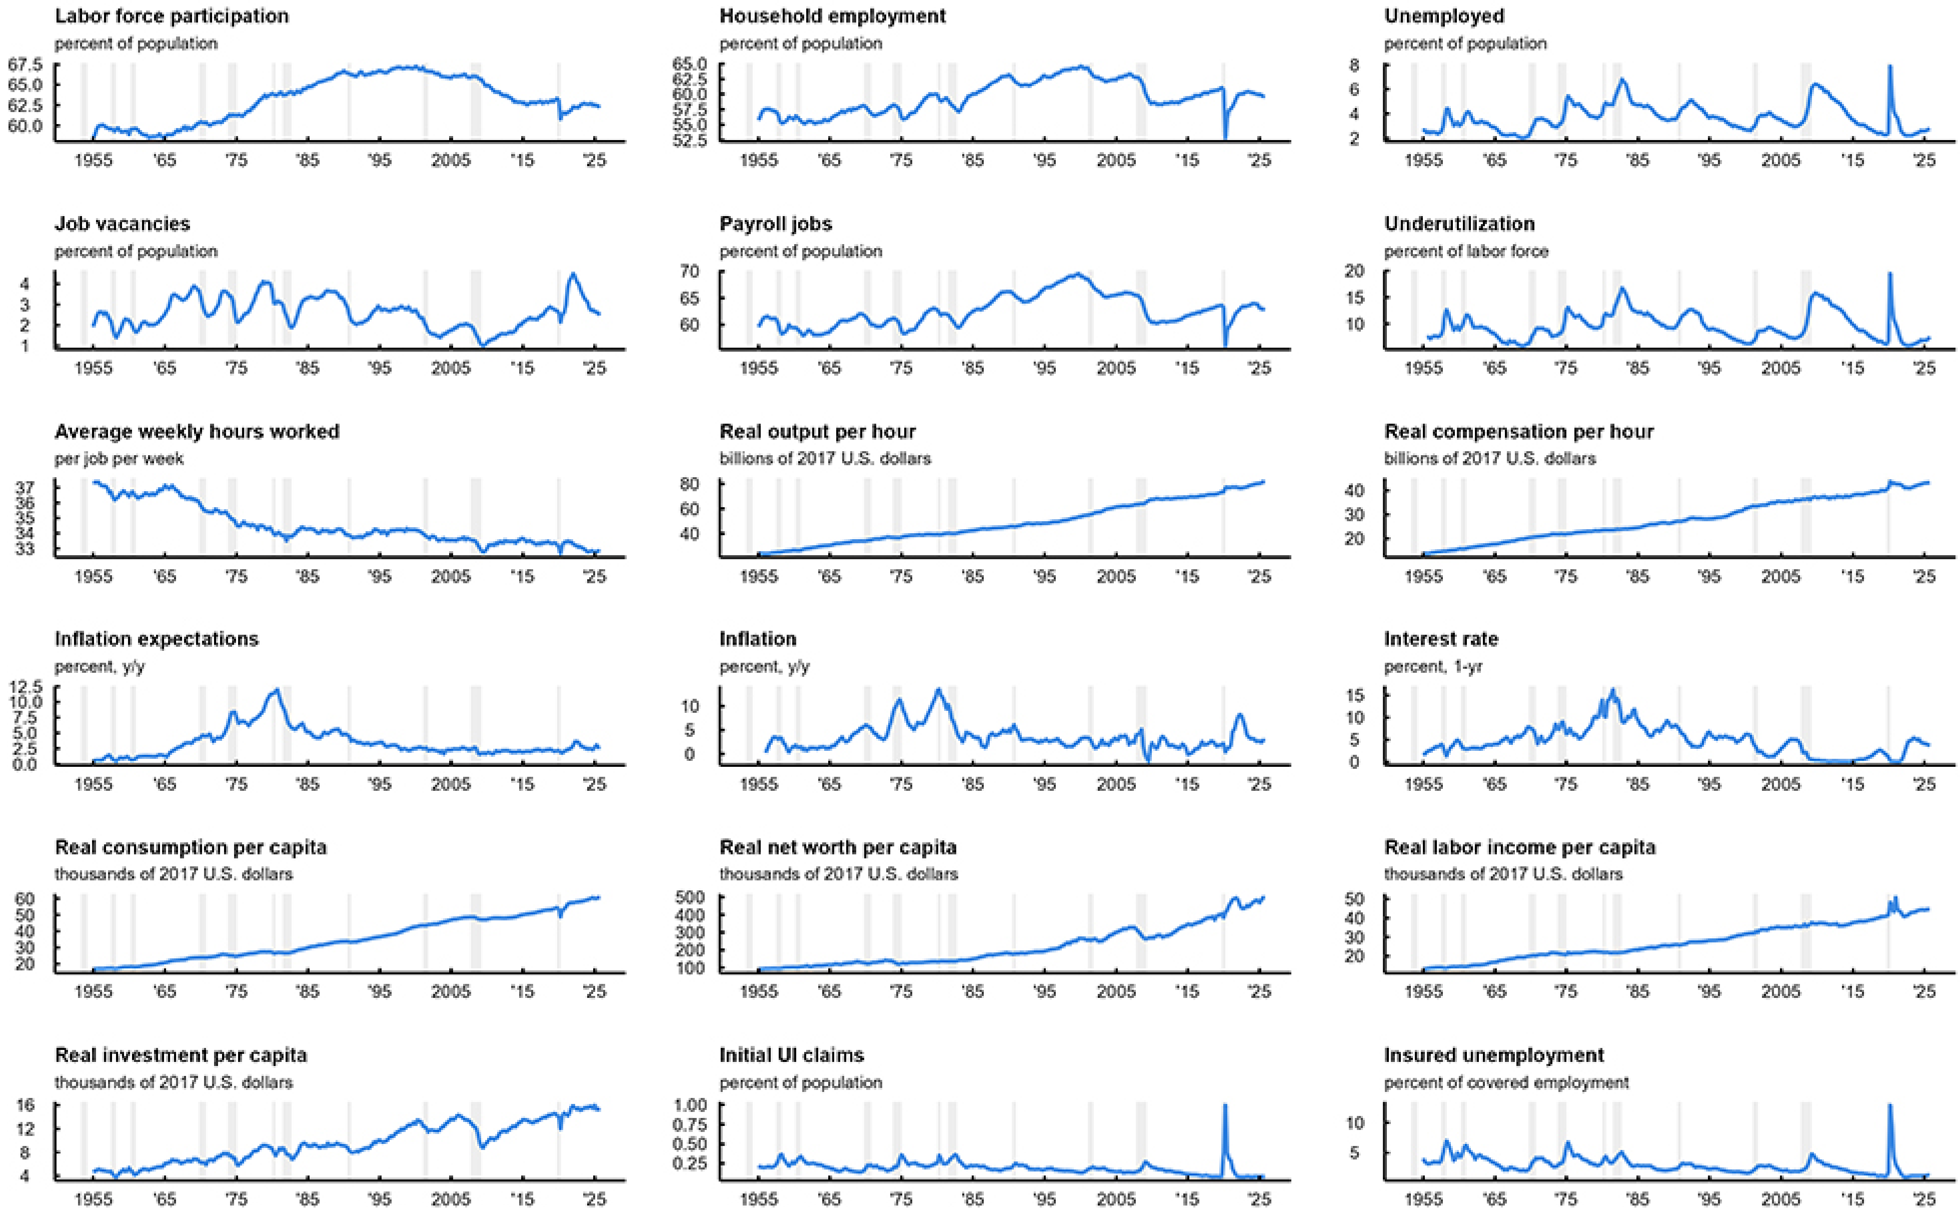

Notes: The figure plots the 18 quarterly time series included in our Bayesian vector autoregression (BVAR) over the sample period 1955–2025. Population is the noninstitutional civilian population aged 16 and older. UI stands for unemployment insurance. The shaded periods correspond to U.S. recessions as defined by the National Bureau of Economic Research.
Sources: Authors' calculations based on data from Haver Analytics.

Given the number of time series (and allowing for a modest number of lags in the VAR), the number of parameters in our model is quite large. *Bayesian shrinkage* methods, implemented via priors, are advantageous in such situations and form the basis of how we efficiently estimate our VAR. These methods have the added benefit of producing forecasts with out-of-sample properties that outperform other common statistical methods (Brave et al., 2019).

To guide the near-term dynamics of our BVAR, we rely upon the Minnesota prior conventions that are commonly used in macroeconomic forecasting (Carriero et al., 2015). Described in detail in the technical appendix covering our BVAR (see appendix 2), they act to tilt our near-term (i.e., one-quarter ahead) forecasts toward a "no-change" forecast. To this prior, we then add the long-run prior framework of Giannone et al. (2019) to discipline our BVAR's longer-run dynamics to match several cointegrating relationships posited by economic theory.

The cointegrating relationships that we impose in our prior (summarized in figure 3) between various labor market indicators capture shared trends shaping long-run growth in the labor force and labor productivity, as well as consumer prices, interest rates and inflation, and household wealth, income, and consumption, among others. The Giannone et al. (2019) long-run priors help to discipline the medium-run dynamics of our forecasts given these shared trends. For further details, see appendix 2.

3. Cointegrating relationships for the long-run priors

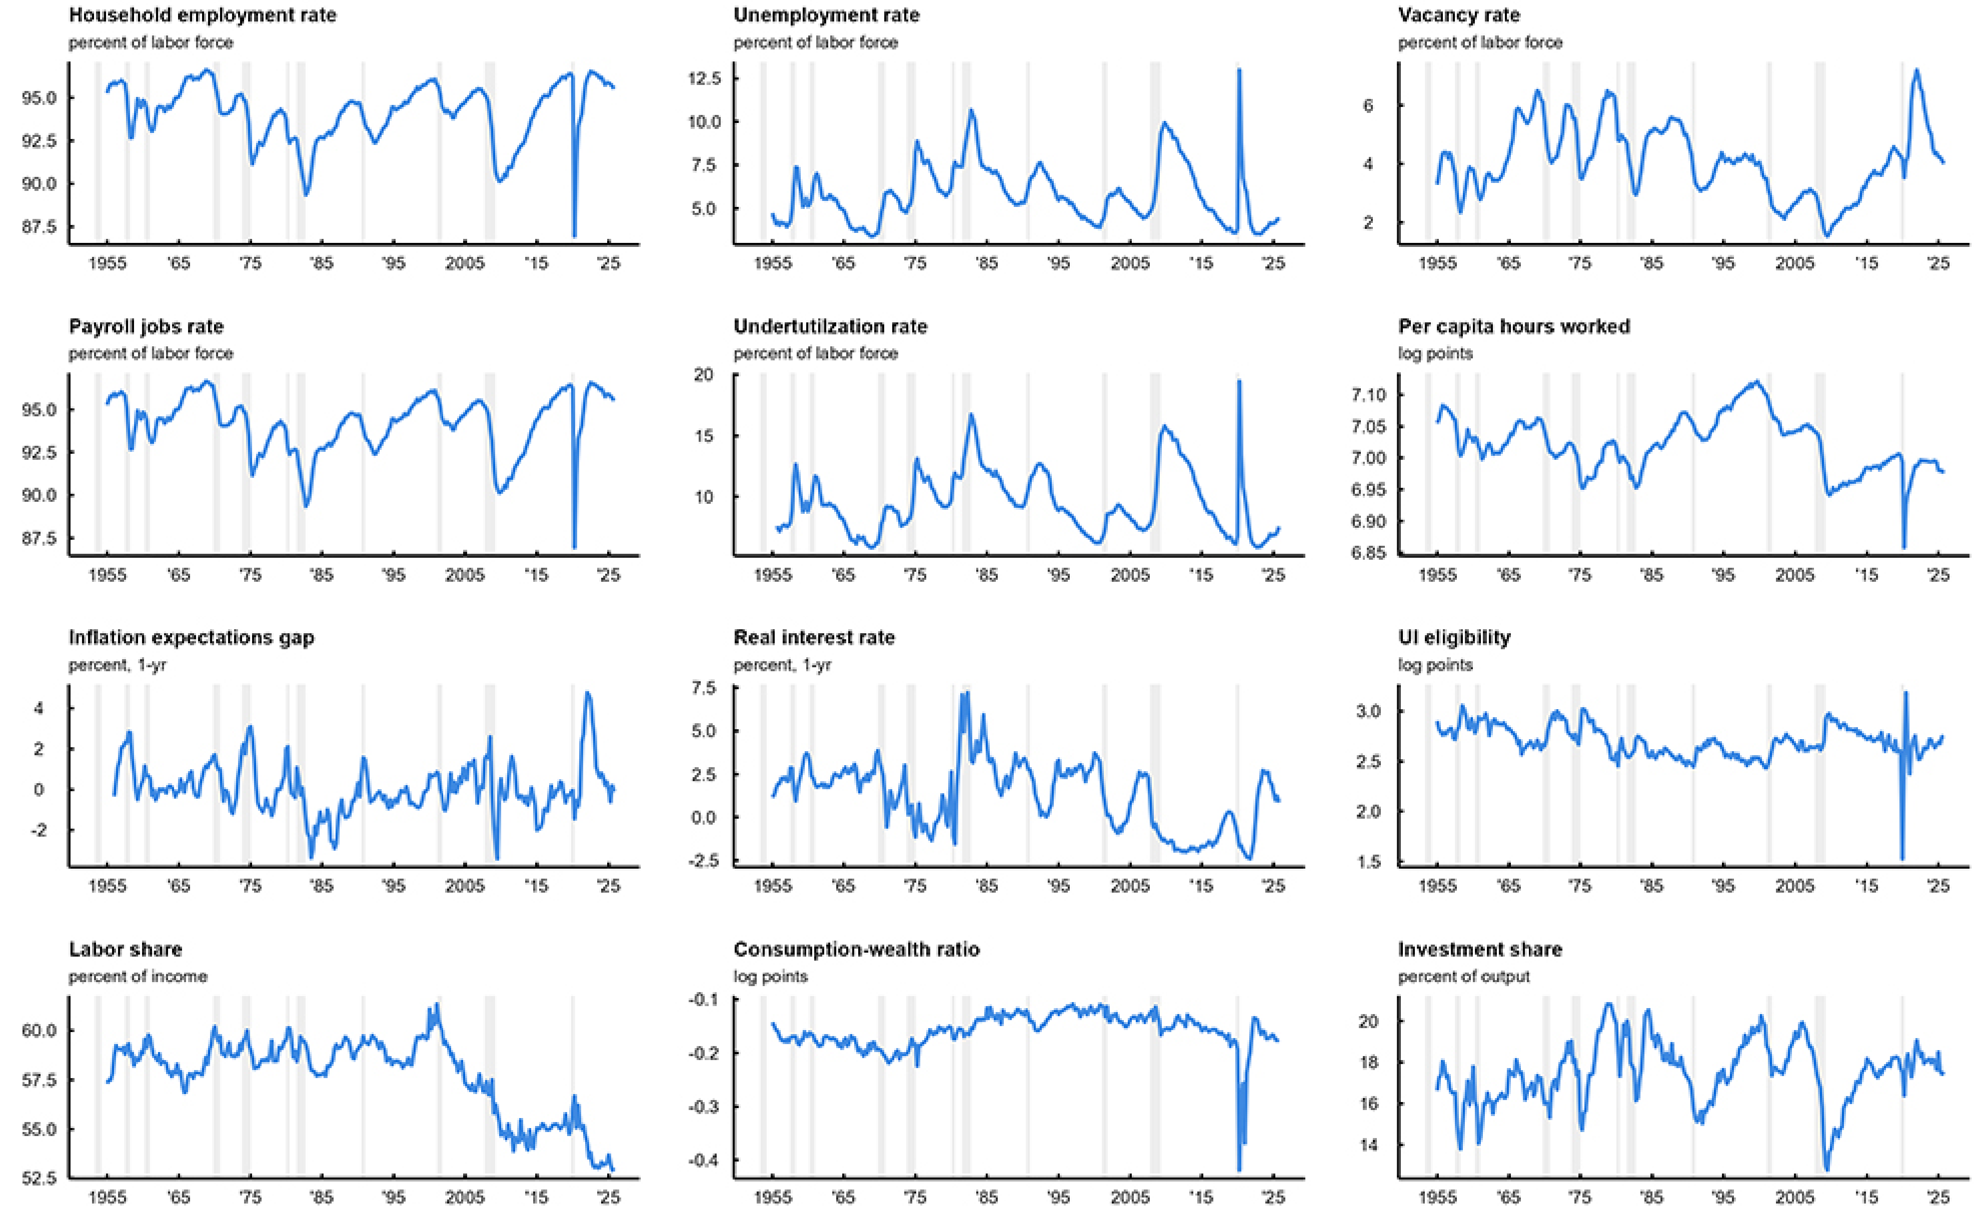

Notes: The figure plots the quarterly time series of cointegrating relationships specified in the long-run priors of our Bayesian vector autoregression (BVAR) over the sample period 1955–2025. UI stands for unemployment insurance. Some series have been converted from log to natural units for ease of viewing. The shaded periods correspond to U.S. recessions as defined by the National Bureau of Economic Research.
Sources: Authors' calculations based on data from Haver Analytics.

Our BVAR is estimated with five lags using Markov chain Monte Carlo (MCMC) methods as described in the documentation supporting the *R* package *BVAR* summarized in Kuschnig and Vashold (2021). To facilitate convergence, we estimate 24 Markov chains of length 125,000 draws with a Metropolis–Hastings step and burn-in (discarded sample) of 25,000 draws and a thin rate of 20 draws (i.e., only every 20th remaining draw is kept) for a grand total of 120,000 draws characterizing the model's posterior distribution combining the likelihood of the data with our priors.

We then construct unconditional forecasts over the coming 12 quarters by simulating the BVAR 120,000 times, each time drawing shocks from the model's posterior distribution that combines the likelihood of the observed data with our theory-based priors describing their long-run properties. We then characterize the degree of uncertainty around the BVAR's *median* forecasts with 70% prediction intervals.

For now, we hold off on discussing our BVAR's forecasts, although the reader may find them shown in figures 7–11 later. Instead, here we focus on the broad features of the forecast that are shaped by our choice of priors. For visual reference, see figure A3 in appendix 2, which shows the median forecasts from our baseline BVAR (Minnesota & LR prior) and compares them with median forecasts from an alternative BVAR (Minnesota & SOC/SUR prior).

The alternative projections in figures A5–A9 in appendix 2 use the *sum-of-coefficients* (SOC) and *single unit root* (SUR) priors. The Giannone et al. (2019) long-run (LR) prior framework generalizes these two alternative priors, tackling on a series-by-series basis the degree of integration and cointegration among our 18 time series that the alternatives address at the system level. Further details can be found in appendix 2.

The median forecasts from our baseline BVAR are often very similar to those obtained from this commonly used alternative BVAR in the near term, but a few important differences related to the long-run prior's treatment of trends and cointegration stand out. Each of these differences can be traced back to the influence of the assumed cointegrating relationships of our long-run prior, which create some additional reversion to the mean that the SOC/SUR priors do not. Some examples are as follows:

- The unemployment and underutilization rates are higher in our baseline, but the vacancy rate is lower, resulting in a lower jobs–workers gap.

- Consistent with a modestly weaker labor market, both employment-to-population and labor force participation measures are lower in our baseline, with initial unemployment insurance claims also higher.
- Growth in real output and hours worked are similar, but with our baseline exhibiting a little higher labor productivity and, thus, also a little higher inflation-adjusted compensation.
- Inflation expectations are very similar in both forecasts, but inflation is a little higher and nominal and real interest rates modestly lower in our baseline (reflecting weaker labor market conditions).
- Consumption is similar in both forecasts, but income and wealth are higher in our baseline forecast and investment is lower.

We take a closer look at some of the apparent effects of our long-run prior by conducting a scenario analysis, as shown in figures 7–11, alluded to earlier. But before describing these results, we first take a closer look at our baseline BVAR forecasts for unemployment, vacancies, and the Beveridge curve.

Labor market slack and supply versus demand influences

As seen in figure 4, our BVAR's median projections for the unemployment rate suggest little further increase in unemployment over the next few years. A careful examination of the historical time series for the unemployment rate demonstrates how unusual such an outcome would be: When the unemployment rate starts to rise, it generally does so fairly quickly and steeply. This is the behavior the Sahm rule[1] is meant to capture. The recent behavior of the unemployment rate, however, is a bit of an exception—with the rate rising right up to the threshold of the Sahm rule but not barreling past it.

4. Unemployment, recessions, and the Sahm rule

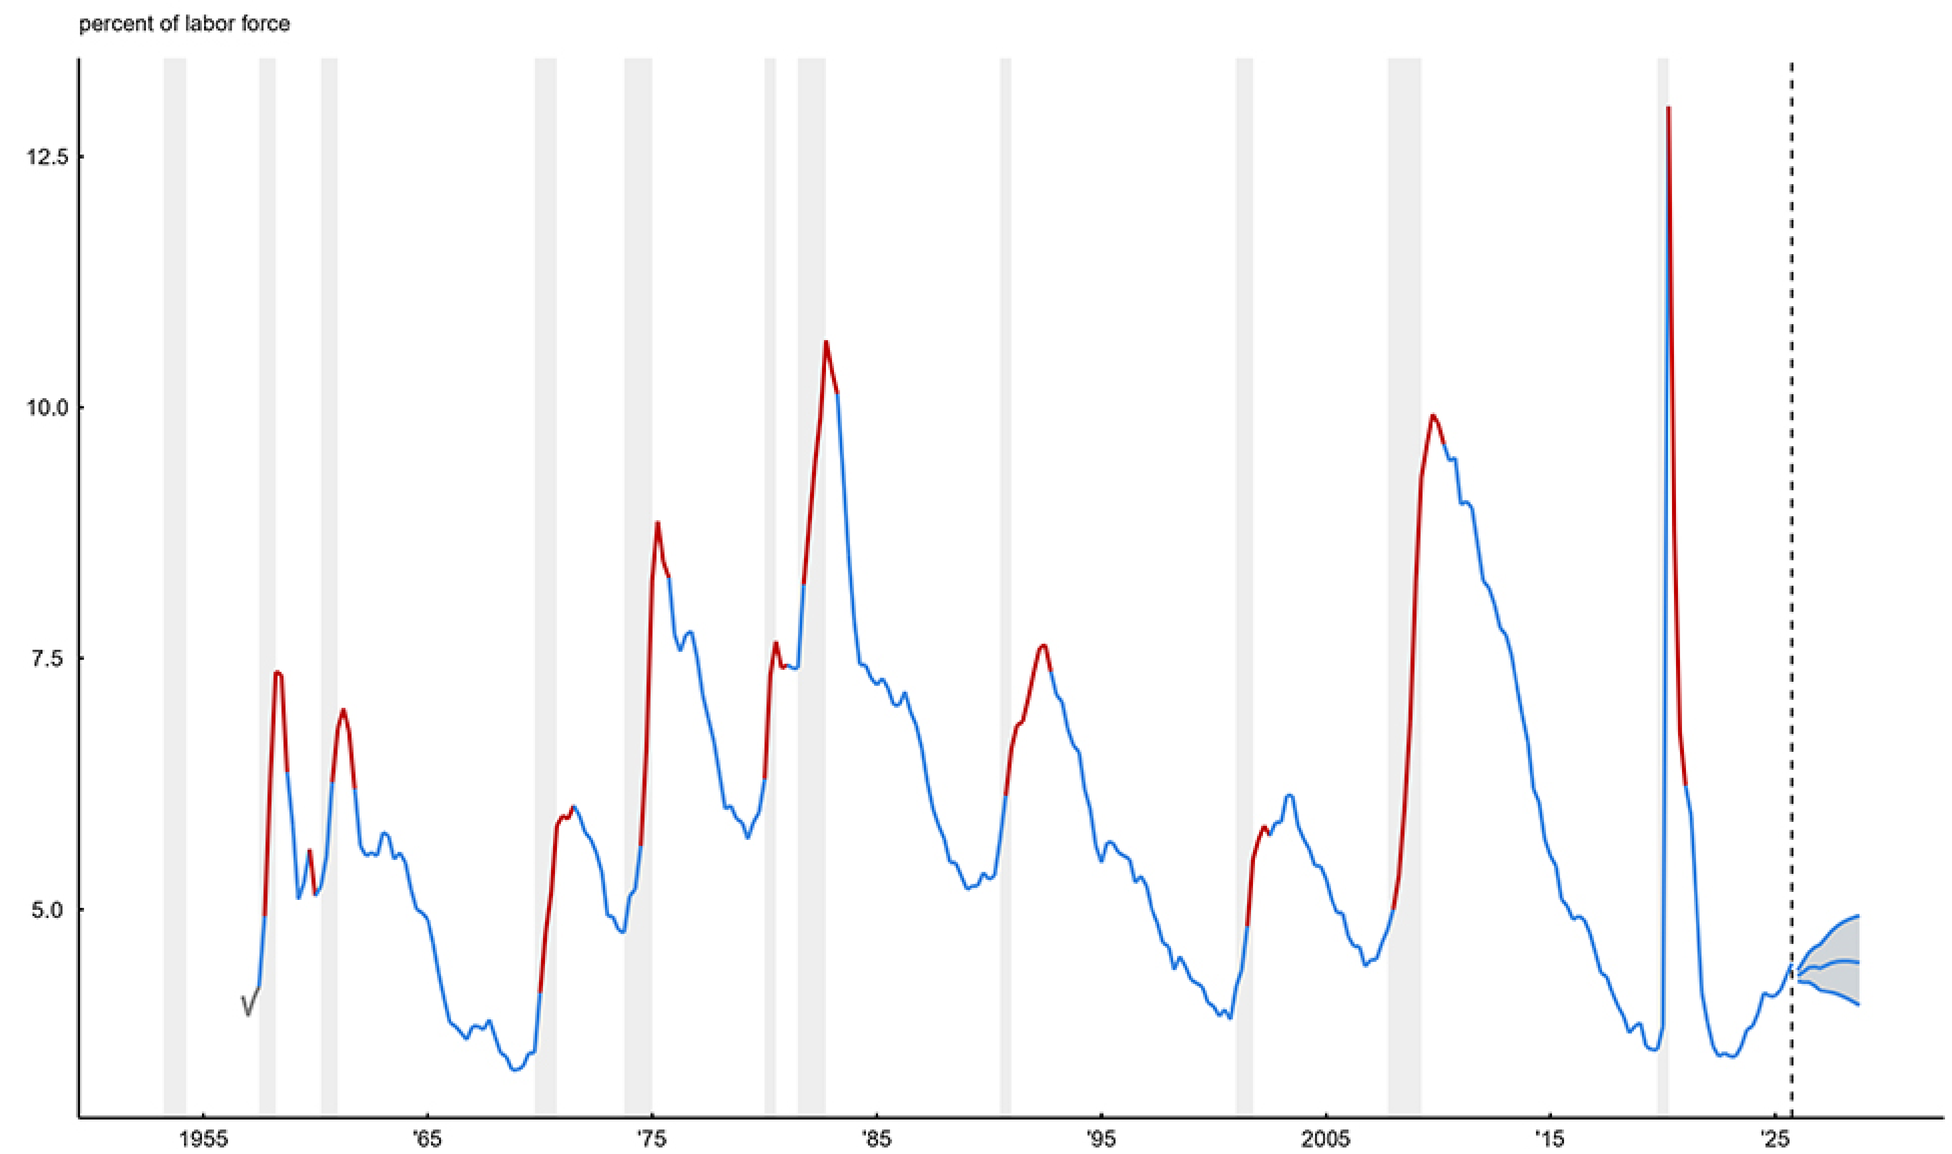

Notes: The figure shows historical values of the unemployment rate (solid blue line to the left of dashed vertical black line) as well as the median projections (middle solid blue line to the right of the dashed vertical black line) for it from our baseline Bayesian vector autoregression (BVAR). The shaded periods correspond to U.S. recessions as defined by the National Bureau of Economic Research, while the red solid lines highlight unemployment rate values for dates where the Sahm rule has been triggered (i.e., a 0.5 increase or more in the unemployment rate from its rolling four-quarter low). The solid lines surrounding the median projections are the upper and lower bounds of their 70% prediction interval (dark gray shading).
Sources: Authors' calculations based on data from Haver Analytics.

To get a better sense of the supply and demand forces driving the recent behavior of the unemployment rate and our BVAR's forecasts for it, we next conduct a simple out-of-sample forecast exercise: We back up the end of our data sample to the first quarter of 2024, 2023, or 2022, each time reestimating the model so that we can compare its forecasts over time to our current forecasts.

Figure 5 displays the forecasted Beveridge curve relating job vacancies (*V*) to the number of unemployed (*U*) as a share of the labor force. For comparison, we include in the figure along with the forecasts (red dots) the five lagged values that enter into the estimated model (black dots) and the subsequent 12 realized values (blue dots) as well as the remaining sample history (gray dots).

5. Forecasting the Beveridge curve

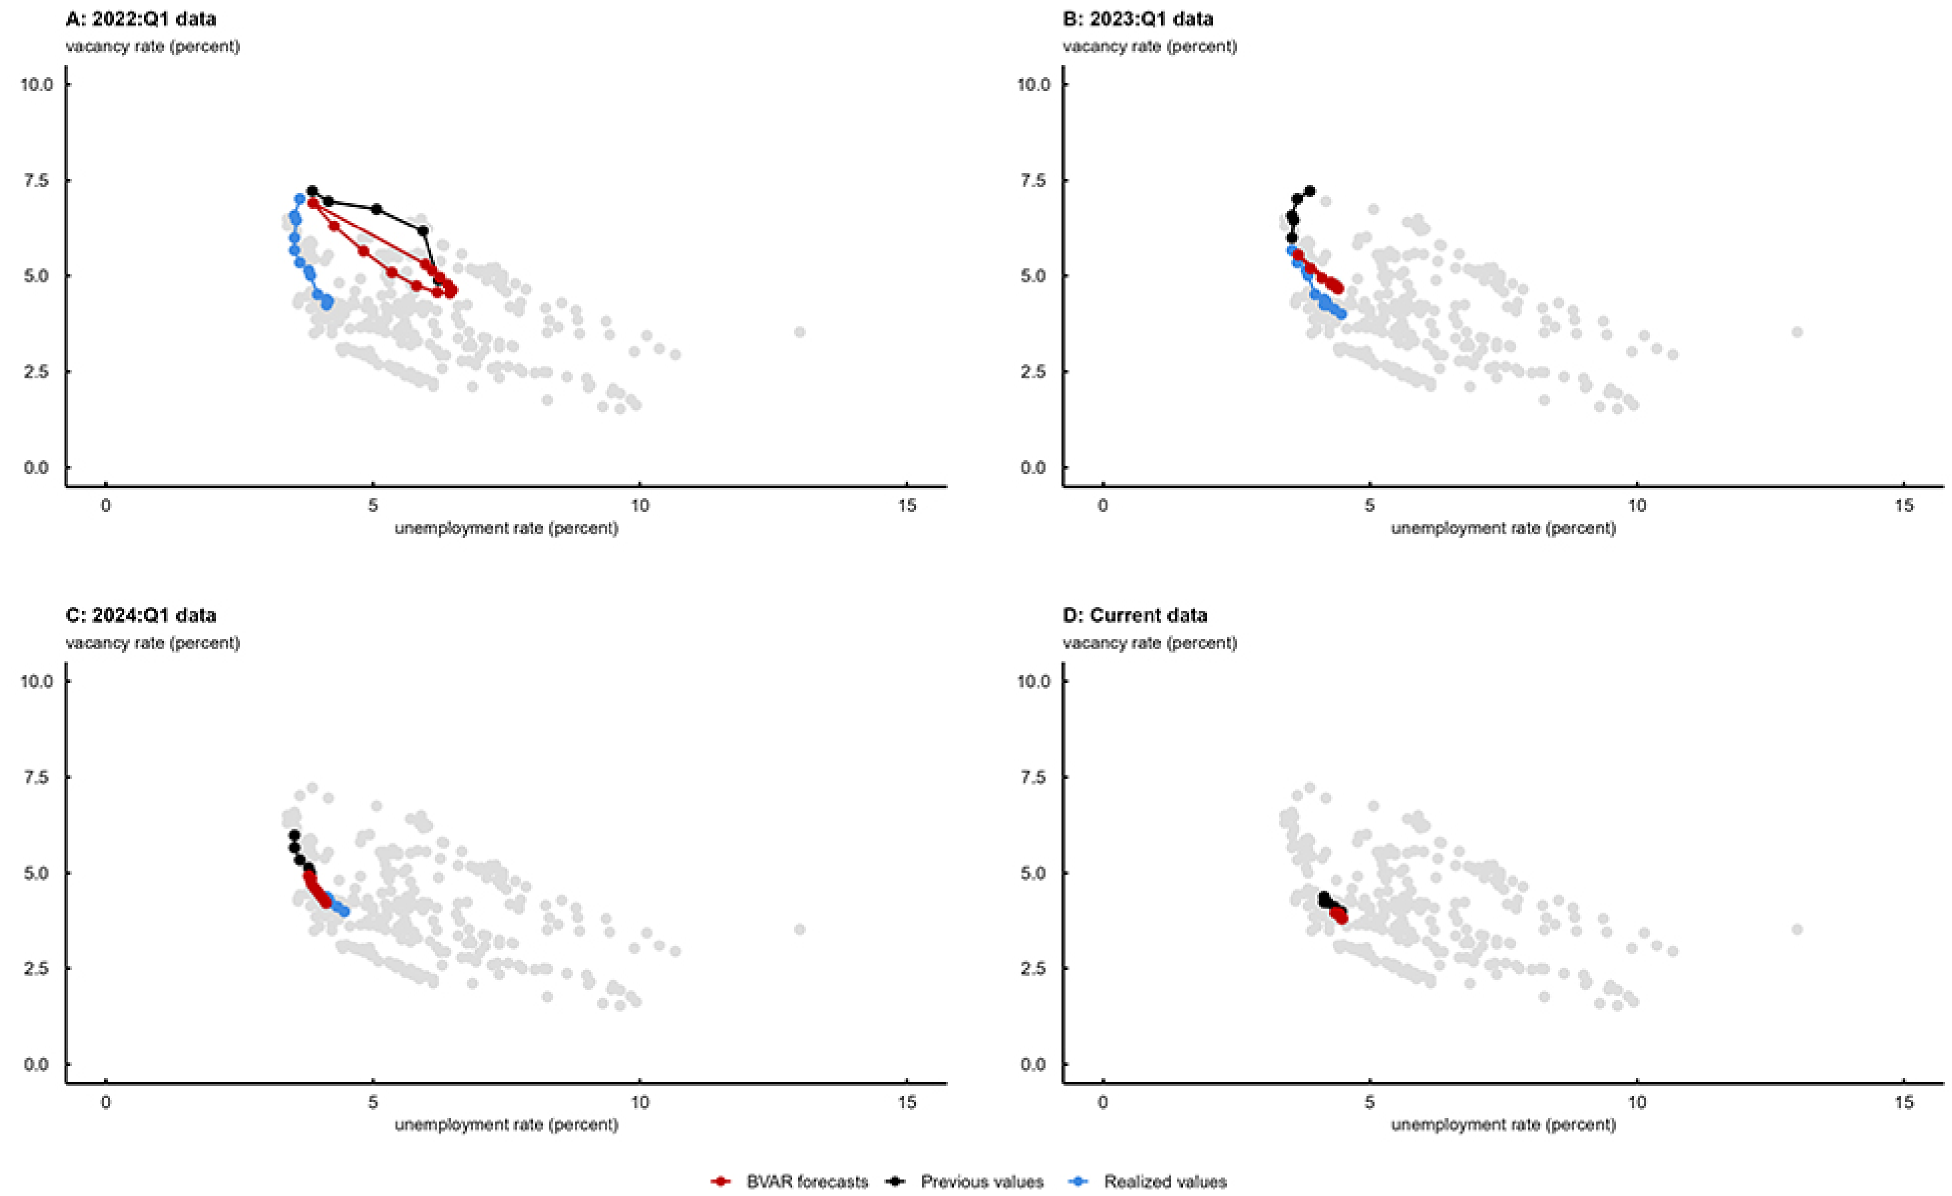

Notes: The figure shows historical values of the unemployment (*U*) and job vacancy (*V*) rates since 1955:Q1 (gray dots). Median projected values from our Bayesian vector autoregression (BVAR) (red dots) made at four different points in time—with data ending in 2022:Q1 (panel A), 2023:Q1 (panel B), 2024:Q1 (panel C), or 2025:Q4 (panel D)—are shown in each panel. The data for the previous five quarters (black dots) along with the actual values for periods that were previously forecasted (blue dots) and the BVAR forecasts (red dots) themselves are connected in the figure to illustrate the actual and predicted Beveridge curves by the model.
Sources: Authors' calculations based on data from Haver Analytics.

The model's forecasts have shifted over time in step with changes in the Beveridge curve. To see this, compare panels A through C of figure 5. The forecasts originating from 2022:Q1 in panel A suggest that declining vacancies would be accompanied by rising unemployment. However, from 2022:Q1 through 2023:Q1, *both* the unemployment and vacancy rates actually declined, with the latter falling to near 6% and the former edging below 4%. This can be seen in panel B, as the black dots make a U-turn over this time period as they transition to the blue dots.

Since 2023:Q1, the differences between the model forecasts and the realized unemployment and vacancy rates have been much smaller. The realized values in panel B of figure 5 show a somewhat steeper Beveridge curve than what the model projected back in 2023:Q1; but by early 2024 (panel C) these differences are mostly gone, with the model's most recent forecasts in panel D suggesting a flatter Beveridge curve.

Figure 5 suggests that not only have we moved along the Beveridge curve during the recovery from the pandemic recession, but the Beveridge curve itself has also shifted. Others have made the same observation, e.g., Barlevy et al. (2024), but one advantage of using our model to analyze these developments is that we can quantify the relative importance of *many* types of shocks that are capable of producing these movements. We do so by translating our BVAR to a Bayesian structural VAR (BSVAR) mirroring the Blanchard and Diamond (1989) model of the Beveridge curve and imposing a large number of sign restrictions to identify multiple shocks to the Beveridge curve that have been examined in the literature.

Here, we follow Ahn and Rudd (2024), Consolo et. al (2023), Schiman (2021), Foroni et. al (2018), and Furlanetto and Groshenny (2016) in identifying shocks that result in either a shift of the Beveridge curve ("shifters") or a movement along it ("sliders"). Shocks that produce a shift in the Beveridge curve result in the unemployment and vacancy rates moving directionally in tandem on impact, while shocks that produce a slide along the Beveridge curve result in the unemployment and vacancy rates changing in opposite directions.

We identify five structural shocks with our BSVAR: 1) labor supply, 2) match efficiency, 3) wage bargaining, 4) aggregate demand, and 5) aggregate supply shocks. These shocks capture changes in the available labor force (labor supply), the matching process of workers to jobs (match efficiency), the wage bargaining process between firms and workers (wage bargaining), and broader macroeconomic conditions (aggregate demand and supply).

A *labor supply (LS) shock* is the only shock we consider that moves labor force participation, per capita jobs and workers, and the unemployment and vacancy rates in the same direction on impact. Our sign restrictions here mirror those in Foroni et al. (2018) and imply that as workers unexpectedly enter the labor force, unemployment temporarily rises. The resulting increase in the supply of available workers puts downward pressure on real wages and hiring costs, leading firms to increase hiring activity and post job vacancies. Over time, the resulting increase in the demand for labor then causes the unemployment rate to fall as vacancies are filled.

In contrast, our remaining shocks capture changes in labor demand. These shocks all move per capita jobs and workers and the unemployment rate in opposite directions, but they differ from each other in the response of the vacancy rate that shifts the Beveridge curve after a *match efficiency (ME) shock (or labor reallocation shock)* and produces a movement along the curve after a *wage bargaining (WB) shock*. The ME and WB shocks are distinct among those that we consider in that labor force participation also falls on impact. Similar to the effects of positive labor supply shocks, real wages and inflation decline as employment rises.

We then disentangle *aggregate activity shocks* that also act as sliders along the Beveridge curve into *aggregate supply (AS)* and *aggregate demand (AD)* types. Both shocks are normalized to increase labor force participation and per capita jobs and workers. What distinguishes them from each other is instead the response of prices and wages on impact. *Aggregate supply shocks* lower inflation and increase real wages (like the other three shocks noted previously), whereas *aggregate demand shocks* increase inflation and, thus, lower real wages. For further details on our sign restrictions, see appendix 3.

To summarize what shocks our BSVAR attributes the recent behavior of the Beveridge curve to, we construct a historical decomposition of the unemployment and vacancy rates for the past four years into median contributions from each of our five individual sign-identified shocks (see figure 6). These decompositions highlight a shifting balance in the shocks that explain movements in the vacancy and unemployment rates—the combination that characterizes the Beveridge curve.

6. Historical decompositions of the unemployment and vacancy rates

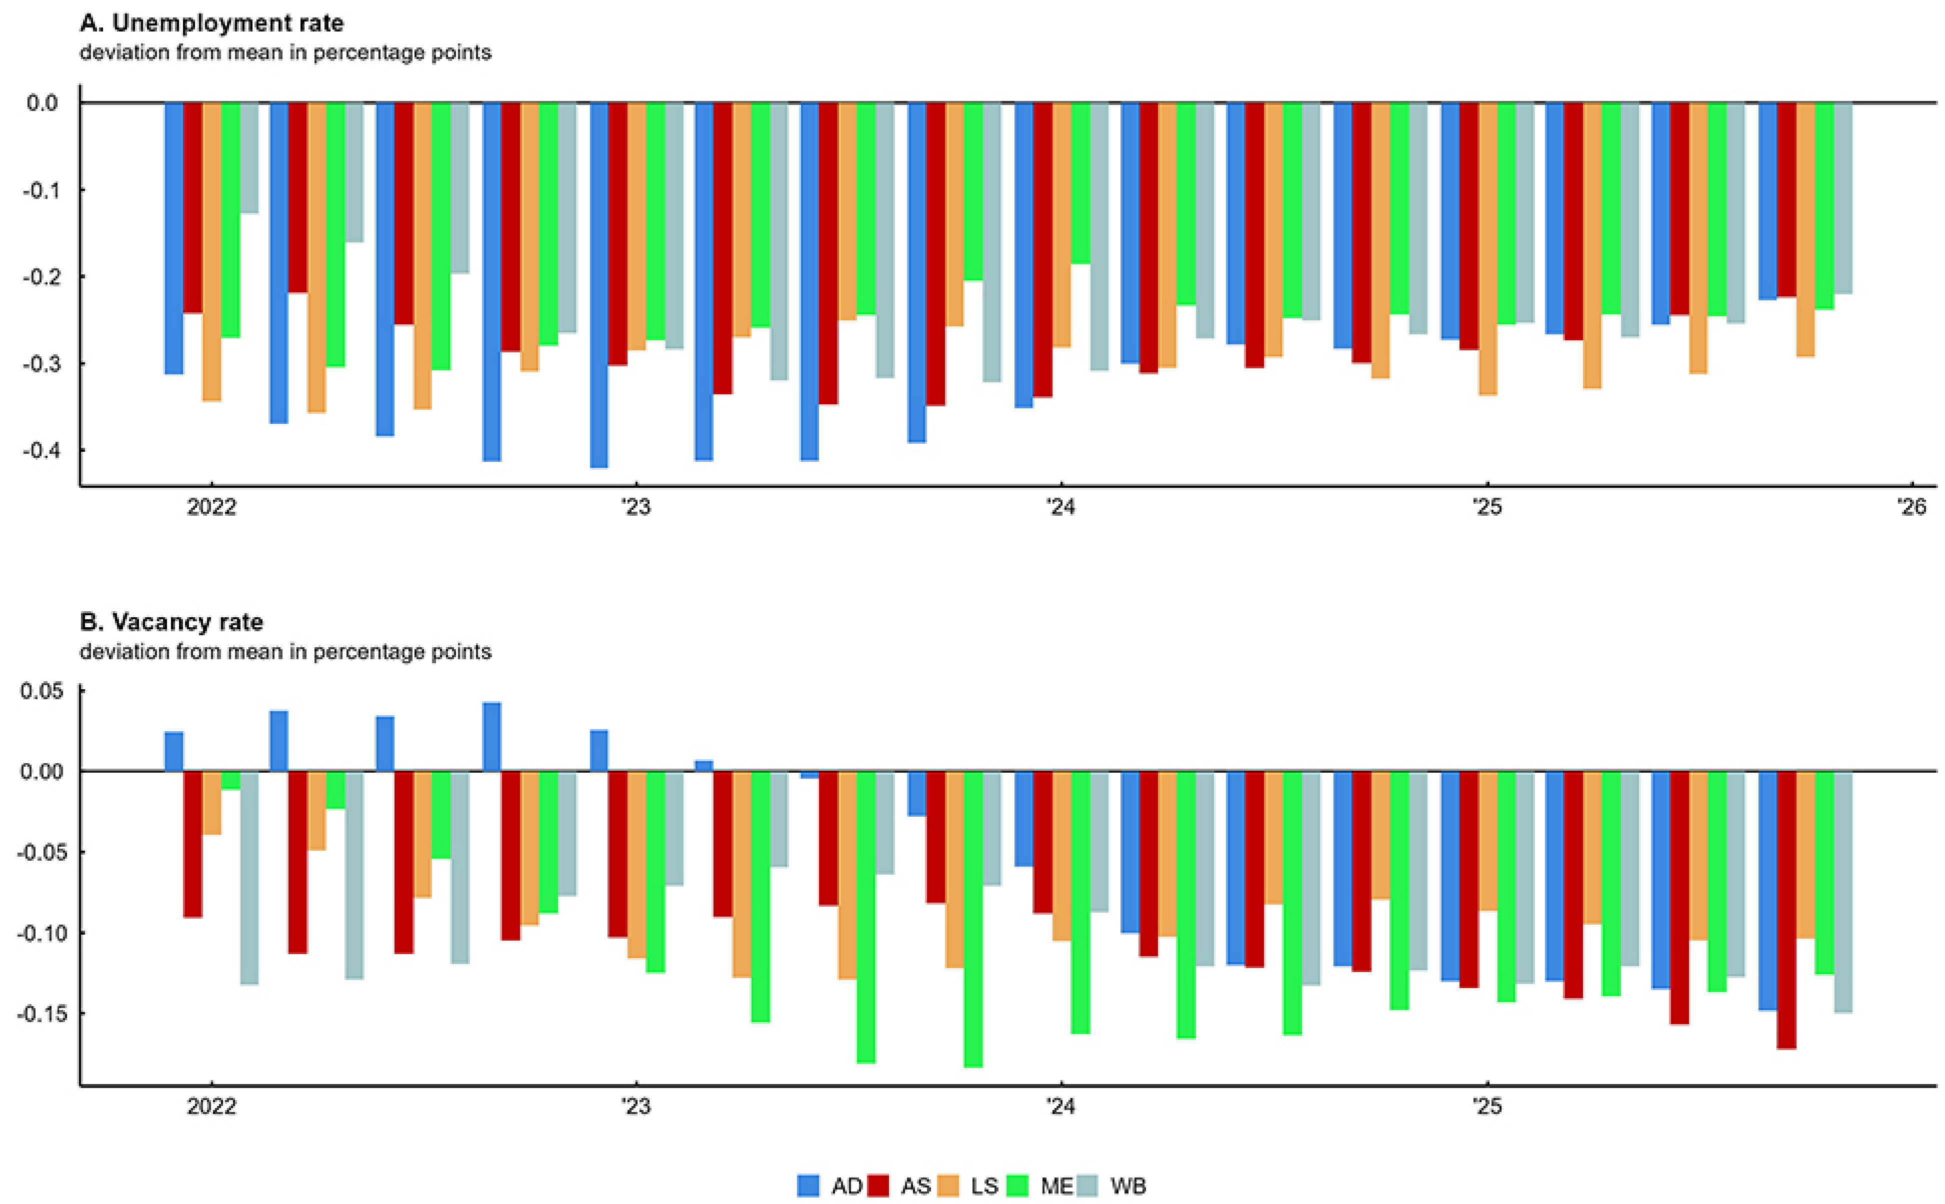

Notes: The figure shows an historical decomposition of the unemployment and vacancy rates to shocks of five types identified via sign restrictions: 1) labor supply (LS), 2) match efficiency (ME), 3) wage bargaining (WB), 4) aggregate supply (AS), and 5) aggregate demand (AD). See the text for further details.
Sources: Authors' calculations based on data from Haver Analytics.

While aggregate demand shocks were a primary driver of historically low unemployment in 2022 and 2023, their effect has been receding over the past two years and has been supplanted by labor supply shocks as the primary contributor to low unemployment. For vacancies, we see a bit of the opposite, with declining vacancies in 2023 and 2024 resulting primarily from wage bargaining and match efficiency shocks but more recent declines stemming almost equally from aggregate demand and aggegate supply shocks.

This shift in shock composition has important implications for the Beveridge curve. In general, the current mix of shocks has anticipated effects on both unemployment and job vacancies that tend to produce a movement along (as opposed to a shift in) the Beveridge curve. This can be seen in figure A11 of appendix 3, which covers the BSVAR, and helps to explain the features of our BVAR's labor market forecast.

Risks to the labor market outlook

Following Crump et. al (2021), we can also use our BVAR to examine potential downside and upside risks to the labor market outlook. We do so by conditioning the BVAR's forecasts on a predetermined path of a relevant variable and backing out the model-implied responses to the other variables in the BVAR that are left unconstrained in this exercise. This is achieved with the *BVAR* package by applying the conditional forecasting algorithm of Waggoner and Zha (1999). Further details can be found in Kuschnig and Vashold (2021).

Figures 7–11 report median (black line) and 70% prediction intervals (gray bands) from our baseline BVAR with the median forecasts (color lines) from four different scenarios: 1) lower labor demand (red), 2) lower labor supply (green), 3) lower household wealth (yellow) and 4) higher investment (blue).

In the following subsections, we describe how each scenario is formed and summarize their results.

Higher unemployment insurance claims

Motivated by figure 6, we first assess the downside risk of a deterioration in labor demand. To investigate this scenario, we condition the BVAR's forecasts on a path for initial unemployment insurance claims consistent with the upper end of the 70% prediction interval shown in figure 7. This scenario (see the red lines in figures 7–11) has initial UI claims quickly rising and then holding at around their mid-2021 levels.

In this case, the unemployment rate rises to around 4.75% by the end of the projection period; while the vacancy rate still falls to about 4%, it falls more rapidly than in the BVAR's baseline forecast. The jobs–workers gap quickly turns negative—reaching about –1% by the end of the projection period (see figure 7). Per capita jobs and hours fall as well, and per capita real gross domestic product (GDP), consumption, and investment either stagnate or decline for a few quarters (see figures 9 and 11).

While such a scenario is less rosy than our BVAR's baseline forecasts, declines in employment and hours of this magnitude would be only mildly recessionary. It is worth noting, however, that the decline is limited in part because the model projects nominal interest rates adjust downward more quickly in this scenario to offset some of the labor market weakness.

Lower labor force participation rate

Figure 6 also suggests that labor supply shocks have played a nontrivial role in promoting both low unemployment and high vacancy rates since 2021. Next, we consider a scenario with much lower labor force participation as a plausible downside risk (see the green lines in figures 7–11) to our BVAR's baseline forecast. We can test this risk by conditioning the BVAR's forecasts on a path for the labor force participation rate—a key labor market indicator associated with movements in labor supply—consistent with the lower end of the 70% prediction interval shown in figure 7.

In this scenario, the unemployment rate also rises more than in the baseline forecast, peaking at a little above 4.5% with the jobs–workers gap near zero throughout the projection period (see figure 7). Per capita real GDP, consumption, and investment, however, are slightly lower than in the baseline forecast (see figures 9 and 11).

Despite the slightly larger increase in the unemployment rate than in our baseline forecast, it is difficult to argue that this scenario is characteristic of a recession. This scenario does lead, however, to both lower inflation and nominal interest rates than in our baseline forecast, accompanied by a decline in inflation expectations.

Lower household wealth

The continued strength of consumer spending in our baseline forecast, supported by rising home and equity prices, is another logical source of risk to the labor market outlook. We examine this risk by conditioning the BVAR's forecasts on a path for per capita real household net worth consistent with the lower end of its 70% prediction interval in figure 11.

The unemployment rate in this scenario also peaks at a little above 4.5%; but vacancies fall by more than in the baseline forecast, with the jobs–workers gap reaching a near identical −1% by the end of the projection period to our lower labor demand scenario. Per capita real output, consumption, and investment as well as nominal and real interest rates (see figures 9 and 10) also evolve very similarly to what we observed in our lower labor demand scenario, suggesting mildly recessionary dynamics.

Higher investment

So far, we've only considered downside risks. Next, we consider an upside risk scenario where the rate of investment picks up to match the upper end of the 70% prediction interval shown in figure 11. To motivate this scenario, we note from figure 6 that aggregate supply shocks are currently playing a nontrivial role. Our higher investment scenario (see the yellows lines in figures 7–11) is meant to capture an acceleration of these trends, perhaps reflecting broader adoption of AI or the possibility that recent expansionary fiscal policy measures could boost gross private domestic investment in the near term.

In this scenario, the unemployment rate gradually begins to fall back to a little above 4%; while the vacancy rate slowly rises to a little above 4% and the jobs–workers gap begins to increase, ending the projection period at about 0.5% (see figure 7). Employment and labor force participation also rise, with per capita hours, real GDP, consumption, labor income, and net worth all above our baseline forecast (see figures 9 and 11). The stronger economy in this scenario leads to broadly similar nominal interest rates compared with those in our baseline forecast, but with real interest rates slightly higher and inflation slightly lower over the forecast horizon (see figure 10).

7. Labor market slack

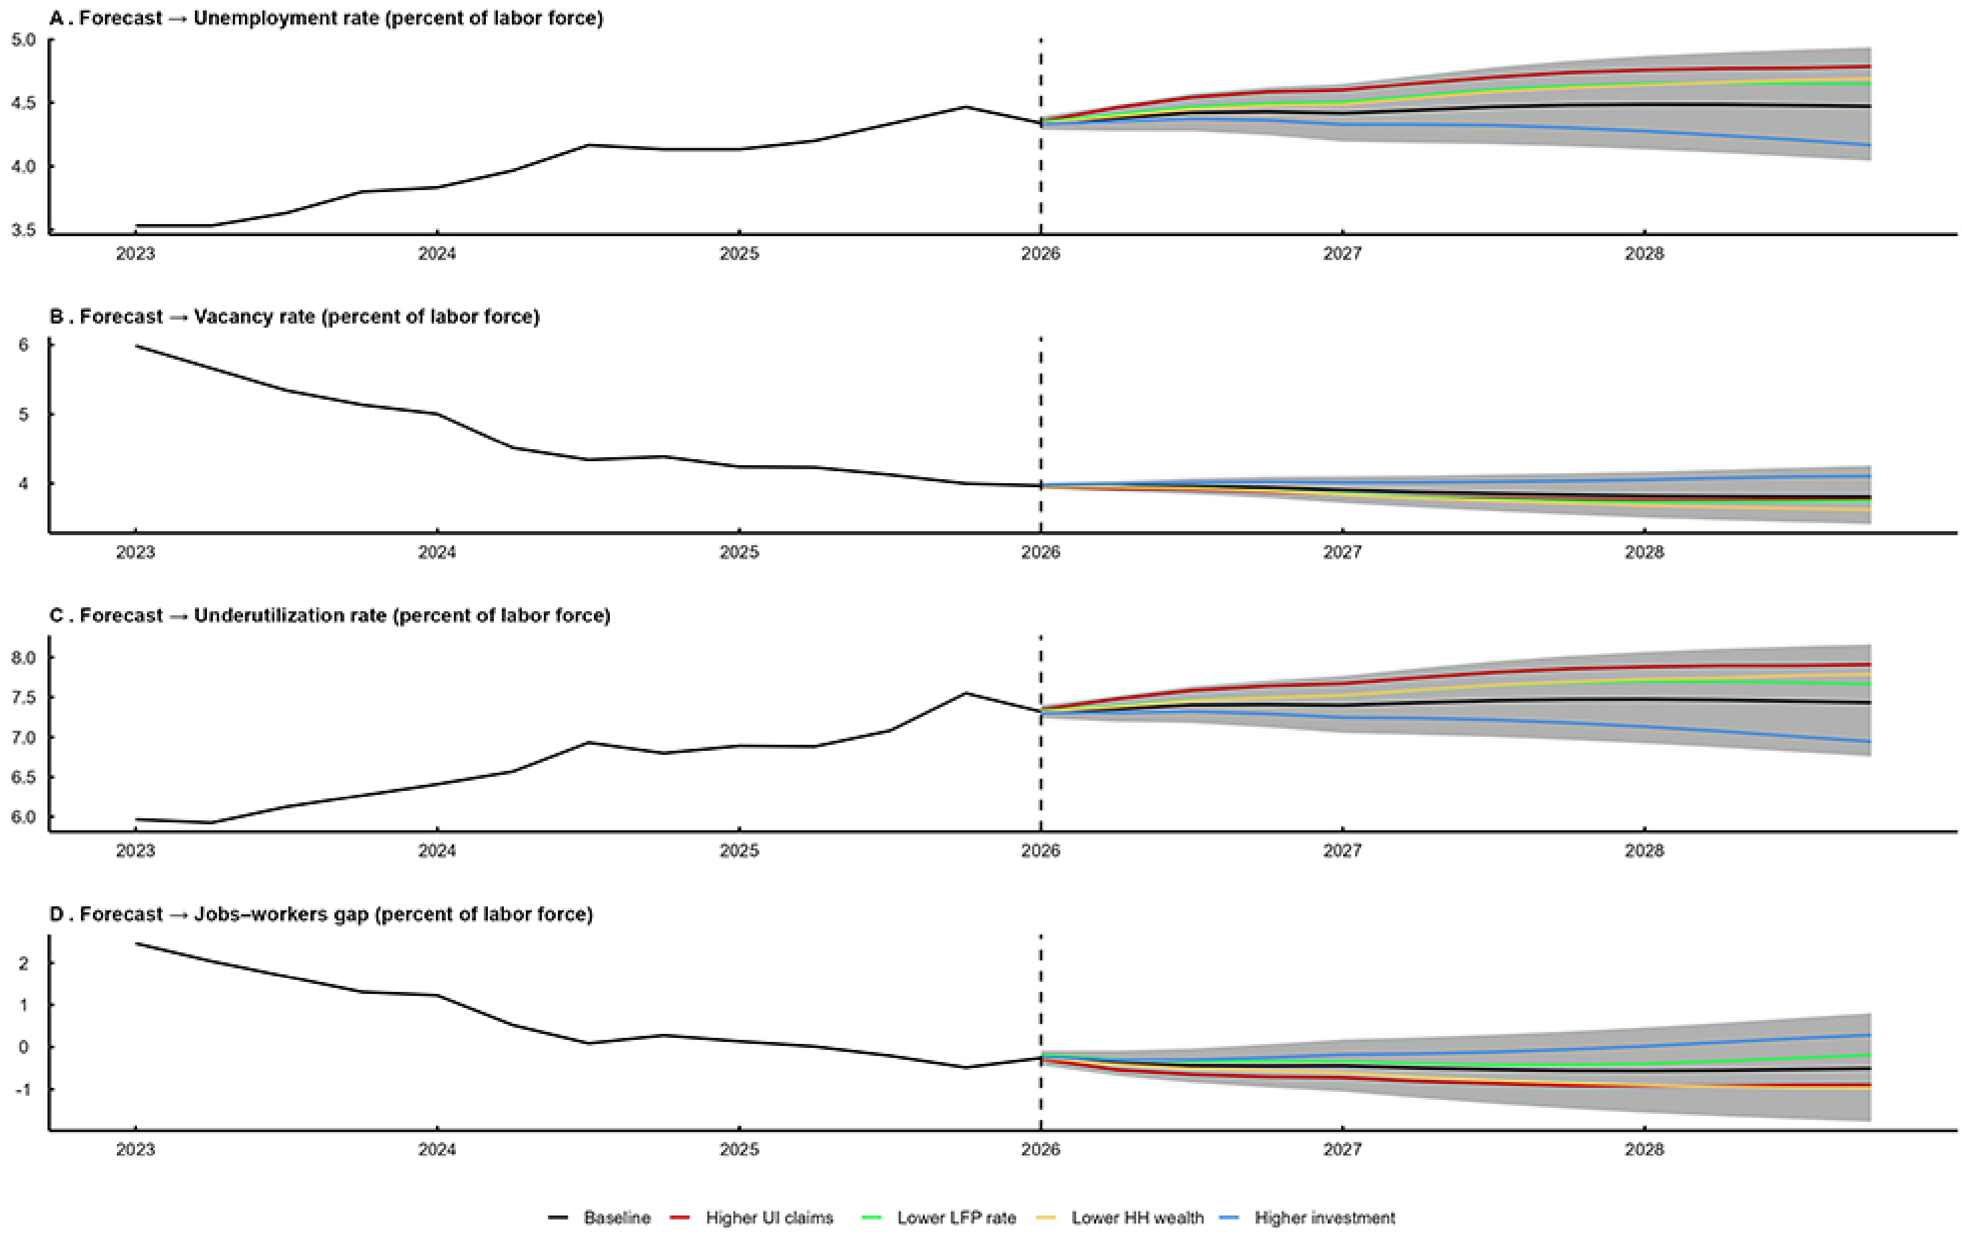

Notes: The figure repeats the median forecasts (black line) and 70% prediction intervals (gray bands) from our baseline Bayesian vector autoregression (BVAR) with Minnesota & long-run priors (baseline forecast). The area to the right of the vertical dashed line in each panel of the figure corresponds with the forecast period beginning in 2026:Q1. The color lines in each panel of the figure (scenario forecast) report the median forecast for a given scenario discussed in the text. In this figure, UI stands for unemployment insurance; LFP, labor force participation; and HH, household.
Sources: Authors' calculations based on data from Haver Analytics.

8. Labor force

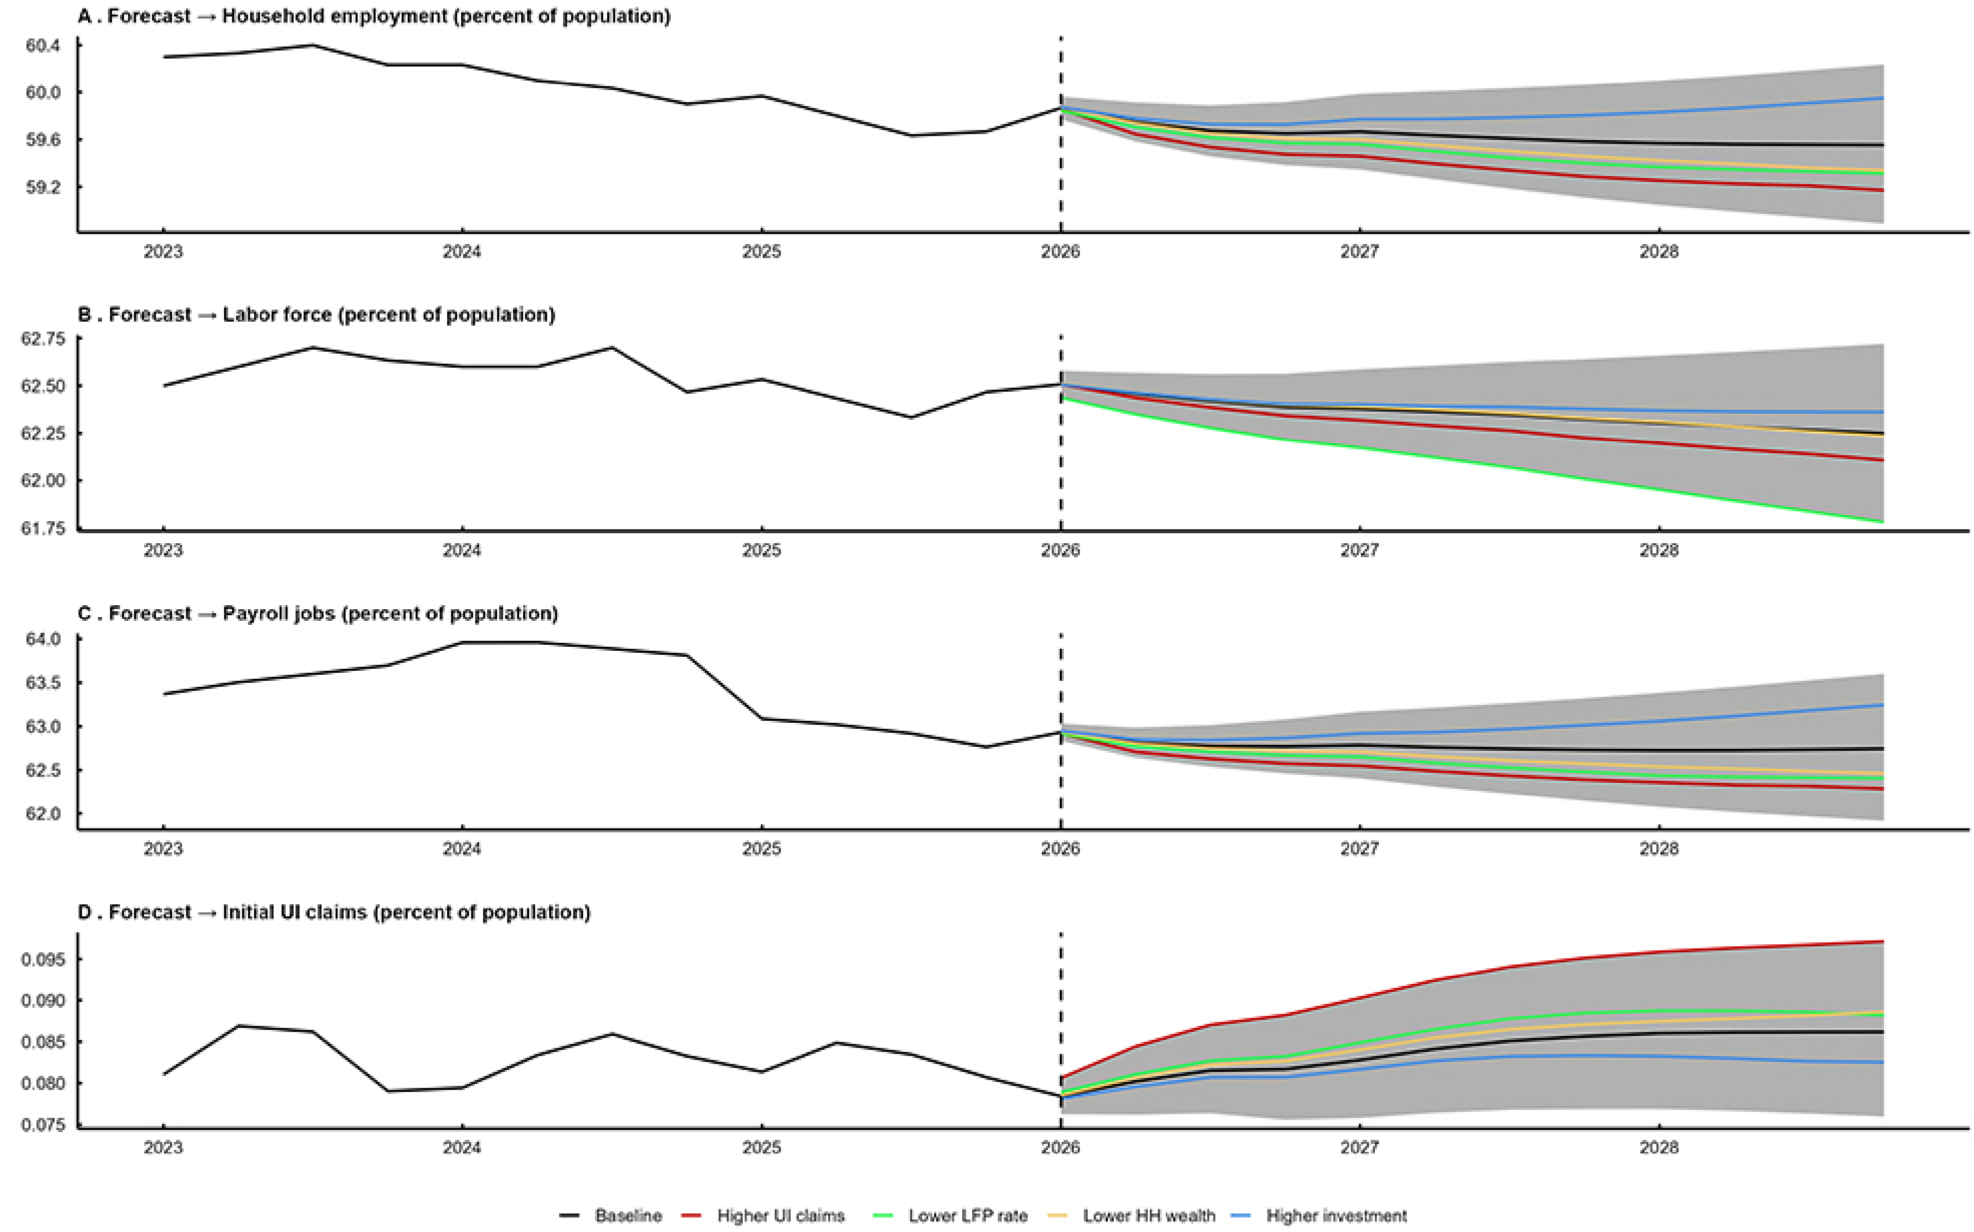

Notes: The figure repeats the median forecasts (black line) and 70% prediction intervals (gray bands) from our baseline Bayesian vector autoregression (BVAR) with Minnesota & long-run priors (baseline forecast). The area to the right of the vertical dashed line in each panel of the figure corresponds with the forecast period beginning in 2026:Q1. The color lines in each panel of the figure (scenario forecast) report the median forecast for a given scenario discussed in the text. In this figure, population stands for the noninstitutional civilian population aged 16 and older; UI, unemployment insurance; LFP, labor force participation; and HH, household.
Sources: Authors' calculations based on data from Haver Analytics.

9. Labor productivity

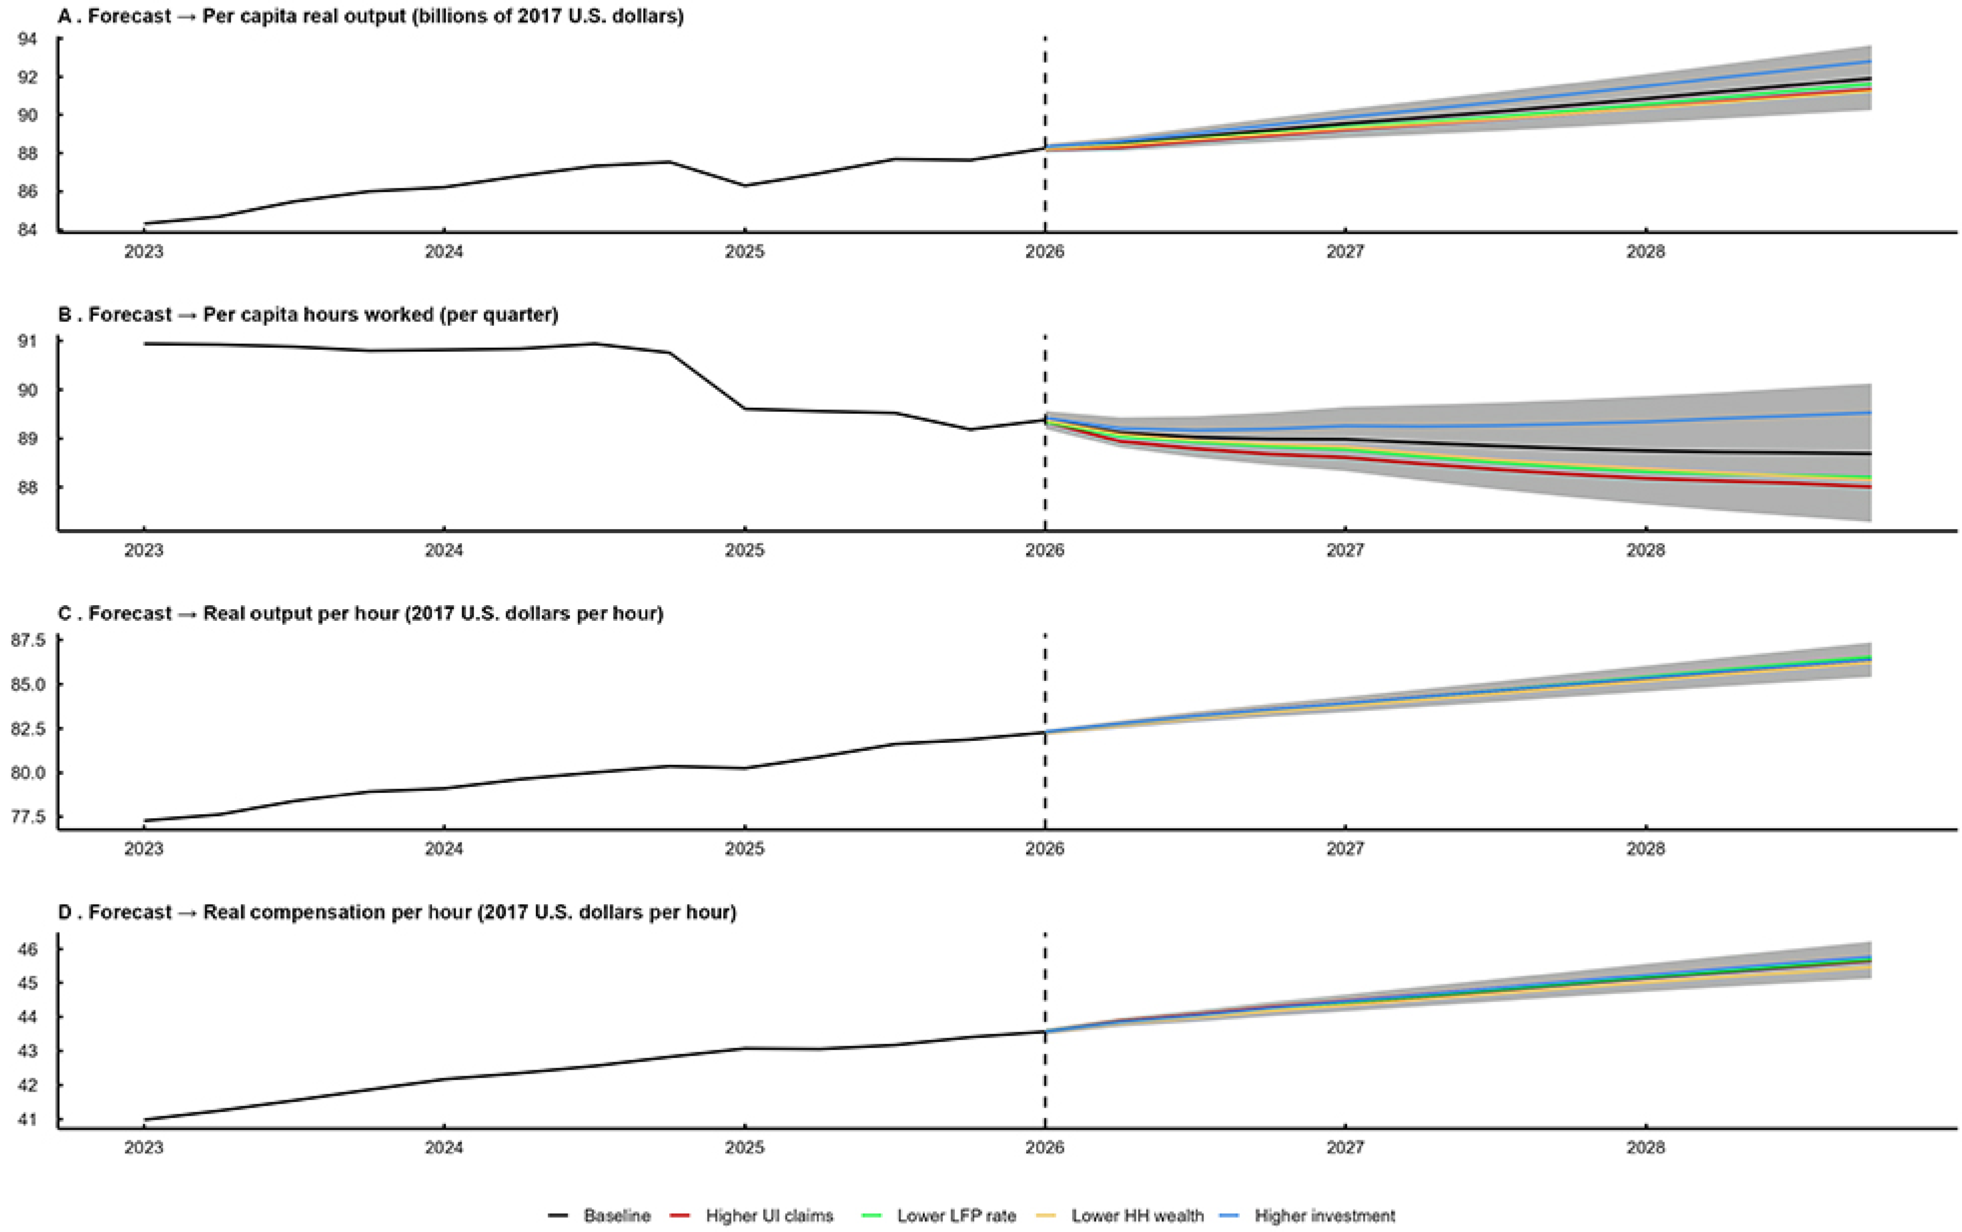

Notes: The figure repeats the median forecasts (black line) and 70% prediction intervals (gray bands) from our baseline Bayesian vector autoregression (BVAR) with Minnesota & long-run priors (baseline forecast). The area to the right of the vertical dashed line in each panel of the figure corresponds with the forecast period beginning in 2026:Q1. The color lines in each panel of the figure (scenario forecast) report the median forecast for a given scenario discussed in the text. In this figure, UI stands for unemployment insurance; LFP, labor force participation; and HH, household.
Sources: Authors' calculations based on data from Haver Analytics.

10. Inflation and interest rates

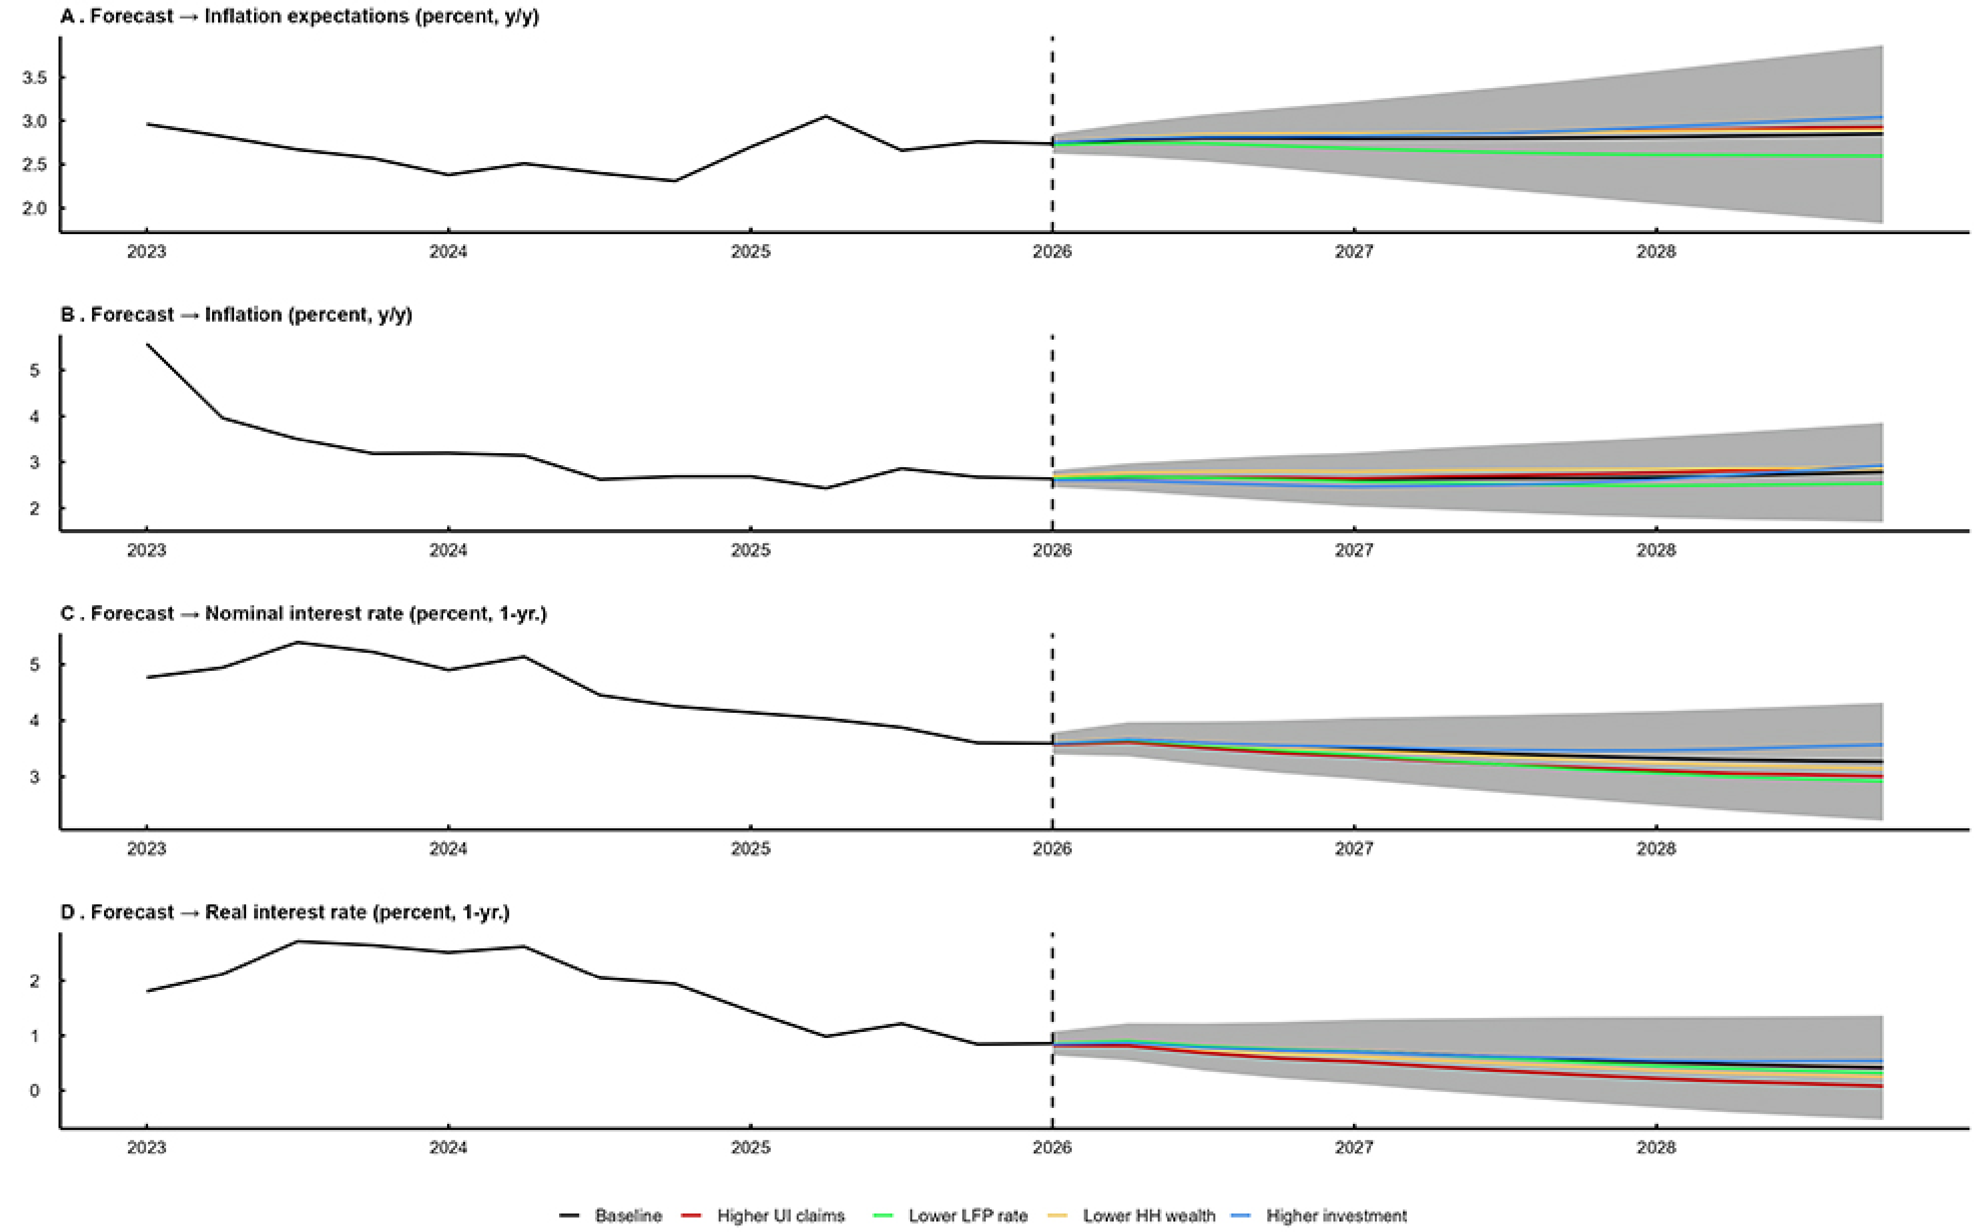

Notes: The figure repeats the median forecasts (black line) and 70% prediction intervals (gray bands) from our baseline Bayesian vector autoregression (BVAR) with Minnesota & long-run priors (baseline forecast). The area to the right of the vertical dashed line in each panel of the figure corresponds with the forecast period beginning in 2026:Q1. The color lines in each panel of the figure (scenario forecast) report the median forecast for a given scenario discussed in the text. In this figure, UI stands for unemployment insurance; LFP, labor force participation; and HH, household.
Sources: Authors' calculations based on data from Haver Analytics.

11. Consumption and investment

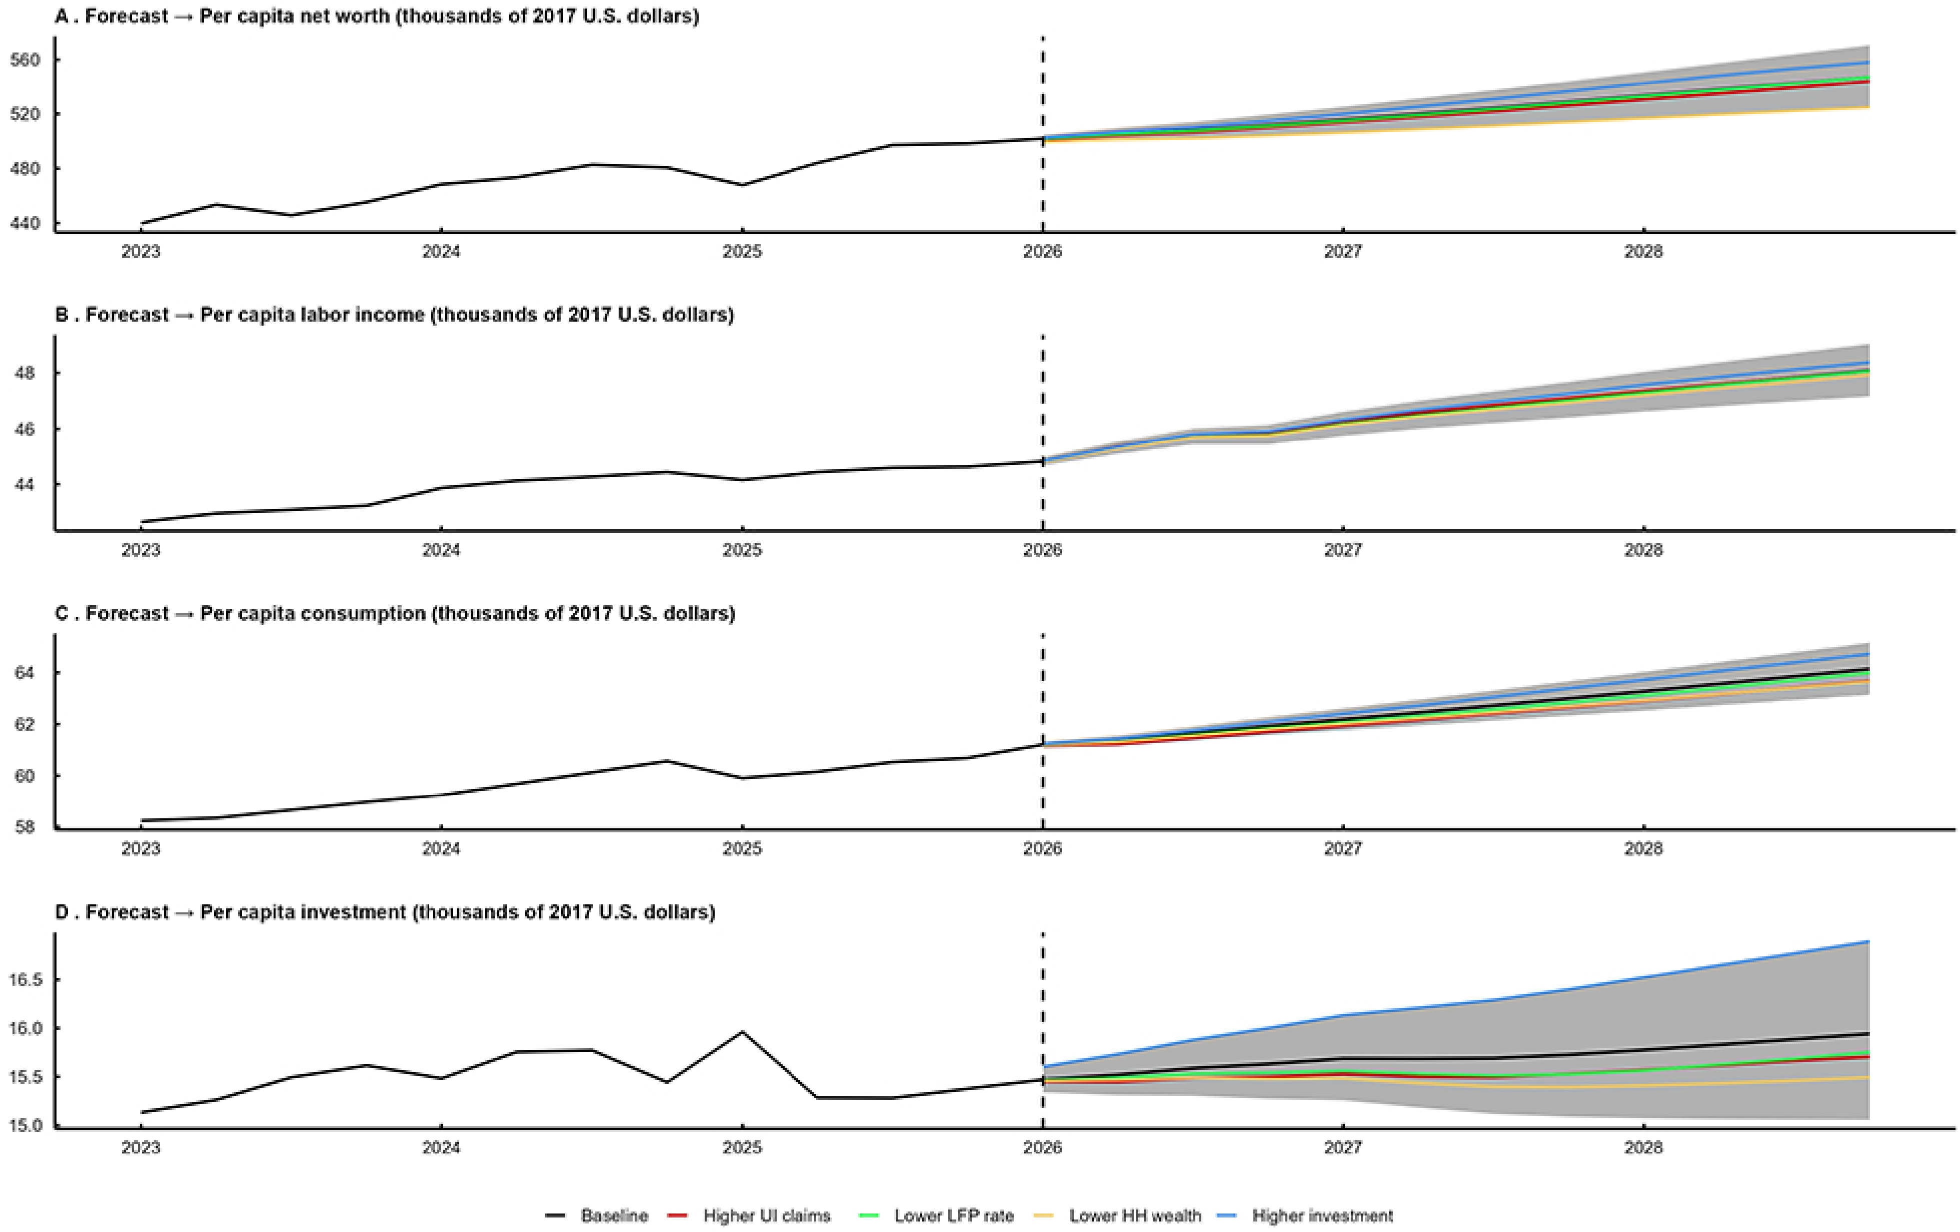

Notes: The figure repeats the median forecasts (black line) and 70% prediction intervals (gray bands) from our baseline Bayesian vector autoregression (BVAR) with Minnesota & long-run priors (baseline forecast). The area to the right of the vertical dashed line in each panel of the figure corresponds with the forecast period beginning in 2026:Q1. The color lines in each panel of the figure (scenario forecast) report the median forecast for a given scenario discussed in the text. In this figure, UI stands for unemployment insurance; LFP, labor force participation; and HH, household.
Sources: Authors' calculations based on data from Haver Analytics.

Conclusion

Our BVAR's baseline forecasts suggest that the labor market is likely to continue slowing in the coming quarters. It is worth restating our word of caution. Like any model, our BVAR is only a *guide* (and not a *crystal ball*). It only allows us to evaluate the likelihood of potential outcomes based on how similar experiences have proceeded in the past and how our prior knowledge about labor market dynamics may be expected to impact those outcomes.

The scenario analyses presented in this article are a way of providing a more robust assessment of the driving features of that forecast. Three of the four scenarios we explored here point to potential adverse risks to the BVAR's forecast for the labor market. In particular, they highlight the upside risk to unemployment from a more substantial softening of either labor demand or supply, as well as of aggregate demand. The fourth scenario, in contrast, highlights the potential downside risks for unemployment from an acceleration in investment growth.

While these scenarios seemed well motivated to us based off of recent experience, still others may be of interest to market participants and policymakers alike. To aid in the process of forecasting near-term labor market conditions and analyzing the supply and demand drivers of the Beveridge curve, we have created an open-source application for download where it is possible for the user to implement scenarios of their own choosing. Further details for the R Shiny app can be found in appendix 4.

Details for Scott A. Brave and Ross Cole are available on their Chicago Fed online profiles. R. Andrew Butters is an associate professor and the Blanche "Peg" Philpott Faculty Fellow at the Kelley School of Business at Indiana University. A former business economist with the Chicago Fed, J. Max Gillet is a PhD in finance candidate at the University of Chicago Booth School of Business.

Notes

[1] The Sahm rule is a real-time indicator developed in Sahm (2019) to signal the beginning of a recession by flagging instances where the current three-month moving average of the unemployment rate is 0.5 percentage points or more above its most recent 12-month minimum.

References

Ahn, Hie Joo, and Jeremy B. Rudd, 2025, "(Re-)connecting inflation and the labor market: A tale of two curves," Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System, No. 2024-050, revised May 2025 (originally issued July 2024). Crossref

Barlevy, Gadi, R. Jason Faberman, Bart Hobijn, and Ayşegül Şahin, 2024, "The shifting reasons for Beveridge curve shifts," *Journal of Economic Perspectives*, Vol. 38, No. 2, Spring, pp. 83–106. Crossref

Barnichon, Regis, 2010, "Building a composite Help-Wanted Index," *Economics Letters*, Vol. 109, No. 3, pp. 175–178. Crossref

Benigno, Pierpaolo, and Gauti B. Eggertsson, 2023, "It's baaack: The surge in inflation in the 2020s and the return of the non-linear Phillips curve," National Bureau of Economic Research, working paper, No. 31197, April. Crossref

Blanchard, Olivier Jean, and Peter A. Diamond, 1989, "The Beveridge curve," *Brookings Papers on Economic Activity*, Vol. 20, No. 1 pp. 1–60, available online.

Brave, Scott A., R. Andrew Butters, and Alejandro Justiniano, 2019, "Forecasting economic activity with mixed frequency BVARs," *International Journal of Forecasting*, Vol. 35, No. 4, October–December, pp. 1692–1707. Crossref

Cajner, Tomaz, and David D. Ratner, 2016, "A cautionary note on the help wanted online data," *FEDS Notes*, Board of Governors of the Federal Reserve System, June 23. Crossref

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino, 2015, "Bayesian VARs: Specification choices and forecast accuracy," *Journal of Applied Econometrics*, Vol. 30, No. 1, January/February, pp. 46–73. Crossref

Cascaldi-Garcia, Danilo, 2025, "Pandemic priors," Board of Governors of the Federal Reserve System, working paper, May, available online.

Chan, Joshua C. C., 2022, "Asymmetric conjugate priors for large Bayesian VARs," *Quantitative Economics*, Vol. 13, No. 3, July, pp. 1145–1169. Crossref

Chan, Joshua C. C., 2021, "Minnesota-type adaptive hierarchical priors for large Bayesian VARs," *International Journal of Forecasting*, Vol. 37, No. 3, July–September, pp. 1212–1226. Crossref

Consolo, Agostino, Claudia Foroni, and Catalina Martínez Hernández, 2023, "A mixed frequency BVAR for the euro area labour market," *Oxford Bulletin of Economics and Statistics*, Vol. 85, No. 5, October, pp. 1048–1082. Crossref

Crump, Richard K., Stefano Eusepi, Domenico Giannone, Eric Qian, and Argia M. Sbordone, 2021, "A large Bayesian VAR of the United States economy," Federal Reserve Bank of New York, staff report, No. 976, August 19, available online.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin, 2013, "The decline of the U.S. labor share," *Brookings Papers on Economic Activity*, Vol. 44, No. 2, Fall, pp. 1–52, available online.

Foroni, Claudia, Francesco Furlanetto, and Antoine Lepetit, 2018, "Labor supply factors and economic fluctuations," *International Economic Review*, Vol. 59, No. 3, August, pp. 1491–1510. Crossref

Francis, Neville, and Valerie A. Ramey, 2009, "Measures of per capita hours and their implications for the technology-hours debate," *Journal of Money, Credit and Banking*, Vol. 41, No. 6, September, pp. 1071–1097. Crossref

Furlanetto, Francesco, and Nicolas Groshenny, 2016, "Mismatch shocks and unemployment during the Great Recession," *Journal of Applied Econometrics*, Vol. 31, No. 7, November/December, pp. 1197–1214. Crossref.

Giacomini, Raffaella, 2013, "The relationship between DSGE and VAR models," in *VAR Models in Macroeconomics—New Developments and Applications: Essays in Honor of Christopher A. Sims*, Thomas B. Fomby, Lutz Kilian, and Anthony Murphy (eds.), Advances in Econometrics, Vol. 32, Leeds, UK: Emerald Group Publishing, pp. 1–26. Crossref

Giannone, Domenico, Michele Lenza, and Giorgio E. Primiceri, 2019, "Priors for the long run," *Journal of the American Statistical Association*, Vol. 114, No. 526, pp. 565–580. Crossref

Giannone, Domenico, Michele Lenza, and Giorgio E. Primiceri, 2015, "Prior selection for vector autoregressions," *Review of Economics and Statistics*, Vol. 97, No. 2, May, pp. 436–451. Crossref

Gittleman, Maury, 2022, "The 'Great Resignation' in perspective," *Monthly Labor Review*, July. Crossref

Kuschnig, Nikolas, and Lukas Vashold, 2021, "BVAR: Bayesian vector autoregressions with hierarchical prior selection in R," *Journal of Statistical Software*, Vol. 100, No. 14, pp. 1–27. Crossref

Lenza, Michele, and Giorgio E. Primiceri, 2022, "How to estimate a vector autoregression after March 2020," *Journal of Applied Econometrics*, Vol. 37, No. 4, June/July, pp. 688–699. Crossref

Lubik, Thomas A., 2021, "Revisiting the Beveridge curve: Why has it shifted so dramatically?," *Economic Brief*, Federal Reserve Bank of Richmond, No. 21-36, October, available online.

O'Trakoun, John, 2022, "Brush up those resumes, dust off that disco ball?," *Macro Minute*, Federal Reserve Bank of Richmond, blog, March 29, available online.

Romer, Christina D., and David H. Romer, 2018, "Phillips lecture—Why some times are different: Macroeconomic policy and the aftermath of financial crises," *Economica*, Vol. 85, No. 337, January, pp. 1–40. Crossref

Rudd, Jeremy, and Karl Whelan, 2002, "A note on the cointegration of consumption, income, and wealth," Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System, No. 2002-53, November, available online.

Sahm, Claudia, 2019, "Direct stimulus payments to individuals," in *Recession Ready: Fiscal Policies to Stabilize the American Economy*, Heather Boushey, Ryan Nunn, and Jay Shambaugh (eds.), Washington, DC: Brookings Institution, pp. 67–92, available online.

Schiman, Stefan, 2021, "Labor supply shocks and the Beveridge curve—Empirical evidence from EU enlargement," *Review of Economic Dynamics*, Vol. 40, April, pp. 108–127. Crossref

Sims, Christopher A., 1980, "Macroeconomics and reality," *Econometrica*, Vol. 48, No. 1, January, pp. 1–48. Crossref

Stock, James H., and Mark W. Watson, 2012, "Disentangling the channels of the 2007–09 recession," *Brookings Papers on Economic Activity*, Vol. 43, No. 1, Spring, pp. 81–135, available online.

Uhlig, Harald, 2017, "Shocks, sign restrictions, and identification," in *Advances in Economics and Econometrics: Eleventh World Congress*, Bo Honoré, Ariel Pakes, Monika Piazzesi, and Larry Samuelson (eds.), Vol. 2, Econometric Society Monographs, Cambridge, UK: Cambridge University Press, pp. 95–127. Crossref

Waggoner, Daniel F., and Tao Zha, 1999, "Conditional forecasts in dynamic multivariate models," *Review of Economics and Statistics*, Vol. 81, No. 4, November, pp. 639–651. Crossref