Technical Appendix for *Economic Perspectives* Article (2026, No. 1): The U.S. Labor Market: A Long-Run Perspective

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

Appendix 1: Data

Figure A1 contains a list of the 18 quarterly time series used in our Bayesian vector autoregression (BVAR) along with their Haver Analytics mnemonics, any necessary data transformations, and their original data sources. The vast majority of data series used in our analysis originate from the U.S. Bureau of Labor Statistics (BLS) or the U.S. Bureau of Economic Analysis (BEA), with only a few exceptions (unemployment insurance claims are from the U.S. Department of Labor (DOL) and the remaining financial and inflation expectations data are from the Federal Reserve System). We used data from all original sources as of March 29, 2026, to estimate the BVAR.

Several data series included in the BVAR are measured in per capita terms. Our population measure used for this purpose is the noninstitutional civilian population aged 16 and older from the BLS. Population controls from the BLS derive from U.S. Census Bureau estimates and are updated on an annual basis. Household employment and unemployment data are from the BLS's *Current Population Survey* (CPS), while job vacancies data are from the BLS's *Job Openings and Labor Turnover Survey* (JOLTS). Unemployment insurance (UI) claims data from the DOL are expressed in either per capita terms for initial claims or as a percent of covered employees by the UI system for continuing claims.

Nominal gross domestic product (GDP), personal consumption expenditures and income, and gross private domestic investment are from the BEA's *National Income and Product Accounts of the United States* (NIPAs). Hours worked, payroll jobs, and compensation are from the BLS's *Total U.S. Economy* estimates. We subtract off contributions to personal income, compensation, hours, and jobs from members of the military using aggregates found in the BEA's NIPAs. Household net worth is from the Federal Reserve Board's *Financial Accounts of the United States*. To deflate nominal output, compensation, and investment, we use the implicit GDP deflator from the BEA. To deflate nominal personal consumption expenditures and income as well as household net worth, we use the implicit personal consumption expenditures (PCE) deflator from the BEA.

Our measure of inflation is the four-quarter log change in the Consumer Price Index (CPI). Our measure of inflation expectations is the one-year-ahead median forecast for the CPI from the *Livingston Survey* prior to 1990:Q2 combined with the similar projection from the *Survey of Professional Forecasters* from this date onward from the Federal Reserve Bank of Philadelphia. The nominal interest rate series is the one-year constant maturity U.S. Treasury bill yield published by the Federal Reserve Board.

A1. Data transformations and original sources

Description	Mnemonic	Trans.	Source
Labor force (% of pop.)	$LP@USECON^{\dagger}/100$	Log	U.S. Bureau of Labor Statistics
Civilian employment (% of pop.)	$LQ@USECON^{\dagger}/100$	Log	U.S. Bureau of Labor Statistics
Vacancies (% of pop.)	$(LJJTLA@USECON^{\dagger}/LTU@USECON^{\dagger}) \times (LP@USECON^{\dagger}/100) \times (LR@USECON^{\dagger}/100)$	Log	U.S. Bureau of Labor Statistics
Unemployed (% of pop.)	$(LR@USECON^{\dagger} \times LP@USECON^{\dagger})/100$	Log	U.S. Bureau of Labor Statistics
Civilian jobs (% of pop.)	$1e6 \times (LXEML@USECON - LXAML@USECON)^{\dagger} /(LN16N@USECON \times 1e3)^{\dagger}$	Log	U.S. Bureau of Labor Statistics
Real output per hour (2017 USD/hr)	$GDP@USNA/(DGDP@USNA/100 \times (LXEHL@USECON - LXAHL@USECON))$	Log	U.S. Bureaus of Economic Analysis and Labor Statistics
Average weekly hours worked (hrs/job/week)	$1e9 \times (LXEHL@USECON - LXAHL@USECON)/(52 \times (LXEML@USECON - LXAML@USECON)^{\dagger} \times 1e6)$	Log	U.S. Bureau of Labor Statistics
Real compensation per hour (2017 USD/hr)	$(LXEFL@USECON - GGDES@USNA) /(DGDP@USNA/100 \times (LXEHL@USECON - LXAHL@USECON))$	Log	U.S. Bureaus of Economic Analysis and Labor Statistics
Per capita real investment (thous. 2017 USD)	$(1e6 \times I@USECON)/(DGDP@USNA/100 \times (LN16N@USECON \times 1e3))$	Log	U.S. Bureaus of Economic Analysis and Labor Statistics
SPF one-year inflation expectations	$ASACX1@SURVEYS/100$	NA	Federal Reserve Bank of Philadelphia
CPI inflation	$Log(PCU@USECON_{t}^{\dagger}/PCU@USECON_{t-4}^{\dagger})$	NA	U.S. Bureau of Labor Statistics
One-year nominal interest rate	$FCM1@USECON/100^{\dagger}$	NA	Federal Reserve Board
Insured unemployment (% of covered employees)	$LIURM@USECON^{\dagger}/100$	Log	U.S. Department of Labor
Initial UI claims (% of pop)	$LICM@USECON^{\dagger}/LN16N@USECON^{\dagger}$	Log	U.S. Department of Labor and U.S. Bureau of Labor Statistics
Underutilization rate (% of labor force)	$(LEPTE@USECON^{\dagger} + LTU@USECON^{\dagger}) /LF@USECON^{\dagger}$	Log	U.S. Bureau of Labor Statistics
Per capita net worth (thous. 2017 USD)	$(1e6 \times (PA15CDA5 - PL15DDG5)@FFUNDS) /(DC@USNA/100 \times (LN16N@USECON \times 1e3))$	Log	Federal Reserve Board and U.S. Bureaus of Economic Analysis and Labor Statistics
Per capita income (thous. 2017 USD)	$(1e6 \times (YCOMPR + YPTP - YPTX - GRCSD)@USECON) /(DC@USNA/100 \times (LN16N@USECON \times 1e3))$	Log	U.S. Bureaus of Economic Analysis and Labor Statistics
Per capita consumption (thous. 2017 USD)	$(1e6 \times C@USECON)/(DC@USNA/100 \times (LN16N@USECON \times 1e3))$	Log	U.S. Bureaus of Economic Analysis and Labor Statistics

† denotes monthly time series that were averaged to obtain quarterly values.

Notes: The figure shows the 18 quarterly time series used in our Bayesian vector autoregression (BVAR) along with their Haver Analytics mnemonics (Mnemonic), any necessary data transformations (Trans.), and their original data sources. SPF stands for *Survey of Professional Forecasters*; CPI, Consumer Price Index; UI, unemployment insurance; and NA, not applicable.

Our vacancies time series is obtained by splicing together two series: 1) data on job openings and the number of unemployed from the BLS and 2) data on help-wanted ads and the number of unemployed from the Conference Board. In doing so, we follow Barnichon (2010). However, while qualitatively similar, our measure differs from the Barnichon measure for two reasons: 1) Barnichon also incorporated information from the Conference Board's Help Wanted OnLine ad index;[1] and 2) we follow the procedure used in the Chicago Fed National Activity Index and splice the vacancy–unemployment (V/U) ratio, rather than splice the vacancy rate. The vacancy rate is then separately derived from the spliced V/U measure. The differences between the two measures are shown in figure A2.

A2. Alternative monthly vacancy–unemployment ratios

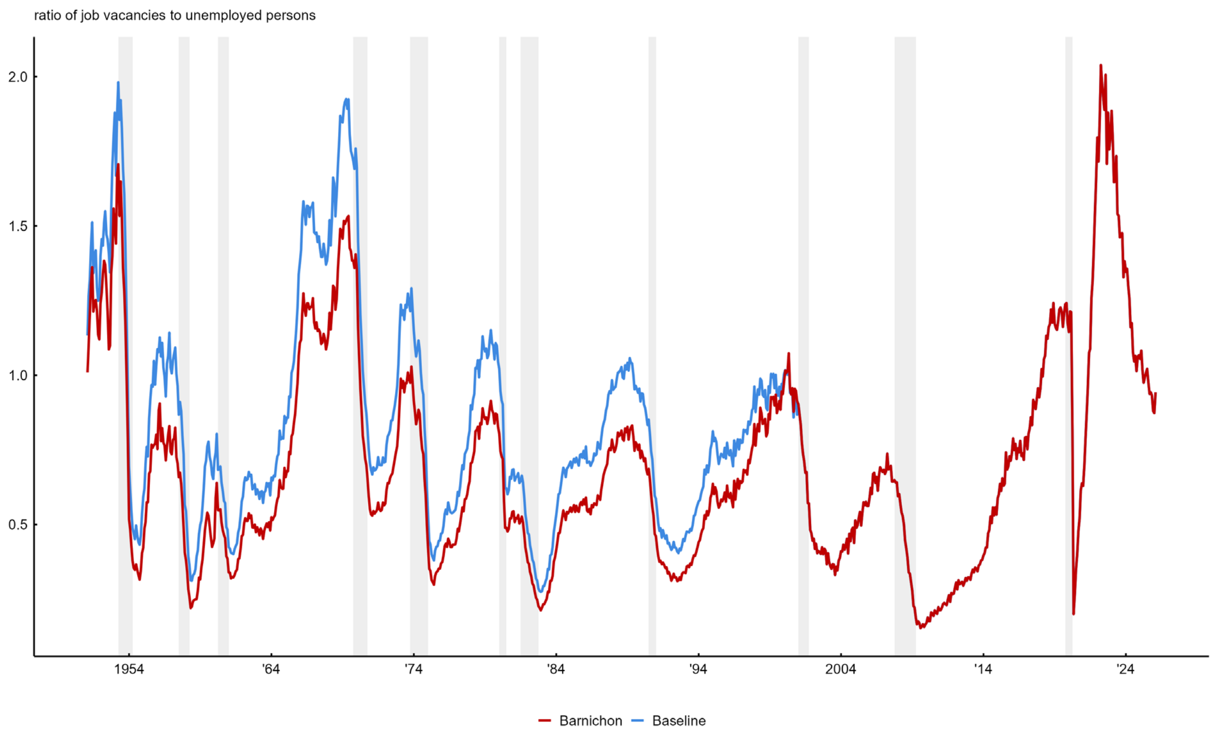

Notes: The figure shows two alternative time series for the U.S. vacancy–unemployment ratio: the one used in this article and referred to as baseline (blue) and an alternative referred to as Barnichon (red) from Barnichon (2010). The shaded periods correspond to U.S. recessions as defined by the National Bureau of Economic Research.

Sources: Authors' calculations based on data from Haver Analytics and Regis Barnichon.

[1] Subsequent research by Cajner and Ratner (2016) questioned the reliability of this information and later versions do not incorporate it—as seen in O'Trakoun (2022) and in the series from Regis Barnichon available for download here.

Appendix 2: BVAR

Our vector autoregression (VAR) takes the following general form (in log levels):

$$y_t = c + A_1 y_{t-1} + \ldots + A_5 y_{t-5} + u_t,$$
$$u_t \sim N(0, \Sigma),$$

where y_t is a 18×1 vector at each time period t , A_i are the five (i.e., five lags) 18×18 coefficient matrices summarizing how any individual data series responds to the ith lag of any other series in the system, and Σ is the covariance matrix of the independently and identically distributed (i.i.d.*)* vector autoregression model error terms u_t, which we assume are independent and identically normally distributed across time periods.

In what follows, we describe how to implement the prior that yields our Bayesian VAR framework.

Minnesota prior

The hierarchical shrinkage prior for our BVAR is the commonly used *Minnesota prior*, which holds that *each* time series exhibits random walk (with or without drift) behavior. This prior acts to *shrink* the VAR coefficients for each time series on its own first lag in A_1 toward a unit root and its further own lags in $A_2, \ldots, A_5$ and other time series' lags in $A_1, \ldots, A_5$ toward zero.

The Minnesota prior distribution for the VAR coefficients $A = (A_1, \ldots, A_5)$ is commonly specified as belonging to the following Normal-Inverse Wishart family

$$\Sigma \sim IW(\Psi, d),$$
$$\beta | \Sigma \sim N(b, \Sigma \otimes \Omega),$$

with β representing the $n \times np$ matrix of VAR coefficients and Σ corresponding to the $n \times n$ variance-covariance matrix of the VAR residuals. The degrees of freedom of the Inverse-Wishart distribution is set to $d = n + 2$ to guarantee the existence of a prior mean for Σ. In addition, Ψ is taken to be a diagonal matrix, with its diagonal vector ψ treated as hyperparameters to be estimated.[2]

Conditional on Σ, the prior for the VAR coefficients has the first and second moments

$$E\left[(A_s)_{ij} | \Sigma\right] = \begin{cases} 1 \text{ if } i = j \text{ and } s = 1 \\ 0 \text{ otherwise} \end{cases},$$
$$cov\left((A_s)_{ij}, (A_r)_{kl} | \Sigma\right) = \begin{cases} \lambda^2 \dfrac{1}{s^\alpha} \dfrac{\Sigma_{ik}}{\psi_j / (d - n - 1)} \text{ if } l = j \text{ and } r = s \\ 0 \text{ otherwise} \end{cases},$$

and is easily mapped into the Normal-Inverse Wishart structure described in this section for b and the Kronecker product $\Sigma \otimes \Omega$.[3]

[2] Chan (2021) contains a nice treatment of this prior in the larger context of the global-local shrinkage priors used in Bayesian linear regression. Lenza and Primiceri (2022) and Cascaldi-Garcia (2025) demonstrate how to extend this structure to include stochastic volatility and/or exogenous variables in order to address the Covid-19 pandemic period.

[3] Chan (2022) and others have sought to relax the symmetry assumption of the traditional Minnesota prior implied by the Kronecker structure. They show that while it is still possible to maintain conjugacy with such an equation-by-equation treatment of $\beta | \Sigma$, the required method of sampling differs from the initial dummy observation method used here.

In addition to λ and α, we treat ψ_j as hyperparameters to be estimated, consistent with the Minnesota prior described in Giannone et al. (2015). Based on past experience with this prior (i.e., Brave et al., 2019), doing so tends to improve near-term forecasting performance for measures of economic activity. However, it also comes with an increase in computational costs. Our approach to estimation described in the main text is tailored to this concern.

Long-run prior

In addition to the *Minnesota prior,* we also impose a set of a *long-run prior* (LR) beliefs on the relationships *between* time series guided by macroeconomic theory that states how various combinations of them are expected to co-move with each other over long periods of time. These combinations in our BVAR are produced by pre-multiplying y_t by the transformation matrix H shown in figure A3.

A3. Long-run prior transformation matrix, H

	L/N	E/N	V/N	U/N	J/N	Y/H	H/J	W/H	D/N	π^e	π	R	CUI/CE	IUI/N	(PT+U)/L	NW/N	I/N	C/N	
L/N	1	1	1	1	1														----> Labor force trend
E/N	-1	1																	----> Employment rate (E/L) is stationary
V/N	-1		1																----> Vacancy rate (V/L) is stationary
U/N	-1			1															----> Unemployment rate (U/L) is stationary
J/N	-1				1														----> Jobs rate (J/L) is stationary
Y/H					1	1	1	1	1										----> Labor productivity trend
H/J					1		1												----> Per capita hours (H/N) is stationary
W/H						-1		1											----> Labor share (W/Y) is stationary
D/N					-1	-1	-1		1										----> Investment share (I/Y) is stationary
π^e										1	1	1							----> Inflation trend
π										-1	1								----> Actual minus expected inflation is stationary
R										-1		1							----> Real interest rate is stationary
CUI/CE													1	1					----> Unemployment insurance take-up trend
IUI/N													-1	1					----> UI eligibility (IUI/CUI * CE/N) is stationary
(PT+U)/L															1				----> Underutilization rate is stationary
NW/N																1	1	1	----> Real consumption trend
I/N																-1	1		----> Income-to-net worth ratio trend
C/N																-0.2	-0.8	1	----> Stationary consumption-wealth ratio

Notes: The figure shows the transformation matrix (*H*) used to create the log-linear combinations in our 18-variable Bayesian vector autoregression (BVAR) representing the invariant long-run prior. The order of variables matches that found in figure A1.

Source: Authors' calculations based on data from Haver Analytics.

The long-run prior transformation matrix H reflects the fact that the time series in our BVAR are known to share a common trend or trends:

- There are two *common trends* among our labor market indicators embodying *long-run growth in the labor force* and the *long-run growth in labor productivity*.
- The first trend is identified from a log-linear combination of the labor force participation rate (L/N), the household employment-to-population ratio (E/N), per capita jobs (J/N), and per capita unemployment and payroll job vacancies (U/N & V/N). Our reference to this particular linear combination as capturing *long-run growth in the labor force* will become apparent shortly.
- The second trend is identified from a log-linear combination of real output and real compensation per hours worked (Y/H & W/H), the workweek (H/J), and per capita payroll jobs and investment (J/N & D/N). Notice that, in logs, this particular linear combination is equivalent to assuming that there exists a jointly shared trend in real output per hour and per capita real compensation and investment (Y/H & W/N & D/N); hence, we have our description of this linear combination as capturing *long-run growth in labor productivity.* In order for this relationship to hold, all three series are deflated by GDP prices.

- Certain combinations of labor market indicators are assumed to be *stationary*, or free of these trends: the household employment and payroll job rates (E/L & J/L); vacancy and unemployment rates (V/L & U/L); per capita hours worked (H/N); labor's share of income (W/Y); and investment's share of output (D/Y). The long-run priors express the belief that these linear combinations reflect *cointegrating* relationships as suggested by macroeconomic theory. Notice that the first four of them are achieved after scaling the variables that make up our first trend by the labor force participation rate (L/N). This is what lends the first trend its interpretation as reflecting *long-run growth in the labor force*.
- We also allow for a *nominal trend* in consumer price inflation (π), one year-ahead inflation expectations (π^e), and the one-year nominal interest rate (R), with the difference in actual and expected inflation ($\pi - \pi^e$) and the real interest rate ($R - \pi^e$) assumed to be stationary.
- In addition, we include in our BVAR the insured unemployment rate (CUI/CE), the share of eligible (or covered) employees claiming unemployment insurance (UI) benefits, and per capita initial unemployment insurance claims (IUI/N). Both series exhibit a slight downward drift that we acknowledge in our long-run prior as being driven by a *secular trend in UI take-up rates.* We also include a labor underutilization rate ($(PT + U)/L$) that in addition to unemployed individuals also includes those working part-time for self-reported economic reasons. This broader unemployment rate is assumed to be stationary as well.
- Finally, we include per capita real household net worth (excluding consumer durables) (NW/N), labor income (net of taxes and transfers) (I/N), and consumption (including consumer durables, C/N). Our long-run prior here follows Rudd and Whelan (2002) in specifying a *stationary consumption–wealth ratio*, with the appropriate price deflator in this case being Personal Consumption Expenditures (PCE) Price Index prices.

To demonstrate the role played by the long-run prior, it is convenient to rewrite the VAR in its "error correction" form, restating it in terms of lags of both (log) levels and (log) differences of y_t as

$$\Delta y_t = c + \Pi y_{t-1} + \Gamma_1 \Delta y_{t-1} + \ldots + \Gamma_4 \Delta y_{t-4} + u_t,$$

where Δ is the first difference operator, $\Pi = (A_1 + \ldots + A_5 - I)$, and $\Gamma_j = -\left(A_{j+1} + \ldots + A_5\right)$ for $j = 1, \ldots, 4$. Then, incorporating the (invertible) transformation matrix H into the rewritten equation yields

$$\Delta y_t = c + \Lambda \tilde{y}_{t-1} + \Gamma_1 \Delta y_{t-1} + \ldots + \Gamma_4 \Delta y_{t-4} + u_t,$$

where $\Lambda = \Pi H^{-1}$ and $\tilde{y}_{t-1} = H y_{t-1}$. Our priors for the long run, thus, embed *a priori* information on Λ from macroeconomic theory about the degree of cointegration among the series in y_t, as captured in the matrix Π and the cointegrating vectors summarized in the matrix H.

Following Giannone et al. (2019), we specify a long-run prior distribution on the loadings Γ_i conditional on H:

$$\Lambda_{\cdot i} | H_{i\cdot}, \Sigma \sim N\left(0, \tilde{\phi}_i(H_{i\cdot})\Sigma\right),$$

where $\Lambda_{\cdot i}$ denotes the ith column of Λ, and $\tilde{\phi}_i(H_{i\cdot})$ is a scalar hyperparameter that depends on the ith row of H. The prior is scaled by the variance of the errors Σ, is independent across i's, and is Gaussian. It implies that the columns of Π are correlated, with a correlation structure that depends on the choice of H and $\tilde{\phi}_i$, the latter of which we parameterize as follows:

$$\tilde{\phi}_i(H_{i\cdot}) = \frac{\phi_i^2}{(H_{i\cdot}\bar{y}_0)^2},$$

where ϕ_i is a scalar hyperparameter for the standard deviation of the prior on the elements of $\Lambda_{\cdot i}$ and $\bar{y}_0 = \frac{1}{p}\sum_{s=1}^{p} y_s$ is a column vector containing the average of the initial p observations of each variable.

For the most part, our long-run priors align with the statistical properties of the data they aim to influence. But this is not always true. For example, the stationarity of both per capita hours worked (per Francis and Ramey, 2009) and the labor share (per Elsby et al., 2013) is questionable. To address this fact, like Giannone et al. (2019), we allow the linear combinations that define per capita hours worked and the labor share to be potentially nonstationary in the prior.

Initial dummy observations

We implement both the Minnesota and long-run priors using the R package BVAR following the examples found in Kuschnig and Vashold (2021). While the long-run priors are not explicitly defined by the BVAR package, the functionality exists for the user to provide them. Conveniently, the long-run priors require only a very modest change be made to the authors' examples provided for the sum-of-coefficients (SOC) and single unit root (SUR) priors. The necessary modifications to be made are described in detail in section 6 of Giannone et al. (2019), and we summarize them in this section.

Following Giannone et al. (2019), an efficient means of simultaneously implementing both the Minnesota and long-run priors described previously is through a set of "initial dummy observations." In principle, both priors can be implemented jointly in our BVAR with five such blocks of dummy observations.

Each of the five blocks of dummy observations are shown here:

$$
\begin{aligned}
Y_d^1 &= \begin{pmatrix} \lambda * diag(\psi_1, \ldots, \psi_{18}) \\ 0_{72\times 18} \end{pmatrix} & X_d^1 &= (\lambda * J * diag(\psi_1, \ldots, \psi_{18}) \quad 0_{90\times 1}), \\
Y_d^2 &= \left(diag(\psi_1, \ldots, \psi_{18})\right) & X_d^2 &= (0_{18\times 91}), \\
Y_d^3 &= (0_{1\times 18}) & X_d^3 &= \left(0_{1\times 90} \quad \frac{1}{\epsilon}\right), \\
Y_d^4 &= (K^1) & X_d^4 &= (1_{1\times 5} \otimes K^1 \quad 0_{18-r\times 1}), \\
Y_d^5 &= (K^2) & X_d^5 &= \left(1_{1\times 5} \otimes K^2 \quad \frac{1}{\phi_{18-r+1}}\right),
\end{aligned}
$$

where $J = diag(1^{0.5\alpha}, \ldots, 5^{0.5\alpha})$, $K^1 = diag\left(\frac{H_{1,\cdot}\bar{y}_0}{\phi_1}, \ldots, \frac{H_{18-r,\cdot}\bar{y}_0}{\phi_{18-r}}\right)[H^{-1}]'_{1:18-r,\cdot}$, $K^2 = [H^{-1}]_{\cdot,18-r+1:18}\frac{H_{18=r+1:18,\cdot}\bar{y}_0}{\phi_{18-r+1}}$ and $\bar{y}_0 = \frac{1}{p}\sum_{s=1}^{5} y_{-s}$ is the average of the initial period's values for each time series in y_t. The first three blocks incorporate the Minnesota prior, while the fourth and fifth blocks incorporate Giannone et al. (2019) rotation-invariant long-run prior, meaning that they vary depending on whether or not they are applied to the nonstationary (fourth block) or stationary (fifth block) log-linear combinations of y_t implied by the rows of *H.*

Following Giannone et al. (2019), we assume that y_t is arranged such that the first $18 - r = 9$ rows of H correspond to the fourth block and the remaining $r = 9$ rows of H corresponds to the fifth block. A key difference in the fourth and fifth blocks is how the intercept terms are handled. The nonstationary combinations defined by the fourth block shrink the impacted intercepts toward zero, while the stationary combinations shrink them toward a finite value in sum at a rate proportional to $\frac{1}{\phi_{18-r+1}}$. To avoid conflict with the Minnesota prior's treatment of the intercepts, we, thus, impose a diffuse prior in the third block by setting $\epsilon = 1e^{-7}$.

We go one step further than the preceding notation and treat separately subsets of the stationary combinations corresponding to the labor market, consumption and investment, and inflation variables. This implies the set of hyperparameters ϕ_i with $i = 1, \ldots, 11$, applying one hyperparameter to each of the nonstationary combinations and another to the sum of the stationary combinations in each subset.

As Giannone et al. (2019) note, there is a close connection between the long-run prior and the SOC and SUR priors when $H = I_{18}$. The primary difference in this case stems from the fact that the long-run prior operates on a variable-by-variable basis, while the SOC and SUR priors are typically expressed as system-wide priors. To impose the SOC and SUR priors instead, redefine the fourth and fifth blocks as

$$Y_d^4 = (K^1) \qquad X_d^4 = (1_{1\times 5} \otimes K^1 \quad 0_{18\times 1}),$$
$$Y_d^5 = (K^2) \qquad X_d^5 = \left(1_{1\times 5} \otimes K^2 \quad \frac{1}{\delta}\right),$$

where $K^1 = diag\left(\frac{\bar{y}_0}{\mu}, \ldots, \frac{\bar{y}_0}{\mu}\right)$ and $K^2 = \frac{\bar{y}'_0}{\delta}$.

Hyperparameters

The *hyperparameters* $(\lambda, \alpha, \psi, \phi)$ govern the degree of shrinkage imposed by the Minnesota and long-run priors in our BVAR. Each establishes an *a priori* belief about how likely the data are to exhibit random walk behavior (λ), how predictive older versus more recent data are (α), and how the residiual variation of each series compares to each other (ψ), as well as the existence of the common trends and cointegrating relationships between indicators (ϕ).

Statistically, the hyperparameters express prior beliefs for the VAR equations on the first autocorrelation coefficients, the constants (or intercepts), the autoregressive lag coefficients, the cross-correlation coefficients, and the order of integration/cointegration of each time series. We follow Giannone et al. (2019) in setting $\alpha = 2$ and specifying hyperpriors for the remaining hyperparameters.

The Gamma hyperpriors for λ and ϕ have modes of 1 with standard deviations of 1 and 0.25, respectively. The hyperprior for ψ has shape and scale parameters of 1.2 and 0.5, respectively. It was calibrated to include the range of default values that are commonly used for ϕ. When we alternatively estimate the SOC and SUR prior hyperparameters μ and δ, we instead use the Gamma hyperpriors of Giannone et al. (2019) with modes and standard deviations of 1 for each.

The hyperpriors for our joint Minnesota and long-run prior are similar to those used by Giannone et al. (2019), but are not identical. In our experience, when estimating ψ in the Minnesota prior block as well, our hyperpriors for λ and ϕ were more appropriately centered and exhibited better Markov chain Monte Carlo (MCMC) convergence. To further aid in convergence, we also employ a Metropolis-Hastings step in the MCMC procedure that makes use of the BVAR package's ability to finely tune the initial acceptance rate and the scale of the Hessian calculation.

Estimated values for the hyperparameters of our baseline model along with convergence diagnostics (Gelman's $\hat{R}$) are shown in figure A4.

A4. Hyperparameter values

	Minnesota & LR prior		Minnesota & SOC/SUR prior	
	Median	$\hat{R}$	Median	$\hat{R}$
Minnesota prior				
λ	2.36	1.02	1.12	1.01
α	2.00		2.00	
μ			0.03	1.00
δ			0.19	1.00
ψ_1	0.01	1.02	0.00	1.01
ψ_2	0.01	1.02	0.01	1.03
ψ_3	0.21	1.01	0.06	1.02
ψ_4	0.11	1.01	0.03	1.01
ψ_5	0.01	1.01	0.01	1.01
ψ_6	0.01	1.01	0.01	1.01
ψ_7	0.01	1.02	0.01	1.01
ψ_8	0.01	1.00	0.01	1.01
ψ_9	0.16	1.01	0.05	1.01
ψ_{10}	0.01	1.02	0.00	1.02
ψ_{11}	0.01	1.01	0.01	1.01
ψ_{12}	0.01	1.01	0.01	1.01
ψ_{13}	0.37	1.01	0.08	1.00
ψ_{14}	0.57	1.01	0.07	1.01
ψ_{15}	0.10	1.02	0.03	1.01
ψ_{16}	0.07	1.01	0.03	1.02
ψ_{17}	0.03	1.01	0.01	1.00
ψ_{18}	0.01	1.01	0.01	1.01
Long-run prior				
ϕ_1	0.27	1.00		
ϕ_2	0.65	1.01		
ϕ_3	1.04	1.01		
ϕ_4	0.91	1.01		
ϕ_5	0.95	1.01		
ϕ_6	0.94	1.01		
ϕ_7	0.85	1.01		
ϕ_8	0.89	1.00		
ϕ_9	0.94	1.01		
ϕ_{10}	0.95	1.01		
ϕ_{11}	0.48	1.00		
Overall		1.05		1.05

Notes: The figure shows the median hyperparameter values and their potential scale reduction factors $\hat{R}$ for varying prior structures. The Markov chain Monte Carlo (MCMC) estimation uses 24 chains each with 125,000 draws with a burn-in of 25,000 and thin rate of 20. The subscripts for the error variance hyperparameters follow the variables listed in the rows of figure A3. The subscripts for the invariance long-run prior hyperparameters follow the relevant blocks of the *H* matrix in figure A3.

Sources: Authors' calculations based on data from Haver Analytics.

Forecasts

Figures A5–A9 feature our baseline BVAR's median unconditional forecasts (Minnesota & LR prior) and 70% prediction intervals along with the median unconditional forecasts for an alternative BVAR specification described mainly in this appendix 2 (Minnesota & SOC/SUR prior).

A5. Labor market slack projections

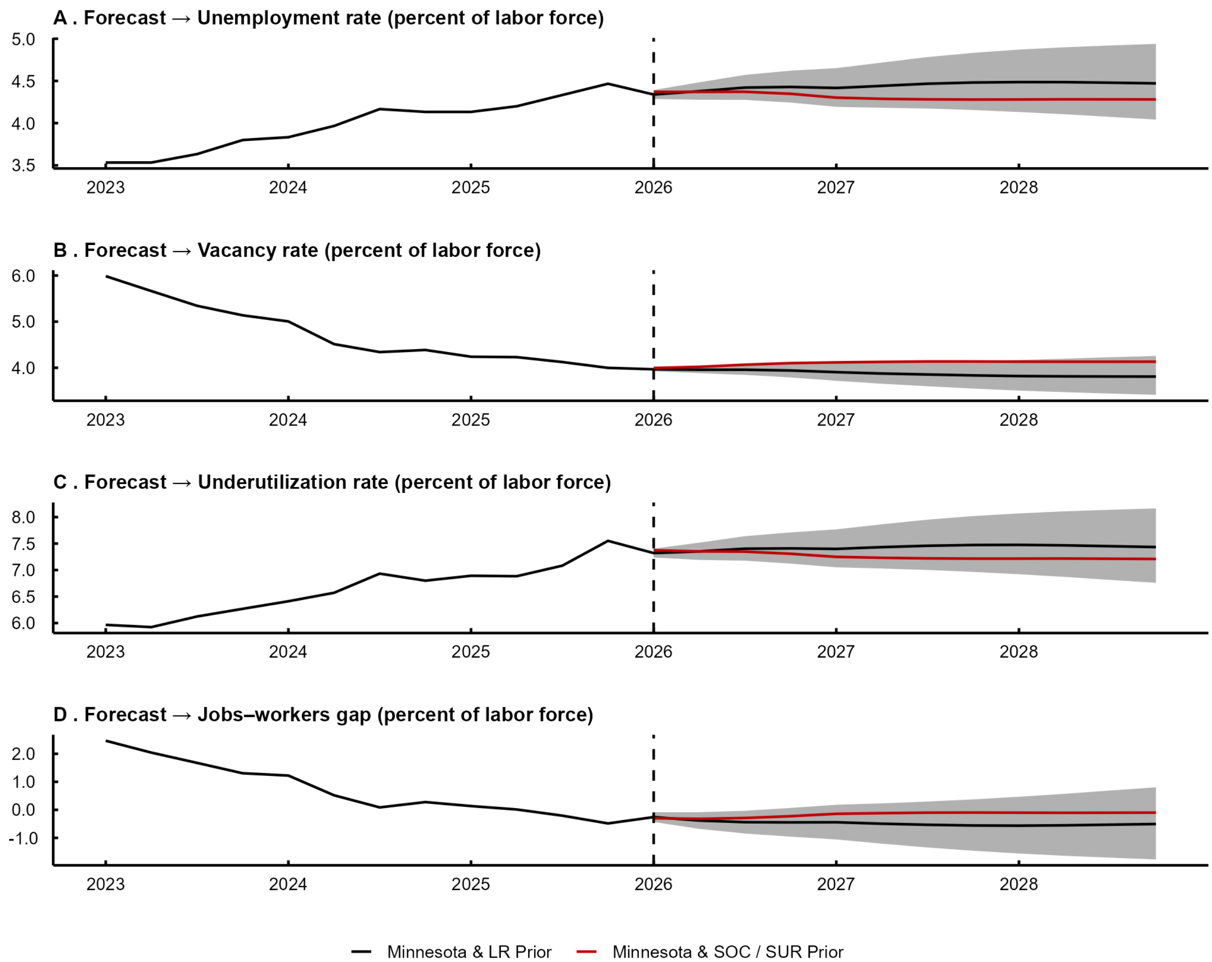

Notes: The figure shows median projections (black lines) and 70% prediction intervals (gray bands) for labor market slack measures (percent of the labor force) from a Bayesian vector autoregression (BVAR) with long-run priors (Minnesota & LR prior). The area to the right of the vertical dashed line in each panel of the figure corresponds with the forecast period beginning in 2026:Q1. The red lines report median projections from a BVAR with alternative priors (Minnesota & SOC/SUR prior).

Sources: Authors' calculations based on data from Haver Analytics.

A6. Labor force projections

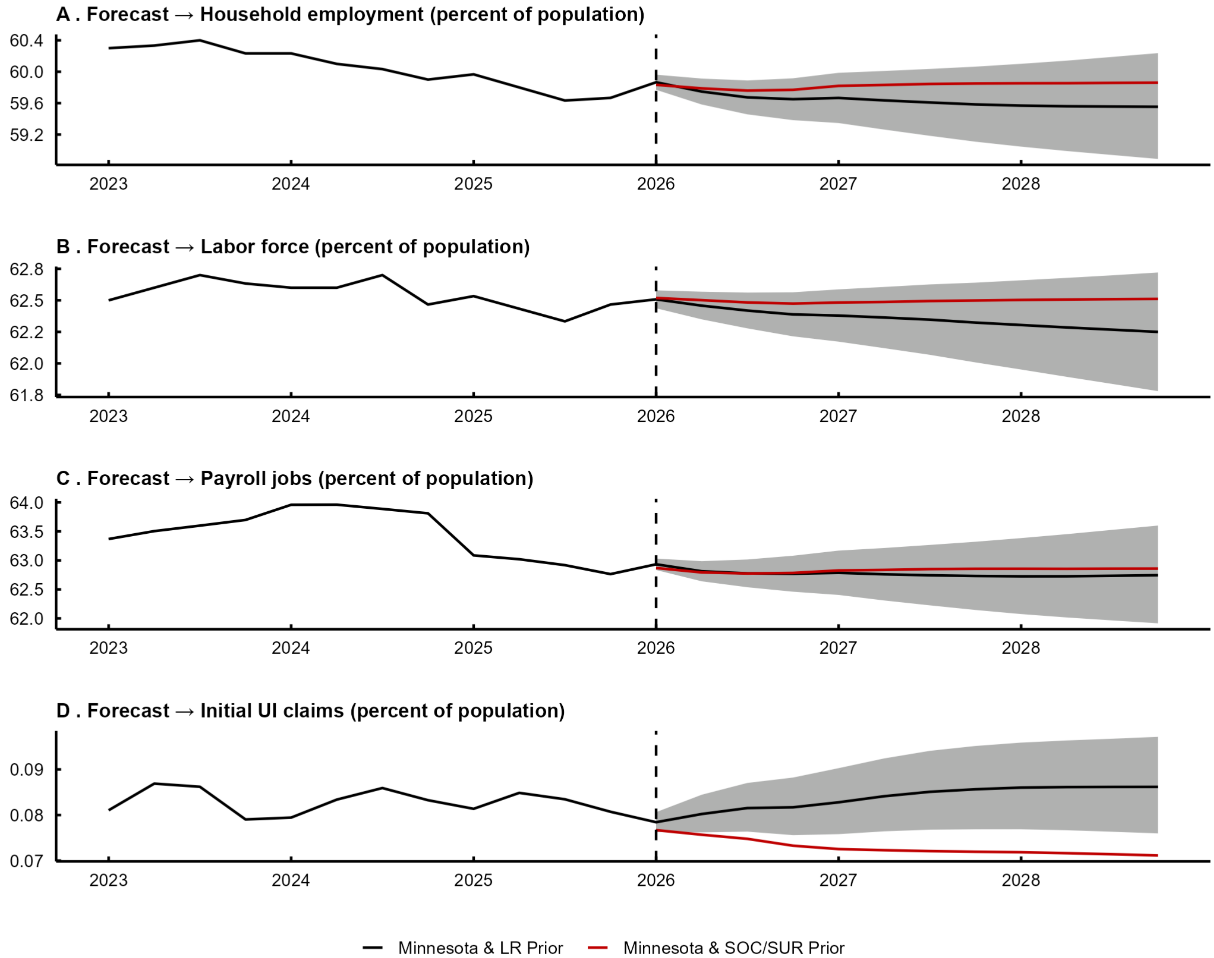

Notes: The figure shows median projections (black lines) and 70% prediction intervals (gray bands) for per capita (percent of the civilian noninstitutional population aged 16 and over) measures of the labor market from a Bayesian vector autoregression (BVAR) with long-run priors (Minnesota & LR prior). 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 red lines report median projections from a BVAR with alternative priors (Minnesota & SOC/SUR prior).

Sources: Authors’ calculations based on data from Haver Analytics.

A7. Labor productivity projections

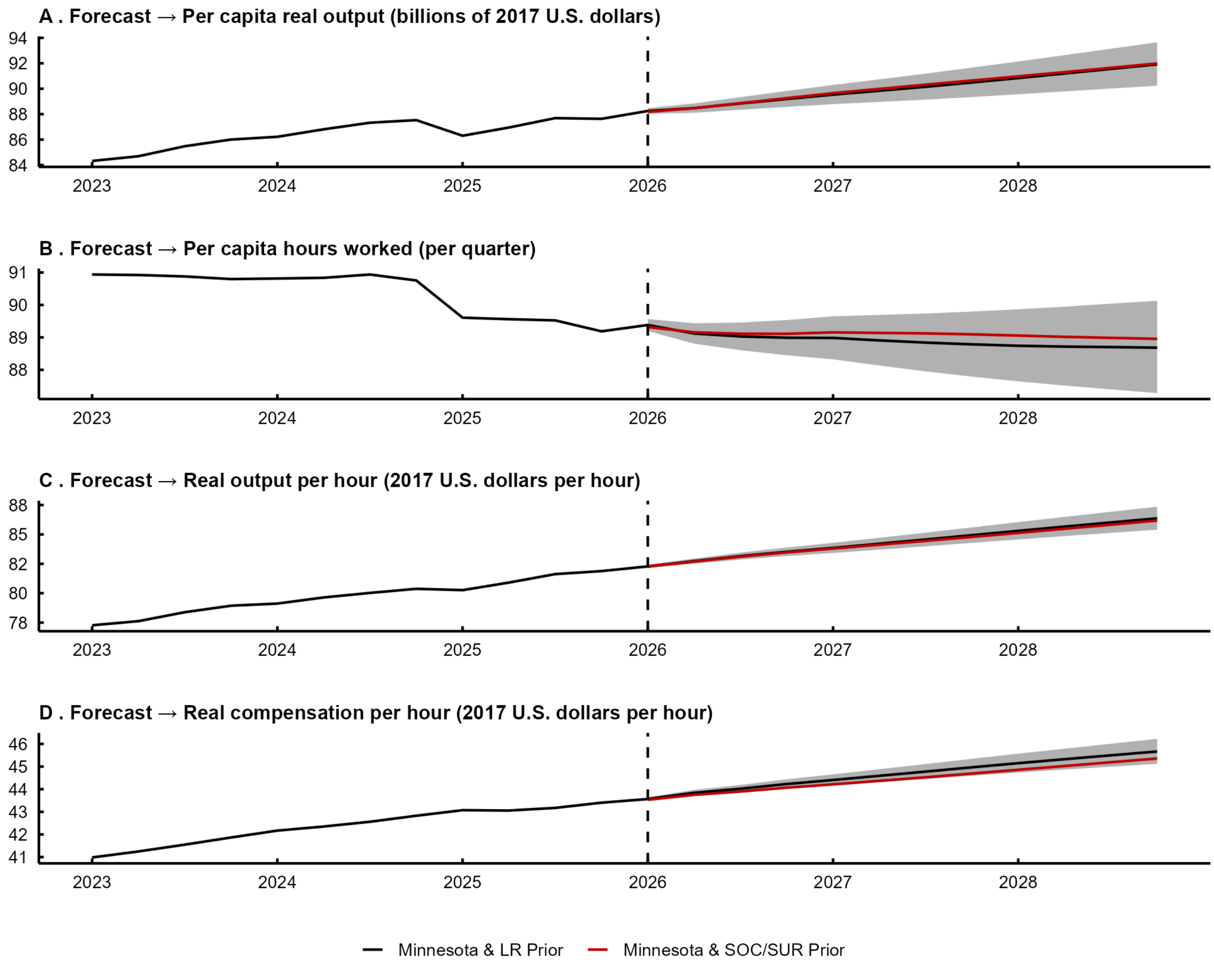

Notes: The figure shows median projections (black lines) and 70% prediction intervals (gray bands) for labor productivity measures from a Bayesian vector autoregression (BVAR) with long-run priors (Minnesota & LR prior). The area to the right of the vertical dashed line in each panel of the figure corresponds with the forecast period beginning in 2026:Q1. The red lines report median projections from a BVAR with alternative priors (Minnesota & SOC/SUR prior).

Sources: Authors' calculations based on data from Haver Analytics.

A8. Inflation and interest rate projections

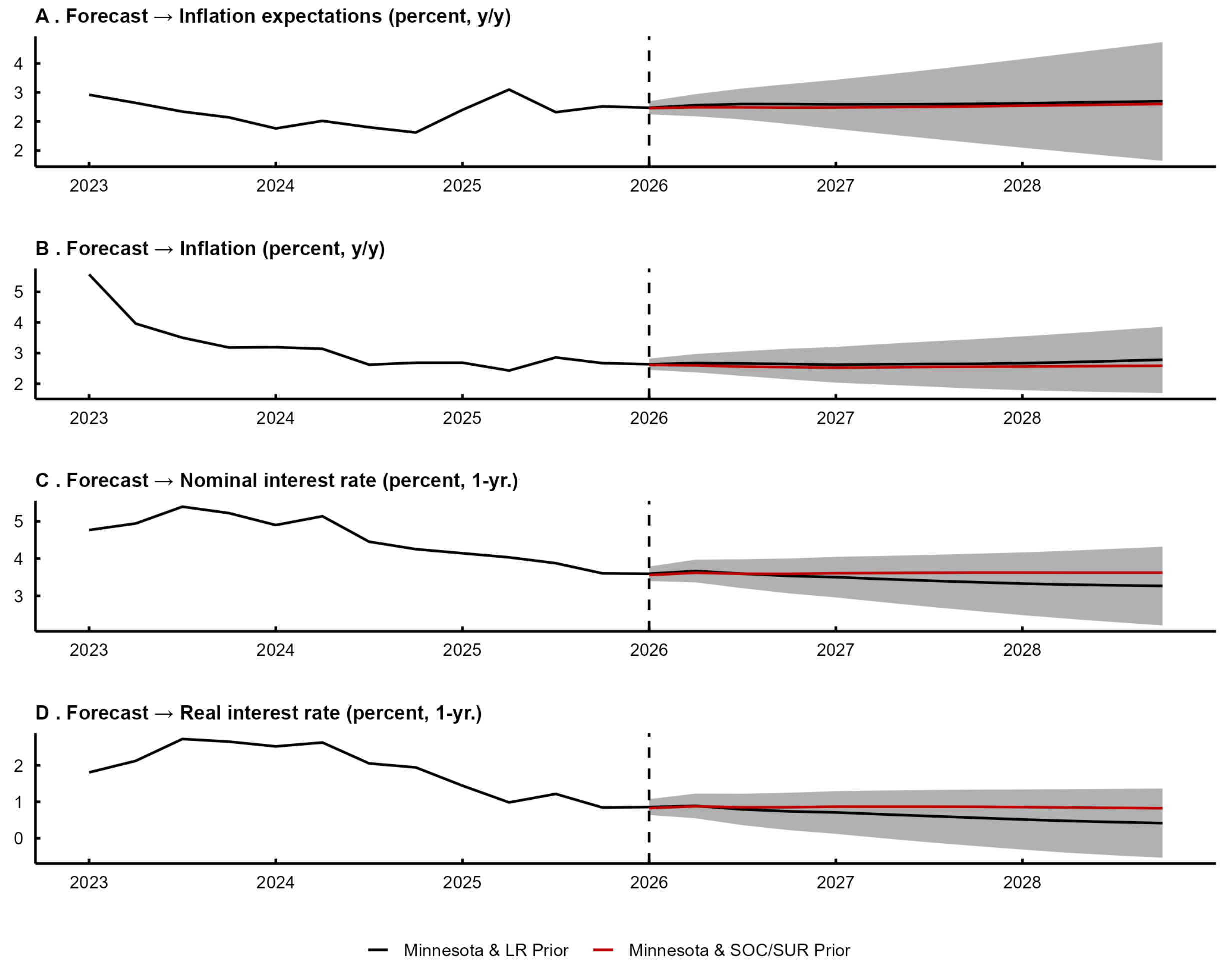

Notes: The figure shows median projections (black lines) and 70% prediction intervals (gray bands) for the inflation and interest rate measures from a Bayesian vector autoregression (BVAR) with long-run priors (Minnesota & LR prior). The area to the right of the vertical dashed line in each panel of the figure corresponds with the forecast period beginning in 2026:Q1. The red lines report median projections from a BVAR with alternative priors (Minnesota & SOC/SUR prior).

Sources: Authors' calculations based on data from Haver Analytics.

A9. Consumption and investment projections

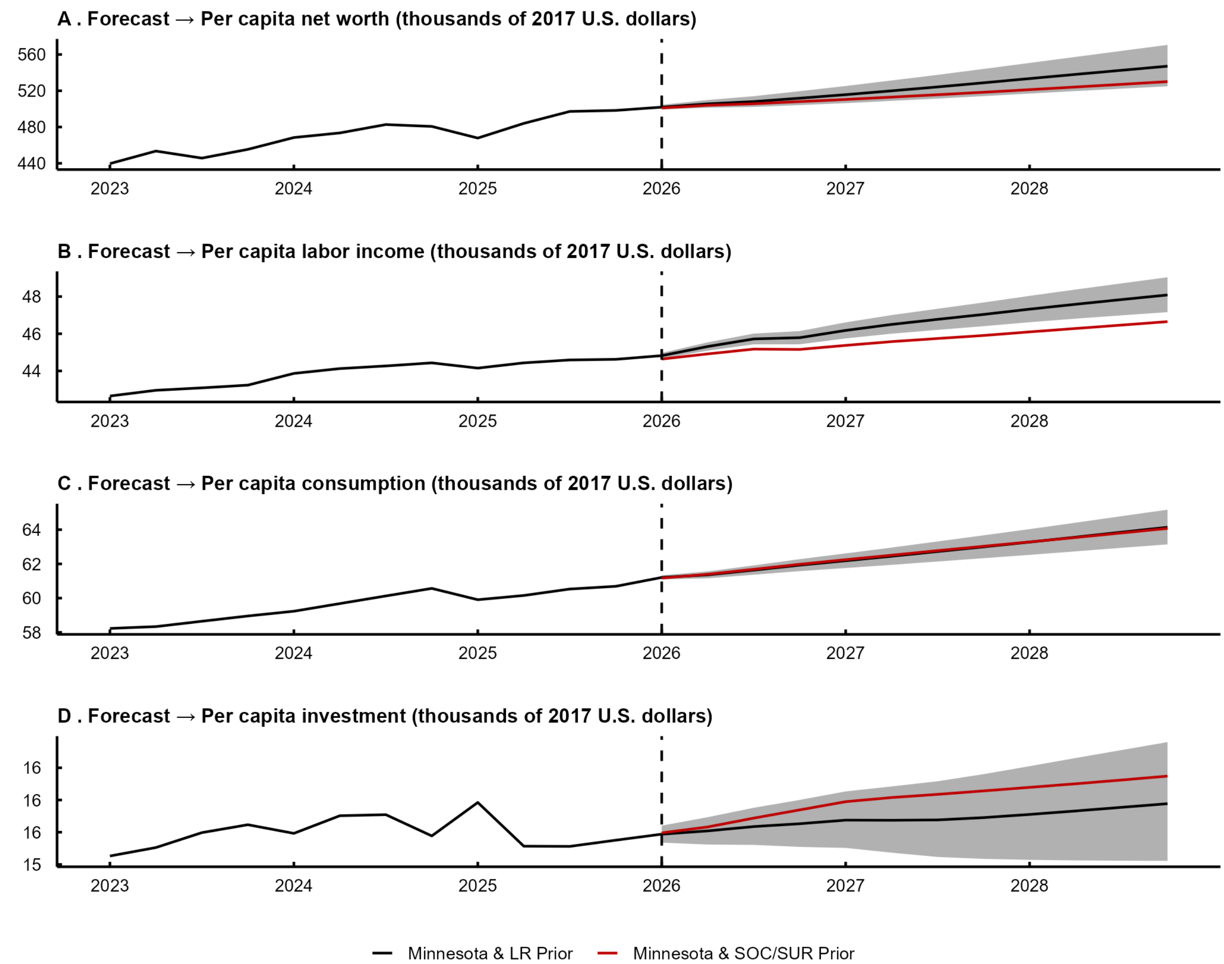

Notes: The figure shows median projections (black lines) and 70% prediction intervals (gray bands) for the consumption, income, and wealth measures from a Bayesian vector autoregression (BVAR) with long-run priors (Minnesota & LR prior). The area to the right of the vertical dashed line in each panel of the figure corresponds with the forecast period beginning in 2026:Q1. The red lines report median projections from a BVAR with alternative priors (Minnesota & SOC/SUR prior).

Sources: Authors' calculations based on data from Haver Analytics.

Appendix 3: BSVAR

Figure A10 summarizes the sign restrictions imposed on the contemporaneous impulse responses in our 18-variable Bayesian structural vector autoregression (BSVAR). We identify only five types of labor supply and demand shocks in this way, leaving the remaining 13 shocks unidentified.

A10. Sign restriction matrix

	LS	ME	WB	AS	AD
L/N	+	+	-	+	+
E/N	+	+	+	+	+
V/N	+	-	+	+	+
U/N	+	-	-	-	-
J/N	+	+	+	+	+
Y/H	NA	NA	NA	NA	NA
H/J	NA	NA	NA	NA	NA
W/H	-	-	-	+	-
D/N	NA	NA	NA	NA	NA
π^e	NA	NA	NA	NA	NA
π	-	-	-	-	+
R	NA	NA	NA	NA	NA
CUI/CE	NA	NA	NA	NA	NA
IUI/N	NA	NA	NA	NA	NA
(PT+U)/L	+	-	-	-	-
NW/N	NA	NA	NA	NA	NA
I/N	NA	NA	NA	NA	NA
C/N	NA	NA	NA	NA	NA

Notes: The figure summarizes the sign restrictions applied to the impulse responses of our 18-variable Bayesian vector autoregression (BVAR). The order of variables (rows) matches that found in figure A1 and borrows the abbreviations from figure A3. Each column reports the signs (+ / – / NA, indicating not applicable) imposed on the contemporaneous impacts for variables from one of five types of shocks: 1) labor supply (LS), 2) match efficiency (ME), 3) wage bargaining (WB), 4) aggregate supply (AS), and 5) aggregate demand (AD).

Figure A11 depicts median impulse responses over a five-year forecast horizon with error bands corresponding to 70% credible intervals. We normalize the sign of the five sign-identified labor supply and demand shocks such that each produces an increase in employment on impact.

A11. Labor supply and demand shocks

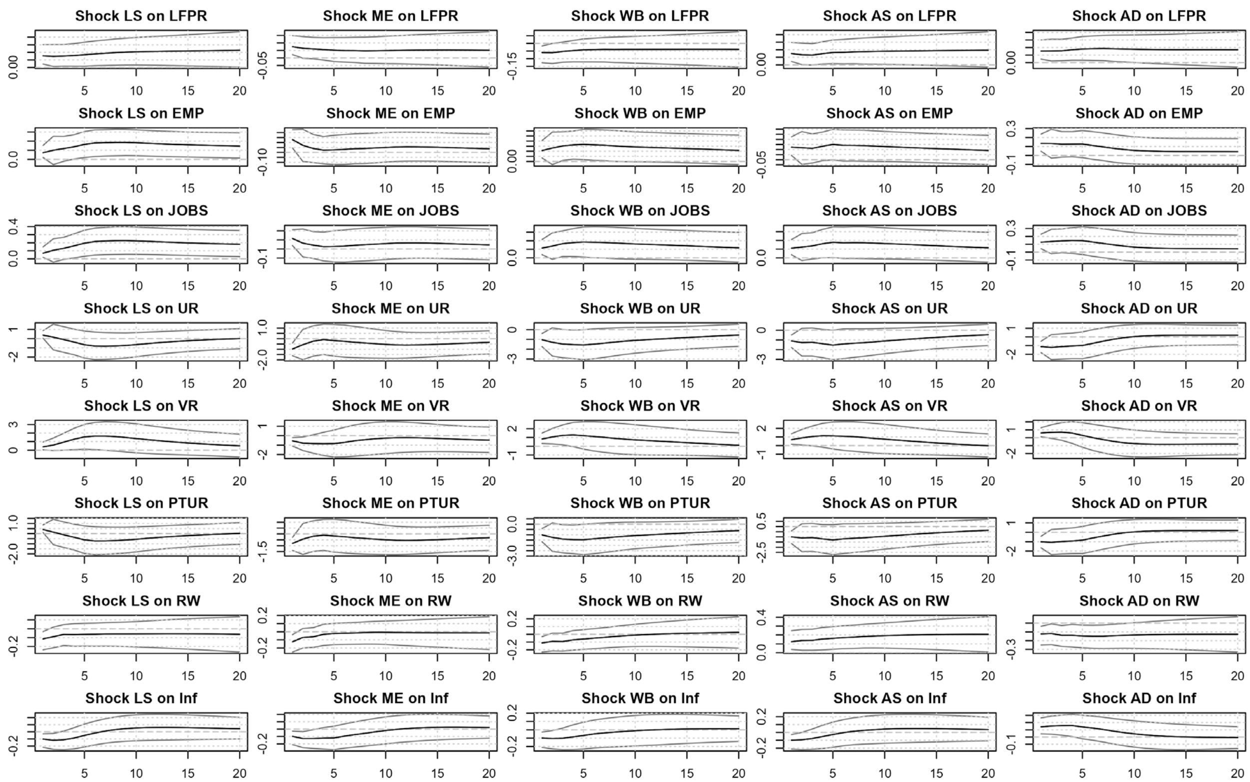

Notes: The figure shows the sign-restricted impulse responses functions for the labor force participation rate (LFPR), employment-to-population ratio (EMP), per capita jobs (JOBS), the unemployment (UR) and vacancy (VR) rates, the underutilization rate (PTUR), real compensation per hour (RW), and consumer price inflation (Inf) 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). The x-axis measures time in quarters, while the y-axis units are percentage points.

Sources: Authors' calculations based on data from Haver Analytics.

Labor supply shocks persistently change both labor force participation and employment, yet their effects on unemployment and vacancies are statistically somewhat uncertain. In the short run, they tend to be a "shifter" of the Beveridge curve with the unemployment rate falling after impact and the vacancy rate rising. However, over time, the unemployment rate also falls while the vacancy rate continues to rise, making them more of a "slider" along the Beveridge curve in the medium run.

In contrast, match efficiency shocks are an unambiguous shifter of the Beveridge curve, with the labor force participation rate and employment rising and the unemployment and vacancy rates falling on impact. The effects of these shocks on all of the variables are considerably more transitory than the effects of labor supply shocks, typically dying out over a period of one to two years.

Wage bargaining shocks, like labor market shocks, lead to persistent changes in employment and the unemployment and vacancy rates, but are accompanied by a short-run decline in labor force participation, making them unique among the five shocks that we consider. Aggregate activity shocks, in contrast, are also accompanied by much larger movements along the Beveridge curve and in employment, but with persistent increases in the labor force participation rate.

Appendix 4: Shiny app

An interactive application that allows the user to further explore our results is included for download with this article. This app was developed in R Shiny. Opening the app.R file in RStudio and clicking “Run App” starts up the Shiny app as long as all dependencies are available locally. This includes the functions in the src/ folder, the R packages called by the library function, and the labor_market_bvar.qs2 file provided with the app.

The app comes packaged with estimated model parameters and historical data that match the article for faster initialization. However, the forecasts shown in figures A5–A9 can be updated to condition on any incoming data and the scenarios shown in figures 7–11 of the main text can be altered as desired.

Welcome!

This tab explains how to use the application.

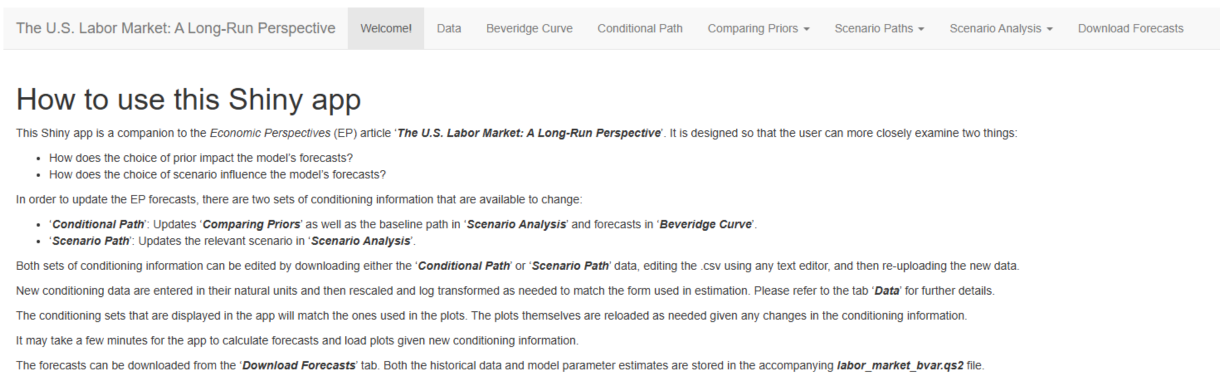

The U.S. Labor Market: A Long-Run Perspective | Welcome! | Data | Beveridge Curve | Conditional Path | Comparing Priors ▾ | Scenario Paths ▾ | Scenario Analysis ▾ | Download Forecasts

How to use this Shiny app

This Shiny app is a companion to the *Economic Perspectives* (EP) article '***The U.S. Labor Market: A Long-Run Perspective***'. It is designed so that the user can more closely examine two things:

- How does the choice of prior impact the model's forecasts?
- How does the choice of scenario influence the model's forecasts?

In order to update the EP forecasts, there are two sets of conditioning information that are available to change:

- '***Conditional Path***': Updates '***Comparing Priors***' as well as the baseline path in '***Scenario Analysis***' and forecasts in '***Beveridge Curve***'.
- '***Scenario Path***': Updates the relevant scenario in '***Scenario Analysis***'.

Both sets of conditioning information can be edited by downloading either the '***Conditional Path***' or '***Scenario Path***' data, editing the .csv using any text editor, and then re-uploading the new data.

New conditioning data are entered in their natural units and then rescaled and log transformed as needed to match the form used in estimation. Please refer to the tab '***Data***' for further details.

The conditioning sets that are displayed in the app will match the ones used in the plots. The plots themselves are reloaded as needed given any changes in the conditioning information.

It may take a few minutes for the app to calculate forecasts and load plots given new conditioning information.

The forecasts can be downloaded from the '***Download Forecasts***' tab. Both the historical data and model parameter estimates are stored in the accompanying ***labor_market_bvar.qs2*** file.

Data

This tab serves as a guide for how to update the data that later forecasts are conditioned upon.

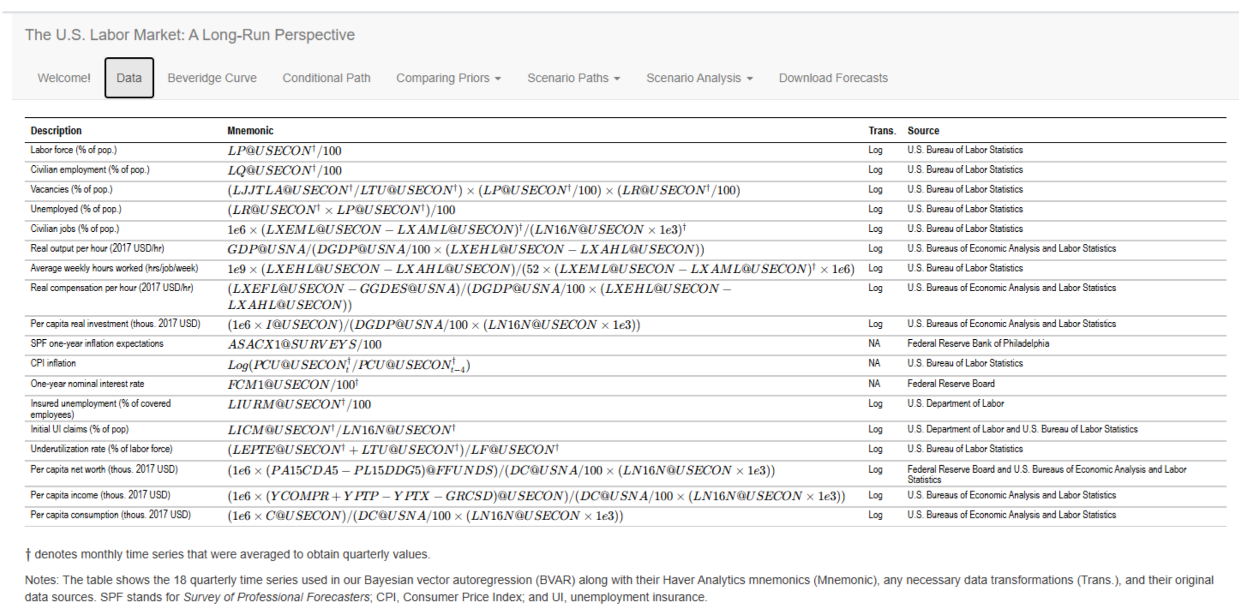

The U.S. Labor Market: A Long-Run Perspective

Welcome! | Data | Beveridge Curve | Conditional Path | Comparing Priors ▾ | Scenario Paths ▾ | Scenario Analysis ▾ | Download Forecasts

Description	Mnemonic	Trans.	Source
Labor force (% of pop.)	$LP@USECON^{\dagger}/100$	Log	U.S. Bureau of Labor Statistics
Civilian employment (% of pop.)	$LQ@USECON^{\dagger}/100$	Log	U.S. Bureau of Labor Statistics
Vacancies (% of pop.)	$(LJJTLA@USECON^{\dagger}/LTU@USECON^{\dagger}) \times (LP@USECON^{\dagger}/100) \times (LR@USECON^{\dagger}/100)$	Log	U.S. Bureau of Labor Statistics
Unemployed (% of pop.)	$(LR@USECON^{\dagger} \times LP@USECON^{\dagger})/100$	Log	U.S. Bureau of Labor Statistics
Civilian jobs (% of pop.)	$1e6 \times (LXEML@USECON - LXAML@USECON)^{\dagger}/(LN16N@USECON \times 1e3)^{\dagger}$	Log	U.S. Bureau of Labor Statistics
Real output per hour (2017 USD/hr)	$GDP@USNA/(DGDP@USNA/100 \times (LXEHL@USECON - LXAHL@USECON))$	Log	U.S. Bureaus of Economic Analysis and Labor Statistics
Average weekly hours worked (hrs/job/week)	$1e9 \times (LXEHL@USECON - LXAHL@USECON)/(52 \times (LXEML@USECON - LXAML@USECON)^{\dagger} \times 1e6)$	Log	U.S. Bureau of Labor Statistics
Real compensation per hour (2017 USD/hr)	$(LXEFL@USECON - GGDES@USNA)/(DGDP@USNA/100 \times (LXEHL@USECON - LXAHL@USECON))$	Log	U.S. Bureaus of Economic Analysis and Labor Statistics
Per capita real investment (thous. 2017 USD)	$(1e6 \times I@USECON)/(DGDP@USNA/100 \times (LN16N@USECON \times 1e3))$	Log	U.S. Bureaus of Economic Analysis and Labor Statistics
SPF one-year inflation expectations	$ASACX1@SURVEYS/100$	NA	Federal Reserve Bank of Philadelphia
CPI inflation	$Log(PCU@USECON_t^{\dagger}/PCU@USECON_{t-4}^{\dagger})$	NA	U.S. Bureau of Labor Statistics
One-year nominal interest rate	$FCM1@USECON/100^{\dagger}$	NA	Federal Reserve Board
Insured unemployment (% of covered employees)	$LIURM@USECON^{\dagger}/100$	Log	U.S. Department of Labor
Initial UI claims (% of pop)	$LICM@USECON^{\dagger}/LN16N@USECON^{\dagger}$	Log	U.S. Department of Labor and U.S. Bureau of Labor Statistics
Underutilization rate (% of labor force)	$(LEPTE@USECON^{\dagger} + LTU@USECON^{\dagger})/LF@USECON^{\dagger}$	Log	U.S. Bureau of Labor Statistics
Per capita net worth (thous. 2017 USD)	$(1e6 \times (PA15CDA5 - PL15DDG5)@FFUNDS)/(DC@USNA/100 \times (LN16N@USECON \times 1e3))$	Log	Federal Reserve Board and U.S. Bureaus of Economic Analysis and Labor Statistics
Per capita income (thous. 2017 USD)	$(1e6 \times (YCOMPR + YPTP - YPTX - GRCSD)@USECON)/(DC@USNA/100 \times (LN16N@USECON \times 1e3))$	Log	U.S. Bureaus of Economic Analysis and Labor Statistics
Per capita consumption (thous. 2017 USD)	$(1e6 \times C@USECON)/(DC@USNA/100 \times (LN16N@USECON \times 1e3))$	Log	U.S. Bureaus of Economic Analysis and Labor Statistics

† denotes monthly time series that were averaged to obtain quarterly values.

Notes: The table shows the 18 quarterly time series used in our Bayesian vector autoregression (BVAR) along with their Haver Analytics mnemonics (Mnemonic), any necessary data transformations (Trans.), and their original data sources. SPF stands for *Survey of Professional Forecasters*; CPI, Consumer Price Index; and UI, unemployment insurance.

Beveridge Curve

This tab shows the unemployment rate versus vacancy rate from the Minnesota and long-run prior model's forecast. The forecasts can change based on the conditional path input.

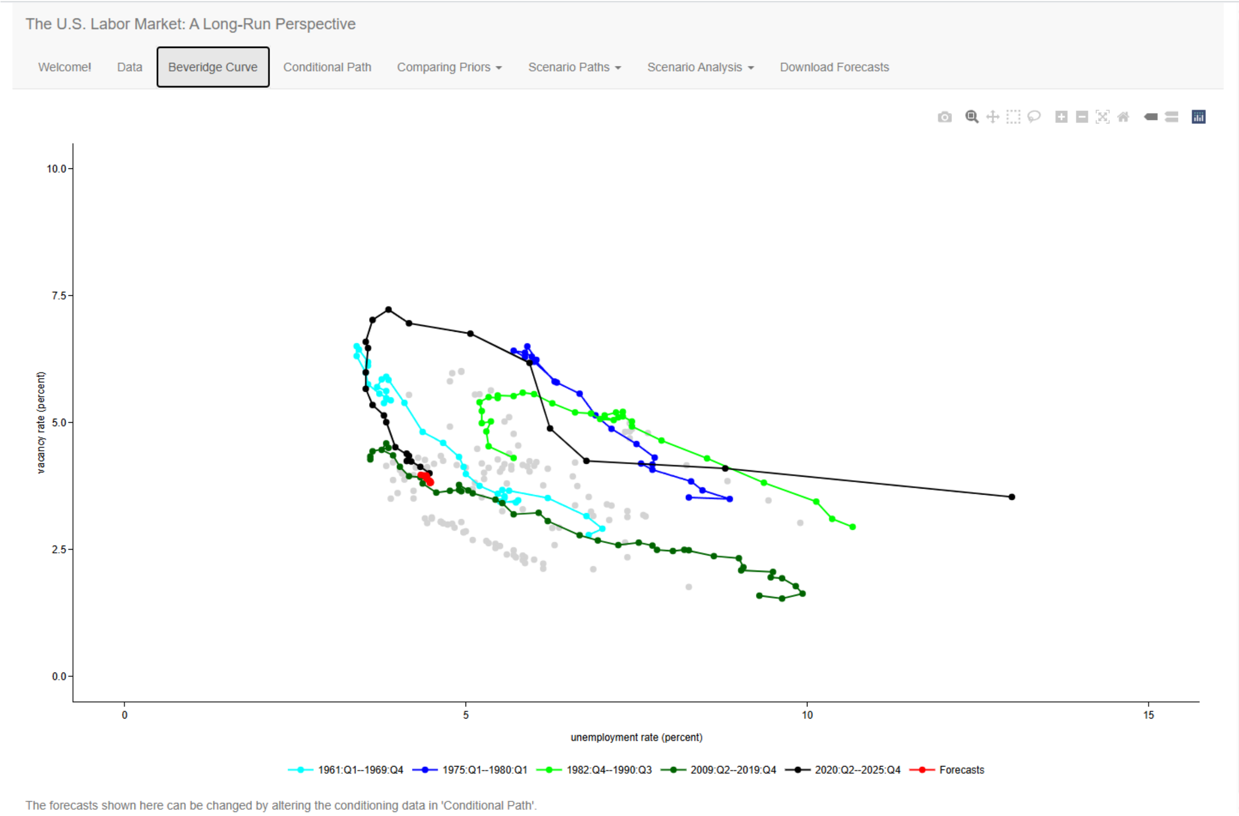

Conditional Path

This tab displays the current conditional path for the baseline model forecast. This path impacts the plots in the tabs called Beveridge Curve, Comparing Priors, and Scenario Analysis. Best practice for changing the path is to use the Download button on this tab to download the path from the article. Then, modify the .csv file in your text editor of choice, and upload the new conditional forecast paths.

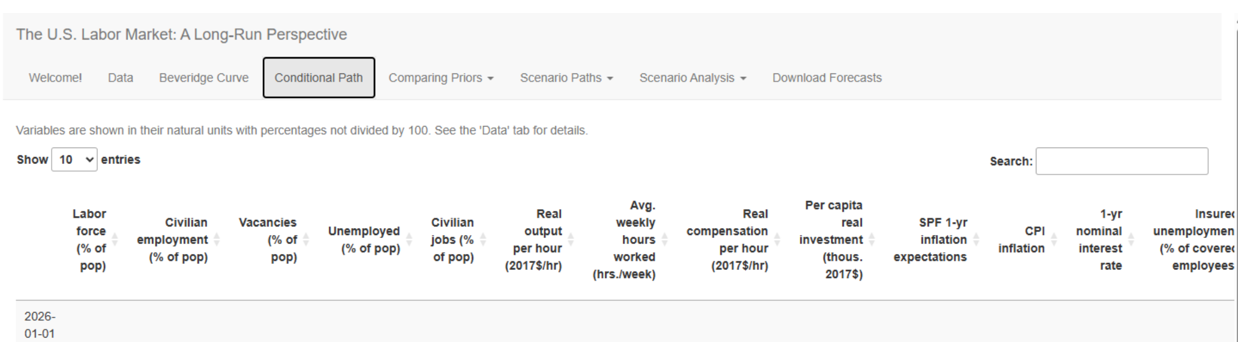

Comparing Priors

This tab displays forecasts based on the Conditional Path inputs under two model specifications: 1) the model with Minnesota and long-run priors and 2) the model with the Minnesota and SOC/SUR priors. These plots will update if the Conditional Path tab is changed.

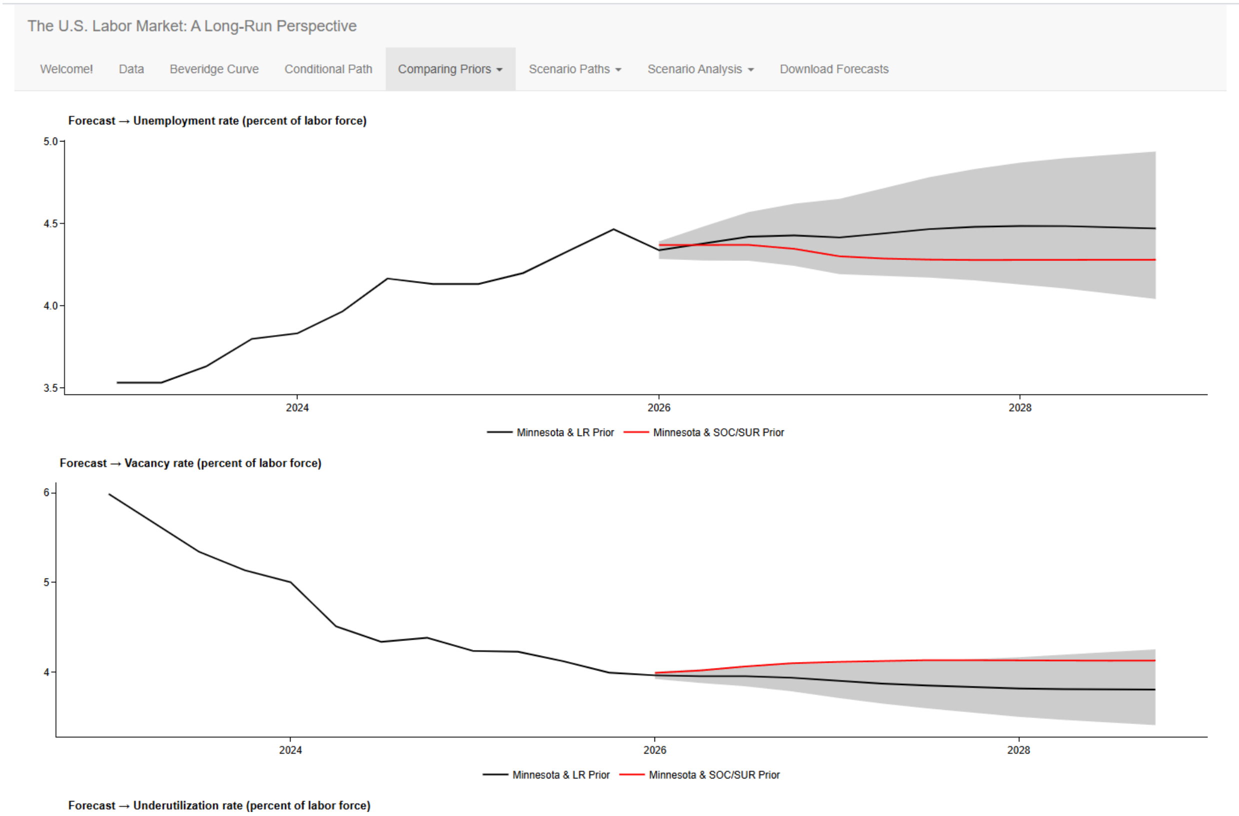

Scenario Path

This tab displays the conditioning paths used to construct the scenario forecasts in the Scenario Analysis tab. Best practice for changing a scenario path is to use the Download button on this tab to download the path from the article. Then, modify the .csv file in your text editor of choice, and upload the new conditional forecasting path for the scenario.

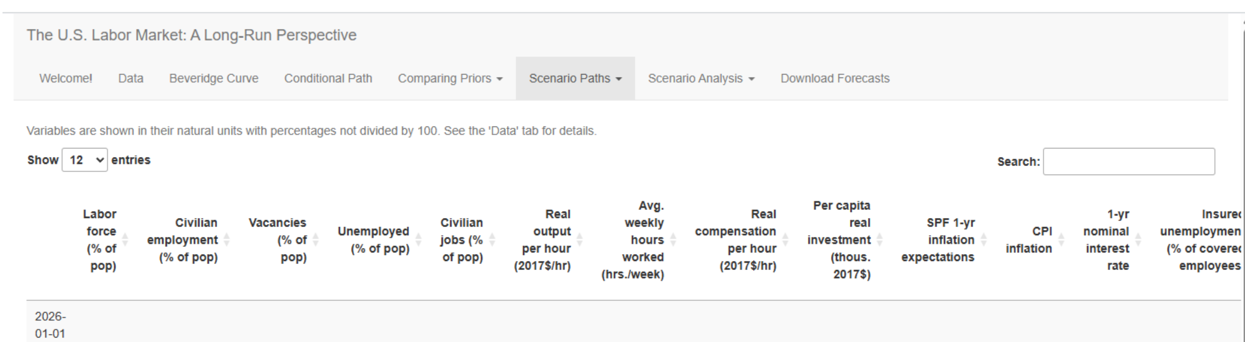

Scenario Analysis

This tab displays conditional forecasts based on the Conditional Path and Scenario Path inputs. These plots will update if either the Conditional Path or Scenario Path inputs are changed.

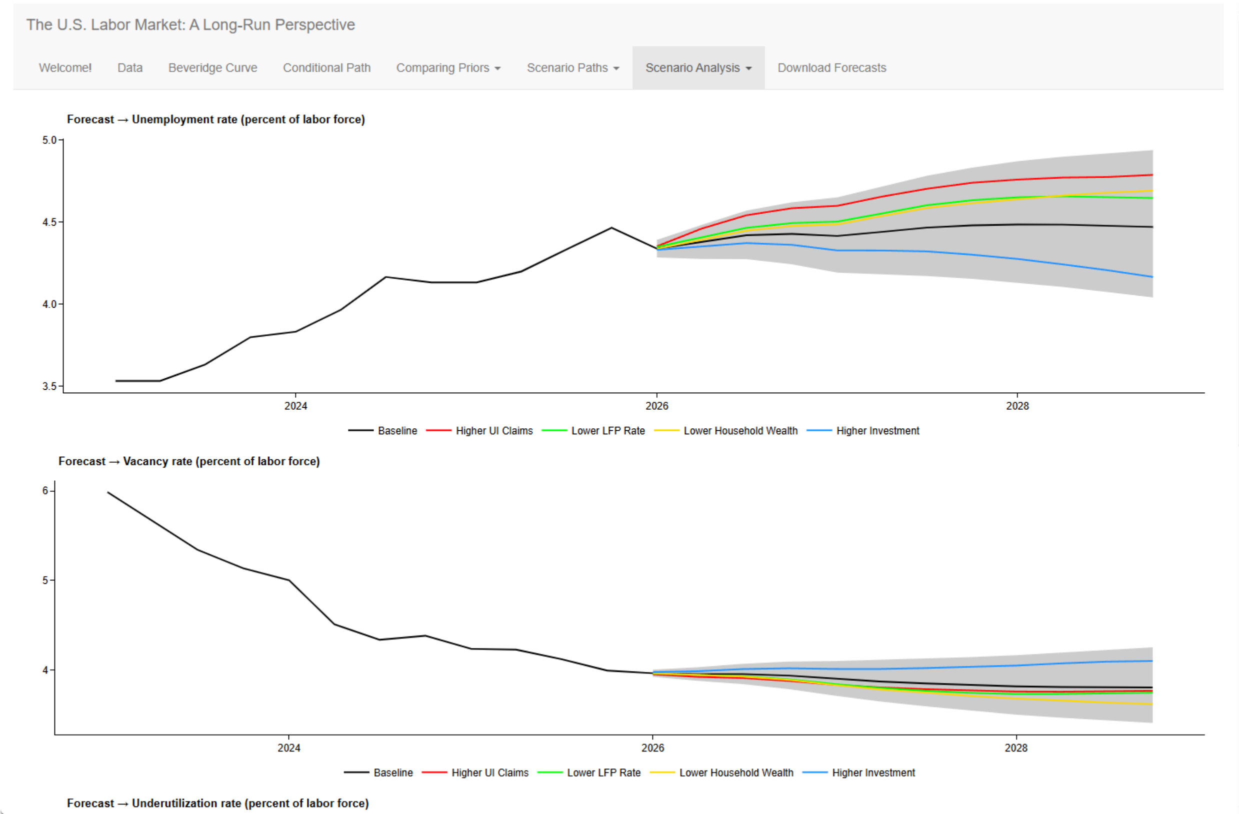

Download Forecasts

Each Download button downloads the 15th, 50th, and 85th percentiles for model variables of interest based on the conditional forecast paths provided in the Conditional Path and Scenario Path tabs.

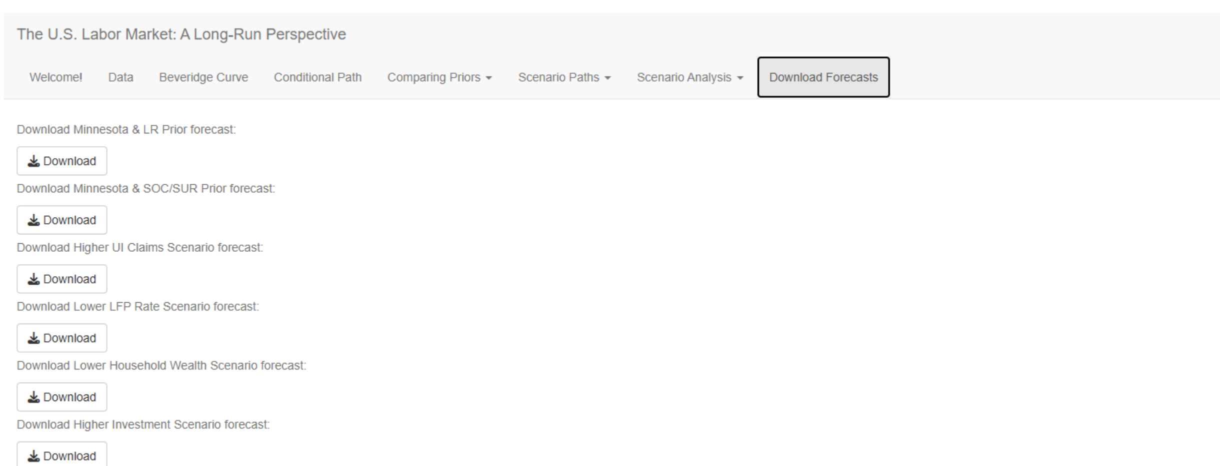

Changing other parts of the model

The app loads a .qs2 file to produce the plots in this application. If the model or its output is modified significantly, the application may no longer be compatible with this article.

Disclaimer

This R Shiny application is covered by the Legal Notices applicable to all information, data, and materials available on the website of the Federal Reserve Bank of Chicago, found here: https://www.chicagofed.org/utilities/legal-notices.