A Large-Scale Bayesian VAR for Disaggregated GDP Components *

Taeyoung Doh† Edward P. Herbst ‡ Alexander Gallin§

Abstract

We introduce a large scale Bayesian VAR for monthly indicators of disaggregated GDP components and estimate it using Bayesian methods. By linking monthly indicators with relevant GDP subcomponents and aggregating them following the national accounting convention, we generate GDP forecasts as well as a few key macroeconomic variables. By accommodating the heterogeneous release schedule of monthly indicators, our model can sequentially updates GDP forecasts at higher (daily or weekly) frequency than feasible based on only aggregate variables while generating nowcasts comparable to the widely used reference GDP tracking model. Our near-term forecasts for aggregate variables are comparable to a medium-scale mixed frequency VAR model but our model performs better at the medium-term horizon.

JEL classification: C32;E01;E27; E37.

Keywords: Bayesian VAR, Disaggregated GDP Components, Nowcasting. Forecasting

1 Introduction

To assess macroeconomic conditions in real time, many GDP nowcasting models combine information from a variety of monthly indicators with GDP components. For instance, GDPNow which is maintained by the Federal Reserve Bank of Atlanta (Higgins (2014)) tracks current quarter GDP

*The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors and the Federal Reserve Bank of Kansas City of the Federal Reserve System or of anyone else associated with the Federal Reserve System. This project started when Taeyoung Doh and Edward P. Herbst were at the Current Macroeconomic Conditions (CMC)section at the Federal Reserve Board and we appreciate help from the CMC staff–especially Jaclyn Lee, Niko Paulson, and Charlene Ramos–during the initial stage of the project. The author(s) thank the Center for Advancement of Data and Research in Economics (CADRE) at the Federal Reserve Bank of Kansas City (FRB-KC) for providing essential "computational/data" resources that contributed to the results reported in this paper. This version: October 2024

†Federal Reserve Bank of Kansas City: Taeyoung.Doh@kc.frb.org

‡Federal Reserve Board: edward.p.herbst@frb.gov

§Federal Reserve Bank of Kansas City:Alex.Gallin@kc.frb.org

by aggregating component estimates such as consumption and residential investment based on "bridge equations" linking quarterly components with monthly indicators released before the official first estimate of GDP by the Bureau of Economic Analysis (BEA). As noted by Klein and Park (1993), the aggregate national account identity explicit in this type of GDP tracking model is an informative economic restriction. By allowing the decomposition of changes in GDP estimates into changes in subcomponent estimates based on surprises in high-frequency indicators, the tracking model facilitates the economic interpretation of data surprises. This is a unique advantage of the tracking model compared to an alternative GDP nowcasting model based on a factor model (e.g., GDP nowcasting model maintained by the Federal Reserve Bank of New York in Baker, Almuzara, O'Keeffe and Sbordone (2023)) in which surprises in individual high-frequency indicators affect GDP estimates only through their loadings on the statistical common component of a variety of indicators typically without any notion of accounting identity restriction. [1] Since factor weights of monthly indicators do not necessarily match expenditure weights of these indicators in GDP, the factor model tends to overweight news of the volatile component such as indicators for durable goods consumption compared to indicators of service consumption although the latter can contribute more to headline GDP given that the expenditure share is greater for services.

However, the current practice of the tracking model where each subcomponent is estimated largely independently of other components is econometrically inefficient since it misses cross-correlations among indicators. GDPNow alleviates this issue by adding the principal component of many indicators into the forecasting equation of an individual indicator. Nonetheless, concerns about the econometric inefficiency still remain because factor weights do not change by the targeted monthly indicator to be forecasted. A fully-specified joint forecasting model of all the high-frequency indicators that respects the accounting identity restriction for GDP components would be desirable. In this paper, we propose a large-scale Bayesian Vector Autoregression (BVAR) model that can achieve this goal.

We include more than a hundred monthly indicators that are directly related to GDP subcomponents as well as five aggregate variables such as the unemployment rate, stock market index, interest rate, core PCE inflation, and the oil future price. While consumption components are available at the monthly frequency, the other twenty-six GDP subcomponents are available only at the quarterly frequency. We approximate the methodology the BEA uses to estimate these quarterly subcomponents based on a selected number of monthly indicators as described in the Handbook of National Income and Product Accounting by running auxiliary regressions of quarterly subcomponents onto monthly indicators aggregated at the quarterly frequency. These auxiliary regressions are used to

[1] Giannone, Reichlin and Small (2008) popularized GDP nowcasting based on an estimated factor model and Antolín-Díaz (2024) describes a state-of-the art nowcasting model of this class that incorporates complex time series patterns such as fat tails. For a more comprehensive review of GDP nowcasting model based on factor models, see Cascaldi-Garcia, Luciani and Modugno (2023).

obtain the likelihood of quarterly subcomponents conditional on available monthly indicators and do not restrict the time series dynamics of monthly indicators. We generate forecasts of aggregated GDP components by calculating the chain-weighted Fisher index as used in the BEA to avoid a distortion from neglecting relative price movements (Whelan (2002)).

Our estimates impose a tight prior on VAR parameters since the number of parameters is quite large given our sample size.[2] We adopt a commonly used Minnesota-type prior but derive the posterior draws of hyperparameters following the idea in Giannone, Lenza and Primiceri (2015) instead of trying to generate posterior draws of VAR conditional on chosen hyperparameters. Our large-scale poses a couple of computational challenges in implementing the hierarchical prior method like Giannone et al. (2015). We use the sequential Monte Carlo method explained in Herbst and Schorfheide (2016) to parallelize the computation for the posterior distribution of hyperparameters. In addition, we obtain posterior draws of VAR coefficients conditional on a draw of hyperparameters using a matrix normal distribution instead of a vectorized multivariate normal distribution, which reduces the computational burden of inverting a large covariance matrix for posterior simulation.

Due to the limitation of obtaining construction-related price indicators, our sample starts from June 2006. We estimate our model using the several subsamples of pre-COVID19 period data and perform pseudo out-of-sample forecasts. Our forecasts for the unemployment rate are comparable to those obtained from a mixed-frequency VAR model of 11 variables in (Schorfheide and Song, 2015, 2021) in the near-term horizon (up to 6 months) but better in the medium term (from 7 months to 24 months). In particular, the MF-VAR model forecasts made at the December of 2017 expect the future weakening in demand with the unemployment rate and inflation moving in the opposite direction. In contrast, our BVAR model captures a decline in the the unemployment rate in the data while predicting low inflation.

Leveraging our model's ability to sequentially update forecasts at higher frequency given the heterogeneous release schedule of monthly indicators, we obtain the nowcasting density of 2024:Q2 GDP growth from the last week of May to the week when the BEA released the first estimate of 2024:Q2 GDP. The median forecast of our model of 2.9 percent is close to the actual first estimate of 2.8 percent. For reference, the Atlanta Fed's GDPNow's comparable tracking estimate was 2.6 percent. The revision in the tracking estimate during the two month period is highly correlated with the revision in the GDPNow estimate with the correlation coefficient of 0.58, indicating that two models tend to interpret information from released indicators in a roughly similar way. However, revisions in our tracking estimate are uncorrelated with those in the nowcasting estimate from a factor model described in Baker et al. (2023). The finding highlights the difference between the

[2]While the prior distribution shrinks the posterior distribution of parameters, it does not rule out the role of many predictors as discussed in Giannone, Lenza and Primiceri (2021).

trakcing model respecting the accounting identify restriction and models which do not respect it.

Related Literature: Our paper is related to the burgeoning literature of incorporating various monthly indicators into macroeconomic forecasting. Giannone et al. (2008) proposed estimating the VAR dynamics of factors extracted from a large number of monthly indicators through the principal component analysis.[3] The estimated VAR model combined with the factor loadings of underlying monthly indicators allows us to update factor estimates in response to surprises in data releases. GDP nowcasts are in turn updated based on the auxiliary regression of linking the current quarter GDP with the estimated factor in the current quarter. The subsequent literature refined this approach by allowing heavy tails in the underlying data (Antolín-Díaz (2024)), or improving the econometric efficiency (see the discussion in Cascaldi-Garcia et al. (2023)), but the basic idea of using statistical factors to connect monthly indicators with GDP forecasts largely remains intact. Our approach is complementary to these factor-based methods. In fact, a large VAR model that incorporates all the underlying data in the factor model as well as GDP estimates can encompass the factor model structure because the factor model is a restricted version of such a large scale VAR model. Since we do not include all the underlying data especially a diffusion index or survey data often used in a factor model, our model does not exactly nest widely used factor models for GDP nowcasting. Nonetheless, our use of the accounting identity restriction on GDP components can provide a more plausible economic interpretation of the aggregate effect of data surprises, which is not feasible in the factor model that ignores the expenditure weight of GDP components.

Our paper is also related to the literature of using a large-scale BVAR for macroeconomic forecasting starting from Bańbura, Giannone and Reichlin (2010). Although BVAR models have been used in the medium-term forecasting for long time, the computational cost of scaling up the model to the extent that is comparable to our model size of more than hundred variables has remained challenging. For example, Crump, Eusepi, Giannone, Qian and Sbordone (2021) estimate a large-scale BVAR model of the U.S. economy using 31 variables but do not include many disaggregated GDP component indicators. Angelini, Lalik, Lenza and Paredes (2019) estimate a similar model for the Euro area using 37 variables. A more elegant mixed-frequency approach that estimates unobserved monthly counterparts of quarterly GDP components using the entire set of VAR observables is even more restricted in scale, typically using no more than 20 variables (e.g., Chan, Poon and Zhu (2023), Schorfheide and Song (2015), and Schorfheide and Song (2021)) because they need additional Kalman smoothing steps to back out estimates of unobserved variables.[4] Also, like

[3] Fitting the VAR for extracted factors is not essential because one can fit a dynamic factor model that estimates factors and parameters governing factor dynamics simultaneously as in a real-time measure of business conditions illustrated by Aruoba, Diebold and Scotti (2009).

[4] Brave, Butters and Justiniano (2019) estimate a larger mixed-frequency VAR model consisting of 30 monthly variables and 7 quarterly variables. Still, our model has almost three times bigger than their model.

factor models, the accounting identity restriction that aggregates GDP components to the headline number is typically not considered as a restriction in this literature. Our large-scale BVAR model for disaggregated GDP components contributes to the literature by illustrating that the real-time estimation of a much larger scale model than currently in practice is feasible and generating forecasts of aggregate variables that respect the accounting identity restriction. In particular, the accounting identity restriction provides a more plausible economic interpretation on the aggregate effect of surprises in monthly indicators for disaggregated components.

2 Econometric Methodology

Our goal is to build a time series framework for tracking and forecasting the U.S. economy at a disaggregated level.[5] Our framework is comprehensive in the sense that it tracks both disaggregated quantities (given by y) and prices (given by p). We focus on disaggregating GDP into $I = 38$ subcomponents, consisting of 12 consumption and 26 non-consumption subcomponents. This level of disaggregation matches the corresponding number in a GDP tracking model maintained by the Federal Reserve Bank of Kansas City (Bae and Doh (2019)) and is comparable to the level considered in GDPNow (Higgins (2014)). Because the non-consumption subcomponents are observed only at a quarterly frequency, we use superscript q, that is y_t^q and p_t^q, to denote the (potentially unobserved) quantities and prices of the subcomponents in month t. It is often convenient to refer to the price-quantity pair as $z_t^q = [y_t^{q'}, p_t^{q'}]'$. Let GDP_t^q and t be the aggregated real quarterly GDP and the final month of that quarter. We calculate the aggregated GDP from prices and real values of subcomponents by the Fisher weighted chain index following the BEA's convention as follows:[6]

$$\text{GDP}_t^q = \text{GDP}_{t-3}^q \sqrt{\frac{\sum_i p_{i,t}^q y_{i,t}^q}{\sum_i p_{i,t}^q y_{i,t-3}^q} \frac{\sum_i p_{i,t-3}^q y_{i,t}^q}{\sum_i p_{i,t-3}^q y_{i,t-3}^q}}. \tag{2.1}$$

Thus, in this disaggregated framework tracking (or multi-step forecasting or analyzing or calculating uncertainty around) GDP requires a forecasting model for the subcomponent quantities and prices. One reason this is difficult is data availability. Other than monthly details of consumption expenditures, GDP subcomponents as well as headline GDP itself are not available at a monthly frequency. The BEA interpolates many non-consumption components at the monthly frequency

[5] Our current work is focused on the reduced-form VAR but we can evaluate policy counterfactuals by adding identifying restrictions of structural shocks. We are working on this.

[6] One exception in this aggregation is changes in inventories, which can be negative. However, when included in gross private investment and headline GDP, inventory investment accounts for only a small portion of the aggregated variable and the BEA includes them in the Fisher weighted chain index. We follows the BEA's convention in incorporating inventories into aggregated forecasts.

using relevant *proxy indicators* for prices and quantities. While these proxy indicators are described in the handbook of national income and product accounting, not every proxy indicator is publicly available data, limiting the level of disaggregation for researchers outside the BEA.

We focus on a set of $n = 105$ monthly indicators that are directly related to GDP subcomponents (inspired by the NIPA handbook) as well as a set of informative aggregate variables such as the unemployment rate, stock market index, interest rate, core PCE inflation, and oil future price. We denote the vector of these monthly indicators at time t by z_t^m. Using $Z_T^m = [z_1^{m\prime}, \ldots, z_T^{m\prime}]$ and similarly for Z_T^q, we conceptualize our time series framework as a joint likelihood for Z_T^m and Z_T^q given model parameters θ, which can be represented by $p(Z_T^m, Z_T^q|\theta)$. This likelihood can be factorized into:

$$p(Z_T^m, Z_T^q|\theta) = p(Z_T^q|Z_T^m, \theta)p(Z_T^m|\theta). \tag{2.2}$$

Thus, the likelihood can be constructed in two steps. A natural likelihood for the subcomponent data given the monthly indicators would be a linear state space model predictive likelihood where the missing subcomponent data is integrated out, *a la* the MF-VAR of Schorfheide and Song (2015). But such an approach can be computationally challenging given the large number of subcomponents. Instead, we interpolate the subcomponent using observed quarterly values and relevant monthly indicators. Such an approach can be grounded in a Bayesian (i.e., likelihood-based) framework, although our current implementation does not include a full Bayesian analysis of the subcomponent data. To construct the second part of the likelihood, we use a large VAR model, which we describe in the next subsection.

2.1 A monthly VAR for Z_t^m

As mentioned above, the vector of monthly indicators z_t^m contains $n = 105$ elements.[7] We assume that they follow a VAR with $p = 12$ lags.

$$z_t^m = \Phi_0 + \Phi_1 z_{t-1}^m + \cdots + \Phi_{12} z_{t-12}^m + u_t, \qquad u_t \sim \mathcal{N}(0, \Sigma). \tag{2.3}$$

Although the likelihood construction is conceptually straightforward, estimation is practically infeasible without shrinkage techniques. This model has 137,970 unique parameters, which is an order of magnitude larger than the number of observations available. To facilitate the use of this model, we combine the likelihood with a prior distribution for the parameters of VAR $\Phi = [\Phi_0, \cdots, \Phi_{12}]'$ and Σ. Following a long tradition in macroeconomics, we use the Normal-

[7] Most indicators are in terms of log-levels except for 8 variables representing changes in inventories, industrial production index, aggregate variables such as PCE inflation, unemployment rate, oil price, and financial market variables.

Inverse-Wishart formulation of the "Minnesota" Prior

$$\Sigma \sim IW(\Psi, \nu) \text{ and } \Phi|\Sigma \sim \mathcal{N}(\bar{\Phi}, \Sigma \otimes \Omega). \tag{2.4}$$

The Minnesota prior is commonly used in Bayesian VAR models to impose structure by assuming each variable is primarily driven by its own past values, while allowing for modest influence from other variables, helping to prevent overfitting and improve forecast accuracy–an overview can be found in Del Negro and Schorfheide (2011). The prior is constructed using a set of hyperparameters to generate dummy observations. Following Giannone et al. (2015), we estimate a subset of these hyperparameters $[\lambda, \delta, \mu]$ and $\psi = \text{diag}(\Psi)$.[8] The parameter λ controls the "effective" overall tightness, with smaller values of λ indicating a higher degree of shrinkage. The parameter μ is known as the co-persistence prior, lower values of which more strongly reflect the belief that when the monthly indicators are all stable at their initial value, they will tend to persistent at that level. Finally, δ governs the "initial observation" prior, lower values of which more strongly reflect the belief that when any given observable is stable around its initial level, it will tend to persistent at that level (regardless of the value of the other variable). These parameters induce correlation in the posterior. We follow Giannone et al. (2015) and estimate these hyperparameters, using the following priors:

$$p(\lambda) = \text{Gamma}(1.64, 0.313), \qquad p(\mu) = \text{Gamma}(2.62, 0.618) \tag{2.5}$$

$$p(\delta) = \text{Gamma}(2.62, 0.618), \qquad \text{and } p(\psi) = \prod_{1}^{n} \text{Uniform}(0, 2). \tag{2.6}$$

We can compute the likelihood (ignoring the initial p observations) of the monthly indicators under these hyperparameters via:

$$p(Z_T^m|\lambda, \mu, \delta, \psi) = \int p(Z_T^m|\Phi, \Sigma) p(\Phi, \Sigma|\lambda, \mu, \delta, \psi) d(\Phi, \Sigma). \tag{2.7}$$

This integral is standard in Bayesian inference and can be solved analytically. Using the prior just described yields a Bayesian model when combined with a likelihood function for the monthly indicators. The likelihood is a complicated function of the hyperparameters, and so we must resort to simulation techniques to elicit the posterior. This is challenging both because the model is large: the vector ψ has 105 parameters, giving us 108 total parameters to estimate. Moreover, the likelihood

[8]The full specification of the Minnesota Prior requires some additional information. We fix these additional parameters: the prior location ("$\bar{y}$") is set using the mean of the first 12 observations $\{z_t^m\}_{t=1}^{12}$. The factor which further scales coefficients of higher lags is set to 1, and the prior on weight on Ψ is set to 1. In the notation of Del Negro and Schorfheide (2011), the hyperparameter vector is given by $[1/\sqrt{\lambda}, 1, 1, 1/\mu, 1/\delta]$. This choice implies the inverse Wishart distribution has $\nu = 107$ degrees of freedom. Note that Ψ is assumed to be a diagonal matrix.

calculation is cumbersome; it requires the inversion of extremely large matrices. For these two reasons the MCMC procedure suggested by Giannone et al. (2015) is impractical. Instead, we use the Sequential Monte Carlo (SMC) method outlined in Herbst and Schorfheide (2016)-and applied to VARs by Bognanni and Herbst (2018)–to estimate the posterior distribution of the hyperparameters. The SMC method employed here involves adjusting an initial set of random draws from the prior distribution using increasingly more information from the likelihood function until the draws reflect the posterior distribution. An important feature of the SMC method relative to MCMC is that it can be easily parallelizable. This allows us to estimate the posterior distribution of the hyperparameters in a reasonable amount of time. Once we have posterior draws of the hyperparameters, we can directly sample posterior draws of Φ and Σ conditional on the hyperparameters.[9]

2.2 Auxiliary regression for z_t^q

The BEA does not release monthly values of non-consumption GDP subcomponents although some of monthly indicators they use to estimate the quarterly values are publicly available. Based on the description in the NIPA handbook, we match a particular non-consumption subcomponent of real GDP ($y_{i,t}^q$) and the price counterpart of it ($p_{i,t}^q$) with a subset of monthly indicators, $z_{y(i),t}^m$ and $z_{p(i),t}^m$. $y(i)$ and $p(i)$ denote the indices of monthly indicators z_t^m that are expected to be relevant for $z_{i,t}^q$ based on the NIPA handbook. We introduce the following auxiliary regression that links the logged levels of $y_{i,t}^q$ and $p_{i,t}^q$ with $z_{y(i),t}^m$ and $z_{p(i),t}^m$, respectively.[10]

$$\begin{aligned} \ln y_{i,t}^q &= \alpha_{y,i} + z_{y(i),t}^m \beta_{y,i} + v_{y_i,t}\,,\; v_{y_i,t} \sim \mathcal{N}(0, \sigma_{y_i}^2), \\ \ln p_{i,t}^q &= \alpha_{p,i} + z_{p(i),t}^m \beta_{p,i} + v_{p_i,t}\,,\; v_{p_i,t} \sim \mathcal{N}(0, \sigma_{p_i}^2). \end{aligned} \tag{2.9}$$

Since we do not observe z_t^q for months outside the end-of-quarter month, we cannot directly use the above equation to estimate parameters. What we observe the following quarterly aggregation of z_t^q, which we call by $x_t^q = (z_t^q + z_{t-1}^q + z_{t-2}^q)$. Suppose that $T_q = [\frac{T}{3}]$ quarters are spanned by T

[9]Sampling $vec(\Phi)$ from a multi-normal distribution is computationally burdensome because we have to obtain the Cholesky decomposition of a very large matrix, $\Sigma \otimes (X^{m\prime}X^m)^{-1}$, whose size is $n(np+1) \times n(np+1)$. In our application, this matrix has a size of 132405 $\times$ 132405. However, we can directly sample Φ instead of $vec(\Phi)$ using the Kronecker-delta structure of the covariance matrix. Bauwens, Lubrano and Richard (2000) provide a formal treatment on a matrix normal random distribution that justifies our sampling method. The parameter is written as:

$$\Phi \sim \hat{\Phi} + C_\Sigma V C_X'. \tag{2.8}$$

where C_Σ is the lower Cholesky decomposition of Σ and C_X is the lower Cholesky decomposition of $(X^{m\prime}X^m)^{-1}$. V is the matrix of $n \times (np+1)$ that consists of standard normal random variables. Since the size of the Cholesky decomposition of Σ and $(X^{m\prime}X^m)$ is much lower than the size of $\Sigma \otimes (X^{m\prime}X^m)^{-1}$, drawing Φ is computationally efficient compared to drawing $vec(\Phi)$.

[10]For inventories, we run the regression of $y_{i,t}^q$.

months in our sample period. There exists a $T \times T_q$ matrix B which converts monthly variables z_t^q into $x_{q,t}$ by

$$x^q = B'z^q = B'(\alpha^q + Qz^m\beta^q + v) = B'\alpha^q + B'Qz^m\beta^q + Bv. \tag{2.10}$$

α^q is the row vector of stacking $\alpha_{y,i}$ and $\alpha_{p,i}$ while Q is the selection matrix matching a particular quarterly subcomponent with a subset of monthly indicators. Similarly, β^q is the matrix of stacking coefficients $\beta_{y,i}$ and $\beta_{p,i}$ and v also stacks error terms in auxiliary regressions. We use the interpolation method explained in Fernandez (1981) to estimate α^q and β^q given z^m and x^q as the solution to the following minimization problem:[11]

$$\begin{aligned} (\hat{z^q}, \hat{\alpha^q}, \hat{\beta^q}) = \text{argmin} \quad & trace(z^q - \alpha^q - Qz^m\beta^q)(z^q - \alpha^q - Qz^m\beta^q)', \\ \text{subject to} \quad & x^q = B'z^q. \end{aligned} \tag{2.11}$$

We represent α^q and β^q as well as variance parameters for v by θ_q. Notice that although only x^q, which is the quarterly sum of z^q is observed, we need to write down the complete likelihood for z^q to make predictions for x^q from monthly indicators.[12]

In principle, we can perform a fully Bayesian analysis for θ by modeling the prior distribution of θ_q. For now, we assume the diffuse prior for α^q and β^q and take the restricted Ordinary Least Squares(OLS) estimator for these parameters and ignore parameter uncertainty at this step.[13]

2.3 Accommodating heterogeneous data release schedule

Monthly indicators are not released at the same time, meaning we may have partial observations for z_t^m. We can partition z_t^m into the observed part $z_t^{m,o}$ and the unobserved part $z_t^{m,n}$. Let's denote the information set spanned by $z_1^m, \cdots, z_{t-1}^m$ by I_{t-1}. The release of $z_t^{m,o}$ expands the information set to $I_{t-1} \cup z_t^{m,o}$. We can update the forecast of $z_t^{m,n}$ conditional on $z_t^{m,o}$ using Kalman filter as follows:[14]

$$E(z_t^{m,n}|I_{t-1} \cup z_t^{m,o}) = E(z_t^{m,n}|I_{t-1}) + \Sigma_{n,o}\Sigma_{o,o}^{-1}(z_t^{m,o} - E(z_t^{m,o}|I_{t-1})), \tag{2.12}$$

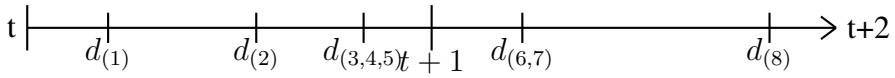
[11] A similar method is used in Chen, Higgins and Zha (2024) to obtain the quarterly Chinese GDP number from the official annual estimate and quarterly indicators.

[12] In practice, we run the auxiliary regression of the growth rate of the quarterly component onto the growth rates of monthly indicators. In this case, we only need to change B matrix to approximate the quarterly growth rate as a weighted sum of monthly growth rates like $\Delta x_t^q = \frac{\Delta z_t^q + 2\Delta z_{t-1}^q + 3\Delta z_{t-2}^q + 2\Delta z_{t-3}^q + \Delta z_{t-4}^q}{3}$.

[13] A fully Bayesian analysis including θ_q is in our future agenda.

[14] When the difference in the release schedule is one month, the formula below is enough to construct the nowcasting data. However, there are roughly one or two weeks early in each month, the maximum gap in the release schedule among indicators is two months. In that case, we sequentially update the forecast of $z_t^{m,n}$ by $E(z_t^{m,n}|I_{t-2} \cup z_{t-1}^{m,o} \cup z_t^{m,o})$.

Figure 1: MONTHLY DATA RELEASE SCHEDULE



where $\Sigma_{n,o}$ is the covariance between $z_t^{m,n}$ and $z_t^{m,o}$ in Σ while $\Sigma_{o,o}$ denotes the block of Σ with respect to $z_t^{m,o}$.

Figure 1 illustrates the heterogeneous data release schedule. $d_{(i)}$ represents a data release event. Each release may provide multiple monthly indicators for use in the BVAR. $d_{(1)}$ corresponds to the release of the employment situation report. $d_{(2)}$ denotes industrial production, retail sales, CPI data release. $d_{(3)}$ represents Treasury outlays which we deflate by PCE index to use as a real spending indicator. $d_{(4)}$ corresponds to PCE details as well as NIPA underlying details, if available. In this week, we receive either the first estimate of the quarter ending in the previous month or the revised estimate of it (second or third estimate) $d_{(6)}$ denotes the release of construction spending which we deflate by the relevant PPI to use as a real spending indicator. $d_{(7)}$ is the international trade data we deflate by import/export Price Indexes. $d_{(8)}$ denotes another monthly update of NIPA underlying details.

After the BEA releases the first estimate of the quarterly GDP, it releases the second and the third estimates for the subsequent two months. Hence, in the final week of each month, we always have one updated quarterly data. We incorporate the most recently released quarterly GDP numbers in our weekly nowcasting and forecasting of GDP components.

2.4 Predictive distribution of GDP growth

From equation (2.1), we know that log GDP growth can be represented by the log of the Fisher weighted sum of disaggregated components.

$$\ln \frac{GDP_t^q}{GDP_{t-3}^q} = \frac{1}{2}(\ln(\sum_i p_{i,t}^q y_{i,t}^q) - \ln(\sum_i p_{i,t}^q y_{i,t-3}^q) + \ln(\sum_i p_{i,t-3}^q y_{i,t}^q) - \ln(\sum_i p_{i,t-3}^q y_{i,t-3}^q)). \tag{2.13}$$

Excluding changes in inventories, $p_{i,t}^q$ and $y_{i,t}^q$ all follow log-normal distributions. Since the product of two log-normal random variables follows a log-normal distribution, equation (2.13) shows that the log GDP growth is the log sum of log normal random variables. However, the sum of log normal random variables does not follow a log normal distribution, which implies that the log GDP growth does not follow a normal distribution even if all the disaggregated components follow log normal distributions. Therefore, the predictive distribution of GDP growth in our model can deviate significantly from a normal distribution although underlying shocks to logged disaggregated

components are normally distributed even without considering parameter uncertainty.[15] Moreover, the degree of the deviation from the normality can be time-varying because the aggregation weights that influence the degree of non-normality are also time-varying. This property is reminiscent of time-varying "growth-at-risk" highlighted by Adrian, Boyarchenko and Giannone (2019) although we allow alternative channels beyond time-varying predictive power of financial conditions on future GDP growth in their paper.

We can decompose GDP growth forecast changes across different vintages into various component contributions using the first-order Taylor approximation to equation (2.13) as follows.

$$\begin{aligned}
\ln \frac{GDP^q_{t,d'}}{GDP^q_{t-3,d'}} - \ln \frac{GDP^q_{t,d}}{GDP^q_{t-3,d}} \approx & \frac{1}{2}\sum_i [\frac{p^q_{i,t,d}y^q_{i,t,d}}{\sum_j p^q_{j,t,d}y^q_{j,t,d}} - \frac{p^q_{i,t,d}y^q_{i,t-3,d}}{\sum_j p_{j,t,d}y^q_{j,t-3,d}}](\ln \frac{p_{i,t,d'}}{p_{i,t,d}}) \\
& + \frac{1}{2}\sum_i [\frac{p^q_{i,t,-3,d}y^q_{i,t,d}}{\sum_j p^q_{j,t-3,d}y^q_{j,t,d}} - \frac{p^q_{i,t-3,d}y^q_{i,t-3,d}}{\sum_j p_{j,t-3,d}y^q_{j,t-3,d}}](\ln \frac{p_{i,t-3,d'}}{p_{i,t-3,d}}) \\
& + \frac{1}{2}\sum_i [\frac{p^q_{i,t,d}y^q_{i,t,d}}{\sum_j p^q_{j,t,d}y^q_{j,t,d}} + \frac{p^q_{i,t-3,d}y^q_{i,t,d}}{\sum_j p_{j,t-3,d}y^q_{j,t,d}}](\ln \frac{y_{i,t,d'}}{y_{i,t,d}}) \\
& - \frac{1}{2}\sum_i [\frac{p^q_{i,t,d}y^q_{i,t-3,d}}{\sum_j p^q_{j,t,d}y^q_{j,t-3,d}} + \frac{p^q_{i,t-3,d}y^q_{i,t-3,d}}{\sum_j p_{j,t-3,d}y^q_{j,t-3,d}}](\ln \frac{y_{i,t-3,d'}}{y_{i,t-3,d}}). \quad (2.14)
\end{aligned}$$

$p^q_{i,t,d}(y^q_{I,t,d})$ and $p^q_{i,t,d'}(y^q_{i,t,d'})$ represent different vintages of same variables with d' corresponding to the later vintage. Since changes in inventories can be negative, this decomposition would be useful for components excluding inventories.

At a given vintage d, we can decompose GDP growth estimates into contribution from various components following the formula used in the BEA's NIPA handbook.

$$\begin{aligned}
\ln \frac{GDP^q_{t,d}}{GDP^q_{t-3,d}} &= \sum_i \frac{[(\frac{p^q_{i,t,d}}{p^F_{t,d}}) + p^q_{i,t-3,d}]}{\sum_j [(\frac{p^q_{j,t,d}}{p^F_{t,d}} + p^q_{j,t-3,d})]}(\ln \frac{y^q_{i,t,d}}{y^q_{i,t-3,d}}), \\
p^F_{t,d} &= \sqrt{\frac{\sum_i p^q_{i,t,d}y^q_{i,t-3,d}}{\sum_i p^q_{i,t-3,d}y^q_{i,t-3,d}} \frac{\sum_i p^q_{i,t,d}y^q_{i,t,d}}{\sum_i p^q_{i,t-3,d}y^q_{i,t,d}}}. \quad (2.15)
\end{aligned}$$

[15] For the mathematical characterization of the tail behavior of the sum of log normal random variables, see Gulisashvili and Tankov (2016). In the appendix, we compare the probability density function of the log sum of two log normal random variables with that of the sum of two normal random variables. The Bayesian predictive density in our model integrates out this non-normal distribution over the posterior distribution of parameters, which can deepen the departure from the normal distribution.

3 Data

We use 105 variables for the monthly BVAR and 26 quarterly real GDP subcomponents as well as 26 quarterly price deflators and 1 employment cost index series available at the quarterly frequency. The monthly BVAR variables consist of 41 disaggregated price indicators, 59 disaggregated real spending indicators, 4 aggregate variables for labor market and financial market conditions, and 1 aggregate core PCE inflation. Among disaggregated real spending indicators, 22 variables from the monthly PCE details and retail sales report are directly related to consumption. Among the rest, 6 variables are related to residential investment, 8 variables to business fixed investment, 5 variables to net exports, 11 variables to inventories, and 7 variables to government spending. All these real spending indicators are obtained by deflating nominal spending data by the relevant price indicator in line with the practice done in the tracking model maintained by the Federal Reserve Bank of Kansas City. Disaggregated price data consist of 14 PCE price details, 10 CPI details, 8 PPI details, 4 indicators from the Census Bureau export/import price index, 2 wage indicators, and 3 housing/construction price indicators.

Our sample data starts from June 2006. For some of disaggregated price data such as Producer Price Index (PPI) for new office building construction, we were not able to find out monthly price indicators that can mimic the variable and this was a hard stop for us.[16] We estimate both the monthly BVAR model and auxiliary regressions using only pre-COVID19 data, ending at December 2019.[17]

4 Empirical Results

4.1 Shrinkage of parameters in the monthly BVAR

While we assume fairly diffuse prior distributions for hyperparameters, their posterior distributions are concentrated around the posterior mode as shown in Figure 2. The finding suggests that our data provide rich information on the direction and magnitude of shrinking VAR parameters.[18]

Our Φ matrix is large with 132,405 elements. Hyperparameters shrink these parameters to make estimation feasible with the limited sample size. Figure 3 $\sim$ 5 show the heatmap of the absolute values of $(\Phi_1, \cdots, \Phi_{12})$ normalized by the posterior standard deviation of each parameter. For Φ_1,

[16] The appendix describes two disaggregated price data we extended by using a large number of other monthly price indicators as predictors.

[17] Incorporating the COVID19 period data in the estimation sample without properly adjusting volatility may generate poor forecasts. Given the scale of our model, incorporating this feature remains computationally challenging but we plan to address this in the ongoing work.

[18] The posterior distribution of Σ is also concentrated relative to the prior distribution although we do not show plots given the large number of hyperparameters. Plots for the diagonal elements of Σ are available upon request.

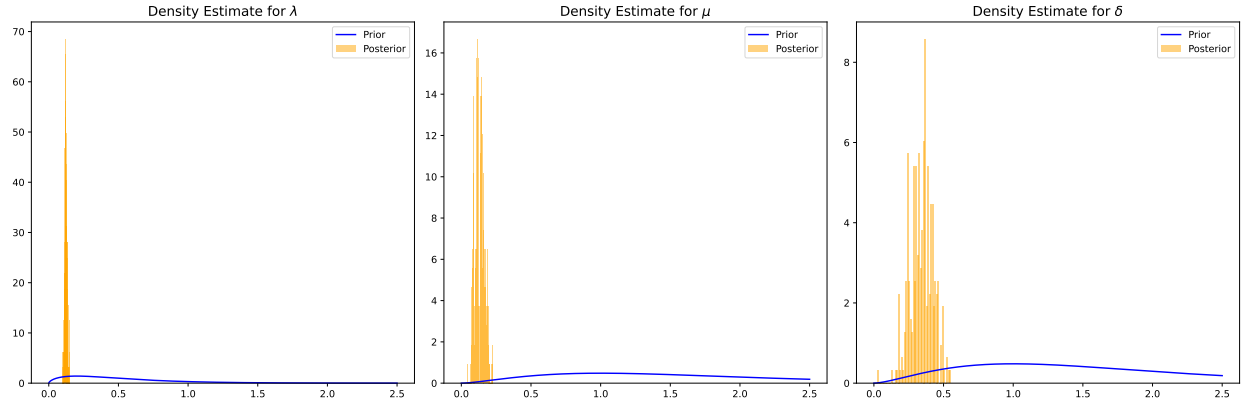


Figure 2: PRIOR AND POSTERIOR DISTRIBUTIONS OF HYPERPARAMETERS

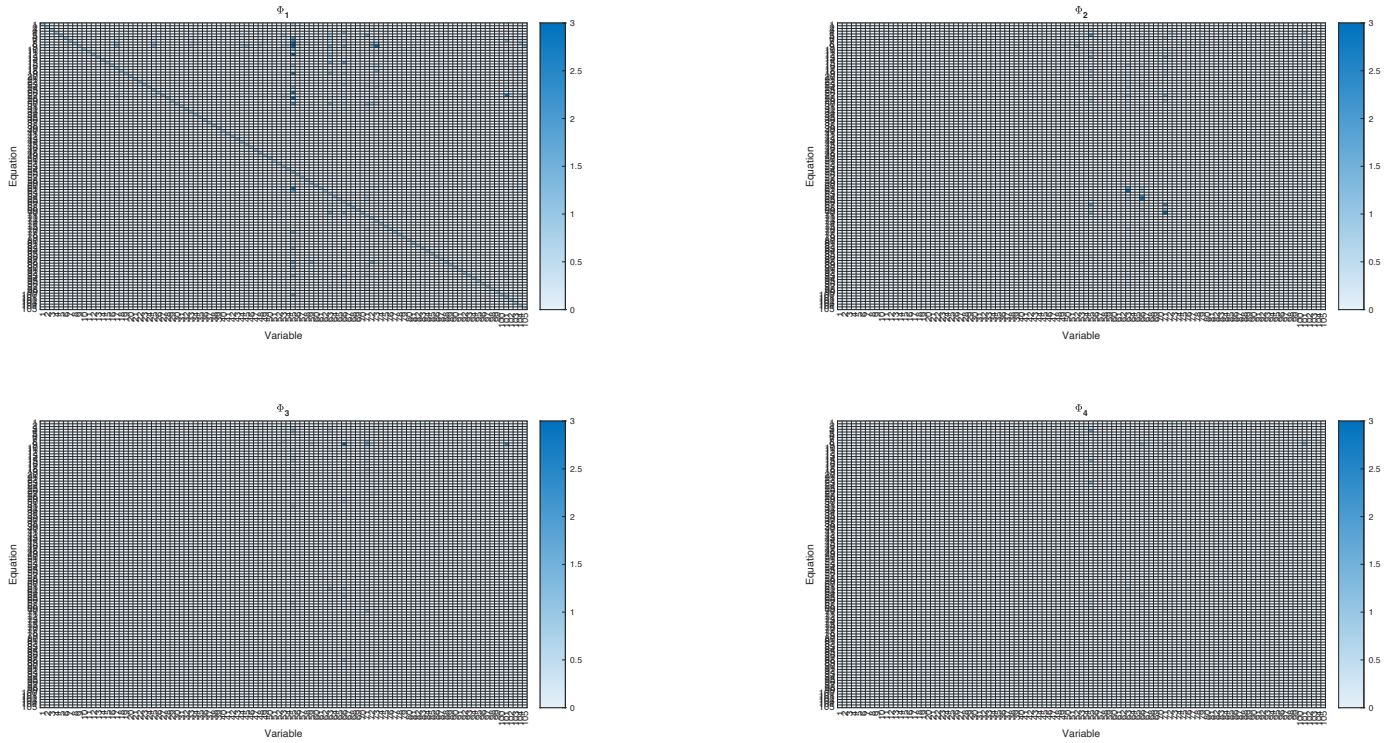


Figure 3: HEATMAP OF Φ_1 - Φ_4

diagonal terms are pretty significant but off-diagonal terms are largely insignificant. However, for coefficients of higher orders, there are pockets of off-diagonal terms that are significant. These patterns of non-negligible cross-variable linkages at higher orders suggest that a simple uni-variate forecasting model may not be able to capture an important feature of many indicators.

We also look the correlations of VAR residuals using the posterior mean of Σ in Figure 6. While most elements are close to 0, we see some correlations are different from 0, again implying the presence of non-negligible cross-equation spillovers.

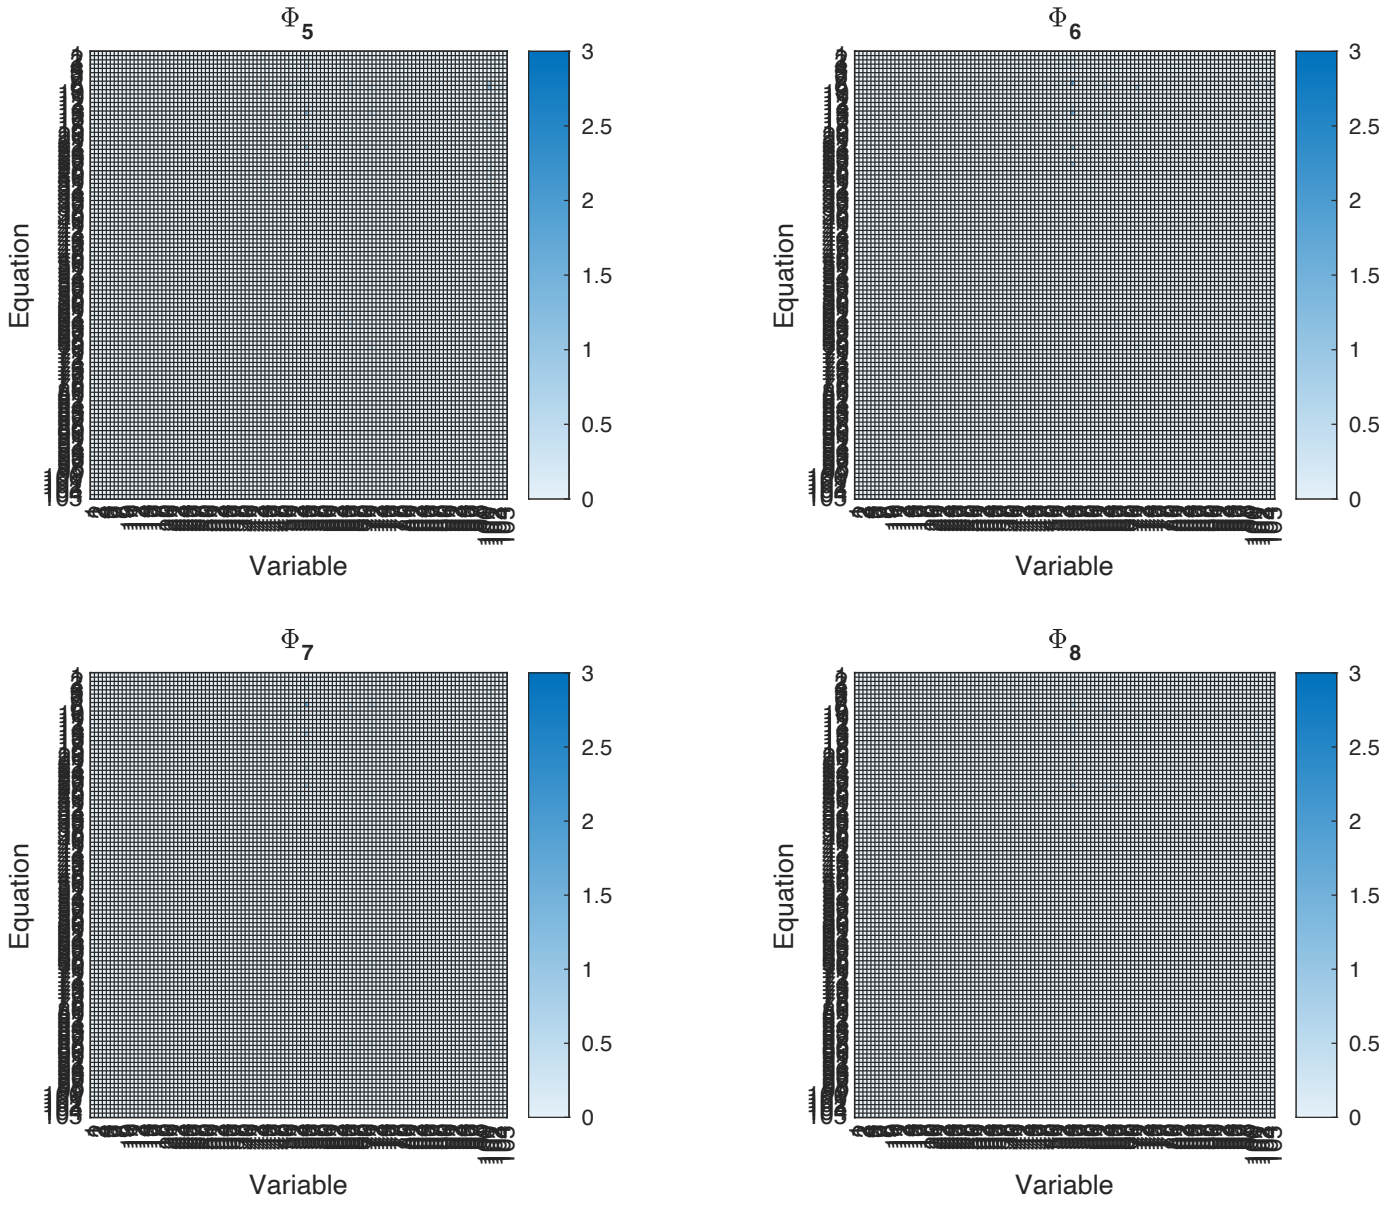


Figure 4: HEATMAP OF Φ_5 AND Φ_8

4.2 Unemployment Rate Forecast Comparison with the Mixed-Frequency VAR in Schorfheide and Song (2021)

We first assess the fit of our model for aggregate variables by comparing our model's pseudo out-of-sample forecasts with those from a benchmark model using only information from aggregate variables. We estimate both our disaggregated component model and the 11-variable mixed-frequency VAR (MF-VAR) model in Schorfheide and Song (2021) using the data from 2006 to 2017

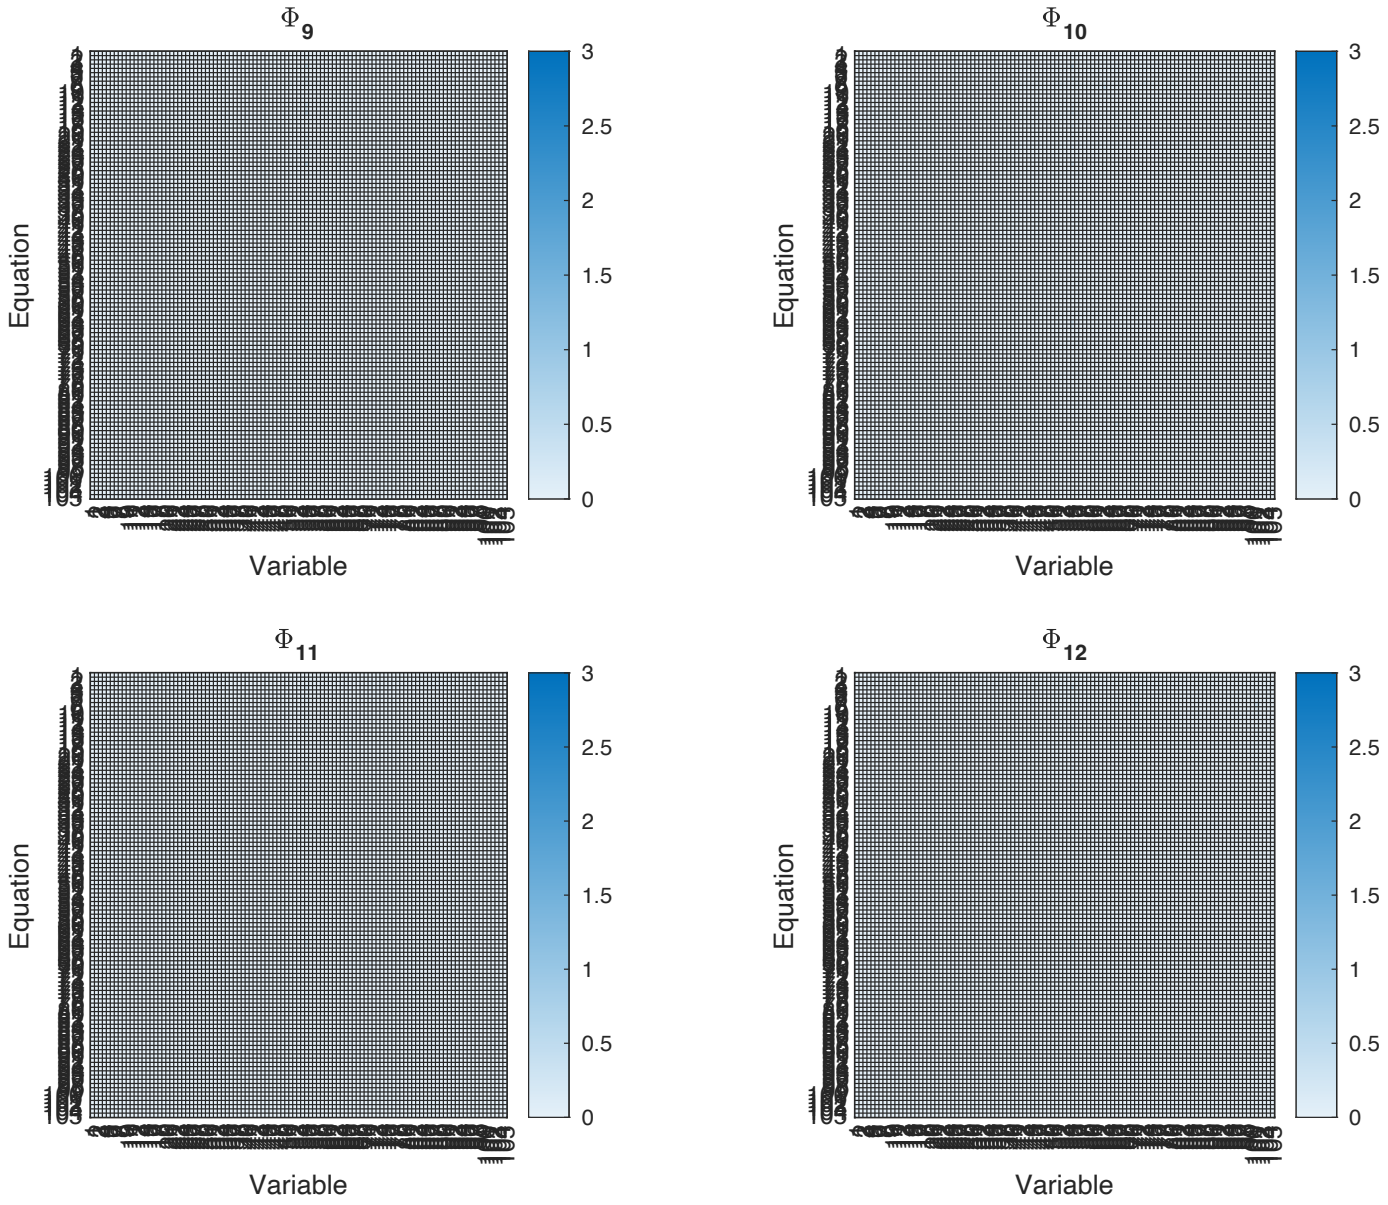


Figure 5: HEATMAP OF Φ_9 AND Φ_{12}

and compare two model's forecast for the unemployment rate.[19] We generate the 95% posterior interval as well as the median forecast from each model's predictive distribution. The predictive distribution incorporates both parameter uncertainty and future shock uncertainty.

Figure 7 shows the unemployment rate forecasts from the two models as well as the realized data. Up to 5 month horizon, model forecasts are nearly identical but forecasts start to diverge from

[19]We estimate the data from June 2006 to December 2017 to estimate our model. For the MF-VAR model, the estimation starts from July 2006 but both models use some of the early sample observations as the pre-sample data to set up prior distributions and this one month difference is unlikely to make any material difference in pseudo out-of-sample forecast for 2018 and 2019.

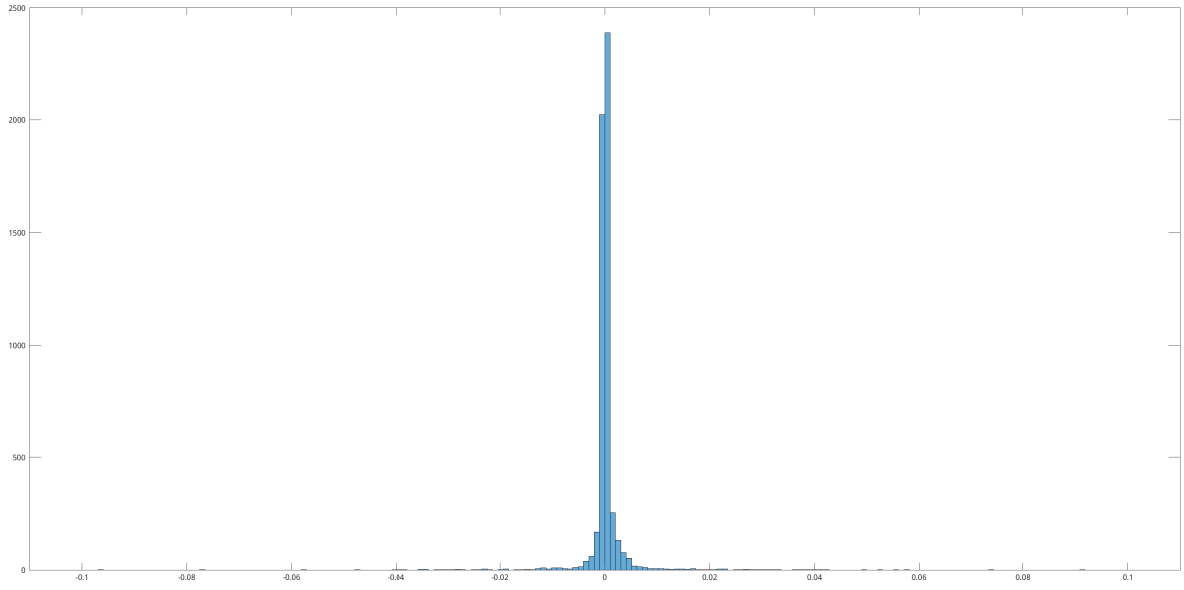


Figure 6: CORRELATIONS OF VAR RESIDUALS

there. Our disaggregated component model's median forecasts do not show that the unemployment rate would rise over the medium-term horizon while the median forecasts from the MF-VAR model do so. The posterior confidence interval is also narrower in our model forecasts, suggesting that it may have a better density forecast. While predicting strong labor market, our model does not forecast high inflation as shown in Figure 8.[20]

4.3 Nowcasting and Forecasting of GDP Growth

Using the posterior draws based on the pre-covid estimation sample, we generate tracking estimate of nowcasts for GDP growth for 2024Q2. We have not compiled our dataset at the weekly frequency before May 30, 2024. We have nine different vintages of our dataset from May 30, 2024 to July 23, 2024 before the BEA releases the first official estimate of GDP on July 25, 2024. At each vintage, we update the predictive density for 2024Q2 headline GDP growth by aggregating estimates of disaggregated GDP components as discussed in Section 2.

Figure 9 shows the nowcasting density for 2024 Q2 GDP growth. At later vintages, nowcasting density is more concentrated around the mode as it needs to fill fewer missing observations to generate GDP estimate. In addition, the shape of the predictive density changes over time non-trivially with the skewness increasing at data releases in the week of June 14, and June 25. This pattern illustrates the departure from the normal distribution as discussed in Section 2. The model's

[20] In contrast, the MF-VAR model's forecasts for the unemployment rate and the consumer price index level are negatively correlated at all the forecast horizons, suggesting that they match relatively low inflation data by projecting high unemployment.

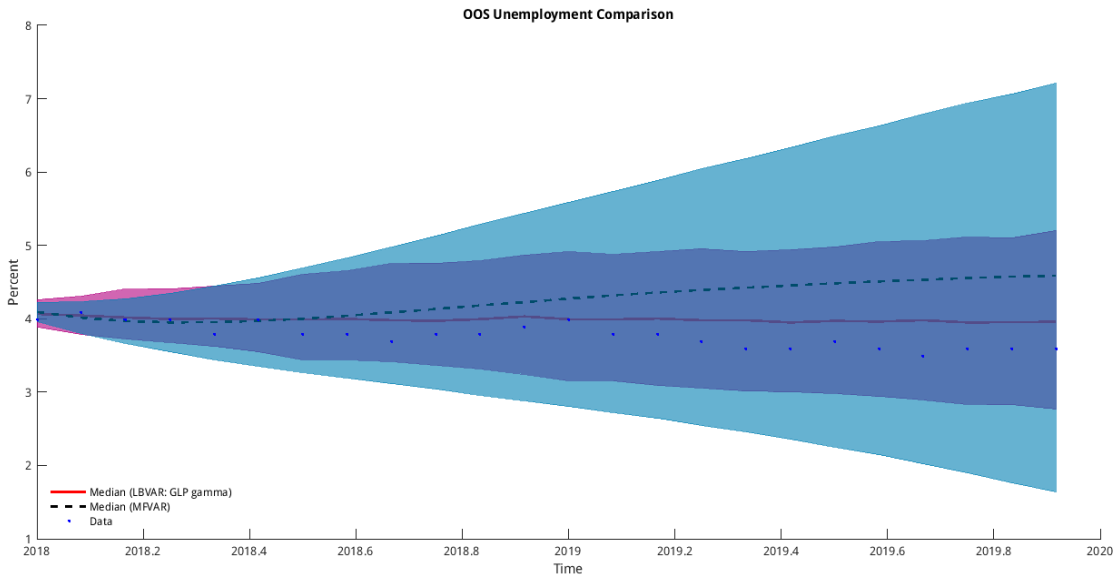


Figure 7: UNEMPLOYMENT FORECAST FOR 2018-2019 USING INFORMATION UP TO DEC. 2017

median estimate of 2024 Q2 GDP growth at the final vintage is 2.9 percent close to both the first estimate of 2.8 percent and the third estimate of 3 percent by the BEA.[21]

Table 1 compares our model's nowcasting estimates at different vintages with models maintained by Federal Reserve Banks of Atlanta, New York, respectively. We pick up these two models because they represent different frameworks. GDPNow from the Federal Reserve Bank of Atlanta takes the "bottom-up" approach similar to ours while the Federal Reserve Bank of New York's model relies on a statistical factor model. Not surprisingly, our model's estimates are closer to GDPNow than the factor model used by the Federal Reserve Bank of New York. In particular, the final vintage estimate of 2024 Q2 GDP growth of 2.6 percent from GDPNow is close to both the official first estimate of 2.8 percent and our final vintage estimate of 2.9 percent. In addition, revisions in nowcast estimates at different vintages are well correlated between our model and GDPNow with the correlation coefficient of 0.62. However, revisions in our model forecasts are nearly uncorrelated with those in the NY Fed's model, again illustrating the difference between the tracking model based on the accounting identity restriction and the factor model. Figure 11 illustrates the discrepancies in the weights of the first principal component on consumption from expenditure shares, which may account for the different interpretation from data surprises between the tracking model and the factor model.

When we decompose the contributions of consumption and investment components across

[21]We provide additional results for tracking of 2024 Q3 GDP growth and 2024 Q4 GDP growth in the appendix.

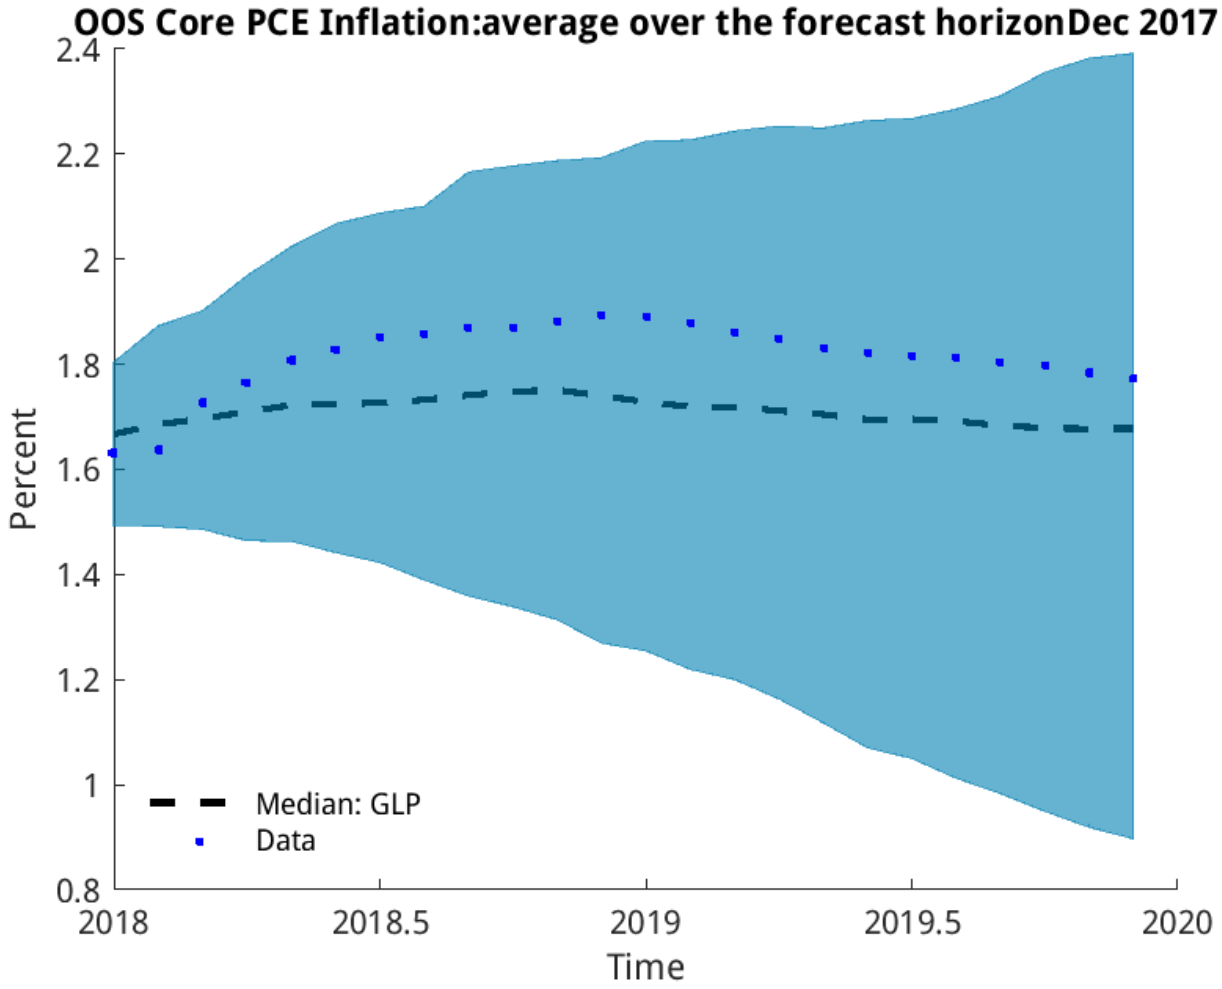


Figure 8: CORE PCE INFLATION FORECAST FOR 2018-2019 USING INFORMATION UP TO DEC. 2017

different data vintages using the formula in equation (2.14) at a more finer level for consumption as shown in Table 2, we see interesting co-movement patterns between consumption components and residential investment. While durable goods consumption is often bundled together with business investment, its' correlation with investment is not particularly higher than services consumption's correlation with investment. In particular, revisions in the contribution of residential investment is highly correlated with food & accommodation services and housing services consumption at least in terms of near-term predictions that matter to GDP nowcasts. Our VAR model uncovers these non-trivial correlations at disaggregated components, which would not be feasible in the simple aggregation of uni-variate forecasting model estimates. For instance, between July 17 and July 23 in Table 2, only existing home sales data (sales volume and price) were updated. While these indicators are directly used only for residential investment, they have non-trivial correlations with

Table 1: 2024 Q2 GDP NOWCAST COMPARISON

Vintage Date	GDPNow	NY Fed Model	BVAR (Median)	25%-75% Interval	5%-95% interval
May 31	2.7	1.76	3.69	[2.5, 4.77]	[0.65, 6.53]
June 7	3.1	1.9	2.99	[1.89, 4.24]	[-0.18. 6.15]
June 14	3,2	1.91	3.97	[2.69, 5.10]	[0.92, 6.72]
June 21	3	1.89	3.83	[2.73, 5.31]	[1.39, 7.12]
June 25	3	1.89	3.81	[2.60 , 5.15]	[0.89 , 6.93]
July 2	1.7	1.93	2.97	[2.05, 3.83]	[0.61, 5.21]
July 9	2	1.79	2.74	[1.91, 3.46]	[0.46, 4.55]
July 17	2.7	1.8	2.91	[1.91, 3.54]	[0.95, 4.49]
July 23	2.6	2.01	2.76	[1.85, 3.61]	[0.91. 4.7]

other indicators in the monthly BVAR, resulting in revisions of consumption nowcasts. Interestingly, business fixed investment components were not affected much but consumption components were revised down similar to residential investment.

One shortcoming in using equation (2.14) is that we cannot apply it to changes in inventories. To understand why nowcasts change in response to data releases beyond consumption and investment, we look into contributions of six major components to headline GDP growth implied by the model. Figure 10 show the revisions in component contributions to GDP growth. At the initial vintage, inventories were expected to drag growth while net exports were projected to positively contribute to growth. But as we have more data on May international trade and inventories, the model reversed the estimates of component contributions. At the final vintage, net exports were expected to drag GDP growth while inventories were expected to be neutral as shown in Figure 12. Nowcasting density of component contributions in Figure 12 suggests that the model forecasts stronger consumption but weaker business fixed investment than the BEA's estimates while roughly matching the private domestic final sales contribution. Except for government spending, official estimates of component contributions are within the plausible range of nowcasting density. The model appears to underestimate the contribution of government spending because it is overconfident about the forecast accuracy. The magnitude of the forecast miss is not particularly large with a 0.1 percentage point downward shift in the contribution of government spending enough to be within the plausible range of the model's nowcasting density.

Besides providing nowcasting estimates at different vintages of data, our model can also generate medium-term forecasts internally consistent with nowcasts. We regard this as a significant advantage of our probabilistic approach to the tracking model compared to GDPNow, which focuses on the point estimate of nowcasting. Our model can sequentially update the predictive density for the medium-term forecast at different vintages of data just as we update the predictive density of nowcasts. Figure 13 shows the predictive density of the one-year average growth rate from 2024:Q2

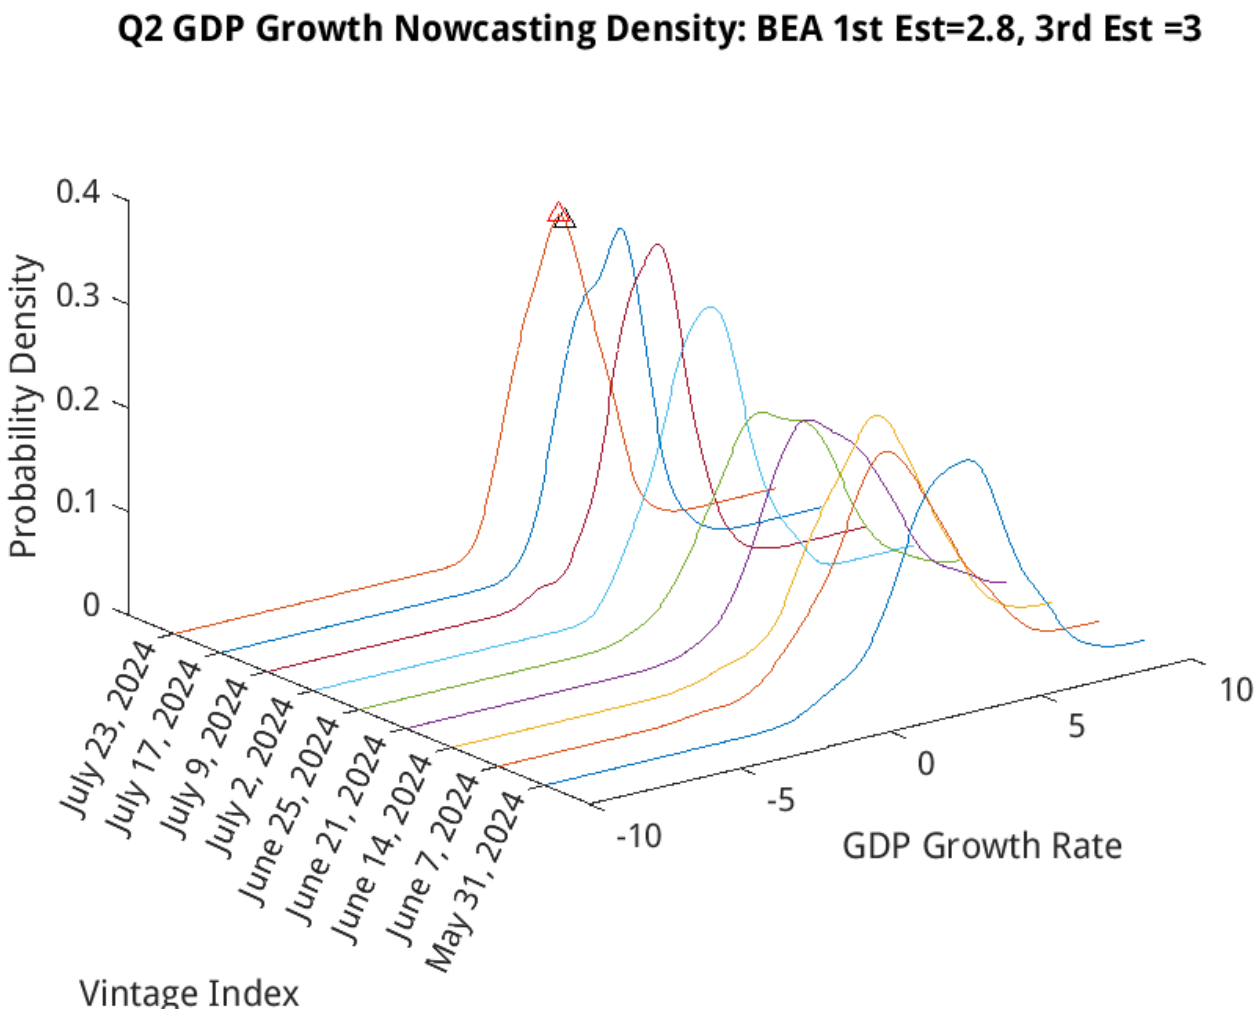


Figure 9: NOWCASTING DENSITY FOR 2024 Q2 HEADLINE GDP GROWTH

Notes: The red triangle mark represents the BEA's 1st estimate while the black triangle mark shows the BEA's 3rd estimate.

to 2025:Q1 across multiple data vintages. Consistent with nowcasts in Figure 9, additional data releases reduce the uncertainty implied in the predictive density. While the model expects near 3 percent for 2024Q2 at the final vintage, the median forecast for the average growth over 12 months is concentrated at the relatively low value of 1.1 percent. In other words, the model sees bigger downside risks from 2024:Q4. In particular, the model expects government spending and net exports will account for most of growth as the contribution from consumption declines as shown in Figure 14. However, compared to the predictive density of nowcasts, the magnitude of uncertainty is quite high with the shape of the predictive density fairly dispersed even at the last vintage of the data.[22]

[22]Of course, this is conditional on model estimates and data available up to July 23, 2024, which is the latest data

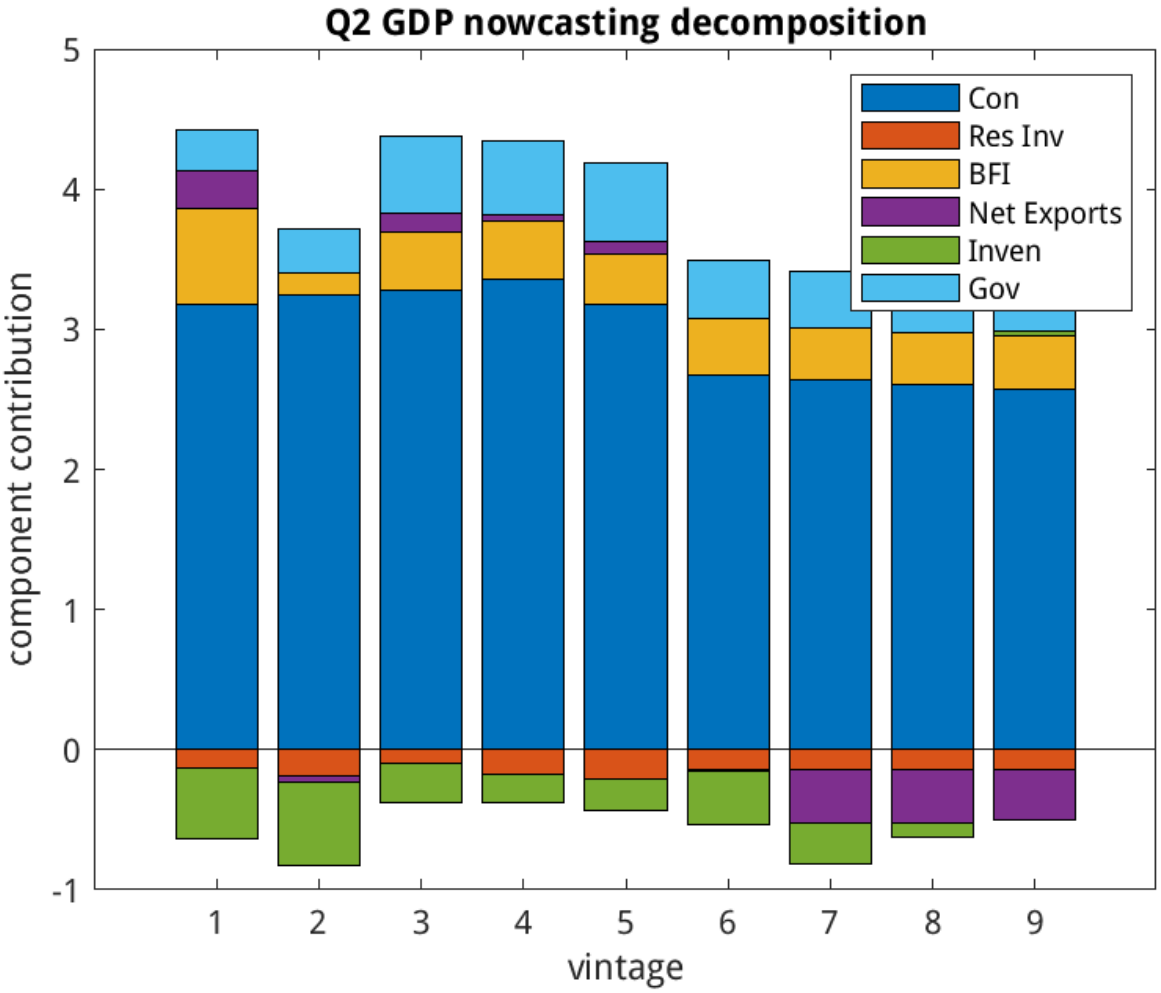


Figure 10: DECOMPOSITION OF 2024 Q2 GDP GROWTH NOWCAST AT DIFFERENT DATA VINTAGES

Notes: Stacked bars represent the median contribution of six components to 2024 Q2 GDP growth. Since we aggregate 38 subcomponent contribution to 6 components, the sum of the median contribution of 6 components is not necessarily equal to the median estimate of headline GDP growth, which is the median of the sum of all the 38 subcomponent contributions. 'Con', 'Res Inv', 'BFI', 'Inven', and 'Gov' denote consumption, residential investment, business fixed investment, changes in inventories, and government spending, respectively.

5 Conclusion

We propose and estimate a large-scale Bayesian VAR for monthly indicators of disaggregated GDP components and generate forecasts of major components of GDP as well as headline GDP while respecting the accounting identity restriction and the chain-weighted index method used by the

vintage for our 2024 Q2 nowcasting. The Federal Reserve started to cut the interest rate from September and signaled further easing down the road. To analyze the impact of policy actions on the medium-term forecast, we need a structural VAR model. We are working on this.

Table 2: 2024 Q2 GDP NOWCAST CHANGE DECOMPOSITION: CONSUMPTION AND INVESTMENT

Vintage Date	Dur. Con.	Nond. Con	Serv. Con	Res. Inv	Nonres. Str.	Eq. Inv	IPP
June 7	0.17	0.03	-0.11	-0.03	-0.25	-0.19	-0.01
June 14	0	0.03	-0.05	0	0.12	0.15	0
June 21	0.03	0.11	-0.09	-0.07	0	-0.04	0
June 25	-0.09	-0.02	-0.08	-0.02	0	-0.05	0
July 2	-0.36	-0.29	0.12	0.06	-0.09	0.2	-0.02
July 9	-0.06	0	0.05	0.02	-0.02	-0.03	-0.01
July 17	-0.09	-0.02	0.03	0	0.03	0.01	0.01
July 23	-0.01	-0.01	-0.07	-0.02	0	0	0

Notes: Numbers in the table represent the median estimate of changes in the contribution of private domestic final sales. 'Dur.. Con', 'Nond. Con','Serv. Con','Res Inv', 'Nonres. Str.', 'Eq. Inv.', and 'IPP' denote durable goods consumption, non-durable goods consumption, services consumption, residential investment, nonresidential structures, equipment investment, and intellectual property product, respectively.

statistical agency. Imposing these restrictions requires forecasts of both disaggregated real spending and price variables, substantially increasing the scale of the model. Given the large number of parameters, estimating the model faces computational challenges especially when we estimate the posterior distribution of the VAR model's hyperparameters using a hierarchical structure as done in Giannone et al. (2015). We address the computational challenge by using the sequential Monte Carlo method, which allows the massive parallelization in the posterior simulation. In addition, we draw VAR parameters conditional on hyperparameters more efficiently by relying on a matrix normal distribution.

Our model generates forecasts for aggregate variables such as the unemployment rate and inflation comparable or superior to a medium-sized mixed-frequency model. In addition, the model can sequentially update GDP nowcast and forecasts in response to incoming data. We compare the model's nowcasts for 2024 Q2 GDP growth at different data vintages with two benchmark GDP nowcasting models: 1) GDPNow by the Federal Reserve Bank of Atlanta, and 2) GDP nowcasting model by the Federal Reserve Bank of New York. Since GDPNow incorporates the accounting identity restriction like our model, our model forecasts are quite comparable to their forecasts but nearly uncorrelated to the GDP nowcasting model by the Federal Reserve Bank of New York, which is based on a factor model not imposing the accounting identity restriction. In spite of the similarity in the point forecast, our model improves GDPNow in some dimensions. First, we complete characterize the probabilistic structure of our forecasts for headline GDP and generate the entire predictive distribution for nowcasts. Second, our method allows medium-term forecasts internally consistent with nowcasts and can allow us to sequentially update the entire forecast distribution of GDP growth at different horizons as we receive incoming data. We highlight how incoming data

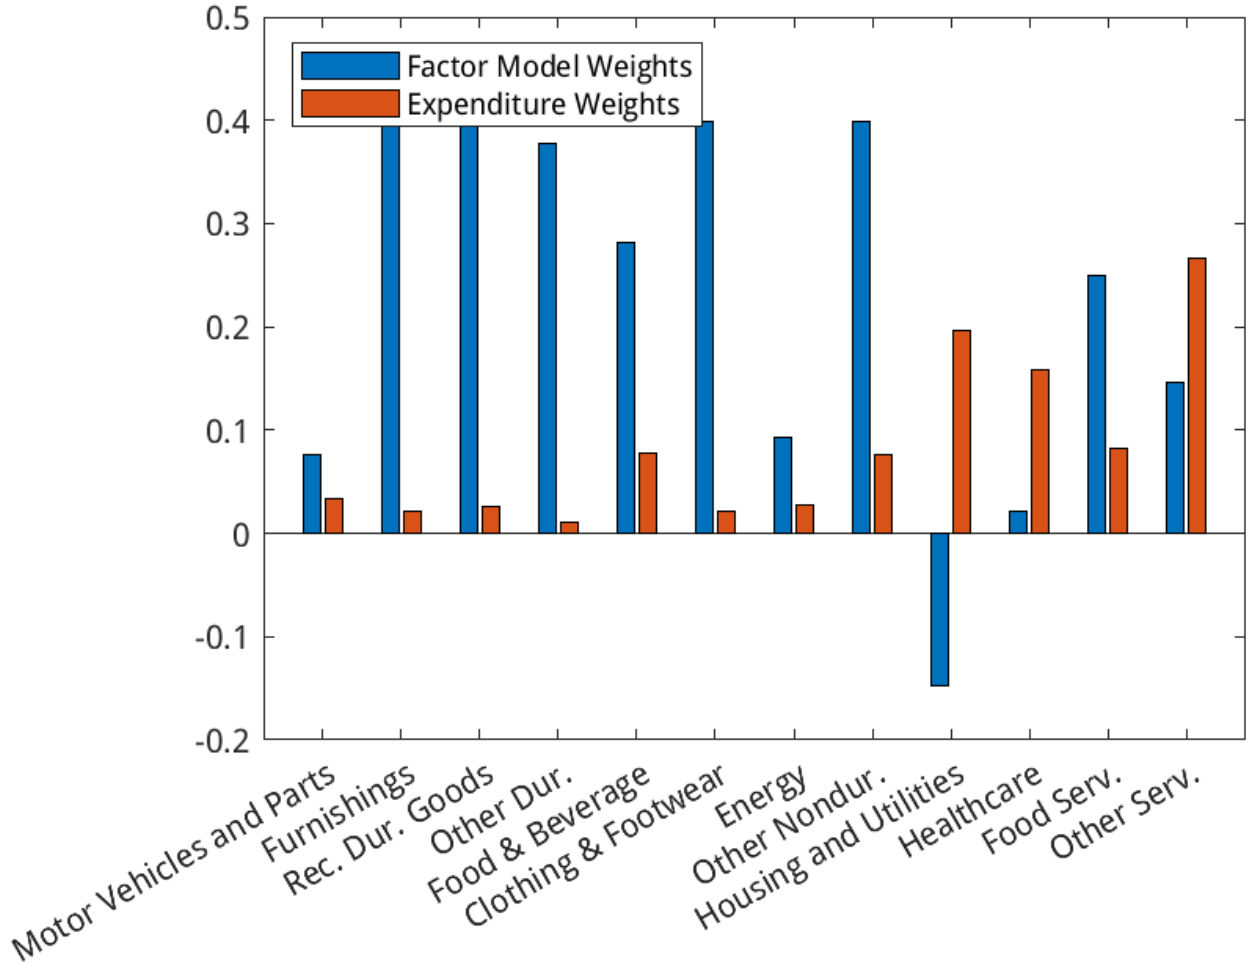


Figure 11: PRINCIPAL COMPONENT WEIGHTS VERSUS EXPENDITURE SHARES: CONSUMPTION COMPONENTS

reduces the uncertainty surrounding nowcasts and medium-term forecasts. Our model also provides the decomposition of headline GDP growth to component contributions, which is useful for creating the forecast narrative.

While the current empirical analysis is restricted to using the reduced-form VAR, the scale and scope of our model can shed new light on the identification of structural shock. A large number of spending indicators whose sensitivity to the interest rate is quite heterogeneous can be used to identify a monetary policy shock. Similarly, we can study how fiscal policies, which target certain sectors of the economy (e.g., infrastructure and manufacturing facilities), propagate into the aggregate economy better than a small-scale model confined to aggregate variables. We expect that the large scale of our model will create new research opportunities.

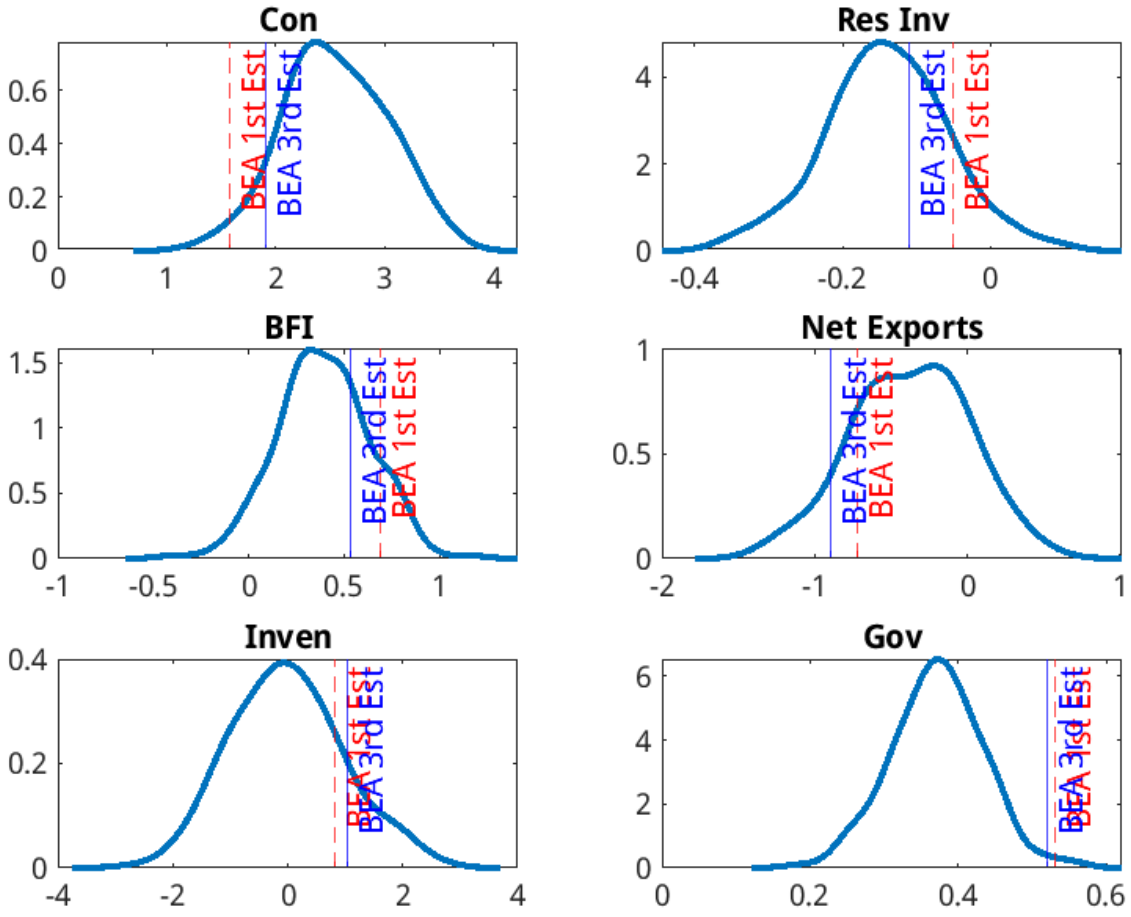


Figure 12: NOWCASTING DENSITY OF COMPONENT CONTRIBUTION TO 2024 Q2 GDP GROWTH AT THE FINAL VINTAGE

Notes: 'Con', 'Res Inv', 'BFI', 'Inven', and 'Gov' denote consumption, residential investment, business fixed investment, changes in inventories, and government spending, respectively.

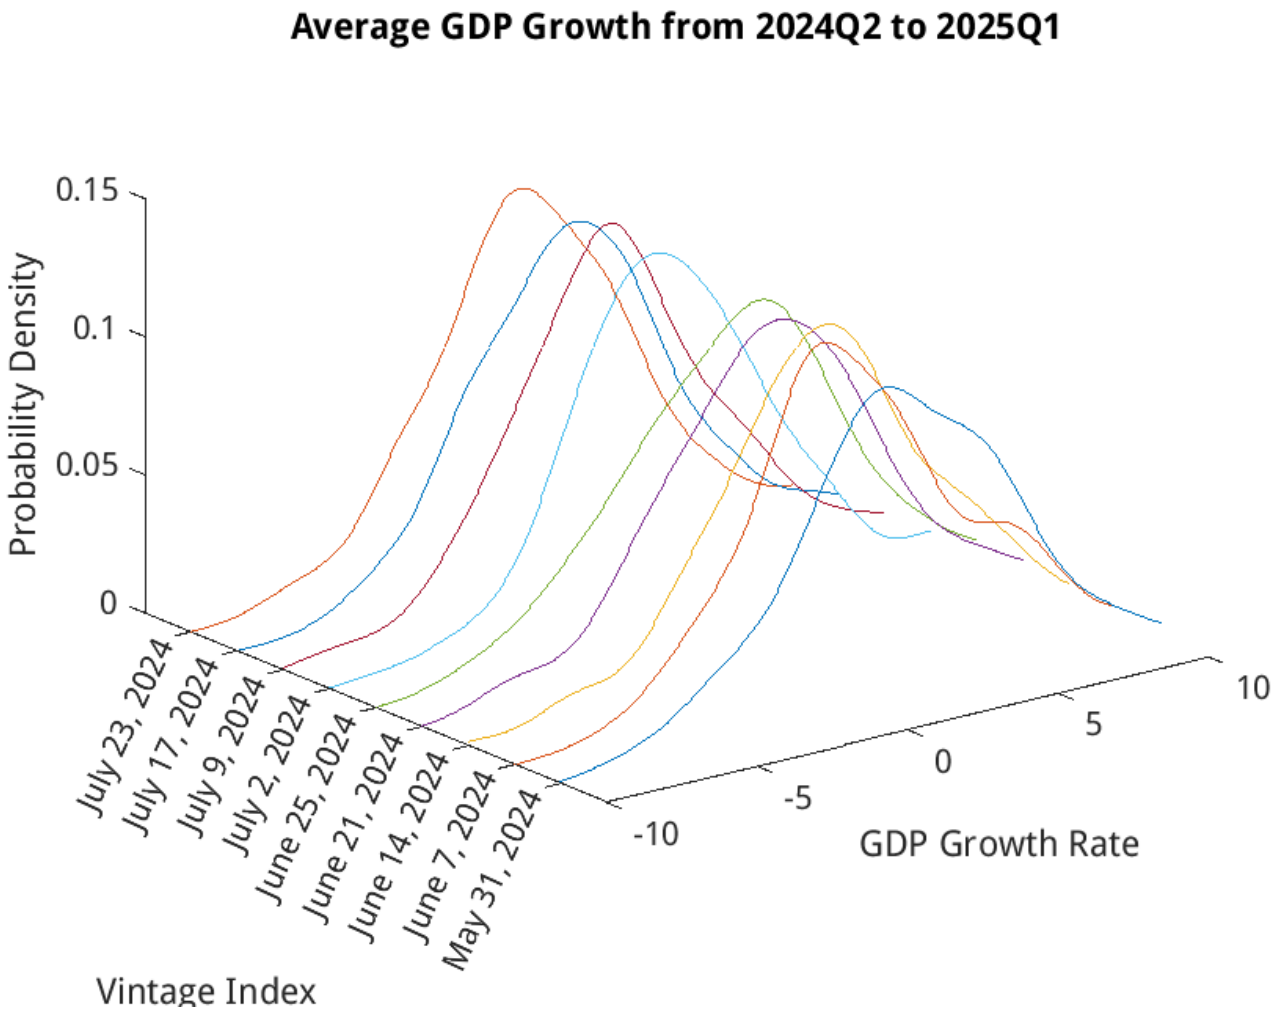


Figure 13: PREDICTIVE DENSITY FOR THE AVERAGE ONE-YEAR GDP GROWTH RATE

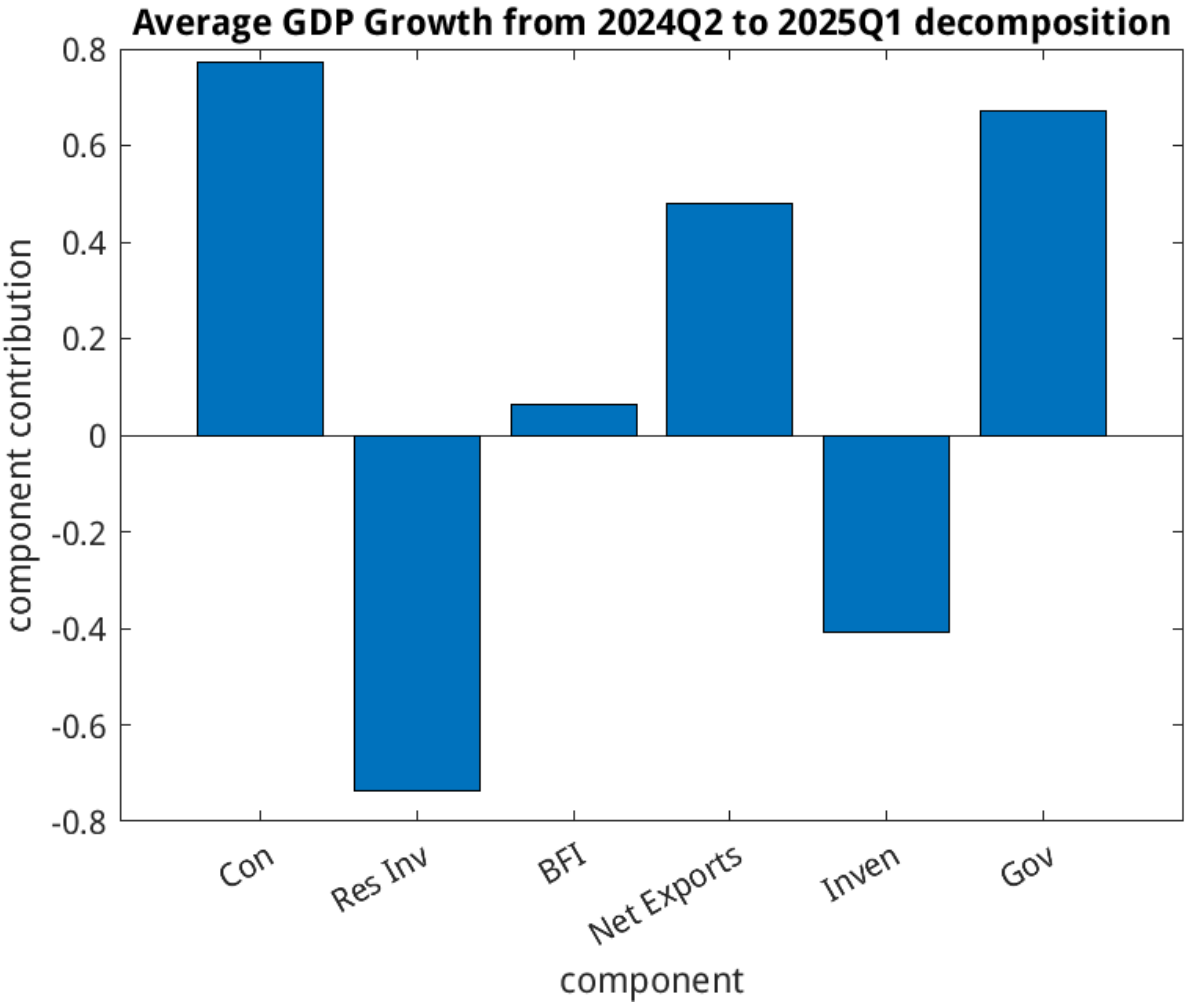


Figure 14: DECOMPOSITION OF THE AVERAGE ONE-YEAR GDP GROWTH RATE

Notes: The decomposition is based on the median forecast of component contributions using the July 23, 2024 data vintage. 'Con', 'Res Inv', 'BFI', 'Inven', and 'Gov' denote consumption, residential investment, business fixed investment, changes in inventories, and government spending, respectively.

References

Adrian, Tobias, Nina Boyarchenko, and Domenico Giannone, "Vulnerable growth," *American Economic Review*, 2019, *109* (4), 1263–1289.

Angelini, Elena, Magdalena Lalik, Michele Lenza, and Joan Paredes, "Mind the gap: A multi-country BVAR benchmark for the Eurosystem projections," *International Journal of Forecasting*, 2019, *35* (4), 1658–1668.

Aruoba, S Borağan, Francis X Diebold, and Chiara Scotti, "Real-time measurement of business conditions," *Journal of Business & Economic Statistics*, 2009, *27* (4), 417–427.

Bae, Jaeheung and Taeyoung Doh, "Tracking US GDP in Real Time," *Economic Review*, 2019, (Q III), 5–19.

Baker, Katie, Martín Almuzara, Hannah O'Keeffe, and Argia M Sbordone, "Reintroducing the New York Fed Staff Nowcast," Technical Report, Federal Reserve Bank of New York 2023.

Bańbura, Marta, Domenico Giannone, and Lucrezia Reichlin, "Large Bayesian vector auto regressions," *Journal of applied Econometrics*, 2010, *25* (1), 71–92.

Bauwens, Luc, Michel Lubrano, and Jean-François Richard, "Bayesian Inference in Dynamic Econometric Models," 2000.

Bognanni, Mark and Edward Herbst, "A sequential Monte Carlo approach to inference in multiple-equation Markov-switching models," *Journal of Applied Econometrics*, 2018, *33* (1), 126–140.

Brave, Scott A, R Andrew Butters, and Alejandro Justiniano, "Forecasting economic activity with mixed frequency BVARs," *International Journal of Forecasting*, 2019, *35* (4), 1692–1707.

Cascaldi-Garcia, Danilo, Matteo Luciani, and Michele Modugno, "Lessons from Nowcasting GDP across the World," 2023.

Chan, Joshua CC, Aubrey Poon, and Dan Zhu, "High-dimensional conditionally Gaussian state space models with missing data," *Journal of Econometrics*, 2023, *236* (1), 105468.

Chen, Kaiji, Patrick Higgins, and Tao Zha, "Constructing quarterly Chinese time series usable for macroeconomic analysis," *Journal of International Money and Finance*, 2024, *143*, 103052.

Crump, Richard K, Stefano Eusepi, Domenico Giannone, Eric Qian, and Argia Sbordone, "A Large Bayesian VAR of the United States Economy," 2021, (976).

Drechsel, Thomas and Ivan Petrella Antolín-Díaz Juan, "Advances in nowcasting economic activity: The role of heterogeneous dynamics and fat tails," *Journal of Econometrics*, 2024, *238* (2), 1–25.

Fernandez, Roque B, "A Methodological Note on the Estimation of Time Series," *The Review of Economics and Statistics*, 1981, *63* (3), 471–476.

Gao, Xin, Hong Xu, and Dong Ye, "Asymptotic Behavior of Tail Density for Sum of Correlated Lognormal Variables," *International Journal of Mathematics and Mathematical Sciences*, 2009, *2009* (1).

Giannone, Domenico, Lucrezia Reichlin, and David Small, "Nowcasting: The real-time informational content of macroeconomic data," *Journal of monetary economics*, 2008, *55* (4), 665–676.

_ , Michele Lenza, and Giorgio E Primiceri, "Prior selection for vector autoregressions," *Review of Economics and Statistics*, 2015, *97* (2), 436–451.

_ , _ , and _ , "Economic predictions with big data: The illusion of sparsity," *Econometrica*, 2021, *89* (5), 2409–2437.

Gulisashvili, Archil and Peter Tankov, "Tail behavior of sums and differences of log-normal random variables," *Bernoulli*, 2016, *22* (1), 444–493.

Herbst, Edward P and Frank Schorfheide, *Bayesian estimation of DSGE models*, Princeton University Press, 2016.

Higgins, Patrick C., "GDPNow: A Model for GDP Nowcasting," Working Paper 2014-7, Federal Reserve Bank of Atlanta 2014.

Klein, L. R. and J. Y. Park, "Economic forecasting at high-frequency intervals," *Journal of Forecasting*, 1993, *12* (3-4), 301–319.

Negro, Marco Del and Frank Schorfheide, "293 Bayesian Macroeconometrics," in "The Oxford Handbook of Bayesian Econometrics," Oxford University Press, 2011.

Schorfheide, Frank and Dongho Song, "Real-time forecasting with a mixed-frequency VAR," *Journal of Business & Economic Statistics*, 2015, *33* (3), 366–380.

_ and _ , "Real-time forecasting with a (standard) mixed-frequency VAR during a pandemic," Technical Report, National Bureau of Economic Research 2021.

Whelan, Karl, "A guide to US chain aggregated NIPA data," *Review of income and wealth*, 2002, *48* (2), 217–233.

Appendix

A. Probability Density Function of the Log Sum of Log Normal Random Variables

Based on the result in Gao, Xu and Ye (2009), we illustrate how the probability density function of the log sum of log normal random variables can be decomposed into the term describing the normal distribution and the term representing the departure from the normality.

Let X_1 and X_2 be two independent standard normal random variables and define z and $\hat{z}$ as follows:

$$\begin{aligned} z &= X_1 + X_2, \\ \hat{z} &= \ln(e^{X_1} + e^{X_2}). \end{aligned}$$

While z follows the normal distribution $\mathcal{N}(0, 2)$, the distribution of $\hat{z}$ is not normal but can be characterized by the following density function as shown in Gao et al. (2009).

$$\begin{aligned} h(\hat{z}) &= \frac{1}{2\pi} \int_0^1 \frac{1}{(1-t)t} e^{-0.5[(\hat{z}+\ln(1-t))^2 + (\hat{z}+\ln t)^2]} dt, \\ &= \frac{1}{\sqrt{\pi}} e^{-\hat{z}^2} \frac{1}{\sqrt{4\pi}} \int_0^1 \frac{1}{(1-t)t} e^{-0.5[(\ln(1-t))^2 + (\ln t)^2 + 2\hat{z}(\ln(1-t)+\ln t)]} dt, \\ &= h_1(\hat{z}) h_2(\hat{z}). \end{aligned}$$

$h_1(\hat{z})$ represent the probability density function of a normal random variable with mean zero and variance equal to 0.5 while $h_2(\hat{z})$ characterize the departure from the normality.

For comparison, the probability density function of z is given by

$$p(z) = \frac{1}{\sqrt{4\pi}} e^{-\frac{z^2}{4}} = \frac{h_1(z)}{2}.$$

B. Extrapolation of CPI Data

In order to extend the balanced data set back to June 2006, two series were extrapolated using multiple regression. The US city average for CPI-U: Household Furnishings and Supplies and CPI-U: Recreation Commodities both begin in January 2010. To extend them back, each were regressed on a number of subordinate series that have observations dating back to December 1999. All data used in the regressions was seasonally adjusted, if any series was not available seasonally

adjusted from Bureau of Labor Statistics (BLS), a standard multiplicative seasonal adjustment function was applied.[23] Table 3 shows the series and their regressors. Note that all series listed are CPI-U US city averages and series used with seasonal adjustment from the BLS are denoted (SA) and those manually adjusted are marked (NSA).

Table 3: CPI REGRESSANDS AND REGRESSORS

Regressand	**Regressors**
Household Furnishings and Supplies (SA)	Other Furniture (SA) Other Household Appliances (SA) Indoor Plants and Flowers (SA) Outdoor Equipment and Supplies (SA) Floor Coverings (NSA) Window Coverings (NSA) Other Household Linens (NSA) Bedroom Furniture (NSA) Living Room, Kitchen and Dining Room Furniture (NSA) Household Major Appliances Clocks, Lamps and Decorator Items (NSA) Dishes and Flatware (NSA) Nonelectric Cookware & Tableware (NSA) Tools, Hardware and Supplies (NSA) Household Cleaning Products (NSA) Household Paper Products (NSA) Misc Household Products (NSA)
Recreation Commodities (SA)	Televisions (SA) Other Video Equipment (SA) Audio Equipment (NSA) Recorded Music & Music Subscriptions (NSA) Pet Food (NSA) Purchases of Pets, Pet Supplies & Accessories (NSA) Sports Vehicles including Bicycles (NSA) Sports Equipment (NSA) Photographic Equipment (NSA) Newspapers and Magazines (NSA) Recreational Books (NSA) Toys, Games, Hobbies & Playground Equipment (NSA) Sewing Machines, Fabric and Supplies (NSA) Music Instruments and Accessories (NSA)

Each of the main series was regressed on the selected regressors and a constant over the overlapping period (January 2010 to March 2024). Then the ordinary least squares (OLS) coefficients were used to extrapolate the two CPI series back to December 1999.

[23]To calculate seasonal factors, we use the earliest six years in the estimation sample as data.

C. Construction of Real Monthly Spending Data

To construct the real monthly indicator data set, all data not directly available in real terms was deflated with a corresponding price index. Table 4 shows the connections between the nominal spending series and their price data. Note that monthly PCE variables are available in real terms and do not appear here. Four series, marked (Q) below, are deflated with quarterly price data. This interpolation is done by assigning the given quarterly value to each of the constituent months (e.g. a value of 125 in Q1 results in January, February, and March being assigned a value of 125). Certain series are seasonally adjusted with the standard seasonal adjustment function used for this paper, they are marked (NSA). One series, CPI-U: Food Away From Home, is seasonally adjusted using the standard seasonal adjustment done by Haver and is marked (NSAH). Finally, two series, CPI-U: Household Furnishings and Supplies and CPI-U: Recreation Commodities have been extrapolated as described in Section 1 of the appendix and are marked (Ext.).

D. Comparison with NIPA Construction

The construction of the spending and price data mimic the structure of data and monthly indicators described by the NIPA handbook. Tables 5 and 6 show examples of the construction for Residential Investment from our model and from the NIPA handbook respectively.

E. Sequential Monte Carlo Procedure

We describe the details of the sequential Monte Carlo procedure that we use to estimate the monthly VAR. We denote the vector of hyperparameters in the monthly VAR by θ^m.

- We use a parallelization in the implementation because ψ is high-dimensional.
- **SMC Step1**: Draw equally weighted M initial particles, $\theta_{1,i}^m$ from the prior distribution $p(\theta^m)$, with $w_{1,i} = \frac{1}{M}$.
- **SMC Step2.1**: For $l = 2, \ldots, L$, calculate new weights for the particles from stage $l-1$ based on $p(Z_{1:T}^m|\theta_{l-1,i})^{\phi_l - \phi_{l-1}}$ where $0 < \phi_1 < \cdots < \phi_L = 1$ are tempering parameters.
- **SMC Step2.2**: Resample with replacements $\hat{\theta}_i^m$ based on normalized weights $w_{l,i} = \frac{p(Z_{1:T}^m|\theta_{l-1,i})^{\phi_l - \phi_{l-1}} w_{l-1,i}}{\frac{1}{M}\sum_{j=1}^M p(Z_{1:T}^m|\theta_{l-1,j})^{\phi_l - \phi_{l-1}} w_{l-1,j}}$.
- **SMC Step2.3**: Mutate $\hat{\theta}_i^m$ using the random-walk Metropolis-Hastings algorithm $K_l(\hat{\theta}_i^m, \zeta_l)$ and obtain $\theta_{l,i}^m$ associated with $w_{l,i}$.
- **SMC Step3**: Use $\theta_{L,i}^m$ and $w_{L,i}$ for posterior simulation.

Table 4: CONSTRUCTION OF REAL MONTHLY DATA

Nominal Spending Series	**Price Index**
Pvt. Construction: Residential: New Single Family (Census)	Single-Family Houses Under Construction: Price Deflator (NSA, Census)
Pvt. Construction: Residential: New Multi-family (Census)	Pvt. Invest.: Multifam. Res. Structures: Chain-type Price Index (Q, BEA)
Ret. Sales: Building Mat., Garden Equip. & Supply Dealers (Census)	PPI: Inputs to Maint. & Repair Const., Goods (NSA, BLS)
Avg. Hourly Earnings of Prod. & Nonsup. Employees: Residential Remodelers (BLS)	ECI: Compensation: Pvt. Industry Workers: Const. (Q, BLS)
Pvt. Construction: Nonresidential (Census)	Constructed Const. Price Index
Mfrs' Shipments: Nondefense Capital Goods excluding Aircraft (Census)	PPI: Capital Equipment (BLS)
Mfrs' Shipments: Nondefense Aircraft & Parts (Census)	PPI Civilian Aircraft (NSA, BLS)
Exports: Goods, BOP Basis (Census)	Export Price Index: All Exports (NSA, BLS)
Imports: Goods, BOP Basis(Census) - Imports: Petroleum End-Use Commodity Category (Census)	Exports of Services: Chain Price Index (Q, BEA)
US Imports of Services (Census)	Import Price Index: Nonpetroluem Imports (NSA, BLS)
Mfrs' Inventories: Durable Goods (Census) + Mfrs' Inventories: Nondurable Goods Industries (Census)	Imports of Services: Chain Price Index (Q, BEA)
Merchant Wholesalers: Inventories: Total (Census)	PPI: Capital Equipment (BLS)
Retail Inventories: Total (Census)	PCE: Chain Price Index (BEA)
Retail Inventories: Motor Vehicle & Parts Dealers (Census)	PCE less Food & Energy: Chain Price Index (BEA)
Ret. Inventories: Total ex. Motor Vehicle & Parts Dealers (Census)	CPI-U: New Vehicles (BLS)
Sum of Federal Outlays for Agriculture, Community/Regional Development, Energy, General Govt., International Affairs, Admin. of Justice, Natural Resources/Environment, General Science/Technology, Transportation (NSA, Treasury)	PCE: Chain Price Index (BEA)
Federal Outlays: National Defense (NSA, Treasury)	PCE: Chain Price Index (BEA)
Value of Total Federal Construction Put in Place (Census)	PCE: Chain Price Index (BEA)
Value of State and Local Construction Put in Place (Census)	PPI: New Office Building Construction (BLS)
Retail Sales: Motor Vehicle & Parts Dealers (Census)	CPI-U: New and Used Motor Vehicles
Ret. Sales: Furniture & Home Furn. & Elec./Appliance Stores (Census)	CPI-U: US City Avg.: Household Furn. and Supplies (Ext., BLS)
Ret. Sales: Sporting Goods, Hobby, Book & Music Stores (Census)	CPI-U: US City Avg.: Recreation Commodities (Ext., BLS)
Retail Sales: Food & Beverage Stores (Census)	CPI-U: Food (BLS)
Retail Sales: Clothing & Accessory Stores (Census)	CPI-U: Apparel (BLS)
Retail Sales: Gasoline Stations (Census)	CPI-U: Energy Commodities (BLS)
Retail Sales: Food Services & Drinking Places (Census)	CPI-U: Food Away From Home (NSAH, BLS)
PCE: Other Services (constructed, BEA)	Other Services Constructed Price Index (constructed, BEA)

We set $M = 3000$ and $L = 500$. In the mutation step, we divide the parameter vector into 12 blocks and apply the random walk Metropolis-Hastings algorithm for each block.

Table 5: STRUCTURE OF INDICATORS AND PRICE DATA FROM THE

Quarterly Component	Monthly Indicator	Price Data
Single Family Structures (BEA)	Pvt Const.: New Single Fam. (Census)	Single Fam. Houses under Const. (Census) PPI: Inputs to Maintenance (BLS) Constructed Const. Price Index (BLS)
Multifamily Structures (BEA)	Pvt. Const.: New Multifam. (Census)	Single Fam. Houses under Const. (census) PPI: Inputs to Maintenance (BLS) Constructed Const. Price Index (BLS)
Broker Commission Costs (BEA)	Existing Single Fam. Sales (NAR) New Single Fam. Sales (Census)	Median Sales Price: Existing Single Fam. Homes (NAR)
Improvements on Res. Structures (BEA)	Ret. Sales: Building Materials, etc. (Census) Avg Hourly Earnings of Residential Remodelers (BLS)	PPI: Inputs to Maintenance (BLS) Avg Hourly Earnings: Const. (BLS)

Table 6: STRUCTURE OF INDICATORS AND PRICE DATA FROM THE NIPA HANDBOOK

Quarterly Component	Monthly Indicator	Price Data
Single Family Structures (BEA)	Value put in place (based on housing starts and average construction costs) (Census)	New Single Fam. Houses Under Const. (Census)
Multifamily Structures (BEA)	Value put in place (Census)	New Multifam. Houses Under Const. (Census)
Broker Commission Costs (BEA)	# of Houses Sold (Census, NAR) Mean Sales Price (Census, NAR) Avg Commission Rate (BEA)	New Single Fam. Houses Under Const. (Census) PPI: Home Maint. & Repair Const. (BLS) ECI: Construction (BLS)
Improvements on Res. Structures (BEA)	Monthly Retail Trade Survey (building supply retail sales, Census) Circana	New Single Fam. Houses Under Const. (Census) PPI: Home Maint. & Repair Const. (BLS) ECI: Construction (BLS)

F. Additional Results for 2024 Q3 GDP and 2024 Q4 GDP Nowcasting

Figure 15 and 16 show the evolution of real GDP growth nowcasting density across different data vintages for the third quarter (Q3) and the fourth quarter (Q4) of 2024, respectively. In both cases, the final vintage median estimates of the model were 2.9% and 2.5% for Q3 and Q4, which were close to the BEA's initial estimate of 2.8% and 2.3%.

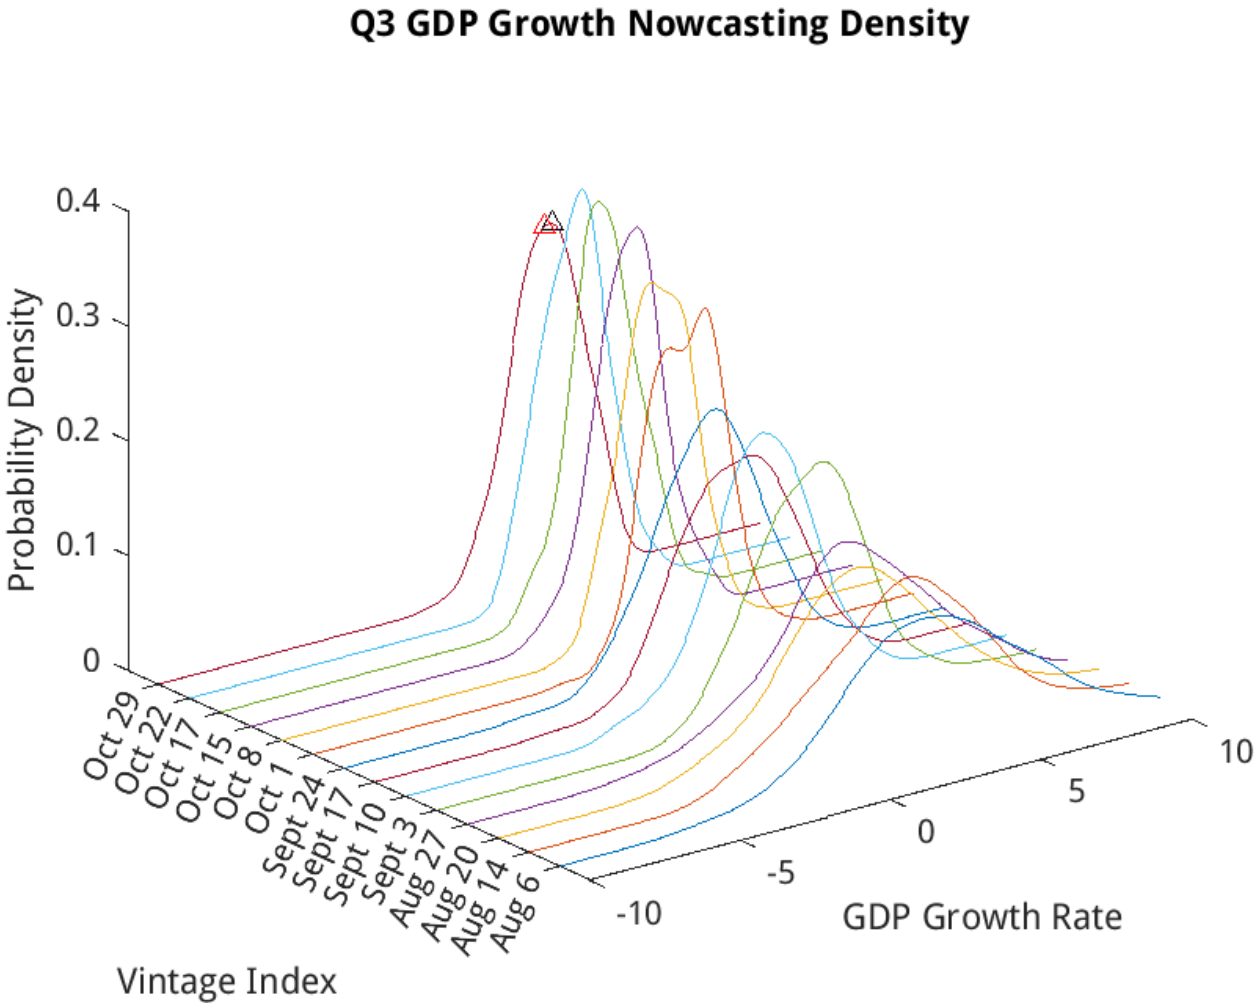


Figure 15: NOWCASTING DENSITY FOR 2024 Q3 HEADLINE GDP GROWTH

Notes: The red triangle mark represents the BEA's 1st estimate while the black triangle mark shows the BEA's 3rd estimate.

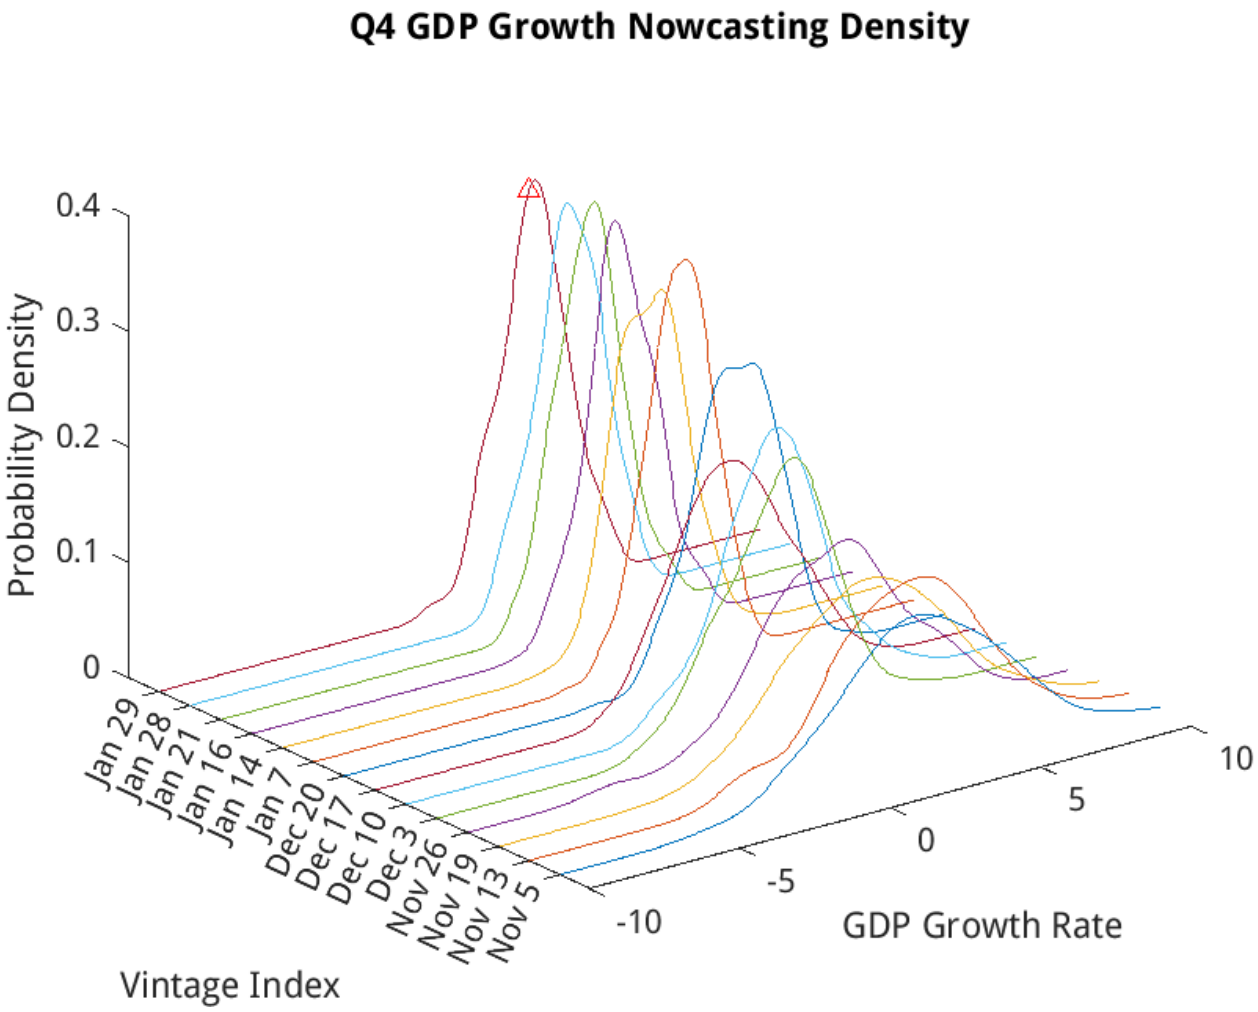


Figure 16: NOWCASTING DENSITY FOR 2024 Q4 HEADLINE GDP GROWTH

Notes: The red triangle mark represents the BEA's 1st estimate.