Technical appendix for
"Looking down the road with ALEX: Forecasting U.S. GDP,"
Chicago Fed Letter, No. 447

An introduction to ALEX: A mixed-frequency Bayesian VAR forecast for U.S. GDP

Scott A. Brave, R. Andrew Butters, and Michael Fogarty

October 20, 2020

ALEX is a mixed-frequency Bayesian vector autoregressive model (MF-BVAR) capturing the co-movement in 107 monthly and quarterly indicators of U.S. economic activity that can be used to forecast the near-term path of U.S. gross domestic product (GDP).[1] ALEX was designed to efficiently process the daily flow of information for the major monthly and quarterly statistical releases that are tracked by public sector and private sector analysts in order to understand the real-time evolution of U.S. GDP. Many of its variables themselves serve as source data used by the U.S. Bureau of Economic Analysis in its GDP calculations.

It is important to note, however, that ALEX is a reduced-form statistical model. It should not be confused with modern structural macroeconometric models, such as dynamic stochastic general equilibrium (DSGE) models. In some cases, these models produce vector autoregressive representations like ALEX, but they embody more than just statistical assumptions about the nature of their data series. Furthermore, ALEX's forecasts are purely model based. No judgment or economic theory is used to adjust them. Even the priors used to estimate ALEX are data driven, as the hyperparameters are selected using empirical Bayesian methods.

Similar models have been used in the past to forecast both U.S. and regional economic activity with varying degrees of success.[2] A predecessor to ALEX, for instance, was shown in Brave, Butters, and Justiniano (2019) to be equally accurate on average in producing forecasts of U.S. real GDP growth zero to two quarters ahead and more accurate three to four quarters ahead in comparison with surveys of private sector forecasts. ALEX incorporates many additional variables relevant for U.S. GDP that are used by forecasters to form more accurate current-quarter real GDP growth forecasts in order to improve upon these earlier results.

In what follows, we describe the state-space framework and priors used to estimate ALEX and demonstrate how point and density forecasts for GDP can be constructed from it.

[1] The name ALEX is a tribute to our dear departed colleague Alejandro Justiniano, or Alex to many of his friends, who was instrumental in its development.

[2] See, e.g., the models used by the Center for Research on the Wisconsin Economy (CROWE) to forecast the U.S. and Wisconsin economies and the Schorfheide and Song model for the U.S.

1 Estimating a mixed-frequency VAR

Here, we explain in detail the construction of the state-space system facilitating the construction of the set of conditional distributions that can be used to estimate a mixed-frequency Bayesian vector autoregressive model, such as ALEX.

1.1 Building the state-space system

We consider an n-dimensional vector y_t of macroeconomic time series of differing frequencies (i.e., some monthly variables and some quarterly variables).[3] Because of the mixed-frequency nature of the series in y_t, not all of the variables within it are observed every period. Therefore, we partition $y_t^{'} = \left[y_t^{q'} \; y_t^{m'} \right]$ such that the first n_q elements collect the vector y_t^q of quarterly variables (such as real GDP), which are observed only once every three periods in a monthly model. In turn, we let y_t^m be made up solely of monthly variables (such as real personal consumption expenditures), with dimension $n_m = n - n_q$.

To describe the monthly dynamics of this system, we let x_t^q denote the monthly latent variables underlying the quarterly series, y_t^q. We combine these latent variables with the indicators observed at a monthly frequency in $x_t^{'} = \left[x_t^{q'} \; x_t^{m'} \right]$. Clearly, each element of x_t^m corresponds to the element of y_t^m when observed. In contrast, some aggregated combination of past x_t^q monthly realizations will equal y_t^q when the quarterly variables are observed. In general, the aggregation for some series i is deterministic and given by

$$y_t^q(i) = G_i(x_t^q(i), x_{t-1}^q(i), ..., x_{t-s}^q(i))$$

for some predetermined horizon s.[4] An example of $G_i(\cdot)$, common for measures of economic activity in levels, is the three-month average of x_t^q, such that

$$y_t^q(i) = \frac{x_t^q(i) + x_{t-1}^q(i) + x_{t-2}^q(i)}{3}. \tag{A1}$$

With the mapping of x_t to y_t determined, the vector x_t and its monthly dynamics are summarized by the vector autoregression of order p given by

$$x_t = c + \phi_1 x_{t-1} + ... + \phi_p x_{t-p} + \epsilon_t, \quad \epsilon_t \sim \text{i.i.d.} N(0, \Sigma), \tag{A2}$$

where each ϕ_l is an n-dimensional square matrix containing the coefficients associated with lag l.[5] The companion form of this monthly VAR together with a measurement equation for y_t delivers the common two-equation state-space system given by

$$y_t = Z_t s_t, \tag{A3}$$
$$s_t = C_t + T_t s_{t-1} + R_t \epsilon_t, \tag{A4}$$

[3] A complete list of the 107 time series used in ALEX can be found in table A1.

[4] We follow the approach of Mariano and Murasawa (2003) and treat the quarterly observations of GDP and its subcomponents as the quarterly *average* of the monthly realizations. This leads to the interpretation that the underlying monthly variable is annualized.

[5] ALEX uses a lag order of three.

with the vector of observables, y_t, defined previously, and the state vector, s_t, defined as

$$s'_t = \left[x'_t, \ldots, x'_{t-p}, \zeta'_t\right],$$

which includes both lags of the time series at the monthly frequency and ζ_t, a vector of accumulators. For all quarterly indicators, the accumulator is defined by equation (A1).

Given the additional variables in the state (the accumulators, ζ_t), the transition matrix is an $np + n_q$ square matrix. In the transition matrix, the entries of the first n rows are the concatenation of the coefficients associated with each lag $\phi = [\phi_1, \phi_2, ..., \phi_p]$. The last n_q rows are made up of two separate components. The first component involves a (time-varying) scaled version of the coefficients associated with the quarterly time series and corresponds to the current monthly contribution to the accumulator series. The second component involves a deterministic series of fractions (e.g., 0, 1/2, and 1/3 for the regular average) that loads onto the lagged value of the accumulator and corresponds to a running total of past contributions of monthly realizations within the current quarter. The remaining entries of this matrix correspond to ones and zeros to preserve the lag structure. The VAR intercepts sit at the top of C_t, and scaled versions of intercepts are in rows associated with each accumulator. The rest of C_t has zeros. Finally, each R_t corresponds to the natural selection matrix, using the same deterministic series of fractions used in T_t augmented to accommodate the additional accumulator variables in the state.[6]

In periods in which all of the indicators are observed, the selection matrix Z_t is composed solely of n selection rows made up of zeros and ones. Specifically, for these periods the Z_t matrix is given by

$$Z_t = \begin{bmatrix} 0 & 0 & \ldots & I_{n_q} \\ 0 & I_{n_m} & \ldots & 0 \end{bmatrix},$$

where the identity matrix in the first n_q rows of Z_t corresponds to the mapping of the accumulators to the quarterly variables and where the identity matrix in the last n_m rows of Z_t corresponds to the mapping of the monthly (base-frequency) time series to their observed counterparts in y_t.

The row dimension of Z_t varies over time because of the changing dimensionality of the observables. For the months in which only monthly time series are observed, the last n_m rows of Z_t will be included. Furthermore, toward the end of the sample, not all of the monthly indicators will be available, depending on their release schedule, and a further subset of these last n_m rows will be used.

1.2 Gibbs sampling procedure

With the model cast in a state-space framework, we can now estimate the full set of parameters and latent states given by $\Theta = \{\phi, c, \Sigma, \{x_t^{latent}, \zeta_t\}_{t=1}^T\}$. With the history of data in the estimation sample through time $t \leq T$ denoted as $Y_{1:t}$, inference on Θ concerns the VAR parameters $\{\phi, c, \Sigma\}$, the latent monthly variables $\{x_t^{latent}\}_{t=1}^T$ (of the quarterly time

[6]For further details on the construction of accumulator variables, see the classic examples in Harvey (1989).

series as well as any missing monthly variables), and the accumulators $\{\zeta\}_{t=1}^{T}$ conditional on $Y_{1:T}$. To conduct inference, Schorfheide and Song (2015) propose a two-block Gibbs sampler that, conditional on a presample $Y_{-p+1:0}$ used to initialize the lags, generates draws from the conditional posterior distributions:

$$P(\phi, c, \Sigma | X_{1:T}, Y_{-p+1:T}) \tag{A5}$$

and

$$P(X_{1:T} | Y_{-p+1:T}, \phi, c, \Sigma), \tag{A6}$$

where we stack $\{x_t^{latent}, \zeta_t\}_{t=1}^{T}$ into the matrix $X_{1:T}$.

The first density, given in equation (A5), is the posterior of the VAR parameters conditional on all data and the latent variables. With a suitable choice of priors, sampling from this distribution is reduced to taking a draw from a straightforward multivariate regression. The second density, given in equation (A6), corresponds to the Kalman smoothed estimates of the latent variables. A draw from this distribution is obtained via the simulation smoother of Durbin and Koopman (2012). Hence, the estimation of ALEX iterates between taking draws from these two conditional posterior distributions.

2 Implementing the Minnesota prior

The use of informed priors (through Bayesian methods) has been shown to improve the forecasting performance of high-dimensional VARs (Litterman, 1986). We appeal to this tradition, and use a modified version of the traditional Minnesota prior when estimating ALEX. An important feature of the Minnesota implementation in our case is that the informed prior over the large dimensional parameter vector is characterized by a very small number of hyperparameters that are readily interpretable (Doan, Litterman, and Sims, 1984). Additionally, a well-developed literature has established values for these hyperparameters that can typically be taken "off the shelf" in practice (e.g., Bańbura, Giannone, and Reichlin, 2010; Carriero, Clark, and Marcellino, 2015; and Giannone, Lenza, and Primiceri, 2015).

In setting the hyperparameters of ALEX, we elect to take an empirical Bayesian approach rather than calibrating them to their traditional values. Taking such an approach requires the maximization (with respect to the hyperparameters) of the marginal data density (MDD), which has been shown to lead to superior out-of-sample forecasts (Geweke, 2001) and in some settings has an analytical form that often further facilitates estimation (Del Negro and Schorfheide, 2011).[7]

2.1 Shrinkage through dummy observations

We address the curse of dimensionality of the VAR parameters (ϕ, Σ) by using the following informative prior distributions on (ϕ, Σ). Generally speaking, the priors we use combine a

[7] In the mixed-frequency setting, a modification is required that utilizes the output from the Gibbs sampler and the modified harmonic mean—see, e.g., Schorfheide and Song (2015) and Brave, Butters, and Justiniano (2019).

slightly modified version of the well-known Minnesota prior (Litterman, 1986) with a set of priors that guide the sum of autoregressive lags as well as the co-persistence of the variables in the model. In particular, the modified Minnesota prior that we use belongs to the normal-inverse Wishart family, to preserve conjugacy, and takes the following form:

$$\Sigma \sim IW(\psi, d),$$

$$\phi|\Sigma \sim N(\Gamma, \Sigma \otimes \Omega).$$

We follow the convention of the Minnesota prior for single-frequency BVARs, so the matrix Γ consists solely of zeros and ones, which has the effect of shrinking the VAR system toward independent random walks or white noise. A small set of hyperparameters (λ_0, λ_1, λ_2, λ_3, and λ_4) then contributes to the characterization of the covariance matrix Ω, and is collected in the hyperparameter vector λ^p. Regarding the prior for Σ, the hyperparameter matrix ψ is assumed to be diagonal, and the degrees of freedom hyperparameter d is chosen such that the prior for Σ is centered at ψ/n, where n is the total number of series in ALEX and ψ_j is the corresponding entry on the main diagonal of ψ. We collect these ψ_j entries into the other component of the hyperparameters, λ^Σ, for which we use empirical Bayesian methods to select.

The hyperparameter λ_0 controls the precision of the prior on the VAR intercepts, or *constant*, c.[8] The overall tightness of the prior is controlled by λ_1, and is subsequently referred to as the *tightness*. As $\lambda_1 \to \infty$, the posterior distribution is dominated by the prior; conversely, as $\lambda_1 \to 0$, the posterior coincides with the ordinary least squares (OLS) estimates of the VAR. The second element of the prior is λ_2, the *decay* hyperparameter, which governs the rate at which coefficients at distant lags are shrunk further toward zero.

Forecast performance has been shown to improve with two additional priors concerning the persistence and co-persistence of the variables in the VAR. These additional priors are designed to prevent initial transients and deterministic components from explaining an implausible share of the long-run variability in the system (Sims and Zha, 1998; and Sims, 2000). The first form of shrinkage is usually known as the *sum of coefficients* prior, and expresses the belief that the sum of own-lag autoregressive coefficients for each individual variable should be one. This is governed by λ_3, with larger values implying (as shown earlier) a tighter prior. The second form of shrinkage is known as the *co-persistence* prior and reflects the belief that if the sum of all VAR coefficients is close to an identity matrix, then the intercepts should be small (or conversely, if the VAR is stationary, then intercepts are not close to zero). The strength with which this prior is imposed is increasing in λ_4. In both cases, a hyperparameter set to zero corresponds to the exclusion of that prior from the system, while approaching infinity corresponds to a system that strictly adheres to the prior.

To operationalize these priors, we use the data augmentation approach often used in the BVAR context (e.g., Bańbura, Giannone, and Reichlin, 2010; and Schorfheide and Song, 2015). The set of dummy observations that implements the forms of shrinkage we consider is given by

[8] Bańbura, Giannone, and Reichlin (2010) set this value to a very small number. Carriero, Clark, and Marcellino (2015) set this hyperparameter to 1. We take a more agnostic stance, instead putting a hyperprior on this hyperparameter and letting the optimization of the MDD include λ_0.

$$
Y_d = \begin{pmatrix} \lambda_1 \mathrm{diag}(\psi_1\delta_1, \ldots, \psi_n\delta_n) \\ 0_{n(p-1)\times n} \\ \cdots \\ \mathrm{diag}(\psi_1, \ldots, \psi_n) \\ \cdots \\ 0_{1\times n} \\ \cdots \\ \lambda_3 \mathrm{diag}(\delta_1\bar{y}_1, \ldots, \delta_n\bar{y}_n) \\ \cdots \\ \lambda_4 \bar{y} \end{pmatrix}, \quad X_d = \begin{pmatrix} \lambda_1 J_p \otimes \mathrm{diag}(\psi_1, \ldots, \psi_n) & 0_{np\times 1} \\ \cdots & \cdots \\ 0_{n\times np} & 0_{n\times 1} \\ \cdots & \cdots \\ 0_{1\times np} & \lambda_0 \\ \cdots & \cdots \\ (1_{1\times p}) \otimes \lambda_3 \mathrm{diag}(\delta_1\bar{y}_1, \ldots, \delta_n\bar{y}_n) & 0_{n\times 1} \\ \cdots & \cdots \\ (1_{1\times p}) \otimes \lambda_4 \bar{y} & \lambda_4 \end{pmatrix}. \tag{A7}
$$

The first block corresponds to the *tightness* and *decay* components of the prior governed by λ_1 and λ_2, respectively; and where $J_p = \mathrm{diag}(1^{\lambda_2}, \ldots, p^{\lambda_2})$, $\bar{y}$ is an n-dimensional vector of sample means, and the n-dimensional vector $\overline{\psi} = (\psi_1, \ldots, \psi_n)'$ has as its i-th element a nonnegative number that is proportional to the residual variance for that series. The series-specific scalars δ_i reflect the centers of the prior for the first-order own-lag autoregressive coefficients and are set to 1. The second block ensures the prior for the residual variances is appropriately centered, and the third block represents the prior for the intercepts with λ_0. The *sum of coefficients* component of the prior is governed by λ_3, where once again the series-specific scalars δ_i correspond to the centers of the prior for the first autoregressive coefficients and are set to 1.[9] Finally, λ_4 controls beliefs regarding the *co-persistence* of the system.

Regarding the covariance matrix, Σ, it can be shown that these dummy observations combined with an improper prior of the form $P(\Sigma) \propto |\Sigma|^{-\frac{n+1}{2}}$ imply an inverse Wishart (IW) prior density centered at $\frac{\psi}{n}$, where the diagonal matrix ψ is fully characterized by the vector $\overline{\psi}$ described previously. This is because the degrees of freedom resulting from multiplying that prior and the "likelihood" of the dummy observations is given by $T^* - k$, where T^* is the number of dummy observations and is equal to $k+n+1+n$ and $k = np+1$. As such, the mean of the IW is equal to

$$
\frac{\psi}{T^* - k - n - 1} = \frac{\psi}{n}.
$$

This suggests caution should be exercised if one wishes to center the prior of the residual variances at presample estimates, which must be scaled by n when using the full complement of priors. It is also noteworthy that if not all four priors are active, then the mean of the implied inverse Wishart prior may not even be well defined.

2.2 Selecting hyperparameters and constructing forecasts

As noted earlier, we take an empirical Bayesian approach to selecting the hyperparameters $(\lambda^p, \lambda^\Sigma)$. Our approach to accomplishing this is to first estimate a bivariate MF-BVAR

[9] Although it is not written this way here, one can separate the centers for the first autoregressive coefficient and the sum of all autoregressive coefficients. We make use of this fact in ALEX for some series where it is more reasonable to center the former at 0 and the latter at 1.

with a related monthly variable that is not included in ALEX for each quarterly variable.[10] Treating the posterior high-frequency estimates of the quarterly variables as well as the other monthly variables of the model as data, we then proceed to optimize the MDD of this generated data set—which is known in closed form, as shown by Del Negro and Schorfheide (2011)—using numerical methods. While feasible to implement in real time, our approach is still more computationally demanding than the common practice of fixing the full set of required hyperparameters to prespecified values. It is, however, more practical than the two-step approach considered in Brave, Butters, and Justiniano (2019), which uses a selection of the optimal hyperparameters from this procedure to characterize the contours that surround them so that an informed grid can be set up to maximize the true MDD in a second step.[11] That said, as in Brave, Butters, and Justiniano (2019), we facilitate the elicitation of hyperparameters by imparting a set of hyperpriors (Giannone, Lenza, and Primiceri, 2015). Table A1 displays the values of the hyperparameters selected by this procedure and used in the estimation of ALEX.

For each iteration of the Gibbs sampler during estimation, forecasts for GDP are generated recursively for the current month up to one year into the future. We report the median of these samples as a *point forecast* and assess uncertainty with the associated *density forecast* by constructing 90% and 70% coverage intervals.[12] This procedure captures three forms of model uncertainty: parameter, state, and shock uncertainty. Parameter uncertainty concerns the estimation of the model's coefficients on past values of its variables. State uncertainty addresses the mixed-frequency nature of the model, incorporating the fact that for quarterly variables we do not directly observe their monthly evolution and must infer it from the model. Shock uncertainty simply deals with the stochastic nature of the model and the way we draw shocks from the model's implied Normal distributions to simulate over the forecast horizon events that are not captured by the model's historical dynamics. Figure A1 shows examples of point and density forecasts for the level of U.S. real GDP using data available on October 20, 2020. For comparison, the figure also shows the consensus forecast and averages of the top ten and bottom ten forecasts from the October 9, 2020, *Blue Chip Economic Indicators* survey. Addressing model uncertainty in ALEX in this way produces very similar results to the level of forecaster disagreement in the *Blue Chip* survey, as shown in the figure.

[10]Each bivariate MF-BVAR uses the same lag order (three) as ALEX and also proper, standard priors.

[11]Brave, Butters, and Justiniano (2019) use the informed grid found in the first step to maximize the MDD according to the modified harmonic mean with respect to the hyperparameter elements within λ^p, whereas the hyperparameters with λ^Σ are held constant from this first step. Given the size of ALEX, we skip this second step where it quickly becomes computationally infeasible to construct the appropriate modified harmonic mean.

[12]Given that the data series in ALEX are approximately normally distributed, the mean and median projections of the model are roughly identical.

References

Bańbura, M., D. Giannone, and L. Reichlin, 2010, "Large Bayesian vector auto regressions," *Journal of Applied Econometrics*, Vol. 25, No. 1, January/February, pp. 71–92. Crossref, https://doi.org/10.1002/jae.1137

Brave, S. A., R. A. Butters, and A. Justiniano, 2019, "Forecasting economic activity with mixed frequency BVARs," *International Journal of Forecasting*, Vol. 35, No. 4, October–December, pp. 1692–1707. Crossref, https://doi.org/10.1016/j.ijforecast.2019.02.010

Carriero, A., T. E. Clark, and M. Marcellino, 2015, "Bayesian VARs: Specification choices and forecast accuracy," *Journal of Applied Econometrics*, Vol. 30, No. 1, January/February, pp. 46–73. Crossref, https://doi.org/10.1002/jae.2315

Del Negro, M., and F. Schorfheide, 2011, "Bayesian macroeconometrics," in *The Oxford Handbook of Bayesian Econometrics*, J. Geweke, G. Koop, and H. van Dijk (eds.), Oxford, UK: Oxford University Press, pp. 293–389. Crossref, https://doi.org/10.1093/oxfordhb/9780199559084.013.0008

Doan, T., R. Litterman, and C. Sims, 1984, "Forecasting and conditional projection using realistic prior distributions," *Econometric Reviews,* Vol. 3, No. 1, pp. 1–100. Crossref, https://doi.org/10.1080/07474938408800053

Durbin, J., and S. J. Koopman, 2012, *Time Series Analysis by State Space Methods*, 2nd ed., Oxford Statistical Science Series, Vol. 38, Oxford, UK: Oxford University Press. Crossref, https://doi.org/10.1093/acprof:oso/9780199641178.001.0001

Geweke, J., 2001, "Bayesian econometrics and forecasting," *Journal of Econometrics*, Vol. 100, No. 1, January, pp. 11–15. Crossref, https://doi.org/10.1016/S0304-4076(00)00046-4

Giannone, D., M. Lenza, and G. E. Primiceri, 2015, "Prior selection for vector autoregressions," *Review of Economics and Statistics*, Vol. 97, No. 2, May, pp. 436–451. Crossref, https://doi.org/10.1162/REST_a_00483

Harvey, A. C., 1989, *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge, UK: Cambridge University Press.

Litterman, R. B., 1986, "Forecasting with Bayesian vector autoregressions—Five years of experience," *Journal of Business & Economic Statistics,* Vol. 4, No. 1, pp. 25–38. Crossref, https://doi.org/10.1080/07350015.1986.10509491

Mariano, R. S., and Y. Murasawa, 2003, "A new coincident index of business cycles based on monthly and quarterly series," *Journal of Applied Econometrics,* Vol. 18, No. 4, July/August, pp. 427–443. Crossref, https://doi.org/10.1002/jae.695

Schorfheide, F., and D. Song, 2015, "Real-time forecasting with a mixed-frequency VAR," *Journal of Business & Economic Statistics,* Vol. 33, No. 3, pp. 366–380.
Crossref, `https://doi.org/10.1080/07350015.2014.954707`

Sims, C. A., 2000, "Using a likelihood perspective to sharpen econometric discourse: Three examples," *Journal of Econometrics,* Vol. 95, No. 2, April, pp. 443–462.
Crossref, `https://doi.org/10.1016/S0304-4076(99)00046-9`

Sims, C. A., and T. Zha, 1998, "Bayesian methods for dynamic multivariate models," *International Economic Review*, Vol. 39, No. 4, November, pp. 949–968.
Crossref, `https://doi.org/10.2307/2527347`

Table A1. Posterior estimates of ALEX hyperparameters

Tightness (λ_1)	Decay (λ_2)	Sum of coefficients (λ_3)	Co-persistence (λ_4)
1.0299	1	24.4934	6.9129

	Series description	Transformation	Frequency	Standard deviation of intercept (ψ_j/λ_0)
	National Income and Product Accounts			
1	Real Gross Domestic Product	LN	Q	0.0004
2	Real Gross Domestic Income	LN	Q	0.0004
3	Real Personal Consumption Expenditures	LN	M	0.0005
4	Real Business Fixed Investment	LN	Q	0.0009
5	Real Residential Investment	LN	Q	0.0032
6	Real Government Consumption Expenditures and Gross Investment	LN	Q	0.0006
7	Real Exports of Goods and Services	LN	Q	0.0014
8	Real Imports of Goods and Services	LN	Q	0.0018
9	Real Private Inventories	LN	Q	0.0002
	Employment, Unemployment and Hours			
10	Total Nonfarm Payroll Employment	LN	M	0.0003
11	Civilian Employment-to-Population Ratio	LV	M	0.0002
12	Labor Force Participation Rate	LV	M	0.0002
13	Civilian Unemployment Rate	LV	M	0.0003
14	Median Duration of Unemployment (Weeks)	LN	M	0.0114
15	Initial Unemployment Insurance Claims (Mils.)	LN	M	0.0055
16	Continuing Unemployment Insurance Claims (Mils.)	LN	M	0.0034
17	Aggregate Weekly Hours of Production and Nonsupervisory Employees on Nonfarm Payrolls	LN	M	0.0006
	Industrial Production (IP) and Capacity Utilization			
18	Total Industry Capacity Utilization	LV	M	0.0004
19	Real Gross Value of Business Equipment	LN	M	0.0019
20	Real Gross Value of Consumer Goods	LN	M	0.0011
21	Real Gross Value of Nonindustrial Supplies	LN	M	0.0009
22	Real Gross Value of Defense and Space Equipment	LN	M	0.003
23	IP: Consumer Energy Products	LN	M	0.0032
24	IP: Electric and Gas Utilities	LN	M	0.0026
25	IP: Computers, Communications Equipment, and Semiconductors	LN	M	0.0024
26	IP: Nonenergy Materials	LN	M	0.0009
27	IP: Motor Vehicle Assemblies	LN	M	0.0101
28	IP: Mining	LN	M	0.0025
29	IP: Oil and Gas Well Drilling	LN	M	0.007
30	Baker Hughes Active Rig Count (Thous.)	LN	M	0.0052
	Sales, Orders, and Inventories			
31	Real Manufacturing and Trade Sales	LN	M	0.0009
32	Real Personal Consumption Expenditures (PCE) Control Group Retail Sales	LN	M	0.0008
33	Real Retail Sales excluding Autos, Building Supplies, and Gas Stations	LN	M	0.0007
34	Real Motor Vehicle Sales	LN	M	0.0047
35	Light Vehicle Sales	LN	M	0.0057
36	Medium and Heavy Truck Sales	LN	M	0.0121
37	Consumer Dollars as a Percent of Auto Sales	LV	M	0.0026
38	Business New Light Vehicle Purchases	LN	M	0.0076
39	Real Retail Sales of Building Materials, Garden Equipment and Supply Dealers	LN	M	0.0037
40	Real Retail Sales of Furniture and Household Appliances	LN	M	0.0027
41	Real Food Service Sales	LN	M	0.0021
42	Real Manufacturers' New Orders of Consumer Goods	LN	M	0.0014
43	Real Manufacturers' Unfilled Orders of Consumer Goods	LN	M	0.0058
44	Real Manufacturers' Shipments of Consumer Goods	LN	M	0.0013
45	Real Manufacturers' New Orders of Construction Materials and Supplies	LN	M	0.0024
46	Real Manufacturers' Unfilled Orders of Construction Materials and Supplies	LN	M	0.0015
47	Real Manufacturers' Shipments of Construction Materials and Supplies	LN	M	0.0021
48	Real Manufacturers' New Orders of Information Technology	LN	M	0.0075
49	Real Manufacturers' Unfilled Orders of Information Technology	LN	M	0.0025
50	Real Manufacturers' Shipments of Information Technology	LN	M	0.004
51	Real Manufacturers' New Orders of Nondefense Capital Goods	LN	M	0.009
52	Real Manufacturers' Unfilled Orders of Nondefense Capital Goods	LN	M	0.0012
53	Real Manufacturers' Shipments of Nondefense Capital Goods	LN	M	0.0025
54	Real Manufacturers' New Orders of Nondefense Capital Goods excluding Aircraft	LN	M	0.0038
55	Real Manufacturers' Unfilled Orders of Nondefense Capital Goods excluding Aircraft	LN	M	0.0009
56	Real Manufacturers' Shipments of Nondefense Capital Goods excluding Aircraft	LN	M	0.0022
57	Real Manufacturers' New Orders of Defense Capital Goods	LN	M	0.0366
58	Real Manufacturers' Unfilled Orders Defense Capital Goods	LN	M	0.0021
59	Real Manufacturers' Shipments of Defense Capital Goods	LN	M	0.0086
60	Real Manufacturing and Trade Inventories	LN	M	0.0005
61	Total Business Inventories-to-Sales Ratio	LV	M	0.002
62	Ward's Domestic Auto Inventories	LN	M	0.0051

LN: Natural Logarithm, LV: Level, M: Monthly, Q: Quarterly

	Series description	Transformation	Frequency	Standard deviation of intercept (ψ_j/λ_0)
	Private Construction and Real Estate			
63	Real Private Nonresidential Construction Spending	LN	M	0.0035
64	Real Private Business Construction Spending	LN	M	0.0045
65	Dodge Report Construction Contracts Square Footage	LN	M	0.0193
66	Completed Housing Units	LN	M	0.0128
67	Real Private New Residential Construction Spending	LN	M	0.0025
68	Single-family Housing Starts	LN	M	0.0107
69	Multi-family Housing Starts	LN	M	0.0342
70	Housing Permits	LN	M	0.0102
71	Shipments of Manufactured Homes	LN	M	0.0094
72	Real Payrolls of Residential Remodelers	LN	M	0.0038
73	Real Value of Home Sales	LN	M	0.0059
74	Months' Supply of Existing Single-family Homes	LV	M	0.0008
75	Pending Home Sales	LN	M	0.0049
	International Trade and Government Spending			
76	Real Balance of Trade in Goods (Census)	LV	M	0.0042
77	Real Exports of Goods (Balance of Payments, or BOP)	LN	M	0.0032
78	Real Exports of Services (BOP)	LN	M	0.0032
79	Real Imports of Goods (BOP)	LN	M	0.0033
80	Real Imports of Services (BOP)	LN	M	0.0028
81	Real Imports of Capital Goods excluding Autos	LN	M	0.0048
82	Real Imports of Motor Vehicles and Parts	LN	M	0.0104
83	Real Imports of Industrial Supplies and Materials	LN	M	0.0077
84	Real Imports of Energy-related Petroleum Products	LN	M	0.0163
85	Real Exports of Agricultural Products	LN	M	0.0107
86	Trade-weighted Exchange Value of the Dollar: Broad Index	LN	M	0.0025
87	Real Federal Outlays from Monthly Treasury Statement	LN	M	0.0155
88	Real Wage and Salaries of Government Employees	LN	M	0.0006
89	Real State and Local Public Construction Spending	LN	M	0.0043
90	Real Federal Public Construction Spending	LN	M	0.0128
	Household and Business Surveys			
91	Number of Households	LN	M	0.0005
92	University of Michigan Survey of Consumers Expectations Index	LV	M	0.0102
93	Institute for Supply Management (ISM) Purchasing Manager's Composite Index	LV	M	0.0025
94	ISM Manufacturing Production Index	LV	M	0.0054
95	ISM Manufacturing New Orders Index	LV	M	0.0055
96	ISM Manufacturing Supplier Deliveries Index	LV	M	0.005
97	ISM Manufacturing Inventories Index	LV	M	0.0063
98	ISM Manufacturing New Export Orders Index	LV	M	0.0049
99	ISM Manufacturing Imports Index	LV	M	0.0045
100	Federal Reserve Bank of Philadelphia Manufacturing Index (ISM basis)	LV	M	0.0054
	Household and Business Balance Sheets			
101	Real Personal Income Less Transfer Payments	LN	M	0.001
102	Real Average Hourly Earnings of Production and Nonsupervisory Employees on Nonfarm Payrolls	LN	M	0.0006
103	Personal Savings Rate	LV	M	0.0012
104	Real Household Net Worth	LN	Q	0.0014
105	Real Private Software Industry Revenue	LN	Q	0.0011
106	Real Healthcare Industry Revenue	LN	Q	0.0012
107	Real Corporate Net Cash Flow	LN	Q	0.0039

LN: Natural Logarithm, LV: Level, M: Monthly, Q: Quarterly

Figure A1. Point and density forecasts for real gross domestic product

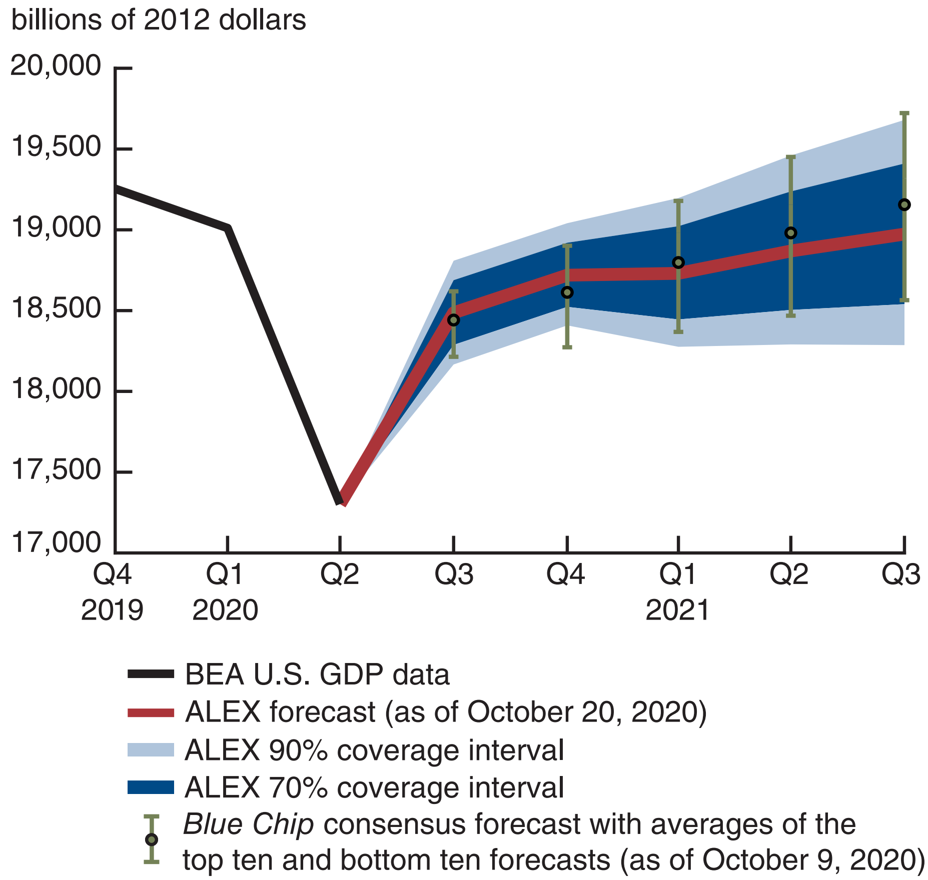

Note: See the text for further details on ALEX and its coverage intervals.
Sources: Authors' calculations; and U.S. Bureau of Economic Analysis (BEA) and Wolters Kluwer's *Blue Chip Economic Indicators* data from Haver Analytics.