Advances in Nowcasting Economic Activity

Juan Antolín-Díaz[1]   Thomas Drechsel[2]   Ivan Petrella[3]

XIII Annual Conference on Real-Time Data Analysis
Banco de España

October 19, 2017

[1]Fulcrum Asset Management
[2]LSE and Centre for Macroeconomics
[3]Warwick Business School and CEPR

This paper

- This paper contributes to the literature on nowcasting economic activity.
- We propose a bayesian dynamic factor model (DFM), which takes seriously the features of the data:
  1. Low-frequency variation in the mean and variance
  2. Heterogeneous responses to common shocks (lead-lags)
  3. Outlier observations and fat tails
  4. Endogenous modeling of seasonality (today as an extension)
- We evaluate performance of the model and new features in a comprehensive out of sample real-time exercise, paying particular attention to density forecasts.
- The project builds on our earlier work: Antolin-Diaz, Drechsel, and Petrella (2017 *ReStat*)

Preview of Results

- The results we show today are for the US
- The proposed improvements are successful at capturing important features of the data
- The real-time nowcasting performance is substantially improved across a variety of metrics (point and density)
  - Capturing trends and SV improves nowcasting performance significantly
  - Heterogeneous dynamics deliver substantial additional improvement
  - Fat tails successfully capture outlier observations in an automated way and improve density forecasts of the monthly variables (*these results are preliminary!*).

Plan of the talk

1. The Model
   - Highlight features of the data
   - Lay out how these are formally captured in the DFM
2. Estimation
   - Brief overview of data and algorithm
3. Setup of Real-Time Evaluation Exercise
4. Evaluation Results
5. Extension: Residual Seasonality
6. Conclusion

The Model

We start from the familiar specification of a DFM
(see, e.g. Giannone, Reichlin, and Small, 2008)

Consider an n-dimensional vector of observables $\mathbf{y}_t$, which follows

$$\begin{aligned} \Delta(\mathbf{y}_t) &= \mathbf{c} + \boldsymbol{\lambda}\mathbf{f}_t + \mathbf{u}_t && (1) \\ (I - \Phi(L))\mathbf{f}_t &= \varepsilon_t && (2) \\ (1 - \rho_i(L))u_{i,t} &= \eta_{i,t}, \qquad i = 1, \dots, n && (3) \end{aligned}$$

$$\begin{aligned} \varepsilon_t &\overset{iid}{\sim} N(0, \boldsymbol{\Sigma}_\varepsilon) && (4) \\ \eta_{i,t} &\overset{iid}{\sim} N(0, \sigma^2_{\eta_i}), \qquad i = 1, \dots, n && (5) \end{aligned}$$

The Model

Why explicitly model low frequency variation?

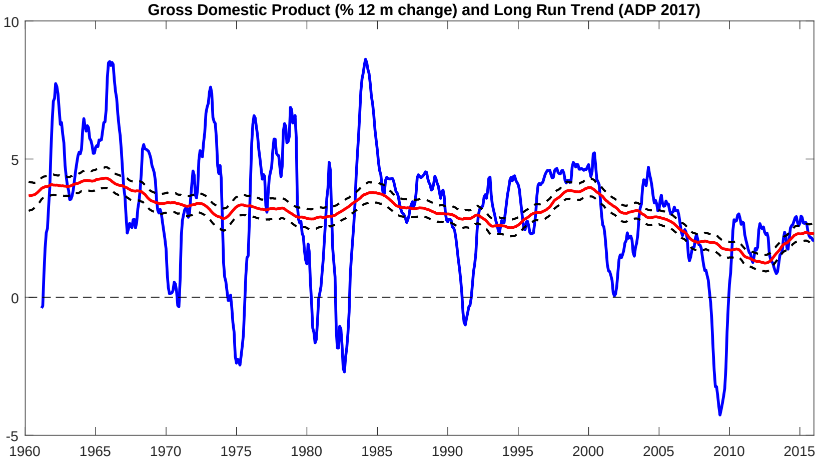

Consider n-dimensional vector of observables $\mathbf{y}_t$, which follows

$$\Delta(\mathbf{y}_t) = \mathbf{c}_t + \boldsymbol{\lambda}\mathbf{f}_t + \mathbf{u}_t, \tag{6}$$

with

$$\mathbf{c}_t = \begin{bmatrix} \mathbf{B} & \mathbf{0} \\ \mathbf{0} & \mathbf{c} \end{bmatrix} \begin{bmatrix} \mathbf{a}_t \\ 1 \end{bmatrix}, \tag{7}$$

and

$$\begin{aligned} (I - \Phi(L))\mathbf{f}_t &= \boldsymbol{\varepsilon}_t, & & (8) \\ (1 - \rho_i(L))u_{i,t} &= \eta_{i,t}, \quad i = 1, \ldots, n & & (9) \end{aligned}$$

The Model

Why model changes in volatility?

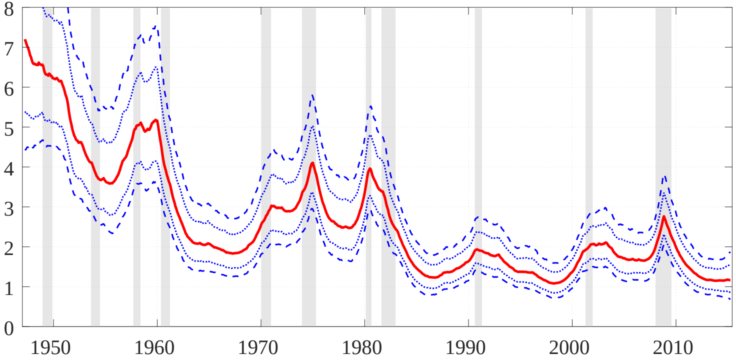

- SV in the common factor captures both secular (McConnell and Perez-Quiros, 2000) and cyclical (e.g., Jurado et al., 2014) movements in volatility.

Consider n-dimensional vector of observables $\mathbf{y}_t$, which follows

$$\Delta(\mathbf{y}_t) = \mathbf{c}_t + \boldsymbol{\lambda}\mathbf{f}_t + \mathbf{u}_t, \tag{10}$$

with

$$\mathbf{c}_t = \begin{bmatrix} \mathbf{B} & \mathbf{0} \\ \mathbf{0} & \mathbf{c} \end{bmatrix} \begin{bmatrix} \mathbf{a}_t \\ 1 \end{bmatrix}, \tag{11}$$

and

$$\begin{aligned} (I - \Phi(L))\mathbf{f}_t &= \sigma_{\varepsilon_t}\boldsymbol{\varepsilon}_t, & & (12) \\ (1 - \rho_i(L))u_{i,t} &= \sigma_{\eta_{i,t}}\eta_{i,t}, \quad i = 1, \ldots, n & & (13) \end{aligned}$$

where the time-varying parameters will be specified as a random walk processes.

The Model

Why allow for heterogeneous dynamics?

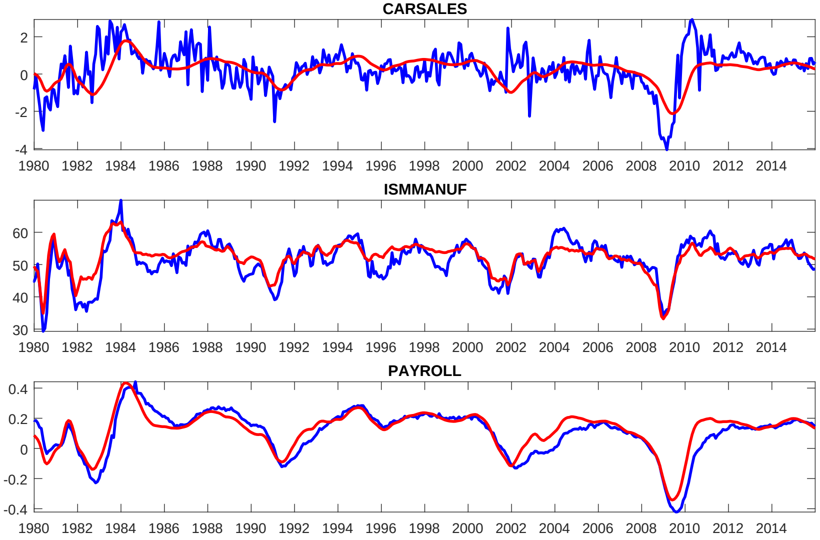

$$\Delta(\mathbf{y}_t) = \mathbf{c}_t + \mathbf{\Lambda}(\mathbf{L})\mathbf{f}_t + \mathbf{u}_t, \tag{14}$$

where $\mathbf{\Lambda}(\mathbf{L})$ contains the loadings on the contemporaneous and lagged common factors.

- Camacho and Perez-Quiros (2010) first noticed that survey data was better aligned with a distributed lag of GDP.
- D'Agostino et al. (2015) show that adding lags improves performance in the context of a small model.

The Model

What do the heterogeneous dynamics achieve?

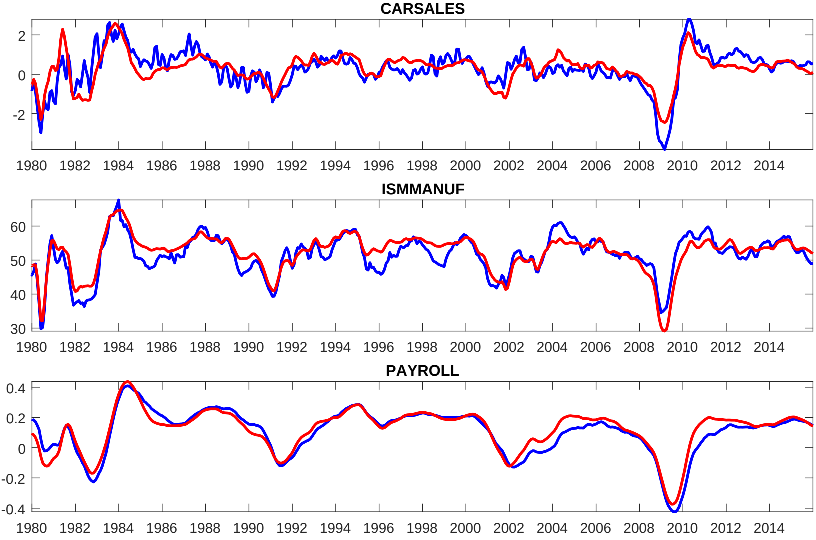

The model

What do the heterogeneous dynamics achieve?

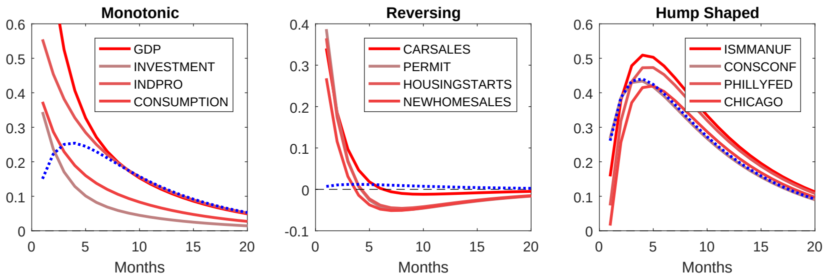

Result: There is substantial heterogeneity in the impulse responses.

The Model

Why model fat tails?

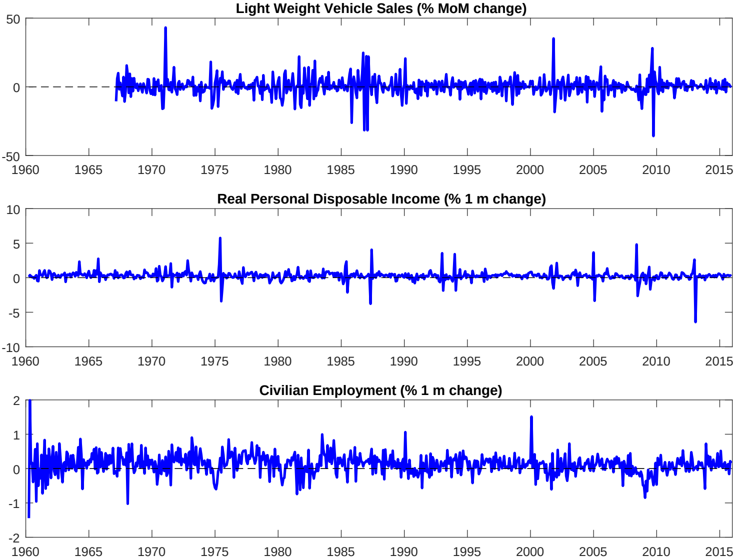

$$\Delta(\mathbf{y}_t - \mathbf{o}_t) = \mathbf{c}_t + \mathbf{\Lambda}(\mathbf{L})\mathbf{f}_t + \mathbf{u}_t, \tag{15}$$

with $\mathbf{o}_t$ as well as the innovations to $\mathbf{u}_t$ following a t-distribution.

The laws of motion of the various components are specified as

$$
\begin{aligned}
(I-\Phi(L))\mathbf{f}_t &= \sigma_{\varepsilon_t}\boldsymbol{\varepsilon}_t, && && (16)\\
(1-\rho_i(L))u_{i,t} &= \sigma_{\eta_{i,t}}\eta_{i,t}, && i=1,\ldots,n && (17)\\
\eta_{i,t} &\overset{iid}{\sim} t(v_{\eta,i}), && i=1,\ldots,n && (18)\\
o_{i,t} &\overset{iid}{\sim} t(v_{o,i}), && i=1,\ldots,n && (19)\\
\varepsilon_t &\overset{iid}{\sim} N(0,\mathbf{I}) && && (20)
\end{aligned}
$$

The degrees of freedom of the t-distributions are estimated jointly with the other parameters of the model.

The model's time-varying parameters are specified to follow driftless random walks:

$$
\begin{aligned}
a_{j,t} &= a_{j,t-1} + v_{a_{j,t}}, & v_{a_{j,t}} &\overset{iid}{\sim} N(0, \omega^2_{a,j}) \;\; j = 1, \dots, r \\
\log \sigma_{\varepsilon_t} &= \log \sigma_{\varepsilon_{t-1}} + v_{\varepsilon,t}, & v_{\varepsilon,t} &\overset{iid}{\sim} N(0, \omega^2_{\varepsilon}) \\
\log \sigma_{\eta_{i,t}} &= \log \sigma_{\eta_{i,t-1}} + v_{\eta_{i,t}}, & v_{\eta_{i,t}} &\overset{iid}{\sim} N(0, \omega^2_{\eta,i}) \;\; i = 1, \dots, n
\end{aligned}
$$

Banbura and Modugno (2014) show in a Gaussian model that the impact of a new release on the nowcasts can be written as a linear function of the *news*:

$$E(y_{k,tk}|\Omega_2) - E(y_{k,tk}|\Omega_1) = \mathbf{w_j}\left(y_{j,tj} - E(y_{j,tj}|\Omega)\right)$$

$$\mathbf{w_j} = \frac{\Lambda_k E\left((f_{t_k} - f_{t_k|\Omega})(f_{t_j} - f_{t_j\Omega})\right)\Lambda_j'}{\Lambda_j E\left((f_{t_j} - f_{t_j|\Omega})(f_{t_j} - f_{t_j\Omega})\right)\Lambda_j' + \sigma^2_{\eta_{j,tj}}}$$

The model

News Decompositions

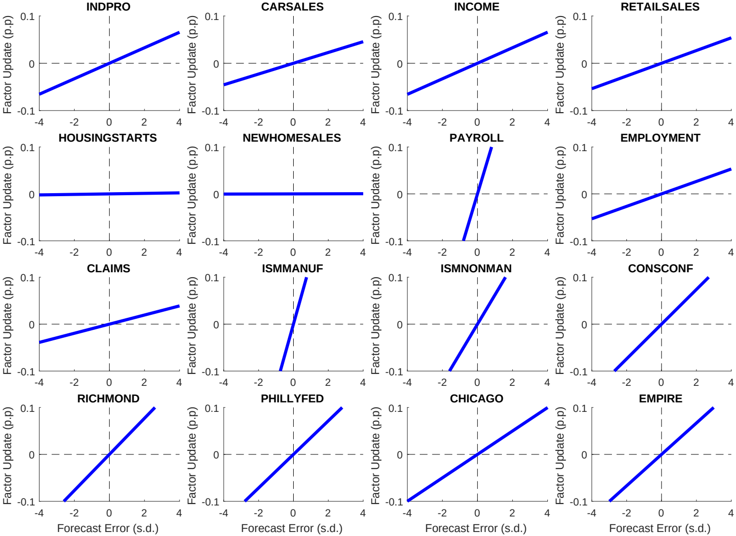

We show that with the Student-t distribution the weights are no longer linear, but depend on the value of the forecast error itself:

$$E(y_{k,tk}|\Omega_2) - E(y_{k,tk}|\Omega_1) = \mathbf{w_j}(y_{j,tj})\,(y_{j,tj} - E(y_{j,tj}|\Omega))$$

$$\mathbf{w_j}(y_{j,tj}) = \frac{\Lambda_k E\Big((f_{t_k} - f_{t_k|\Omega})(f_{t_j} - f_{t_j\Omega})\Big)\Lambda_j'}{\Lambda_j E\Big((f_{t_j} - f_{t_j|\Omega})(f_{t_j} - f_{t_j\Omega})\Big)\Lambda_j' + \sigma^2_{\eta_{j,tj}}\delta_{j,tj}}$$

$$\delta_{j,tj} = (((y_{j,tj} - E(y_{j,tj}|\Omega))^2/\sigma^2_{\eta_{j,tj}} + v_{o,j})/(v_{o,i} + 1)$$

Large errors are discounted as outlier observations containing less information.

The model

What do the fat tails achieve?

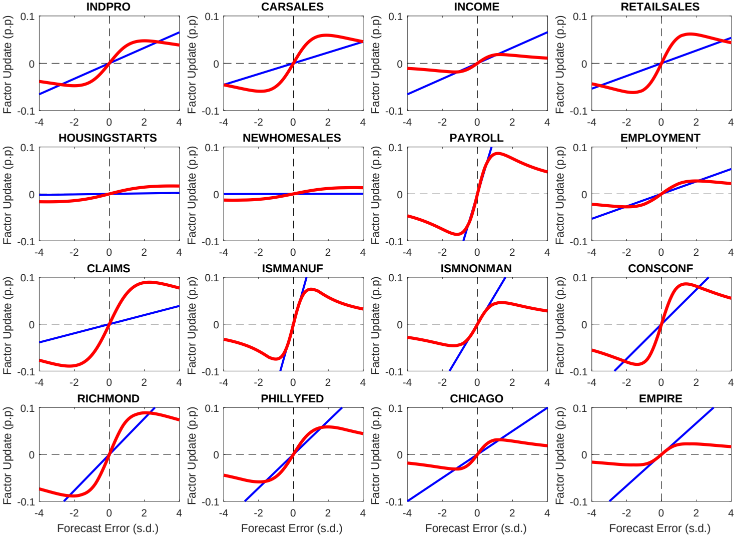

Estimation

Estimation

Data set

	Frequency
Hard Indicators	
GDP (Chained $)	Q
GDI (Chained $)	Q
Consumption (Chained $, Non Dur. + Serv.)	Q
Investment (Chained $, Fixed + Cons. Dur.)	Q
Aggregate Hours Worked (Total Economy)	Q
Industrial Production	M
Payroll Employment (Establishment Survey)	M
Real Retail Sales Food Services	M
Real Personal Income less Transfer Payments	M
New Orders of Capital Goods	M
Light Weight Vehicle Sales	M
Real Exports of Goods	M
Real Imports of Goods	M
Building Permits	M
Housing Starts	M
New Home Sales	M
Civilian Employment (Household Survey)	M
Unemployed	M
Initial Claims for Unemployment Insurance	M
Soft Indicators	
Markit Manufacturing PMI	M
ISM Manufacturing PMI	M
ISM Non-manufacturing PMI	M
Conference Board: Consumer Confidence	M
University of Michigan: Consumer Sentiment	M
Richmond Fed Mfg Survey	M
Philadelphia Fed Business Outlook	M
Chicago PMI	M
NFIB: Small Business Optimism Index	M
Empire State Manufacturing Survey	M

Estimation

Algorithm (brief overview)

- Model is specified at monthly frequency. Observed growth rates of quarterly variables are related to the unobserved monthly growth rate using a weighted mean (see Mariano and Murasawa, 2003)
- We use a hierarchical implementation of a Gibbs Sampler algorithm (Moench, Ng, and Potter, 2013) which iterates between a small DFM on the outlier adjusted data and the univariate measurement equations. This leads to large computational gains due to parallelisation of this step.
- SVs are sampled following Kim et al. (1998).

- Number of lags in polynomials $\boldsymbol{\Lambda}(\mathbf{L})$, $\phi(L)$, and $\rho(L)$: Set to $m = 1$, $p = 2$, and $q = 2$
- "Minnesota"-style priors applied to coefficients in $\boldsymbol{\Lambda}(L)$, $\phi(L)$ and $\rho_i(L)$. Details
- Variance on priors set to $\frac{\tau}{h^2}$, where τ governs tightness of prior, and h ranges over lag numbers $1:p$, $1:q$, $1:m+1$.
- Following D'Agostino et al. (2015), we set $\tau = 0.2$, a value which is standard in the Bayesian VAR literature.
- Shrink ω_a^2, ω_ε^2 and $\omega_{\eta,i}^2$ towards zero (standard DFM). For ω_a^2 set IG prior with one d.f. and scale 1e-3. For ω_ϵ^2 and $\omega_{\eta,i}^2$ set IG prior with one d.f. and scale 1e-4 (see Primiceri, 2005).

Real-time out of sample evaluation

Real-time out of sample evaluation

Details of data base construction

- We construct a real-time data base for the US and other G7 countries: Germany, France, Italy, Canada, UK, and Japan.
- We consider 2744 data vintages since 11 January 2000 to 31 December 2015.
- For each vintage, sample start is Jan 1960, appending missing observations to any series which starts after that date.
- Sources: (1) ALFRED, (2) Bloomberg, (3) Haver Analytics, (4) OECD Main Economic Indicators.
- Use appropriate deflators for nominal-only vintages.
- Splice data for series with methodological changes.
- Apply seasonal adjustment in real time for survey data.

Real-time out of sample evaluation

Implementation of the exercise

- The model is fully re-estimated every time new data is released/revised. The out of sample exercise starts in January 2000 and ends in December 2015. For the US, on average there is a data release on 15 different dates every month. This means 2744 vintages of data.
- Thanks to efficient implementation of the code, it takes just 2 hours to run 8000 iterations of the Gibbs sampler. But this would mean 7 months of computations!
- Made feasible by using Amazon Web Services cloud computing platform.

Evaluation Results

Evaluation Results

What we show

- In the following slides:
  - Model 0 = Baseline DFM
  - Model 1 = Trend + SV
  - Model 2 = Trend + SV + heterogeneous dynamics
  - Model 3 = Trend + SV + heterogeneous dynamics + fat tails
- We will consider the following metrics
  1. RMSE / MAE
  2. Probability Integral Transform (PITs)
  3. Log Score
  4. Continuous Rank Probability Score (CRPS)
- Everything is for US GDP (1st release)

Results

Forecasts vs. actual over time

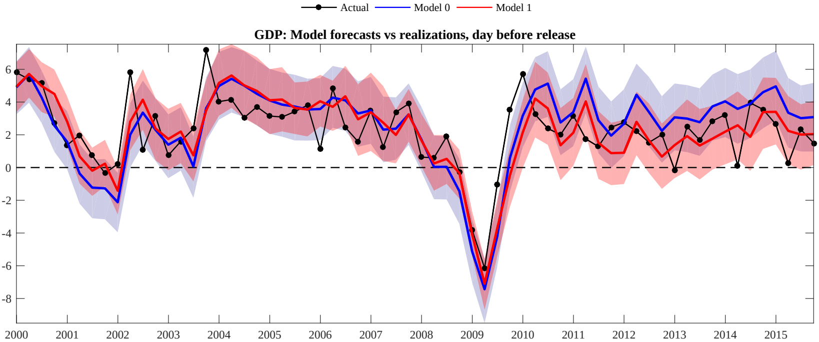

- The addition of the long run trend eliminates the upward bias in GDP forecasts after the crisis...

Results

Forecasts errors over time

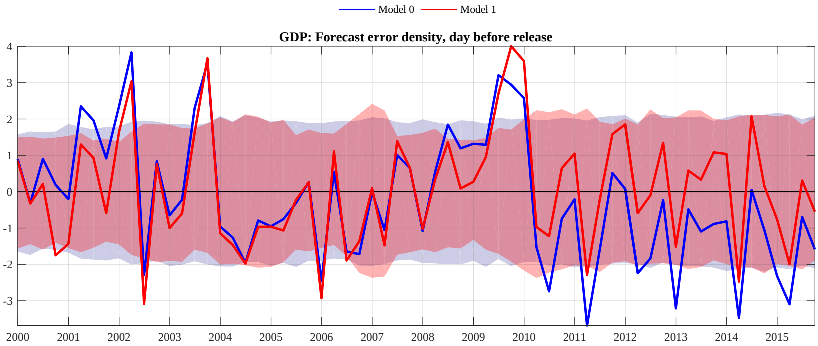

- ...and SV leads to better calibrated density forecasts.

Results

Forecasts vs. actual over time

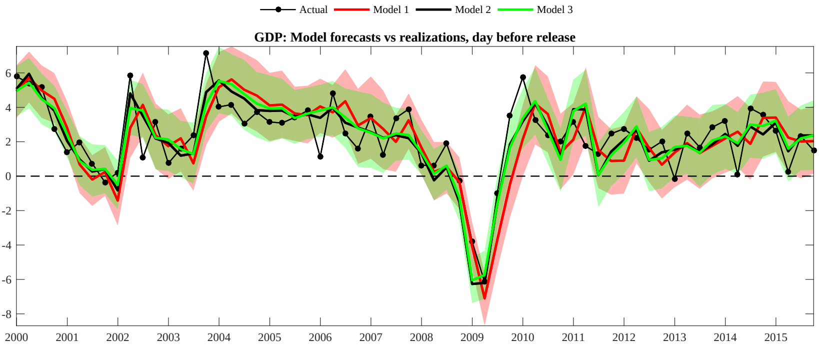

- Heterogeneous dynamics capture recoveries more accurately

Results

GDP: Root mean squared error (RMSE) across horizons

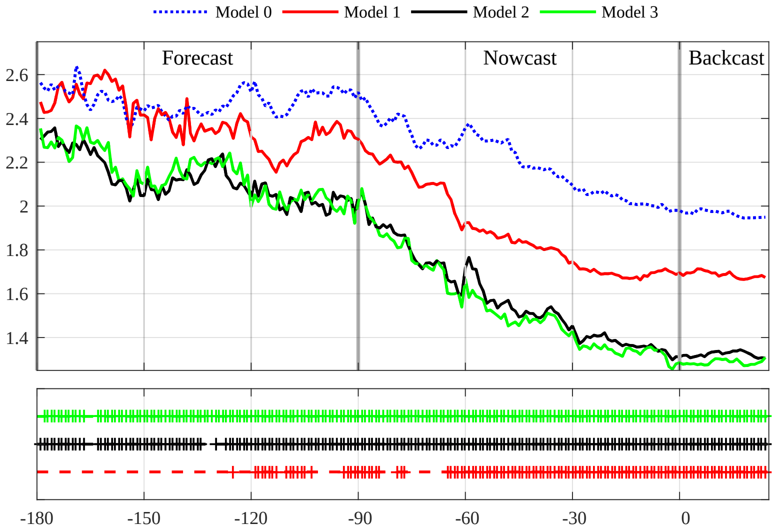

Note: Results with MAE very similar.

Results

Root mean squared error (RMSE), Rolling

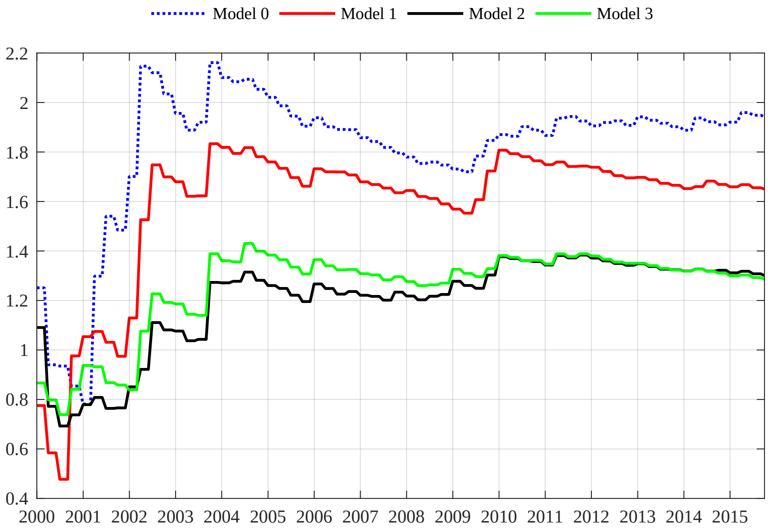

Note: Results with MAE very similar.

Results

Logscore

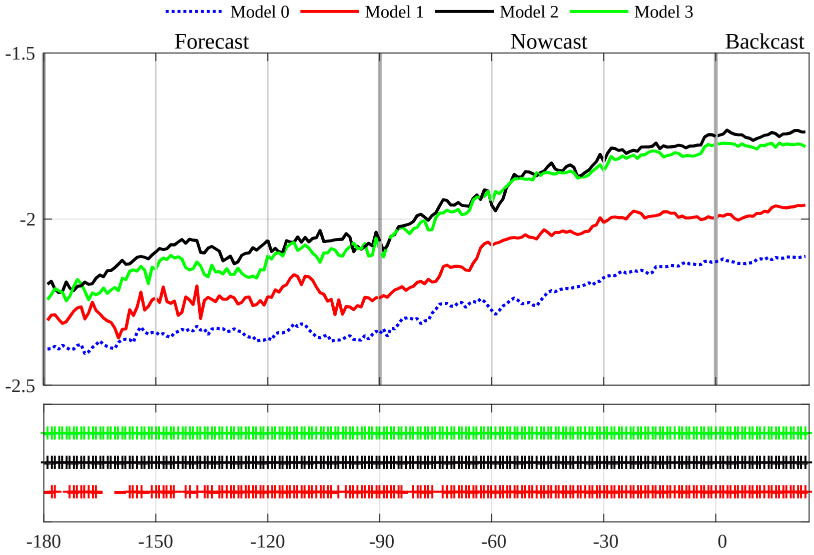

Results

CRPS

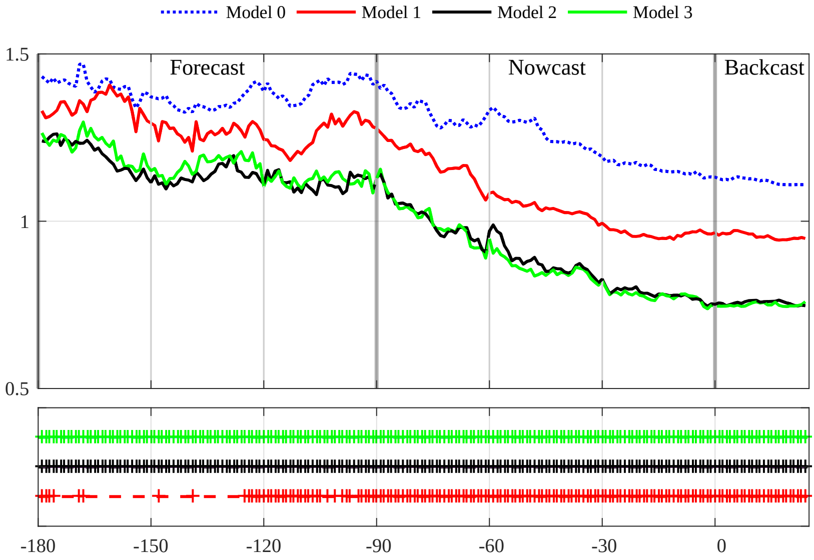

Results

Monthly Variables

	RMSE				LogScore				CRPS			
	M0	M1	M2	M3	M0	M1	M2	M3	M0	M1	M2	M3
INDPRO	0.55	1	**0.97**	**0.98**	-0.91	1.23	1	**0.93****	0.31	**0.99**	**0.97**	**0.97**
NEWORDERS	2.96	1.02	**0.98****	1.01	-2.74	**0.93*****	**0.92*****	**0.95*****	1.82	**0.93*****	**0.9*****	**0.94*****
CARSALES	7.18	**0.99**	1	**0.99**	-3.77	1.24	**0.91***	**0.89***	3.75	**0.99**	1	**0.98**
INCOME	0.83	1	**0.99****	1	-2.14	–	–	1.48	0.38	**0.95*****	**0.92*****	**0.96****
RETAILSALES	0.99	**0.98***	**0.98**	**0.96***	-2.08	1.77	1.94	**0.72***	0.5	**0.91*****	**0.93*****	**0.91*****
EXPORTS	2.51	**0.98****	**0.99****	**0.98****	-2.65	**0.89*****	**0.89*****	**0.88*****	1.62	**0.87*****	**0.88*****	**0.86*****
IMPORTS	2.53	1.03	1.02	1.04	-2.7	**0.92****	**0.9*****	**0.9*****	1.67	**0.88*****	**0.87*****	**0.88*****
PERMIT	5.88	1	1.01	1.02	-3.2	1	1	1	3.32	1.01	1.01	1.01
HOUSINGSTARTS	8.22	1	1	1.01	-3.53	**0.99**	1	**0.99**	4.6	**0.99**	1	1
NEWHOMESALES	8.36	1.01	1.01	1.01	-3.59	1.02	1	**0.99**	4.62	1.02	1.01	1.01
PAYROLL	0.12	**0.96***	**0.87*****	**0.89*****	0.63	1.35	1.49	1.35	0.07	**0.9*****	**0.82*****	**0.85*****
EMPLOYMENT	0.26	1	**0.99**	**0.99**	-0.07	1.93	1.54	**0.72**	0.14	1	**0.99**	**0.99**

- Trend and SV deliver significant improvement
- Substantial additional improvement from heterogeneous dynamics
- Fat tails improve density forecasts of monthly variables

Extension: Residual Seasonality

Modeling residual seasonality

Motivation

- Standard practice: Take seasonally adjusted data directly from statistical agency or adjust using off-the-shelf package outside of model (see e.g. Camacho, Lovcha, Perez-Quiros, 2012).
- Recent data has sparked debate on whether standard methods are vulnerable to "residual seasonality" (see Rudebusch et al., 2015, Stark, 2015, Wright, 2013).
- Not studied in detail: Impact of seasonal adjustment errors in the presence of "ragged edge", i.e. the fact that in real time only a few observations are available at the end of the sample

Modify observation equation to include $\mathbf{s}_t$:

$$\Delta(\mathbf{y}_t - \mathbf{s}_t - \mathbf{o}_t) = \mathbf{c}_t + \mathbf{\Lambda}(\mathbf{L})\mathbf{f}_t + \mathbf{u}_t,$$

where

$$s_{i,t} = -\sum_{j=1}^{11} s_{t-j} + \sigma_{\varepsilon_s,t}\varepsilon_t^s, \quad i = 1, \ldots, n$$

and

$$\log \sigma^s_{\eta_{i,t}} = \log \sigma^s_{\eta_{i,t-1}} + v^s_{\eta_{i,t}}, \quad v^s_{\eta_{i,t}} \overset{iid}{\sim} N(0, \omega^2_{\eta,s,i}) \quad i = 1, \ldots, n$$

Modeling residual seasonality

Results for GDP and GDI

Posterior Estimate of Time-Varying Seasonal Effect by Quarter

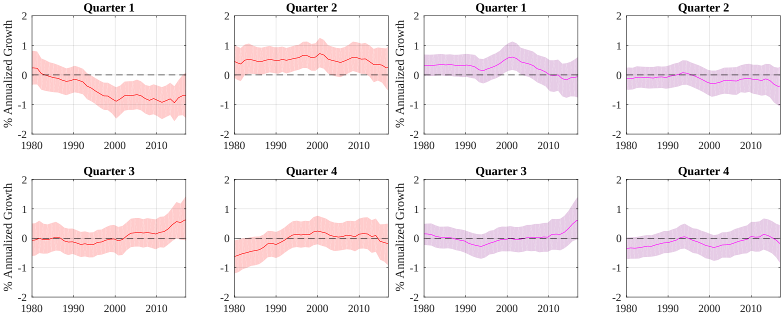

Modeling residual seasonality

Results for GDP and GDI

Fan Chart with seasonal effect

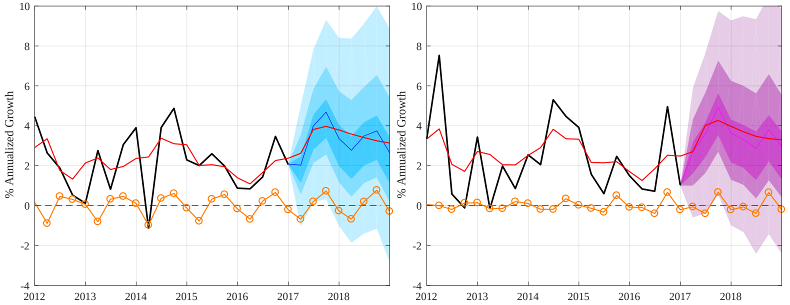

Note: In both panels, the black line represents the actual data. The red line is the underlying measure of economic activity extracted from the DFM. The residual seasonality is shown by the orange line. The blue line is the forecast with the fan charts capturing the uncertainty in the forecast horizon.

Conclusion

Conclusion

Summary of our contribution

- We propose a bayesian DFM, which incorporates:
  1. Low-frequency variation in the mean and variance
  2. Heterogeneous responses to common shocks
  3. Outlier observations and fat tails
  4. Endogenous modeling of seasonality (extension)

- The real-time nowcasting performance is substantially improved across a variety of metrics
  - Capturing trends and SV improves nowcasting performance significantly
  - Heterogeneous dynamics deliver substantial additional improvement
  - Fat tails improve density forecasts of monthly variables (more to be done)

References

Antolin-Diaz, J., T. Drechsel, and I. Petrella (2017): "Tracking the slowdown in long-run GDP growth," *Review of Economics and Statistics*, 99.

Banbura, M. and M. Modugno (2014): "Maximum Likelihood Estimation of Factor Models on Datasets with Arbitrary Pattern of Missing Data," *Journal of Applied Econometrics*, 29, 133–160.

Camacho, M. and G. Perez-Quiros (2010): "Introducing the euro-sting: Short-term indicator of euro area growth," *Journal of Applied Econometrics*, 25, 663–694.

D'Agostino, A., D. Giannone, M. Lenza, and M. Modugno (2015): "Nowcasting Business Cycles: A Bayesian Approach to Dynamic Heterogeneous Factor Models," Tech. rep.

Giannone, D., L. Reichlin, and D. Small (2008): "Nowcasting: The real-time informational content of macroeconomic data," *Journal of Monetary Economics*, 55, 665–676.

Jurado, K., S. C. Ludvigson, and S. Ng (2014): "Measuring Uncertainty," *American Economic Review, Forthcoming*.

McConnell, M. M. and G. Perez-Quiros (2000): "Output Fluctuations in the United States: What Has Changed since the Early 1980's?" *American Economic Review*, 90, 1464–1476.

Moench, E., S. Ng, and S. Potter (2013): "Dynamic hierarchical factor models," *Review of Economics and Statistics*, 95, 1811–1817.

Appendix Slides

Appendix

Details on Priors

- For AR coefficients of factor dynamics, $\phi(L)$, prior mean is set to 0.9 for first lag, and zero in subsequent lags. Reflects a belief that factor captures highly persistent but stationary business cycle process.
- For factor loadings, $\mathbf{\Lambda}(L)$, prior mean is set to 1 for first lag, and zero in subsequent lags. Shrinks model towards factor being the cross-sectional average, see D'Agostino et al. (2015).
- For AR coefficients of idiosyncratic components, $\rho_i(L)$ prior is set to zero for all lags, shrinking model towards no serial correlation in $u_{i,t}$.

Back to main slides