The Macroeconomy as a Random Forest

Philippe Goulet Coulombe

goulet_coulombe.philippe@uqam.ca

Université du Québec à Montréal

June 15, 2021

Final Destination

Modeling *flexibly* macro relationships without assuming what flexible means first. Take something fundamental: a Phillips' curve.

$$u_t^{\text{gap}} \rightarrow \pi_t$$

The statistical characterization of "$\rightarrow$" has forecasting, policy and theoretical (!) implications. Better get it right.

One way out is getting "$\rightarrow$" from off-the-shelf nonparametric Machine Learning (ML) techniques. **But:**

- Likely too flexible and wildly inefficient for the short *time series* we have.
- No obvious parameter(s) to look at — interpretation is fuzzy.

Another is assuming $\pi_t = \beta_t u_t^{\text{gap}} + \text{stuff}_t$. **But:**

- Rigid
- In-sample fit notoriously don't translate in out-of-sample gains.

Solution: ***Generalized*** **Time-Varying Parameters** via Random Forests.

(Machine) Learning β_t's

- I propose *Macroeconomic Random Forests* (MRF): fix the linear part X_t and let the coefficients β_t vary trough time according to a Random Forest.

$$\begin{tikzcd} S_t \arrow[r, "\mathcal{F}_{\text{MRF}}"] \arrow[rr, bend right, "\mathcal{F}_{\text{RF}}"'] & \beta_t \arrow[r, "X_t\beta_t"] & y_t \end{tikzcd}$$

- MRF is nice "meeting halfway"

  $\Rightarrow$ Brings macro closer to ML by squashing many popular nonlinearities (structural change/breaks, thresholds, regime-switching, etc.) into an arbitrarily large S_t, handled easily by RF.

  $\Rightarrow$ The core output are β_t's, *Generalized Time-Varying Parameters* (GTVPs):

$$y_t = X_t\beta_t + \epsilon_t, \qquad \beta_t = \mathcal{F}(S_t)$$

  $\Leftarrow$ Brings ML closer to macro by adapting RF to the reality of economic time series. MRF $\succ$ RF if the linear part is pervasive (like in a (V)AR).

Forecasting around 2008

What do forecasts look like for UR change? $\rightarrow R^2_{OOS}$ 80% for $h = 1$

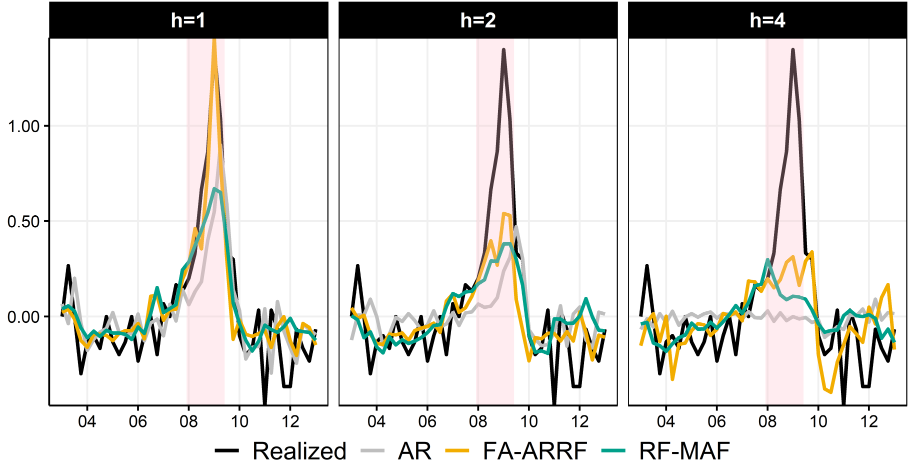

GTVPs of the one-quarter ahead UR forecast

$$\Delta UR_{t+1} = \mu_t + \phi_t^1 y_t + \phi_t^2 y_{t-1} + \gamma_t^1 F_t^1 + \gamma_t^2 F_t^2 + e_{t+1}.$$

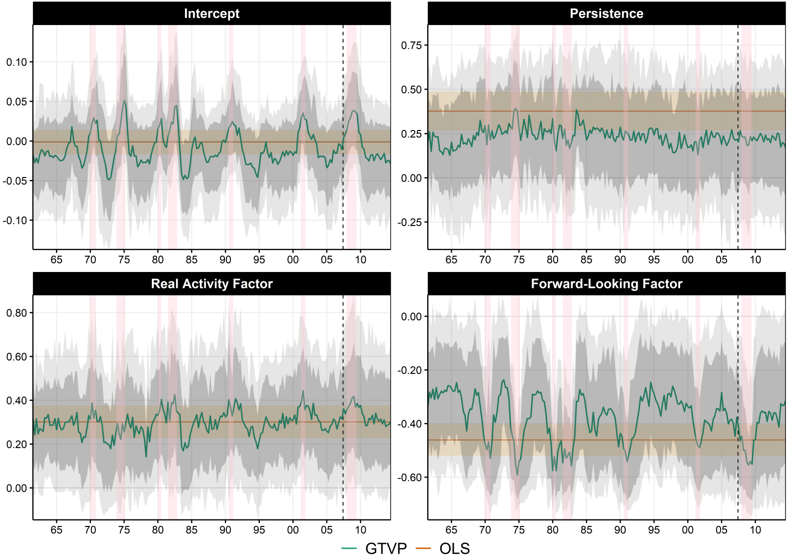

Figure: GTVPs of the one-quarter ahead UR forecast. The grey bands are the 68% and 90% credible region. The pale orange region is the OLS coefficient $\pm$ one standard error. The vertical dotted blue line is the end of the training sample. Pink shading corresponds to NBER recessions.

A Phillips' Curve

À la (Blanchard et al., 2015) and many others

$$\pi_t = \mu_t + \beta_{1,t}\hat{\pi}_t^{SR} + \beta_{2,t}u_t^{GAP} + \beta_{3,t}\pi_t^{IMP} + \varepsilon_t$$

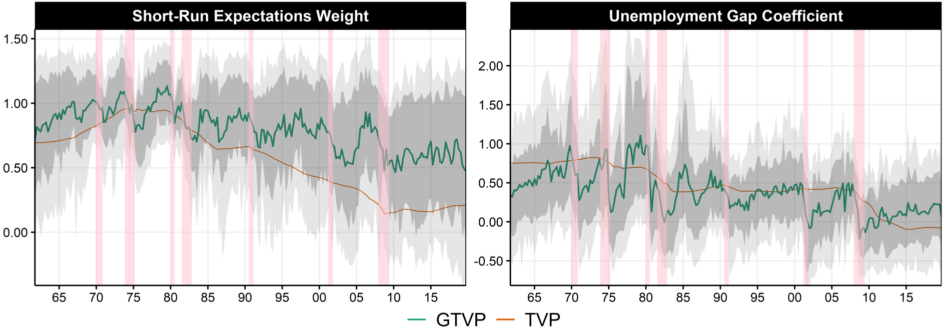

"Conclusion"

I proposed a new time series model that

1. works;
2. is interpretable;
3. is highly versatile;
4. is off-the-shelf (`R package` is available);

Extensions/applications:

- VARs
- Conditional CAPM
- HAR volatility
- Arctic Sea Ice
- DSGEs?
- Anything goes
- I'm personally working on a deep learning version.

Try it with your favorite X_t *today*!

Under Pressure

Employment Cost Index