Financial indicators and density forecasts for US output and inflation

Piergiorgio Alessandri* and Haroon Mumtaz§

*Banca d'Italia and §Queen Mary, University of London.
The presentation does not reflect the official views of Banca d'Italia.

Closing conference of the MaRs network
European Central Bank
23 June 2014

Introduction

Questions

- Are financial indicators useful in forecasting output and inflation?
- Does the answer depend on what kind of **events** the forecaster is interested in predicting? (central case/bad scenarios)
- Does the answer depend on what kind of **models** the forecaster relies on? (linear/nonlinear)
- Was the Great Recession predictable on the basis of real-time financial information?

Introduction

Answers and conjectures

1. Yes (with qualifications)
2. Yes: financial info is particularly useful in predicting "tail outcomes" and recessions.
3. Yes: nonlinear models are better because they account for changes in the size and impact of financial shocks.
4. No idea

The paper in a nutshell (1)
Approach

We cast the analysis as a density prediction problem:

$$m\left(y_t, f_t; \theta\right) \longrightarrow pdf_t\left(y_{t+k}\right)$$

where

- m is an econometric model
- y_t is a macroeconomic variable (industrial production, inflation).
- f_t is a financial indicator (*Financial Condition Index*).

The focus is on (i) role of f_t and (ii) comparison between linear VARs and Threshold VARs (calm/crisis periods).

The paper in a nutshell (2)

Results

1. f_t improves the quality of the predictive densities.
2. TAR generates better densities than VAR
3. TAR could have anticipated (up to a point...) the Great Recession.

Broader implications:

- Predictive distributions are useful to study the finance-macro nexus

The paper in a nutshell (2)

Results

1. f_t improves the quality of the predictive densities.
2. TAR generates better densities than VAR
3. TAR could have anticipated (up to a point...) the Great Recession.

Broader implications:

- Predictive distributions are useful to study the finance-macro nexus
- Non-linearities matter and can be used in forecasting

Literature (1)

1. **Forecasting with financial indicators:** Stock-Watson (2003, 2012), Gilchrist-Yankov-Zakrajšek (2009, 2012); Ng-Wright (2013), ... Emphasis on point forecasts and linear models.
2. **Density forecasting in macro:** eg. Clark (2011). No specific analysis of the role of financial factors.
3. **Early warnings:** Borio-Lowe (2002), Barro-Ursua (2009), Lo Duca-Peltonen (2011). Low frequency data and arbitrary definition of "crises".

This paper

Contributes to (2); proposes density forecasting as a generalisation of (1) and a link between (1) and (3)

Literature (2)

4 **GE models with financial shocks** (Gertler-Kiyotaki 2010; Jermann-Quadrini 2012; Kiyotaki-Moore 2012; Liu-Wang-Zha 2013; ...); and with occasionally binding credit constraints (Bianchi 2012; Bianchi-Mendoza 2011; Guerrieri-Iacoviello 2013).

5 **Empirical models with financial thresholds** (McCallum 1991; Balke 2004; Guerrieri-Iacoviello 2013). Emphasis on impulse-response analysis.

Bottomline: financial shocks matter, and may have different implications in good and bad (credit-constrained) times.

This paper

Studies the nonlinearity modelled in (4) and documented in (5) from a forecasting perspective (see toy PE model in the paper)

- Data
- Models
- Simulating and evaluating distributions
- Results
- Conclusions

Data

US data, March 1973 – August 2012.

y_t : Industrial Production growth
π_t : CPI inflation
r_t : Fed Funds rate
f_t : Financial Conditions Index

FCI is a dynamic factor constructed from an unbalanced panel of 100 mixed-frequency indicators of financial activity (Brave & Butters 2012; Chicago Fed):

- Real time
- Very broad coverage: money, debt and equity markets, financial sector leverage, ...

Financial Condition Index

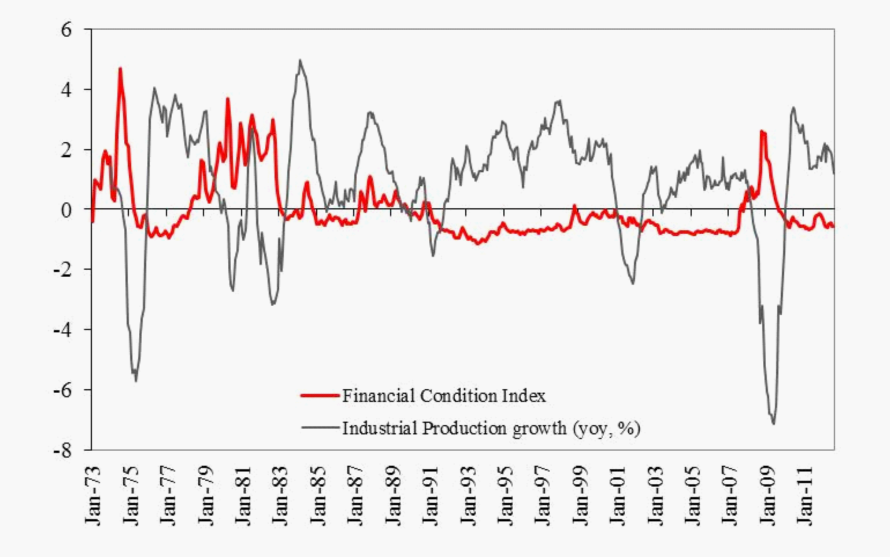

- Data
- **Models**
- Simulating and evaluating distributions
- Results
- Conclusions

Linear vector autoregression (VAR)

$$Y_t = c + \sum_{j=1}^{P} B_j Y_{t-j} + \Omega^{1/2} e_t, \quad e_t \sim N(0, I) \tag{1}$$

We set $P = 13$ and study two specifications

- $VAR^{\S}$: $\quad Y_t = (y_t, \pi_t, r_t)$
- VAR: $\quad Y_t = (y_t, \pi_t, r_t, f_t)$

Natural conjugate prior (N, IW) as in e.g. Banbura-Giannone-Reichlin (JAE, 2010). All variables are treated as independent AR(1) processes: $Y_t = c + \Gamma Y_{t-1} + \Sigma e_t$, where Γ and Σ are diagonal.

Threshold vector autoregression (TAR)

$$
\begin{aligned}
Y_t &= c_{S_t} + \sum_{j=1}^{P} B_{S_t,j} Y_{t-j} + \Omega_{S_t}^{1/2} e_t, \quad e_t \sim N(0, I) && (2)\\
S_t &= \{0, 1\} && (3)\\
S_t &= 1 \iff f_{t-d} \leq f^* && (4)
\end{aligned}
$$

where $Y_t = (y_t, \pi_t, r_t, f_t)$. Note f_t drives the transitions across regimes.

Symmetric natural conjugate prior for the two regimes, as in the VAR.
Agnostic prior for (f^*, d):

$$
f^* \sim N\left(\frac{\Sigma_t f_t}{T}, \bar{k}\right)
$$
$$
d \sim U\{1, ..., 13\}
$$

Estimation

- Bayesian approach
- All priors are *deliberately* uninformative and a-theoretical.
- VAR posterior is known analytically (Banbura et al., 2010).
- TAR posterior can be simulated by Gibbs sampling (Chen-Lee, 1995)
- For each estimation we use 20,000 Gibbs sampling draws and discard the first 15,000

Estimation results (1)

Financial regimes in US history

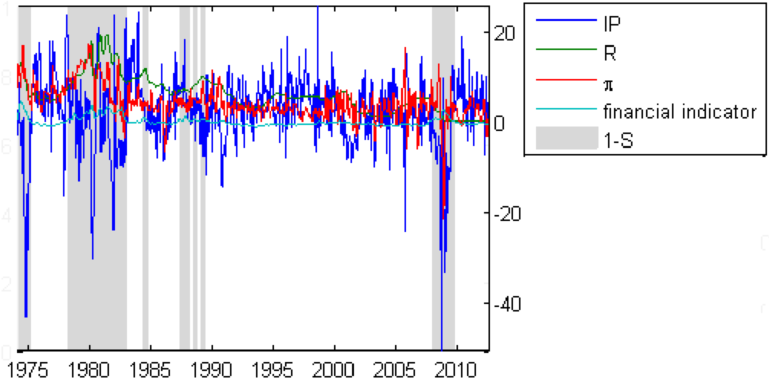

$(1 - \hat{S}_t) = 1 \Leftrightarrow f_{t-d} > f^* \Leftrightarrow$ financial distress/binding credit constraints

Estimation results (2)

The impact of a one-standard-deviation financial shock

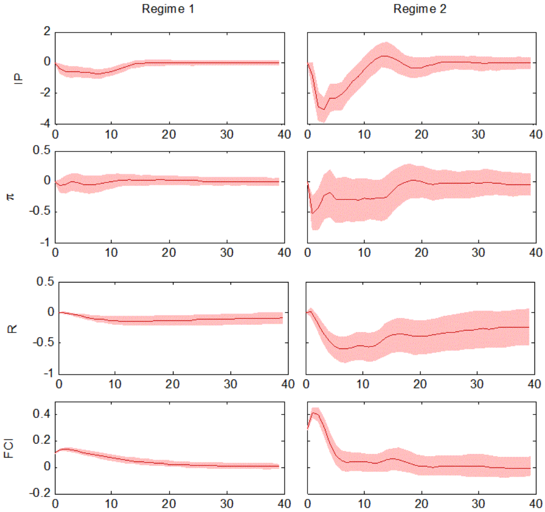

- Data
- Models
- **Simulating and evaluating distributions**
- Results
- Conclusions

Generating the predictive densities (1)

Simulation strategy

Collect a model's parameters into Θ_t.The k-periods ahead predictive density is:

$$p\left(Y_{t+k} \,|\, Y_t\right) = \int p\left(Y_{t+K} \,|\, Y_t, \Theta_{t+k}\right) p\left(\Theta_{t+k} \,|\, Y_t, \Theta_t\right) p\left(\Theta_t \,|\, Y_t\right) d\Theta$$

To simulate the PD:

1. draw $\Theta_t$from the time-t estimate of the posterior (3rd term)
2. simulate forward any time-varying parameters (2nd term)
3. use Θ_{t+k} to simulate paths for Y_{t+k} (1st term).

Implementation:

- Recursive, starting from 1973.03-1983.04 sample
- At each step, the models are simulated up to 12 months ahead.
- This gives 354 density forecasts $p_t^m(Y_{t+k})$ per model m.

Evaluation:

- Calibration diagnostics (skipped for brevity)
- Log-scores: $LS_t^m = \log p_t^m(Y_{t+k}^o)$
- LS high $\Longleftrightarrow$ model m attaches high likelihood *ex-ante* to the actual data Y_{t+k}^o

- Data
- Models
- Simulating and evaluating distributions
- **Results**
- Conclusions

Results

Overview

- f_t yields large improvements in LS for most variables and horizons.

Results
Overview

- f_t yields large improvements in LS for most variables and horizons.
- For industrial production, the improvement is huge around the Great Recession.

Results

Overview

- f_t yields large improvements in LS for most variables and horizons.
- For industrial production, the improvement is huge around the Great Recession.
- TAR produces much better distributions than the linear VARs (higher LS).

Results
Overview

- f_t yields large improvements in LS for most variables and horizons.
- For industrial production, the improvement is huge around the Great Recession.
- TAR produces much better distributions than the linear VARs (higher LS).
- It also produces noisier central forecasts (higher RMSE). So which model should a central bank use?

Results (1) Point vs density forecasting

Average Root Mean Square Errors and Log-Scores

		RMSE				LS			
		1M	3M	6M	12M	1M	3M	6M	12M
VAR§	y	5.604	6.465	6.804	7.019	-3.674	-3.338	-3.418	-3.948
	r	0.167*	0.357	0.598	0.985	-0.675	-1.380	-1.754	-2.118
	π	2.078	2.607*	2.812*	3.077*	-2.584	-2.658	-2.266	-2.137
	f	–	–	–	–	–	–	–	–
VAR	y	5.446*	6.166*	6.558*	6.912*	-3.553	-3.156	-3.032	-2.964
	r	0.177	0.365	0.602	0.989	-0.645	-1.357	-1.723	-2.101
	π	2.067*	2.620	2.839	3.115	-2.583	-2.550	-2.339	-2.171
	f	0.102*	0.197	0.289	0.386	0.135	-0.649	-0.957	-1.130
TAR	y	5.491	6.187	6.594	6.934	-3.491*	-3.152*	-3.005*	-2.885*
	r	0.167	0.338*	0.555*	0.943*	0.022*	-0.778*	-1.364*	-1.999*
	π	2.115	2.667	2.864	3.116	-2.503*	-2.415*	-2.195*	-2.080*
	f	0.104	0.190*	0.271*	0.367*	0.496*	-0.122*	-0.431*	-0.717*

* denotes best model for each criterion/variable/horizon

Results (2) Log-scores

LS for 12M-ahead industrial production growth

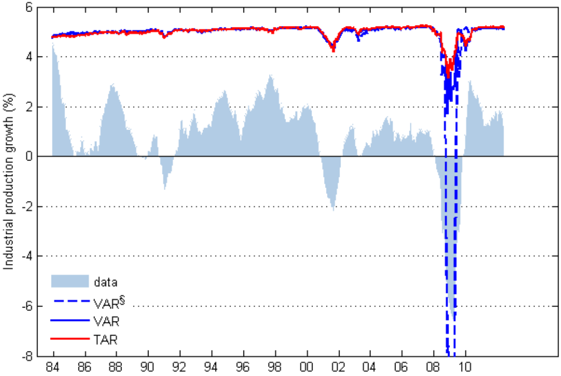

- *VAR*§ is hopeless around the Great Recession
- Recessions are generally hard to predict
- *TAR* beats *VAR* around the Great Recession

Results (3) Bayes factors

Which model should be deemed to be the best as of time t?

Cumulative log predictive Bayes factor: $\Sigma \log \left(LS^{m_1} / LS^{m_2} \right)$

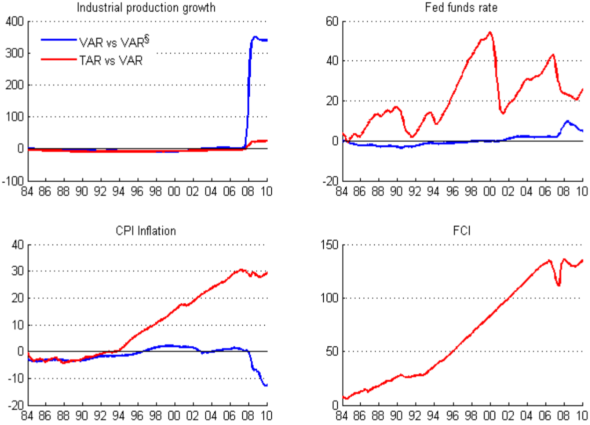

Results (4) Predictive densities and early warnings

Is the signal strong enough to trigger policy actions?

Ex-ante recession probability: $prob_t^m \left(\Sigma_{h=1}^{12} y_{t+h} < \mathbf{0}\right)$

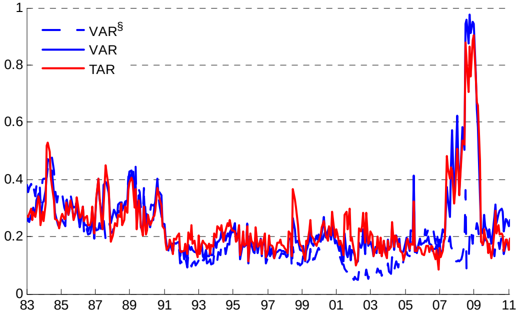

VAR/TAR virtually identical: all that matters is the presence of FCI

Results (4) Predictive densities and early warnings

Is the signal strong enough to trigger policy actions?

Ex-ante "great recession" probability: $prob_t^m \left(\Sigma_{h=1}^{12} y_{t+h} < \textbf{-20\%}\right)$

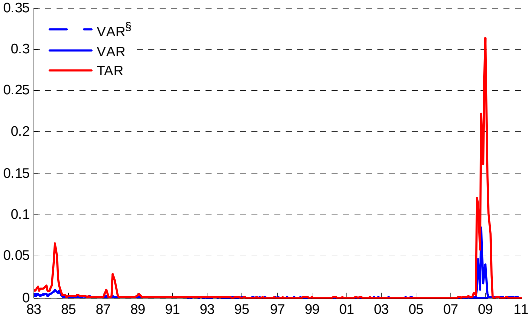

TAR anticipates a more severe downturn.

Robustness

- Data: "excess bond premium" (Gilchrist and Zakrajšek, 2012) instead of Financial Condition Index.

  $\longrightarrow$ Similar qualitative results.

- Models: rolling VAR, Markov-switching VAR with transition probabilities that depend on FCI.

  $\longrightarrow$ Both dominated by TAR.

- Data
- Models
- Simulating and evaluating distributions
- Results
- **Conclusions**

Conclusions

1. **Predictive distributions** are a better tool than point forecasts to study the predictive power of financial indicators.

Conclusions

1. **Predictive distributions** are a better tool than point forecasts to study the predictive power of financial indicators.
2. **Financial indicators** yield large improvements in distributions across most variables/horizons.

Conclusions

1. **Predictive distributions** are a better tool than point forecasts to study the predictive power of financial indicators.
2. **Financial indicators** yield large improvements in distributions across most variables/horizons.
3. **Non-linearities** matter: TAR gives better density forecasts than a VAR. (But it may loose on RMSE – no model is perfect. So think about forecaster's risk preferences.)

Conclusions

1. **Predictive distributions** are a better tool than point forecasts to study the predictive power of financial indicators.
2. **Financial indicators** yield large improvements in distributions across most variables/horizons.
3. **Non-linearities** matter: TAR gives better density forecasts than a VAR. (But it may loose on RMSE – no model is perfect. So think about forecaster's risk preferences.)
4. **Great Recession:** essentially unpredictable, but less so for a TAR with finance-driven regimes.

Next steps

- Work out distributional implications of credit constraints in a (more) general equilibrium model.
- Think formally about risk preferences and model selection.
- Refine priors on good/bad regimes
- More robustness (sample, prior hyperparameters, ...)

Thanks!