## Economic predictions with big data: The illusion of sparsity

Domenico Giannone
New York Fed

Michele Lenza
European Central Bank

Giorgio Primiceri
Northwestern University
$10^{\text {th }}$ ECB Workshop on Forecasting Techniques
Frankfurt, June 19, 2018

## Predictive modeling with big data

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i i d} \mathcal{N}\left(0, \sigma^{2}\right)
$$

## Predictive modeling with big data

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{\overline{i d}} \widetilde{\mathcal{N}}\left(0, \sigma^{2}\right)
$$

- Big data: large $k$


## Predictive modeling with big data

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i \overline{i d}} \mathcal{N}\left(0, \sigma^{2}\right)
$$

- Big data: large $k$
- Standard inference (ML or flat prior) is a bad idea
$>$ Proliferation of parameters
$>$ High estimation uncertainty
$>$ Overfitting and imprecise out-of-sample forecasting / poor external validity


## Predictive modeling with big data

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i i d} \mathcal{N}\left(0, \sigma^{2}\right)
$$

- Big data: large $k$
- Standard inference (ML or flat prior) is a bad idea
$>$ Proliferation of parameters
$>$ High estimation uncertainty
$>$ Overfitting and imprecise out-of-sample forecasting / poor external validity
$\Rightarrow$ Methods to address curse of dimensionality (Ng, 2013, CHL, 2017)
$>$ Sparse modeling e.g. hand picking, Lasso regression
$>$ Dense modeling e.g. Ridge regression, Factor models


## This paper: Sparse or dense modeling?

## This paper: Sparse or dense modeling?

- Answer is an empirical matter
> Study a variety of predictive problems in macro, micro and finance


## This paper: Sparse or dense modeling?

- Answer is an empirical matter
> Study a variety of predictive problems in macro, micro and finance

■ Popular techniques not suitable to answer the question
$>$ Sparsity/density often assumed
$>$ A small set of predictors might be selected simply to reduce estimation error, even if the model is not sparse

## This paper: Sparse or dense modeling?

- Answer is an empirical matter
> Study a variety of predictive problems in macro, micro and finance

■ Popular techniques not suitable to answer the question
> Sparsity/density often assumed
> A small set of predictors might be selected simply to reduce estimation error, even if the model is not sparse

- Our predictive model
$>$ sparsity, without assuming it
$>$ shrinkage, to give a chance to large models
> Bayesian inference on sparsity and shrinkage


## Main results

1. No clear pattern of sparsity
> Posterior not concentrated on a single sparse model, but on a wide set

## Main results

1. No clear pattern of sparsity
> Posterior not concentrated on a single sparse model, but on a wide set
2. More sparsity emerges only if very tight prior favoring small models

## Main results

1. No clear pattern of sparsity
> Posterior not concentrated on a single sparse model, but on a wide set
2. More sparsity emerges only if very tight prior favoring small models $\pm$

The illusion of sparsity

## Outline

■ The predictive model

- Applications to macro, micro and finance
- Sparse or dense modeling?
> Exploring the posterior


## Outline

- The predictive model
- Applications to macro, micro and finance

■ Sparse or dense modeling?
> Exploring the posterior

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i \overline{i d}} \mathcal{N}\left(0, \sigma^{2}\right)
$$

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i \overline{i d}} \widetilde{\mathcal{N}}\left(0, \sigma^{2}\right)
$$

- Prior

$$
\beta_{i} \mid \sigma^{2}, \gamma^{2}, q \underset{\text { iid }}{\tilde{u}}\left\{\begin{array}{l}
0 \quad \text { with probability } \quad 1-q
\end{array}\right.
$$

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i \overline{i d}} \mathcal{N}\left(0, \sigma^{2}\right)
$$

- Prior

$$
\beta_{i} \mid \sigma^{2}, \gamma^{2}, q \underset{\text { iid }}{\sim}\left\{\begin{array}{cl}
\mathcal{N}\left(0, \sigma^{2} \gamma^{2}\right) & \text { with probability } \\
0 & \text { with probability } 1-q
\end{array}\right.
$$

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i \overline{i d}} \mathcal{N}\left(0, \sigma^{2}\right)
$$

- Prior

$$
\beta_{i} \mid \sigma^{2}, \gamma^{2}, q \underset{\text { iid }}{\sim}\left\{\begin{array}{cl}
\mathcal{N}\left(0, \sigma^{2} \gamma^{2}\right) & \text { with probability } \\
0 & \text { with probability } \\
0 & 1-q
\end{array}\right.
$$

## Probability of inclusion, controls size

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{\overline{i d}} \widetilde{\mathcal{N}}\left(0, \sigma^{2}\right)
$$

- Prior



## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i \overline{i d}} \mathcal{N}\left(0, \sigma^{2}\right)
$$

- Prior

$$
\beta_{i} \mid \sigma^{2}, \gamma^{2}, q \underset{\text { iud }}{\sim}\left\{\begin{array}{cl}
\mathcal{N}\left(0, \sigma^{2} \gamma^{2}\right) & \text { with probability } \\
0 & \text { with probability } 1-q
\end{array}\right.
$$

■ "Spike-and-slab" prior
> Mitchel and Beauchamp (1988)
> Vast literature on Bayesian Model Averaging and Variable Selection

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i \overline{i d}} \mathcal{N}\left(0, \sigma^{2}\right)
$$

- Prior

$$
\beta_{i} \mid \sigma^{2}, \gamma^{2}, q \underset{\text { iid }}{\sim}\left\{\begin{array}{cl}
\mathcal{N}\left(0, \sigma^{2} \gamma^{2}\right) & \text { with probability } \\
0 & \text { with probability } \\
0 & 1-q
\end{array}\right.
$$

■ "Spike-and-slab" prior
> Mitchel and Beauchamp (1988)
$>$ Vast literature on Bayesian Model Averaging and Variable Selection
> This paper: inference on $q$ and $\gamma^{2}$

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i i d} \mathcal{N}\left(0, \sigma^{2}\right)
$$

- Prior

$$
\beta_{i} \mid \sigma^{2}, \gamma^{2}, q \underset{\text { iid }}{\sim}\left\{\begin{array}{cc}
\mathcal{N}\left(0, \sigma^{2} \gamma^{2}\right) & \text { with probability } \\
0 & \text { with probability } \\
0 & 1-q
\end{array}\right.
$$

■ Hyperpriors

$$
q \sim \mathcal{B}(a, b)
$$

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i \overline{i d}} \mathcal{N}\left(0, \sigma^{2}\right)
$$

- Prior

$$
\beta_{i} \mid \sigma^{2}, \gamma^{2}, q \underset{i i d}{\sim}\left\{\begin{array}{cc}
\mathcal{N}\left(0, \sigma^{2} \gamma^{2}\right) & \text { with probability } \\
0 & \text { with probability } \\
0 & 1-q
\end{array}\right.
$$

■ Hyperpriors

$$
q \sim \mathcal{B}(a, b), \quad \frac{q k \operatorname{var}(x) \gamma^{2}}{q k \operatorname{var}(x) \gamma^{2}+1} \equiv R^{2}
$$

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i \overline{i d}} \mathcal{N}\left(0, \sigma^{2}\right)
$$

- Prior

$$
\beta_{i} \mid \sigma^{2}, \gamma^{2}, q \underset{\text { iid }}{\sim}\left\{\begin{array}{cl}
\mathcal{N}\left(0, \sigma^{2} \gamma^{2}\right) & \text { with probability } \\
0 & \text { with probability } \\
0 & 1-q
\end{array}\right.
$$

■ Hyperpriors

$$
q \sim \mathcal{B}(a, b), \quad \searrow \frac{q k \operatorname{var}(x) \gamma^{2}}{q k \operatorname{var}(x) \gamma^{2}+1} \equiv R^{2}
$$

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i \overline{i d}} \mathcal{N}\left(0, \sigma^{2}\right)
$$

- Prior

$$
\beta_{i} \mid \sigma^{2}, \gamma^{2}, q \underset{i i d}{\sim}\left\{\begin{array}{cc}
\mathcal{N}\left(0, \sigma^{2} \gamma^{2}\right) & \text { with probability } \\
0 & \text { with probability } \\
0 & 1-q
\end{array}\right.
$$

■ Hyperpriors

$$
q \sim \mathcal{B}(a, b), \quad \frac{q k \operatorname{var}(x) \gamma^{2}}{q k \operatorname{var}(x) \gamma^{2}+1} \equiv R^{2} \sim \mathcal{B}(A, B)
$$

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i \overline{i d}} \mathcal{N}\left(0, \sigma^{2}\right)
$$

- Prior

$$
\beta_{i} \mid \sigma^{2}, \gamma^{2}, q \underset{i i d}{\sim}\left\{\begin{array}{cc}
\mathcal{N}\left(0, \sigma^{2} \gamma^{2}\right) & \text { with probability } \\
0 & \text { with probability } \\
0 & 1-q
\end{array}\right.
$$

■ Hyperpriors

$$
q \sim \mathcal{B}(1,1), \quad \frac{q k \operatorname{var}(x) \gamma^{2}}{q k \operatorname{var}(x) \gamma^{2}+1} \equiv R^{2} \sim \mathcal{B}(1,1)
$$

The implied joint prior on $q$ and $\gamma^{2}$


## Outline

■ The predictive model

- Applications to macro, micro and finance

■ Sparse or dense modeling?
> Exploring the posterior

## Economic applications

- Macro
- Forecasting industrial production with many macro predictors
$>$ The determinants of economic growth in a cross-section of countries
- Finance
> Prediction of the US aggregate equity premium over time
$>$ Explaining the cross-section of equity returns across firms
- Micro
> Understanding the decline in crime rates in US states during the 1990s
> The determinants of government takings of private properties in US judicial circuits
- Some references:
> Stock-Watson (2002a and b), Barro-Lee (1994), Sala-i-Martin et al. (2004), Welch-Goyal (2008), Freyberger et al. (2017), Donohue-Levitt (2001), Chen-Yeh (2012), Belloni et al. (2011, 2012, 2014).


## Economic applications

## Macro 1

Macro 2

Finance 1 US equity premium

Countries average growth 1960-1985

Growth rate of US Industrial Prod.

Micro 1

Micro 2

Stock returns of US firms

Crime rate in US states

Eminent domain judicial decisions

130 lagged macro and financial indicators

60 country charact' socio-econ, inst.

16 lagged macro and financial indicators

## Sample

659 time-series obs. Feb. 60-Dec. 14

90 cross-section obs.

58 time-series obs. 1948-2015
$\approx 1400 \mathrm{k}$ panel obs. Jul. 63-Dec. 15, $\approx 2 k$ firms

476 panel obs.
Jan. 86-Dec. 97, 48 states
312 panel obs.
1975-2008, circuits

## Outline

■ The predictive model

- Applications to macro, micro and finance
- Sparse or dense modeling?
> Exploring the posterior


## Exploring the posterior

1. No clear pattern of sparsity
> Posterior not concentrated on a single sparse model, but on a wide set
2. More sparsity emerges only if very tight prior favoring small models

## Exploring the posterior

0. Inclusion probability and shrinkage are complements, but imperfect
1. No clear pattern of sparsity
> Posterior not concentrated on a single sparse model, but on a wide set
2. More sparsity emerges only if very tight prior favoring small models

## Exploring the posterior

0. Inclusion probability and shrinkage are complements, but imperfect
1. No clear pattern of sparsity
> Posterior not concentrated on a single sparse model, but on a wide set
2. More sparsity emerges only if very tight prior favoring small models

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{\overline{i d}} \widetilde{\mathcal{N}}\left(0, \sigma^{2}\right)
$$

- Prior



## Probability of inclusion and prior variance ( $q$ and $\gamma^{2}$ )








## Probability of inclusion and prior variance ( $q$ and $\gamma^{2}$ )



Finance 1






## Probability of inclusion and prior variance ( $q$ and $\gamma^{2}$ )



Finance 1


Micro 1



Finance 2



## Posterior probability of inclusion: $p(q \mid Y)$




Finance 1





## Exploring the posterior

0. Inclusion probability and shrinkage are complements, but imperfect
1. No clear pattern of sparsity
> Posterior not concentrated on a single sparse model, but on a wide set
2. More sparsity emerges only if very tight prior favoring small models

## Posterior probability of inclusion: $p(q \mid Y)$




Finance 1





## Patterns of sparsity:

## Probability of inclusion of each coefficient



## Exploring the posterior

0. Inclusion probability and shrinkage are complements, but imperfect
1. No clear pattern of sparsity
> Posterior not concentrated on a single sparse model, but on a wide set

- Predictors rarely systematically excluded
- Model uncertainty is pervasive
- Best predictions not with single model, but mixture of many (BMA)

2. More sparsity emerges only if very tight prior favoring small models

## Exploring the posterior

0. Inclusion probability and shrinkage are complements, but imperfect
1. No clear pattern of sparsity
> Posterior not concentrated on a single sparse model, but on a wide set
2. More sparsity emerges only if very tight prior favoring small models

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i \overline{i d}} \mathcal{N}\left(0, \sigma^{2}\right)
$$

- Prior

$$
\beta_{i} \mid \sigma^{2}, \gamma^{2}, q \underset{i i d}{\sim}\left\{\begin{array}{cc}
\mathcal{N}\left(0, \sigma^{2} \gamma^{2}\right) & \text { with probability } \\
0 & \text { with probability } \\
0 & 1-q
\end{array}\right.
$$

■ Hyperpriors

$$
q \sim \mathcal{B}(1,1), \quad \frac{q k \operatorname{var}(x) \gamma^{2}}{q k \operatorname{var}(x) \gamma^{2}+1} \equiv R^{2} \sim \mathcal{B}(1,1)
$$

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i \overline{i d}} \mathcal{N}\left(0, \sigma^{2}\right)
$$

- Prior

$$
\beta_{i} \mid \sigma^{2}, \gamma^{2}, q \underset{i i d}{\sim}\left\{\begin{array}{cc}
\mathcal{N}\left(0, \sigma^{2} \gamma^{2}\right) & \text { with probability } \\
0 & \text { with probability } \\
0 & 1-q
\end{array}\right.
$$

■ Hyperpriors

$$
q \sim \mathcal{B}(1, k), \quad \frac{q k \operatorname{var}(x) \gamma^{2}}{q k \operatorname{var}(x) \gamma^{2}+1} \equiv R^{2} \sim \mathcal{B}(1,1)
$$

## The predictive model

$$
y_{t}=\beta_{1} x_{1 t}+\cdots+\beta_{k} x_{k t}+\varepsilon_{t}, \varepsilon_{t} \widetilde{i \overline{i d}} \widetilde{\mathcal{N}}\left(0, \sigma^{2}\right)
$$

- Prior

$$
\beta_{i} \mid \sigma^{2}, \gamma^{2}, q \underset{\text { iid }}{\sim}\left\{\begin{array}{cl}
\mathcal{N}\left(0, \sigma^{2} \gamma^{2}\right) & \text { with probability } \\
0 & \text { with probability } \\
0 & 1-q
\end{array}\right.
$$

Castillo et al. (2015, Annals)
■ Hyperpriors

$$
q \sim \mathcal{B}(1, k), \quad \frac{q k \operatorname{var}(x) \gamma^{2}}{q k \operatorname{var}(x) \gamma^{2}+1} \equiv R^{2} \sim \mathcal{B}(1,1)
$$

## Patterns of sparsity with a flat prior on $q$

## Probability of inclusion of each coefficient



Macro 2



Micro 2


## Patterns of sparsity with a tight prior on low $q$

## Probability of inclusion of each coefficient




## Patterns of sparsity with a flat prior on $q$

## Probability of inclusion of each coefficient, given $q$

Macro 1: flat hyperprior


Macro 1: flat hyperprior


## Patterns of sparsity with a tight prior on $q$

## Probability of inclusion of each coefficient, given $q$




## Baseline hyperprior: flat on $q$




Finance 1



Micro 1



## Alternative hyperprior: tight on low $q$



Finance 1


Micro 1



## Posterior of $R^{2}$ with a flat and a tight prior on $q$



Finance 1


Micro 1



Finance 2



## Summing up

0. Inclusion probability and shrinkage are complements, but imperfect
1. No clear pattern of sparsity
> Posterior not concentrated on a single sparse model, but on a wide set
2. More sparsity emerges only if very tight prior favoring small models $\sqrt{\square}$

## The illusion of sparsity

## The prior distribution

- Alternative representation

$$
\begin{gathered}
\beta_{i} \mid \sigma^{2}, \gamma^{2}, q \underset{\text { iid }}{\sim} \mathcal{N}\left(0, \sigma^{2} \gamma^{2} z_{i}\right) \\
z_{i} \mid q_{\text {iid }}^{\sim} \operatorname{Bernoulli}(q)
\end{gathered}
$$

- Relation with other popular shrinkage methods
> Ridge: $q=1$
$>$ Lasso: $z_{i} \tilde{\text { iid }}$ Exponential
$>$ Lava: $z_{i} \underset{\text { iid }}{\text { Shifted Exponential }}$
$>$ Horse shoe: $z_{i} \tilde{\text { iid }}$ Half Cauchy
$>$ Elastic net: $z_{i} \widetilde{\text { iid }}$ transformation of a truncated Gamma
- None admits a sparse representation of with positive probability


## Bayesian interpretation of various shrinkage methods



## $p(q \mid Y)$ as a measure of predictive accuracy

- Posterior of $q$

$$
p(q \mid Y) \propto p(Y \mid q) \cdot p(q)
$$

## $p(q \mid Y)$ as a measure of predictive accuracy

- Posterior of $q$

$$
p(q \mid Y) \propto p(Y \mid q)
$$

## $p(q \mid Y)$ as a measure of predictive accuracy

- Posterior of $q$

$$
p(q \mid Y) \propto p(Y \mid q)=\prod_{t}^{T} p\left(y_{t} \mid y^{t-1}, q\right)
$$

## $p(q \mid Y)$ as a measure of predictive accuracy

- Posterior of $q$

$$
p(q \mid Y) \propto p(Y \mid q)=\prod_{t}^{T} p\left(y_{t} \mid y^{t-1}, q\right)
$$

predictive score

## $p(q \mid Y)$ as a measure of predictive accuracy

- Posterior of $q$

$$
p(q \mid Y) \propto p(Y \mid q)=\prod_{t}^{T} p\left(y_{t} \mid y^{t-1}, q\right)
$$

$\Rightarrow$ Average log-predictive score

$$
\frac{1}{T} \sum_{t}^{T} \log p\left(y_{t} \mid y^{t-1}, q\right)=\frac{1}{T} \log p(q \mid Y)+\text { constant }
$$

## Average log-predictive score, relative to best fitting model








## Probability of inclusion $(q)$ and $R^{2}$



