# Enhanced Bayesian Neural Networks for Macroeconomics and Finance 

Niko Hauzenberger ${ }^{1}$, Florian Huber ${ }^{1}$, Karin Klieber ${ }^{2}$, Massimiliano Marcellino ${ }^{3}$<br>${ }^{1}$ University of Salzburg, ${ }^{2}$ Oesterreichische Nationalbank, ${ }^{3}$ Bocconi University


*The content of these slides reflects the views of the authors and not necessarily those of the OeNB or

## Karin Klieber

Monetary Policy Section
Oesterreichische Nationalbank
karin.klieber@oenb.at

## Motivation (1)

US inflation


US-UK exchange rate


US industrial production



## Motivation (2)

- Practitioners and academics rely on large datasets to form forecasts, tailor policies or improve decisions, often particularly relevant in problematic times
- Regularization-based techniques as popular way to overcome the curse of dimensionality (Belmonte et al., 2014; Bhattacharya and Dunson, 2011; Carvalho et al., 2010; Griffin and Brown, 2013; Huber and Pfarrhofer, 2021)
- However, common assumption of linearity often remains
- Two major questions arise:
- How to model the relationship between a response and a large set of covariates?
- How to safeguard against overfitting?


## What we do (1)

## Methodological

- Apply neural networks as a device for learning any function under relatively few assumptions (Hornik et al., 1989), yet, performance of NNs in macro forecasting rather bad (Makridakis S., 2018)
- Use Bayesian methods to determine the adequate network structure
- Allowing for either sparse or dense datasets (Giannone et al., 2021)
- Combining alternative activation functions (Agostinelli et al., 2014; Karlik and Olgac, 2011)
- Selecting the number of neurons via shrinkage (Bhattacharya and Dunson, 2011; Carvalho et al., 2010)
- Introducing heteroscedasticity in the error terms (Kastner and Frühwirth-Schnatter, 2014)
- Design efficient MCMC algorithm for estimation of the resulting BNN even with large set of regressors


## What we do (2)

## Empirical

- Show that our approach works well in simulations
- Apply BNNs to a set of prominent macro and finance applications
- Conduct a thorough forecasting exercise
- Explore the degree of non-linearities


## Econometric framework

A general non-linear regression

$$
y_{t}=\boldsymbol{x}_{t}^{\prime} \gamma+f\left(\boldsymbol{x}_{t}\right)+\varepsilon_{t}, \quad \varepsilon_{t} \sim \mathcal{N}\left(0, \sigma_{t}^{2}\right)
$$

- $y_{t}$ denotes a scalar time series,
- $\boldsymbol{x}_{t}$ a set of $K$ covariates,
- $\gamma$ a vector of $K$ (linear) coefficients,
- $f: \mathbb{R}^{K} \rightarrow \mathbb{R}$ a function of unknown (non-linear) form
- How to specify $f$ ?


## A shallow Bayesian Neural Network

A neural network with a single hidden layer

$$
f\left(\boldsymbol{x}_{t}\right) \approx \sum_{q=1}^{Q} \beta_{q} h_{q}\left(\boldsymbol{x}_{t}^{\prime} \kappa_{q}+\zeta_{q}\right)
$$

- $\boldsymbol{\beta}=\left(\beta_{1}, \ldots, \beta_{Q}\right)^{\prime}$ denotes a $Q \times 1$ vector of factor loadings
- $Q$ the number of neurons
- $h(\bullet)$ a non-linear activation function
- $\kappa=\left(\kappa_{1}, \ldots, \kappa_{Q}\right)$ a $K \times Q$ matrix of non-linear coefficients
- $\boldsymbol{\zeta}=\left(\zeta_{1}, \ldots, \zeta_{Q}\right)^{\prime}$ a $Q \times 1$ vector of bias terms


From Neural Networks to Bayesian Neural Networks


## Stochastic model selection

- Selecting the number of neurons ( $Q$ )
- Multiplicative Gamma Process (MGP) prior (Bhattacharya and Dunson, 2011)

$$
\beta_{j} \sim \mathcal{N}\left(0, \phi_{j}^{-1}\right), \quad \phi_{j}=\prod_{q=1}^{Q} \delta_{q}, \quad \delta_{1} \sim \mathcal{G}\left(a_{1}, 1\right), \quad \delta_{l} \sim \mathcal{G}\left(a_{2}, 1\right), \text { for } q>1
$$

- Shrinking the weighting coefficients ( $\kappa$ )
- Column-wise horseshoe prior (Carvalho et al., 2010)

$$
\kappa_{q j} \sim \mathcal{N}\left(0, \lambda_{q}^{2} \phi_{q j}^{2}\right), \quad \lambda_{q} \sim \mathcal{C}^{+}(0,1), \quad \phi_{q j} \sim \mathcal{C}^{+}(0,1)
$$

- Choosing between (four) activation functions ( $h$ )
- Introduce latent discrete random variable $\delta_{q} \in[1,4]$

$$
\operatorname{Prob}\left(\delta_{q}=j\right)=\underline{\omega}_{q j}=\frac{1}{4}
$$

- Shrink $\gamma$ with the horseshoe prior (Carvalho et al., 2010)


## Set of activation functions



## Illustration of the activation functions

## Effect of money growth on inflation


linear


sigmoid



## Posterior simulation

- Coefficients $\gamma$ and $\beta$ are obtained jointly from a standard multivariate Gaussian posterior
- Hyperparameters for MGP prior are obtained through simple Gibbs updating steps (Bhattacharya and Dunson, 2011)
- $p\left(\kappa_{q} \mid \bullet\right)$, for $q=1, \ldots, Q$ :
- If the corresponding scaling parameter of the MGP prior exceeds a threshold very close to zero, we sample $\kappa_{q}$ using HMC step (Neal et al., 2011)
- Otherwise, $\kappa_{q}$ is obtained by drawing from the prior
- $h_{q}(x)$ is simulated from a multinomial distribution resulting in a mixture of activation functions
- $\left\{\log \left(\sigma_{t}^{2} \mid \bullet\right)\right\}_{t=1}^{T}$ is simulated using the algorithm proposed in Kastner and Frühwirth-Schnatter (2014)


## Simulation study

- We illustrate our approach through different DGPs and vary
- the functional form: linear vs highly nonlinear (neural network model)
- the variance specification: homoskedastic vs heteroskedastic
- the size of the model: large (60) vs small (30) number of regressors
- the parameter space: sparse vs dense
- We show that our approach works well for both DGPs
- We clearly outperform the linear model for the highly nonlinear DGPs
- We can capture the linear DGP without overfitting
- Details are given in the appendix

Forecasting setup (1)

|  | Dependent variable | Set of predictors | Sample | Range | - | Hold-out | Source |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Macro A | A.1) Industrial production <br> A.2) Inflation <br> A.3) Employment | Large (120 economic \& financial variables) | Monthly data for the US | $\begin{aligned} & \text { 1960M1 to } \\ & \text { 2020M12 } \end{aligned}$ | one-step-ahead | $\begin{aligned} & \text { 2000M1 to } \\ & \text { 2020M12 } \end{aligned}$ | McCracken and Ng (2016) |
| Macro B | Average economic growth rate | 60 country-specific characteristics | Cross-section | 90 countries | randomly sampled (100 times) | 45 countries | Barro and Lee (1994) |
| Macro C | USD/GBP exchange rate returns (q०q) | 20 exchange rate determinants | Quarterly data for the US and UK | $\begin{aligned} & \text { 1990Q1 to } \\ & \text { 2019Q4 } \end{aligned}$ | one-step- and four-steps-ahead | $\begin{aligned} & \text { 2000Q1 to } \\ & \text { 2019Q4 } \end{aligned}$ | Wright (2008); Rossi (2013) |
| Finance | Equity premium | 16 economic \& financial variables | Annual data for the US | 1948 to 2020 | one-year-ahead | 1965 to 2020 | Welch and Goyal (2008) |

## Forecasting setup (2)

- Set of competing models:
- Bayesian neural network (BNN)
- BNN with neuron-specific activation functions (BNN-NS)
- Bayesian linear regression model with the horseshoe shrinkage prior and stochastic volatility (benchmark)
- Bayesian neural network by backpropagation, labeled BNN-BP (Blundell et al., 2015) Details
- Bayesian additive regression trees, labeled BART (Chipman et al., 2010; Huber et al., 2020; Huber and Rossini, 2022) Details
- Evaluation metrics:
- Root mean squared error (RMSE) for point forecasts
- Log predictive likelihood (LPL) for density forecasts


## Summary of findings

- Superior forecasting performance
- Our BNN offers substantial improvements in density forecasting performance
- We find competitive performance when interest centers on point forecasts
- Controlling for non-linear relations is of particular importance during recessionary periods of the business cycle
- For Macro A, we see the highest gains during the Global Financial Crisis and the COVID-19 pandemic
- For Macro C and Finance, the BNN yields superior density forecasting performance during the Global Financial Crisis
- In the cross section, the BNN handles extraordinarily low and high growth rates best
- Good performance in terms of density forecasts is often accompanied by larger in-sample fit for extreme observations
- Deep BNN yields comparable results Dealis


## Out-of-sample predictive accuracy - relative LPL



GDP per capita growth


Industrial production



USD/GBP exchange rate


Employment


S\&P 500 excess returns


- BNN-NS BNN-BP - BART


## Non-linearities in the sample



## Effective number of neurons $Q^{*}$



## Relationship in-sample fit and out-of-sample predictability



## Closing remarks

- We have developed a non-parametric regression model based on Bayesian Neural Networks
- Our framework allows to remain agnostic on the form of the network
- We use popular techniques from the Bayesian literature to determine the adequate network structure
- In a broad set of macro and finance applications we show the superior forecasting performance of our approach


## References I

Agostinelli, F., M. Hoffman, P. Sadowski, AND P. Baldi (2014): "Learning activation functions to improve deep neural networks," arXiv preprint arXiv:1412.6830.
Barro, R. J., AND J.-W. Lee (1994): "Sources of economic growth," in Carnegie-Rochester conference series on public policy, vol. 40, pp. 1-46. Elsevier.
Belmonte, M. A., G. Koop, AND D. Korobilis (2014): "Hierarchical shrinkage in time-varying parameter models," Journal of Forecasting, 33(1), 80-94.
Bhattacharya, A., AND D. B. Dunson (2011): "Sparse Bayesian infinite factor models," Biometrika, pp. 291-306.
Blundell, C., J. Cornebise, K. Kavukcuoglu, and D. Wierstra (2015): "Weight uncertainty in neural network," in International Conference on Machine Learning, pp. 1613-1622. PMLR.

Carvalho, C. M., N. G. Polson, AND J. G. Scott (2010): "The horseshoe estimator for sparse signals," Biometrika, 97(2), 465-480.
Chipman, H. A., E. I. George, And R. E. McCulloch (2010): "BART: Bayesian additive regression trees," The Annals of Applied Statistics, 4(1), 266-298.

## References II

Giannone, D., M. Lenza, AND G. E. Primiceri (2021): "Economic predictions with big data: The illusion of sparsity," Econometrica, 89(5), 2409-2437.
Griffin, J. E., AND P. J. Brown (2013): "Some priors for sparse regression modelling," Bayesian Analysis, 8(3), 691-702.
Hoffman, M. D., A. Gelman, et AL. (2014): "The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo.," Journal of Machine Learning Research, 15(1), 1593-1623.
Hornik, K., M. Stinchcombe, AND H. White (1989): "Multilayer feedforward networks are universal approximators," Neural Networks, 2(5), 359-366.
Huber, F., G. Koop, L. Onorante, M. Pfarrhofer, and J. Schreiner (2020): "Nowcasting in a pandemic using non-parametric mixed frequency VARs," Journal of Econometrics, (forthcoming).
Huber, F., and M. Pfarrhofer (2021): "Dynamic shrinkage in time-varying parameter stochastic volatility in mean models," Journal of Applied Econometrics, 36(2), 262-270.
Huber, F., AND L. Rossini (2022): "Inference in Bayesian additive vector autoregressive tree models," The Annals of Applied Statistics, 16(1), 104-123.

## References III

KaRlik, B., AND A. V. Olgac (2011): "Performance analysis of various activation functions in generalized MLP architectures of neural networks," International Journal of Artificial Intelligence and Expert Systems, 1(4), 111-122.
Kastner, G., AND S. Frühwirth-Schnatter (2014): "Ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC estimation of stochastic volatility models," Computational Statistics \& Data Analysis, 76, 408-423.
MAKALIC, E., AND D. F. Schmidt (2015): "A simple sampler for the horseshoe estimator," IEEE Signal Processing Letters, 23(1), 179-182.
MAKRIDAKIS S., Spiliotis E., A. V. (2018): "Statistical and machine learning forecasting methods: Concerns and ways forward," PLoS One, 13.

McCracken, M. W., AND S. NG (2016): "FRED-MD: A monthly database for macroeconomic research," Journal of Business \& Economic Statistics, 34(4), 574-589.
Neal, R. M., et Al. (2011): "MCMC using Hamiltonian dynamics," Handbook of Markov Chain Monte Carlo, 2(11), 2.

Rossi, B. (2013): "Exchange rate predictability," Journal of Economic Literature, 51(4), 1063-1119.

## References IV

WELCh, I., AND A. Goyal (2008): "A comprehensive look at the empirical performance of equity premium prediction," The Review of Financial Studies, 21(4), 1455-1508.

Wright, J. H. (2008): "Bayesian model averaging and exchange rate forecasts," Journal of Econometrics, 146(2), 329-341.

## APPENDIX

Bayesian neural learning of non-linearities in macroeconomics and finance

## Deep BNN

A deep BNN with $L$ hidden layers implies that:

$$
\begin{aligned}
f\left(\boldsymbol{x}_{t}\right) & \approx\left(h^{(L)}\left(\boldsymbol{\kappa}^{(L)^{\prime}} \hat{\boldsymbol{x}}_{t}^{(L-1)}+\boldsymbol{\zeta}^{(L)}\right)\right) \boldsymbol{\beta}, \quad \text { with } \\
\hat{\boldsymbol{x}}_{t}^{(I)} & =h^{(I)}\left(\boldsymbol{\kappa}^{(I)^{\prime}} \hat{\boldsymbol{x}}_{t}^{(I-1)}+\boldsymbol{\zeta}^{(I)}\right), \quad \text { for } 1 \leq I \leq L-1 .
\end{aligned}
$$

- $I=\{1, \ldots, L\}$ denotes the number of layers
- $h^{(I)}(\bullet)$ refers to a layer-specific (nonlinear) activation function common to all neurons $\boldsymbol{\kappa}^{(/)}$and bias terms $\boldsymbol{\zeta}^{(/)}$in the respective layer
$-\hat{\boldsymbol{x}}_{t}^{(0)}=\hat{\boldsymbol{x}}_{t}$ for the first layer (i.e., $I=1$ )


## Results for the Deep BNN

| Application |  | Deep BNN |  |
| :--- | :--- | :---: | :---: |
| with 3 hidden layers |  |  |  |
| Macro A | Inflation <br> Industrial production <br> Employment |  |  |
| Macro B | 0.00 <br> -0.02 <br> 0.03 |  |  |
| Macro C | -0.38 |  |  |
| Finance | 0.02 |  |  |

Note: The table shows log predictive likelihoods (LPLs) relative to the best performing shallow BNN. Results are averaged across the hold-out.

## Simulation study co back

$$
\begin{aligned}
y_{t} & =f\left(\boldsymbol{x}_{t}^{\prime} \boldsymbol{\kappa}_{\text {true }}\right)^{\prime} \boldsymbol{\beta}_{\text {true }}+v_{t}, \quad v_{t} \sim \mathcal{N}\left(0, \sigma_{t, \text { true }}^{2}\right), \quad \boldsymbol{\kappa}_{\text {true }}=\boldsymbol{I}_{K}, \\
x_{j t} & \sim \mathcal{N}(0,1), \text { for } j=1, \ldots, K \text { and } K \in[30,60],
\end{aligned}
$$

- $\beta_{j, \text { true }} \sim N\left(0, c^{2}\right)$
- sparse DGP with $10 \%$ active neurons: $c=0.5^{2}$
- dense DGP with $90 \%$ active neurons: $c=0.04$
- 2 different variance specifications:
- homoskedastic: $\sigma_{t, \text { true }}^{2}=0.1$ for all t
- heteroskedastic: $\sigma_{t, \text { true }}^{2}=0.1 \exp \left(\eta_{t}\right), \eta_{t} \sim \mathcal{N}\left(0,0.1^{2}\right)$
- $f: R^{K} \mapsto R^{Q}$ defined as
- linear
- a shallow neural network with a single hidden layer and a randomly selected activation function for each neuron


## Synthetic: Forecast performance for 100 hold-out observations (estimated with SV)

| K | Sparsity | Noise | Non-linear DGP |  |  | Linear DGP |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | BNN | BNN-NS | Linear model | BNN | BNN-NS | Linear model |
| 30 | Dense | hetero | 1.00 | 0.93 | 0.51 | 1.01 | 1.01 | 0.43 |
|  |  | homo | 0.99 | 0.86 | 0.41 | 1.02 | 1.02 | 0.32 |
|  | Sparse | hetero | 0.99 | 0.80 | 0.98 | 0.99 | 0.99 | 0.49 |
|  |  | homo | 0.98 | 0.72 | 1.00 | 1.01 | 1.01 | 0.35 |
| 60 | Dense | hetero | 1.00 | 0.89 | 0.61 | 1.00 | 1.00 | 0.42 |
|  |  | homo | 1.01 | 0.84 | 0.54 | 1.01 | 1.01 | 0.33 |
|  | Sparse | hetero | 0.99 | 0.92 | 1.48 | 0.98 | 0.99 | 0.51 |
|  |  | homo | 1.00 | 0.94 | 1.37 | 1.01 | 1.00 | 0.38 |

Note: The table shows root mean squared errors (RMSEs) relative to the benchmark linear model. In bold we mark the best performing model for each case. The grey shaded area gives the actual RMSE scores of the benchmark. Results are averaged across the hold-out.

## Synthetic: Forecast performance for 100 hold-out observations (estimated without SV)

| K | Sparsity | Noise | Non-linear DGP |  |  | Linear DGP |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | BNN | BNN-NS | Linear model | BNN | BNN-NS | Linear model |
| 30 | Dense | hetero | 1.01 | 0.94 | 0.51 | 1.02 | 1.02 | 0.43 |
|  |  | homo | 0.98 | 0.84 | 0.41 | 1.02 | 1.02 | 0.32 |
|  | Sparse | hetero | 1.01 | 0.84 | 0.98 | 1.02 | 1.01 | 0.49 |
|  |  | homo | 1.02 | 0.72 | 1.01 | 1.00 | 1.00 | 0.35 |
| 60 | Dense | hetero | 1.01 | 0.87 | 0.61 | 1.02 | 1.02 | 0.43 |
|  |  | homo | 1.01 | 0.85 | 0.53 | 1.01 | 1.01 | 0.33 |
|  | Sparse | hetero | 1.00 | 0.93 | 1.48 | 1.01 | 1.01 | 0.51 |
|  |  | homo | 1.01 | 0.95 | 1.38 | 1.01 | 1.01 | 0.38 |

Note: The table shows root mean squared errors (RMSEs) relative to the benchmark linear model. In bold we mark the best performing model for each case. The grey shaded area gives the actual RMSE scores of the benchmark. Results are averaged across the hold-out.

## Macro A: Forecast performance across 252 hold-out observations co back

| Covariates | Model |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | BART | BNN | BNN-NS | BNN-BP | Linear model |
| Inflation |  |  |  |  |  |
| Large | 0.97 | $0.94{ }^{* * *}$ | $0.94{ }^{* * *}$ | 1.03 | $0.94{ }^{* * *}$ |
|  | (-0.03) | (0.09***) | (0.08***) | (-0.12***) | (0.08***) |
| PCA | 1.07** | 1.00 | 1.01 | 1.04** | 1.16 |
|  | $\left(-0.20{ }^{* * *}\right)$ | (0.00) | $\left(-0.04{ }^{* * *}\right)$ | (-0.11***) | (-1.45) |
| Industrial production |  |  |  |  |  |
| Large | 0.88 | 0.97*** | $1.02^{* *}$ | 0.93 | 0.98*** |
|  | (0.08) | (0.14***) | (0.17**) | (-0.41*) | (0.12***) |
| PCA | 0.89 | 1.00 | 1.01 | 0.94 | 1.75 |
|  | (-0.01) | (0.05**) | (0.07*) | (-0.46) | (-1.33) |
| Employment |  |  |  |  |  |
| Large | 1.04 | 1.00 |  | 1.01 |  |
|  | (0.11) | (0.14**) | (0.14) | (-0.87) | (0.10*) |
| PCA | 1.03 | 0.99 | 1.00 | 1.01 | 3.50 |
|  | (-0.19*) | (0.07) | (-0.01) | (-0.59) | (-1.88) |

[^0]
## Macro B: Forecast performance across 45 hold-out countries and 100 replications

| Covariates | Model |  |  |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: |
|  | BART | BNN | BNN-NS | BNN-BP | Linear model |
| Kitchen sink | $0.95^{* * *}$ | 1.01 | 1.01 | $1.04^{* * *}$ | 5.24 |
|  | $\left(0.06^{* * *}\right)$ | $\left(0.82^{* * *}\right)$ | $\left(0.95^{* * *}\right)$ | $\left(-0.18^{* * *}\right)$ | $(-4.40)$ |

Note: The table shows root mean squared errors (RMSEs), and average log predictive likelihoods (LPLs) in parentheses, relative to the linear benchmark. In bold we mark the best performing model for each case. The grey shaded area gives the actual RMSE and LPL scores of our benchmark (linear model). Results are averaged across the hold-out.

## Macro C: Forecast performance across 80 hold-out observations co back

| Covariates | Multivariate models |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | BART | BNN | BNN-NS | BNN-BP | Linear model |
| one-quarter ahead |  |  |  |  |  |
| All fundamentals | $\begin{gathered} 1.03 \\ (-0.02) \end{gathered}$ | $\begin{gathered} 1.02 \\ (0.01) \end{gathered}$ | $\begin{gathered} 1.02 \\ (0.02) \end{gathered}$ | $\begin{gathered} 1.03 \\ \left(-0.09^{*}\right) \end{gathered}$ | $\begin{gathered} 1.01 \\ (-0.01) \end{gathered}$ |
| Kitchen sink | $\begin{gathered} 1.04 \\ (-0.01) \end{gathered}$ | $\begin{gathered} 0.97 \\ (0.03) \end{gathered}$ | $\begin{gathered} 0.99 \\ (0.03) \end{gathered}$ | $\begin{gathered} 1.02 \\ (-0.10) \end{gathered}$ | $\begin{gathered} 0.95 \\ (-1.29) \end{gathered}$ |
| one-year-ahead |  |  |  |  |  |
| All fundamentals | $\begin{gathered} 1.04 \\ (0.00) \end{gathered}$ | $\begin{gathered} 1.00 \\ (0.02) \end{gathered}$ | $\begin{gathered} 1.00 \\ (0.00) \end{gathered}$ | $\begin{gathered} 1.01 \\ \left(-0.09^{\star \star}\right) \end{gathered}$ | $\begin{gathered} 1.00 \\ (-0.01) \end{gathered}$ |
| Kitchen sink | $\begin{gathered} 1.03 \\ (0.06) \end{gathered}$ | $\begin{gathered} 1.01 \\ (0.01) \end{gathered}$ | $\begin{gathered} 1.02 \\ (0.01) \end{gathered}$ | $\begin{aligned} & 0.98^{\star \star} \\ & (-0.09) \end{aligned}$ | $\begin{gathered} 0.96 \\ (-1.29) \end{gathered}$ |

Finance: Forecast performance across 56 hold-out observations comat

| Covariates | Multivariate models |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | BART | BNN | BNN-NS | BNN-BP | Linear model |
| Kitchen sink | $\begin{gathered} 1.00^{*} \\ (0.01) \end{gathered}$ | $\begin{gathered} 1.00 \\ (0.00) \end{gathered}$ | $\begin{gathered} 0.99 \\ (0.01) \end{gathered}$ | $\begin{gathered} 1.03 \\ (-0.04) \end{gathered}$ | $\begin{gathered} 1.07 \\ (-1.46) \end{gathered}$ |
|  | Univariate models |  |  |  |  |
|  | BART |  | BNN-NS | BNN-BP | Linear model |
| Dividend price ratio | $\begin{gathered} 1.01 \\ (-0.04) \end{gathered}$ |  | $\begin{gathered} 0.99 \\ (0.00) \end{gathered}$ | $\begin{gathered} 0.99 \\ (0.00) \end{gathered}$ | $\begin{gathered} 0.99 \\ (-0.01) \end{gathered}$ |
| Dividend yield | $\begin{gathered} 1.01 \\ (-0.02) \end{gathered}$ |  | $\begin{gathered} 0.99 \\ (0.00) \end{gathered}$ | $\begin{gathered} 1.01 \\ (-0.03) \end{gathered}$ | $\begin{gathered} 0.99 \\ (0.00) \end{gathered}$ |
| Inflation | $\begin{gathered} 1.02 \\ (-0.02) \end{gathered}$ |  | $\begin{gathered} 0.99 \\ (0.01) \end{gathered}$ | $\begin{gathered} 1.00 \\ (0.00) \end{gathered}$ | $\begin{gathered} 1.00 \\ (0.00) \end{gathered}$ |
| Term spread | $\begin{gathered} 1.00 \\ (-0.01) \end{gathered}$ |  | $\begin{gathered} 1.01 \\ (0.00) \end{gathered}$ | $\begin{gathered} 0.99 \\ (0.00) \end{gathered}$ | $\begin{gathered} 0.99 \\ (-0.01) \end{gathered}$ |

[^1] (linear model). Results are averaged across the hold-outs

Full conditional posterior distribution (1) comas

- Sample $\boldsymbol{\theta}=\left(\boldsymbol{\gamma}^{\prime}, \boldsymbol{\beta}^{\prime}\right)^{\prime}$

$$
\boldsymbol{\theta} \mid \cdot \sim \mathcal{N}\left(\overline{\boldsymbol{\theta}}, \overline{\boldsymbol{V}}_{\boldsymbol{\theta}}\right),
$$

with

$$
\begin{array}{r}
\overline{\boldsymbol{V}}_{\boldsymbol{\theta}}=\left(\tilde{\boldsymbol{x}}^{\prime} \boldsymbol{\Sigma}^{-1} \tilde{\boldsymbol{x}}+\underline{\boldsymbol{V}}_{\boldsymbol{\theta}}^{-1}\right)^{-1}, \\
\overline{\boldsymbol{\theta}}=\overline{\boldsymbol{V}}_{\boldsymbol{\theta}} \tilde{\boldsymbol{x}}^{\prime} \boldsymbol{\Sigma}^{-1} \boldsymbol{y} .
\end{array}
$$

- $\tilde{\boldsymbol{x}}=\left(\tilde{\boldsymbol{x}}_{1}^{\prime}, \ldots, \tilde{\boldsymbol{x}}_{T}^{\prime}\right)^{\prime}$ as a $(K+Q) \times T$ matrix of neurons with element $\tilde{\boldsymbol{x}}_{t}=\left(\boldsymbol{x}_{t}^{\prime}, h_{1}\left(\boldsymbol{x}_{t}^{\prime} \kappa_{1}+\zeta_{1}\right), \ldots, h_{Q}\left(\boldsymbol{x}_{t}^{\prime} \kappa_{Q}+\zeta_{Q}\right)\right)^{\prime}$
- $\boldsymbol{\Sigma}=\operatorname{diag}\left(\sigma_{1}^{2}, \ldots, \sigma_{T}^{2}\right)$ as a $T \times T$ matrix capturing the variances
- $\underline{\boldsymbol{V}}_{\boldsymbol{\theta}}=\operatorname{diag}\left(\phi_{\gamma}^{-1}, \phi_{\boldsymbol{\beta}}^{-1}\right)$ where $\phi_{\gamma}^{-1}=\left(\phi_{\gamma_{1}}^{-1}, \ldots, \phi_{\gamma_{K}}^{-1}\right)^{\prime}$ as the $K$ prior variances for the constant coefficients and $\phi_{\boldsymbol{\beta}}^{-1}=\left(\phi_{\beta_{1}}^{-1}, \ldots, \phi_{\beta_{Q}}^{-1}\right)^{\prime}$ for the non-linear coefficients

Full conditional posterior distribution (2) comas

- The prior on $\gamma$ is Normal of the form:

$$
\gamma_{j} \sim \mathcal{N}\left(0, \phi_{\gamma_{j}}^{-1}\right), \quad \phi_{\gamma_{j}}^{-1}=\lambda_{\gamma}^{2} \varphi_{\gamma_{j}}^{2}, \quad \text { for } j=1, \ldots, K .
$$

with the global and local shrinkage parameters, $\lambda_{\gamma}^{2}$ and $\varphi_{\gamma_{j}}^{2}$ (HS prior in hierarchical representation of Makalic and Schmidt (2015)):

$$
\begin{array}{r}
\varphi_{\gamma_{j}}^{2} \left\lvert\, \bullet \sim \mathcal{G}^{-1}\left(1, c_{\gamma_{j}}^{-1}+\frac{\gamma_{j}^{2}}{2 \lambda_{\gamma}^{2}}\right)\right., \\
\lambda_{\gamma}^{2} \left\lvert\, \bullet \sim \mathcal{G}^{-1}\left(\frac{K+1}{2}, d_{\gamma}^{-1}+\sum_{j=1}^{K} \frac{\gamma_{j}^{2}}{2 \varphi_{\gamma_{j}}^{2}}\right)\right., \\
c_{\gamma_{j}} \mid \bullet \sim \mathcal{G}^{-1}\left(1,1+\varphi_{\gamma_{j}}^{-2}\right), \\
d_{\gamma} \mid \bullet \sim \mathcal{G}^{-1}\left(1,1+\lambda_{\gamma}^{-2}\right) .
\end{array}
$$

Full conditional posterior distribution (3) comat

- We sample the hyperparameters associated with the MGP prior on $\beta$ from inverse Gamma distributions:

$$
\begin{array}{r}
\delta_{1} \sim \mathcal{G}^{-1}\left(a_{1}+\frac{Q}{2}, 1+\frac{1}{2} \sum_{q=1}^{Q}\left(\phi_{\beta_{q}} \beta_{q}^{2}\right)\right), \\
\delta_{r} \sim \mathcal{G}^{-1}\left(a_{2}+\frac{Q-r-1}{2}, 1+\frac{1}{2} \sum_{q=1}^{Q}\left(\phi_{\beta_{q}} \beta_{q}^{2}\right)\right) .
\end{array}
$$

## Full conditional posterior distribution (4) comat

- Hamiltonian Monte Carlo (HMC) within Gibbs step to draw $\kappa_{q}(q=1, \ldots, Q)$ : Let $\boldsymbol{r}_{q}$ denote an auxiliary moment variable, where $\boldsymbol{r}_{q} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$ and $\mathcal{L}\left(\boldsymbol{\kappa}_{q}\right)=\log p\left(\boldsymbol{\kappa}_{q} \mid \bullet\right)$ the log conditional posterior density of $\boldsymbol{\kappa}_{q}$. The fictitious Hamiltonian system of the conditional posterior density of $\kappa_{q}$ is then given by

$$
\mathcal{H}\left(\boldsymbol{\kappa}_{q}, \boldsymbol{r}_{q}\right)=-\mathcal{L}\left(\boldsymbol{\kappa}_{q}\right)+\frac{1}{2} \boldsymbol{r}_{q}^{\prime} \boldsymbol{r}_{q} .
$$

We simulate the Hamiltonian dynamics via the leapfrog integrator with proposals:

$$
\begin{aligned}
\boldsymbol{r}_{q}^{* *} & =\boldsymbol{r}_{q}^{(a)}+\frac{\epsilon}{2} \nabla_{\boldsymbol{\kappa}_{q}} \mathcal{L}\left(\boldsymbol{\kappa}_{q}^{(a)}\right), \\
\boldsymbol{\kappa}_{q}^{*} & =\boldsymbol{\kappa}_{q}^{(a)}+\epsilon \boldsymbol{r}_{q}^{* *}, \\
\boldsymbol{r}_{q}^{*} & =\boldsymbol{r}_{q}^{* *}+\frac{\epsilon}{2} \nabla_{\boldsymbol{\kappa}_{q}} \mathcal{L}\left(\boldsymbol{\kappa}_{q}^{*}\right),
\end{aligned}
$$

where $\boldsymbol{\kappa}_{q}^{(a)}$ and $\boldsymbol{r}_{q}^{(a)}$ denote the previously accepted values, $\nabla_{\kappa_{q}} \mathcal{L}\left(\boldsymbol{\kappa}_{q}\right)$ is the gradient of the log conditional posterior, $\epsilon$ is the discrete step size to generate a full-step proposal for $\boldsymbol{\kappa}_{q}$ (i.e., $\boldsymbol{\kappa}_{q}^{*}$ ) and half-step updates for the momentum $\boldsymbol{r}_{q}\left(\boldsymbol{r}_{q}^{* *}\right.$ for the first half-step update and $\boldsymbol{r}_{q}^{*}$ for the final proposal).

## Full conditional posterior distribution (5) comat

- We repeat the leapfrog method in $n=1, \ldots, N$ steps. To tune $N$ and $\epsilon$ we use the No U-Turn Sampler (NUTS) as in Hoffman et al. (2014).

Finally, we evaluate the proposed and previously accepted values by means of a Metropolis accept/reject step and determine the acceptance probability $\eta_{q}$ for proposed $\kappa_{q}^{*}$ :

$$
\eta_{q}=\min \left(1, \frac{\exp \left(\mathcal{L}\left(\boldsymbol{\kappa}_{q}^{*}\right)-\frac{1}{2} \boldsymbol{r}_{q}^{\prime *} \boldsymbol{r}_{q}^{*}\right)}{\exp \left(\mathcal{L}\left(\boldsymbol{\kappa}_{q}^{(a)}\right)-\frac{1}{2} \boldsymbol{r}_{q}^{\prime(a)} \boldsymbol{r}_{q}^{(a)}\right)}\right) .
$$

This step remains conceptually similar for the deep BNN. In this case, we only need to iterate through all layers and also take into account the fact that the composite function of $L$ layers has different implications on the gradients of the layer-specific neurons.

Full conditional posterior distribution (6)

- Column-wise horseshoe prior on the elements of $\kappa_{q}$ :

$$
\begin{aligned}
\varphi_{\kappa_{j q}}^{2} \left\lvert\, \bullet \sim \mathcal{G}^{-1}\left(1, c_{\kappa_{j} q}^{-1}+\frac{\kappa_{j q}^{2}}{2 \lambda_{\kappa_{q}}^{2}}\right)\right., \\
\lambda_{\kappa_{q}}^{2} \left\lvert\, \bullet \sim \mathcal{G}^{-1}\left(\frac{K+1}{2}, d_{\kappa_{q}}^{-1}+\sum_{j=1}^{K} \frac{\kappa_{j q}^{2}}{2 \varphi_{\kappa_{j q}}^{2}}\right)\right., \\
c_{\kappa_{j q}}^{2} \mid \bullet \sim \mathcal{G}^{-1}\left(1,1+\varphi_{\kappa_{j q}}^{-2}\right), \\
d_{\kappa_{q}} \mid \bullet \sim \mathcal{G}^{-1}\left(1,1+\lambda_{\kappa_{q}}^{-2}\right) .
\end{aligned}
$$

## Full conditional posterior distribution (7) œ.

- Choose activation function $h_{q}$ by drawing $\delta_{q}$ from a multinomial distribution:

$$
\operatorname{Pr}\left(\delta_{q}=m \mid \bullet\right) \propto \omega_{q m} \times \exp \left\{-\frac{1}{2}\left(\left(\hat{\boldsymbol{y}}_{q}-\boldsymbol{\mu}_{q}^{(m)}\right)^{\prime} \boldsymbol{\Sigma}^{-1}\left(\hat{\boldsymbol{y}}_{q}-\boldsymbol{\mu}_{q}^{(m)}\right)\right)\right\}, \quad \text { for } m=1, \ldots, 4,
$$

where $\boldsymbol{\mu}_{q}^{(m)}=\left(\mu_{1 q}^{(m)}, \ldots, \mu_{T q}^{(m)}\right)^{\prime}$ with elements $\mu_{t q}^{(m)}=\beta_{q} h_{q}^{(m)}\left(\boldsymbol{x}_{t}^{\prime} \kappa_{q}+\zeta_{q}\right)$.

## Bayesian Neural Network by backpropagation comas

- Variational inference scheme for learning the posterior distribution on the weights of a neural network (Blundell et al., 2015)
- Maximize the log-likelihood subject to a Kullback-Leibler complexity term on the parameters (with reparameterization trick and stochastic gradient descent)
- Prior on the weights as a scale mixture of two Gaussian densities with zero mean but differing variances ( $\sigma_{1}^{2}=3$ and $\sigma_{2}^{2}=0.0025$ )
- Cross validation exercise:
- Macro A: time series split with 24 months
- Macro B: random split in 20 replications
- Macro C: time series split with 12 quarters
- Finance: time series split with 10 years
- 1000 epochs
- MSE loss function
- ADAM optimizer
- learning rate 0.014


## Bayesian additive regression trees

- Approximate $f$ using Bayesian additive regression trees (Chipman et al., 2010)
- Sum over a number of $Z$ regression trees $\left(g_{z}\right)$ :

$$
f\left(\boldsymbol{x}_{t}\right) \approx \sum_{z=1}^{z} g_{z}\left(\boldsymbol{x}_{t} \mid \mathcal{T}_{z}, \boldsymbol{\rho}_{z}\right)
$$

- Tree structure $\mathcal{T}_{z}$
- Terminal node parameter $\rho_{z}$
- $Z=250$
- Prior on the tree structure built upon a tree-generating stochastic process
- Determining the probability that a given node is nonterminal
- Selection of variables used in a splitting rule (to spawn left and right children nodes)
- Terminal node parameter: conjugate Gaussian prior distribution with data-based prior variance


## BART illustration coback




[^0]:    Note: The table shows root mean squared errors (RMSEs), and average log predictive likelihoods (LPLs) in parentheses, relative to the linear benchmark. In bold we mark the best performing model for each case. The grey shaded area gives the actual RMSE and LPL scores of our benchmark (linear model). Results are averaged across the hold-out.

[^1]:    Note: The table shows RMSEs with average LPLs in parentheses. The grey shaded area gives the actual RMSE and LPL scores of our benchmark

