Machine learning for financial crisis prediction and the construction of a coherent narrative

Andreas Joseph

Bank of England[1], King's College London (DAFM)

29. PKES Annual Workshop

29. May 2019, London

[1] Joint work with Kristina Bluwstein, Marcus Buckmann, Miao Kang, Sujit Kapadia and Özgür Şimşek. Disclaimer: The expressed views are my own and not necessarily those of the Bank of England (BoE) or its committees. All errors are ours.

Main findings

- State-of-the-art **machine learning models outperform benchmark logit** in out-of-sample prediction tests $\Rightarrow$ higher model confidence
- Main model drivers: Global and domestic **credit and the slope of the yield curve** $\Rightarrow$ black box opened
- **Yield curve new and interesting:** bank profitability, financial fragility and risk perceptions associated with a low/negative slope.
- Overall, results point to a **Minsky-type narrative for the built-up of financial crises** (in line with PK perspective):
  1. Shared low risk perception
  2. Optimism: high investment and credit growth
  3. Disconnect between real and financial factors (slowing consumption)

What is Machine Learning (ML)?

- Statistical toolbox of **non-linear** models mostly originating from computer science ("PC-like"), **statistical learning** probably better name
- **Supervised**: Universal approximators (ANN, SVM, random forests, etc.), focus on prediction, less (causal) inference
- **Unsupervised**: General clustering techniques (k-means, autoencoders, GANs, etc.)
- **Reinforcement**: "agent-based" modelling to navigate general complex environments (e.g. Alpha-Go). Maybe the best path AGI.

Everything today is about <u>supervised learning</u>, i.e. minimising

$$\mathbb{E}_x\left[||y - \hat{f}(\theta)\|_p\right]$$

Artificial Neural Network (ANN)

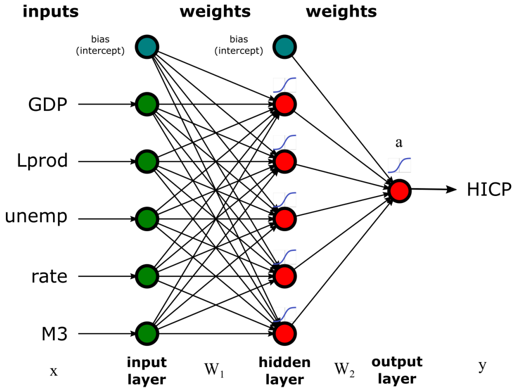

$\hat{y} = a_2(\cdot a_1(x \cdot W_1) \cdot W_2)$, see SWP 674

a: activation function

- **"Origin of the black box:"** Size of weight matrices W's not pre-determined (non-parametric model).
- But source of good performance (key AI driver)
- Shapley regressions (SWP 784, Joseph, 2019) provide rigorous framework to address above problem.

Pro's & Con's of ML relative to econometric approach

Advantages

- Often higher accuracy ⇒ higher confidence in predictions
- Learn *unknown* functional form from data ⇒ lower risk of misspecification
- Return richer information set ⇒ (potentially) better informed decisions

Disadvantages

- Higher complexity ⇒ more difficult to understand and to communicate ("black box critique")
- Less guarantees (convergence, overfitting, etc.) ⇒ more robustness tests needed
- Often larger data requirement ⇒ more constraints on applicability

Jordà-Schularick-Taylor Macrohistory Database

Observations

- 17 developed countries, annual data between 1870 and 2016
- 92 crisis, 2407 non-crisis observations
- 24 potential indicators

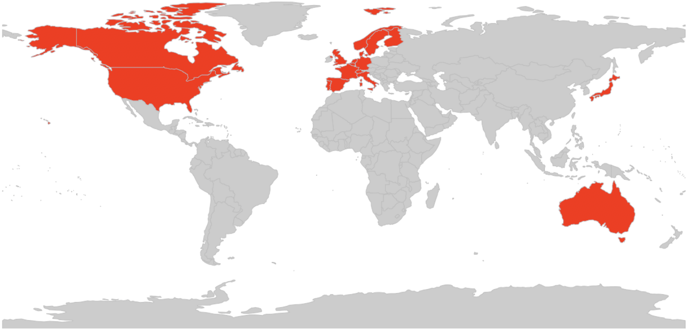

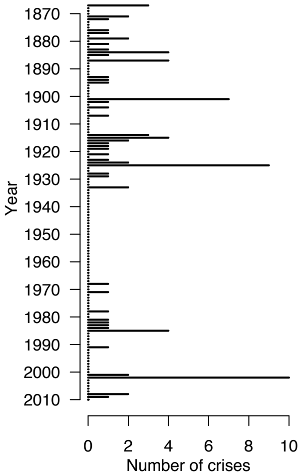

Jordà-Schularick-Taylor Macrohistory Database

Observations

- 17 developed countries, annual data between 1870 and 2016
- 92 crisis, 2407 non-crisis observations
- 24 potential indicators

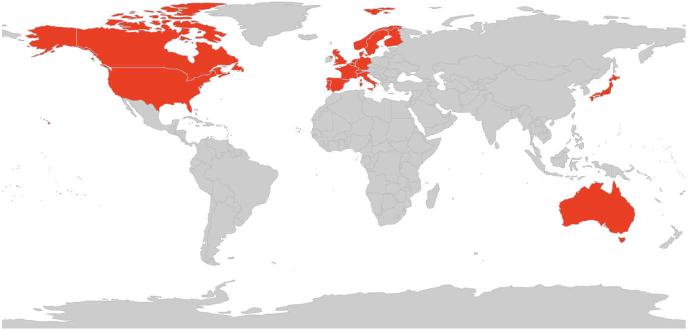

Subset of variables we use

- Non-financial credit
- Rates, yield curve
- Debt service ratio
- Current account balance
- Stock Prices
- CPI
- Consumption
- Investment
- Broad money
- Public debt

We capture $\geq$ 50% of world output most of the time

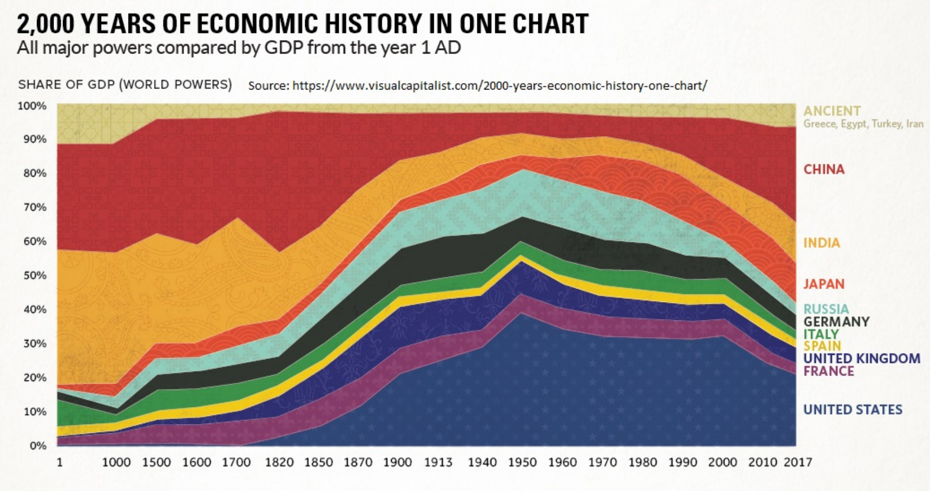

Problem dimensions

Transformation selection (e.g. Percentage changes vs. ratio changes)

$\downarrow$

Variable selection (e.g. Long- short term rate vs. long and short term rate)

$\downarrow$

Horizon of change (growth rates/ratio changes of 1,..,5 years)

$\downarrow$

Landing zone (Predict 1 or 1–2 years before a crisis)

$\downarrow$

Model selection (Horse race of machine learning models)

Baseline empirical approach

Data

- Make data stationary
  - 2-year ratio change: $\frac{credit_{1980}}{GDP_{1980}} - \frac{credit_{1978}}{GDP_{1978}}$
  - 2-year growth rates: $\frac{stocks_{1980} - stocks_{1978}}{stocks_{1978}}$
- Predict one and two years in advance of a crisis
- Exclude actual crisis observation and five following years (post-crisis bias)
- Exclude world wars and 1933–1938
- Global variables: Mean across all countries excluding country of interest

Modelling

- Nested cross-validation & forecasting evaluation
- Bootstrapped & averaged models (bagging)

Out-of-sample performance in the ROC space

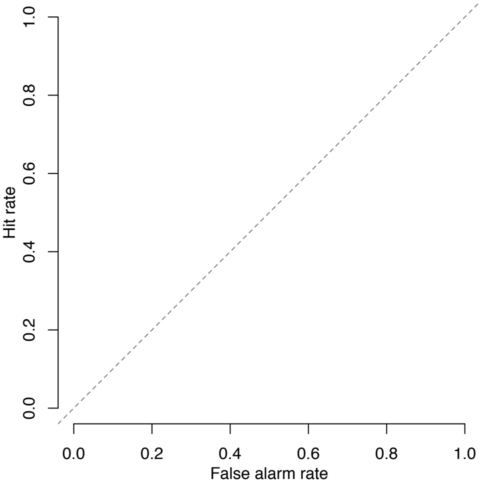

Linear baseline

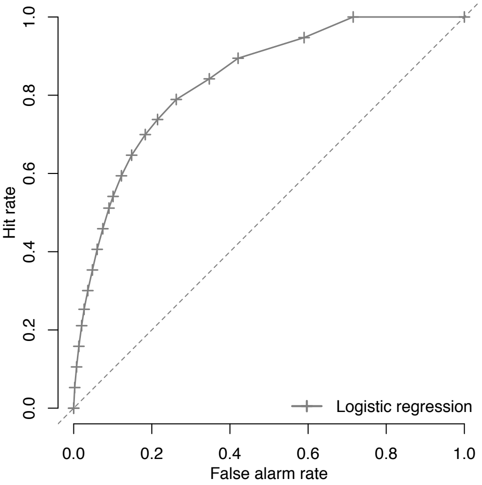

+ Decision trees

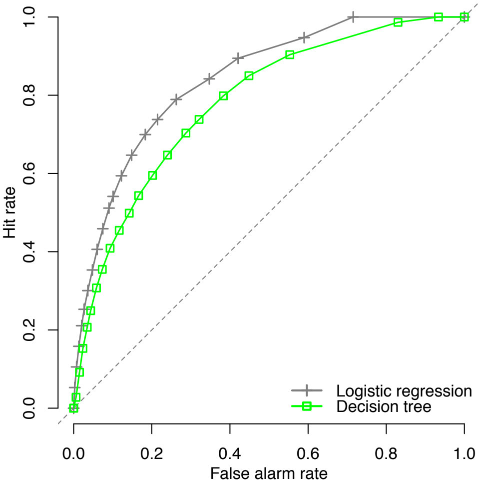

+ Neural network

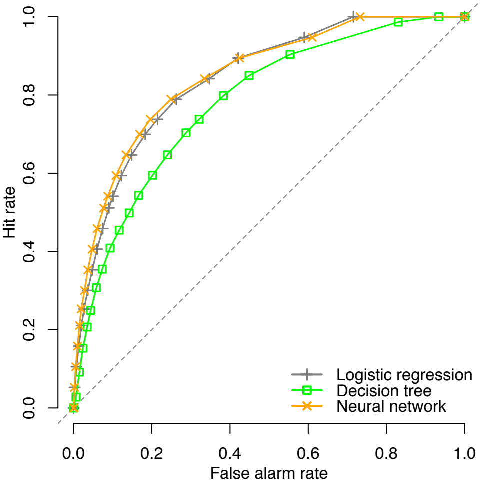

+ SVM

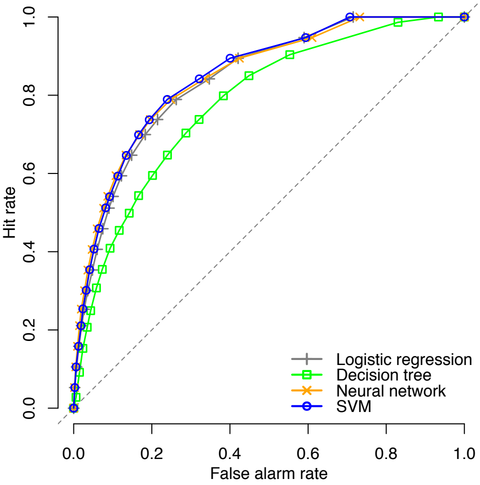

Hit rate
0.0
0.2
0.4
0.6
0.8
1.0
False alarm rate
0.0
0.2
0.4
0.6
0.8
1.0
Logistic regression
Decision tree
Neural network
SVM

+ Random forest

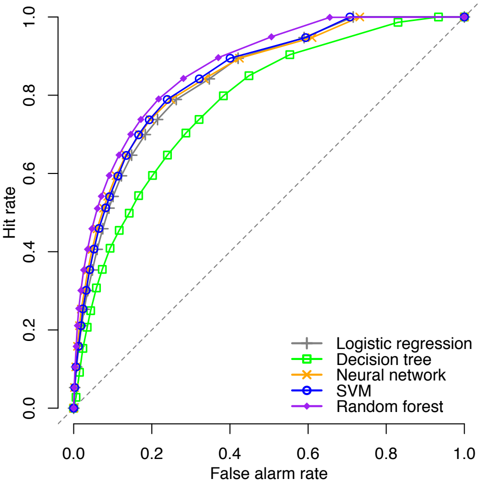

The winner is: Extremely randomized trees

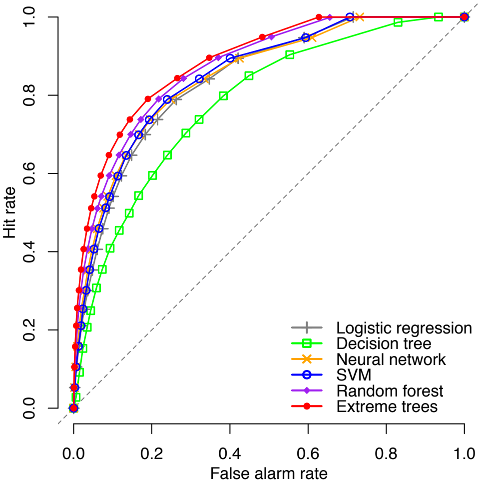

Area under the curve (AUC) performance

Extreme trees	**0.870**
Random forest	0.855
SVM	0.832
Neural net	0.829
Logistic regression	0.822
Decision tree	0.759

100 replications of 5-fold cross-validation.

Standard errors not shown but consistently below 0.002.

Detour: Shapley values in cooperative game theory

- How much does a player A contribute a collective payoff f obtained by a group of N? (Shapley, 1953).
- Observe payoff of the group with and without player A.
- Contribution depends on the other players in the game.
- All possible coalitions S need to be evaluated.

$$\phi_A = \sum_{S \subseteq N \setminus A} \frac{|S|!(|N| - |S| - 1)!}{|N|!} [f(S \cup \{A\} - f(S)] \tag{1}$$

Nice math properties but computationally complex

Intuitive example: stealing apples together

- Three siblings (strong [S], tall [T] & smart [M]) set off to nick some apples A (pay-off) from the neighbour's tree
- For each sibling, sum over marginal contribution to coalitions of one and two
- So, the Shapley value of the strong sibling is then:

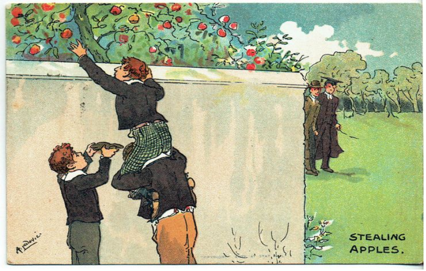

Source: 6oxgangsavenueedinburgh

$$\phi_S = \frac{1}{6}[A(T,S) - A(T)] + \frac{1}{6}[A(M,S) - A(M)] + \frac{1}{3}[A(T,M,S) - A(T,M)] \quad (2)$$

Shapley values applied to our problem

	Game Theory	Machine Learning
N	Players	Predictors
$\hat{f}/\hat{y}$	Collective payoff	Predicted value for one observation
S	Coalition	Predictors used for prediction
Source	Shapley (1953)	Strumbelj and Kononenko (2010)
		Lundberg and Lee (2017b)

Model Shapley decomposition based on (1): $\Phi^S\left(\hat{f}(x_{ik})\right) \equiv \phi_0 + \sum_{k=1}^{m} \phi_{ik}$.

See Strumbelj and Kononenko (2010); Lundberg and Lee (2017a) for details.

See also: bankunderground.co.uk/opening-the-machine-learning-black-box

Model explanations using Shapley decompositions: high agreement

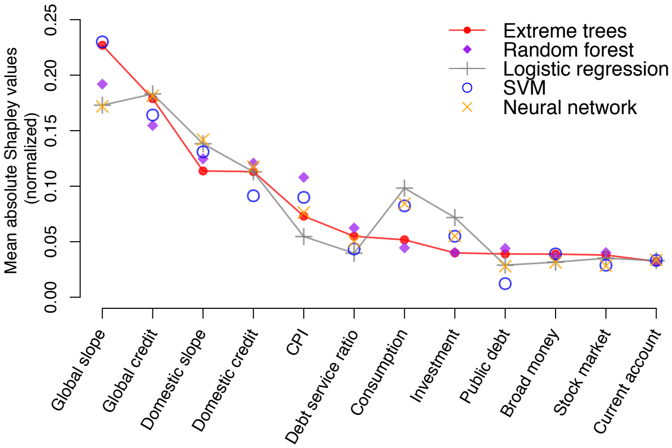

Key indicators:

- Domestic credit (Schularick and Taylor, 2012; Aikman et al., 2013)
- Global credit (Alessi and Detken, 2011; Cesa-Bianchi et al., 2018)
- Domestic slope (Babeckỳ et al., 2014; Joy et al., 2017)
- Global slope (new finding)

It's the short end of the yield curve

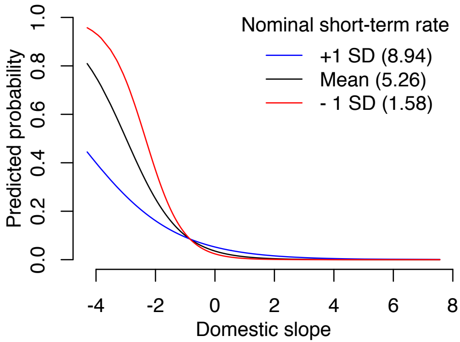

Effect of interaction of domestic slope of the yield curve and nominal short-term interest rate in logistic regression.

Non-linearity of extreme trees

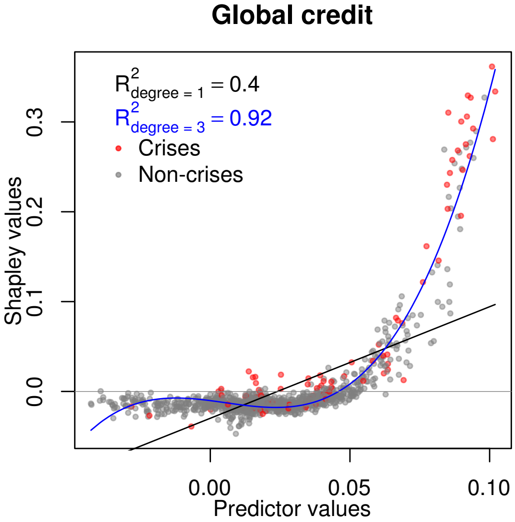

- ML models identify strong non-linearities
- Importantly, these are not known a priori
- Directions of associations match those in the linear model

Extreme trees model Shapley value decomposition

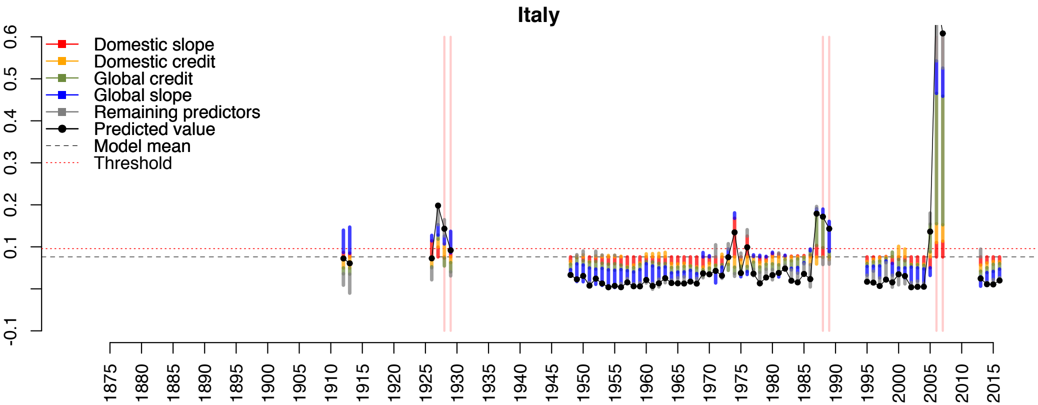

Extreme trees model Shapley value decomposition

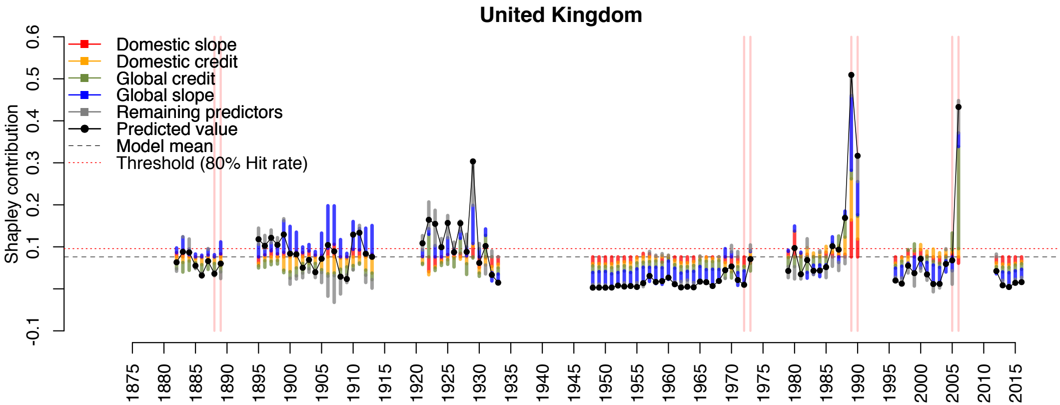

Shapley regression for econometric analysis (SWP 784)

$$\hat{y} \; = \; P[y_{crisis}|x] \; = \; Logit\big(\phi_0 + \hat{\beta}^S \, \Phi_{ML}(x)\big) \tag{3}$$

The Shapley values $\Phi_{ML}(x_k)$ are interpreted as ML model-based transformations of variable x_k.

(Shapley) regression table for extreme trees

Variable	Shapley regression Direction	Share	α-lvl	p	Logit regression Coeff.	α-lvl	p
Global slope	-	0.23	***	0.000	-0.61	***	0.000
Global credit	+	0.18	***	0.000	0.67	***	0.000
Domestic slope	-	0.11	***	0.000	-0.58	***	0.000
Domestic credit	+	0.11	***	0.000	0.43	***	0.002
CPI	-	0.07	***	0.002	-0.24		0.160
Debt service ratio	+	0.05		0.244	0.16		0.347
Consumption	-	0.05	**	0.027	-0.42	***	0.003
Investment	+	0.04	***	0.005	0.32	**	0.016
Public debt	-	0.04		0.295	-0.03		0.845
Broad money	+	0.04	*	0.081	0.04		0.817
Stock market	-	0.04	**	0.020	-0.13		0.451
Current account	-	0.03		0.296	-0.08		0.525

Extreme trees variable contributions for predicting financial crisis. Bootstrap clustered SE, α-level: *: 10%, **: 5%, ***: 1%, n-obs: 2499.

Other non-linearities: Interactions (e.g. slope and credit)

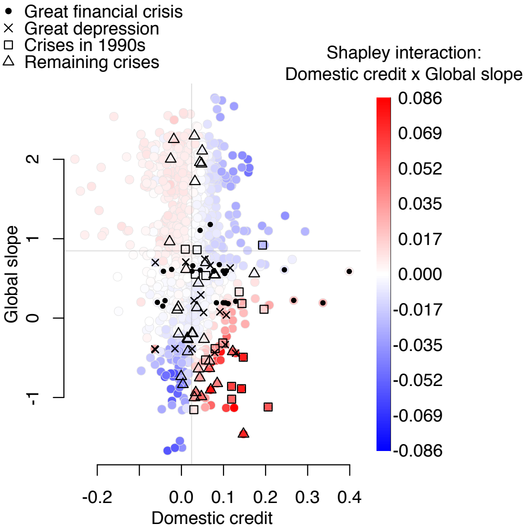

- High credit growth x negative slope of the yield curve increases predicted value
- Credit booms are more dangerous during or when expecting global economic slowdowns
- Additional important signals: Opposing signs from consumption and investment

Shapley regression results for single interaction terms with global variables

Interaction	Direction	Share	α-lvl	p-values
Global slope x Global credit	-	0.06	***	0.002
Global slope x Domestic slope	+	0.03		0.169
Global slope x Domestic credit	-	**0.07**	***	0.004
Global slope x Investment	-	0.04	***	0.000
Global slope x Consumption	+	0.03	*	0.058
Global slope x CPI	+	0.04	***	0.003
Global slope x Stock market	-	0.03		0.185
Global credit x Domestic credit	+	0.03	*	0.083
Global credit x Domestic slope	-	0.03	**	0.027
Global credit x Investment	+	0.02	**	0.036
Global credit x CPI	-	0.04	***	0.001
Global credit x Consumption	-	0.03	***	0.002
Global credit x Stock market	+	0.03	**	0.014

Extreme trees interaction terms, α-level: *: 10%, **: 5%, ***: 1%, n-obs: 2499.

Wrap-up

- **ML outperforms benchmark logit** in out-of-sample financial crisis prediction
- Main common model drivers: Global and domestic **credit and the slop of the yield curve**
- Interactions point to a **Minsky-type narrative for the built-up of a crisis**
- Shapley regressions provide a well-defined **framework for statistical inference on ML models**
- ML modelling results can be **communicated similar to standard regression** outputs

The End: THX - Q & A

Robustness checks (I)

Setups	Crises	Extreme trees	Random forest	Logit regression	SVM	Neural net	Decision tree
Baseline	93	0.84	0.83	0.80	0.79	0.79	0.73
TESTING TRANSFORMATIONS							
Growth rates only	93	0.78	0.77	0.74	0.71	0.72	0.68
Hamilton filter	87	0.82	0.83	0.79	0.78	0.80	0.75
*	87	0.84	0.83	0.80	0.77	0.78	0.76
ADDING VARIABLES							
Nominal rates	93	0.83	0.82	0.80	0.78	0.77	0.73
Real rates	93	0.82	0.82	0.80	0.78	0.79	0.75
Loans by sector	50	0.85	0.84	0.84	0.77	0.82	0.78
*	50	0.87	0.86	0.84	0.76	0.81	0.79
House prices	81	0.86	0.84	0.80	0.78	0.78	0.76
*	81	0.85	0.84	0.80	0.77	0.79	0.76

Robustness checks (II)

Setups	Crises	Extreme trees	Random forest	Logit regression	SVM	Neural net	Decision tree
Baseline	93	0.84	0.83	0.80	0.79	0.79	0.73
CHANGING THE HORIZON							
1 year	93	0.81	0.81	0.80	0.78	0.78	0.71
*	93	0.85	0.83	0.80	0.78	0.79	0.74
3 years	90	0.83	0.83	0.80	0.78	0.77	0.74
*	90	0.84	0.83	0.80	0.79	0.79	0.73
4 years	88	0.86	0.85	0.79	0.80	0.78	0.76
*	88	0.84	0.83	0.80	0.78	0.79	0.75
5 years	87	0.85	0.84	0.79	0.80	0.77	0.75
*	87	0.84	0.83	0.80	0.78	0.79	0.76
PREDICT ONE YEAR BEFORE CRISIS							
	48	0.85	0.81	0.81	0.79	0.80	0.72

Change of Shapley values over time

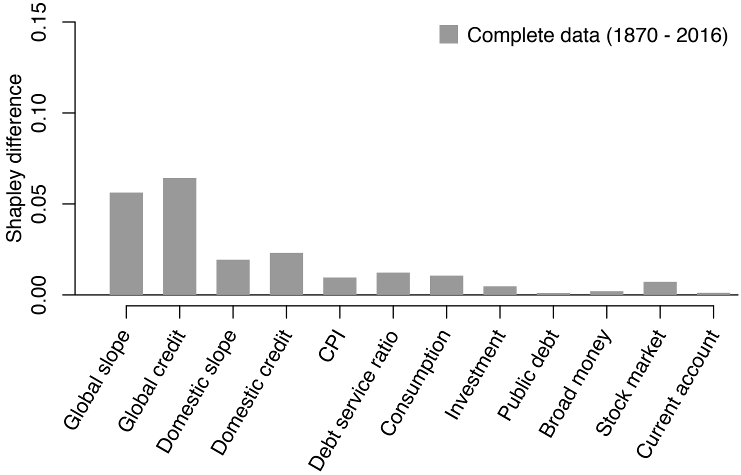

Change of Shapley values over time

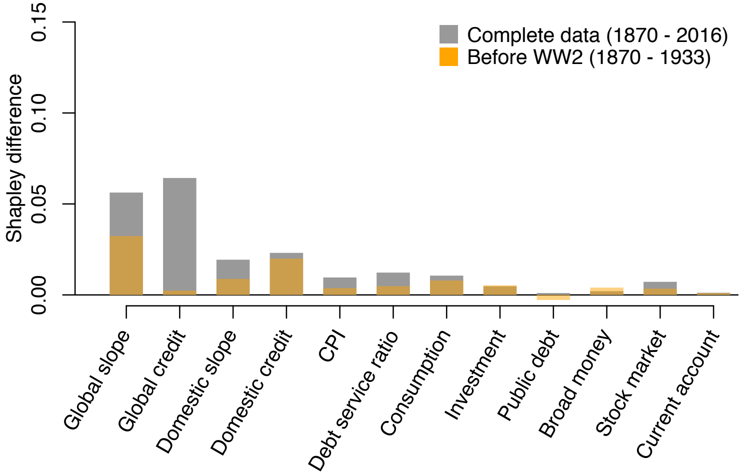

Change of Shapley values over time

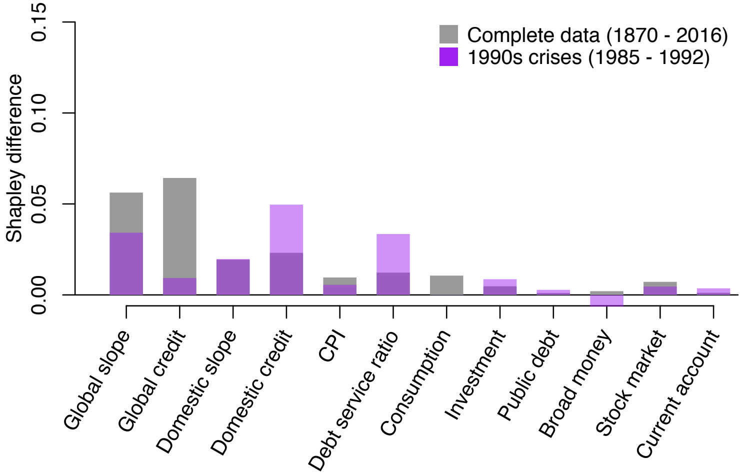

Change of Shapley values over time

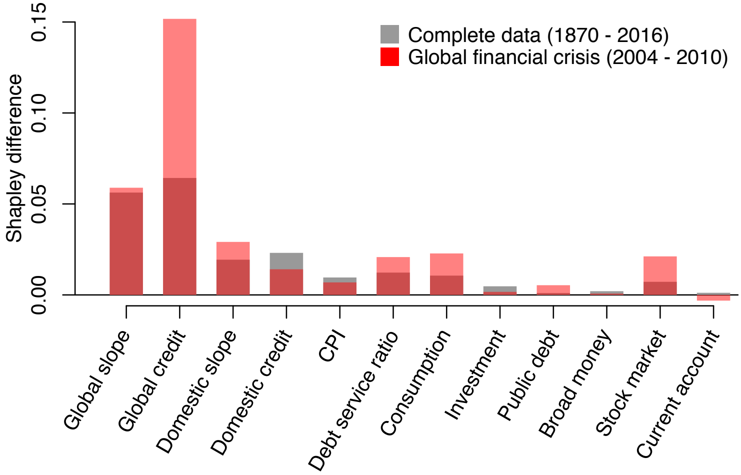

Related literature

1. **Conceptual**: Ethical, safety, privacy, <u>communication</u>

   [Lipton (2016); Miller (2017); Doshi-Velez and Kim (2017)]

2. **Technical (CS)**: Descriptive model decompositions

   [Ribeiro et al. (2016); Shrikumar et al. (2017); Lundberg and Lee (2017b)]

3. **Technical (econ)**: ML-based statistical inference

   [Chernozhukov et al. (2018); Wager and Athey (2015); Mullainathan and Spiess (2017)]

The current work addresses (1), builds on (2) and is complimentary to (3).

References I

Aikman, D., Haldane, A. G., and Nelson, B. D. (2013). Curbing the Credit Cycle. *The Economic Journal*, 125(585):1072–1109.

Alessi, L. and Detken, C. (2011). Quasi real time early warning indicators for costly asset price boom/bust cycles: A role for global liquidity. *European Journal of Political Economy*, 27(3):520–533.

Babeckỳ, J., Havranek, T., Mateju, J., Rusnák, M., Smidkova, K., and Vasicek, B. (2014). Banking, debt, and currency crises in developed countries: Stylized facts and early warning indicators. *Journal of Financial Stability*, 15:1–17.

Cesa-Bianchi, A., Martin, F. E., and Thwaites, G. (2018). Foreign booms, domestic busts: The global dimension of banking crises. *Journal of Financial Intermediation*.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68.

Doshi-Velez, F. and Kim, B. (2017). Towards A Rigorous Science of Interpretable Machine Learning. *ArXiv e-prints: 1702.08608*.

References II

Joy, M., Rusnák, M., Šmídková, K., and Vašíček, B. (2017). Banking and currency crises: Differential diagnostics for developed countries. *International Journal of Finance & Economics*, 22(1):44–67.

Lipton, Z. C. (2016). The mythos of model interpretability. *CoRR*, abs/1606.03490.

Lundberg, S. M. and Lee, S.-I. (2017a). A Unified Approach to Interpreting Model Predictions. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems 30*, pages 4765–4774. Curran Associates, Inc.

Lundberg, S. M. and Lee, S.-I. (2017b). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, pages 4765–4774.

Miller, T. (2017). Explanation in Artificial Intelligence: Insights from the Social Sciences. *ArXiv e-prints: 1706.07269*.

Mullainathan, S. and Spiess, J. (2017). Machine learning: An applied econometric approach. *Journal of Economic Perspectives*, 31(2):87–106.

Ribeiro, M., Singh, S., and Guestrin, C. (2016). "why should I trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD*, pages 1135–1144.

References III

Schularick, M. and Taylor, A. M. (2012). Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870-2008. *American Economic Review*, 102(2):1029–61.

Shapley, L. S. (1953). A value for n-person games. *Contributions to the Theory of Games*, 2(28):307–317.

Shrikumar, A., Greenside, P., and Anshul, K. (2017). Learning important features through propagating activation differences. *CoRR*, abs/1704.02685.

Strumbelj, E. and Kononenko, I. (2010). An efficient explanation of individual classifications using game theory. *The Journal of Machine Learning Research*, 11:1–18.

Wager, S. and Athey, S. (2015). Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *ArXiv e-prints: 1510.04342*. 1510.04342.