Understanding Recessions with Machine Learning and AI

Rickard Nyman and Paul Ormerod

Algorithmic Economics Ltd and Department of Computer Science, UCL

Presentation to Bank of England/DAFM conference on Modelling with Big Data and Machine Learning, November 2018

Overview

- The recession that wasn't
- The recession that was
- Semi-supervised ML in text classification
- Time series analysis with ML


London Feel Good Factor

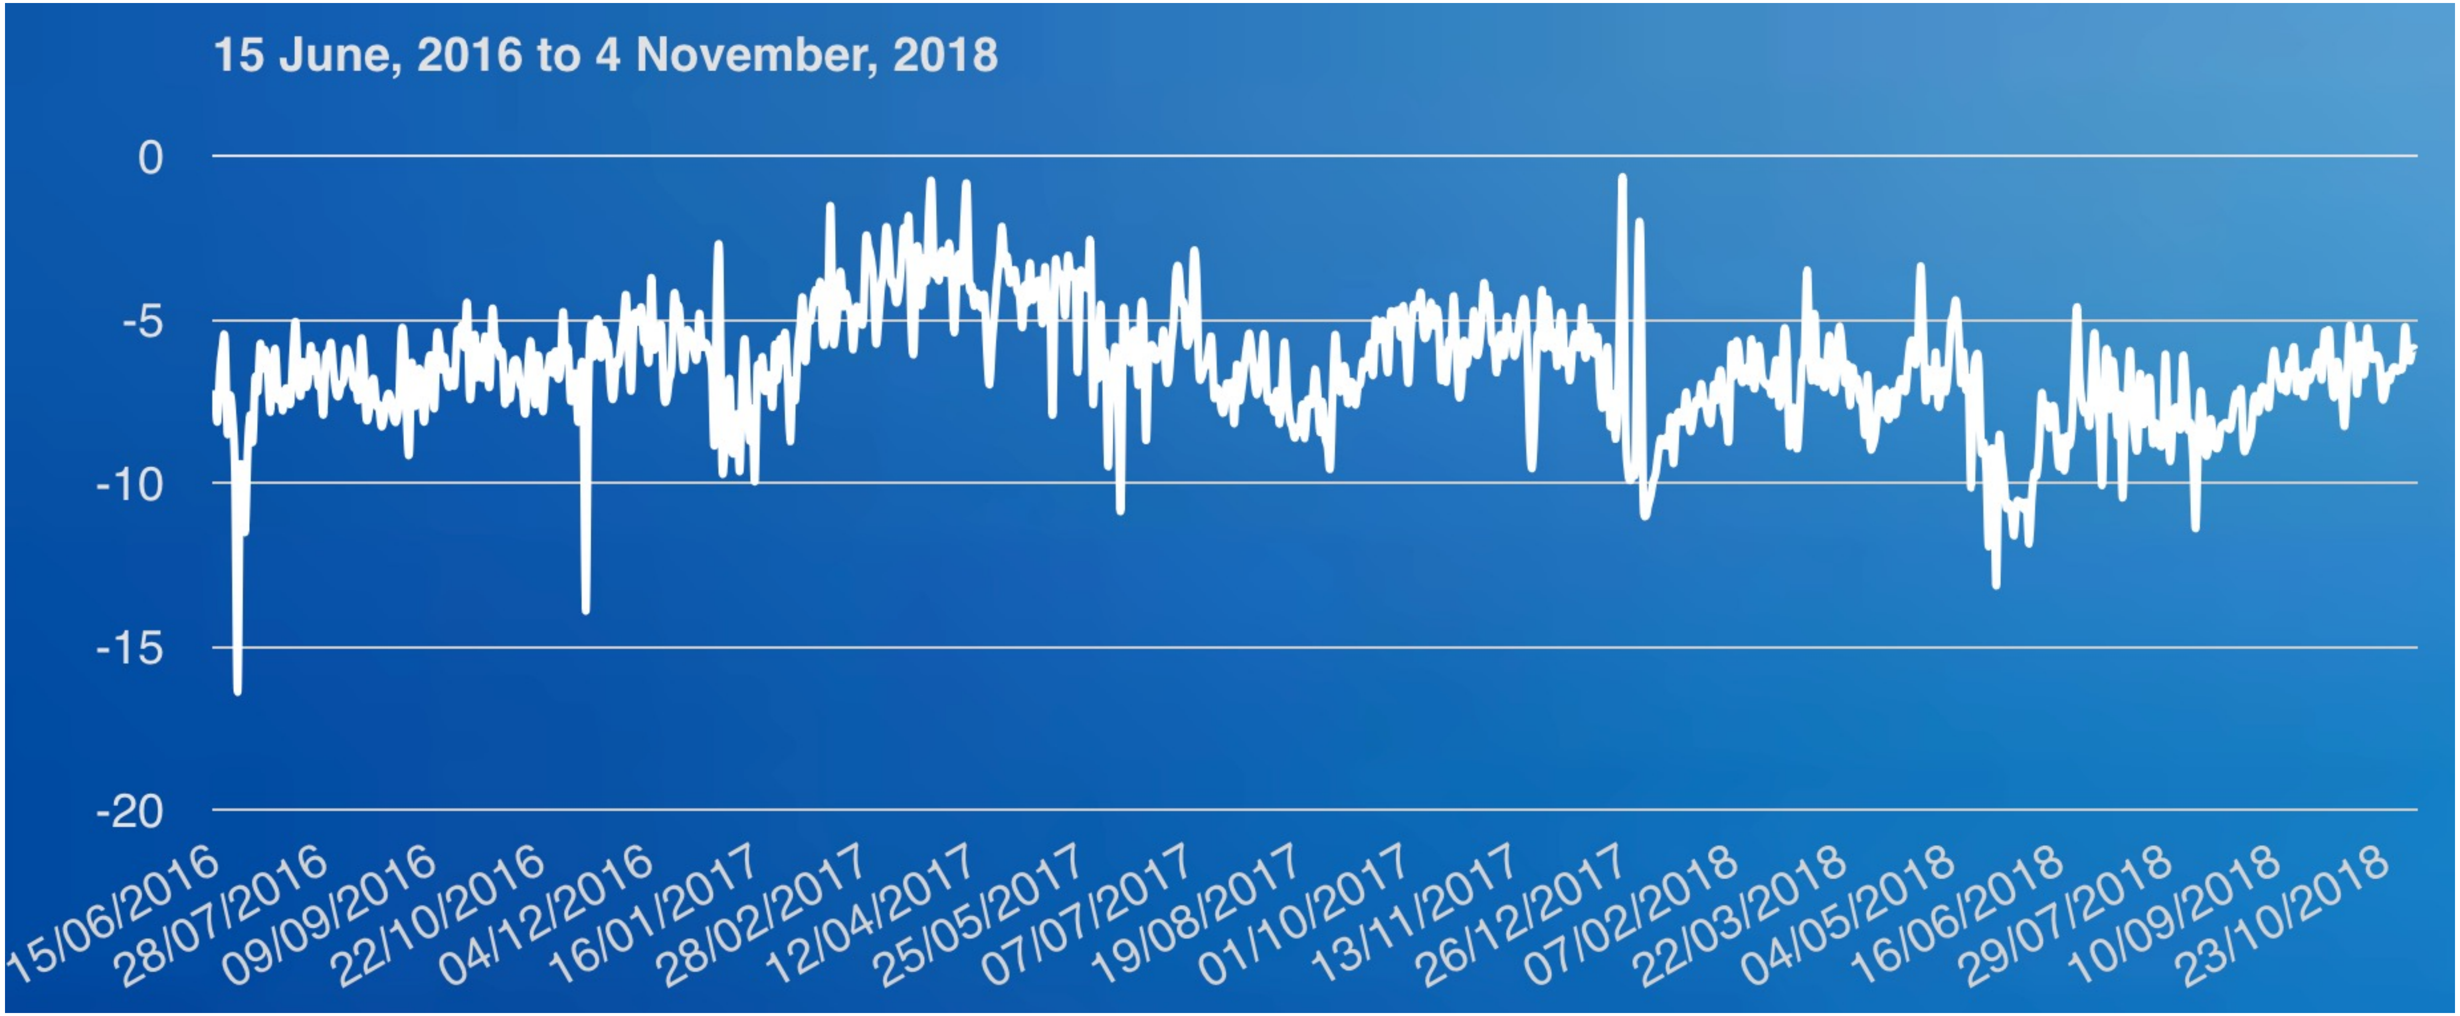

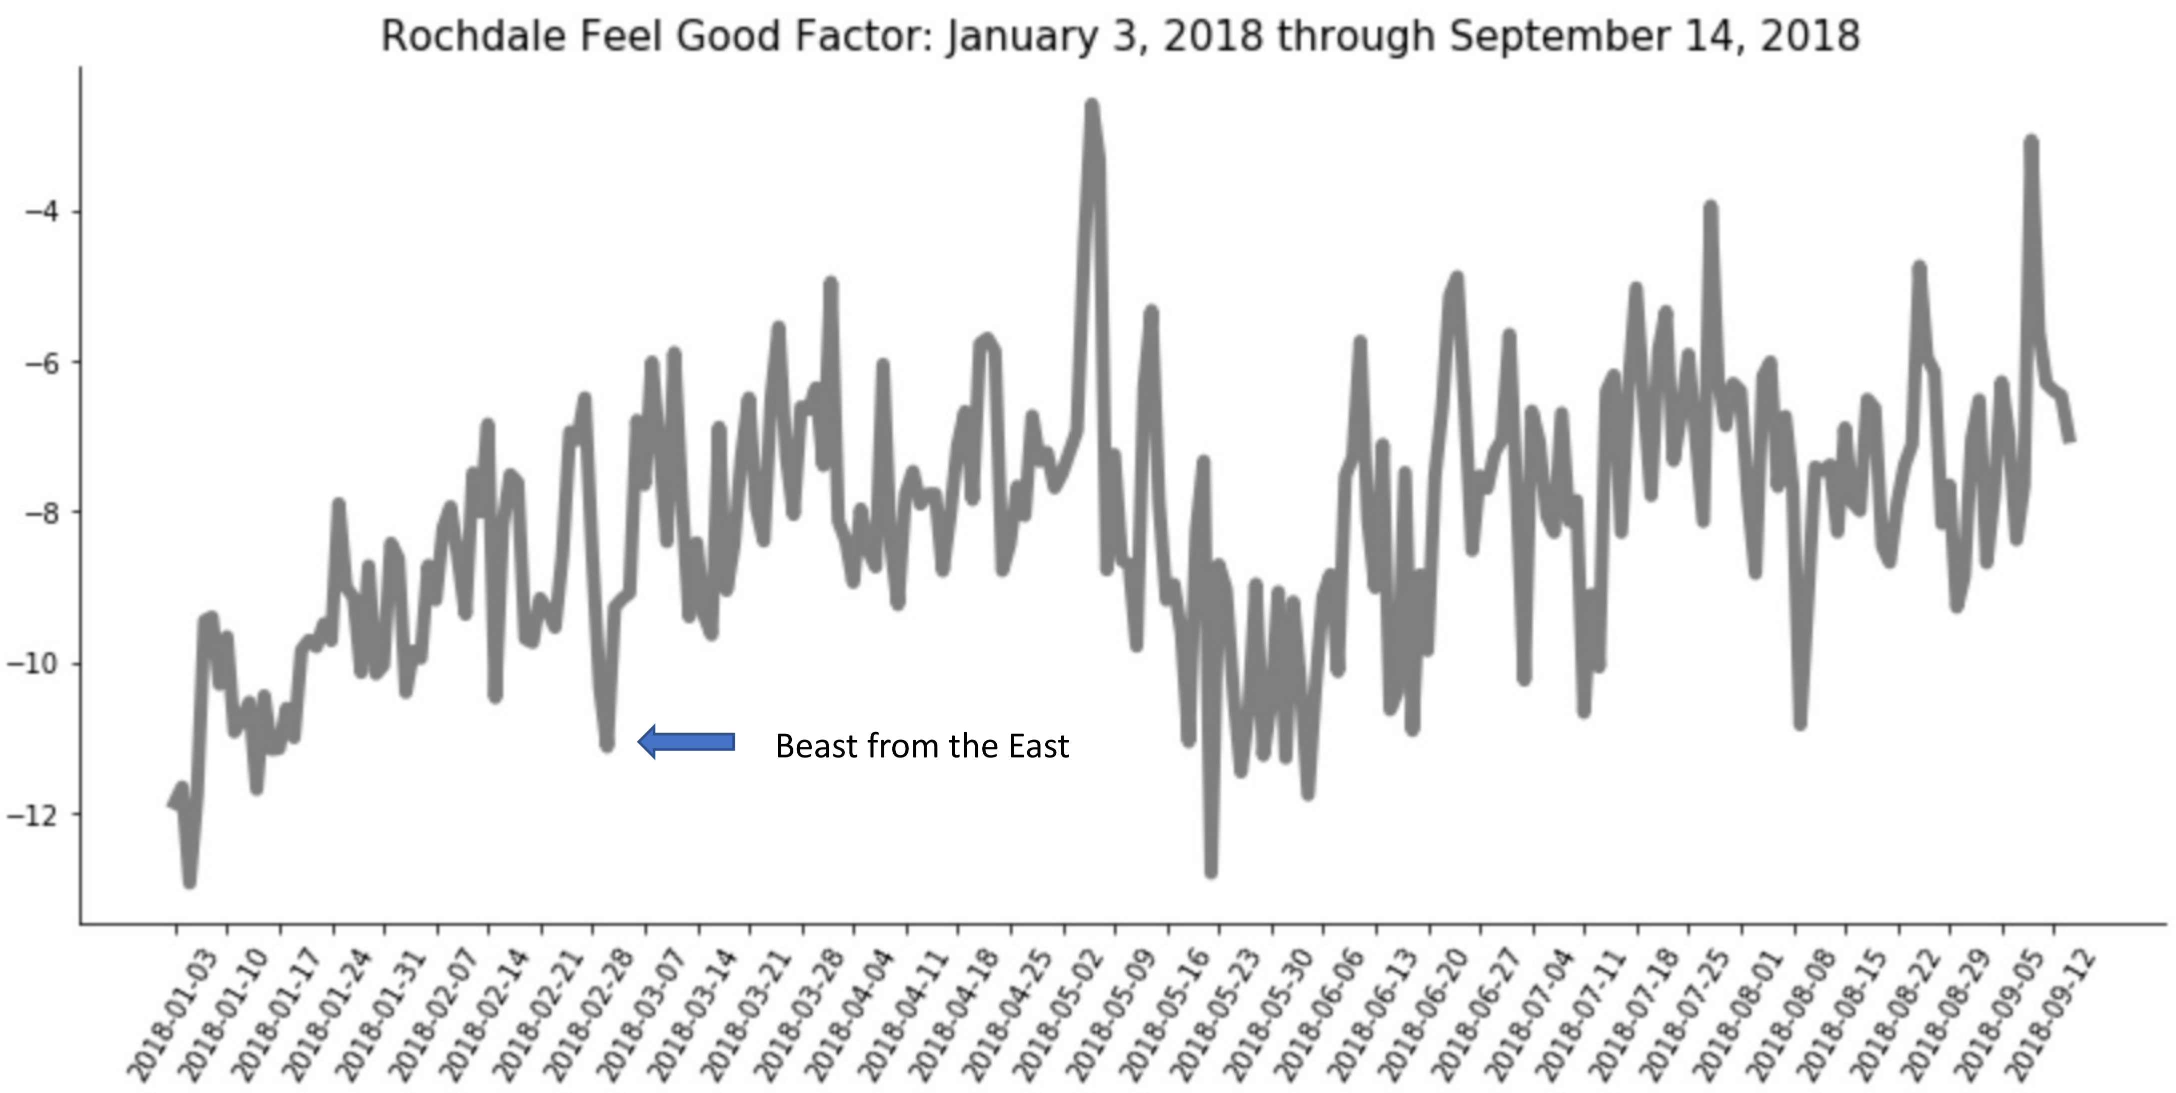

Rochdale Feel Good Factor: January 3, 2018 through September 14, 2018
−4
−6
−8
−10
−12
2018-01-03
2018-01-10
2018-01-17
2018-01-24
2018-01-31
2018-02-07
2018-02-14
2018-02-21
2018-02-28
2018-03-07
2018-03-14
2018-03-21
2018-03-28
2018-04-04
2018-04-11
2018-04-18
2018-04-25
2018-05-02
2018-05-09
2018-05-16
2018-05-23
2018-05-30
2018-06-06
2018-06-13
2018-06-20
2018-06-27
2018-07-04
2018-07-11
2018-07-18
2018-07-25
2018-08-01
2018-08-08
2018-08-15
2018-08-22
2018-08-29
2018-09-05
2018-09-12
Beast from the East

The forecasting record

- Even at the beginning of 2008, the economic recession of 2008/09 was not being predicted.
- The failure to predict recessions is a persistent theme in economic forecasting.
- The Survey of Professional Forecasters (SPF) provides data on predictions made for the growth of total output, GDP, in the United States for one through four quarters ahead, going back to 1968.
- The SPF is a benchmark to judge ML performance

Table 1 Regressions of GDP quarter on quarter growth, annualised rate, per cent, on SPF forecasts one quarter ahead

	Dependent variable (GDP quarter on quarter growth, annualised rate, per cent)			
Sample Period	1970Q2 – 2010Q4		1990Q2 – 2010Q4	
	Most Recent Estimate	Third Estimate	Most Recent Estimate	Third Estimate
	(1)	(2)	(3)	(4)
SPF	0.962	1.059	1.146	1.221
	(0.131)	(0.126)	(0.244)	(0.218)
Constant	0.404	-0.102	-0.341	-0.438
	(0.420)	(0.372)	(0.664)	(0.593)
Adj R2	0.245	0.300	0.205	0.272
Res Std. Error	3.033	2.910	2.362	2.110

Potential limits to predictability

- Ormerod and Mounfield (*Physica A 2000)* “Random matrix theory and the failure of macroeconomic forecasting”
- Can be thought of as a signal processing technique which identifies the true amount of information in the data compared to noise
- Data series dominated by noise
- We conjecture it is the high dimensional nature of the problem relative to the available number of observations. Empirically, the data will appear to contain large amounts of “noise” in such situations
- There may therefore be upper limits to how accurate forecasts might be
- We could not improve on the one-quarter ahead SPF track record

SPF Four Quarters Ahead

- Four quarters ahead, the mean SPF prediction has *never* been for negative growth over the entire 1970Q2 – 2018Q3 period
- Regressions of the mean forecast in the SPF against the out-turn have at best very little and often no explanatory power (across different sample periods)

Table 2 Regressions of GDP quarter on quarter growth, annualised rate, per cent, on SPF forecasts four quarters ahead

	Dependent variable (GDP quarter on quarter growth, annualised rate, per cent)			
Sample Period	1970Q2 – 2010Q4		1990Q2 – 2010Q4	
	Most Recent Estimate (1)	Third Estimate (2)	Most Recent Estimate (3)	Third Estimate (4)
SPF	0.713	0.985	-0.407	0.171
	(0.287)	(0.281)	(0.544)	(0.509)
Constant	0.691	-0.417	3.689	2.132
	(0.942)	(0.922)	(1.570)	(1.429)
Adj R2	0.031	0.065	-0.005	-0.011
Res Std. Error	3.436	3.362	2.657	2.485

A ML approach

- Fernandez-Delgado et al. (*Journal of Machine Learning Research* 2014) compare 179 classification algorithms from 17 “families” such as Bayesian, neural networks, logistic and multinomial regression.
- They examine their performance on 121 data sets in the University of California at Irvine machine learning repository. This repository is in standard use in machine learning research. The data sets are what we would think of as cross-sectional
- The random forest family of algorithms achieves the best results. Bayesian and logistic regression algorithms, “are not competitive at all”. (p.3175)
- The limited amount of work in time-series analysis suggests that RF is the approach to use
- Because of the limited number of observations, we have not been able to apply modern neural network techniques

The analysis

- We compare OLS and RF out-of-sample predictions
- We use data 1970Q2-1989Q2 to train the model and then predict 1990Q2 (i.e. data available in 1990Q2)
- We roll this forward quarter by quarter, ending with a model trained 1970Q2-2009Q4 to predict 2010Q4
- We used the default values for the various options available for inputs in the random forest algorithm. In other words, we did not attempt in any way to optimise the accuracy of the predictions by trying different combinations of input parameters to the algorithm (we obtained very similar results using both R and Python)
- We used a (very) limited set of explanatory variables, using current and one lagged value of each, and carried out minimal pruning. We only excluded variables which made no difference to the forecasting performance.

The data

- Third estimate of GDP
- Treasury bill rate; 10 year government bond yield; quarterly percentage change in S and P 500; household debt as percent of GDP; non-financial corporate debt as percent of GDP; public debt as percent of GDP
- Leaving out the data on the bond yield and public debt made no difference to the forecasting performance (in fact it improved it slightly)

Actual regressed on the 4-quarter ahead forecasts, 1990Q2-2010Q4, average of 100 predictions

Table 3 Regressions of GDP quarter on quarter growth, annualised rate, per cent, third estimate on random forest predictions made four quarters previously

	Dependent variable: GDP quarter on quarter growth, annualised rate, per cent, third estimate
Prediction	0.827
	(0.157)
Constant	0.304
	(0.5)
Adjusted R2	0.246
Residual Std. Error	2.146

Plot of the actual 3rd estimate GDP growth and 4-quarter ahead forecasts, 1990Q2-2010Q4, average of 100 predictions

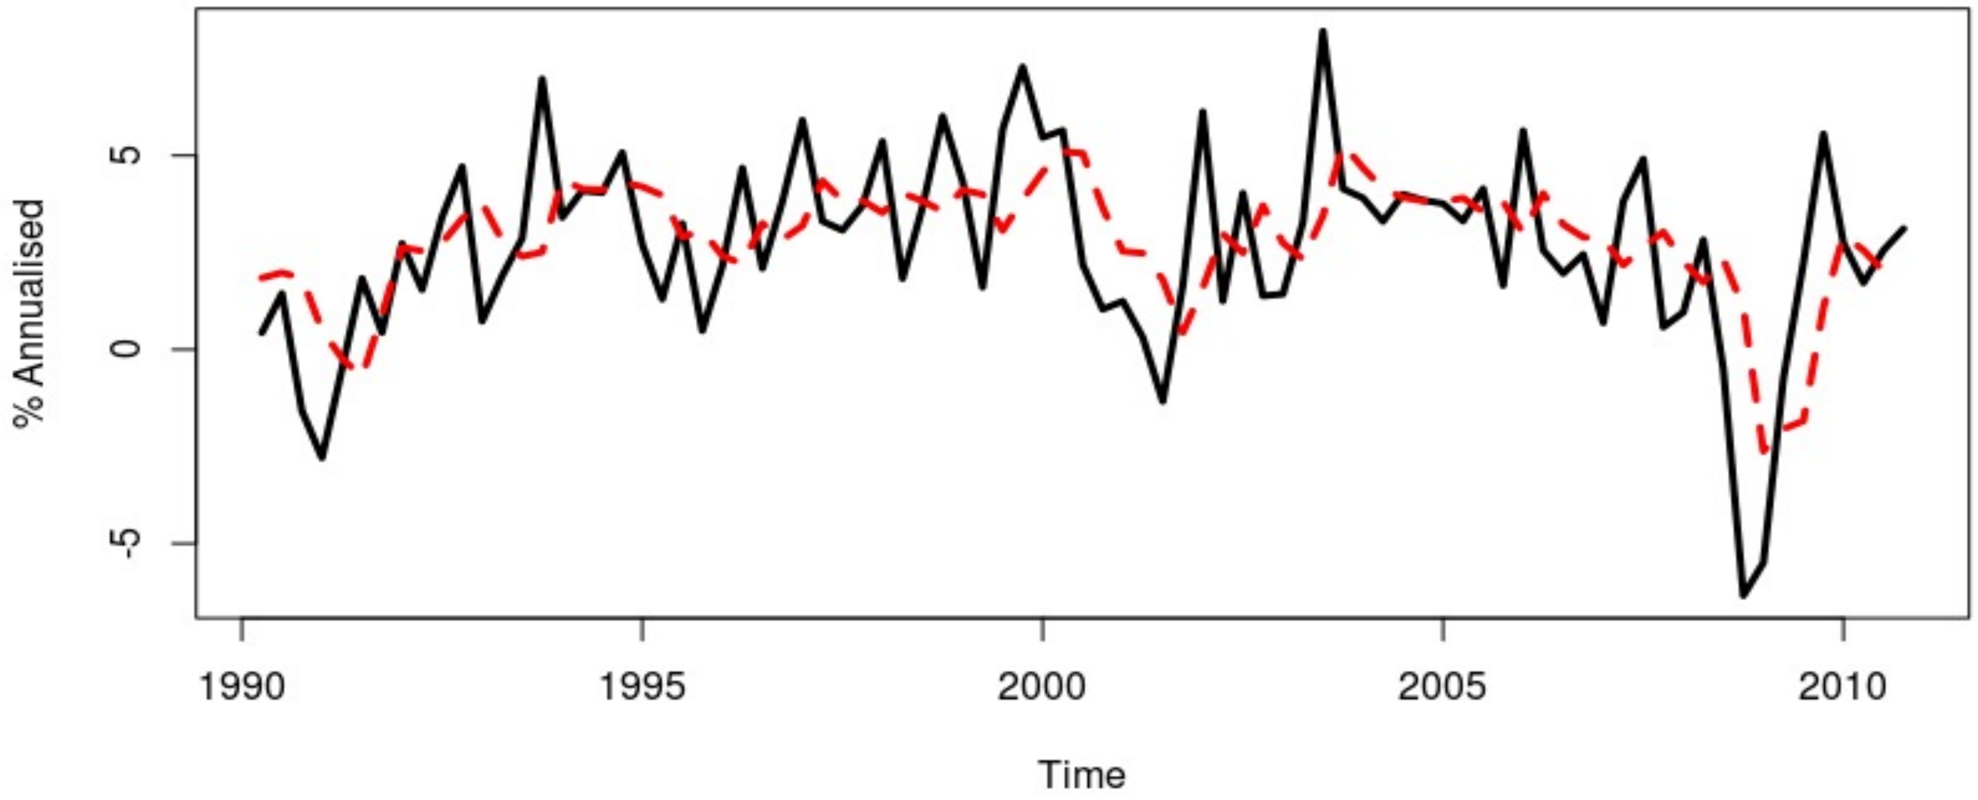

Period	Actual annualised third estimate quarter on quarter real GDP growth	Random forest predictions made four quarters previously
2008Q1	0.96	2.27
2008Q2	2.83	1.73
2008Q3	-0.51	2.32
2008Q4	-6.34	1.08
2009Q1	-5.49	-2.62
2009Q2	-0.74	-2.05
2009Q3	2.24	-1.84
2009Q4	5.55	1.09

Table 4 Adjusted R squared values of Regressions of GDP quarter on quarter growth, annualised rate, per cent, third estimate on random forest predictions made four quarters previously, 1990Q2-2010Q4

All variables included	0.245
Omitting change in share prices	0.226
Omitting household debt to GDP	0.195
Omitting corporate debt to GDP	0.255
Omitting Treasury Bill rate	0.233
Omitting household and corporate debt to GDP	0.119

UK results

- Identical approach
- No SPF equivalent available, though we know the UK forecasting performance is very similar to that of the US
- Same variables (UK based obviously)
- We left out 10 year bond and Treasury bill variables from the “base” model

Regressions of GDP quarter on quarter growth, annualised rate, per cent, third estimate on random forest predictions made four quarters previously

	Dependent variable: GDP quarter on quarter growth, annualised rate, per cent, third estimate
Prediction	0.803
	(0.136)
Constant	0.395
	(0.300)
Adjusted R2	0.294
Residual Std. Error	1.945

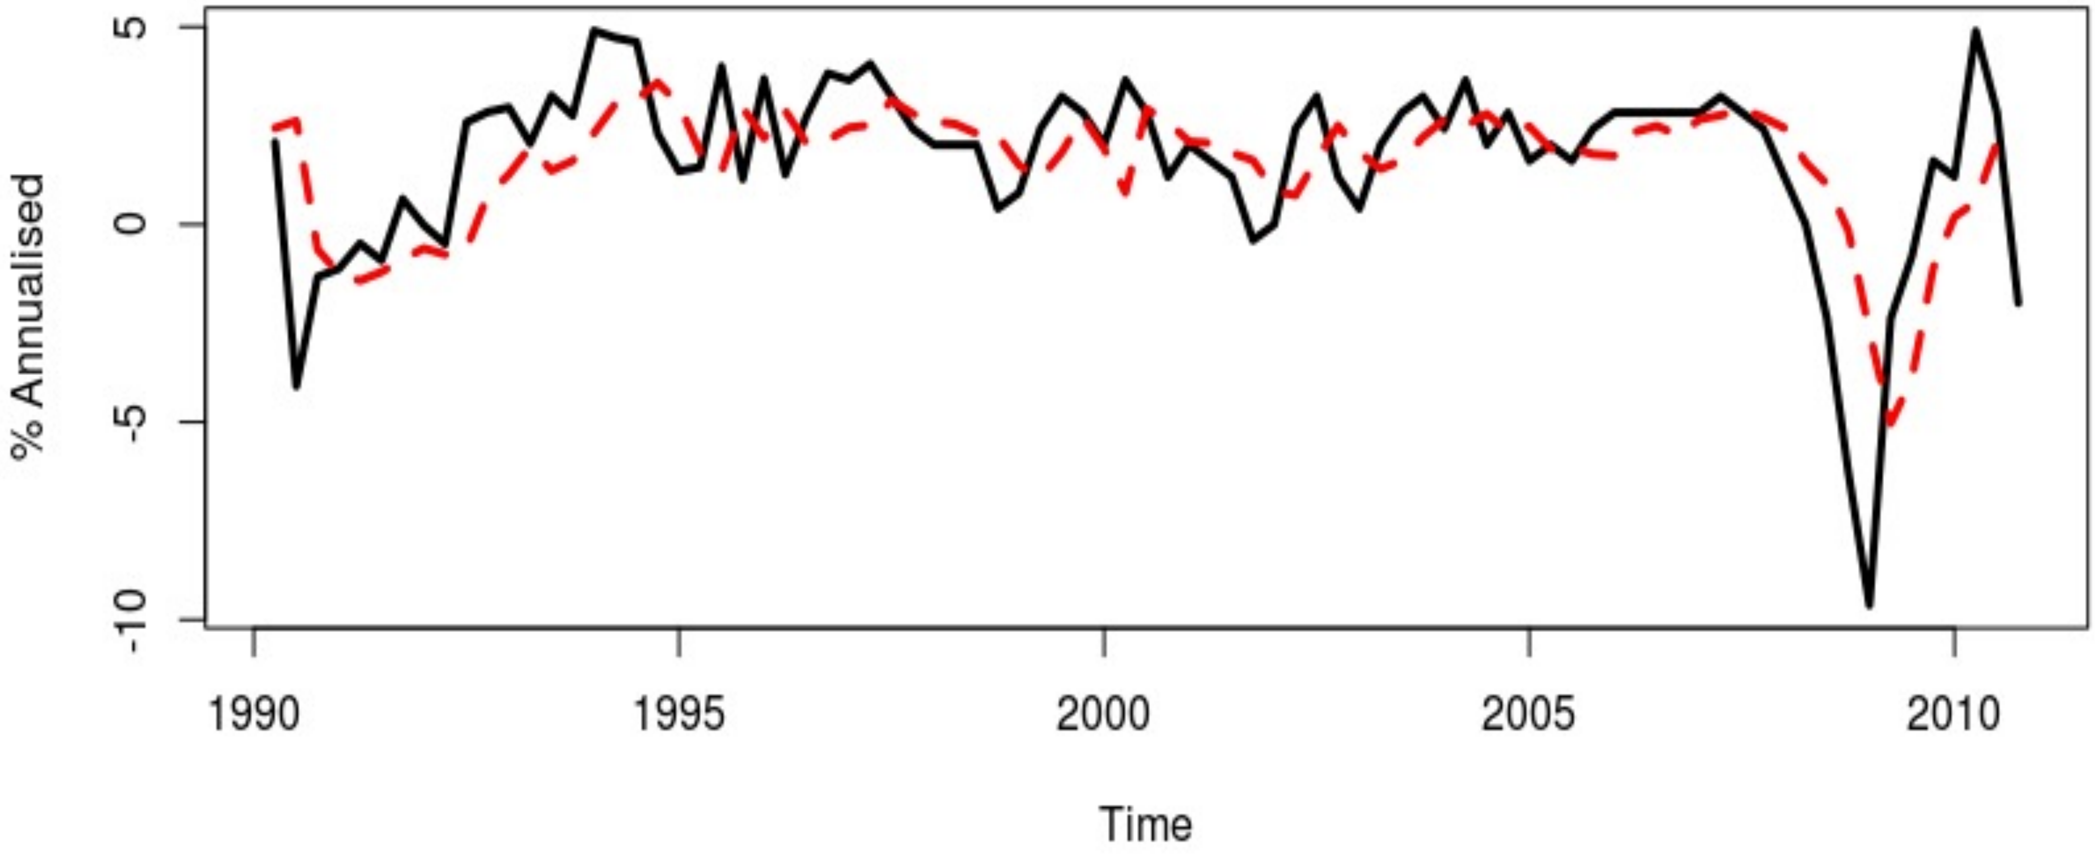

5
0
-5
-10
% Annualised
1990
1995
2000
2005
2010
Time

Period	Actual annualised third estimate quarter on quarter real GDP growth	Random forest predictions made four quarters previously
2008Q1	1.21	2.44
2008Q2	0.00	1.56
2008Q3	-2.38	1.02
2008Q4	-6.25	-0.15
2009Q1	-9.63	-2.71
2009Q2	-2.38	-5.03
2009Q3	-0.80	-3.89
2009Q4	1.61	-1.10

Adjusted R squared values of Regressions of GDP quarter on quarter growth, annualised rate, per cent, third estimate on random forest predictions made four quarters previously, 1990Q2-2010Q4

All variables included	0.293
Omitting change in share prices	0.265
Omitting household debt to GDP	0.258
Omitting corporate debt to GDP	0.296
Omitting public debt to GDP	0.274
Omitting household and corporate debt to GDP	0.092