Non-linear relation between industrial production and business surveys data

Bruno, Giancarlo

ISAE (Institute for studies and economic analisys) Rome (Italy)

September 2009

Non-linear relation between industrial production and business surveys data

by

Giancarlo Bruno
ISAE, piazza dell'Indipendenza, 4, 00185 Roma, Italia
e-mail: g.bruno@isae.it

Working paper n. 119
September 2009

The Series "*Documenti di Lavoro" of the Istituto di Studi e Analisi Economica* - Institute for Studies and Economic Analyses (ISAE) hosts the preliminary results of the research projects carried out within ISAE. The diffusion of the papers is subject to the favourable opinion of an anonymous referee, whom we would like to thank. The opinions expressed are merely the Authors' own and in no way involve the ISAE responsability.

The series is meant for experts and policy-makers with the aim of submitting proposals and raising suggestions and criticism.

La serie "Documenti di Lavoro" dell'Istituto di Studi e Analisi Economica ospita i risultati preliminari di ricerche predisposte all'interno dell'ISAE. La diffusione delle ricerche è autorizzata previo il parere favorevole di un anonimo esperto della materia che qui si ringrazia. Le opinioni espresse nei "Documenti di Lavoro" riflettono esclusivamente il pensiero degli autori e non impegnano la responsabilità dell'Ente.

La serie è destinata agli esperti e agli operatori di politica economica, al fine di formulare proposte e suscitare suggerimenti o critiche.

Stampato presso la sede dell'Istituto

ISAE - Piazza dell'Indipendenza, 4 - 00185 Roma.
Tel. +39-06444821; www.isae.it

ABSTRACT

In this paper I compare different models, a linear and a non-linear one, for forecasting industrial production by means of some related indicators. I claim that the difficulties associated with the correct identification of a non-linear model could be a possible cause of the often observed worse performance of non-linear models with respect to linear ones observed in the empirical literature. To cope with this issue I use a non-linear non-parametric model. The results are promising, as the forecasting performance shows a clear improvement over the linear parametric model.

Keywords*:* Forecasting, Business Surveys, Non-linear time-series models, Non-parametric models.

JEL classification: C22, C53.

CONTENTS

1. INTRODUCTION 5
2. LITERATURE 5
3. MODEL 7
3.1 Estimation 8
3.2 Bandwidth and lag length selection 9
4. DATA 10
5. EMPIRICAL FRAMEWORK 12
6. RESULTS 13
7 CONCLUSIONS 18
References 19
Appendix 21

1 INTRODUCTION

This paper[1] deals with the modelling of the relation between Italian industrial production and some leading indicators. In particular the focusing is on the forecasting performance of a non-linear non-parametric model vs. a linear one. The results are analysed within a set of forecasting performance indicators and show a superiority of the non-linear non-parametric model.

The relevance of the problem at hand should be self-evident to all practitioners used to provide short-term forecasts of, e.g., GDP: actually, obtaining a good forecast of industrial production is often the most important step, because it is by far the most relevant indicator about short-term development of the economy.

The rationale for this exercise is that there is no theoretical or practical reason why the relation between the industrial production index and the selected indicators should be linear; nevertheless, it is well documented in the long stream of literature on non-linear time-series modelling, that correct specification of non-linear models can be a very difficult task and, I would add, once obtained would anyway be affected by stability problems (in the sense of temporal stability of the model) even more formidable than linear ones. Having said that, a natural alternative could be represented by a non-linear non-parametric model. The non-parametric feature could be useful to overcome the identification issue involved with specifying a particular non-linear model. This comes at a cost: estimation consistency rates are slower than those obtainable for a correctly specified parametric model, either linear or non-linear. Anyway, I think that this issue is much more overlooked than it should be, at least in a forecasting context. In fact, forecasting models are much more susceptible to be miss-specified, as their construction must take into consideration issues such as data availability and timeliness, which greatly limit the opportunity to build a correctly specified model.

2 LITERATURE

There is a considerable literature in modelling the industrial production index, both in the univariate and in the multivariate framework.

[1] I wish to thank T. Proietti for his precious advise, F. Peracchi and G. Cubadda for their valuable comments.

Moreover, in both cases, linear as well as non-linear specifications have been employed.

As far as the univariate framework is considered, in many cases a simple linear model shows a better performance over non-linear ones. As an example, Siliverstovs and Dijk (2003) compare linear autoregressive (AR), linear AR with breaks, threshold autoregressive (TAR), self-exciting autoregressive (SETAR) and Markov-switching autoregressive (MS-AR) models, in terms of point, interval and density forecasts. They found that linear AR outperforms the other models when point forecasts are considered, although MS-AR model are more accurate for interval and density forecasts. The study was conducted on seasonally adjusted data: however it is well known that the use of such data with AR models is at least questionable.[2] Moreover, seasonal adjustment implies revisions in the data, which in order to be properly accounted for in the forecast evaluation, would need the use of different vintages of data. In the end, the raw data might well be the final target to forecast. Indeed, the issue of correctly treating the seasonality in monthly industrial production has received attention too, leading also in this case to detect and model non-linearities. As an example, Osborn and Matas-Mir (2003) observe the non-linearity emerging from the interactions of seasonal and business cycle fluctuations, finding a reduction in seasonality in the upper regime of the business cycle. A similar kind of non-linearity was found also in the Italian case by Proietti (1998) and Bruno and Lupi (2004). Franses and van Dijk (2005) consider different seasonal models and conclude that simpler models for seasonality yield better point forecasts for short horizons, while more elaborate models perform better for longer horizons.

Öcal (2000) compares a smooth-transition autoregressive (STAR) model vs. a linear AR within a set of macroeconomic variables. In particular, for the industrial production he finds that the best model is a three regimes STAR, even though no statistically significant differences are found between linear and non-linear models.

On the other hand, the usefulness of non-linear models seems stronger in the case of multivariate models. As in Bradley and Jansen (2004), who consider a STAR model against a linear AR model for joint industrial production and stock returns. Some improvement in

[2] Seasonal adjustment procedures most of times cause a zero in the spectral density of the adjusted data at seasonal frequencies, so that such series do not possess an invertible representation.

industrial production forecast is observed for the non-linear model over the linear one.

Venetis et al. (2004) compare a linear autoregressive distributed lag (ADL) model with a TAR model using industrial production and term spread; the non-linear modelling obtains in part better results. Also Jagric (2003) finds that neural network approach helps improving forecasting industrial production by means of a leading indicator over a linear model, while Simpson et al. (2001), who consider a linear ADL and a MS model with a leading indicator, find that one-step ahead forecasts produced by the linear model are better.

Huh (1998) exploits asymmetry in the relation between industrial production and an index of financial markets conditions. One-step ahead forecasts from a linear and a MS model are compared, with the latter performing significantly better.

In the Italian case, Marchetti and Parigi (2000) find evidence of non-linearity in the relation between industrial production and electricity consumption, which they represent with a STAR model. Anyway, the best forecasting performance is obtained with a linear model.

3 MODEL

In this work I consider two set of variables:

- I denote with X_t the variable of interest to be forecast (stationary transformation of the industrial production index);
- I denote with Z_t the related indicator, which is chosen in turn from a group of three variables better described in section 5.

In the empirical exercise, which is fully described in section 5, it is necessary to consider the different timing with which the two variables are released. In this case the related indicators are released about 45 days before the industrial production index, so when the latter is available for month t, the former are available at least up to month $t+1$. Therefore, if a relation is found between X_t and Z_{t-d}, with $d \geq 0$, it is possible to forecast X_t up to $d+1$ step-ahead.

The non-linear non-parametric model used is a functional coefficient regression (FCR) model:

$$X_t = a_1(Z_{t-d})X_{t-1} + \ldots + a_p(Z_{t-d})X_{t-p} + \varepsilon_t, \quad t = p+1, \ldots, T \quad (1)$$

where ε_t is a martingale difference process and $\{X_t, \ldots, X_{t-p}\}$ is a strictly stationary $\beta-$mixing process.

Model (1) is non-parametric in the sense that the functional form of the coefficients $a_i(\cdot)$ is not specified. It is derived from the state-dependent model introduced by Priestley (1980). Cai et al. (2000a) and Cai et al. (2000b) address the issue of estimation, bandwidth selection and testing. The main justification for using such a model is that the autoregressive coefficients are varying, and depend, in a rather smooth way, on the state of the leading indicator Z_t at a certain lag d.

This kind of model has some appealing features, in that it nests the linear AR model, as well as some popular non-linear parametric models, such as threshold autoregressive (TAR), exponential autoregressive (EXPAR) and smooth transition autoregressive (STAR) models. Therefore, it is sufficiently general to handle many kinds of non-linearities often found in macroeconomic time series, while reducing considerably the problem of model complexity: the unknown functions, in fact, depend only on one variable in this set-up.

Moreover, it has a nice interpretation, as the coefficients depend on the "state" of the variable Z_{t-d} in a smooth way, differently from what happens in the TAR model, where the autoregressive parameters shift discontinuously following the discrete number of states associated to the variable Z_{t-d}.

3.1 Estimation

The estimation of the unknown functions $a(\cdot)$ of model (1) can be carried out approximating them locally with a polynomial of suitable order. In particular, considering a first order polynomial $a_i(u)$ can be approximated as follows around x:

$$a_i(u) \approx a_i(x) + a_i'(u - x) \equiv \alpha_i + \beta_i(u - x). \tag{2}$$

In order to carry out this locally and denoting with $U_t = Z_{t-d}$, one has to minimize the following expression with respect to $\{\alpha_i, \beta_i, i = 1, \ldots p\}$:

$$\sum_{t=p+1}^{T} \left\{ X_t - \sum_{i=1}^{p} \left[\alpha_i + \beta_i(U_t - u)\right] X_{t-p} \right\}^2 K\left(\frac{U_t - u}{h}\right) \tag{3}$$

where $K(\cdot)$ is a non-negative weight function which downweight observations far from the point u while the parameter h is a smoothing constant, generally called *bandwidth*. This parameter represents how much "local" the estimator is, that is the width of the interval around a point u which is used for estimating the coefficient functions at that point.

Denote with $\boldsymbol{X}$ the $T \times 2n$ matrix whose t row is: $(X_{t-1}, \dots, X_{t-p}, X_{t-1}(U_t - u), \dots, X_{t-p}(U_t - u))$, with $\boldsymbol{Y}$ the $T \times 1$ vector $(X_1, \dots, X_T)$, with $\boldsymbol{W}$ the $T \times T$ matrix with t diagonal element equal to $h^{-1}K\left(\frac{U_t - u}{h}\right)$ and 0 elsewhere. The minimization problem (3) has the following solution:

$$\hat{\boldsymbol{\beta}} = (\boldsymbol{X}'\boldsymbol{W}\boldsymbol{X})^{-1}\boldsymbol{X}'\boldsymbol{W}\boldsymbol{Y} \tag{4}$$

where $\boldsymbol{\beta} = (\alpha_1, \dots, \alpha_p, \beta_1, \dots, \beta_p)$. Therefore, the first p elements of $\hat{\boldsymbol{\beta}}$, which I denote with $\{\hat{\alpha}_i\}_{i=1,\dots,p}$ are the local linear estimate of the functional coefficients $\{a_i(u)\}_{i=1,\dots,p}$.

3.2 Bandwidth and lag length selection

In order to carry out the estimation of the functional coefficients also the values of h and p must be estimated from the data. [3]

A form of cross validation has been used to select both these quantities, following Cai et al. (2000a). In particular, denoting with Q and m two integers such that $Qm < T$; the first Q sub-series of length $T - qm$ are used $(q = 1, \dots, Q)$ to estimate the model and then the one-step forecasting errors are computed for the next m points of the series.

For a given h (bandwidth) and p define the average prediction error for the single sub-series:

$$APE_q(h, p) = \frac{1}{m} \sum_{t=T-qm+1}^{T-qm+m} \left[X_t - \sum_{i=1}^{p} \hat{\alpha}_{i,q}(Z_{t-d}) X_{t-i} \right]^2, q = 1, \dots, Q. \tag{5}$$

where $\hat{\alpha}_{i,q}(u)$ is the coefficient estimated using the observations $\{1, \dots, T - qm\}$. For example, when $q = 2$ the set of data $\{1, \dots, T - 2m\}$ is used for getting the estimated coefficients $\hat{\alpha}_{i,2}(u)$, and the

[3] In principle this is valid also for d, the delay parameter of the leading indicator; nevertheless for the purpose of this paper this is not necessary, as will be better illustrated in section 5.

subsequent set of data $\{T-2m+1, \ldots, T-m\}$ is used for calculating the prediction error (5). Following Cai et al. (2000b), Q is taken equal to 4 and m equal to $T/10$.

Moreover, define the quantity $APE(h, p)$, which averages over the all the sub-series considered:

$$APE(h, p) = Q^{-1} \sum_{q=1}^{Q} APE_q(h, p). \tag{6}$$

The values of p and h are then chosen so that (6) is minimized.

In the procedure described above the bandwidth is maintained fixed over the support of u. An alternative approach, so called k-nearest neighbour (k-nn), consists, instead, in taking a fixed number of observation around a given value of u, leading to a bandwidth which is not constant over the support of u. Though I applied also this method for estimating the functional coefficients $a_i(\cdot)$, I do not show the results here, which were not as good as in the fixed bandwidth case.

4 DATA

I carry out this exercise using Italian data, in particular:

- industrial production index (X_t), which is published monthly by ISTAT, the national statistical office;

- survey results (Z_t) on: production trend (PT), production level (PL), order books (OB), released monthly by ISAE, a state owned economic research institute.

The industrial production index is considered in its working-days adjusted form[4]. Moreover, stationarity is achieved through log transformation and seasonal differentiation.

The leading indicators are produced through a survey where industrial entrepreneurs are asked many questions. Among them: the production trend (PT) in their firm during the following 3-4 months; the answer can be "increasing", "decreasing", "stationary". Individual

[4] This is a minor point in our view, as the published working days adjusted series is obtained by means of the procedure TRAMO-SEATS, which is equivalent to apply a linear transformation to the original series. As the TRAMO-SEATS specifications are publicly available, switching between raw and working-days adjusted data is straightforward.

Figure 1: Italian Industrial production industry - total industry excluding construction - left panel: index base 2005=100, right panel: seasonal difference of logs.

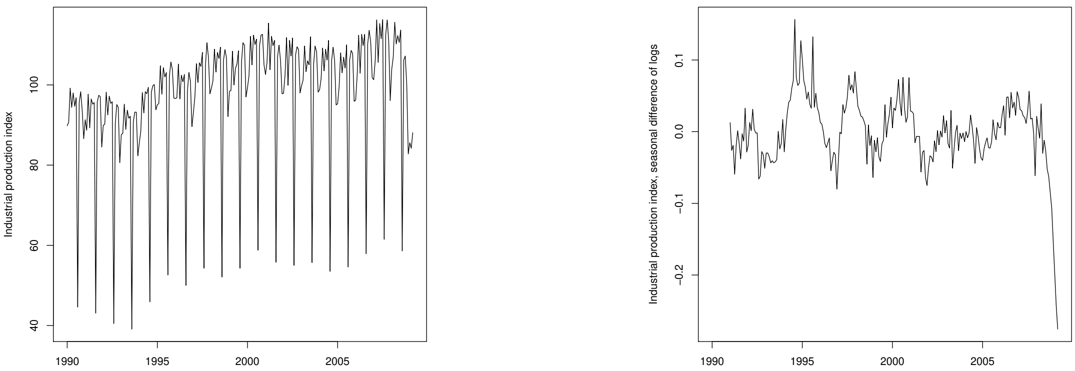

results are then suitably aggregated to provide shares attributable to the different answers for the whole manufacturing sector. Here we follow the usual technique of quantifying those results with the so called *balance*, i.e. the difference between “increasing” and “decreasing” aggregate answers. Another question is the current production level (PL); the answers can be “high”, “normal”, “low”, and the results are aggregated as before to obtain a balance between “high” and “low” answers. The same holds for order books (OB). In principle it is possible to conjecture that PT is a leading series, as well as OB, while PL should be a coincident indicator. In practice this is not always so clear cut.

The survey series obtained are significantly affected by seasonality; anyway the use of the seasonal difference here could be questionable, as the series can be hardly thought of as being seasonally integrated (actually the series are bounded).[5] Therefore I removed the seasonality by taking a 12-term asymmetric moving average. This is preferable in my opinion to the alternative of removing the seasonality by more elaborate filtering methods, like X-12 or TRAMO-SEATS, because they imply a revision pattern, which I would like to avoid as much as possibile in a forecasting exercise.

[5] Indeed, Pappalardo (1998) shows that in most cases business surveys data can be characterized by a stationary seasonality.

Figure 2: Survey results – balances – 12-term moving average.

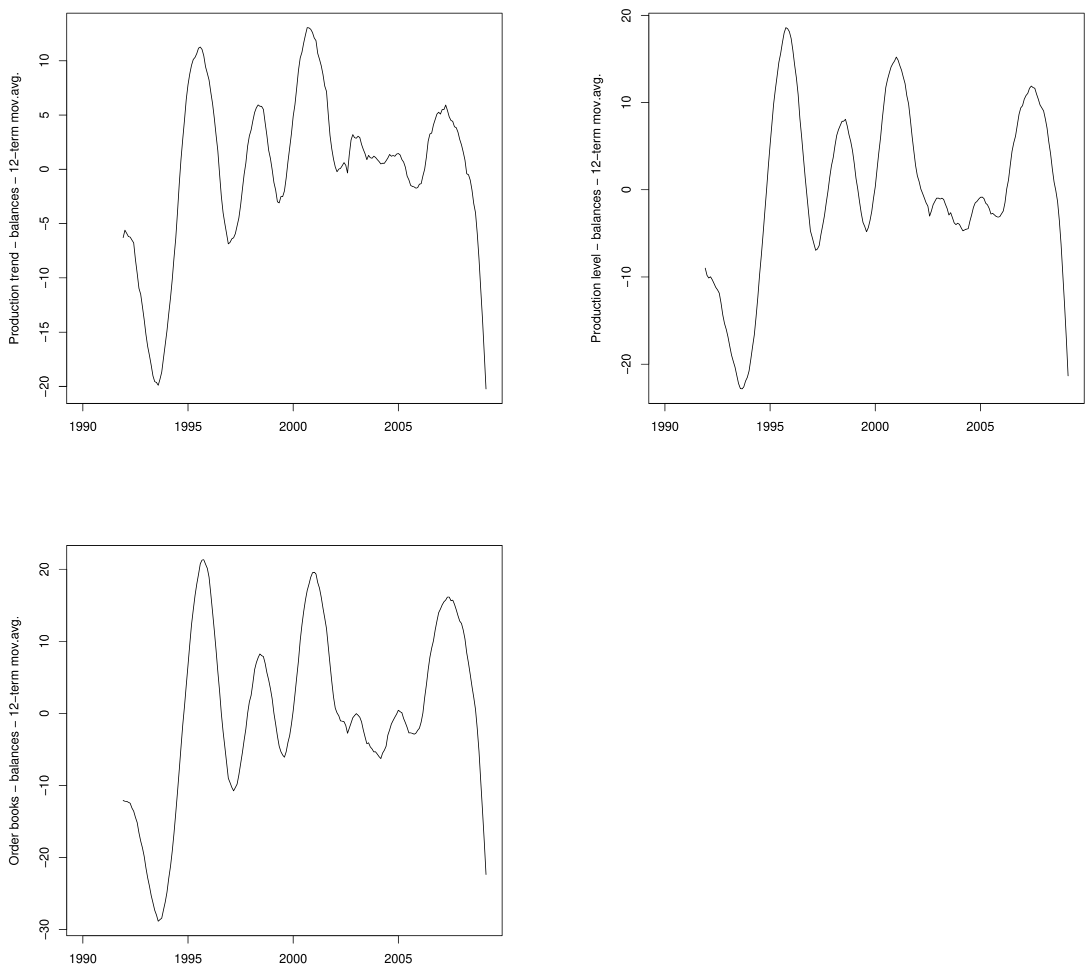

5 EMPIRICAL FRAMEWORK

The results of the FCR model are compared to those stemming from an autoregressive distributed lag (ADL) model. The latter can be considered a linear benchmark for those who are seeking to forecast a variable by means of another variable. In our case the ADL model takes the following form:

$$X_t = \sum_{i=1}^{p} \alpha_i X_{t-i} + \sum_{j=d}^{q} \beta_j Z_{t-j} + \varepsilon_t \qquad d \geq 0. \tag{7}$$

Obviously also in this case the order p and q must be chosen in some way, usually by some likelihood-based criterion.

I choose not to establish a value for d, the delay with which the indicator variable enters the relation with the industrial production, neither in the FCR nor in the ADL model. Indeed, my purpose is to compare the FCR with the ADL model and I compare this performance separately for different values of d, i.e. $d = 0, \ldots, 5$. Given the release timing of X_t and Z_t, the latter is always available with a lead of at least one month; therefore, for every d, forecasts can be generated up to $d+1$ step-ahead.

Once allowed for the data loss due to differentiation and lag creation, I have a database of 196 monthly observations. Identification of significant lags for the models were carried out on the first 145 observations; the last 48 were used for forecast evaluation. Obviously a possibly different set of lags was selected for every lag d with which the leading indicators enter the relationships (1) and (7). More specifically, in the case of the FCR model, for every value of d between 0 and 5 a different model was identified, by means of the cross validation criterion (6); the set of lags considered was $\{1, \ldots, p, 12\}$. The seasonal lag was always included, while the different values of $1 \leq p \leq 6$ were considered.

The ADL model was identified for every d considering a general model of the form:

$$X_t = \sum_{i=1}^{12} \alpha_i X_{t-i} + \sum_{j=d}^{12} \beta_j Z_{t-j} + \varepsilon_t \quad d \in \{0, \ldots, 5\}. \tag{8}$$

and selecting a subset of regressors by means of the BIC criterion in a stepwise regression.

6 RESULTS

Forecasting performance was evaluated with reference to some usual indicators. In particular, denoting with X_t the true observation of the variable X at time t and with $\hat{X}_{st}$ the s-step ahead forecast for X_t, and with $1, \ldots, \tau$ the interval of evaluation, I calculated the following measures:

- mean error (ME): $\frac{1}{\tau} \sum_{t=1}^{\tau} \left(X_t - \hat{X}_{st} \right)$;
- mean absolute error (MAE): $\frac{1}{\tau} \sum_{t=1}^{\tau} \left| X_t - \hat{X}_{st} \right|$;
- root mean squared error (RMSE): $\sqrt{\frac{1}{\tau} \sum_{t=1}^{\tau} \left(X_t - \hat{X}_{st} \right)^2}$;

- median error (MedE): $Med\left\{X_t - \hat{X}_{st}\right\}_{t=1,\dots,\tau}$;
- median absolute error (MedAE): $Med\left\{\left|X_t - \hat{X}_{st}\right|\right\}_{t=1,\dots,\tau}$.

In table 1 I summarize the main results obtained using the RMSE. In particular the ratio of FCR root mean squared forecasting error over that of the ADL model is given, so that a value less than 1 in the table means a better performance of the FCR model.

Moreover the statistical significance of the results obtained was assessed by means of the variant to the Diebold-Mariano test proposed by Harvey et al. (1998). Let us denote with e_{it} the forecasting errors stemming from model i at time t, then when comparing τ forecasts stemming from two competing models i and j the Diebold-Mariano statistics is:

$$DM = \frac{\tau^{-1}\sum_{t=1}^{\tau}[g(e_{it}) - g(e_{jt})]}{\sqrt{\tau^{-1}2\pi \hat{f_d(0)}}} \tag{9}$$

where $f_d(0)$ is a consistent estimate of the spectral density of $\tau^{-1}\sum_{t=1}^{\tau}[g(e_{it}) - g(e_{jt})]$ at frequency 0. The variant of the test proposed by Harvey et al. (1998) considers also the forecasting horizon s:

$$DM^* = \left[\frac{\tau + 1 - 2s + \tau^{-1}s(s-1)}{\tau}\right]^{1/2} DM. \tag{10}$$

The authors propose to compare such a statistic with the Student t distribution with $\tau - 1$ degrees of freedom. In this paper I consider the function $g(\cdot) = |\cdot|$.

The results of table 1 show an impressive improvement in forecasting performance of the FCR model with respect to the linear one, especially at lower level of d and at shortest horizons. The best improvement is perhaps achieved when PT is used as the state variable; in the short horizons and for values of d less than 4 the forecasting performance improvement is generally around 20%. Moreover, at 1-step ahead horizon the results are always significant at conventional values. Only in one case the ADL model outperforms the FCR one. In the case of PL, the improvement is slightly less pronounced, even though it is still true that the non-linear non-parametric model always outperforms the linear one; the improvement is significant, according to the DM test, especially for values of d larger than 2. When OB is used as the state variable, the results are less favourable to the FCR model, but the general pattern is similar to those observed when PL is used.

Table 1: Ratio of FCR/ADL RMSE by state variable value (d) and forecasting horizon (s)

	s=1	s=2	s=3	s=4	s=5	s=6
PT						
d=0	0.798**					
d=1	0.819*	0.816				
d=2	0.784**	0.783	0.921			
d=3	0.798**	0.794	0.902	1.045		
d=4	0.918**	0.929*	0.914	0.960	0.957	
d=5	0.957*	0.964	0.933	0.967	0.963	0.964
PL						
	s=1	s=2	s=3	s=4	s=5	s=6
d=0	0.760					
d=1	0.884	0.926				
d=2	0.921	0.970	0.998			
d=3	0.928**	0.943**	0.914**	0.951**		
d=4	0.961*	0.968**	0.932**	0.965*	0.965**	
d=5	0.976	0.976*	0.941**	0.973*	0.969**	0.961**
OB						
	s=1	s=2	s=3	s=4	s=5	s=6
d=0	0.866					
d=1	0.892	0.947				
d=2	0.927	0.962	1.070##			
d=3	0.968	0.999	0.999	1.068		
d=4	0.954**	0.954**	0.933**	0.980	1.010	
d=5	0.958*	0.953	0.937*	0.982	1.003	1.012

* denotes FCR forecast are better than ADL forecasts at 10% confidence level, ** at 5%. # is used when ADL forecasts are significantly better.

Another evaluation criteria employed is the fraction of corrected directional forecasts, defined as:

$$\frac{1}{\tau}\sum_{t=1}^{\tau}\mathbb{I}_{(X_t-X_{t-s})(\hat{X}_{st}-X_{t-s})=1}.$$

In table 2 I show for each indicator and each value of delay d and forecasting horizon s, the ratio between the fraction of correct directional forecasts of FCR model over those stemming from the ADL model. A value larger than 1 means that the FCR model performs better than the ADL one. For each state variable considered (PT, PL, and OB) there are 21 possibile comparisons. The results mimic those obtained with the RMSE comparison: when PT is taken as the state variable in 10 cases out of 21 the FCR model shows a better performance, in 3 cases it is the same, in 8 cases the ADL obtains better results. Turning to PL the superiority of the FCR model is sharper, with 12 cases favouring it and 5 cases favouring the ADL, the remaining 4 being equal. Only when OB is taken as the indicator, the performance of FCR model is worst than the ADL, with 9 cases favouring the first against 10 for the latter, and two cases being equal.

Table 2: Ratio of FCR/ADL correct directional forecasts by state variable value (d) and forecasting horizon (s)

			PT			
	s=1	s=2	s=3	s=4	s=5	s=6
d=1	1.115					
d=2	1.115	1.103				
d=3	1.120	1.103	1.034			
d=4	0.900	1.000	1.000	1.032		
d=5	0.963	0.935	0.931	1.000	1.028	
d=6	0.963	0.935	0.931	1.030	1.028	0.944
			PL			
	s=1	s=2	s=3	s=4	s=5	s=6
d=1	0.867					
d=2	0.966	0.939				
d=3	0.964	0.829	1.038			
d=4	1.000	1.074	1.080	1.033		
d=5	1.000	1.074	1.080	1.067	1.057	
d=6	1.000	1.111	1.080	1.100	1.057	1.000
			OB			
	s=1	s=2	s=3	s=4	s=5	s=6
d=1	0.931					
d=2	1.000	0.909				
d=3	0.929	0.967	0.867			
d=4	0.963	1.000	1.083	1.067		
d=5	0.929	0.967	1.040	1.103	1.088	
d=6	0.964	1.036	0.960	1.143	1.057	1.029

7 CONCLUSIONS

A non-linear non-parametric framework has been used to model the relationship between industrial production and some related indicators. The model has a nice interpretation, as the related indicators are directly interpretable as indicators of the business cycle state.

Forecasting errors up to 6-step-ahead, as compared to a baseline autoregressive distributed lag model, show in general a valuable reduction in magnitude using the non-linear non-parametric model, by some common measures. Moreover, the differences reported are sometimes statistically significant according to the Diebold-Mariano test.

Further elaborations could include:

- a deeper analysis of the forecasting performance results, developing indicators more suited to the case at hand;
- the calculation of a direct multi-step forecast (as in Harvill and Ray (2005));
- the calculation of density forecasts;
- as far as the modelling is concerned, the calculation of generalised impulse response function (GIRF) could give some more insights into the dynamic properties of the model.

REFERENCES

Bradley, M. D. and Jansen, D. W. (2004). Forecasting with a nonlinear dynamic model of stock returns and industrial production. *International Journal of Forecasting*, 20(2):321–342.

Bruno, G. and Lupi, C. (2004). Forecasting industrial production and the early detection of turning points. *Empirical Economics*, 29(3):647–671.

Cai, Z., Fan, J., and Li, R. (2000a). Efficient estimation and inferences for varying-coefficient models. *Journal of the American Statistical Association*, 95(451):888–902.

Cai, Z., Fan, J., and Yao, Q. (2000b). Functional-coefficient regression models for nonlinear time series. *Journal of the American Statistical Association*, 95(451):941–956.

Fan, J. and Yao, Q. (2003). *Nonlinear Time Series: Nonparametric and Parametric Methods*. Springer, New York.

Franses, P. H. and van Dijk, D. (2005). The forecasting performance of various models for seasonality and nonlinearity for quarterly industrial production. *International Journal of Forecasting*, 21(1):87–102.

Harvey, D. I., Leybourne, S. J., and Newbold, P. (1998). Tests for forecast encompassing. *Journal of Business and Economic Statistics*, 16:254–259.

Harvill, J. L. and Ray, B. K. (2005). A note on multi-step forecasting with functional coefficient autoregressive models. *International Journal of Forecasting*, 21:717–727.

Huh, C. (1998). Forecasting industrial production using models with business cycle asymmetry. *Economic Review*, (1):29–41.

Jagric, T. (2003). A nonlinear approach to forecasting with leading economic indicators. *Studies in Nonlinear Dynamics & Econometrics*, 7(2):1135–1135.

Marchetti, D. J. and Parigi, G. (2000). Energy consumption, survey data and the prediction of industrial production in Italy: a comparison and combination of different models. *Journal of Forecasting*, 19(5):419–440.

Öcal, N. (2000). Nonlinear models for U.K. macroeconomic time series. *Studies in Nonlinear Dynamics & Econometrics*, 4(3):123–135.

Osborn, D. R. and Matas-Mir, A. (2003). The extent of seasonal/business cycle interactions in European industrial production. Discussion Paper 38, The University of Manchester.

Pappalardo, C. (1998). La stagionalità nelle serie ISCO. Rassegna di lavori dell'ISCO 3, ISCO.

Priestley, M. B. (1980). State-dependent models: A general approach to nonlinear time series analysis. *Journal of Time Series Analysis*, 1:47–71.

Proietti, T. (1998). Seasonal heteroscedasticity and trends. *Journal of Forecasting*, 17:1–17.

Siliverstovs, B. and Dijk, D. V. (2003). Forecasting industrial production with linear, nonlinear and structural change models. Econometric Institute Report 321, Erasmus University Rotterdam, Econometric Institute.

Simpson, P. W., Osborn, D. R., and Sensier, M. (2001). Forecasting UK industrial production over the business cycle. *Journal of Forecasting*, 20(6):405–24.

Teräsvirta, T., Lin, C., and Granger, C. (1993). Power of the neural network linearity test. *Journal of Time Series Analysis*, 14:209–220.

Venetis, I. A., Peel, D. A., and Paya, I. (2004). Asymmetry in the link between the yield spread and industrial production: threshold effects and forecasting. *Journal of Forecasting*, 23(5):373–384.

A ESTIMATION EXAMPLE

In this appendix some estimation results are shown for a particular FCR model. This particular model is taken just as an example, which I think is useful to shed some light also on the use of such models as descriptive devices. The example reported refers to the use of PL as the state variable, with d=1. In this case the lags selected are: 1 to 5, 12.

The first step, once the model has been defined, is to get an estimate of the bandwidth to be used. In figure 3 I report the plot with the results from the cross-validation criterion adopted.

Figure 3: Average prediction error (APE) vs. bandwidth

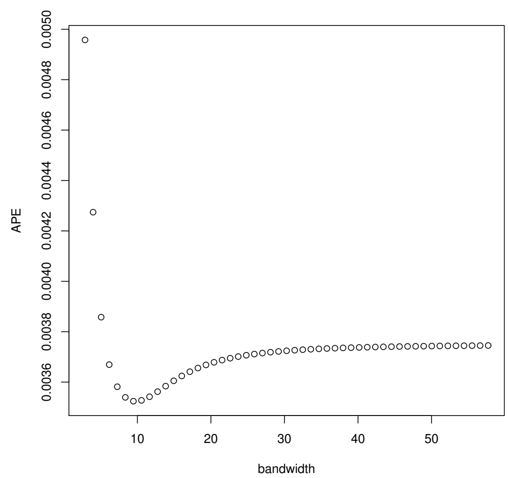

A plot of the six coefficient functions is shown in figure 4; it is possible to see that they show a considerable variability.

In the following plots I try to summarize these results, taking advantage that for a *given* value of the state variable, the model is a linear AR model.

In particular, in figure 5 the sum of the coefficients functions is reported in a scatter against the value of the state variable. In a linear AR model the sum of the coefficients is often considered as an indicator of persistence. In this case the plot shows that for central values of the state variable persistence is lower, while it is higher for low values.

Another way to summarize the results is shown in figure 6. In this case I consider three values of the state variable (the first, second and third quartile) and consider the spectral densities associated to the three linear AR models implied by those values. It is possibile to appreciate the increasing persistence associated with diminishing values of the state variable.

The previous evidence could be interpreted as a possible sign of asymmetry over the business cycle.

Figure 4: Coefficient functions for the model with PL_{t-1} as the state variable.

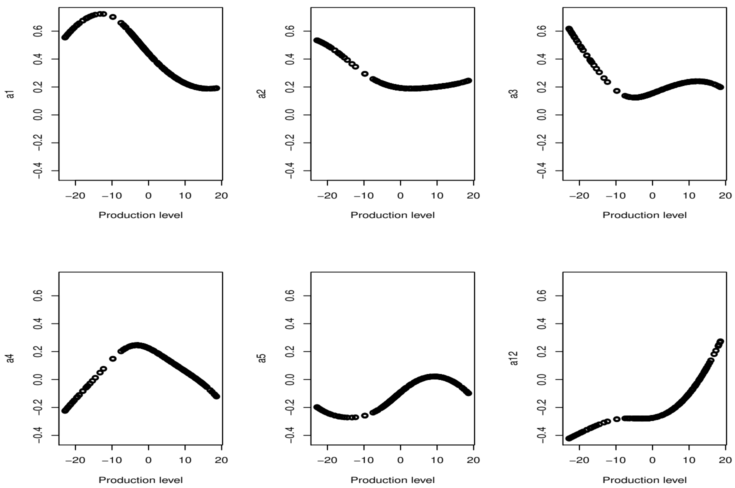

Figure 5: Sum of coefficient functions for the model with PL_{t-1} as the state variable.

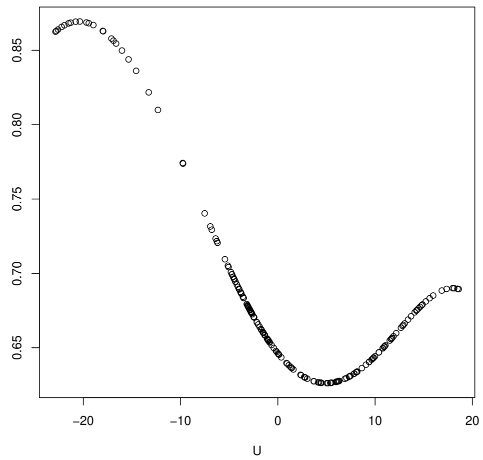

Figure 6: Spectral density of the three autoregressive models associated with the first, second and third quartile of the state variable.

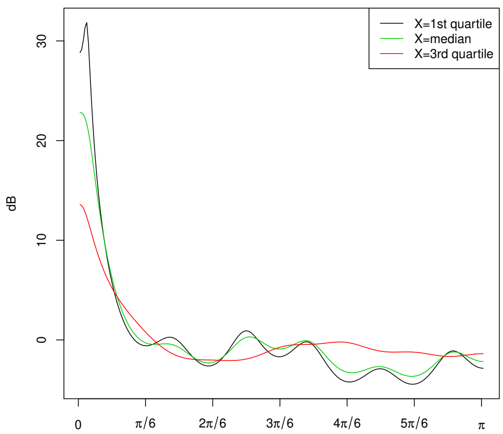

The residuals of the model can be checked to see if the are white noise. In table 3 some test statistics about the residuals are reported. Normality tests reject the null hypothesis mainly due to 4 outlying observations; the same statistics do not show departure from normality when these outliers are removed. Runs test and Ljiung-Box test cannot reject the null of white noise residuals. Moreover they appear to be linear at conventional levels according to the tests proposed by Teräsvirta and White (Teräsvirta et al., 1993).

Table 3: Main residual diagnostic for the FCR model with LP_{t-1}.

Test	p-value
Jarque-Bera:	0.00
Shapiro-Wilk:	0.00
Jarque-Bera(*):	0.28
Shapiro-Wilk(*):	0.16
Runs:	0.20
Teräsvirta nn:	0.09
White nn:	0.06
Ljiung-Box (1):	0.61
Ljiung-Box (3):	0.76
Ljiung-Box (12):	0.27

(*) with 4 outliers removed

In the end, following Fan and Yao (2003) a test can be carried out to verify the null hypothesis of a linear AR(p) model against that of a FCR model. The test proposed is based on the residual sum of squares of the two models:

$$T_{n,6} = \frac{T-p}{2}(RSS_0 - RSS_1)/RSS_1 \qquad (11)$$

where RSS_0 is the residuals sum of squares of the AR(p) model and RSS_1 that of the FCR model. The distribution of the test can be found bootstrapping the residuals from the FCR model. In the case of LP_{t-1} the test rejected the null of linearity with a p-value of 0.001 and 1000 bootstrap replications.

Working Papers available:		
n. 75/07	R. BASILE	Intra-distribution dynamics of regional per-capita income in Europe: evidence from alternative conditional density estimators
n. 76/07	M. BOVI	National Accounts, Fiscal Rules and Fiscal Policy Mind the Hidden Gaps
n. 77/07	L. CROSILLA S. LEPROUX	Leading indicators on construction and retail trade sectors based on ISAE survey data
n. 78/07	R. CERQUETI M. COSTANTINI	Non parametric Fractional Cointegration Analysis
n. 79/07	R. DE SANTIS C. VICARELLI	The “deeper” and the “wider” EU strategies of trade integration
n. 80/07	S. de NARDIS R. DE SANTIS C. VICARELLI	The Euro’s Effects on Trade in a Dynamic Setting
n. 81/07	M. BOVI R. DELL’ANNO	The Changing Nature of the OECD Shadow Economy
n. 82/07	C. DE LUCIA	Did the FED Inflate a Housing Price Bubble? A Cointegration Analysis between the 1980s and the 1990s
n. 83/07	T. CESARONI	Inspecting the cyclical properties of the Italian Manufacturing Business survey data
n. 84/07	M. MALGARINI	Inventories and business cycle volatility: an analysis based on ISAE survey data
n. 85/07	D. MARCHESI	The Rule Incentives that Rule Civil Justice
n. 86/07	M. COSTANTINI S. de NARDIS	Estimates of Structural Changes in the Wage Equation: Some Evidence for Italy
n. 87/07	R. BASILE M. MANTUANO	La concentrazione geografica dell’industria in Italia: 1971-2001
n. 88/07	S. de NARDIS R. DE SANTIS C. VICARELLI	The single currency’s effects on Eurozone sectoral trade: winners and losers?
n. 89/07	B.M. MARTELLI G. ROCCHETTII	Cyclical features of the ISAE business services series

Working Papers available:

n. 90/08	M. MALGARINI	Quantitative inflation perceptions and expectations of Italian Consumers
n. 91/08	P. L. SCANDIZZO M. VENTURA	Contingent valuation of natural resources: a case study for Sicily
n. 92/08	F. FULLONE B.M. MARTELLI	Re-thinking the ISAE Consumer Survey Processing Procedure
n. 93/08	M. BOVI P. CLAEYS	Treasury v dodgers. A tale of fiscal consolidation and tax evasion
n. 94/08	R. DI BIASE	Aliquote di imposta sul lavoro dipendente: analisi per figure tipo e con dati campionari
n. 95/08	M. BOVI	The “Psycho-analysis” of Common People’s Forecast Errors. Evidence from European Consumer Surveys
n. 96/08	F. BUSATO A. GIRARDI A. ARGENTIERO	Technology and non-technology shocks in a two-sector economy
n. 97/08	A. GIRARDI	The Informational Content of Trades on the EuroMTS Platform
n. 98/08	G. BRUNO	Forecasting Using Functional Coefficients Autoregressive Models
n. 99/08	A. MAJOCCHI A. ZATTI	Land Use, Congestion and Urban Management
n. 100/08	A. MAJOCCHI	Theories of Fiscal Federalism and the European Experience
n. 101/08	S. de NARDIS C. PAPPALARDO C. VICARELLI	The Euro adoption’s impact on extensive and intensive margins of trade: the Italian case
n. 102/08	A. GIRARDI P.PAESANI	Structural Reforms and Fiscal Discipline in Europe
n. 103/08	S. TENAGLIA M. VENTURA	Valuing environmental patents legal protection when data is not availableby
n. 104/08	P. L. SCANDIZZO M. VENTURA	A model of public and private partnership through concession contracts

Working Papers available:

n. 105/08	M. BOSCHI A. GIRARDI	The contribution of domestic, regional and international factors to Latin America's business cycle
n. 106/08	T. CESARONI	Economic integration and industrial sector fluctuations: evidence from Italy
n. 107/08	G. BOTTONE	Human Capital: an Institutional Economics point of view
n. 108/09	T. CESARONI M. MALGARINI L. MACCINI	Business cycle stylized facts and inventory behaviour: new evidence for the Euro area
n. 109/09	G. BOTTONE	Education in Italy: is there any return?
n. 110/09	S. de NARDIS C. PAPPALARDO	Export, Productivity and Product Switching: the case of Italian Manufacturing Firms
n. 111/09	M. BOVI R. CERQUETI	Why is the Tax Evasion so Persistent?
n. 112/09	B. ANASTASIA M. MANCINI U. TRIVELLATO	Il sostegno al reddito dei disoccupati: note sullo stato dell'arte. Tra riformismo strisciante,inerzie dell'impianto categoriale e incerti orizzonti di *flexicurity*
n. 113/09	A. ARGENTIERO	Some New Evidence on the Role of Collateral: Lazy Banks or Diligent Banks?
n. 114/09	M. FIORAMANTI	Estimation and Decomposition of Total Factor Productivity Growth in the EU Manufacturing Sector: a Stochastic Frontier Approach
n. 115/09	E. DE ANGELIS C. PAPPALARDO	String Matching Algorithms. An Application to ISAE and ISTAT Firms' Registers
n. 116/09	L. CROSILLA S. LEPROUX M. MALGARINI F. SPINELLI	Factor based Composite Indicators for the Italian Economy
n. 117/09	G. BOTTONE	A new notion of progress: Institutional quality
n. 118/09	F. DI NICOLA	Tassazione e Sostegno del Reddito Familiare