Forecasting Spanish GDP Growth During the Debt Crisis

Group 5

April 11, 2013

1 Motivation

Economic growth forecasting plays an important role in economic policy decisionmaking as well as its potential outcome. Within the context of the Eurozone sovereign debt crisis, understanding the growth trajectory of the monetary union as a whole–as well as, and perhaps more importantly, that of individual member states–is a crucial element in the crisis resolution framework.

In this study, we use a principal component estimator (PCE) to forecast the quarter-to-quarter growth in the Spanish GDP and compare this to an ARMA(4,2) model. We choose to limit the sample size to the period 1977:2 to 2012:4 to keep as many variables as possible in the analysis. We extend the PCE by allowing for non-linearities in two different ways. Allowing for non-linearities in the principal component factors does not improve the forecast, however, allowing for non-linearities in the lagged GDP growth improves the forecast during crisis times.

Forecasts serve as an instrument in formulating appropriate fiscal polices across the Eurozone. The ongoing sovereign debt crisis has ushered in a wave of austerity measures intended to combat the monetary union's inflated deficit and debt levels. Countries are particularly aiming to bring down their deficit-to-GDP ratios and debt-to-GDP ratios to the respective 3% and 60% thresholds stipulated in the European Stability and Growth Pact. Budgetary policy therefore needs to be very precise in setting cuts that will achieve these targets. Being foreward looking, however, these policy measures rest on an appropriate growth estimate for the quarters and year ahead. Forecasting is consequently sought to deliver a reliable indicator of economic growth upon which fiscal policy can be built.

Announcements of forecasts–as well as whether an economy meets the forecasted economic growth level–further provide important signals to market investors and consumers as they make consequential decisions within the economy. Financial investors may use forecasts in their decisionmaking as they provide an outlook for the future growth possibilities of an economy–and therefore the attractiveness of investing there. They also are used to speculate ECB policies like interest rate setting as the central bank has been looked upon as a dampener of the crisis with the ability to continue cutting interest rates. Forecasts may also contribute to the outlook that is presented to the citizens and thus consumers in the economy, thereby impacting consumer and business confidence figures that influence investment and consumption trajectories. During the Eurozone crisis in particular, numerous forecasts that were subsequently not met provided negative signals that can influence the very fragile and volatile development of market confidence. Eurozone policymakers have been continuously setting out to gain credibility in the face of financial markets, considered crucial to bringing down the soaring interest rate spreads that feed back into their economic growth trajectories with the threat of high debt and low growth cycles or "bad equilibria."

This study presents a methodological approach to forecasting quarterly growth in Spain for 2013. At this stage in the crisis, Spain is viewed through an especially critical lens as it has yet to show clear results for a regain in growth and competitiveness since its recent housing crisis boom–even as recently as dicussed in this week's report by the European Commission (2013). The fourth largest economy in the Eurozone, Spain has kept markets wondering whether or not it will need to apply for an aid package or access to extraordinary ECB measures like the Outright Monetary Transactions mechanism. A forecast of economic growth below or similar to the one of last year is a potential trigger for further challenges to a regain of confidence in not only Spain's recovery, but in the credibility of Eurozone policymakers' crisis resolution strategy more broadly. The forecast is furthermore a key element in Spain's budgetary plans.

1.1 Methodological overview

To optimize forecasting effectiveness, especially with regard to the use of information, a variety of methods have been developed to incorporate large quantities of data in the process of macroeconomic prediction. Working within a data-rich environment would imply running into the so-called curse of dimensionality if traditional time series models try to use all predictors at once (Liebermann, 2012). Two broad strands of methods have been developed, one of which selects the most important variables (such as with automated model selection procedures) and subsequently uses them in traditional forecasting models. A second category uses all variables rather than selecting the optimal ones, and includes methods like factor-based models, partial-least squares, and factor combination (Groen and Kapetanios, 2009). Our study falls into the latter approach as we aim to use as much information available as possible in an efficient manner.

To reduce the dimension of our expansive data set and organize it in a more simple and interpretable manner, we employ factor analysis that extracts a number of common factors using principal components from the data set to use in forecasting. Stock and Watson (2002a,b) popularized the use of principal components to estimate factor models in particularly a data-rich environment (Groen and Kapetanios, 2009). To use this method we first address the issue of missing observations and transform our variables into stationary ones by taking first-differences of any non-stationary variables.

To address non-linearity, we combine the use of principal components with an additively separable semiparametric model rather than solely use a basic linear h-step ahead predctive regression. We employ a Sieve estimation as our econometric framework as it allows for a different treatment of tail regions than the one a linear framework lends itself to. This is motivated by the switching methodology used in Auerbach and Gorodnichenko (2012) to estimate state-dependent fiscal multipliers. In their paper, the possibility that the size of fiscal multipliers varies during the business cycle motivates the type of method chosen to estimate spending multipliers. Similarly, in our approach the possibility that changes to an economic variable such as government expenditure can in times of crisis have a more significant impact on the growth forecast motivates our use of a non-linear econometric setup.

Compared to the use of linear methods, our approach may limit the chances of, for example, over-estimating economic growth potential in times of crisis. As explained in the previous section, this would help policymakers make more realistic fiscal policy plans that offset debt and promote further economic growth. With respect to the political election cycle, politicians may especially exploit forecasts that are in a sense over-optimistic to assure voters that the economic outlook is ameliorating over the course of their term, so constructing a more sensitive forecast for a crisis period may be especially beneficial. Finally, the announcement of an over-optimistic forecast that needs to be adjusted because tail events contribute to lower growth carry costly signals (for example with interest rate spikes) when markets re-asses their view of the economy.

To assess the performance of our forecasting methodology, we compare it with that of a simple ARMA (4,2) model.

2 Data

The data provided in the case consist of quarterly data from the period 1970:1 to 2010:4 for a long list of variables. Before analyzing the data some data transformations are required. The dependent variable, quarterly growth rate of real GDP, is created by deflating nominal GDP and taking first differences.

Next, we disregard variables that are too similar to other variables in the data set. For instance, we exclude variables in local currency when they also appear in US dollar. When variables are listed in both values and volumes we exclude the volumes as we also include the deflators as regressors. In total we disregard 20 variables.

Third, we test for stationarity of the variables using the Augmented Dickey-Fuller unit root test and take first differences of all non-stationary variables. When we take first difference we loose the information about the levels of the variables which could potentially be important for forecasting. Alternatively, we could have explored whether some of the variables were trend stationary in order to keep the levels in the model along with a trend, but due to time constraints we chose not to pursue this further. Once taking time differences we further tested the stationarity of the first differences in case any of the variables would show signs of I(2)-ness. We found the first difference of fixed capital formation in housing to be non-stationary, however, as this variable is only available after 1999 we do not include it in the analysis (see below).

Fourth, we standardize all regressors by subtracting the sample mean and dividing by the standard deviation. Finally, we need to consider the missing values in the variables. This is done in the next section.

3 Choice of observations and variables

Having transformed our data, we now need to decide how we handle the (some what) many missing values in our data set. There is a difficult trade-off when deciding how to reduce the data set when handling the problem of missing values.[1] On one side, we wish to keep as many time-periods as possible (i.e. have a long sample) in order to obtain as much power as possible in our estimation and have a well calibrated model. On the other hand, there are quite a few variables that have missing observations in the initial periods. Thus, keeping more years in our sample will entail loosing more variables in the analysis. Therefore, the somewhat naive method of deleting all variables that have missing values in any of the periods will be very costly in terms of loosing information from important variables throughout the sample. Further, keeping the observations in the initial part of the sample (at the cost of deleting many variables) does not contain important information with respect to forecasting in the later part in the sample. Put differently, knowing the values of the variables in the early 70's is not crucial to forecasting economic growth in the the new millennium compared to having more variables providing important information throughout the sample.[2]

Therefore, it is crucial when deciding how to balance this trade-off between many years and many variables in the sample that one uses economic theory as a guide to distinguish between key variables that are essential for forecasting GDP, and variables that are less costly to neglect from the analysis and thereby avoiding restricting the sample length further.

We have chosen to cut of the sample from 1977:2. By doing so, we keep the following variables that would otherwise have been excluded; CB, CBD, CBG-DRR, NTR, NTRD, PMGSX, PMNW, PXGSX, PXNW, RPMGS, RPXGS, LF, UN, UNR and IRS. We will now briefly highlight some of the variables that we think are important for forecasting economic growth thus motivating

[1] There are of cause many other ways to handle the missing data besides limiting the data set. One could try to impute the data, seek other sources of data (A lot of the variables are key macroeconomic data that are published by several different institutions and one could imagine that some of the data could be found elsewhere) etc. Due to time constraints these options does are not pursued.

[2] These considerations about the naive approach of deleting all variables that have missing values for any of the periods was confirmed when we ran our model on such a data set. We have reported this output in Figure **6** in the appendix. This figure clearly, when you compare them to our results below, confirm that handling the missing values by deleting all variables that contain them is undesirable.

our decision to cut the sample at 1977:Q2. By choosing this sample we keep information about:

- The current account balance: it may not be possible to a priory determine weather the current account balance is necessarily good or bad for growth [3], but it can reflect underlying economic trends that are important for economic growth. Thus, the variables can capture important underlying economic conditions.
- Prices and volumes for imports and export: Keeping the information about the prices of different import and export categories as well as information about export performance could be an important determinant for economic growth (and thus for forecasting economic growth) by, among other things, providing important information about the competitiveness.
- The labour market: Keeping the variables labour force, unemployment level as well as the unemployment rate is obviously important for determining economic growth as they together describe the labour market[4]
- The short term interest rate: Keeping the short term interest rate is very important when forecasting economic growth. The level of the short term interest rate is important for the financing cost of firms, investment as well as for expectation formation which all are key factors effecting economic growth.

By choosing the effective sample to begin in 1977:2 we also excluded 4 variables from the analysis; IHV, CPIH, PCOREH and CPIH_YTYPCT. The housing variable (IHV) could potentially be a very important factor in our analysis as the current deep crises that Spain is in can to some extent be explained by the large housing bobble and the following downturn in the housing market (and thereby the resulting negative consequences for the banking sector and the economy as a whole). However, as we only have observations for the housing market from 2000:1 and onwards it would leave us with an extremely short sample thus hampering the power of our estimations. We therefore choose to exclude this variable from our analysis. Furthermore, excluding the variables describing inflation could also be a problem for our analysis. Again, CPI is not observed before 1992:1 thus it would be very costly in terms of lost years to include it in our sample. Furthermore, we believe that the inclusion of the many deflators (which are available for the full sample) will to a large extent capture the movements of prices.

Lastly, we have also excluded HRS, CBRD and CPIDR as they had missing values in the very end of the sample. However, this is not a serious drawback for our analysis as a lot of the information contained in these variables are captured by the remaining variables. Further, these variables are not key determinants of economic growth.

[3] http://www.imf.org/external/pubs/ft/fandd/basics/current.htm

[4] One needs all the variables in the analysis to fully describe the changes in the labour force . For example a falling labour force will not be accounted for in the unemployment rate.

4 The model

Having described our data and having selected a subsample of the data we now wish to set up and describe the model we are using. The fundamental problem, as described by the case maker, is the large number of parameters relative to the degrees of freedom. All the methods for forecasting in a rich-data environment fundamentally seeks to reduce the dimensions of the data keeping as much of the variation and information as possible. We have chosen to use the principle component analysis. Fundamentally, this analysis seeks to describe the variance-covariance structure of the variables by using linear combinations of these variables. The goal is to describe the variation sufficiently using q principle components thereby facilitating forecasting in the rich-data environment.

4.1 Linear Principal Component Estimator (PCE)

The first specification we examine is a linear Principal Component estimator. We consider the model,

$$y_{t+h} = \sum_{h=0}^{p} \alpha_h y_{t-h} + \sum_{j=1}^{q} \beta_j f_j + \varepsilon_t,$$

where f_j is the j'th factor that comes out of a Principal Component Analysis (for this, we use the `pca` procedure in Stata). We choose to include all factors corresponding to eigenvalues larger than unity, which turns out to give us $q = 11$ factors. For the AR part of the PCE model, we choose a lag length of $p = 4$ based on what we learned from our macroeconometric work on the American GDP quarterly data.

4.2 Non-linear Principal Component Estimator (NLPCE)

In this section, we outline the econometric methodology we apply to getting nonlinearities into our forecast model. The intuitive motivation for why we do this can be understood in two ways. Firstly, there is an economic motivation for nonlinearities in the dynamics over time as we have discussed earlier. Secondly, we can graphically illustrate the importance by considering the simple AR(1) model of the development in GDP. This is done in figure 1.

The figure shows a scatter plot of $\Delta \log y_t$ and $\Delta \log y_{t-1}$, where the time unit is a quarter as well as the linear OLS prediction and two linear splines. Each section of the splines covers the same amount of data. Clearly, we see that there seems to be a different coefficient in the center part of the data from the outer parts of the data. In any linear specifications, the estimated coefficient is somehow a mean of these "local coefficients", or rather, of the local derivative over the range of $\Delta \log y_{t-1}$ for whatever true functional form governs the data.

What we propose is the following

$$y_{t+h} = \sum_{h=0}^{p} \varphi_h(y_{t-h}) + \sum_{j=0}^{q} \psi_j(x_{t-j}) + \varepsilon_t.$$

Figure 1: Illustrating nonlinearities in the AR(1)

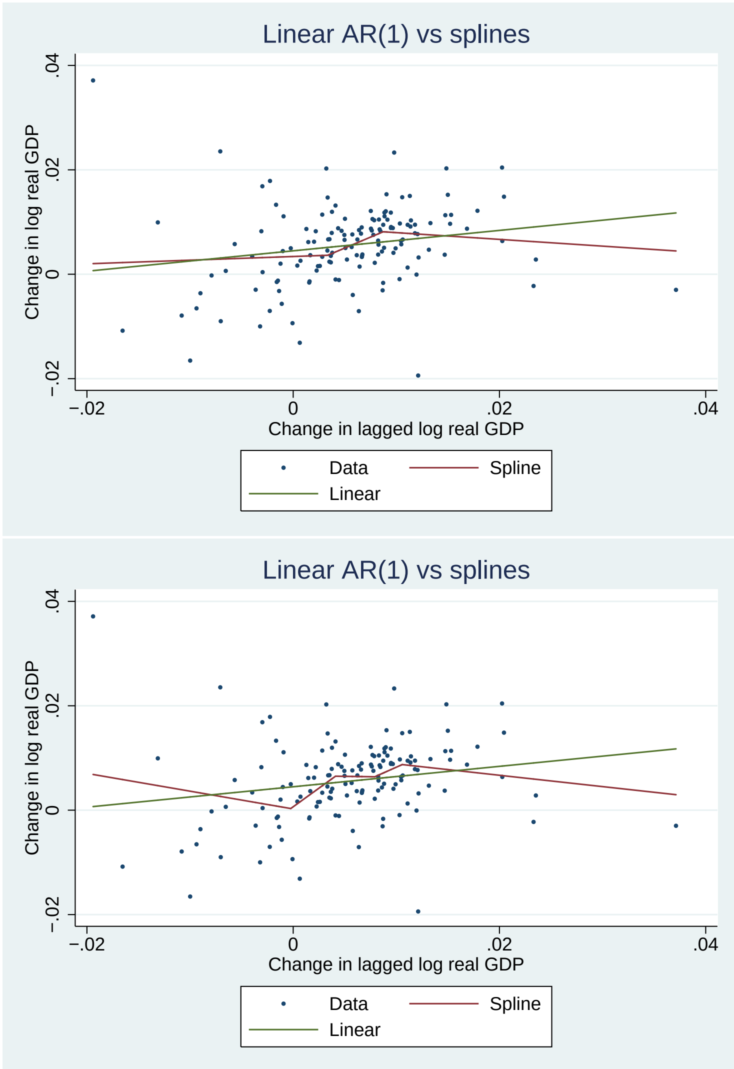

It is the *additively separable semiparametric model*.

As we discussed in the introduction, modelling nonlinearities is particularly importent when the focus is on predicting crises, because we want are looking for predictors of explosive changes in GDP growth, such as the large fall in growth in 2008 as experienced in Spain. Hence, we want to allow for different effects of the variables in the tails. Before we proceed we want to stress that we are perfectly aware that semiparametric techniques are ill-suited at getting good estimates of tail behaviour in small samples, leading to the risk of over-fitting.

The econometric framework we'll use is called *Sieve estimation*. Essentially, we approximate

$$\varphi_h(z) \cong \sum_{k=1}^{K} \varphi_{hk}(z),$$

where $\varphi_{hk}(\cdot)$ is a sequence of base functions and K is the Sieve dimension, which is sort of a bandwidth parameter, calibrated to the data. We have worked with both $K = 3$ and $K = 5$ but found that there were insufficient observations to support the larger dimension. Our choice of Sieve will be a piece-wise linear spline, but if time had allowed it, we would have preferred a more sophisticated polynomial local spline. for the grid placement, we use a data adaptive method based on uniformly placed grid points over the empirical distribution of the given variable. If more time had been given, we would have also preferred to try working with a local linear regression (not a spline, i.e. a global approximation method) which is known to be better at handling tail behaviour than for example the Nadaraya-Watson estimator.

In the implementation, we found it infeasible to allow all variables to be fully nonlinear, so we ended up using two different specifications; In the first specification, we ended up letting $\varphi_h(y) = \alpha_h y$ for all h and $\psi_j(x) = \beta_j x$ for $j = 5, 6, ..., 11$. That is, we only modelled non-linearities in the first 4 factors. In the second, we only let $\varphi_0(\cdot)$ be non-linear and used linear coefficients on all other varibles.

4.3 Inference

An important aspect of the inference is that we are in essence doing a 2-step esetimation. In the first step, we estimate the factors from the Principal Component Analysis by MLE, and in the second step, we insert those in a simple linear regression. Hence, we should correct the standard errors used in the second step.

Inference on Sieve-based estimators is fairly recent, but new research by X. Chen of Yale University indicates that the big virtue of Sieve estimators is that once one has chosen the Sieve dimension, K, the econometrician proceeds with inference as if he were dealing with a *parametric* model. Hence, we just use the t-values from the default output of our regression.

5 Results

5.1 Principal Component Estimator (PCE)

We are now interested in comparing our forecast based on the PCE to both the actual data series and a naive forecast based on an ARMA(4,2) model. We can inspect the three data series graphically in Figure 2 both in the full period and in the 2003:1-2012:4. In particular in the most recent period shown in the lower panel it is clear that the PCE forecast improves a lot on the naive forecast from the ARMA model. The latter does not manage to predict the strong dip in 2008, while the PCE actually does quite a good job in predicting the crisis.

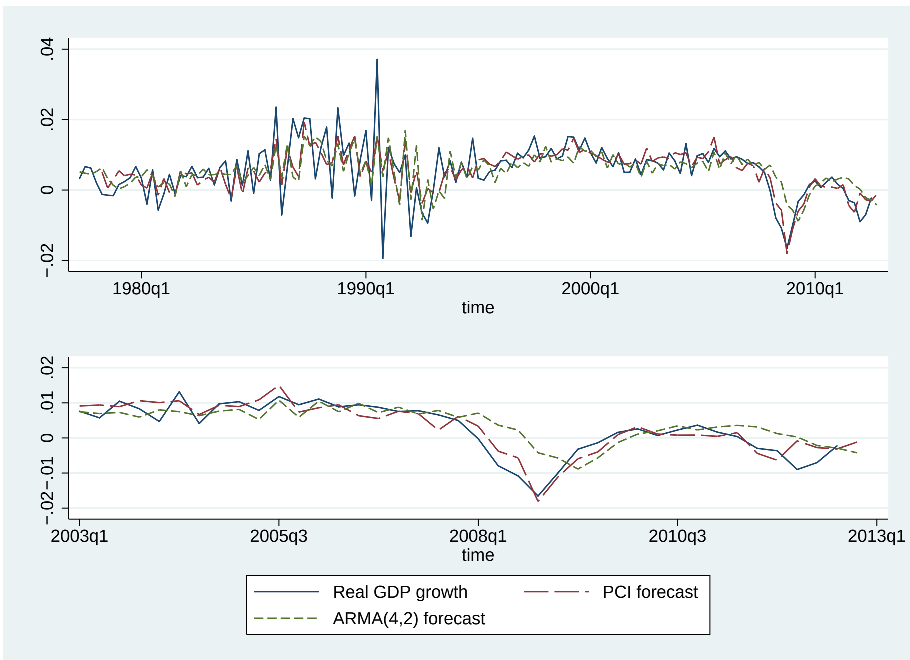

Figure 2: Actual and forecasted real GDP growth for 1977:2-2012:4 (top panel) and 2003:1-2012:4 (bottom panel).

Moving on from the graphical inspection we assess the quality of the prediction by comparing the Root Mean Square Error (RMSE) and Mean Absolut Error (MAE) of the ARMA model and the PCA. Table 1 shows the RMSE for the full period and the most recent period. In both periods the PCE model has a lower RMSE than the ARMA implying that it the PCE is better in describing the actual growth in real GDP.

Using the RMSE as a criterion of the forecast accuracy has been criticized by Armstrong (2001). He notes that the forecast error variance is vulnerable to outliers, because the difference between the actual data and the forecast is

Table 1: RMSE from Benchmark Model and Alternatives.

	1977:2-2012:4		2003:1-2012:4	
	RMSE	MAE	RMSE	MAE
ARMA	.00607	.00440	.00476	.00346
PCE	.00551	.00378	.00287	.00236
NL PCE (factors)	.00532	.00375	.00316	.00251
NL PCE (lag)	.00542	.00360	.00277	.00222

squared. Hence, if the data being forecast contain large outliers, other outliers that are less vulnerable to outliers would be preferred. Examples could be absolute error, sum of absolute errors or mean absolute deviation.

5.2 Out-of-sample Prediction Using the Linear PCE

In Figure 3 we show the out of sample predictions for 2013 using th linear PCE. Note that the graph shows the full predictions for 2013 using each of the steps. To get our best predictions for each quarter, use the 1-step prediction for the first, the two-step for the second and so forth. That way, all the data is used for each prediction.

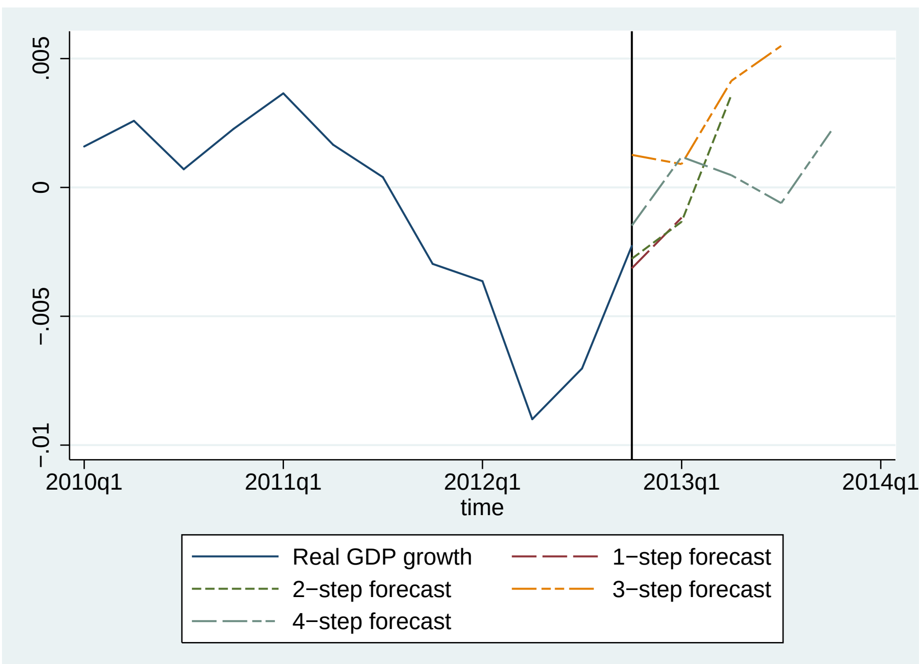

Figure 3: Out of sample predictions

5.3 Non-linear Principal Component Estimator (NLPCE)

Now, we present the results for the non-linear PCE (NLPCE), as described earlier. We have worked with two different variations of of the model; In one (Figure 4) we allow factors 1 through 4 to enter non-linearly and in the other (Figure 5), we allow the first lag of GDP growth to enter non-linearly. The RMSE and MAE are shown in Table 1.

We have done two NLPCE analyses, one with non-linearity in the factors and one with non-linearity in the first lag. We call them NLPCE-F and NLPCE-L here for brevity. We see that going from the linear PCE to NLPCE in the overall sample brings down both the RMSE and MAE. However, when we restrict focus to the period 2003:1–2012:4, the NLPCE-F actually performs worse than the linear PCE. We conclude that this could be due to over-fitting. We conclude that the NLPCE-F performs better in normal times than crisis times, so it is not our preferred model.

For the NLPCE-L, where non-linearity works through the lag, we see from the graph that it follows the slump in 2008 extremely smoothly, both on the way down and back up. We see this particularly clearly in the RMSE results, confirming what the graph shows.

We must stress, however, that we have not checked whether the different forecasts are significantly different from each other, which we conjecture they are not. We are merely stating conclusions about the point estimates. Given more time, a natural extension would be to get standard errors added and potentially eventually take into account noise from the first stage estimation of the factors.

This study applies principal component analysis in both a linear and non-linear framework to forecast quarterly GDP in Spain for 2013. In order to access the forecast accuracy we use both graphical inspection and the RMSE. However, because the RMSE has been criticized in the literature we also compare the MAE to make sure that our conclusions are not affected by large outliers. We saw that the basic principal component estimation performed well compared to a benchmark ARMA model. In particular we saw that it vastly outperforemed the ARMA model during the crisis years 2008 and 2011.

W

When we compare the linear and the non-linear PCE in factors

When we applied the non-linear PCE approach it appears to performs slightly worse compared to the linear principal component analysis in times of normal economic activity, but seems to slightly outperform during crisis-times. However, we do not yet provide significance measures because of time constraints.

With respect to a further sensitivity analysis, we many benefit from comparing the results from using principal components with an application of a partial least squares (PLS) regression. In their comparison of PLS with other forecasting techniques employed in data-rich environments, Groen and Kapetanios (2008) argue that PLS addresses the critique of PC models that the factor carrying a high forecasting power may be dominated by other factors as PC targets the best fit of the entire data set rather than exclusively focusing on the target

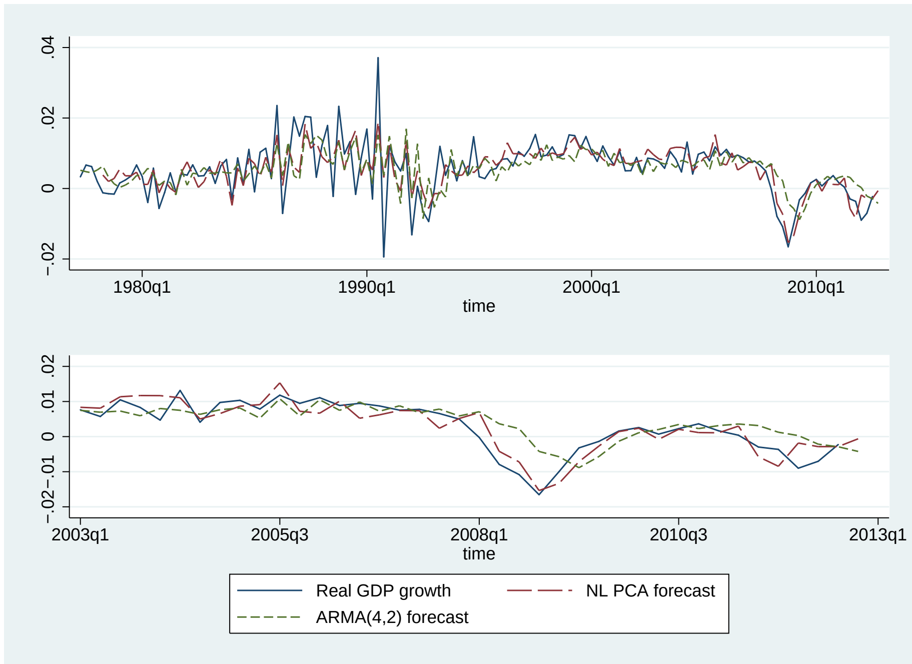

Figure 4: Actual and forecasted real GDP growth for 1977:2-2012:4 (top panel) and 2003:1-2012:4 (bottom panel).

variable. Alternatively, PLS focuses on the target forecast variable by providing "linear, orthogonal combinations of the predictor variables such that the linear combinations maximize the covariance between the target forecast variable and each of the common components constructed from the predictor variables" (Groen and Kapetanios 2008, p.2-3). Complementing our methods with this approach may therefore provide a more robust forecast.

The forecast implies X for the EZ crisis and Spain?

References

Armstrong, J. S. (2001): "Evaluating forecasting methods," *INTERNATIONAL SERIES IN OPERATIONS RESEARCH AND MANAGEMENT SCIENCE*, pp. 443–472.

Auerbach, A. J., and Y. Gorodnichenko (2012): "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, 4(2), 1–27.

European Commission, t. (2013): "In-depth review for Spain," Discussion paper, Commission Staff Working Document, 10 April.

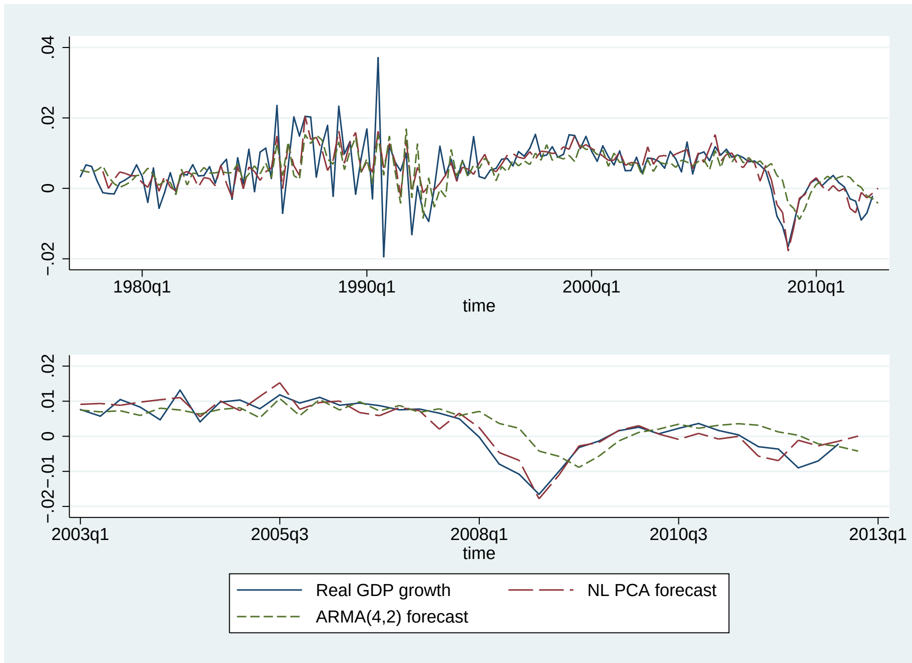

Figure 5: Actual and forecasted real GDP growth for 1977:2-2012:4 (top panel) and 2003:1-2012:4 (bottom panel).

Groen, J., and G. Kapetanios (2009): "Revisiting useful approaches to data-rich macroeconomic forecasting," *FRB of New York Staff Report*, (327).

Liebermann, J. (2012): "Real-time forecasting in a data-rich environment," Discussion paper, Central Bank of Ireland Research Technical Paper 07/RT/12.

Stock, J. H., and M. W. Watson (2002a): "Forecasting using principal components from a large number of predictors," *Journal of the American statistical association*, 97(460), 1167–1179.

——— (2002b): "Macroeconomic forecasting using diffusion indexes," *Journal of Business & Economic Statistics*, 20(2), 147–162.

6 Appendix

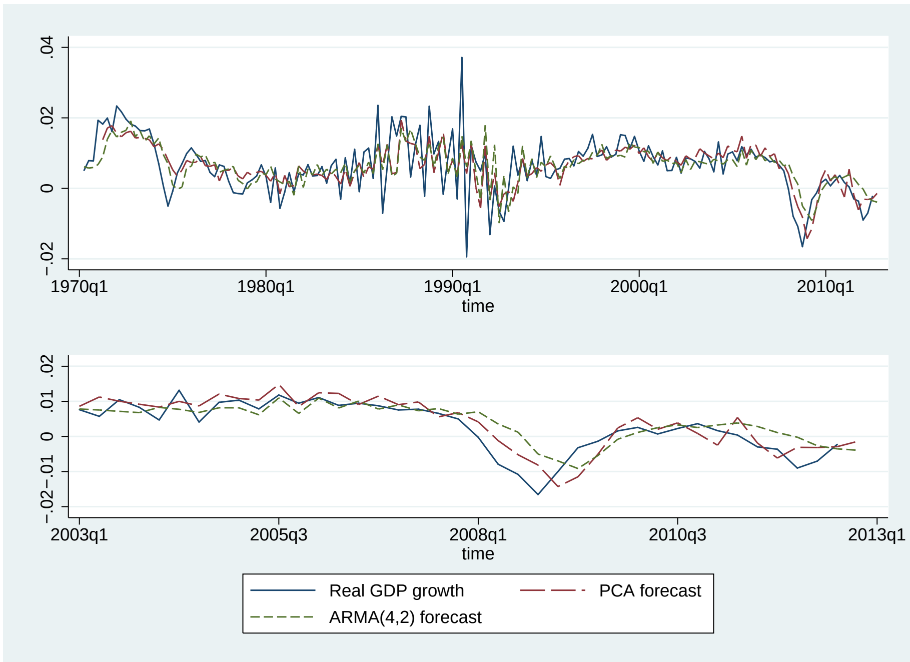

Figure 6: Actual and forecasted real GDP growth for 1970:2-2012:4 (top panel) and 2003:1-2012:4 (bottom panel) with the full number of time periods but a limited number of variables.