Secret Ingredients?

The Informational Contribution of Factors to Standard VAR Analysis*

Ryan Chahrour†

September 15, 2010

Abstract

Recent work with Factor-Augmented Vector Autoregression (FAVAR) suggests that standard VAR analysis can be improved by incorporating the information in a large number of macroeconomic time series. I examine what new information FAVAR factors contribute. Using a sparse modification to principal components, I find that 1) extracted factors and their impulse responses have a natural economic interpretation and 2) a particular small-scale VAR specification closely reproduces the results of standard FAVARs on US data. My results suggest that three leading economic indicators - private payroll employment, the NAPM purchasing managers' index, and housing starts - substantially capture the extra information introduced by FAVAR.

*Thank you to Özge Akinci, Mohammed Bouaddi, Mehmet Caner, Marc Giannoni, Emi Nakamura, Serena Ng, Alexei Onatski, Stephanie Schmitt-Grohé, Dalibor Stanović, Jón Steinsson, Martín Uribe and seminar participants at the 2010 Canadian Economics Association Annual Meeting for very helpful conversations. Original version circulated in November, 2008; revised August, 2010.

†Ph.D Candidate in Economics, Columbia University, 420 W. 118^{th} Street, New York, NY 10027. Contact at rc2374@columbia.edu.

1 Introduction

In recent years, factor models have gained prominence in empirical macroeconomic research. Principal components analysis (PCA) and its close relative, dynamic principal components, are natural ways to reduce the dimensionality of hundreds of macroeconomic or financial time series while preserving, practitioners hope, much of the information contained in those series. The computational convenience of PCA, in particular, has led to an explosion of applications summarizing large sets of macroeconomic time series data with relatively few extracted factors.

Among these methods, factor-augmented vector autoregression (FAVAR), introduced by Bernanke, Boivin, and Eliasz (2005), is especially popular. The Bernanke, Boivin and Eliasz (BBE) method proposes estimating a vector autoregression directly on principal components-estimated factors. Impulse responses are then generated by projecting variables onto the space spanned by the estimated factors. Despite its many advantages (increased degrees of freedom, estimable impulse responses for a huge set of variables), a drawback of FAVAR methodology is that the factors, and therefore the estimated system, have no clear economic interpretation. Additionally, little research has examined the nature of the "new information" incorporated by FAVARs. While FAVARs give macroeconomists theoretically plausible results, why they do so is an open question.

In this paper, I argue that the "extra" information captured by the principal components factors used in FAVAR is effectively spanned by a small set of observable time series. To do this, I use a modification of standard principal components with two key features: 1) it generates factors which depend on a small subset of the variables used in estimation, and 2) it incorporates prior information about variable groupings. This procedure, which I call grouped sparse principal components, or GPCA, is based on the sparse principal components estimator of Zou, Hastie, and Tibshirani (2006) and makes use of the growing statistics literature on $\mathcal{L}1$-norm constrained (lasso) regression.

Because of 1), GPCA-estimated factors capture a smaller portion of variance in the data than do standard PCA factors. Because of 2), each factor is estimated as a linear combination of variables from only one or very few of the *a priori* groups. If the effect of 2) is strong enough and 1) does not eliminate relevant variation in the factors, the factors used in FAVAR can be associated with particular economic concepts, thereby lending interpretation to the system estimated in a FAVAR.

To further my primary claim, I show that from the perspective of dynamics GPCA-estimated factors are essentially identical to standard PCA factors. Furthermore, I show that factors estimated by GPCA are indeed readily interpreted - as production, price, consumption, unemployment, and stock factors. I interpret an additional factor, which corresponds to private payroll employment, the NAPM purchasing managers index, and aggregate housing starts, as a factor of "leading indicators." I show how to use the GPCA factors as a guide to specifying a standard VAR. This specification generates monetary impulse responses that are remarkably close to those generated by a FAVAR estimated on standard PCA factors. Combined, these results lead to my conclusion that the variance captured by PCA factors is easily summarized by very few series.

2 FAVAR Methodology

I briefly review the basic econometric model, and discuss the advantages of FAVAR. This area is nicely surveyed by Stock and Watson (2005b); recent examples include Boivin, Giannoni, and Mihov (2009). In section (2.1), I argue that associating the factors used in estimation with economic concepts may be crucial for the correct identification of structural shocks, a claim which stands in contrast to the current practice of identification in FAVAR models.

Let x_t be an $N \times 1$ vector of observed time series data. Suppose that observations x_t are generated by a factor model, such that

$$x_t = \Lambda f_t + \epsilon_t \tag{1}$$

where f_t is a $K \times 1$ vector of (potentially unobservable) factors, $K << N$. Let ϵ_t be an $N \times 1$ vector of series-specific idiosyncratic shocks. Let X be the $T \times N$ matrix of stacked row vectors, $x_1', x_2', ..., x_t'$ and define the matrix F in a similar manner. I assume that X is transformed to be mean-zero and stationary, with columns of equal variance. Finally, suppose that the dynamics of the factors are given by a structural VAR

$$f_t = \psi(L) f_{t-1} + A^{-1} \zeta_t \tag{2}$$

with the covariance of the structural shocks, $E[\zeta_t \zeta_t']$, assumed to be the identity matrix.

Estimation of equations (1) and (2) typically follows a two-step process. In the first step, estimates of the factors, $\hat{f}_t$, are generated using principal components or a modification thereof. Equation (2) is then estimated using $\hat{f}_t$ in place of f_t. A variety of techniques (discussed below) can then be used to identify the structural matrix A and the corresponding impulse responses.

Proponents of FAVAR argue that its ability to incorporate a larger information set (for example, of the scale used by monetary policy makers) improves its ability to identify a monetary policy shock. In particular, a robust finding is that FAVAR impulse responses demonstrate a significantly reduced price puzzle, the counter-intuitive response, common in standard VARs, of an increase in prices in response to a contractionary monetary policy shock.

FAVARs are also believed to be more easily accorded with theory because they bypass the need to commit to a particular correspondence between theoretical variables and observed series. Furthermore, it is possible to extract an impulse response for any variable in X by projecting it on the driving factors $\hat{F}$. For these reasons, FAVAR has been adopted widely in applied work.

2.1 Approaches to Identification

When latent factors must be estimated, identification of A is rarely straightforward and there does not appear to be a consensus on how to proceed. One strategy, described most generally by Stock and Watson (2005b), combines equations (1) and (2) to consider restrictions on the moving-average representation for the observed series:

$$x_t = \mathbf{B}(L) \zeta_t + \epsilon_t \tag{3}$$

The main advantage of this approach is that it requires no economic assumptions be made on the factors themselves. The main disadvantage is that it requires the analyst to define at least $K(K-1)$ restriction assumptions on the impulse responses of the "informational" variables in X. Yet, each of these series could, in principle, depend on *any combination* of the underlying factors in the economy. Thus, imposing the needed restrictions requires the analyst to ascribe concrete theoretical interpretations to individual variables, thereby negating (at least partially) FAVAR's advantage in avoiding strong assumptions that link observed variables and underlying theoretical concepts.

The alternative approach to identification is to make assumptions on the dynamics of the factors themselves. Of course, doing this requires that at least some factors be identified with economic concepts. Boivin, Giannoni, and Mihov (2009), for example, assume that the stance of monetary policy is measured perfectly by the federal funds rate (FYFF) and impose it as an element of f_t. They then identify A with the assumption that all other (unspecified) factors in the economy respond with a lag to a monetary shock. This approach avoids making strong assumptions on the factor structure of the observed variables (for all but the interest rate), but the validity of the identifying assumptions clearly depends on the economic content embodied in the remaining factors. This issue is difficult to address, however, without specific interpretations for all of the factors.

3 Factor Estimation

Consider the least-squares estimator of equation (1)

$$(\hat{\Lambda}, \hat{\beta}) = \underset{\Lambda,\beta}{\operatorname{argmin}} \sum_{t=1}^{T} ||x_t - \Lambda\beta x_t||^2 \tag{4}$$

where the estimated factors, $\hat{f}_t$, are defined as the linear combinations of the data given by $\hat{\beta}x_t$. The solution to the minimization problem in equation (4) is clearly not unique, since for any nonsingular matrix H, the coefficient matrices $\tilde{\Lambda} \equiv \hat{\Lambda}H^{-1}$ and $\tilde{\beta} \equiv H\hat{\beta}$ give identical residuals.

Identification of the "rotation" matrix H requires K^2 restrictions on Λ and β. The choice of these restrictions, however, does not affect the space that is spanned by the estimated factors.

When the researcher has no interest in the economic content of factors, estimation typically proceeds via standard principle component analysis (PCA). To achieve identification, PCA imposes the (purely statistical) restrictions that $\hat{\Lambda}'\hat{\Lambda}$ is the identity and $\hat{F}'\hat{F}$ is diagonal. The principal components estimator is therefore given by the solution to

$$(\hat{\Lambda}^{PCA}, \hat{\beta}^{PCA}) = \underset{\Lambda,\beta}{\operatorname{argmin}} \sum_{t=1}^{T} ||x_t - \Lambda\beta x_t||^2 \tag{5}$$

subject to

$$\Lambda'\Lambda = I_K \tag{6}$$

$$\beta X' X \beta' = D_K \tag{7}$$

where I_K is the identity matrix of dimension K and D_K is a diagonal square matrix of the same size.

The space spanned by PCA-estimated factors is consistent under standard assumptions (Chamberlain and Rothschild, 1983; Stock and Watson, 2002a). Other methods of factor estimation are also typically based on PCA. In particular, Boivin, Giannoni, and Mojon (2008) use an iterative approach that allows them to impose that particular factors correspond to given observed series, while the remaining factors are estimated by PCA on the space of X not spanned by the observed factors. In the iterative principal components case, only the unobserved columns of $\hat{F}$ are mutually orthogonal.

A disadvantage of standard PCA is that the PCA loadings, the columns of $\hat{\beta}^{PCA}$, are typically nonzero for every series in X. Therefore, though the extracted factors have convenient properties for estimation, they are difficult or impossible to interpret. As discussed above, this poses particular challenges for the researcher seeking to identify structural shocks. A sparse representation of the factors may help the analyst hone in on the key series driving the FAVAR results and better understand the system being estimated.

3.1 Grouped Principal Components Estimator

To address the difficulty in interpreting factors, I suggest an alternative estimator which relaxes the orthogonality restriction on $\hat{F}$ and replaces it with a constraint on the $\mathcal{L}1$-norm of the columns of $\hat{\beta}$. The estimator that I propose here, which I call Grouped Principal Components (GPCA), is a modification of the Sparse Principal Components (SPCA) estimator proposed by Zou, Hastie, and Tibshirani (2006). The key feature of both GPCA and SPCA is that they generate factors for which many of the loadings in $\hat{\beta}$ are *exactly* zero. Because the estimated factors load on relatively few variables, they are more easily linked to particular economic concepts and, thus, are more easily interpreted.

Zou, Hastie, and Tibshirani (2006) show that the PCA problem can be rewritten, replacing the constraint on the factors $\hat{F}$ with a constraint on the $\mathcal{L}2$-norm (ridge penalty) of the rows of β. The SPCA estimator of Zou, Hastie, and Tibshirani (2006), in turn, augments this version of the PCA problem with a penalty on the $\mathcal{L}1$-norm (lasso penalty) of the rows of β. Thus the SPCA estimator is given by the solution to

$$(\hat{\Lambda}^{SPCA}, \hat{\beta}^{SPCA}) = \underset{\Lambda,\beta}{\operatorname{argmin}} \sum_{t=1}^{T} ||x_t - \Lambda\beta x_t||^2 + \lambda_2 \sum_{k=1}^{K} ||\beta_k||^2 + \sum_{k=1}^{K} \lambda_{1,k} ||\beta_k||_1 \tag{8}$$

$$\text{subject to } \Lambda'\Lambda = I_K$$

As is typical in other lasso-type regression problems, larger penalty parameters $\lambda_{1,k}$ yield an increasing proportion of factor loadings that are exactly equal to zero.

When data have natural groupings, I argue that the Zou, Hastie, and Tibshirani (2006) procedure may be improved by incorporating the grouping information in the factor estimation stage. Imagine the econometrician has a prior that separates the data matrix X into G non-intersecting groups, so that $X = \{X^1, X^2, ..., X^G\}$, each group with $N_1, N_2, ..., N_G$ columns respectively. Further, let $\overrightarrow{g}$ be the vector of column indexes of the variables in group g. Finally, let e.g. $\hat{\beta}^{PCA}_{k,l}$ be the k^{th} row and l^{th} column (columns, if l is a vector) of $\hat{\beta}^{PCA}$.

Then, for each factor k, and each group g, generate the weights

$$\hat{w}_{k,g} = \frac{1}{N_g} \sum_{l \in \overrightarrow{g}} \left| \hat{\beta}^{PCA}_{k,l} \right| \tag{9}$$

In words, $\hat{w}_{k,g}$ is the mean of the $\mathcal{L}1$-normed loadings of PCA for group g and factor k.

The GPCA estimator is then given by the solution to

$$(\hat{\Lambda}^{GPCA}, \hat{\beta}^{GPCA}) = \underset{\Lambda,\beta}{\operatorname{argmin}} \sum_{t=1}^{T} ||x_t - \Lambda\beta x_t||^2 + \lambda_2 \sum_{k=1}^{K} \sum_{g=1}^{G} \left|\left| \frac{\beta_{k,\overrightarrow{g}}}{\hat{w}_{k,g}} \right|\right|^2 + \sum_{k=1}^{K} \sum_{g=1}^{G} \lambda_{1,k} \left|\left| \frac{\beta_{k,\overrightarrow{g}}}{\hat{w}_{k,g}} \right|\right|_1 \tag{10}$$

$$\text{subject to } \Lambda'\Lambda = I_K$$

GPCA is thus a two-step estimation procedure very similar to other "weighted" procedures. Weights are generated from simple PCA loadings, and used in the second-step weighted estimation problem. Note that SPCA is a special case of GPCA where all variables are assumed to be part of a single group.

The logic of this approach is as follows: if variables of the same group are known to have similar economic content, then variation in a particular variable from that group is likely to correspond to variation in the factor(s) that most heavily influence that group as a whole. The PCA step gives an initial estimate, $\hat{w}_{k,g}$, of this influence. In the second step, variables from groups with high loadings in the first stage face a smaller penalty, and are therefore more likely to be selected. As with other two-step estimators, this estimator can be iterated to convergence (see appendix (A.1).)

While I do find factors that load on only one or very few groups of variables, nothing in the procedure itself prevents factors from loading on variables in many or even all groups. Such a case will only occur, however, if broadly-based loadings are necessary to capture a significant portion of the variance in the panel. It is hoped, and indeed I find, that this procedure gives factor estimates that preserve the information contained in the standard PCA factors, while facilitating factor interpretation.

Recently, Leng and Wang (2009) have independently developed a similar modification of principal components, which they call "General Adaptive Sparse Principal Components." The estimator I propose here can be seen as a special case of the estimator in their paper, although they do not consider the implications of such grouping information. They do consider the theoretical properties of the estimator in some detail, and I refer the reader to their paper for certain asymptotic results.

A remaining issue is the selection of the size of the lasso-penalty. If it is small, then the resulting factors will be very close to PCA; if the penalty is too large, the estimated factors may miss crucial information. In appendix A, I describe a procedure for parameterizing the sparsity of this estimator using a single parameter, which I call κ. Details on the construction and implementation of the estimator can also be found in the appendix. There, I also show how to extend the procedure to impose a factor or factors that are observed, as has often been done with the interest rate in the FAVAR literature.

4 Data

The results in this paper are based on an updated and expanded version of the BBE dataset. My dataset includes 139 US variables, with monthly observations from 1959:M1 through 2007:M12. Non-stationary variables are transformed to growth rates via log-differencing. The data fall into 14 groups corresponding to their economic significance: output, unemployment, employment, earnings, housing, inventories & orders, stock prices, exchange rates, interest rates, spreads, money and credit, survey expectations, prices, and consumption. I use these concepts to derive my groups for GPCA estimation. Note that other *a priori* groupings are possible. In particular, variables could be grouped by the type of transformation that has been performed. Appendix B gives the variable names and groupings, along with details on the variables sources.

As a robustness check, I perform the same routines on two overlapping 30 year subsamples of the data. The early subsample spans from 1959:M1 through 1989:M12 and the later subsample, from 1977M1 to 2007M12.

5 Results

As a baseline, figure (1) shows impulse responses for industrial production, PCE prices, CPI prices, unemployment, and the federal funds rate to an identified single standard deviation shock to monetary policy for a VAR estimated on industrial production (IPS10), CPI inflation (PUNEW), and the federal funds rate (FYFF), along with the corresponding traditional FAVAR responses.[1] The 3-variable VAR responses are substantially different from the FAVAR-based responses. In particular, prices increase in response to a contractionary monetary shock over the entire period, demonstrating a strong "price puzzle."

To generate impulse responses, I must select identification procedures for both the small-scale VAR and the FAVAR. For the three variable VAR, I use a standard choleski ordering with the interest rate last. For the factor-augmented VAR, I use the Stock and Watson (2005a) procedure described above to impose restrictions on the $\mathbf{B}(\mathbf{0})$ matrix, namely that manufacturing industrial production, the PCE deflator, PCE consumption excluding food, the long term unemployment rate, and total private employment (IPS43, PCE-P, PCC, LHU15, and CES0500) cannot respond contemporaneously to the monetary policy shock. I choose this approach because the most popular alternative, to impose FYFF as a factor and order it last, implies implausible restrictions as I argue below. Of course, I could place restrictions on any number of variables with the Stock and Watson approach, and there is nothing to ensure that any particular variable corresponds directly to the underlying theoretical concept. In this case, however, I will establish a tight relationship between these variables and the estimated factors, thereby rationalizing the choice to place theoretical restrictions directly on these variables.

[1] Note that BBE report impulse responses in terms of standard deviation units, while I report responses in the original scale of the variables. Additionally, responses for variables that were transformed to log-difference for estimation are given here in terms of cumulative responses.

Finally, it may be surprising that the 3-variable VAR shown in figure (1) displays responses for PCE prices and unemployment, in addition to the variables included in the VAR. Note, however, that estimating this standard VAR is equivalent to a FAVAR, with the restriction that the factors $\hat{F}$ span the space given by [IPS10, PUNEW, FYFF]. Thus, the procedure for estimating numerous responses in FAVAR is also feasible for standard VAR specifications.

5.1 GPCA Factors and Estimated Responses

Table (1) displays results for the cumulative percent of explained variance (PEV) for factors estimated via PCA, SPCA, and the proposed GPCA procedure. For all methods, I assume there are seven factors. I study robustness to the number of factors in a separate subsection.

A few observations are worthwhile. First, the the table shows that, while the first seven standard principal components explain 55% of the variation in the dataset, PEV for GPCA falls with κ (the parameter that governs the tradeoff between sparse factor loadings and capturing variance) but only a very small amount, to around 53% when κ is 1. Second, note that total variance captured by the SPCA tends to fall faster than for GPCA as the sparsity parameter increases. Finally, the strict ordering of PEV contributions of each additional PCA factor does not hold for the SPCA or GPCA case. Instead, the PEV contributions are more evenly "spread out" among the factors, reflecting the algorithm's ability to rotate the factors in a way that achieves sparsity without greatly impacting *total* PEV.

Table (2) gives the number of non-zero loadings for each factor for the same values of κ. While PCA loads positively (in absolute value) on all series for all factors, the number of non-zero loadings per-factor is greatly reduced using the sparse methods. Note, however, that the number of series with positive loadings for at least one of the seven factors could be relatively large. Although each factor loads on 18 or fewer variables for GPCA with $\kappa = 1$, among all factors there are 61 series with positive loadings; no series is repeated in two factors. Thus, in principle, FAVAR could still be incorporating information on far more variables than typically incorporated in a standard VAR.

Despite the changes in factor loadings, however, estimated impulse responses are largely unaffected. Figure (2) compares the responses generated by the standard procedure to those generated by the GPCA factors (along with the small-scale VAR described in section (5.3).) Dashed-dot lines giving the 90% confidence interval for the PCA-based FAVAR.[2] The responses are strikingly close to those generated by the original procedure. In particular, both the PCE-P and PUNEW responses are essentially unchanged. For other variables (e.g. unemployment) responses deviate somewhat more, however, the qualitative implications do not change for any of the variables I have studied. The impulse responses generated using the GPCA-based factors are difficult to distinguish from those generated using PCA factors.

5.2 Interpreting Factors

Can the GPCA factors be interpreted? Figure (3) shows PCA loadings for each factor and group, in terms of the fraction of the $\mathcal{L}1$-norm of total loadings for each factor. As usual, PCA loads on all groups, making it difficult to derive a correspondence with any particular economic concept(s). Figure (4) shows the much more sparse group loadings for GPCA. Factor one is clearly related to measures of real output, factor two to prices, factor three to unemployment (with some dependence on interest rate spreads), factor four to interest rates, factor five to stock prices, and factor seven to measures of aggregate consumption. Factor six remains somewhat ambiguous, with loadings on variables in the housing, prices and employment groups.

The first row of table (4) displays the variable with the highest correlation to each factor, along with the correlation coefficient. The variables corresponding to each factor, in order, are manufacturing industrial production (IPS43), the PCE price index (PCE-P), unemployment 15+ weeks (LHU15), the federal funds effective rate (FYFF), the Standard and Poor's stock index (FSPCOM), total aggregate payrolls (CES0500), and personal consumption expenditure excluding food (PCC). For all but factor six, the maximum correlation is greater than .9. For

[2]This bootstrap takes into account the fact that factors themselves are generated regressors and accounts for the bias in OLS estimation, following Kilian (1998).

reference, the average maximum correlation for standard PCA factors is .60. Not surprisingly, the most correlated variable in each case comes from the most heavily loaded variable grouping. This table suggests that most of the information in the GPCA factors is available in just a few prominent series. As noted above, however, factor six requires special attention.

These particular interpretations of the factors are consistent with previous work in both the VAR and factor model literature. Five of the factors, namely production, prices, consumption, the federal funds interest rate, and unemployment (although not always this measure), have long been included in VARs designed to identify monetary policy shocks. Today, stock prices (alternatively, a commodity price) are also a commonplace addition to empirical specifications. One interpretation of factor six is that it consists of leading indicators. In addition to private payroll employment, the factor is highly correlated with the purchasing managers index (PMI), a composite index of several survey indicators collected from purchasing managers by the *Institute for Supply Management* and housing starts (HSFR), from the inventories/orders and housing groups, respectively. In their forecasting exercises, Bai and Ng (2008) consistently find PMI, HSFR, and (to a lesser extent) CES0500 to have strong predictive power for inflation at horizons between one month and two years. Armah and Swanson (2008) similarly find HSFR, along with various stock market series, to be good factor proxies in their forecasting exercises. My results could be interpreted as further support for an emerging consensus on the most "information rich" series in US data.

5.3 Small-Scale VARs

Given the close association of each factor with a distinct economic concept, it may be possible to proxy for the unobserved factors with a single observed variable. The correlations in table (4) suggest a seven variable VAR in [IPS43, PCE-P, PCC, LHU14, CES0500, FFR, S&P]. The GPCA loadings, however, show that factor six contains important variation that cannot be captured by any single series. To address this issue, I instead append the most highly correlated variables from *each* of the groups with positive loadings for the factor: PMI, HSFR, and CES0500. Combined, these three variables explain 99% of the variation in the sixth factor, although individually they have R-squares of less than .8.

The impulse responses for the nine-variable specification are shown by the light-blue line in figure (2). Once again, the responses are quite similar to those from the canonical FAVAR. Notably, the price response remains quite similar to the original FAVAR results. The production responses are remarkably similar as well. The GPCA factor interpretations lead directly to a small-scale VAR with properties that are very close to the original PCA-based and to the GPCA-based FAVAR.

5.4 How Important is Grouping?

In principle, the interpretability of factors could be achieved by any sufficiently sparse modification of PCA. A natural question, then, is how important is the addition of grouping to the results? Table (3) shows the number of nonzero loading for the various estimation methods. Without imposing the groups (simple SPCA), the factors nearly always load on more groups *ex post*. This is especially true for low values of κ. Still, the standard SPCA factors are clearly far more sparse than PCA, while spanning essentially the same space.

In practice, I find that my primary result - that impulse responses based on are sparsely-estimated factors closely replicate PCA-based responses - is significantly more robust when the groups are imposed. For small κ, both SPCA and GPCA closely match the standard FAVAR results, although factor interpretation for SPCA is somewhat more difficult. As κ increases, the SPCA factors begin to correspond to very few groups as well, but the performance vis-a-vis GPCA in reproducing the original FAVAR results deteriorates. Without grouping, the main points of this paper still stand, although the results are less stark. I take this as evidence that the role of variable grouping in factor estimation warrants further attention at both theoretical and applied levels.

5.5 Identification with Interpretable Factors

So far, I have identified shocks to monetary policy using the MA representation of the observable series, rather than placing restrictions on the dynamics of the factors themselves. With an interpretation for the factors, however, it is possible to consider restrictions on the dynamics of the factors themselves. Boivin, Giannoni, and Mihov (2009) impose that the the federal funds rate is among the factors, and use a choleski decomposition ordering that variable last. Yet, if the factor interpretations suggested here are correct (recall from table (1) that the sparse factors span essentially the same space as the PCA-based factors), the standard recursive identification assumption on the shocks is not likely to hold. In particular, stock prices and the leading indicators may respond contemporaneously to interest rate shocks. Unlike in the standard PCA case, we can easily reorder the GPCA factors so that the factors linked with the stock and leading indicator concepts are ordered after the interest rate. We can consider other identification schemes on the factors themselves, as well, such as those using restrictions on long-run impact of shocks or error-variance decompositions. Furthermore, because the economic concepts behind the factors are clear, it is economically interesting to examine the factor impulse responses directly.

To demonstrate the implications of the identification choice, I consider three alternative identification assumptions. Identification assumption one (ID 1) corresponds to the standard (and I argue implausible) assumption that the federal funds rate is order last in the FAVAR. Identification assumption two (ID 2) also consists of a choleski ordering, but places federal funds rate before the set of leading indications and the stock market. Finally, identification assumption 3 (ID 3) relaxes the strict ordering assumption, and permits the interest rate to respond contemporaneously to shocks to the leading indicator factor, and the leading indicator factor to respond contemporaneously to stock market shocks. In order to achieve identification for ID 3, I impose two long-run restrictions, namely that monetary policy cannot affect long-run output, and stock price shocks cannot affect long-run unemployment.

Figure (5) shows the factor impulse responses to a unit monetary policy shock, under the three alternative identification assumptions. For these figures, I have estimated the factors with GPCA with $\kappa = 1$, with the additional restriction that the fourth factor is exactly equal the federal funds rate. This is a relatively modest imposition, given the close correlation between factor four and the federal funds rate documented above. The appendix shows a simple modification of the GPCA procedure which incorporates this restriction.[3] Since the factors represent linear combinations of variables with potentially different natural scales, I give the responses in terms of standard deviation units.

The figure shows that the qualitative responses are the same for all three identification schemes. However, ID 2 shows a stronger price response and slightly larger output response in the initial periods. ID 3, in turn, shows the most speedy price response, and a somewhat shorter period of decreased output. Figure (6) demonstrates a similar pattern for the observable variable impulse responses, based on the same three identification schemes. Overall, this exercise suggests that the features of FAVAR responses are robust to a range of identification schemes. Identification assumptions that deviate from the standard choleski ordering approach, however, can yield very different impulse responses in some cases.

5.6 Robustness

I perform the same exercises on the two subsamples described in the data section. Figure (7) reproduces impulse responses for the same series over the early sample period, where the baseline is now a standard FAVAR estimated on the early subsample. Once again, impulse responses are theoretically reasonable and quite similar to those estimated by standard FAVAR. The interpretation of the factors, including the leading indicator factor which loads on employment, housing, and inventories, are unchanged. In fact, table (4) shows that factor proxies are nearly identical to the full sample, with the exception of PMI replacing CES0500 as the best proxy for the leading indication factor and the unemployment rate (LHUR) rather than the long-term unemployment

[3] This approach corresponds closely to the iterative PCA approach of Boivin, Giannoni, and Mihov (2009). In fact, when $\lambda_{1,k} = 0$ the total variance captured by iterative PCA remains slightly greater than in the GPCA method used here. A simple rotation of the Boivin, Giannoni, and Mihov (2009) factors show that the restrictions implied by the iterative PCA method are a strict subset of the restrictions imposed by the unpenalized GPCA estimator.

rate (LHUR15) corresponding most closely to the unemployment factor. Figure (7) shows that the identical specification, [IPS43, PCE-P, PCC, LHU14, CES0500, FFR, PMI, HSFR, S&P], for the small-scale VAR remains quite close to the standard FAVAR results.

For the latter subsample, factor interpretation is remarkably robust, although the ordering of factors shifts significantly. Table (4) shows that the price factor now corresponds most closely to CPI-excluding shelter (PUXHS). The unemployment rate again replaces the the long-term unemployment level, but all other proxy variables match the full sample. In contrast to the early sample, however, impulse responses vary substantially across specifications.[4] The error bands demonstrate, however, that the point estimates come with substantially more uncertainty as well. These results suggest that the dynamics of the full sample period are being driven largely by the experience in the early part of the sample.

The results for both datasets are robust to a wide range of values for the sparsity parameter, κ. My experimentation suggests that reasonable impulses *and* interpretability are maintained for κ approximately between .1 and 2. Smaller values for κ give impulse responses close to FAVAR, but with imprecise interpretation of the factors, while larger values for κ cause the responses to diverge greatly from FAVAR.

Finally, I examine the results for FAVARs with four through eight factors. I find FAVAR impulse responses vary somewhat with different numbers of factors (see the figures in BBE, for example), but the impact of using GPCA-estimated factors is uniformly small. Furthermore, while the groups represented in each factor tend to bunch with fewer factors allowed in estimation (for example, the consumption and stock factors merge when only six factors are allowed in estimation) the "leading indicator" factor, which loads on employment, housing, and inventories, is present regardless of the number factors selected.

[4]Impulse response point estimates are also remarkably sensitive to the identification scheme used.

6 Factor Rotations

As noted earlier, standard principal components can estimate only a rotation of the true factors. Recently, Boivin, Giannoni, and Stevanović (2010) and Bai and Ng (2010) have suggested using theoretical restrictions in order to recover an estimate of the true (unrotated) factors and facilitate factor interpretation. Here I compare my results to the interpretations implied by PCA rotation methods.

In order to recover economically meaningful factors, Bai and Ng (2010) consider two alternatives to the standard PCA restrictions. Let Λ_K be the first K rows of the matrix of factor loadings in equation (1). The first set of alternative restrictions, which they call PC2, imposes that $\hat{F}'\hat{F} = I_K$ and $\hat{\Lambda}_K$ is lower triangular. The rotated estimates of the factors is then $\hat{F}^{PC2} = \hat{F}^{PCA}(\hat{H}^{PC2})'$, where $\hat{H}^{PC2} = \hat{Q}$ comes from the LQ-decomposition of $\hat{\Lambda}_K^{PCA}$. Implementing this rotation requires the econometrician to select a subset of K observable variables on which to make a recursive assumption, namely, that there is a variable (chosen to be ordered first) which depends only on the first factor, another (chosen to be ordered second) which depends only on the first two factors, and so on.

A second alternative, which Bai and Ng call PC3, leaves $\hat{F}'\hat{F}$ unrestricted, and instead imposes that $\hat{\Lambda}_K = I_K$. In this case, $\hat{H}^{PC3} = \Lambda_K^{PCA}$. This scheme amounts to assuming that there exist series which are "pure" indicators for each of the true factors. This assumption certainly seems strong. However, the results using GPCA so far suggest that this is, in fact, not terribly far from the truth.

Finally, Boivin, Giannoni, and Stevanović (2010) use the structural identification assumptions from the full FAVAR specification in equations (1) and (2) to pin down the rotation matrix. In particular, they restrict the impact response to the structural shocks of the first K variables, given by the top $K \times K$ block of the matrix $\mathbf{B}(0)$, to be lower-triangular exactly I do to identify monetary policy shocks above. Simple algebra shows that this restriction is identical to recursive assumption on Λ in PC2. Rather than impose orthogonal factors, however, this approach imposes the restriction that $\hat{H}^{BGS}(\hat{H}^{BGS})' = \Sigma_e$, where Σ_e is the covariance matrix of the reduced form residuals of the VAR estimated on the PCA factors. Crucially, their approach implies the structural assumption that the matrix A^{-1} is the identity.[5]

Table (5) shows the variable of maximum correlation for each factor, for GPCA factors and each of the three alternative PCA rotation schemes. The rotation techniques require an *a priori* choice on the ordering of the first K variables. Obviously many specifications are possible, even given a rich set of theoretical restrictions. In order to give the rotation procedures the best chance at generating interpretable factors, I use the results from GPCA and set the first seven variables to [IPS43, PCE-P, PCC, LHU14, CES0500, FFR, S&P]. With a few exceptions, the results corroborate the interpretations established above. However, the proxy correlations for the rotation techniques are somewhat lower on average.

I interpret these results as suggesting that rotation methods are not well suited, alone, to establish strong links between estimated factors and observable series. The set of possible restrictions is immense, and the resulting interpretations are less clear cut than with the GPCA factors. In general, the fact that GPCA factors span less of the data may be of some concern. In the case of US data, however, I have shown that the loss of explanatory power is negligible, even for very sparsely estimated factors.

7 Conclusions

I draw two main conclusions from the exercises above. First, sparsely estimated factors can provide economic meaning to the FAVAR system given in equation (2). GPCA furthers that interpretation by incorporating natural variable groups in the factor-estimation step. In US data, the GPCA-based factors correspond quite well to production, price, unemployment, consumption, interest rate, and stock prices. The final factor consists of a set of well-known leading indicators.

Second, the informational benefits of incorporating estimated factors in VAR estimation is

[5]To see this, consider estimating equation (2) using factors rotated by a matrix H^*. Then, the matrix recovering the structural shocks, $\hat{H}^{BGS}$, is an estimate of AH^*, which corresponds to the true rotation of the factors only when A is the identity.

moderate. Nearly identical responses can be achieved even when factors are quite sparsely estimated. To the extent that FAVAR does incorporate new information with respect to traditional VARs, this addition depends on a relatively small subset of the data series used in the classic FAVAR implementation. In particular, in US data, it appears that the addition of the purchasing managers' index, housing starts, and private payroll employment to an otherwise standard VAR specification closely and robustly replicates the canonical FAVAR results.

References

Armah, N. A. C., and N. R. Swanson (2008). "Seeing Inside the Black Box: Using Diffusion Index Methodology to Construct Factor Proxies in Largescale Macroeconomic Time Series Environments," *SSRN eLibrary.*

Bai, J., and S. Ng (2008). "Forecasting Economic Time Series Using Targeted Predictors," *Journal of Econometrics*, 146(2), 304 – 317, Honoring the research contributions of Charles R. Nelson.

——— (2010). "Principal Components Estimation and Identification of the Factors," Manuscript.

Belviso, F., and F. Milani (2006). "Structural Factor-Augmented VARs (SFAVARs) and the Effects of Monetary Policy.," *B.E. Journal of Macroeconomics: Topics in Macroeconomics*, 6(3), 1 – 44.

Bernanke, B. S., and J. Boivin (2003). "Monetary Policy in a Data Rich Environmet," *Journal of Monetary Economics*, 50(3), 525–546.

Bernanke, B. S., J. Boivin, and P. Eliasz (2005). "Measuring the Effects of Monetary Policy: a Factor-Augmented Vector Autoregressive (FAVAR) Approach," *Quarterly Journal of Economics*, 120(1), 387–422.

Boivin, J., M. P. Giannoni, and I. Mihov (2009). "Sticky Prices and Monetary Policy: Evidence From Disaggregated U.S. Data," *American Economic Review*, 99(1), 350–84.

Boivin, J., M. P. Giannoni, and B. Mojon (2008). "How Has the Euro Changed the Monetary Transmission?," NBER Working Papers 14190, National Bureau of Economic Research, Inc.

Boivin, J., M. P. Giannoni, and D. Stevanović (2010). "Dynamic Effects of a Credit Shock in a Data-Rich Environment," Manuscript.

Boivin, J., and S. Ng (2006). "Are More Data Always Better For Factor Analysis?," *Journal of Econometrics*, 132(1), 169 – 194.

Chamberlain, G., and M. Rothschild (1983). "Arbitrage, Factor Structure, and Mean-Variance Analysis On Large Asset Markets," *Econometrica*, 51(5), 1281–1304.

Kilian, L. (1998). "Small-Sample Confidence Intervals For Impulse Response Functions," *The Review of Economics and Statistics*, 80(2), 218–230.

Leng, C., and H. Wang (2009). "On General Adaptive Sparse Principal Component Analysis," *Journal of Computational and Graphical Statistics*, 18, 201–215.

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 and Economic Statistics*, 20(2), 147–162.

——— (2005a). "An Empirical Comparison of Methods For Forecasting With Many Predictors," *manuscript.*

——— (2005b). "Implications of Dynamic Factor Models For VAR Analysis," *manuscript.*

Zou, H., and T. Hastie (2005). "Regularization and Variable Selection Via the Elastic Net," *Journal of the Royal Statistical Society*, 67(2), 301–320.

Zou, H., T. Hastie, and R. Tibshirani (2006). "Sparse Principal Components Analysis," *Journal of Computational and Graphical Statistics*, 15(2), 265–286.

Table 1: Percent explained variance (PEV) of estimated factors.

		$\kappa = 0.1$		$\kappa = 1$	
	PCA	**SPCA**	**GPCA**	**SPCA**	**GPCA**
Factor 1	0.18	0.13	0.13	0.07	0.13
Factor 2	0.32	0.25	0.24	0.13	0.25
Factor 3	0.39	0.33	0.32	0.21	0.33
Factor 4	0.44	0.40	0.40	0.33	0.40
Factor 5	0.48	0.46	0.46	0.37	0.43
Factor 6	0.51	0.50	0.50	0.47	0.49
Factor 7	0.55	0.54	0.53	0.50	0.53

Table 2: Number of series with non-zero loadings in $\hat{\beta}$.

		$\kappa = 0.1$		$\kappa = 1$	
	PCA	**SPCA**	**GPCA**	**SPCA**	**GPCA**
Factor 1	139	5	11	1	7
Factor 2	139	10	7	1	10
Factor 3	139	12	19	5	11
Factor 4	139	21	32	11	4
Factor 5	139	30	30	5	5
Factor 6	139	49	5	18	18
Factor 7	139	14	13	4	6

Table 3: Number of groups with non-zero loadings in $\hat{\beta}$.

		$\kappa = 0.1$		$\kappa = 1$	
	PCA	**SPCA**	**GPCA**	**SPCA**	**GPCA**
Factor 1	14	1	1	1	1
Factor 2	14	4	2	1	1
Factor 3	14	6	4	3	2
Factor 4	14	5	3	2	1
Factor 5	14	9	5	1	1
Factor 6	14	11	1	4	3
Factor 7	14	4	3	1	1

Table 4: Proxies for GPCA Factors (correlation).

	Full Sample	**Early Sub-sample**	**Late Sub-sample**
Factor 1	IPS43 (0.98)	IPS43 (0.98)	LHUR (0.98)
Factor 2	PPCE (0.96)	PPCE (0.97)	PUXHS (0.97)
Factor 3	LHU15 (0.93)	LHUR (0.88)	FYFF (0.99)
Factor 4	FYFF (0.99)	PMI (0.86)	CES0500 (0.92)
Factor 5	FSPCOM (0.98)	FYFF (0.99)	IPS43 (0.98)
Factor 6	CES0500 (0.88)	PCC (0.96)	PCC (0.95)
Factor 7	PCC (0.97)	FSPCOM (0.99)	FSPCOM (0.98)

Table 5: Proxies for GPCA and rotated PCA factors (correlation).

	Grouped PCA	**Bai & Ng:PC2**	**Bai & Ng:PC3**	**Boivin et al.**
Factor 1	IPS43 (0.98)	IPS43 (0.97)	IPS43 (0.97)	IPS43 (0.97)
Factor 2	PPCE (0.96)	PPCE (0.95)	PPCE (0.95)	PPCE (0.95)
Factor 3	LHU15 (0.93)	PCC (0.82)	PCC (0.86)	PCC (0.79)
Factor 4	FYFF (0.99)	CES4000 (0.64)	CES0500 (0.90)	CES0800 (0.74)
Factor 5	FSPCOM (0.98)	LHU15 (0.88)	LHU26 (0.94)	LHU26 (0.93)
Factor 6	CES0500 (0.88)	FYFF (0.75)	FYFF (0.98)	FYAAAC (0.92)
Factor 7	PCC (0.97)	CES155 (0.65)	PMI (0.88)	FYBAAC (0.89)

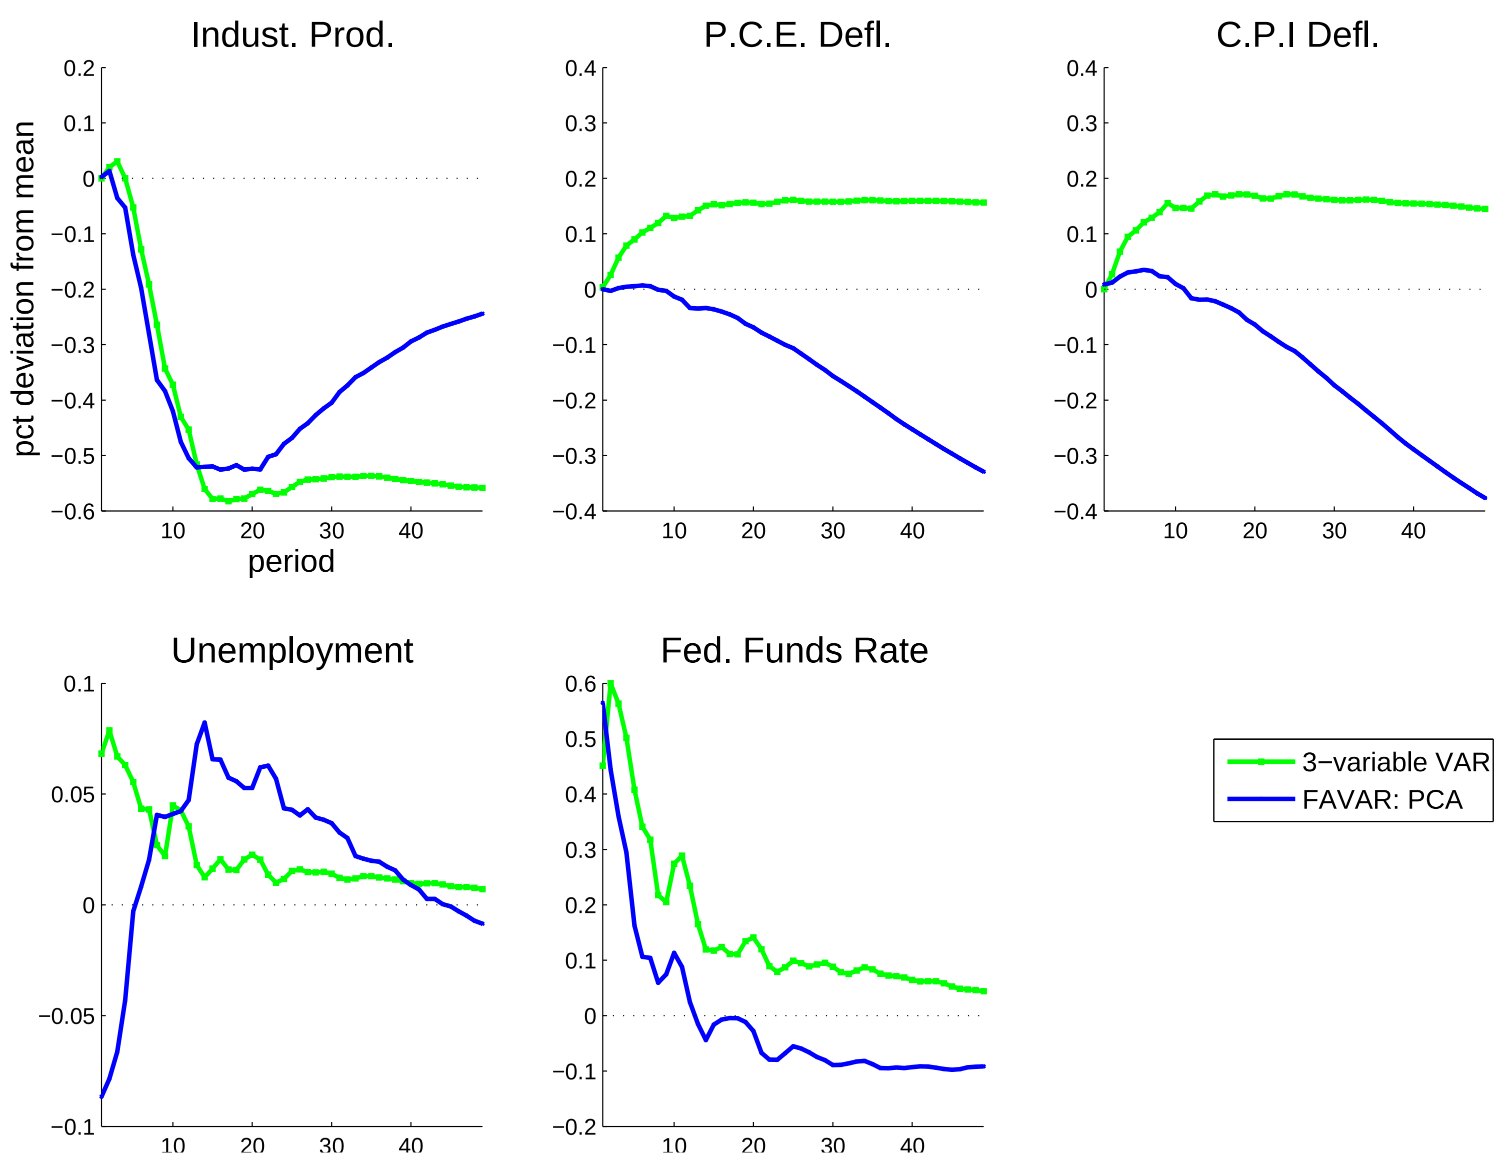

Figure 1: Full Sample - observable variable impulse responses to a single standard deviation shock to monetary policy for PCA-based FAVAR vs VAR in [IP, CPI, FFR]

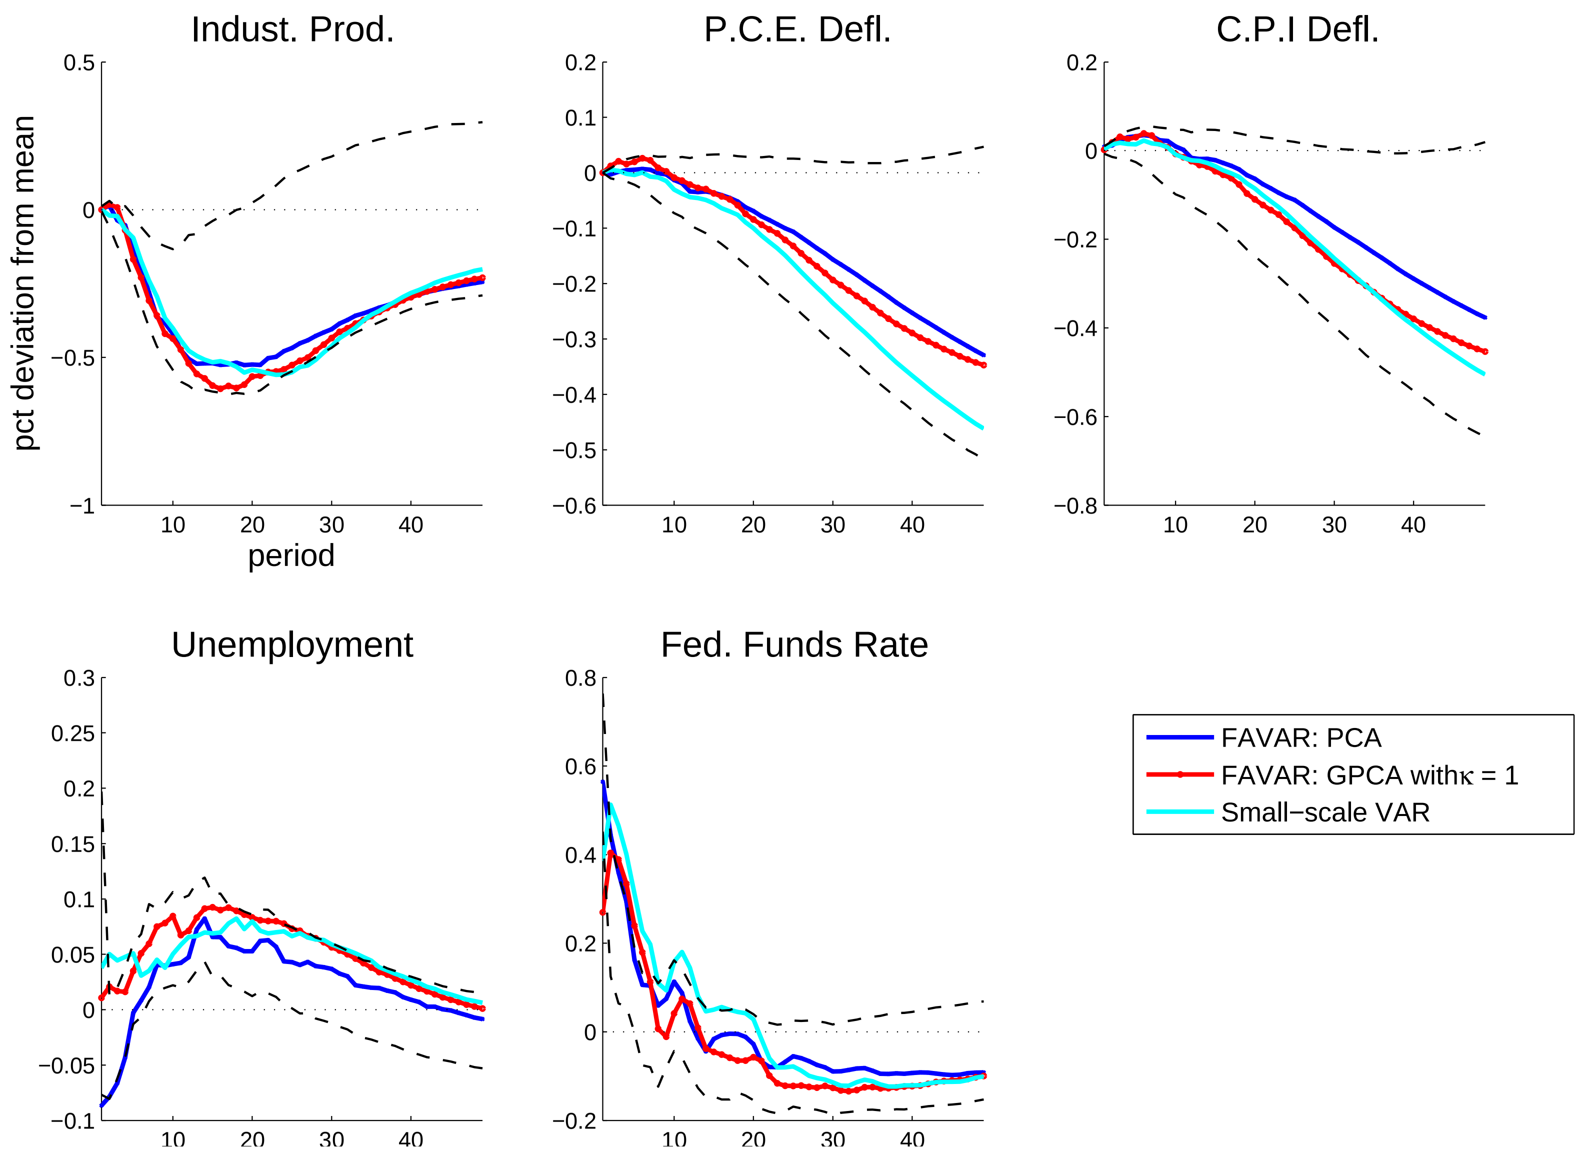

Figure 2: Full Sample - observable variable impulse responses to a single standard deviation shock to monetary policy for PCA-based FAVAR vs GPCA FAVAR and small-scale VAR.

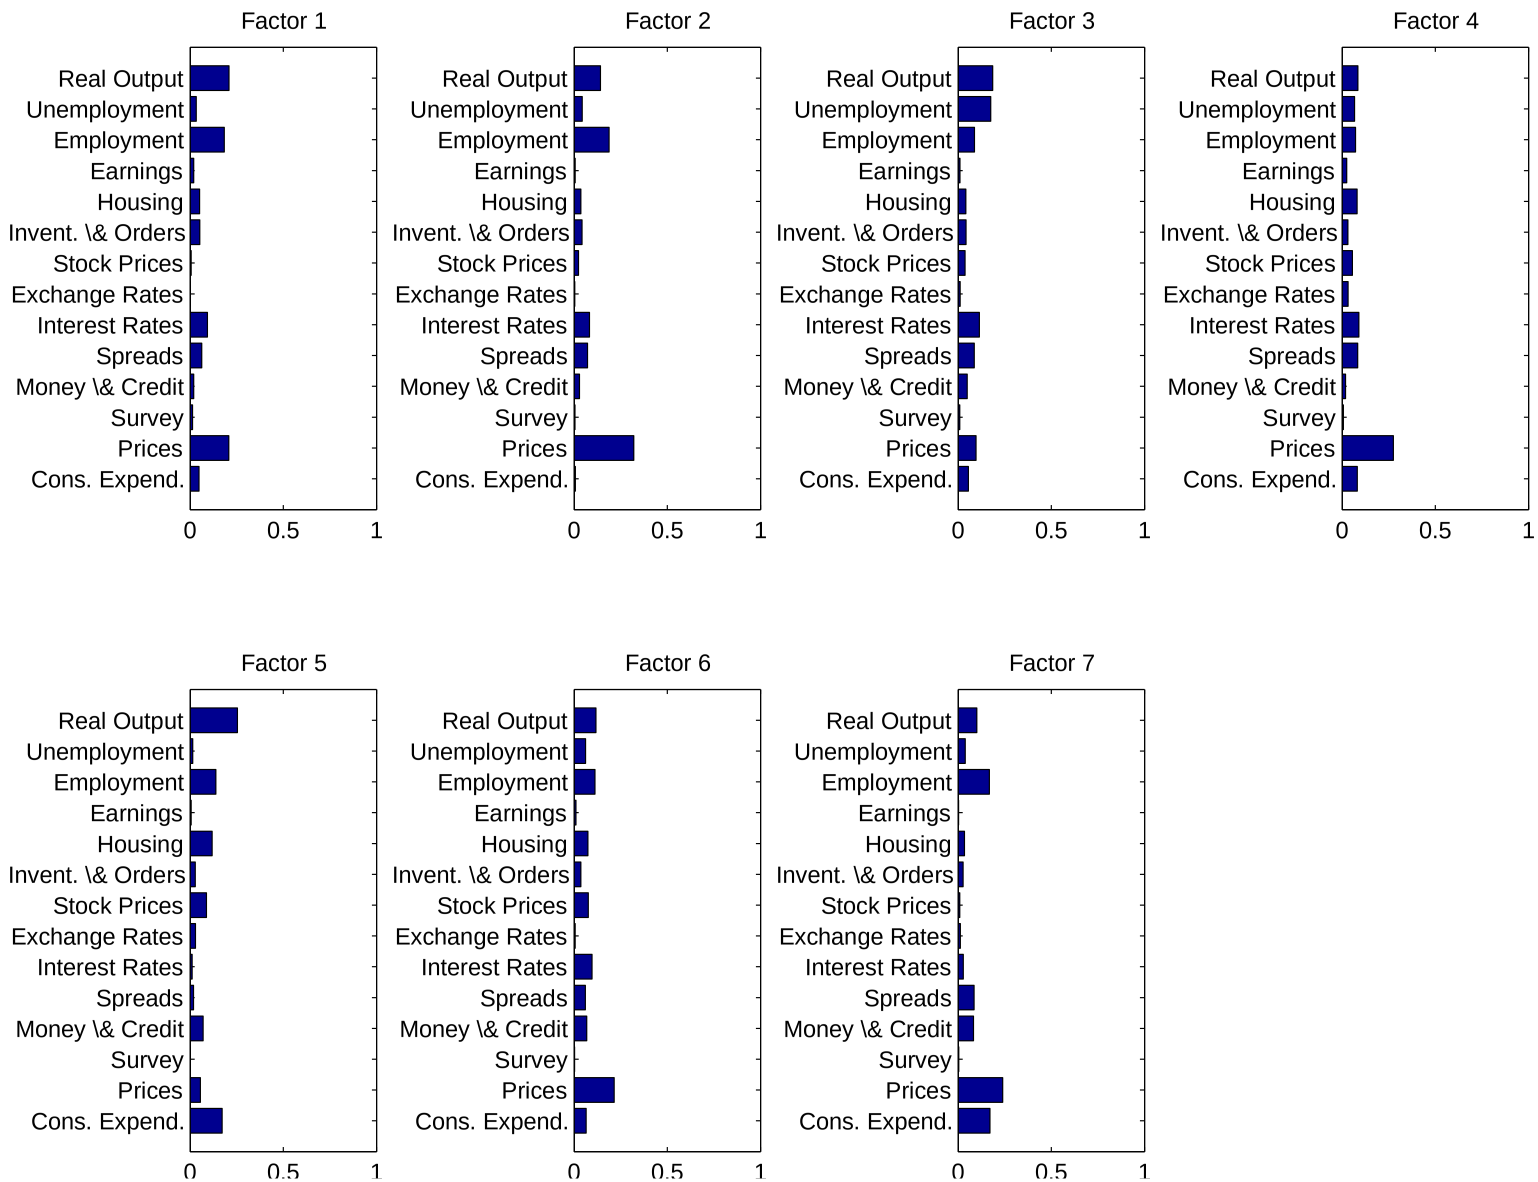

Figure 3: Unconstrained relative loadings by group (fraction of $\mathcal{L}1$-norm)

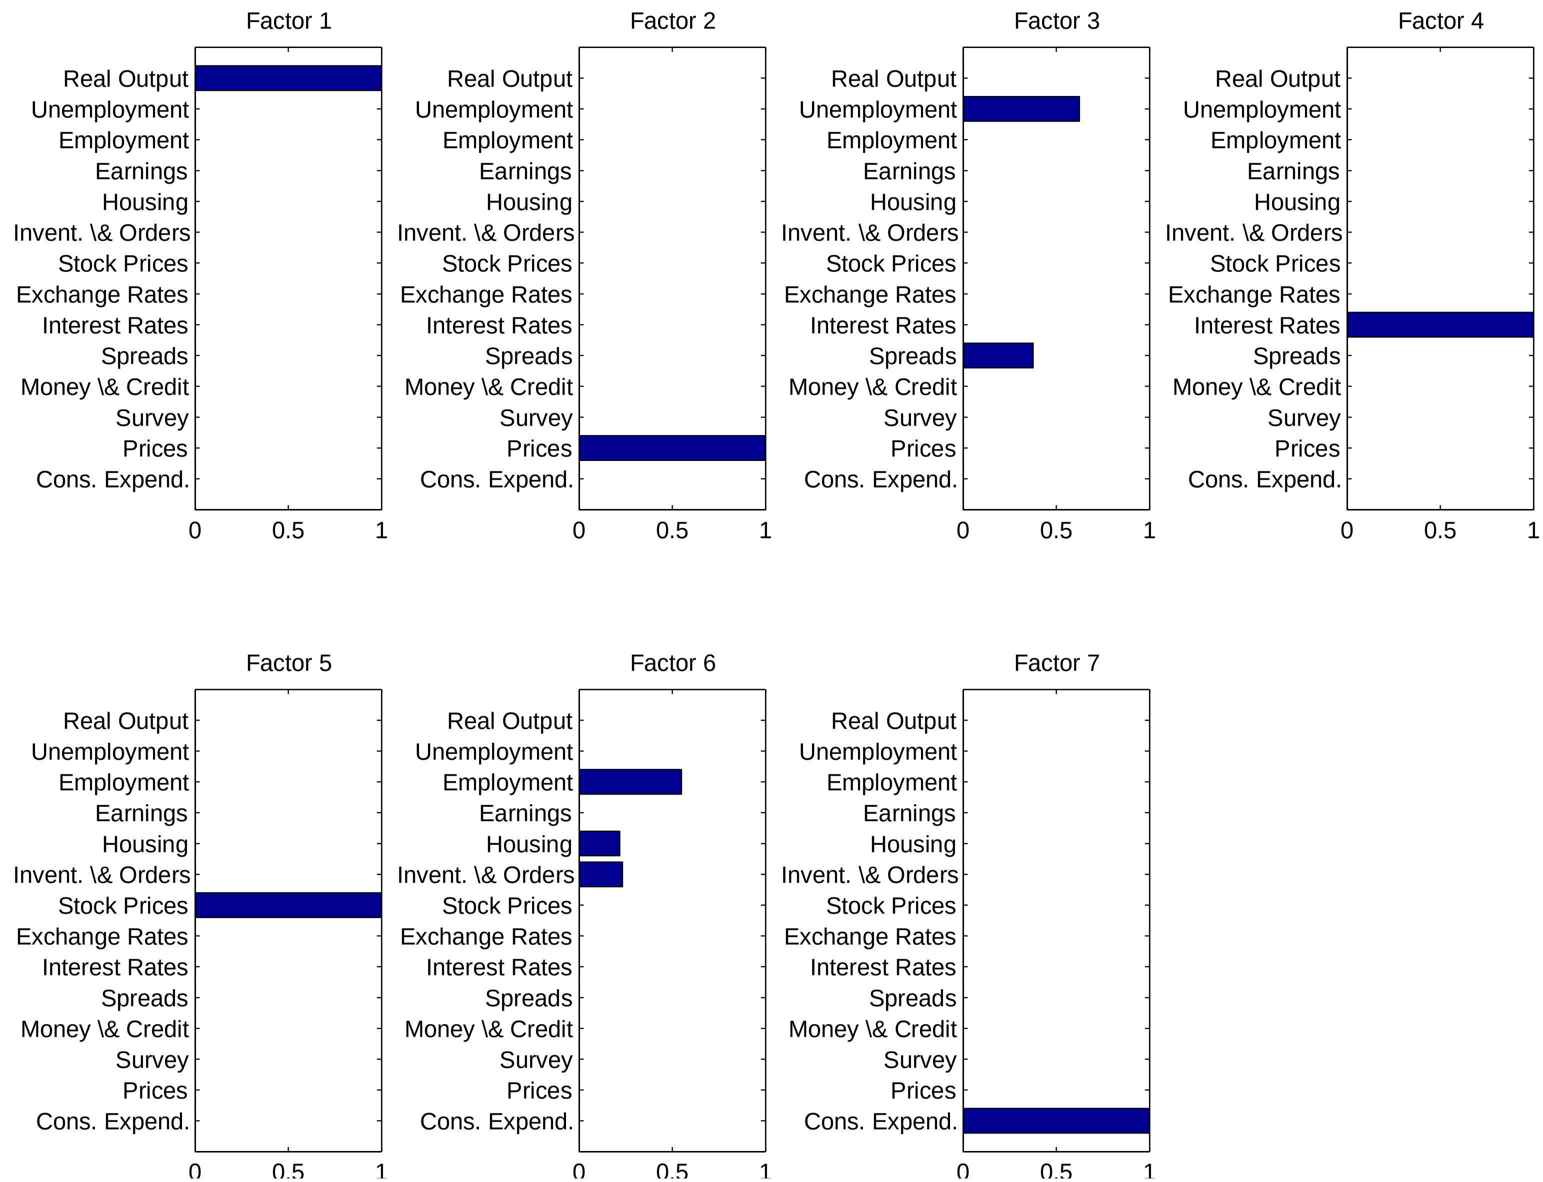

Figure 4: GPCA-constrained relative loadings by group (fraction of $\mathcal{L}1$-norm norm, $\kappa = 1$)

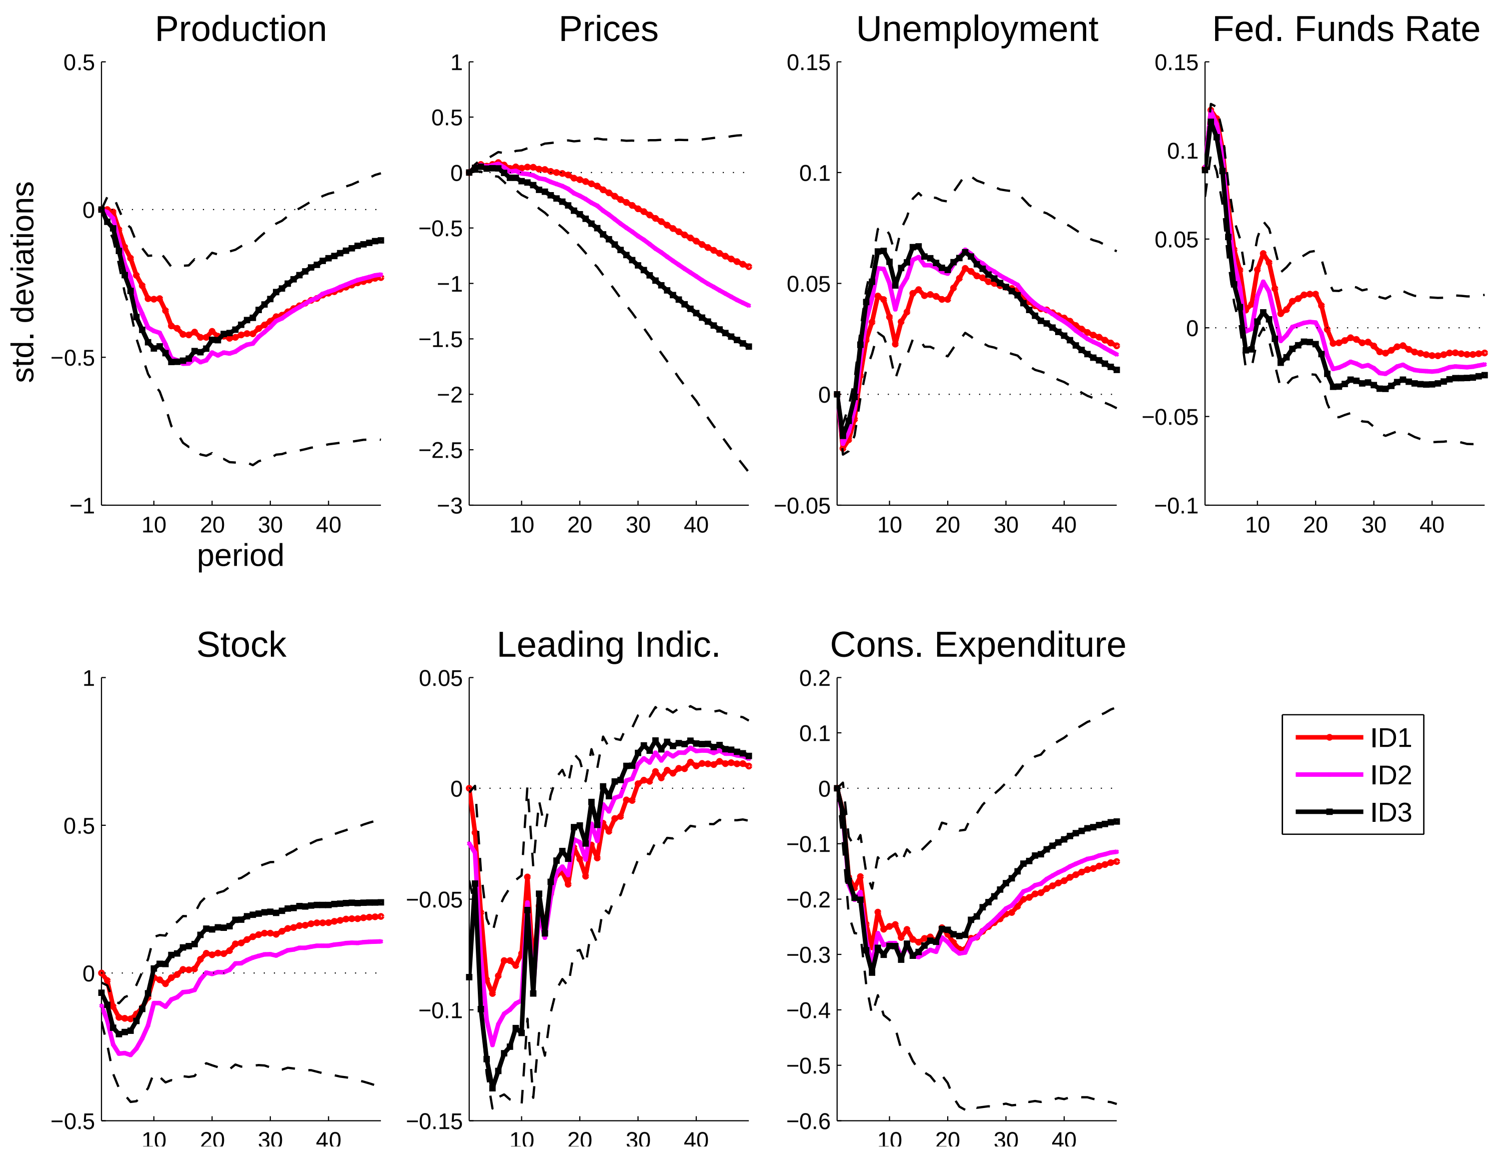

Figure 5: Full Sample - Factor impulse responses to a single standard deviation shock to monetary policy, under alternative identification schemes when monetary policy is assumed to be fully observable. Responses in standard deviation units.

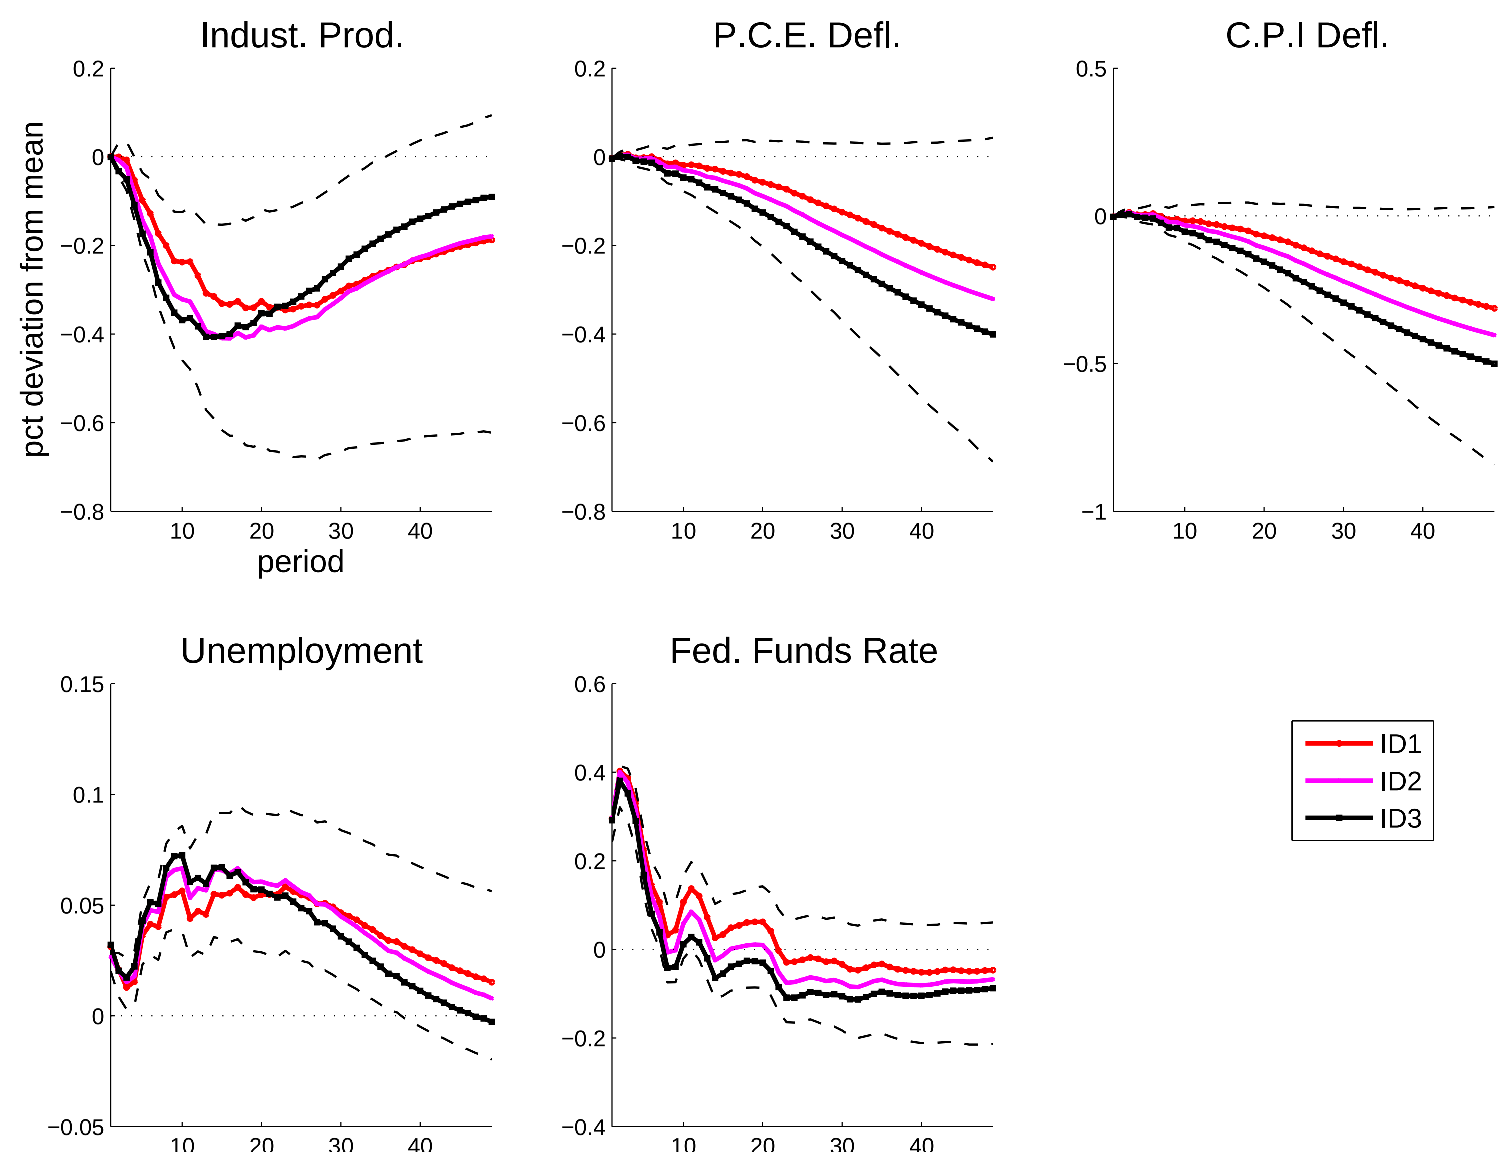

Figure 6: Full Sample - observable variable impulse responses to a single standard deviation shock to monetary policy, under alternative identification schemes when monetary policy is assumed to be fully observable.

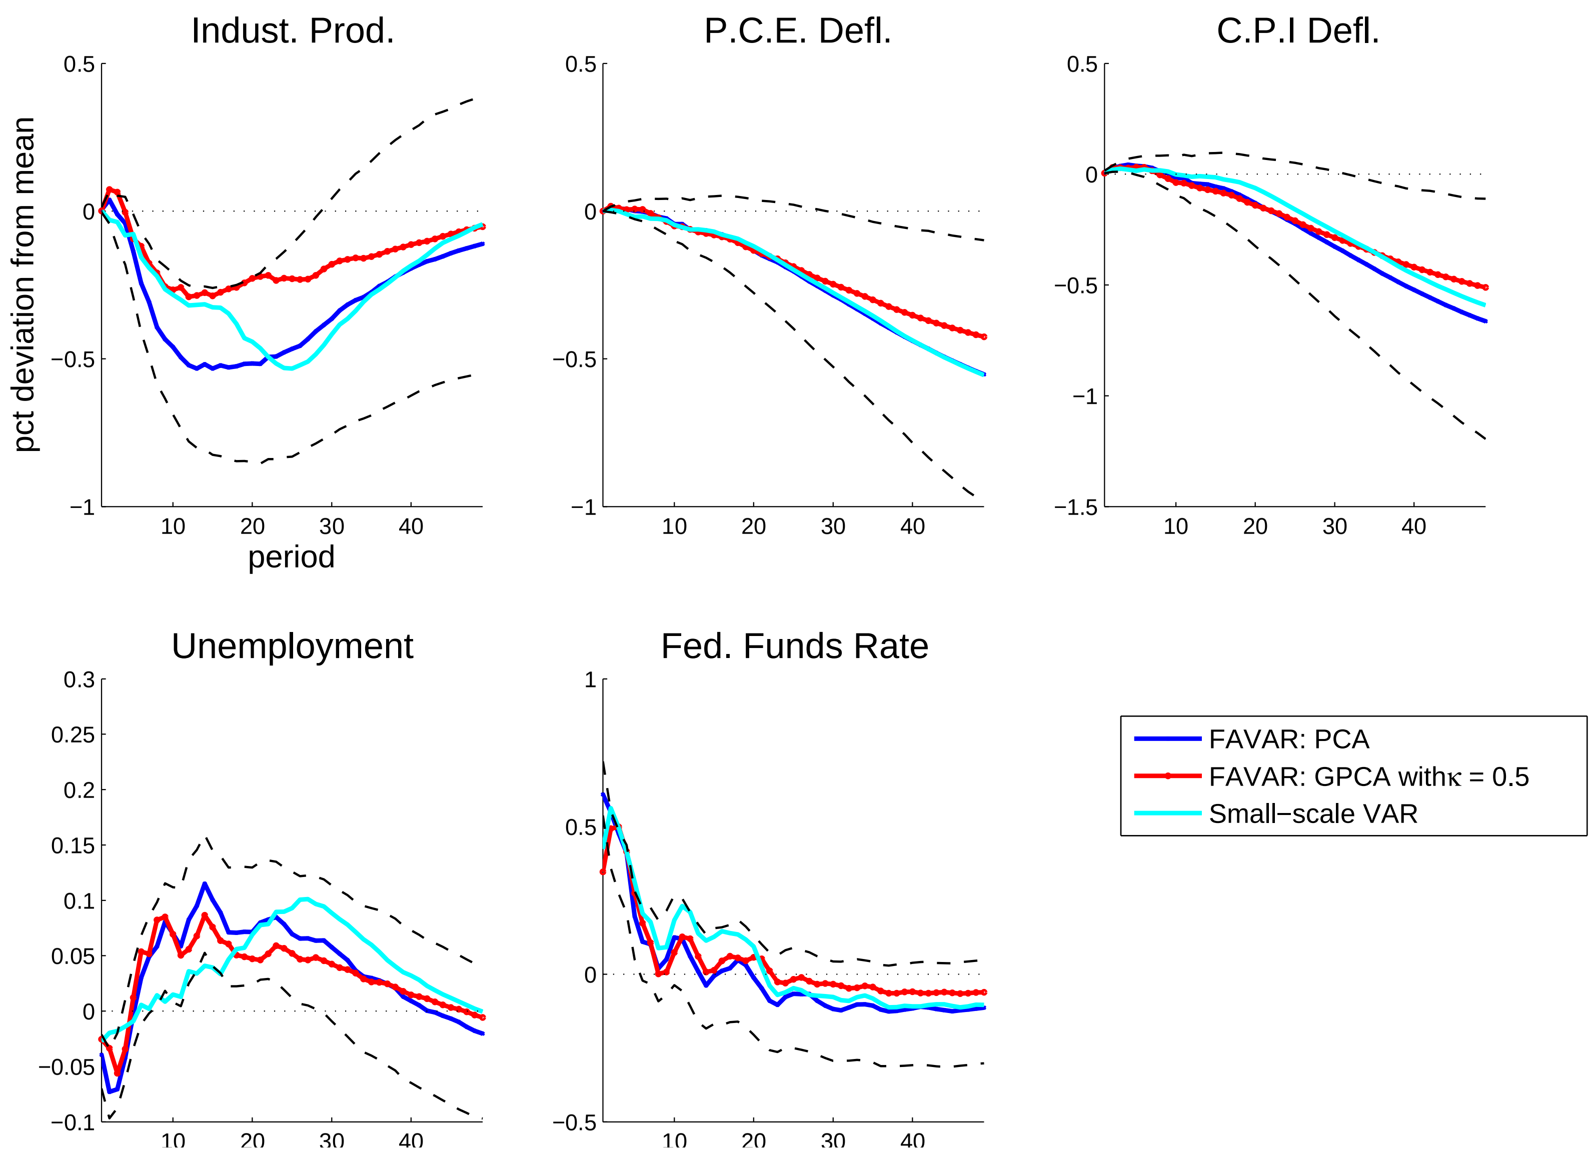

Figure 7: Early Subsample - observable variable impulse responses to a single standard deviation shock to monetary policy for PCA-based FAVAR vs GPCA FAVAR and small-scale VAR.

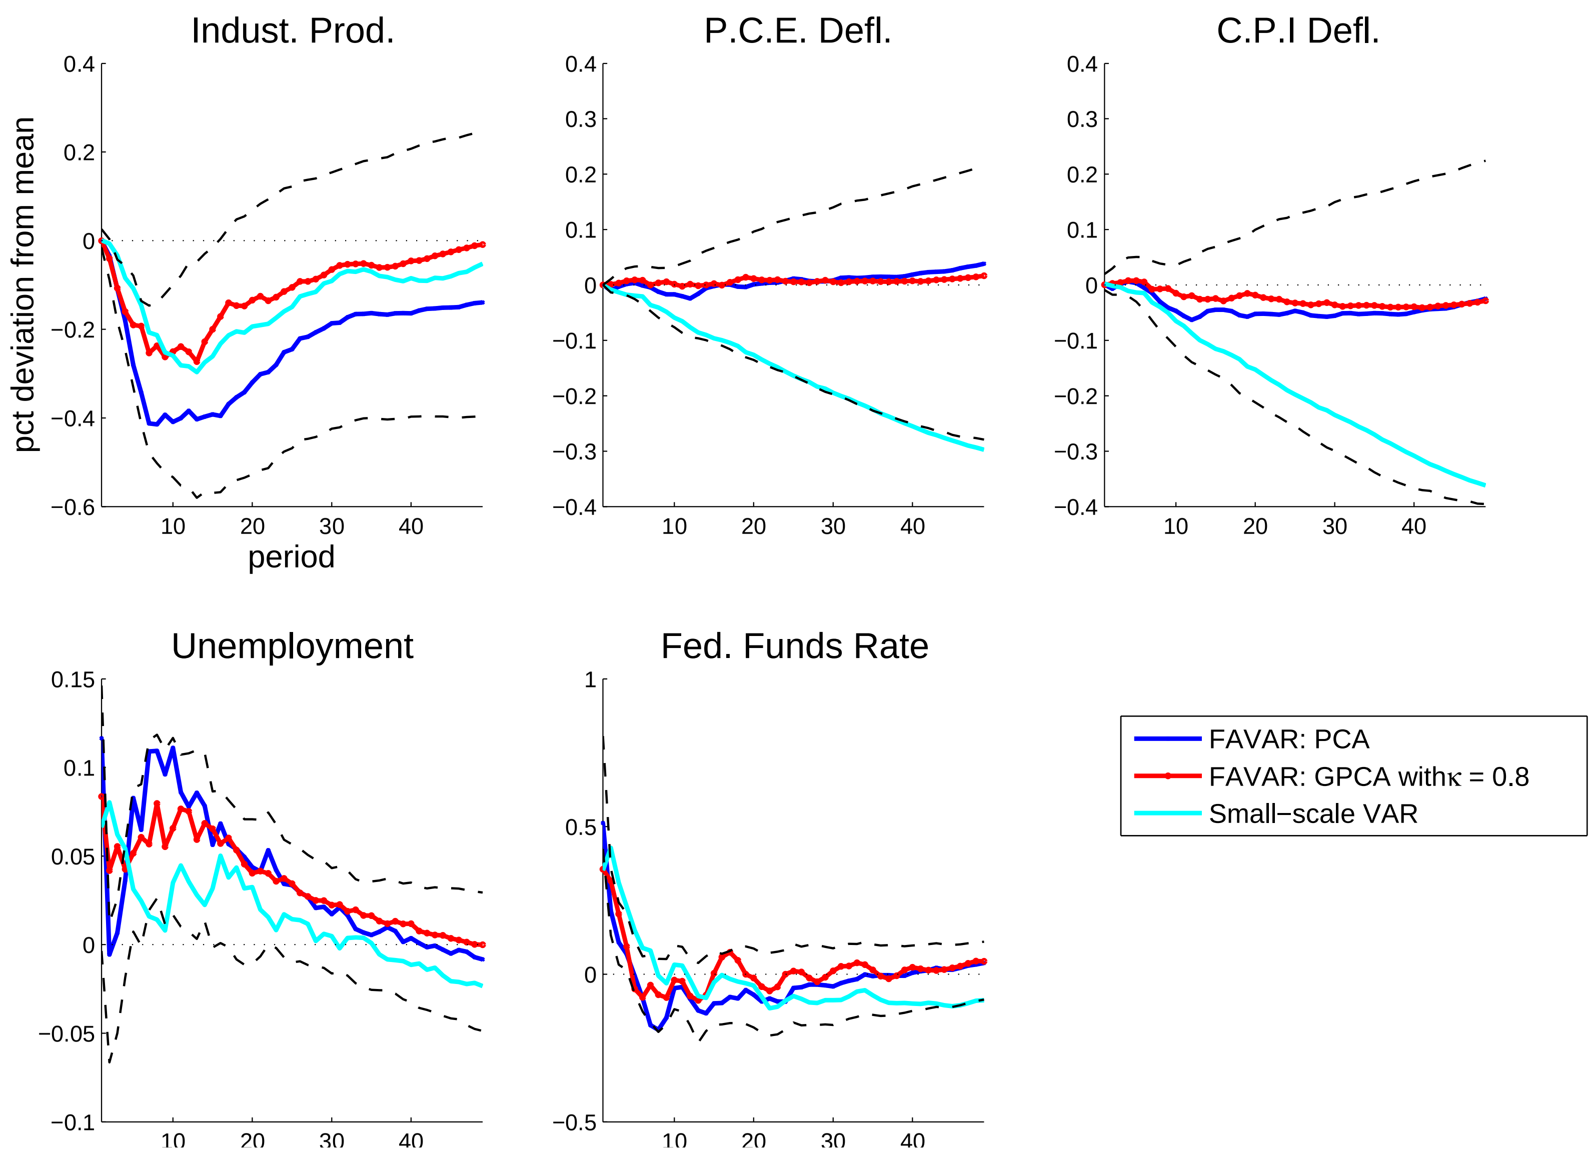

Figure 8: Late Subsample - observable variable impulse responses to a single standard deviation shock to monetary policy for PCA-based FAVAR vs GPCA FAVAR and small-scale VAR.

A Appendix

A.1 Computing GPCA factors

Estimation of the GPCA loadings is a rather simple modification of the algorithm used by Zou, Hastie, and Tibshirani (2006) to solve for their SPCA loadings. An implementation of the Elastic-EN algorithm for solving elastic net problems is a pre-requisite for this algorithm. As for SPCA, estimating loadings for K factors is not equivalent to estimating $K+1$ factors and retaining only the first K of those factors.

GPCA Algorithm:

1. Solve for the standard PCA loadings and generate the weights as given above.
2. Initialize matrix $A^0 \equiv [\alpha_1^0, ..., \alpha_K^0]$ as the loadings of the first K PCA factors.
3. Fixing A^0, solve

$$\beta_k^0 = \underset{\beta_k}{\operatorname{argmin}} \left|\left|X\alpha_k^0 - X\beta_k\right|\right|^2 + \lambda_2 \sum_{g=1}^{G} \left|\left|\frac{\beta_{k,\vec{g}}}{\hat{w}_{k,g}}\right|\right|^2 + \lambda_{1,k} \sum_{g=1}^{G} \left|\left|\frac{\beta_{k,\vec{g}}}{\hat{w}_{k,g}}\right|\right|_1 \tag{A.1}$$

4. Given $\beta^0 = [\beta_1^0, ..., \beta_K^0]$, compute $A^1 = UV'$ where (U, V') are from the SVD of $X'X\beta = UDV'$.
5. (Optional) Update group weights using A^1.
6. Repeat steps 3-5 until convergence.

A few comments are in order. Step 3 requires solving an "adaptive" elastic net problem. Since every elastic net problem can be written as a lasso problem (Zou and Hastie, 2005), this is just an application of the adaptive lasso. Zou, Hastie, and Tibshirani (2006) note that only the covariance matrix (or correlation matrix, for standardized data) is needed to solve the SPCA problem. This modification applies here as well. In simulations I find that step 5, the optional updating of weights, can greatly increase the group-wise sparsity of $\hat{\beta}$. Since PCA recovers asymptotically only the span of the factors (and not the individual factors themselves), the first-stage weights may be large for many groups, even if the series in each group load on only one factor. The results in this paper always include this step, although they are only slightly affected by its inclusion.

As a final comment, note that it is possible to fix particular columns of β and proceed with the above procedure. In this manner, one can impose (for example) that a particular factor corresponds to a particular observed series, which is how the estimator is implemented in section (5.5).

A.2 Tuning

In order to implement the above algorithm, I must select a single value for λ_2 as well as a vector λ_1 of length K, the number of factors to be estimated. The choice of λ_2 affects the numerical performance of the algorithm, but this is its only role. The choice of lasso penalties $\lambda_{1,k}$, however, governs the tradeoff the between sparsity and capturing maximum variance. This will have an important effect on the performance of the estimated factors. As noted by Zou and Hastie (2005), these tuning parameters are not the only possible ones: we could tune by λ_2 and the $\mathcal{L}$-1 norm of the estimated coefficients (t), the fraction of the $\mathcal{L}$1-norm (s) , or the number of non-zero loadings (η).

A conclusive criterion for picking $\lambda_{1,k}$ is difficult to develop, and indeed none is suggested by Zou, Hastie, and Tibshirani (2006). The task is made somewhat easier, however, by the fact that the initial iteration of the the Elastic-EN algorithm provides the entire solution path for all $\lambda_{1,k}$. Zou, Hastie, and Tibshirani (2006) refers to this step as the *direct sparse approximation* (DSPCA) of principal component k, because it is equivalent to a lasso-type regression of the factors on the dataset X. In order to achieve some discipline on the choice of $\lambda_{1,k}$, I adopt an information-type criterion that formalizes the visual procedure used by Zou, Hastie, and Tibshirani (2006).

Suppose I have estimated the $k-1$ factor loadings. For the k^{th} factor, the LARS-EN algorithm generates a the piecewise-linear function of sparse loadings, $\hat{\beta}_k(\lambda_{1,k})$, as a function of $\lambda_{1,k}$, the $\mathcal{L}1$-norm penalty multiplier. Then, let $PEV(\lambda_{1,k})$ measure the percent of total variance explained by the first k factors, $PEVM(\lambda_{i,k})$ measure the marginal PEV contribution of adding the factor $X\hat{\beta}_k(\lambda_{1,k})'$ to set of $k-1$ (already fixed) factors. Along this path, I can also track the number of non-zero loadings in $\hat{\beta}(\lambda_{1,k})$, denoted by $nz(\lambda_{1,k})$.

My criterion is to pick $\lambda^*_{1,k}$ such that

$$\lambda^*_{1,k} = \underset{\lambda_{1,k}}{\text{argmax}} \quad \log\left(\frac{PEVM(\lambda_{1,k})}{PEVM(0)}\right) - \kappa\frac{nz(\lambda_{1,k})}{nz(0)} \tag{A.2}$$

In words, $\lambda^*_{1,k}$ is selected so as to maximize the log of the fraction of the unconstrained variance contribution minus a constant times the proportion of variables with non-zero loadings. The piecewise-linear nature of the problem means this is a discrete optimization over a finite number of values. The functional form in ((A.2)) captures the idea that the information spanned in the first few steps of the estimation is more important than marginal contributions later on. It also precludes choosing a degenerate factor with all zero loadings.

Leng and Wang (2009) suggest an alternative information-type criterion, which can be adapted in a similar way. For this criterion, choose $\lambda^*_{1,k}$ according the criterion

$$\lambda^*_{1,k} = \underset{\lambda_{1,k}}{\text{argmin}} \quad \left[\alpha^0_k - \hat{\beta}_k(\lambda_{1,k})\right]' \hat{\Sigma}_x \left[\alpha^0_k - \hat{\beta}_k(\lambda_{1,k})\right] + \kappa n_z(\lambda_{1,k}) \tag{A.3}$$

where $\hat{\Sigma}_x$ is the sample variance-covariance matrix of the data and $n_z(\lambda_{1,k})$ is the number of non-zero elements in $\hat{\beta}_k(\lambda_{1,k})$.

In both cases, the constant κ represents the econometrician's choice regarding the priority of capturing variance versus maintaining the sparsity of the of the solution. When κ is zero, the solution to the GPCA problem is given by standard PCA; when κ is very large, the estimation will yield factors which load on only one variable.

These two approaches have some important differences. First, unlike the Leng and Wang (2009) version, my criterion makes explicit reference to how well the sparsely-estimated factors span the observed data, which is the key concern in the FAVAR setting. Secondly, because my criterion is normalized by the total PEV contribution of each factor, it tends to generate factors with similar degrees of sparsity for 1 ... K, whereas the Leng and Wang (2009) criterion typically generates a first factor that is not very sparse, a second factor that is more sparse, etc. For these reasons, I use my criterion for the baseline results in this paper, although the two generate similar results in most cases.

A.3 Measuring Percent Explained Variance

SPCA and GPCA-estimated factors are generally not orthogonal, complicating the calculation of the percent of explained variance (PEV), a common summary statistic for PCA. Zou, Hastie, and Tibshirani (2006) offer one metric that corresponds to PCA in the unrestricted case. However, their suggestion depends on the scaling of factors. In this paper, I use the following scale-invariant alternative statistic.

Let $\hat{F} = X\hat{\beta}$ be a set of estimated factors, and let $\hat{P}$ be the matrix projecting the data X onto the space spanned by the factors. Then the percent of explained variance is equal to

$$PEV \equiv 1 - \min_{\Gamma} \frac{\left|\left|X - X\hat{\beta}\Gamma'\right|\right|^2}{||X||^2} = \frac{||\hat{P}X||^2}{||X||^2} = \frac{tr(X'X\hat{\beta}(\hat{\beta}'X'X\hat{\beta})^{-1}\hat{\beta}'X'X)}{tr(X'X)} \tag{A.4}$$

B Data

The data were downloaded from the Global Insight online data service on July 14, 2010, from the BASIC, US CENTRAL, and DRI international databases. Spread variables were computed by the author (AC = author's calculation). Personal consumption expenditure levels and price indexes were downloaded directly from the Bureau of Economic Analysis on July 19, 2010. Transformation codes are (1) no transformation, (4) log-level, and (5) log-change. All variable descriptions and pneumonics are from the original sources.

Variable	Code	Description	Trans.	Source
Real Output				
1	IPS10	INDUSTRIAL PRODUCTION INDEX - TOTAL INDEX	5	BASIC
2	IPS11	INDUSTRIAL PRODUCTION INDEX - PRODUCTS, TOTAL	5	BASIC
3	IPS299	INDUSTRIAL PRODUCTION INDEX - FINAL PRODUCTS	5	BASIC
4	IPS12	INDUSTRIAL PRODUCTION INDEX - CONSUMER GOODS	5	BASIC
5	IPS13	INDUSTRIAL PRODUCTION INDEX - DURABLE CONSUMER GOODS	5	BASIC
6	IPS18	INDUSTRIAL PRODUCTION INDEX - NONDURABLE CONSUMER GOODS	5	BASIC
7	IPS25	INDUSTRIAL PRODUCTION INDEX - BUSINESS EQUIPMENT	5	BASIC
8	IPS30	INDUSTRIAL PRODUCTION INDEX - CONSTRUCTION SUPPLIES	5	BASIC
9	IPS32	INDUSTRIAL PRODUCTION INDEX - MATERIALS	5	BASIC
10	IPS34	INDUSTRIAL PRODUCTION INDEX - DURABLE GOODS MATERIALS	5	BASIC
11	IPS38	INDUSTRIAL PRODUCTION INDEX - NONDURABLE GOODS MATERIALS	5	BASIC
12	IPS43	INDUSTRIAL PRODUCTION INDEX - MANUFACTURING (SIC)	5	BASIC
13	IPS306	INDUSTRIAL PRODUCTION INDEX - FUELS	5	BASIC
14	IPS307	INDUSTRIAL PRODUCTION INDEX - RESIDENTIAL UTILITIES	5	BASIC
15	IPS308	INDUSTRIAL PRODUCTION INDEX - EQUIPMENT, TOTAL	5	BASIC
16	IPS311	INDUSTRIAL PRODUCTION INDEX - OIL AND GAS WELL DRILLING AN	5	BASIC
17	IPS316	INDUSTRIAL PRODUCTION INDEX - BASIC METALS	5	BASIC
18	GVP21	GROSS VALUE OF PRODUCT - INTERMEDIATE PRODUCTS–GROSS VALUE	5	BASIC
19	IPS45	INDUSTRIAL PRODUCTION INDEX - DURABLE MANUFACTURING (NAICS)	5	BASIC
20	IPS57	INDUSTRIAL PRODUCTION INDEX - NONDURABLE MANUFACTURING (NAI	5	BASIC
21	UTL11	CAPACITY UTILIZATION - MANUFACTURING (SIC)	1	BASIC
22	PMP	NAPM PRODUCTION INDEX (PERCENT)	1	BASIC
Unemployment				
23	LHUR	UNEMPLOYMENT RATE: ALL WORKERS, 16 YEARS & OVER (%,SA)	1	BASIC
24	LHU680	UNEMPLOY.BY DURATION: AVERAGE(MEAN)DURATION IN WEEKS (SA)	1	BASIC
25	LHU5	UNEMPLOY.BY DURATION: PERSONS UNEMPL.LESS THAN 5 WKS (THOUS.	1	BASIC
26	LHU14	UNEMPLOY.BY DURATION: PERSONS UNEMPL.5 TO 14 WKS (THOUS.,SA)	1	BASIC
27	LHU15	UNEMPLOY.BY DURATION: PERSONS UNEMPL.15 WKS + (THOUS.,SA)	1	BASIC
28	LHU26	UNEMPLOY.BY DURATION: PERSONS UNEMPL.15 TO 26 WKS (THOUS.,SA	1	BASIC
29	LHELX	EMPLOYMENT: RATIO; HELP-WANTED ADS:NO. UNEMPLOYED CLF	4	BASIC
Employment				
30	LHEL	INDEX OF HELP-WANTED ADVERTISING IN NEWSPAPERS (1967=100;SA)	5	BASIC
31	LHEM	CIVILIAN LABOR FORCE: EMPLOYED, TOTAL (THOUS.,SA)	5	BASIC
32	LHNAG	CIVILIAN LABOR FORCE: EMPLOYED, NONAGRIC.INDUSTRIES (THOUS.,	5	BASIC
33	CES0000	ALL EMPL: TOT NFARM,	5	USCEN
34	CES0500	ALL EMPL: TOT PRIV,	5	USCEN
35	CES0600	ALL EMPL: GDS PRODUCING,	5	USCEN
36	CES1000	ALL EMPL: MINING AND LOGGING	5	USCEN
37	CES2000	ALL EMPL: CONSTR,	5	USCEN
38	CES3000	ALL EMPL: MFG,SA (M)	5	USCEN
39	CES3100	ALL EMPL: DUR GDS,	5	USCEN
40	CES3200	ALL EMPL: NON-DURABLES,SA (M)	5	USCEN
41	CES0700	ALL EMPL: SVC	5	USCEN
42	CES4000	ALL EMPL: TRADE,TRANSPORT,&UTILITIES,SA (M)	5	USCEN
43	CES4200	ALL EMPL: RETAIL TRADE,SA (M)	5	USCEN
44	CES4142	ALL EMPL: WHOLESALE TRADE,WHOLESALE TRADE,SA (M)	5	USCEN
45	CES5500	ALL EMPL: FIN ACTIVITIES,SA (M)	5	USCEN
46	CES0800	ALL EMPL: PRIV SVC,	5	USCEN
47	CES9000	ALL EMPL: GOVT,	5	USCEN
48	PMEMP	NAPM EMPLOYMENT INDEX (PERCENT)	1	BASIC
49	CES155	AVG WKLY OVERTIME HOURS, PROD WRKRS, NONFARM - MFG	1	BASIC
50	CES154	AVG WKLY HOURS, PROD WRKRS, NONFARM - MFG	1	BASIC
Earnings				

Continued on next page ...

Variable	Code	Description	Trans.	Source
51	YPR	PERS INCOME CH 2000 $,SA-US	5	USCEN
52	A0M051	PERS INCOME LESS TRSF PMT (AR BIL. CHAIN 2000 $),SA-US	5	USCEN
Housing				
53	HSFR	HOUSING STARTS:NONFARM(1947-58);TOTAL FARM&NONFARM(1959-)(TH	4	BASIC
54	HSNE	HOUSING STARTS:NORTHEAST (THOUS.U.)S.A.	4	BASIC
55	HSMW	HOUSING STARTS:MIDWEST(THOUS.U.)S.A.	4	BASIC
56	HSSOU	HOUSING STARTS:SOUTH (THOUS.U.)S.A.	4	BASIC
57	HSWST	HOUSING STARTS:WEST (THOUS.U.)S.A.	4	BASIC
58	HSBR	HOUSING AUTHORIZED: TOTAL NEW PRIV HOUSING UNITS (THOUS.,SAA	4	BASIC
59	HMOB	MOBILE HOMES: MANUFACTURERS' SHIPMENTS (THOUS.OF UNITS,SAAR)	4	BASIC
Invent. & Orders				
60	PMI	PURCHASING MANAGERS' INDEX (SA)	1	BASIC
61	PMNV	NAPM INVENTORIES INDEX (PERCENT)	1	BASIC
62	PMNO	NAPM NEW ORDERS INDEX (PERCENT)	1	BASIC
63	PMDEL	NAPM VENDOR DELIVERIES INDEX (PERCENT)	1	BASIC
64	MOCMQ	NEW ORDERS (NET) - CONSUMER GOODS & MATERIALS, 1996 DOLLARS	5	BASIC
65	MSONDQ	NEW ORDERS, NONDEFENSE CAPITAL GOODS, IN 1996 DOLLARS (BCI)	5	BASIC
Stock Prices				
66	FSPCOM	S&P'S COMMON STOCK PRICE INDEX: COMPOSITE (1941-43=10)	5	BASIC
67	FSPIN	S&P'S COMMON STOCK PRICE INDEX: INDUSTRIALS (1941-43=10)	5	BASIC
68	FSDXP	S&P'S COMPOSITE COMMON STOCK: DIVIDEND YIELD (% PER ANNUM)	5	BASIC
69	FSPXE	S&P'S COMPOSITE COMMON STOCK: PRICE-EARNINGS RATIO (%,NSA)	5	BASIC
70	FSDJ	COMMON STOCK PRICES: DOW JONES INDUSTRIAL AVERAGE	5	BASIC
Exchange Rates				
71	RX@SZ	EXCHANGE RATE - SWISS FRANCS PER U.S. DOLLAR - SWITZERLAND	5	DRIIN
72	RX@JP	EXCHANGE RATE - YEN PER US DOLLAR - JAPAN	5	DRIIN
73	RX@UK	EXCHANGE RATE - U.K. POUNDS PER U.S. DOLLAR - UNITED KINGDOM	5	DRIIN
74	RX@CN	EXCHANGE RATE - CANADIAN DOLLAR PER US DOLLAR - CANADA	5	DRIIN
Interest Rates				
75	FYFF	INTEREST RATE: FEDERAL FUNDS (EFFECTIVE) (% PER ANNUM,NSA)	1	BASIC
76	FYGM3	INTEREST RATE: U.S.TREASURY BILLS,SEC MKT,3-MO.(% PER ANN,NS	1	BASIC
77	FYGM6	INTEREST RATE: U.S.TREASURY BILLS,SEC MKT,6-MO.(% PER ANN,NS	1	BASIC
78	FYGT1	INTEREST RATE: U.S.TREASURY CONST MATURITIES,1-YR.(% PER ANN	1	BASIC
79	FYGT5	INTEREST RATE: U.S.TREASURY CONST MATURITIES,5-YR.(% PER ANN	1	BASIC
80	FYGT10	INTEREST RATE: U.S.TREASURY CONST MATURITIES,10-YR.(% PER AN	1	BASIC
81	FYAAAC	BOND YIELD: MOODY'S AAA CORPORATE (% PER ANNUM)	1	BASIC
82	FYBAAC	BOND YIELD: MOODY'S BAA CORPORATE (% PER ANNUM)	1	BASIC
Spreads				
83	SFYG	SPREAD: FYGM3-FYFF	1	AC
84	SFYG	SPREAD: FYGM6-FYFF	1	AC
85	SFYG	SPREAD: FYGT1-FYFF	1	AC
86	SFYG	SPREAD: FYGT5-FYFF	1	AC
87	SFYGT	SPREAD: FYGT10-FYFF	1	AC
88	SFYA	SPREAD: FYAAA-FYFF	1	AC
89	SFYB	SPREAD: FYBAA-FYFF	1	AC
Money & Credit				
90	FM1	MONEY STOCK: M1(CURR,TRAV.CKS,DEM DEP,OTHER CK'ABLE DEP)(BIL	5	BASIC
91	FM2	MONEY STOCK:M2(M1+O'NITE RPS,EURO$,G/P&B/D MMMFS&SAV&SM TIME	5	BASIC
92	FMFBA	MONETARY BASE, ADJ FOR RESERVE REQUIREMENT CHANGES(MIL$,SA)	5	BASIC
93	FMRRA	DEPOSITORY INST RESERVES:TOTAL,ADJ FOR RESERVE REQ CHGS(MIL$	5	BASIC
94	FMRNBA	DEPOSITORY INST RESERVES:NONBORROWED,ADJ RES REQ CHGS(MIL$,S	5	BASIC
95	FCLBMC	WKLY RP LG COM'L BANKS:NET CHANGE COM'L & INDUS LOANS(BIL$,S	1	BASIC
96	CCINRV	CONSUMER CREDIT OUTSTANDING - NONREVOLVING(G19)	5	BASIC
97	A0M101	COML&IND LOANS OUTST,SA-US	5	USCEN
Survey				
98	U0M083	LEADING INDEX COMP: INDEX OF CONSUMER EXPECT,NSA-US	1	USCEN
Prices				
99	PMCP	NAPM COMMODITY PRICES INDEX (PERCENT)	1	BASIC
100	JNS@CRB	SPOT MKT PRICE INDEX-ALL COMMO (22) (CRB),NSA-US	5	USCEN
101	PWFSA	PRODUCER PRICE INDEX: FINISHED GOODS (82=100,SA)	5	BASIC
102	PWFCSA	PRODUCER PRICE INDEX:FINISHED CONSUMER GOODS (82=100,SA)	5	BASIC
103	PWIMSA	PRODUCER PRICE INDEX:INTERMED MAT.SUPPLIES & COMPONENTS(82=1	5	BASIC
104	PWCMSA	PRODUCER PRICE INDEX:CRUDE MATERIALS (82=100,SA)	5	BASIC
105	PUNEW	CPI-U: ALL ITEMS (82-84=100,SA)	5	BASIC
106	PU83	CPI-U: APPAREL & UPKEEP (82-84=100,SA)	5	BASIC
107	PU84	CPI-U: TRANSPORTATION (82-84=100,SA)	5	BASIC
108	PU85	CPI-U: MEDICAL CARE (82-84=100,SA)	5	BASIC
109	PUC	CPI-U: COMMODITIES (82-84=100,SA)	5	BASIC
110	PUCD	CPI-U: DURABLES (82-84=100,SA)	5	BASIC
111	PUXF	CPI-U: ALL ITEMS LESS FOOD (82-84=100,SA)	5	BASIC
112	PUXHS	CPI-U: ALL ITEMS LESS SHELTER (82-84=100,SA)	5	BASIC
113	PUXM	CPI-U: ALL ITEMS LESS MIDICAL CARE (82-84=100,SA)	5	BASIC
114	PCIUSER	CPI-U: SVC,SA-US	5	USCEN
115	PPCE	PRICE INDEX: PERSONAL CONSUMPTION EXPENDITURES	5	B.E.A.
116	PPHC	PRICE INDEX: HOUSEHOLD CONSUMPTION EXPENDITURES	5	B.E.A.
117	PPCC	PRICE INDEX: PERSONAL CONSUMPTION EXPENDITURES EXCLUDING FOO	5	B.E.A.
118	PDUR	PRICE INDEX: DURABLE GOODS	5	B.E.A.
119	PCAF	PRICE INDEX: CLOTHING, FOOTWEAR, AND RELATED SERVICES	5	B.E.A.
120	PHUF	PRICE INDEX: HOUSING, UTILITIES, AND FUELS	5	B.E.A.
121	PFHH	PRICE INDEX: FURNISHINGS, HOUSEHOLD EQUIPMENT, AND ROUTINE H	5	B.E.A.

Continued on next page ...

Variable	Code	Description	Trans.	Source
122	PHLT	PRICE INDEX: HEALTH	5	B.E.A.
123	PRRL	PRICE INDEX: RECREATION	5	B.E.A.
124	POIS	PRICE INDEX: OTHER GOODS AND SERVICES	5	B.E.A.
125	PFSA	PRICE INDEX: FOOD SERVICES AND ACCOMMODATIONS	5	B.E.A.
126	PTRN	PRICE INDEX: TRANSPORTATION	5	B.E.A.
127	PCMC	PRICE INDEX: COMMUNICATION	5	B.E.A.
128	PEDU	PRICE INDEX: EDUCATION	5	B.E.A.
129	PFXA	PRICE INDEX: FOOD AND BEVERAGES PURCHASED FOR OFF-PREMISES C	5	B.E.A.
Cons. Expend.				
130	PCE	EXPENDITURE: PERSONAL CONSUMPTION EXPENDITURES	5	B.E.A.
131	PHC	EXPENDITURE: HOUSEHOLD CONSUMPTION EXPENDITURES	5	B.E.A.
132	PCC	EXPENDITURE: PERSONAL CONSUMPTION EXPENDITURES EXCLUDING FOO	5	B.E.A.
133	DUR	EXPENDITURE: DURABLE GOODS	5	B.E.A.
134	CAF	EXPENDITURE: CLOTHING, FOOTWEAR, AND RELATED SERVICES	5	B.E.A.
135	HUF	EXPENDITURE: HOUSING, UTILITIES, AND FUELS	5	B.E.A.
136	FHH	EXPENDITURE: FURNISHINGS, HOUSEHOLD EQUIPMENT, AND ROUTINE H	5	B.E.A.
137	HLT	EXPENDITURE: HEALTH	5	B.E.A.
138	RRL	EXPENDITURE: RECREATION	5	B.E.A.
139	OIS	EXPENDITURE: OTHER GOODS AND SERVICES	5	B.E.A.