What moves real GNP?*

Harald Uhlig
Humboldt University Berlin
CentER, Tilburg University
and CEPR

PRELIMINARY
COMMENTS WELCOME

First draft: March 8th, 1999
This revision: April 8, 2003

Abstract

This paper aims at identifying the main shocks, which cause movements in real GNP. It does so by searching for two shocks in the context of a VAR model, which explain the majority of the k-step ahead prediction error variances in real GNP for horizons between 0 and 5 years.

We find that two shocks can typically explain more than 90% of the variance at all horizons for real GNP. While one shock looks like a productivity shock in the line of the real business cycle literature, the other one seems to be wage-push or inflationary shock, unrelated to consumption or government spending and not induced by monetary policy. While the first shock can be viewed as a "supply shock", the second shock does not have an obvious "demand shock" interpretation.

*This paper grew out of a joint project, called "What are the main macroeconomic forces", with Lucrezia Reichlin, and owes a substantial intellectual debt to her input and ideas and to the conversations we had on this topic. I am grateful for comments in particular by Peter Kennen as well as by seminar audiences at a NASEF conference on Hydra, at UCL and in Mainz. The address of the author is Prof. Harald Uhlig, Humboldt University, Wirtschaftswissenschaftliche Fakultät, Spandauer Str. 1, 10178 Berlin, GERMANY. e-mail: uhlig@wiwi.hu-berlin.de, fax: +49-30-2093 5934, home page http://www.wiwi.hu-berlin.de/wpol/

1 Introduction

This paper aims at identifying the main shocks, which cause movements in real GNP. It does so by searching for a limited number of shocks, which explain the majority of the movements in a given set of time series.

The idea that a limited number of shocks are sufficient to explain most of the movements in a possibly large set of macroeconomic aggregates, comes from the well-documented observation that business cycles are characterized by a strong comovement between the main macroeconomic time series at business cycle frequencies.

Theorists who want to construct small, insightful macroeconomic models therefore find it desirable to concentrate the analysis on a few shocks and their propagation rather than to capture all of the movements within one theory.

Thus, for example, many papers of the real business cycle school concentrate their analysis solely on the consequences of productivity shocks, whereas papers which focus on monetary analysis often only model monetary shocks plus perhaps a "demand" and "supply" shock. The implicit hope is that these models provide insights into the main dynamic mechanisms at work in an economy, which explain perhaps not all, but at least a substantial part of the movement we see in the data. The fraction of the variation explained in the data by these shocks could be seen as a measure of relevance of these theories, and are therefore the focus of some investigations. For example, it is claimed that productivity shocks explain about 70 percent of the movements in GDP, while monetary shocks appear to account for fairly little in the GDP movements.

But rather than postulating a particular type of shock and investigating it theoretically in the hope of explaining much of what one sees, one might want to proceed the other way around and construct theories in the light of empirical evidence of what actually moves macroeconomic time series. If there were a few shocks, which indeed empirically explained the majority of the time series movements, and which at the same time can be given an appealing theoretical interpretation, constructing theories focussing on these ought to be a fruitful exercise. On the other hand, if many shocks are empirically needed to explain a substantial fraction of aggregate movements in the data, there would be little hope to explain most of it with a "one-shock-fits-all" macroeconomic theory. In sum, before constructing a macroeconomic theory, the theorist may seek to find empirical insight into the question: "what are the main macroeconomic forces?"

This paper aims at providing some answers to this question, and to provide tools to study this question in greater detail. This paper is largely atheoretical: it should be viewed as "food for thought" for theorists. As a consequence, we give at most a tentative economic interpretation to our findings. While this paper has a methodological component, we have tried to separate the discussions of our empirical results from the more methodological discussion in the hope that the results and their interpretation can be appreciated and understood by a broader audience with a minimum of methodological details. Obviously, the understanding of the results gets deeper when understanding exactly how they were obtained, so the paper also spends a substantial fraction of space on detailing the methodology.

The paper is organized as follows. Section 2 provides a guide to the methodology employed, concentrating on a heuristic rather than a detailed description and providing only what is needed to appreciate the results. They are then presented in section 3 for a two-shock analysis and in section 4, when focussing on a single shock. After that, the paper discusses the methodology in detail. Section 5 summarizes some VAR basics. Section 6 provides the details on a general approach, which allows us to make the question in the title of the paper precise, and to identify shocks which are the main movers. Finally, section 7 concludes.

2 A brief guide to the methodology and some results

Our approach works as follows. Using a vector autoregression, we seek to find "shocks" which may be responsible for most of the movements in GNP, and study their propagation. We are interested in particular in the movements up to several years after the shock, i.e., our focus will not

be restricted to explaining as much as possible of the instantaneous movements in GNP.

We estimate a core VAR, using seven variables at quarterly frequencies, a constant and six lags. 14 further variables are added individually as periphery variables: while lagged and contemporaneous variables of the core variables as well as own lags are right-hand-side regressors, they do not in turn show up as regressors for the core variables or for other periphery variables. We use a Bayesian methodology for inference.

The data is from 1964 to 2001, has been obtained from the Federal Reserve Bank of St. Louis Web Site, and is in detail:

1. Real GNP (GNPC96)
2. labor productivity (OPHNFB)
3. real hourly wages (COMPRNFB) divided by labor productivity (OPHNFB)
4. Oil Price (OILPRICE)
5. CPI Inflation (calculated from CPIAUCSL)
6. Goverm. Spending (GCEC1)
7. Fed. Funds Rate (FEDFUNDS)

for the core and

1. real nondurable consumption (PCENDC96)
2. Private Investment (FPIC1)
3. Hours worked (AWHI)
4. Industrial Production (INDPRO)
5. Capacity Utilization (CUMFG)

6. S&P 500, cumulated returns (TRSP500)
7. 10yr bond rate (GS10)
8. PPI Inflation (calculated from PFCGEF)
9. Nonborred Reserves (BOGNONBR)
10. M1 (M1SL)
11. M3 (M3SL)
12. Exchange Rate (TWEXBMTH)
13. Exports (EXPGSC1)
14. Imports (IMPGSC1)

for the periphery. The variables have been given the obvious transformations: most variables have been used in logarithmic form. Where variables are already in percent, they have not been transformed further. Table 1 provides a list of the standard deviations of the VAR-MLE one-step ahead prediction errors of the transformed variables, all expressed in percent (i.e., a standard deviation of 0.02 for a variable used in logarithmic form is noted as 2%, etc..).

The idea here is to restrict the core VAR to a reasonably small list of variables in order to avoid the obvious degrees-of-freedom problem, that the number of to-be-estimated coefficients grows with the square of the number of variables. We have chosen as core variables a list which are prime suspects as movers of business cycles in macroeconomic reasoning and theories:

1. **Real GNP.** Real GNP is included as the focus variable of interest.

2. **Labor productivity.** Real business cycle theory has focussed on productivity shocks as the key driving force for business cycles: for that reason, labor productivity has been included in the core. The real business cycle theories focus on total factor productivity rather than labor productivity: the difference lies in accounting for movements in capital. However, in practically all of these theories, capital is predetermined, so that there is no difference between the unpredictable surprise in total factor productivity or labor productivity. Furthermore, since a series for labor productivity was readily available, including it seemed more natural than constructing a series for total factor productivity.

3. **real hourly wages divided by labor productivity.** High unemployment is frequently blamed on wages exceeding productivity, and wage restraint e.g. in union renegotiations is offered as a cure. Therefore, fluctuations in employment and GNP might have their cause in movements in the wedge between wages and labor productivity.

4. **Oil Price.** Oil price changes have often been argued to be a major cause of recessions, and are the second key factor in determining production costs, next to the wage-productivity gap.

5. **CPI Inflation.** While the wage-productivity gap and oil prices focus on price movements on the input side in the production process, CPI inflation focusses on price movements on the output side. If a large fraction of firms decides to suddenly raise their prices, this can lead to a dampening in demand and a business cycle slow down.

6. **Goverm. Spending.** Fiscal policy still is high on the list of causes for business cycle movements for many macroeconomists, despite the fact that government spending is fairly acyclical.

7. **Federal Funds Rate.** Monetary policy has often been blamed for causing economic fluctuations, such as the recessions at the beginning of the 1980s, despite considerable amount of research over the last decade, which has argued, that monetary policy shocks are unlikely to be a major contributor.

A number of other theories - e.g. sunspot shocks and aggregate increasing returns to scale - would imply certain comovements between some of these key variables. Furthermore, some prominent theories would be hard to fit into a clear pattern here. In particular, explanations relying on export-led growth or on autonomous exchange rate movements as a key force in driving the business cycle, might not have particularly clear implications with respect to the variables above. Much of this can be investigated with the help of the periphery variables, however. E.g., if indeed export-led growth is a key motor behind business cycle movements, we should see those shocks that drive GNP movements to also have concurrent and large effects on exports, when including that variable as periphery.

Consider now a VAR in the core variables or the core variable plus a periphery variable, and consider a shock, given by som identification procedure. One can then calculate the impulse responses for this shock as well as the variance of the k-step ahead prediction error which is due to future occurences of this

shock during the next k periods. What we will focus on is the sum of these variances for GNP. We shall consider the sum running from $k = 0$ to $k = 19$, i.e., 5 years ahead, which should cover short- as well as medium-run movements. We also investigate only the sum of variances for the first year, $k = 0$ to $k = 3$, thereby investigating in particular the causes for rather immediate fluctuations, as well as the variances for years three to five, $k = 8$ to $k = 19$.

The core idea of our approach now is the following. Rather than identify some shock as, say, a monetary policy shock, a productivity shock or a fiscal policy shock, and investigate its contribution to the variance of the $k-$step ahead prediction error of GNP, we will identify one or two shocks, which explain as much as possible of the chosen sum of the variances for GNP. I.e., what we are looking for are the major shocks or the main macroeconomic forces, driving GNP. The method is related to the study of principal components, but is not a straight-forward application: a bit of math to figure out, which matrix is the relevant one for calculating the principal components is required. The details are in section 6.

Obviously, there is a degree of arbitrariness here. While it is perhaps clear that we should focus on the explainable variance of GNP movements, and while we think the choices made are reasonable ones, one may consider shorter or longer horizons, or focus on explaining as much as possible of the variance at a particular horizon or calculate the variance of some filtered response. All this is possible in principle with this method.

When looking for two major shocks, a crucial point needs to be cleared up. Identifying two shocks will amount to finding the shocks corresponding to the two largest eigenvalues of some matrix. It may seem natural to identify the first shock as the one which corresponds to the first eigenvalue and the second shock as the one corresponding to the second eigenvalue. That would be misleading, however. Our aim is to find the two shocks which generate the *total* of that variation: disentangling them into individually meaningful shocks is still an exercise on its own.

We solve this problem by providing results for several possible decompositions. Roughly speaking, if some vector a_1 represents the shock corresponding to the first eigenvalue and a_2 represents the shock corresponding to the second eigenvalue, then all pairs $[a(\theta), a(\theta+\pi/2)]$, where

$$a(\theta) = \cos(\theta * \frac{2\pi}{360})a_1 + \sin(\theta * \frac{2\pi}{360})a_2, \tag{1}$$

represent legitimate (orthogonal) decompositions as well, still explaining in total all the variation measured by these first two eigenvalues (the factor $2\pi/360$ is there to measure θ in degrees rather than fractions of 2π).

When identifying two shocks, we therefore present results for

$$\theta \in \{0; 30; 60; 90; 120; 150\}$$

rather than just for a_1 and a_2. The legitimate pairings arising from (orthogonal) shocks would then be $\theta \in \{0; 90\}$, $\theta \in \{30; 120\}$ and $\theta \in \{60; 150\}$: further "rotations" by 30 degrees only result in repetitions (possibly with a flipped sign). The shock a_1 corresponds to $\theta = 0$ and the shock a_2 corresponds to $\theta = 90$.

3 Identifying two main shocks: results.

There are three sets of results. The first set of results is for two shocks, explaining as much as possible of the GNP variance during the first five years after the shock. The results are shown in figures 7 to 12 and described below. The second and third set of results are for one shock only, and described in section 4.

We first investigate two shocks, explaining as much as possible of the sum of the variances of the k-step ahead prediction errors for GNP, where $k = 0$ to $k = 19$, i.e., for the first five years after the shock. The results are shown in figures 7 to 12. Six figures are necessary for the two-shock case in order to investigate the various rotations, as described above. Figures 7 and 10 contain the results for the seven core variables, whereas the other figures contain the results for the 14 periphery variables. The variance decompositions are summarized in tables 2 and 3.

3.1 Variance decompositions

In figures 7 to 9, the fraction of the k-step ahead prediction error for the various variables which can be explained by both shocks together (first column) or by a single shock (columns 2 to 7) is shown. Columns 2 and 3 contain the orthogonal pairing for $\theta = 0$ and $\theta = 90$, as explained above, and show the fraction of variance at each horizon by each of the two shocks individually. Likewise, columns 4 and 5 contain the orthogonal pairing $\theta = 30$ and $\theta = 150$ and columns 6 and 7 contain the orthogonal pairing $\theta = 60$ and $\theta = 150$.

Each figure contains four lines. The three curved lines denote the median together with error bands denoting the 16% and the 84% quantile. The forth, horizontal line is drawn at the value of 50% for the fraction explained, and simply serves as a visual aid for comparison. Thus, if e.g. the median is above the horizontal 50% line everywhere, this means that (at the median estimate), the shock(s) can explain at least as much as 50% of the k-step ahead prediction error variance for any horizon k up to five years. The upper range in all graphs is set at 100%, again for aiding visual comparison.

Examine the first row, providing the variance decomposition for GNP. As one can see, the two shocks together explain vastly more than 50%: closer inspection of the figure reveals that the median estimate is never below 80% after the first half year, and mostly stays close to or above 90%. This means, that two shocks really are enough to explain practically all of the GNP movements, except very immediately after the shock and - to a much lesser extent - two years after the shock. Examining the split of the total explained variance on each individual shock, depending on θ, we see that the pairing $\theta = 0$ and $\theta = 90$ is particularly intriguing: there, the $\theta = 0$-shock explains most of the variation at longer horizons, whereas the $\theta = 90$ shock explains most of the variation at shorter horizons. We shall focus on this pairing in the following discussion, and call them the "medium-run" and the "short-run" shock. We also perform a one-shock analysis for the short run in subsection 4.1 and the medium run in subsection 4.2.

The first column of table 2 provides summary for all the other variables, using both shocks together (i.e. for the "Total" column in the figures 7 to 9). A "1" in this table indicates the quarter, where the median line is above the 50% comparison line for a particular quarter.

In light of the figures and this first column, we find that

1. The two shocks together explain easily more than the majority of the fluctuations at nearly all prediction horizons for real GNP, labor productivity, real nondurable consumption, private investment, hours worked, industrial production and imports.

2. They explain around half of the fluctuations for most of the prediction horizons for the wage-productivity gap, CPI inflation, capacity utilization and the S&P 500 composite returns.

3. With a one- to two-year lag, these shocks explain around half of the fluctuations of government spending, the federal funds rate and M3 and, to a lesser extent, the oil price, the 10-year bond rate and exports.

4. They never explain more than half of the fluctuations for PPI inflation, nonborrowed reserves, M1 and the exchange rate.

Some important conclusions can already be drawn here.

1. The list of variables, for which the two shocks explain the majority of the fluctuations, are pretty much the variables typically emphasized by real business cycle theorists. The one notable addition (at least compared to the closed-economy RBC literature) are imports.

2. This perspective does not change, when taking into account the next list of variables, for which around half of the fluctuations are explained for most horizons, although these variables indicate which modification of the benchmark real business cycle theory might be particularly suitable to account for the facts shown here. E.g. CPI inflation is largely driven by these shocks as well, as is capacity utilitization. A substantial amount of the movement in the wage-productivity gap is explained by these shocks too: presumably, this gap is an important part of the story.

3. The variables which play a large role in the literature on monetary policy shocks are either explained only with a lag (like the federal funds rate) or not in a substantial way at all, in particular, nonborrowed reserves, M1 and the exchange rate. This is consistent with the prevalent view that the Federal Reserve Bank does not role its dice when deciding on monetary policy and thus, that monetary policy shocks contribute little to business cycle movements. This still leaves open the question, why so little of the M1 movements or the movements in nonborrowed reserves are explained by these two shocks. A considerable degree of monetary neutrality would be one avenue to answer it.

4. For other variables which have received considerable attention as movers of the business cycle, like oil prices or government spending, most of their fluctations are explained with a lag. Again, this suggests that these variables react to the state of the business cycle rather than causing them.

While the first column of table 2 showed the variance decomposition results for both shocks, the second column summarizes the findings with respect to the medium-run shock, $\theta = 0$. Here, the interpretation of a productivity shock emphasized by the real business cycle literature, suggests itself even more strongly: this shock alone explains less than 50% for all monetary and fiscal variables, except M3. Interestingly, it also explains less than 50% of the capacity utilization movements.

The short-run shock, $\theta = 90$, explains far less of the variance of GNP than the medium-run shock: the median line crossed the 50% comparison line only briefly in the third figure of the top row in figure 7. Therefore, table 3 lists, where that shock explains more than 20% of the prediction error variance: the results are in the second column, whereas the first column provides the results for the first shock, when also comparing it to 20% rather than 50% as above.

The short-run shock does not lend itself easily to an immediate interpretation. Monetary policy might have something to do with it, as the short-run shock seems to explain fluctuations in the federal funds rate and nonborrowed reserves more easily than the medium-run shock: however, both variables still only seem to react with a lag. Thus, something else is probably the cause for this shock. Interestingly, the short-run shock seems to more easily account for capacity utilization fluctuations than the long-run shock, and accounts more easily for export fluctuations at the short horizon. Furthermore, the short-run shock seems to have little to nothing to do with consumption fluctuations or government spending. Thus, while the medium-term shock could also be called a "supply shock", it would be hard to call the short-run shock a "demand shock".

In evaluating these two shocks, one needs to keep in mind, that the short-run shock only explains on average about a third of the GNP fluctuations during the two-year horizon following the shock and 15% on average during the entire five year horizon, the medium-run shock not only explains 70% on average and nearly 90% during years three to five after the shock, it also explains around 43% on average during the first two years after the shock. I.e., the medium-run shock is a considerably larger contributor to the causes of GNP fluctuations than is the short-run shock, and even during the first two years, it contributes even more. Only when focussing on the first year alone does the table turn: there, the medium-run shock explains as much as 27% on average, while the short-run shock explains nearly 45%.

3.2 Impulse responses

So far, we have only investigated variance decompositions. What is important is to check whether the interpretations above hold up to further scrutiny, when examining impulse responses. For example, the real business cycle interpretion would fall apart, should productivity and GNP move in opposite directions.

Figures 10 to 12 show the impulse responses for the various orthogonal pairings. The three curved lines denote the 16%, the 50% (median) and the 84% quantile. The three straight horizontal lines denote the zero response plus minus one standard deviation of the VAR-MLE one-step ahead prediction errors of the series for comparison. For the non-core variables, the prediction error is calculated, using also the contemporaneous core variables as regressors: as a result, the core shocks can have an impact on non-core variables exceeding their one-standard deviation band. The scale is the same in each row.

There is a certain arbitrariness in assigning a sign to the impulse responses to a shock: we have chosen to multiply all the impulse responses with the sign of the GNP response two quarters after the shock, in order to provide meaningful pictures.

We concentrate on the description of the pairing $\theta = 0$ and $\theta = 90$, with the caveat, that choosing this pairing is not a result of the empirical analysis, but a hopefully sensible interpretation. We refer to the estimates along the median impulse response, which - as the figures show - are obviously subject to considerable standard deviations.

3.2.1 The medium-run shock, $\theta = 0$.

With regards to the medium-term $\theta = 0$ shock, we observe the following:

1. Productivity jumps by about half a standard deviation of the MLE one step ahead prediction error, i.e. by about 0.25%, and gradually and monotically climbs to a

plateau of a full standard deviation of an overall increase by 0.5% within three years. In other words, this looks like a persistent productivity shock or some otherwise induced persistent jump in productivity, where some marginal improvements follow the initial larger jump.

2. On impact, the wage-productivity gap jumps down by 0.25%, i.e. when productivity jumps, real wages do not move at all. The wage-productivity gap is closed rather monotonistically but very gradually: even five years later, there is still a gap of somewhere around 0.09%. Fitting an AR(1) to the median response shows that about 5% of the gap is eliminated each quarter.

3. real GNP grows steadily from somewhat less than 0.2% initially to 0.8% eventually after five years, compared to the no-shock scenario. Closer inspection of the median response (and subject to the usual caveats regarding standard errors), it seems that the first growth phase to an increase of about 0.5% takes about one year, after which real GNP reaches a plateau for another year, before growing again. Private consumption does nearly the same. Government spending, private investment, hours worked and industrial production all follow pretty much the same pattern, but on different scales: government spending grows eventually only by 0.5%, and hours worked grows by 0.6% to 0.7%, whereas private investment grows to 2% at its peak four years after the shock. Capacity utilization behaves somewhat erratically around a plateau 0.2% above its no-shock value.

4. CPI inflation is lower by a quarter of a percent on average during the first three years after the shock, and rises subsequently to a level 0.15% above its initial value five years after the shock. The oil price is lower by 1.3% on average, but nothing clear happens to PPI inflation.

5. The federal funds rate, while somewhat erratic, falls by 0.15% to 0.2% on average. The decline of the ten-year bond rate is somewhat less: 0.1% on average and 0.2% at its bottom two years after the shock. The stock market jumps up permantly by 3%. M1 and M3 both increase by approximately 1% eventually. Nothing clear happens to nonborrowed reserves.

6. The dollar appreciates gradually towards a peak of about 1% two years after the shock. While exports pick up only gradually three years after the shock, eventually climbing by more than 1%, imports react more swiftly, climbing towards a plateau about 2% above its inital level one year after the shock.

This looks very much like a productivity shock as in the real business cycle literature, although there are some obvious subtleties. For example, there is the rather persistent wage-productivity gap. The increase in GNP is gradual. Ten-year bond rates and short-term interest rates decline rather than rise. Available theories may be in need of modification in order to explain these facts.

3.2.2 The short-run shock, $\theta = 90$.

The second "short-run" $\theta = 90$ shock has the following features:

1. real GNP increases by as much as 0.5% one half year after the shock, then the response turns around and becomes negative three years later, dropping eventually by as much as 0.2%. Real nondurable consumption, private investment, hours worked, industrial production, capacity utilization, imports, and exports all show a similarly short-lived cyclical pattern. In particular, while hours worked do not react on impact for the medium-run shock above, they jump up immediately by nearly 0.2% for the short-run shock, reaching a peak of more than 0.4% half a year after the shock, before declining to -0.4% four to five years after the shock.

2. While nothing obvious happens to labor productivity, except for a very brief inital upward blip of about 0.1% during the first year, the wage-productivity gap turns positive after the first year, at a level of approximately 0.1% above its no-shock value.

3. Government spending rises by 0.3% two to three years after the shock, before returning back to (nearly) zero.

4. CPI inflation increases by 0.3% two years after the shock, before turning around, eventually dropping 0.2% below its no-shock level. This pattern is preceeded by a rather swift rise in the federal funds rate by more than 0.3turning negative four years after the shock. Ten-year bond rates rise too by about 0.1% on average. Nothing clear happens to the stock market. Whereas there is no clear pattern in M1, nonborrowed reserves fall by as much as 1.8% at its bottom one year after the shock, while M3 rises to a plateau 0.3% above its no-shock value two years after the shock. PPI inflation drops by nearly 0.8% one year after the shock, before levelling off again. The exchange rate shows no clear pattern.

5. The oil price drops by 2.5% during the first half year following the shock, but then rises swiftly to nearly 3.5% above its no-shock value two years after the shock, before levelling off.

This shock is clearly harder to interpret. It may be consistent with the following story. Inflationary pressures may be building up due to misjudged productivity signals and unjustified wage increases. The increased wage pressure is passed on by the firms to the consumer, with the Federal Reserve Bank anticipating and fighting the CPI inflation to come, balancing this against the negative effects on the economy. While this leads to a tighter supply for nonborrowed reserves, savings demands for M3-vehicles increases due to higher interest rates. The economy turns south both because of wages which are too high as well as the increase in interest rates.

4 Analyzing one shock only.

In this section, the results are described, when analyzing one shock only. The second set of results, shown in figure 6 and described in subsection 4.1 focus on short run movements and one shock, explaining as much as possible of the GNP variance during the first year after the shock, while the third set of results, shown in figure5 and described in subsection 4.2 focus on medium run movements and one shock, explaining as much as possible of the GNP variance during years three to five after the shock.

4.1 One shock: short run movements

In figure 6, we have tried to isolate the short-run shock by searching for a single shock, maximizing its contribution to explaining the variation of real GNP during the first year after the shock. That shock explains a substantial amount of fluctuation during the first year for real GNP, investment, hours worked, industrial production and imports, and, to a lesser extent, labor productivity, real nondurable consumption, capacity utilization, government spending (for years two to three after the shock) and M3 (after the first year).

The impulse responses show a pattern which appears to be a bit of a mixture between the patterns of the medium-run and the short-run shock above. The wage-productivity gap, for example, starts with a negative blip and then quickly returns to zero rather than turning positive as was the case for the short-run shock above. Likewise, real GNP, investment, hours worked, industrial production, capacity utilization and imports merely return to zero after their initial positive blip during their first year. Productivity is increased persistently by about half a standard deviation. Consumption is permanently increased by one (non-core prediction error) standard deviation above its no-shock level. There no longer is a clear pattern in CPI inflation, while the federal funds rate is still shown to be above its no-shock path during the first two years following the shock. We see the pattern documented above for the monetary aggregates, with nonborrowed reserves falling, M3 rising and M1 not showing a particularly distinctive pattern, although it now has a rising tendency.

We conclude that it is hard to truly isolate the short-term shock, when identifying just one shock. This should not surprise: given the numbers provided at the end of subsection 3.1, the shock isolated here should be expected to be something like two-fifth of the medium-run shock "mixed" with three-fifth of the short-run shock.

4.2 One shock: medium run movements

In figure 5, we have concentrated the analysis on the medium-run shock by searching for a single shock, maximizing its contribution to explaining the variation of real GNP during years three to five after the shock.

This exercise is a successful one: the patterns here are very similar to the patterns to the $\theta = 0$-shock disussed in the two-shock analysis, and need no further comment. We conclude that these may be rather robust features of the data, and that this shock is the major force moving GNP. The simplest explanation seems to be a productivity shock much as what has been proposed by real business cycle theorists, although standard theories are in need of some modifications to account for the all the key features here.

5 VAR Basics

We now turn to the nuts and bolts of our method. To fix notation and to introduce terms and expressions useful later on, when performing a particular principal components analysis, we need to start from describing vector autoregressions and impulse response analysis.

A VAR is given by

$$Y_t = c + B_{(1)}Y_{t-1} + B_{(2)}Y_{t-2} + \ldots + B_{(l)}Y_{t-l} + u_t, t = 1, \ldots, T \tag{2}$$

or, more compactly,

$$c + B(L)Y_t = u_t \tag{3}$$

where Y_t is a $m \times 1$ (covariance stationary) vector of data at date $t = 1 - l, \ldots, T$, $B_{(i)}$ are coefficient matrices of size $m \times m$ and u_t is the one-step ahead prediction error with variance-covariance matrix Σ. For ease of notation, we set $c = 0$ wlog. Furthermore, in case a non-core variable is added to the core VAR, we suppose that it is the last variable in Y_t: in that case, all coefficients in $B_{(1)}$ to $B_{(l)}$ in all other rows are zero, if they are the regression coefficient on the non-core variable.

We wish to decompose the m one-step ahead prediction errors in u_t into m mutually orthogonal innovations, and normalized to be of variance 1. Usually, this is done to get at economically meaningful or fundamental innovations. Here, our aim is simply to do proper innovation accounting as well as to find shocks which explain the majority of some variance to be specified.

What is needed is to find a matrix A such that $u_t = Av_t$ and so that $E[v_t v_t'] = I$. The j-th column of A (or its negative) then represents the immediate impact on all variables of the j-th innovation in v, one standard error in size. A needs to satisfy the restriction

$$\Sigma = E[u_t u_t'] = AE[v_t v_t']A' = AA' \tag{4}$$

We shall be interested only in a limited number of shocks. Given the discussion above, this amounts to finding a submatrix A_1 of A. The following definitions are useful

Definition 1 *The vector $a \in \mathbf{R}^m$ is called an* **impulse vector***, iff there is some matrix A, so that $AA' = \Sigma$ and so that a is a column of A.*

Definition 2 *The matrix $m \times p$ matrix $A_1 \in \mathbf{R}^m$ is called an* **impulse matrix** *of size p, iff there is some matrix A, so that $AA' = \Sigma$ and so that $A = [A_1, A_2]$ for some other matrix A_2.*

In particular, an impulse vector is an impulse matrix of size 1, and every column of an impulse matrix of size $p \geq 1$ is an impulse vector.

It is easy to characterize all impulse matrices A_1 as follows, see Faust (1999), Uhlig (1999) or Mountford-Uhlig (2001). For A_1, find A_2, so that $A = [A_1, A_2]$ satisfies $AA' = \Sigma$. Let $\tilde{A}$ be some other decomposition of Σ, $\tilde{A}\tilde{A}' = \Sigma$: the Choleksy decomposition is particularly useful here. Then, there must be an orthogonal matrix Q, i.e. an $m \times m$ matrix Q satisfying $QQ' = I$, so that $A = \tilde{A}Q$. Conversely, for any orthogonal matrix Q, partition Q into $Q = [Q_1, Q_2]$. Then, $A_1 = \tilde{A}Q_1$ is an impulse matrix.

If the VAR is stationary, write its inversion as

$$Y_t = C(L)u_t$$

where

$$C(L) = \sum_{l=0}^{\infty} C_l L^l$$

Let $\tilde{A}$ be the Cholesky decomposition[1] $\Sigma = \tilde{A}\tilde{A}'$, $\tilde{A}$ lower triangular, of the variance-covariance matrix Σ. One can obtain the impulse responses $\tilde{R}(L)$ to the shocks identified in this decomposition as

$$\tilde{R}(L) = C(L)\tilde{A}.$$

I.e.,

$$Y_t = \tilde{R}(L)\tilde{v}_t$$

where

$$u_t = \tilde{A}\tilde{v}_t$$

To find the responses to some other decomposition $AA' = \Sigma$, find the orthogonal matrix Q satisfying $A = \tilde{A}Q$, as explained above. Then, the impulse responses to the shocks identified in A are given by

$$R(L) = \tilde{R}(L)Q \tag{5}$$

With these results, one can calculate the variance-covariance matrix of the k-step ahead prediction error, and decompose it into m parts, with one part for each impulse vector in the decomposition A as follows. The k-step ahead prediction error of Y_{t+k}, given all the data up to and including $t-1$, is given by

$$e_{t+k}(k) = \sum_{l=0}^{k} \tilde{R}_l Q v_{t+k-l}$$

Note that $e_t(0) = u_t$. The variance-covariance matrix of $e_{t+k}(k)$ is

$$\Sigma(k) = \sum_{l=0}^{k} \tilde{R}_l \tilde{R}_l' \tag{6}$$

Note that $\Sigma(0) = \Sigma$. $\Sigma(k)$ can now be decomposed as

$$\Sigma(k) = \sum_{j=1}^{m} \Sigma(k, j) \tag{7}$$

[1] Any other decomposition would work equally well. We have chosen the Cholesky decomposition here to fix ideas, and because it is particularly easy to compute.

with

$$\Sigma(k,j) = \sum_{l=0}^{k} (\tilde{R}_l q_j)(\tilde{R}_l q_j)' \tag{8}$$

These relationships facilitate the search for impulse vectors or impulse matrices in some of our methods.

6 Explaining as much as possible of the variance

This method aims at finding an impulse vector a or an impulse matrix A_1, so that this impulse vector or matrix explains as much as possible of the sum of the k-step ahead prediction error variance of some variable i, say, for prediction horizons $k = \underline{k} \leq \bar{k}$. Focus on the case of a single impulse vector and a VAR, containing only the core variables, first. Formally then, the task is to explain as much as possible of the variance

$$\sigma^2(\underline{k}, \bar{k}) = \sum_{k=\underline{k}}^{\bar{k}} \Sigma(k)_{ii}$$

with a single impulse vector. Using the notation above, this amounts to finding a vector q_1 of unit length, maximizing the sum over $(\Sigma(k,1))_{ii}, k = \underline{k}, \ldots, \bar{k}$. Let $E_{(ii)}$ be a matrix only with zeros, except for $(E_{(ii)})_{ii} = 1$. Calculate

$$\begin{aligned}
\sigma^2(\underline{k}, \bar{k}; q_1) &= \sum_{k=\underline{k}}^{\bar{k}} \sum_{l=0}^{k} \left((\tilde{R}_l q_1)(\tilde{R}_l q_1)' \right)_{ii} \\
&= \sum_{k=\underline{k}}^{\bar{k}} \sum_{l=0}^{k} \text{ trace } \left(E_{(ii)} (\tilde{R}_l q_1)(\tilde{R}_l q_1)' \right) \\
&= q_1' S q_1
\end{aligned}$$

where

$$\begin{aligned}
S &= \sum_{k=\underline{k}}^{\bar{k}} \sum_{l=0}^{k} \tilde{R}_l' E_{(ii)} \tilde{R}_l \\
&= \sum_{l=0}^{\bar{k}} (\bar{k} + 1 - \max(\underline{k}, l)) \tilde{R}_l' E_{(ii)} \tilde{R}_l \\
&= \sum_{l=0}^{\bar{k}} (\bar{k} + 1 - \max(\underline{k}, l)) \tilde{R}_{l,i}' \tilde{R}_{l,i}
\end{aligned}$$

where $\tilde{R}_{l,i}$ denotes row i of $\tilde{R}_l$, i.e., the response of variable i. Note the difference to (6): the transpose is now on the *first*, not the second R_l.

The maximization problem subject to the side constraint $q_1' q_1 = 1$ can be written as a Lagrangian,

$$L(q_1, \theta) = q_1' S q_1 - \lambda(q_1' q_1 - 1)$$

with the first-order condition

$$Sq_1 = \lambda q_1 \tag{9}$$

and the side constraint. From this equation, we see that the solution q_1 is an eigenvector of S with eigenvalue λ. We also see that

$$\sigma^2(\underline{k}, \bar{k}; q_1) = \lambda$$

Thus, to maximize this variance, we need to find the eigenvector with the maximal eigenvalue λ, i.e., we need to find the first principal component. The impulse vector a is then found as

$$a = \tilde{A} q_1 \tag{10}$$

as before.

To include a non-core variable in the VAR, use q_1 as defined above, and append it with an extra 0 as its last entry,

$$\hat{q} = \begin{bmatrix} q_1 \\ 0 \end{bmatrix}$$

Let $\tilde{A}$ be the Cholesky decomposition for the core VAR, containing m variables, and let $\hat{A}$ be the Cholesky decomposition for the VAR, including the non-core variable as its last row. Note that

$$\hat{A} = \begin{bmatrix} \tilde{A} & 0_{m,1} \\ \hat{A}_{m+1,1:m} & \hat{A}_{m+1,m+1} \end{bmatrix}$$

where

$$\hat{A}_{m+1,1:m} = \begin{bmatrix} \hat{A}_{m+1,1} \dots \hat{A}_{m+1,m} \end{bmatrix}$$

The impulse vector $\hat{a}$ for the VAR containing the non-core variable is given by

$$\hat{a} = \hat{A} \hat{q} \tag{11}$$

It follows that

$$\hat{a} = \begin{bmatrix} a \\ \hat{A}_{m+1,1:m} q_1 \end{bmatrix}$$

A more intuitive way to put this is that the impulse response for the non-core variables is obtained by regressing the non-core variables not only on lagged but also on contemporaneous values of the core variables, and then substituting in the vector a for the contemporaneous values, when calculating the impulse response. This is, in fact, how the calculations have been done in practice.

This all can easily be generalized to p vectors: set Q_1 to be the matrix of the unit-length eigenvectors of S corresponding to the p largest eigenvalues. Equivalently, Q_1 is the matrix of the first p principal components, normalized to unit length. The impulse matrix is given by $A_1 = \tilde{A} Q_1$.

Note that if $\bar{k} - \underline{k} = p - 1$ or smaller, then

$$\sum_{j=1}^{p} \sigma^2(\underline{k}, \bar{k}; q_j) = \sigma^2(\underline{k}, \bar{k})$$

i.e., it is possible to explain all of this variance. To see this, note that the matrix S will be of no more than rank p in that case.

7 Conclusions

This paper has aimed at providing empirical insight into identifying the main forces moving real GNP. It did so by searching for one or two shocks, which explain the majority of the movements in GNP. More precisely, we have sought to explain as much as possible of the sum of the k-step ahead prediction error variances for GNP over a horizon of up to 5 years with two shocks or as much as possible of the sum over the first year or years three to five with just one shock. The methodology is explained in detail in the second half of the paper.

We have applied this method to a set of 21 macroeconomic variables, where we have used seven variables in a core VAR and added 14 variables as periphery. We found that two shocks are sufficient to explain around 90% of the k-step ahead prediction error for real GNP at all but the very shortest horizons between 0 and five years. A particularly appealing interpretation is to separate these two shocks into a medium-run shock and a short-run shock.

The medium-run shock explains on average 70% of the prediction error variance in GNP. Its impulse responses suggest a productivity shock interpretation much in the line of the real business cycle literature, although the details of the patterns make modifications to standard real business cycle theories necessary in order to account for the salient facts. For example, there is the rather persistent wage-productivity gap. The increase in GNP is gradual. Ten-year bond rates and short-term interest rates decline rather than rise. This shock can also identified well in isolation, when searching for a single shock explaining as much as possible of the fluctuations three to five years after the shock.

The short-run shock has the feature that GNP blips up shortly before turning negative, and is harder to interpret. It may be due to wage pressure coming from misjudged productivity developments, subsequent inflationary pressures, and attempts by the Federal Reserve to fight this inflation through raising interest rates. The short-run shock turns out to be difficult to isolate when searching for just a single shock, since the medium-run shock also contributes a considerable amount of variation even at shorter horizons.

In sum, while one shock looks like a productivity shock in the line of the real business cycle literature, the other one seems to be wage-push or inflationary shock, unrelated to consumption or government spending and not induced by monetary policy. While the first shock can be viewed as a "supply shock", the second shock does not have an obvious "demand shock" interpretation.

References

[1] No references yet except for one "placeholder"! To be completed in the next draft.

[2] Uhlig, Harald (1999), "What are the effects of monetary policy on output? Results from an agnostic identification procedure," draft, Tilburg University.

Tables and Figures

Real GNP	0.50%
lab.productivity	0.49%
wages minus prod	0.51%
Goverm. Spending	0.65%
Fed. Funds Rate	0.63%
Inflation (CPI)	0.35%
Oil Price	9.27%
real nondur.cons	0.29%
Priv. Investment	0.79%
Hours worked	0.17%
Ind.Production	0.51%
Capacity Utiliz.	0.44%
S&P 500	3.63%
10yr bond rate	0.30%
Inflation (PPI)	1.08%
Nonborr.Reserves	2.10%
M1	0.53%
M3	0.28%
Exchange Rate	1.32%
Exports	1.30%
Imports	1.23%

Table 1: *Standard deviations of the MLE one-step ahead prediction errors. For the core variables, Real GNP up to Oil Price, the one-step ahead prediction error is calculated, using a VAR in the lagged values of the core variables as regressors. For the other, non-core variables, the one-step ahead prediction error is calculated, using lagged and contemporaneous core variables and own lags.*

	two shocks (> 50% of var.)					$\theta = 0$, (> 50% of var.)				
Year after shock:	0	1	2	3	4	0	1	2	3	4
Variable:										
Real GNP	o111	1111	1111	1111	1111	ooo1	1111	1111	1111	1111
lab.productivity	o111	1111	1111	1111	1111	oo11	1111	1111	1111	1111
wages div. by prod.	oo11	1111	1ooo	1111	11oo	oo11	1ooo	oooo	oooo	oooo
Goverm. Spending	oooo	ooo1	1111	1111	1111	oooo	oooo	oooo	oooo	oooo
Fed. Funds Rate	oooo	1111	1111	1ooo	oooo	oooo	oooo	oooo	oooo	oooo
Inflation (CPI)	o111	1111	1111	oooo	ooo1	oooo	oooo	oooo	oooo	oooo
Oil Price	oooo	oooo	1111	11oo	oooo	oooo	oooo	oooo	oooo	oooo
real nondur.cons	o111	1111	1111	1111	1111	oo11	1111	1111	1111	1111
Priv. Investment	o111	1111	1111	1111	1111	oooo	1111	1111	1111	1111
Hours worked	oo11	111o	1111	1111	1111	oooo	oooo	oo11	1111	1111
Ind.Production	oo11	1111	1111	1111	1111	oooo	oooo	o111	1111	1111
Capacity Utiliz.	oo11	1ooo	oo11	1111	1111	oooo	oooo	oooo	oooo	oooo
S&P 500	o1o1	1111	1111	111o	1ooo	ooo1	1111	oo11	11oo	oooo
1oyr bond rate	oooo	oooo	1111	1111	oooo	oooo	oooo	oooo	oooo	oooo
Inflation (PPI)	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo
Nonborr.Reserves	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo
M1	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo
M3	oooo	oooo	1111	1111	1111	oooo	oooo	oooo	1111	1111
Exchange Rate	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo
Exports	oooo	oooo	oooo	oo11	1111	oooo	oooo	oooo	ooo1	1111
Imports	oo11	1111	1111	1111	1111	ooo1	1111	1111	1111	1111

Table 2: *Two-shock analysis: shown are the horizons (quarters after the shock, grouped into years), when the two shocks moving real GNP also together explain more than 50% of the k-step-ahead prediction variance of the variable indicated and when the "medium-run shock" ($\theta = 0$) alone explains more than 50%, using the median of the posterior for this statistic.*

	$\theta = 0$, (> 20% of var.)					$\theta = 90$, (> 20% of var.)				
Year after shock:	0	1	2	3	4	0	1	2	3	4
Variable:										
Real GNP	oo11	1111	1111	1111	1111	1111	111o	oooo	oooo	oooo
lab.productivity	1111	1111	1111	1111	1111	oooo	oooo	oooo	oooo	oooo
wages div. by prod.	1111	1111	1111	1111	1ooo	oooo	oooo	oooo	oooo	oooo
Goverm. Spending	ooo1	1oo1	1111	1111	1111	oooo	ooo1	oooo	oooo	oooo
Fed. Funds Rate	oooo	o11o	1111	1ooo	oooo	oo11	1111	1o1o	oooo	oooo
Inflation (CPI)	1111	1111	1111	1ooo	oooo	oooo	o111	11oo	oooo	oooo
Oil Price	oooo	oooo	o111	1ooo	oooo	oooo	oo11	11oo	oooo	oooo
real nondur.cons	o111	1111	1111	1111	1111	oooo	oooo	oooo	oooo	oooo
Priv. Investment	oo11	1111	1111	1111	1111	o111	11oo	oooo	oooo	oooo
Hours worked	ooo1	1111	1111	1111	1111	o111	111o	oooo	ooo1	111o
Ind.Production	ooo1	1111	1111	1111	1111	o111	11oo	oooo	oooo	oooo
Capacity Utiliz.	oooo	oooo	oooo	1111	11oo	o111	11oo	ooo1	1111	1111
S&P 500	o111	1111	1111	1111	1111	oooo	oooo	oooo	oooo	oooo
1oyr bond rate	oooo	ooo1	1111	1111	oooo	oooo	oooo	ooo1	11oo	oooo
Inflation (PPI)	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo
Nonborr.Reserves	oooo	oooo	oooo	oooo	oooo	oo11	1oo1	oooo	oooo	oooo
M1	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo	oooo
M3	ooo1	1111	1111	1111	1111	oooo	oooo	oooo	oooo	oooo
Exchange Rate	oooo	1111	11oo	oooo	oooo	oooo	oooo	oooo	oooo	oooo
Exports	oooo	oooo	oooo	o111	1111	oo11	o111	o1oo	oooo	oooo
Imports	oo11	1111	1111	1111	1111	o11o	oooo	oooo	oooo	oooo

Table 3: *Two shock analysis: shown are the horizons (quarters after the shock, grouped into years), when the "medium-run shock" ($\theta = 0$) and the "short-run shock" ($\theta = 90$) each explain more than 20% of the k-step-ahead prediction variance of the variable indicated, using the median estimate.*

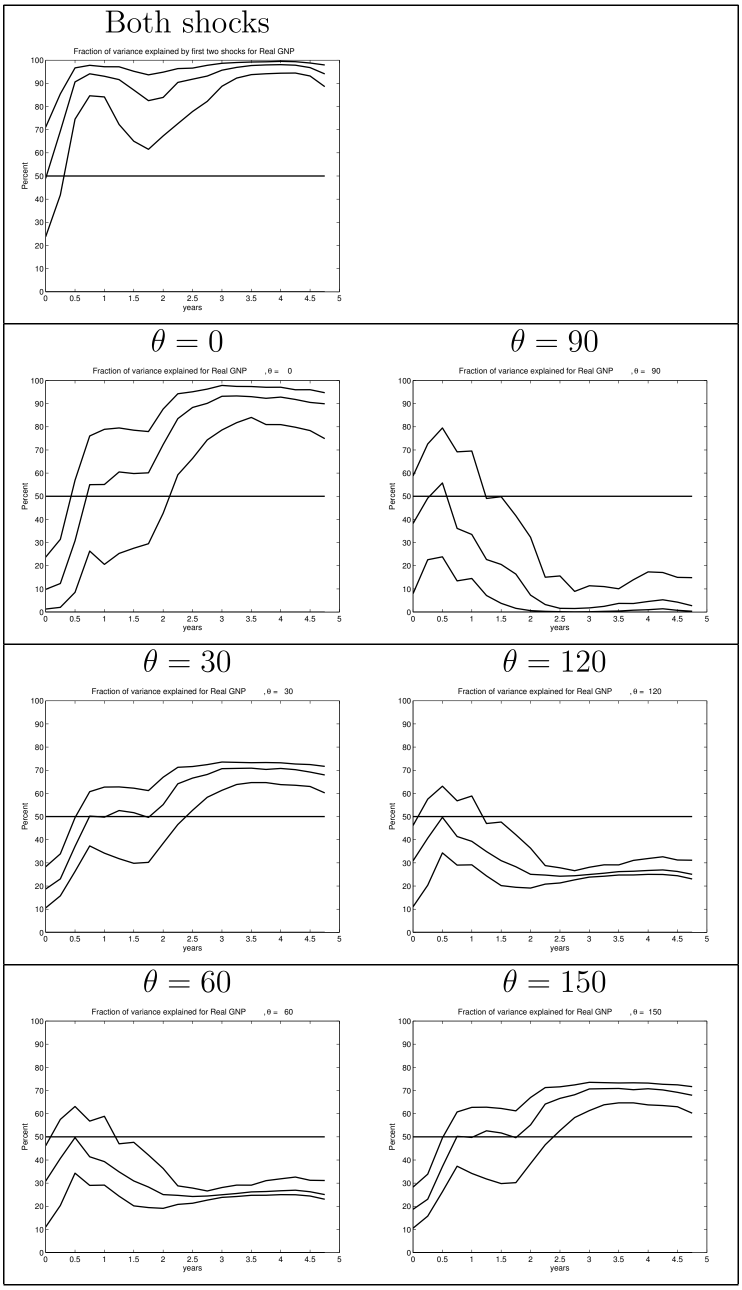

Figure 1 *Variance decompositions for output, using three orthogonal pairings. Horizontal scale: 0 to 5 years after the shock. Lines are 16 percent quartile, median, 84 percent quartile. The horizontal line is at 50% for visual comparison.*

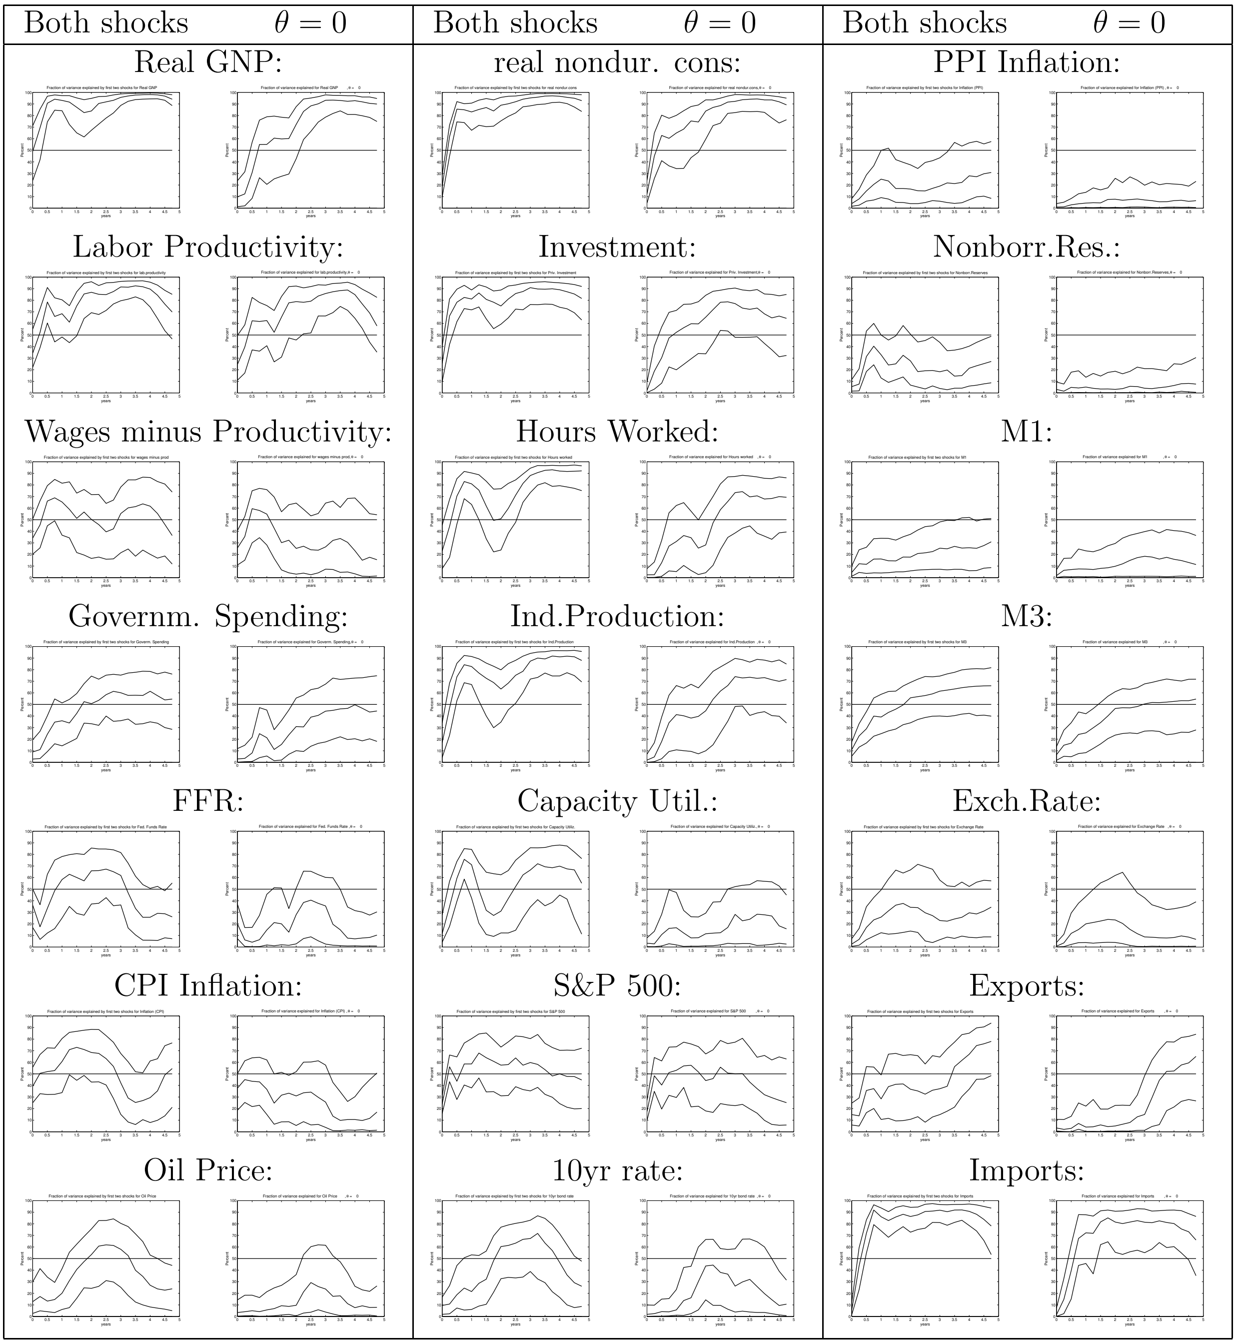

Figure 2 *Variance decompositions for both shocks together as well as the the $\theta = 0$ shock. The left column contains the core variables, the other two columns the periphery variables. Horizontal scale: 0 to 5 years after the shock. Lines are 16 percent quartile, median, 84 percent quartile. The horizontal line is at 50% for visual comparison.*

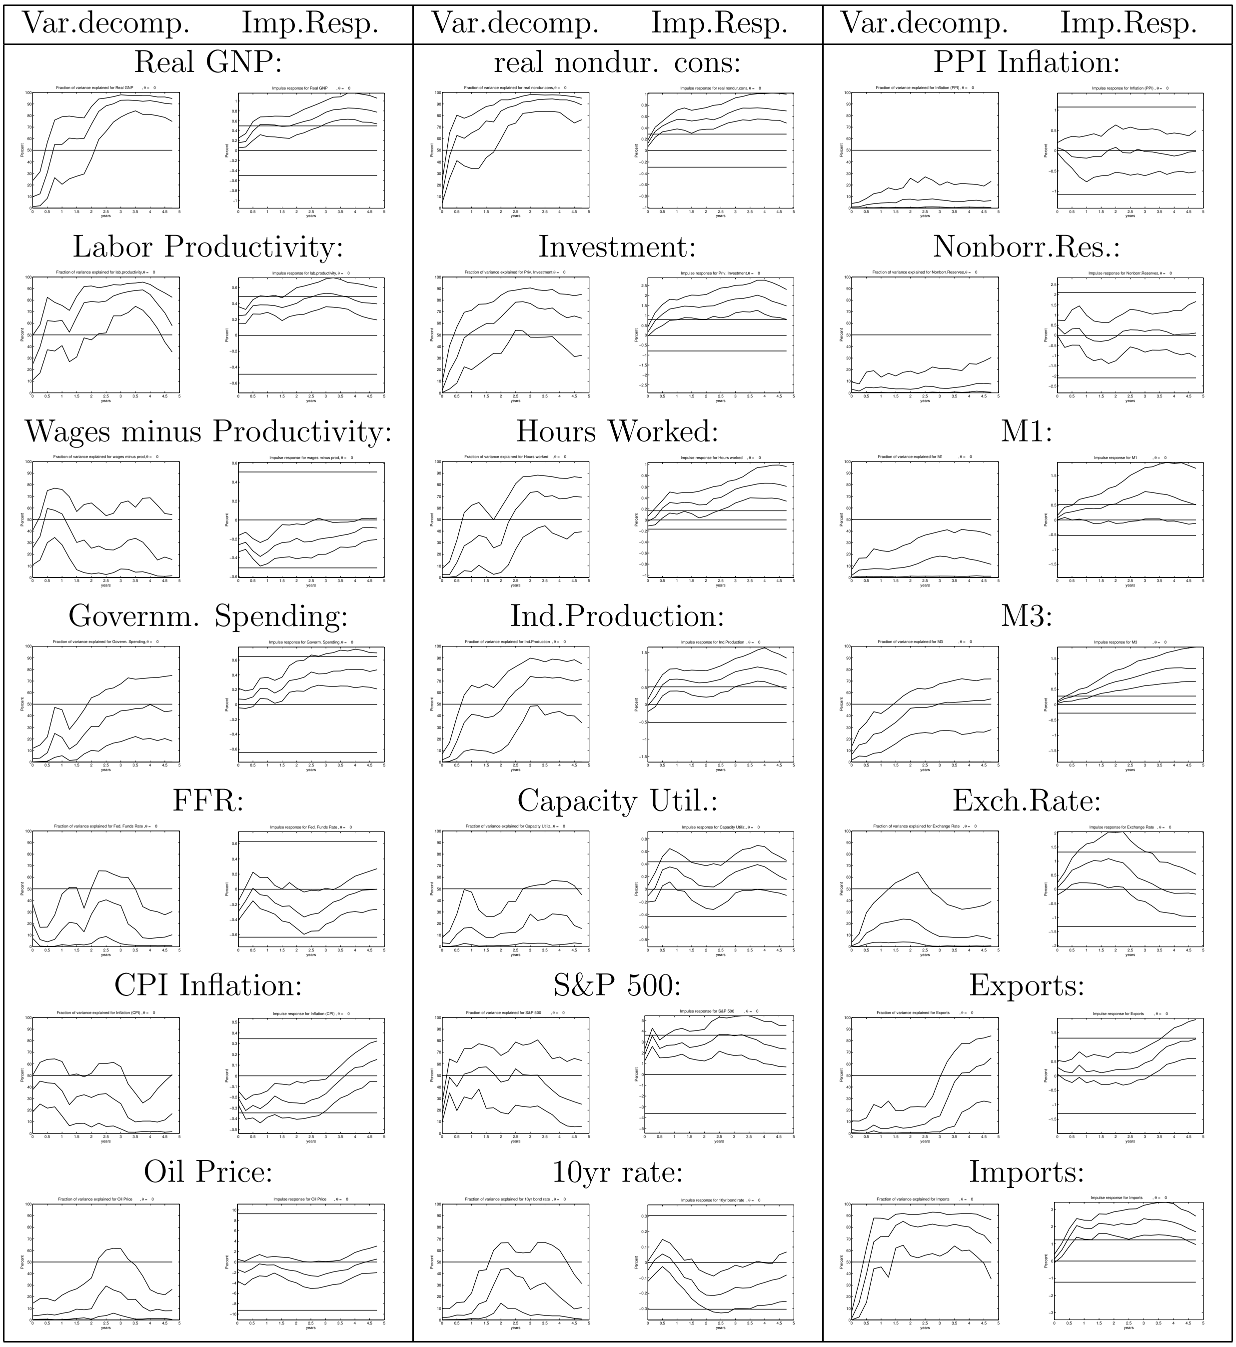

Figure 3 *Results for the medium-run shock, $\theta = 0$. Shown are variance decompositions and impulse responses, a pair for each variable. The left column contains the core variables, the other two columns the periphery variables. Horizontal scale: 0 to 5 years after the shock. Lines are 16 percent quartile, median, 84 percent quartile. The horizontal line for the variance decompositions is at 50% for visual comparison. The straight lines for the impulse responses are the zero response plus minus the standard deviation of the VAR-MLE one-step ahead prediction errors for comparison.*

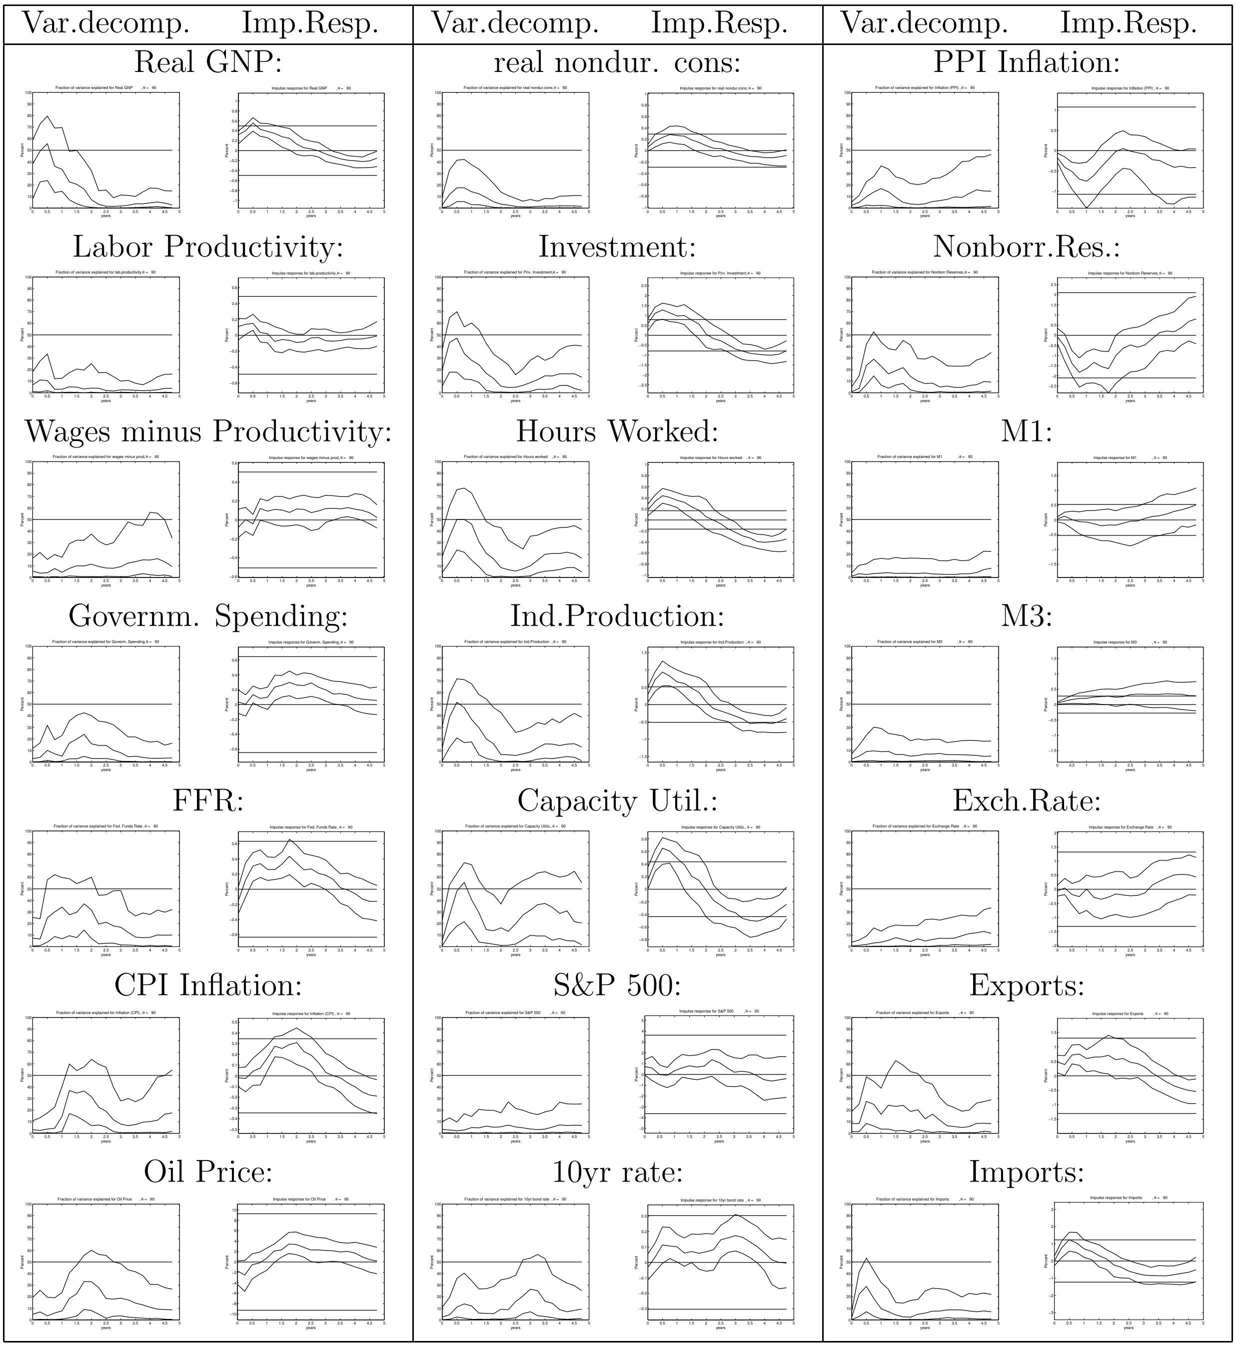

Figure 4 *Results for the short-run shock, $\theta = 90$. Shown are variance decompositions and impulse responses, a pair for each variable. The left column contains the core variables, the other two columns the periphery variables. Horizontal scale: 0 to 5 years after the shock. Lines are 16 percent quartile, median, 84 percent quartile. The horizontal line for the variance decompositions is at 50% for visual comparison. The straight lines for the impulse responses are the zero response plus minus the standard deviation of the VAR-MLE one-step ahead prediction errors for comparison.*

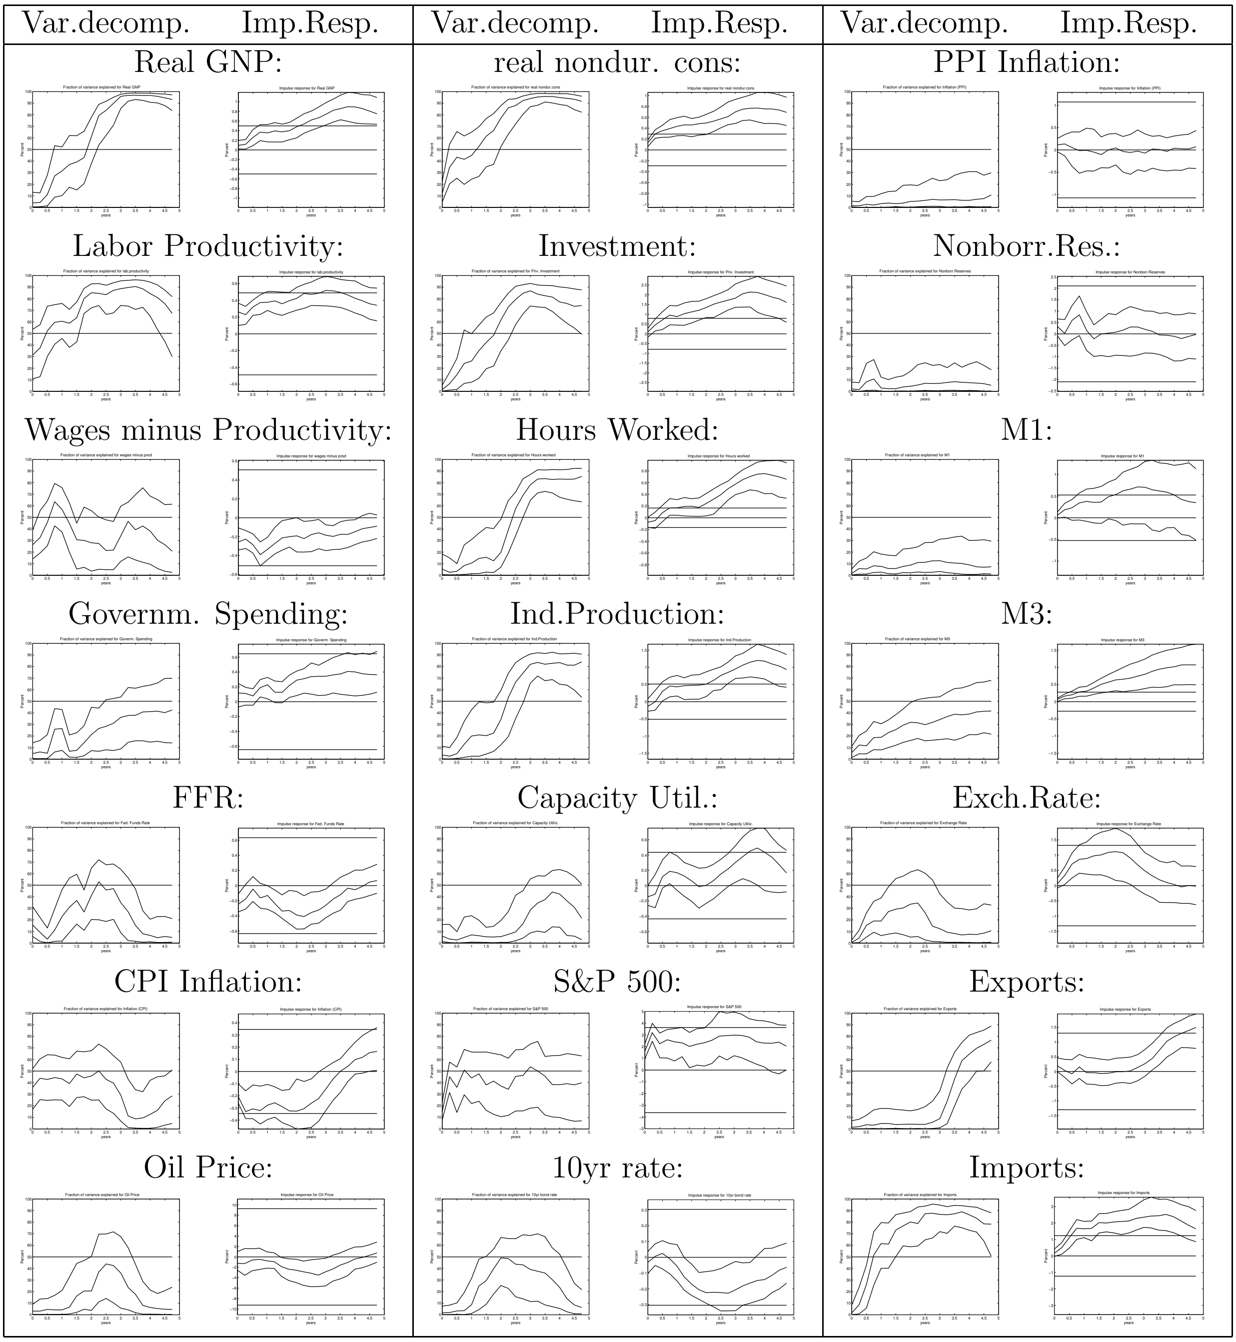

Figure 5 *One principal component, maximizing the explained variance of real GNP three to five years after the shock. Shown are variance decompositions and impulse responses, a pair for each variable. The left column contains the core variables, the other two columns the periphery variables. Horizontal scale: 0 to 5 years after the shock. Lines are 16 percent quartile, median, 84 percent quartile. The horizontal line for the variance decompositions is at 50% for visual comparison. The straight lines for the impulse responses are the zero response plus minus the standard deviation of the VAR-MLE one-step ahead prediction errors for comparison.*

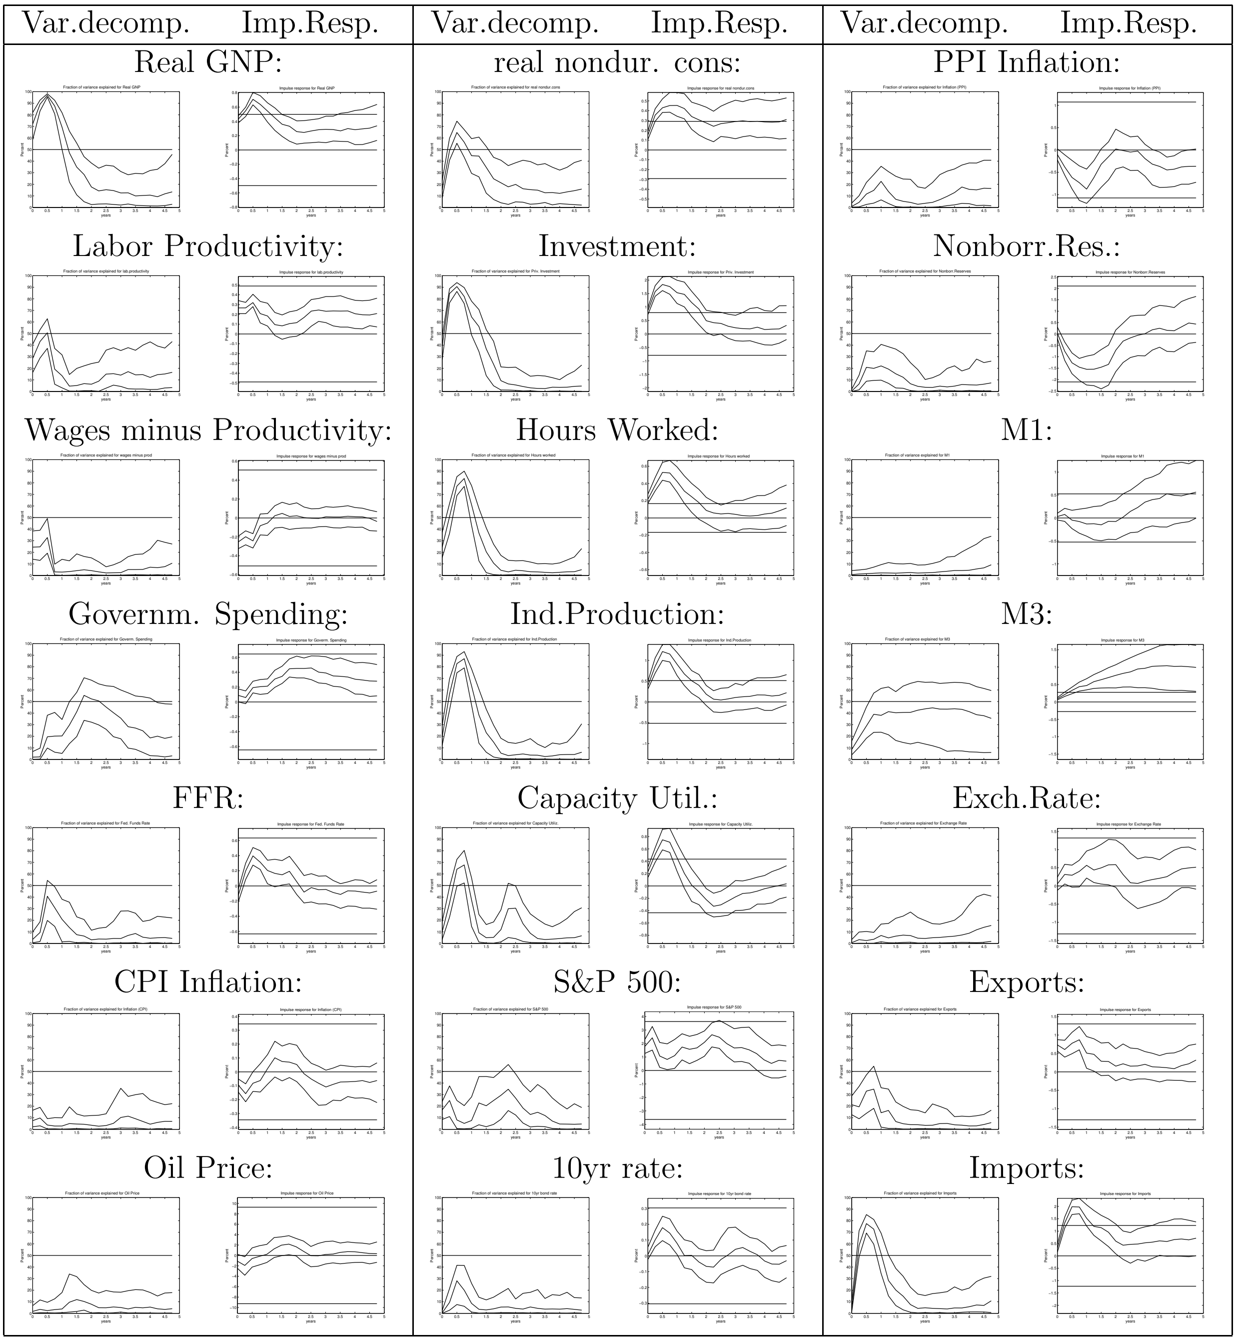

Figure 6 *One principal component, maximizing the explained variance of real GNP during the first year after the shock. Shown are variance decompositions and impulse responses, a pair for each variable. The left column contains the core variables, the other two columns the periphery variables. Horizontal scale: 0 to 5 years after the shock. Lines are 16 percent quartile, median, 84 percent quartile. The horizontal line for the variance decompositions is at 50% for visual comparison. The straight lines for the impulse responses are the zero response plus minus the standard deviation of the VAR-MLE one-step ahead prediction errors for comparison.*

Technical Appendix

A Estimation and Inference

For convenience, we collect here the main tools for estimation and inference. Stack the core VAR system (2) as

$$\mathbf{Y} = \mathbf{X}\mathbf{B} + \mathbf{u} \tag{12}$$

where $X_t = [1, Y'_{t-1}, Y'_{t-2}, \ldots, Y'_{t-l}]'$, $\mathbf{Y} = [Y_1, \ldots, Y_T]'$, $\mathbf{X} = [X_1, \ldots, X_T]'$, $\mathbf{u} = [u_1, \ldots, u_T]'$ and $\mathbf{B} = [B_{(1)}, \ldots, B_{(l)}]'$. To compute the impulse response to an impulse vector a, let $\mathbf{a} = [a', 0_{1,m(l-1)}]'$ as well as

$$\Gamma = \begin{bmatrix} \multicolumn{2}{c}{\mathbf{B}} \\ I_{m(l-1)} & 0_{m(l-1),m} \end{bmatrix}$$

and compute $r_{k,j} = (\Gamma^k \mathbf{a})_j, k = 0, 1, 2, \ldots$ to get the response of variable j at horizon k. The variance of the k-step ahead forecast error due to an impulse vector a is obtained by simply squaring its impulse responses. Summing again over all a_j, where a_j is the j-th column of some matrix A with $AA' = \Sigma$ delivers the total variance of the k-step ahead forecast error.

We assume that the u_t's are independent and normally distributed. The MLE for $(\mathbf{B}, \Sigma)$ is given by

$$\hat{B} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y},\ \hat{\Sigma} = \frac{1}{T}(\mathbf{Y} - \mathbf{X}\hat{B})'(\mathbf{Y} - \mathbf{X}\hat{B}) \tag{13}$$

Our prior and posterior for $(\mathbf{B}, \Sigma)$ belongs to the

Normal-Wishart family, whose properties are further discussed in Uhlig (1994). A proper Normal-Wishart distribution is parameterized by a "mean coefficient" matrix $\bar{B}$ of size $k \times m$, a positive definite "mean covariance" matrix S of size $m \times m$ as well as a positive definite matrix N of size $l \times l$ and a "degrees of freedom" real number $\nu \geq 0$ to describe the uncertainty about $(\mathbf{B}, \Sigma)$ around $(\bar{B}, S)$. The Normal-Wishart distribution specifies, that Σ^{-1} follows a Wishart distribution[2] $\mathcal{W}_m(S^{-1}/\nu, \nu)$ with $E[\Sigma^{-1}] = S^{-1}$, and that, conditionally on Σ, the coefficient matrix in its columnwise vectorized form, $\mathrm{vec}(\mathbf{B})$ follows a Normal distribution $\mathcal{N}(\mathrm{vec}(\bar{B}), \Sigma \otimes N^{-1})$.

Proposition 1 on p. 670 in Uhlig (1994) states, that if the prior is described by $\bar{B}_0$, N_0, S_0 and ν_0, then the posterior is desribed by $\bar{B}_T$, N_T, S_T and ν_T, where

$$\begin{aligned}
\nu_T &= T + \nu_0 \\
N_T &= N_0 + \mathbf{X}'\mathbf{X} \\
\bar{B}_T &= N_T^{-1}(N_0\bar{B}_0 + \mathbf{X}'\mathbf{X}\hat{B}) \\
S_T &= \frac{\nu_0}{\nu_T}S_0 + \frac{T}{\nu_T}\hat{\Sigma} + \frac{1}{\nu_T}(\hat{B} - \bar{B}_0)'N_0 N_T^{-1}\mathbf{X}'\mathbf{X}(\hat{B} - \bar{B}_0)
\end{aligned}$$

[2]To draw from this distribution, use e.g. $\Sigma = (R * R')^{-1}$, where R is a $m \times \nu$ matrix with each column an independent draw from a Normal distribution $\mathcal{N}(0, S^{-1}/\nu)$ with mean zero and variance-covariance matrix S^{-1}.

We use a "weak" prior, and use $N_0 = 0$, $\nu_0 = 0$, S_0 and $\bar{B}_0$ arbitrary. Then, $\bar{B}_T = \hat{B}$, $S_T = \hat{\Sigma}$, $\nu_T = T$, $N_T = \mathbf{X}'\mathbf{X}$, which is also the form of the posterior used in the RATS manual for drawing error bands, see example 10.1 in Doan (1992).

For the non-core variables, one simply needs to estimate a regression of these variables on their own lags and contemporaneous and lagged values of the core variables. This is a standard regression exercise in Bayesian econometrics. Assuming the prediction error to be following an inverse gamma distribution (which is the same as a one-dimensional inverse-Wishart distribution), one can essentially follow the algebra above. In sum, first take a draw from the posterior of the core VAR. Conditionally on this draw, the coefficients for the non-core variables are Normal-gamma distributed, including the coefficients on the contemporaneous core variables, using formulas which are just appropriately rewritten versions of the formulas above.

No attempt has been made to impose more specific prior knowledge such as the "no change forecast" of the Minnesota prior, see Doan, Litterman and Sims (1984), special treatments of roots near unity, see the discussion in Sims and Uhlig (1991) as well as Uhlig (1994), or to impose the more sophisticated priors of Leeper, Sims and Zha (1996) or Sims and Zha (1996). Also, we have not experimented with regime-switching as in Bernanke and Mihov (1996a,b) or with stochastic volatility as in Uhlig (1997).

To draw inferences, take n draws from the VAR posterior and, for each of these draws find the impulse vector a according to one of the methods stated in the body of the paper. Calculate the impulse responses as well as the fraction of variances attributable to a, see the appendix to Uhlig (1999) for further details. Finally, error bands for e.g. the k-step ahead impulse response are calculated by sorting the results from all draws for the k-step ahead impulse response, etc..

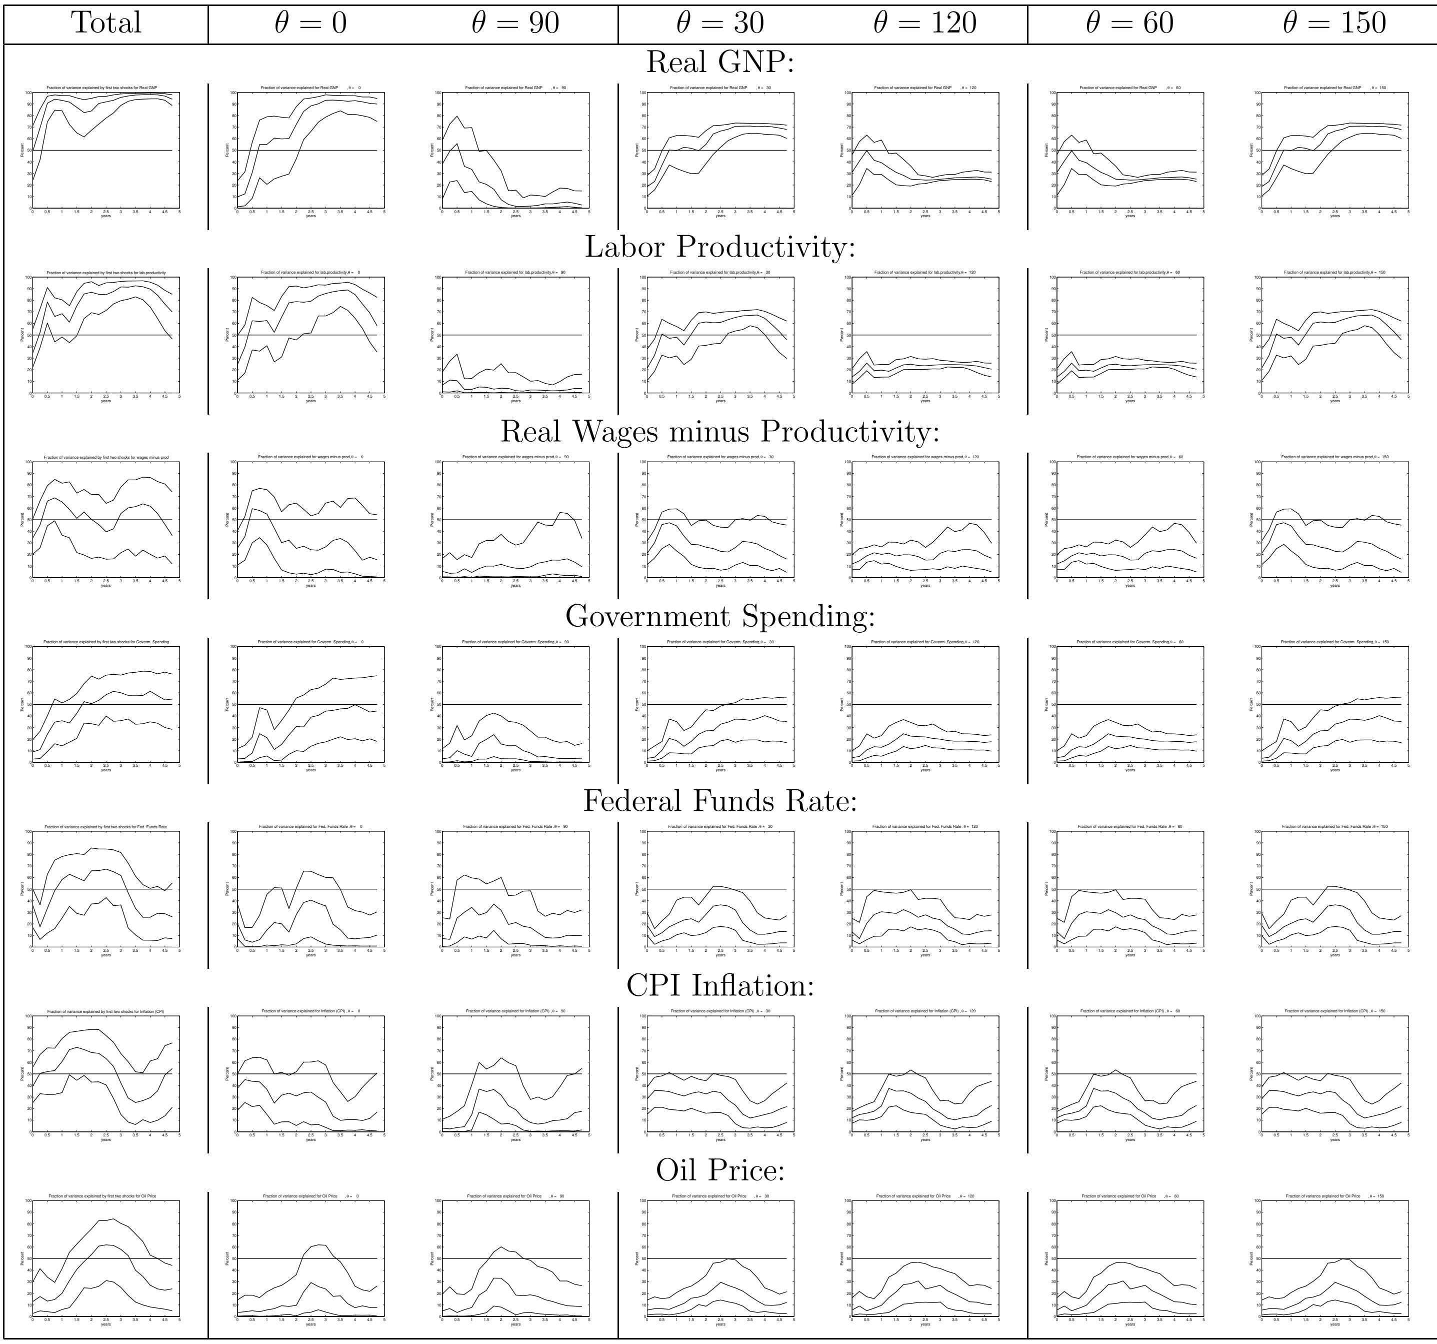

Figure 7 *Variance explained, part 1: core variables. First column: total. Other columns: for the pairings ($\theta = 0$, $\theta = 90$), ($\theta = 30$, $\theta = 120$) or ($\theta = 60$, $\theta = 150$). Horizontal scale: 0 to 5 years after the shock. Lines are 16 percent quartile, median, 84 percent quartile. The horizontal line for the variance decompositions is at 50% for visual comparison.*

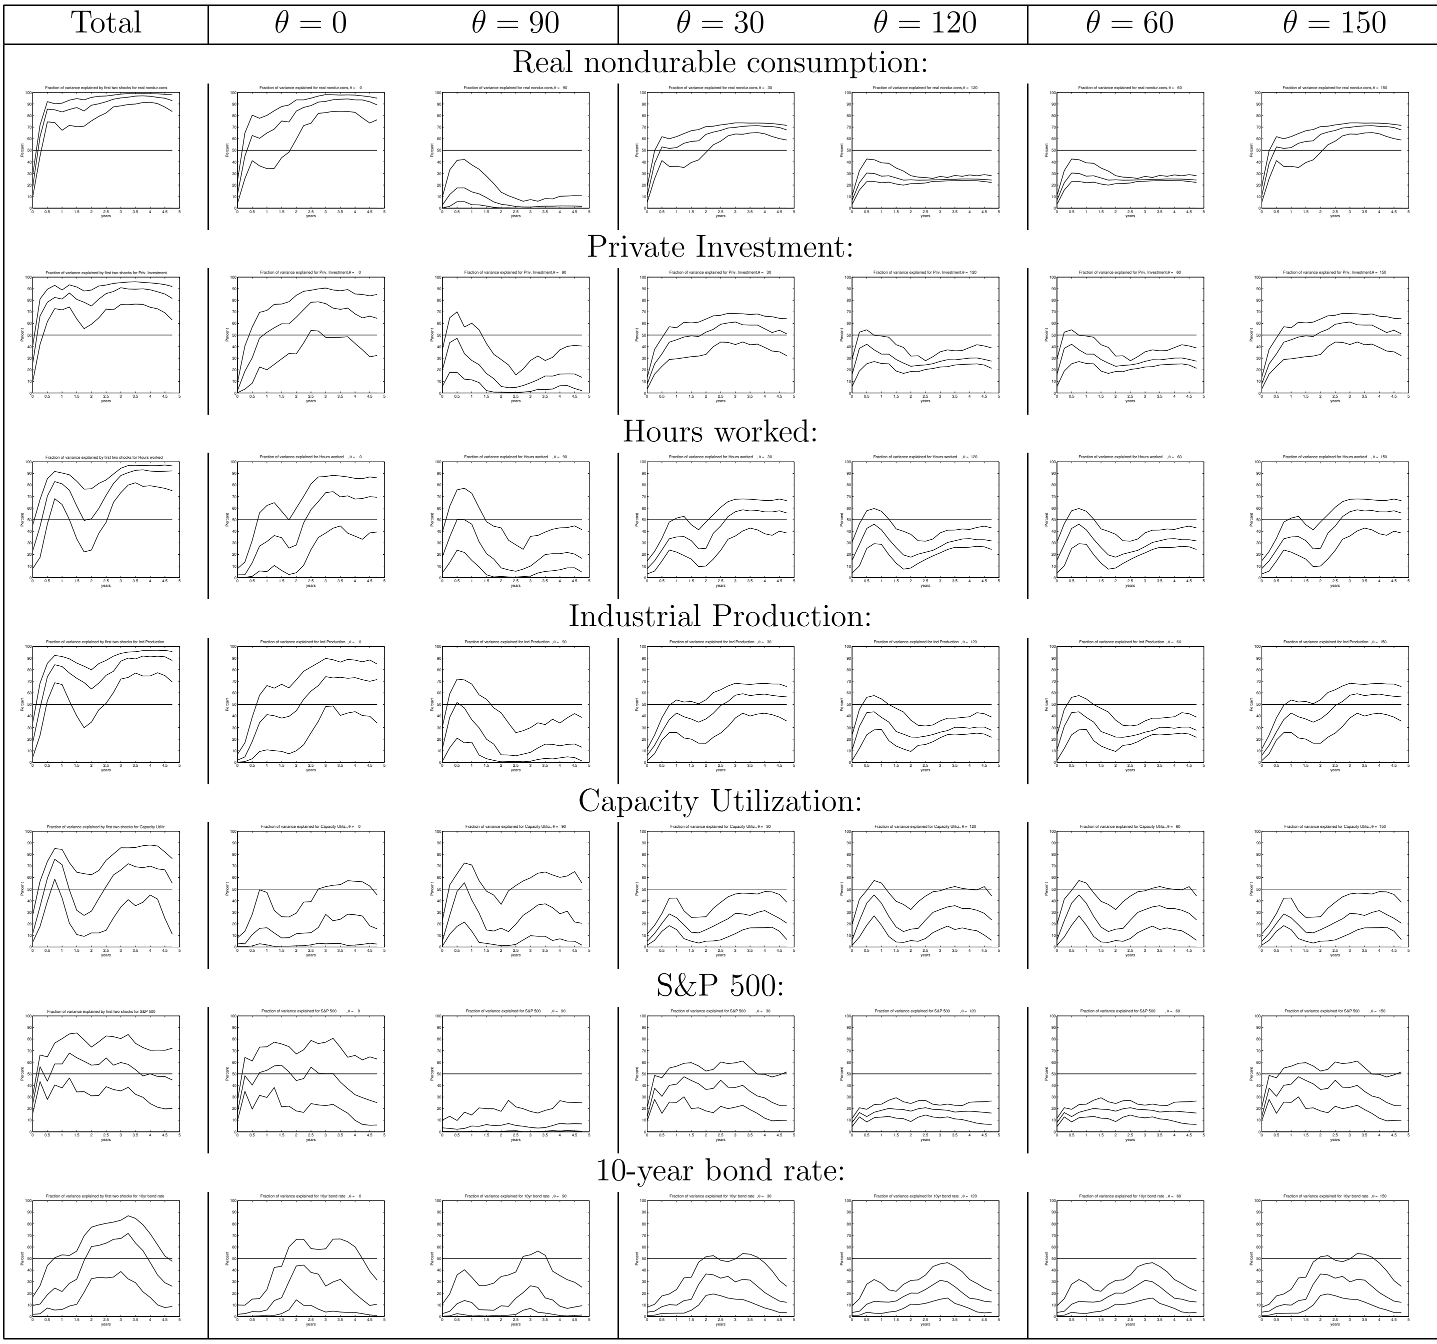

Figure 8 *Variance explained, part 2: periphery variables. First column: total. Other columns: for the pairings ($\theta = 0$, $\theta = 90$), ($\theta = 30$, $\theta = 120$) or ($\theta = 60$, $\theta = 150$). Horizontal scale: 0 to 5 years after the shock. Lines are 16 percent quartile, median, 84 percent quartile. The horizontal line for the variance decompositions is at 50% for visual comparison.*

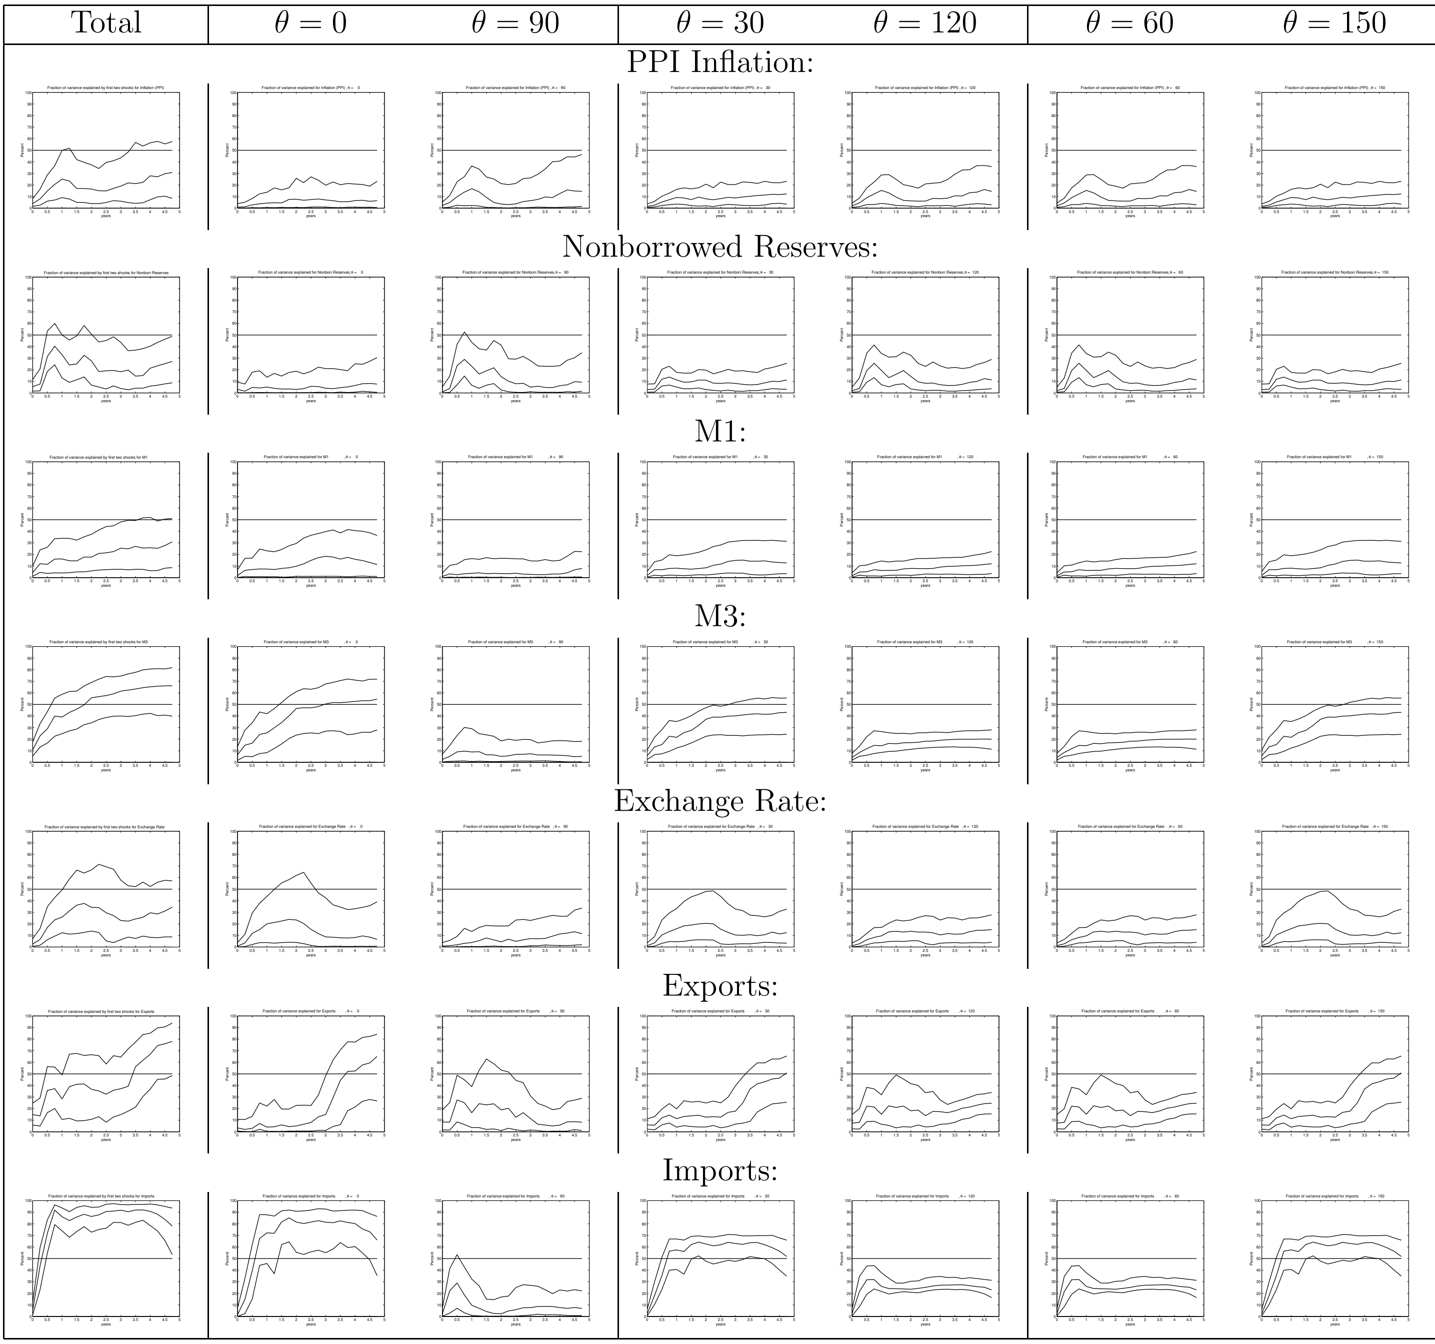

Figure 9 *Variance explained, part 3: periphery variables. First column: total. Other columns: for the pairings ($\theta = 0$, $\theta = 90$), ($\theta = 30$, $\theta = 120$) or ($\theta = 60$, $\theta = 150$). Horizontal scale: 0 to 5 years after the shock. Lines are 16 percent quartile, median, 84 percent quartile. The horizontal line for the variance decompositions is at 50% for visual comparison. The horizontal line for the variance decompositions is at 50% for visual comparison.*

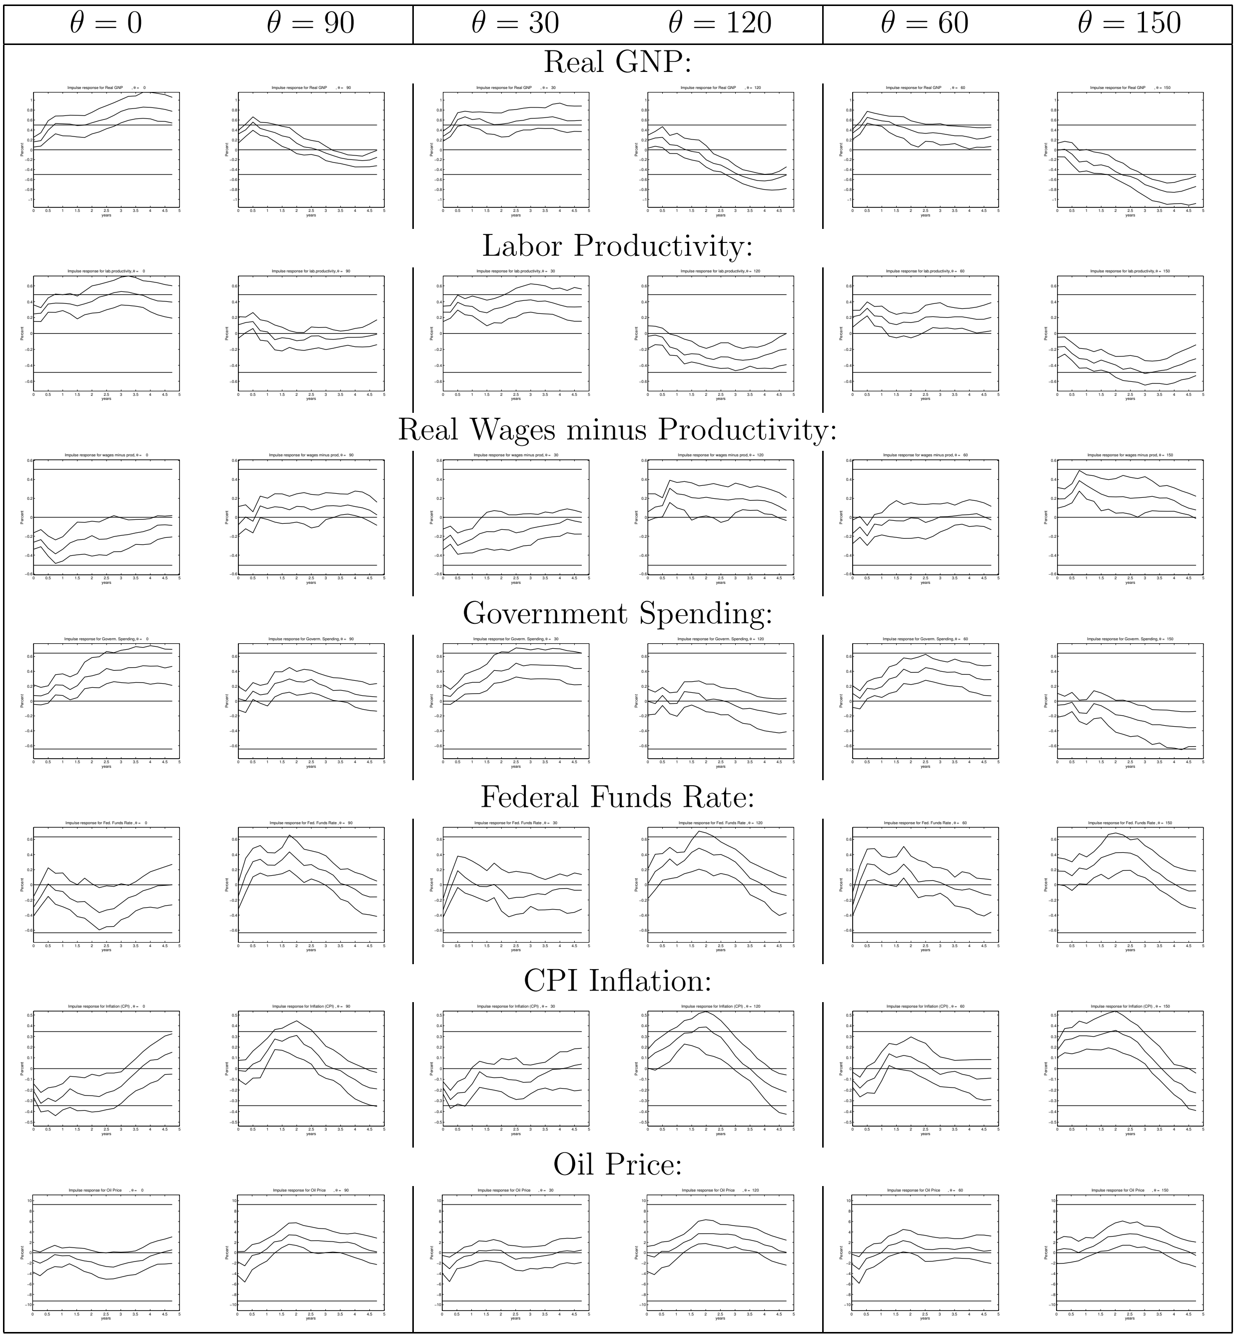

Figure 10 *Impulse responses, part 1: core variables. The columns are pairings ($\theta = 0$, $\theta = 90$), ($\theta = 30$, $\theta = 120$) or ($\theta = 60$, $\theta = 150$). Columns 1 and 2 provide a pair of orthogonal shocks, likewise columns 3 and 4 and columns 5 and 6. Horizontal scale: 0 to 5 years after the shock. Lines are 16 percent quartile, median, 84 percent quartile. The straight lines are the zero response plus minus the standard deviation of the VAR-MLE one-step ahead prediction errors for comparison.*

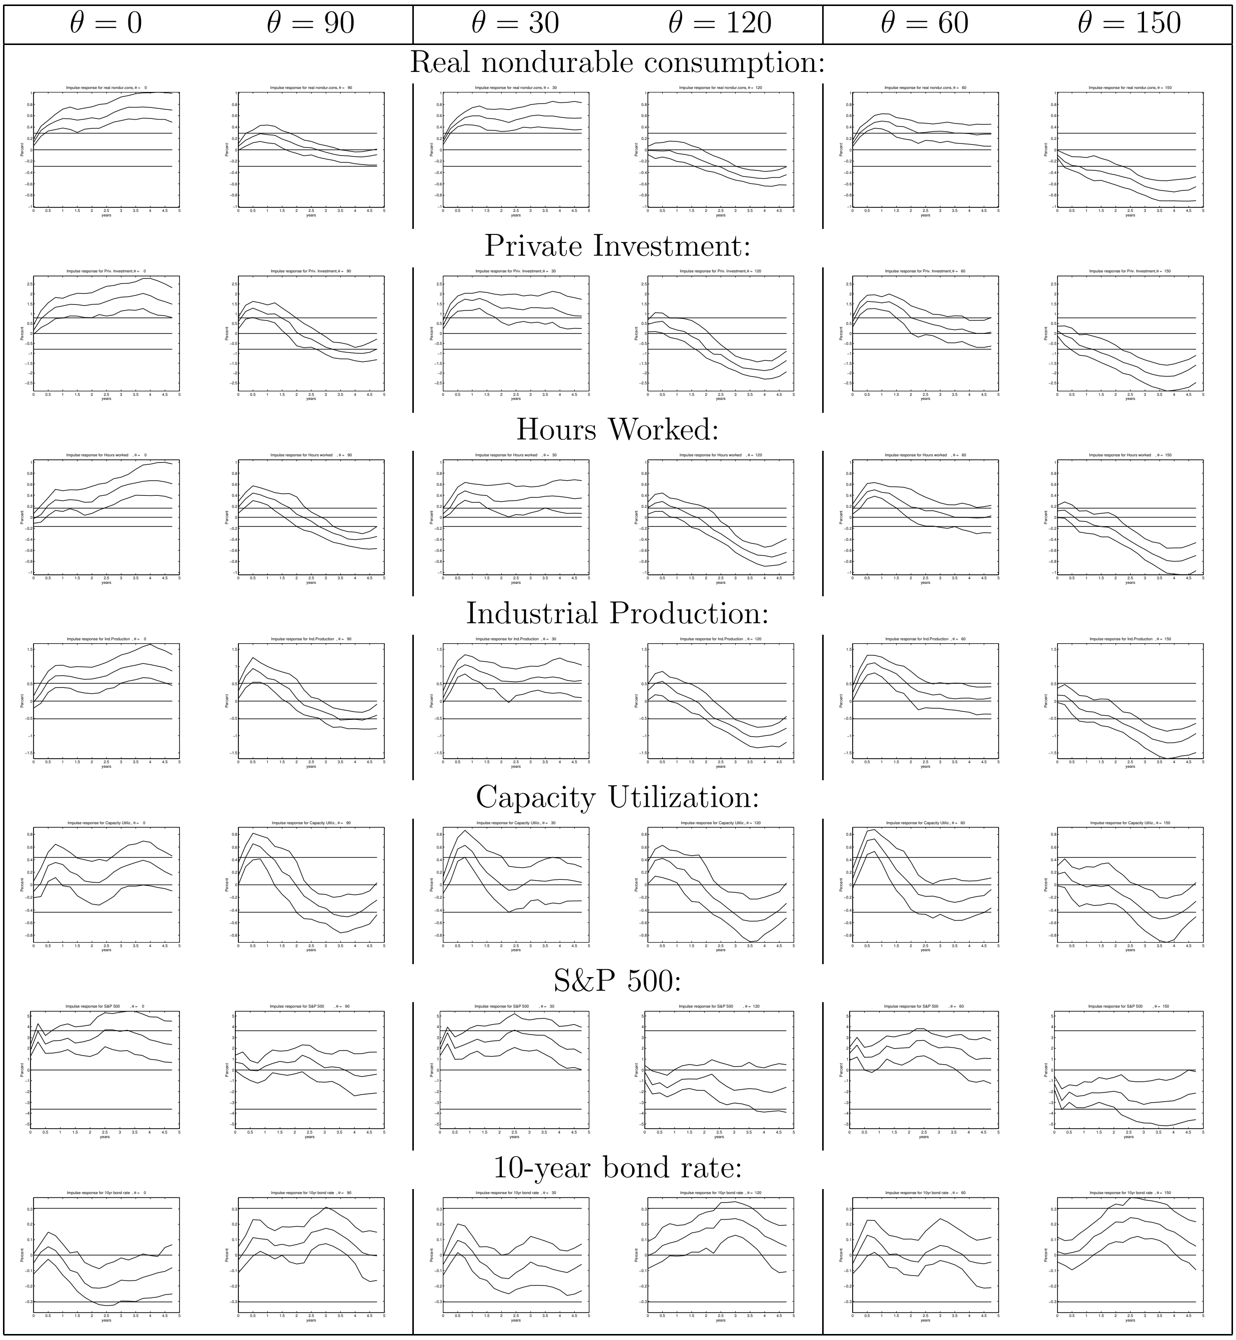

Figure 11 *Impulse responses, part 2: periphery variables. The columns are pairings ($\theta = 0$, $\theta = 90$), ($\theta = 30$, $\theta = 120$) or ($\theta = 60$, $\theta = 150$). Columns 1 and 2 provide a pair of orthogonal shocks, likewise columns 3 and 4 and columns 5 and 6. Horizontal scale: 0 to 5 years after the shock. Lines are 16 percent quartile, median, 84 percent quartile. The straight lines are the zero response plus minus the standard deviation of the VAR-MLE one-step ahead prediction errors for comparison.*

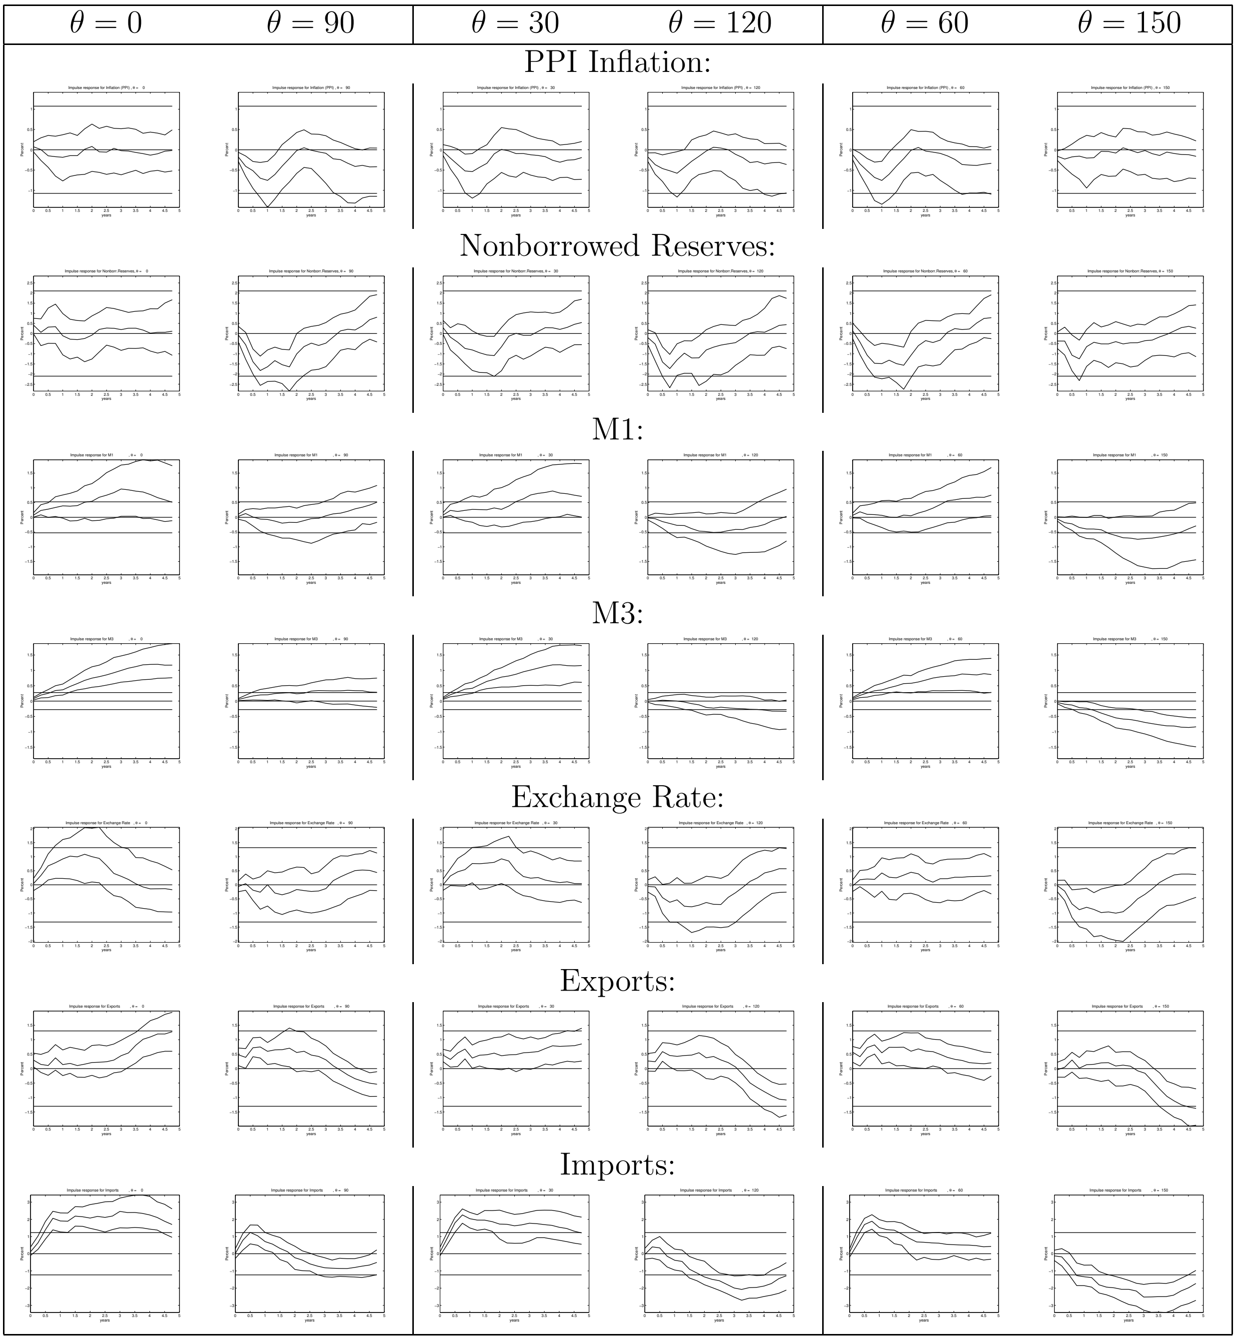

Figure 12 *Impulse responses, part 3: periphery variables. The columns are pairings ($\theta = 0$, $\theta = 90$), ($\theta = 30$, $\theta = 120$) or ($\theta = 60$, $\theta = 150$). Columns 1 and 2 provide a pair of orthogonal shocks, likewise columns 3 and 4 and columns 5 and 6. Horizontal scale: 0 to 5 years after the shock. Lines are 16 percent quartile, median, 84 percent quartile. The straight lines are the zero response plus minus the standard deviation of the VAR-MLE one-step ahead prediction errors for comparison.*