Distilling the main driver of business cycles

Gideon Magnus

Senior Quantitative Analyst
Morningstar, Inc.

Preliminary — not to be quoted

July 15, 2016

Corresponding author:
Gideon Magnus
Morningstar, Inc.
22 W Washington Street
Chicago, IL 60602
USA
Phone: 312-696-6454
e-mail: gideonmagnus@gmail.com

Abstract This paper presents a new approach for empirically distilling the main driver of business cycles. We estimate a dynamic factor model with a large collection of U.S. macroeconomic and financial indicators. We then identify a factor that concentrates predictive power for real per capita GDP and study the dynamic effects of a typical shock to this factor, thereby tracing out a typical business cycle. The cycle lasts about three years and is characterized by increases in many measures of economic activity, including output, consumption, productivity, investment, and profits. Utilization of labor and capital increases, and we find that people are more willing to take on investment risk, bidding up the prices of risky assets, and accepting lower expected excess returns. Our results lend tentative support to the idea that productivity shocks are the main driver of business cycle fluctuations.

Keywords Business cycles · News · Dynamic factor models

JEL Classification E3 · C32 · D8 · E44

1 Introduction

What are the main forces that drive fluctuations in economic output? Finding the answer to this fundamental question is not easy (Cochrane, 1994; Ramey, 2015). Identifying a causal relationship requires a source of plausibly exogenous variation. But few (if any) observable macroeconomic variables meet this qualification. The predominant approach has therefore been to assume that exogenous (i.e. unexpected) events correspond with statistical innovations, called shocks. The expected effects of a shock reflect the information conveyed by that shock. For example, one may find that when government spending is unexpectedly high, subsequent unemployment is expected to be low. This does not necessarily mean that the low unemployment is *caused* by the higher spending of the government. It may well be that, for example, higher than average spending today is predictive of lower than average spending in the future. In order to identify causal relationships one needs additional structure, preferably formed by economic theory.

A prime example of a theoretical restriction is the fact that, for a broad range of models, technology is the only variable that can have a permanent (i.e. long-run) effect on labor productivity. Galí (1999) provides an example of the application of this restriction and he concludes that technology cannot be a key driver of business cycle fluctuations, since labor hours worked (under his identification method) tend to respond negatively to a positive technology shock. This is an important result, as it challenges the influential belief that technological fluctuations play a key part in business cycle fluctuations (Kydland and Prescott, 1982; Long and Plosser, 1983; Cooley and Prescott, 1995; King and Rebelo, 1999). Galí's ideas did not go unchallenged; see e.g. Uhlig (2004a), Fisher (2004), Christiano, Eichenbaum, and Evans (2005), Francis and Ramey (2005, 2006), Michelacci and Lopez-Salido (2007), and Francis et al. (2014).

Rather than starting with a particular theory one may also work the other way around. This idea was advocated by Uhlig (2004b), and it is the approach adopted in the current paper. Thus, we first distill statistically a factor with a strong dynamic effect on output, and then we examine to which degree the effects of this factor are consistent with various theories.

Since we are trying to understand dynamic patterns, we start from the idea that a key driver of output should have *persistent* effects. A shock to a key driving variable should reveal new and important information about future output. Our identification method bundles all predictive information for future real per capita GDP.

We study the effects of a given shock in a dynamic factor model. Our shock effectively captures *all* relevant (i.e. predictive) information contained in the variables. We include thirty-three variables, each of which depends on a number of latent factors that follow a first-order normal vector autoregression (VAR). Our variables include measures of GDP, unemployment, investment, consumption, sales, orders, and inflation, as well as financial market indicators, such as stock returns, a stock earnings yield, and three yield curve metrics.

Our main findings can be summarized as follows. A typical business cycle lasts about three years, and it boosts economic output (GDP) as well as a host of other measures of economic activity like consumption, productivity, investment, profits, and orders. The effect of the shock diminishes quickly. In fact, after about one year the effect on GDP growth and many other variables is negative. Subsequently, the variables rebound back to their long-run means. During a cycle, the utilization of labor and capital increases (but with a delay) and consumer confidence is higher than its long-run mean.

The yield curve is higher and flatter during an economic boom, while corporate credit spreads are tighter. This suggests that during downturns people bid up safe bonds, while requiring higher returns on risky bonds. Stock returns lead the cycle, which is not surprising since unexpected good news about future economic conditions should be instantaneously reflected in higher stock prices. In fact, we find that higher stock prices not only reflect higher future earnings but also lower future excess returns. In general, during upturns people appear to be more willing to take on investment risk on both stocks and bonds, while during recessions people require higher returns for bearing financial risk.

We also examine whether our results are consistent with various theories, and tentatively conclude that our results support the 'traditional' theory of business cycle fluctuations being driven chiefly by shocks to productivity.

Our paper is related to a recent literature in which shocks are identified by their predictive power. This body of research has come to be known as the 'news-driven business cycle' literature; see Beaudry and Portier (2014) for an overview.

Uhlig (2004b) estimates a 'core' (seven variables) and a 'periphery' (fourteen variables) VAR. The core variables are predictive of the periphery variables, but not vice versa. He then identifies two shocks that together maximize the predictive power for GNP over a five-year horizon. The first shock appears to resemble a technology shock, while the second is more in the nature of a 'wage-push' or inflationary shock. Our approach is related to Uhlig's, but differs from it in a number of ways. First, we use a much larger collection of variables with a variety of frequencies and historical availabilities. Second, we allow all variables to have predictive power, not just a subset. For example, stock returns are not one of Uhlig's core variables, but we allow them to have predictive power. Third, while Uhlig maximizes predictive power over a five-year horizon, our shock is maximally predictive for one particular future point in time only. It turns out, however, that our shock is also highly predictive for horizons other than its main target.

A natural candidate for locating shocks that predict future economic developments is stock prices, since they (presumably) quickly reflect new information about expected future productivity growth. Beaudry and Portier (2006) do exactly this, using two- to four-variable VARs, and they find evidence that the revelation of this information in itself causes aggregate fluctuations, well before actual technological improvements (as measured by total factor productivity) materialize.

Barsky and Sims (2011) argue that there may well be many variables besides stock prices that are predictive of future total factor productivity (TFP) and GDP. They estimate a seven-variable VAR and identify a shock that has no effect on contemporaneous TFP, but concentrates predictive power for future TFP over a twelve-year horizon. Their shock leads to higher consumption, but lower GDP, hours, and investment. They argue that although these results are compatible with macroeconomic models with news shocks, the co-movements are not typical in U.S. time series, and so they conclude that their identification scheme (and the theory behind it) is not supported by the data.

Of course, seven variables may not be sufficient to capture the full extent of new information revealed in each period. Instead of expanding the number of variables, Forni et al. (2013) propose to identify shocks by assuming that observed behavior in a *subsequent* time period fully reveals the new information received by agents in a previous time period; see also Lorenzoni (2009).

Like in our approach, Forni et al. (2009) and Forni, Gambetti, and Sala (2014) use a relatively large collection of variables to study the effects of predictive shocks. But unlike our approach they use a complete panel of quarterly data (monthly data are aggregated to quarterly), while we use data with a variety of frequencies and historical availabilities. Also, we estimate a latent factor model, while they construct factors as linear functions of the variables, using principal components of the sample variance matrix. Most importantly, Forni et al. (2009) identify shocks by assuming that it is the only shock to have a long-run effect on per capita output (GDP), while Forni, Gambetti, and Sala (2014) focus instead on TFP and look at two types of identification. Our method of shock identification is rather different and will be explained in Section 3.

The main advantage of dynamic factor models (DFMs) is that they describe the joint dynamics of a large number of related variables in only a few dimensions. But DFMs are also ideally suited to handle missing data. This is important for us because, as is typical in this area of study, many data are missing due to variation in the point in time at which data collection began, variation in measurement frequency, and variation in the point in time at which new data are released.

The literature on DFMs in macroeconomics has largely focused on their applicability to forecasting. Indeed, a key benefit is that they naturally provide inference on *all* missing data points — past, present, and future. Inference on present-time variables is sometimes called 'nowcasting', and our approach draws on recent papers focusing on this concept; see Forni et al. (2000), Bai (2003), Bernanke and Boivin (2003), Ghysels, Santa-Clara, and Valkanov (2004), Evans (2005), Giannone, Reichlin, and Small (2008), Aruoba, Diebold, and Scotti (2009), Faust and Wright (2009), Bańbura and Modugno (2010), Bańbura, Giannone, and Reichlin (2010), Schorfheide and Song (2012), Doz, Giannone, and Reichlin (2012), Bańbura et al. (2013), and Hindrayanto, Koopman, and de Winter (2014).

Nowcasts and forecasts can be updated whenever new data are released, and thus allow us to closely monitor the U.S. economy in real time, gauge its current state, and provide better-informed forecasts.

The plan of this paper is as follows. In Section 2 we present and discuss the model. In Section 3 we show how we identify the model and construct a typical business cycle. In Section 4 we discuss our main results, and Section 5 concludes. There are three appendices. Appendix A discusses how we deal with quarterly observations in a monthly setup, Appendix B describes the estimation procedure, and Appendix C lists our data sources.

2 The model

We consider a vector y_t containing J variables observed in month t, where $t = 1, \dots, T$. The vector y_t depends linearly on K latent factors collected in the vector f_t, and these factors follow a mean-zero normal vector autoregression. The model is thus given by

$$y_t = \alpha + \Lambda f_t + \varepsilon_t, \qquad \varepsilon_t \sim \mathrm{N}_J(0, R),$$

where the variance matrix R is diagonal and

$$f_t = A f_{t-1} + u_t, \qquad u_t \sim \mathrm{N}_K(0, V).$$

Assuming that R is diagonal is without loss of generality. The model could for instance include a number of factors that are uncorrelated over time, representing correlated but completely transient shocks. Or, a factor could exclusively affect a single observable variable, and no others.

Although this is a relatively simple and well-known model, there are various peculiarities and underlying assumptions that we wish to highlight. First, we assume a stationary distribution, which implies that α represents the long-run mean of y_t. To enforce this assumption we impose an upper bound on the modulus of the eigenvalues of A; see Appendix B for details.

Second, we assume that some of the components in α are equal to each other. The variables in this subset all measure nominal per capita growth rates, and if we would not make this assumption then the model would predict that, in the long run, the ratio of (say) consumption to GDP would become either infinitely large or small. The assumption that long-run real growth rates are non-zero is of course unrealistic and should not be taken literally. These growth rates should be interpreted as average growth for the recent past and not too distant future.

Third, as with any latent factor model, we must discuss (the lack of) identification. This arises because for any $K \times K$ non-singular matrix S we can express the model equivalently as

$$y_t = \alpha + \Lambda^* f_t^* + \varepsilon_t, \qquad \varepsilon_t \sim \mathrm{N}_J(0, R),$$
$$f_t^* = A^* f_{t-1}^* + u_t^*, \qquad u_t^* \sim \mathrm{N}_K(0, V^*),$$

where

$$\Lambda^* = \Lambda S^{-1}, \quad f_t^* = S f_t, \quad A^* = S A S^{-1}, \quad u_t^* = S u_t, \quad V^* = S V S'.$$

The transformed model implies the same data-generation process for y_t and is therefore equivalent to the original one. Identification requires fixing a subset of the parameters in such a way that applying any choice of S (except $S = I_K$) would modify those fixed parameters. For example, setting $V = I_K$ does not identify the model, since we can choose any orthogonal matrix S and retain $V^* = S I_K S' = I_K = V$. In contrast, setting $V = I_K$ *and* also selecting a non-singular $K \times K$ matrix of model parameters Ξ does identify the model, because now there is an associated *unique* orthogonal matrix Q_Ξ such that

$$\Xi' = Q_\Xi L_\Xi',$$

where L_Ξ is also unique, lower triangular, and with positive diagonal elements (this is the *QR decomposition*). Thus, if we set $V = I_K$ and fix Ξ to be lower triangular with positive diagonal elements, then the only transformation of the model that retains these two features is $S = I_K$.

Our identification method follows this line of reasoning. Specifically, we estimate the model using the normalization $V = I_K$ (i.e. we estimate a non-identified model). Then, after estimation, we identify the model using a QR decomposition on a particular Ξ matrix. Our cycle construction method (discussed below in Section 3) forms the basis for this identification scheme.

In addition to the above three concerns, there are two data issues. The first issue is that our model is formulated in terms of nominal variables, but that we are typically more interested in real variables, particularly real growth rates. We do not inflation-adjust the data, since the release dates for (new) inflation data do not coincide with release dates of our nominal variables. Since we measure all growth rates in logarithms, we can easily create real (i.e. inflation-adjusted) variables, as follows:

$$y_{jt}^{real} = y_{jt}^{nominal} - \text{inflation}_t.$$

The statistical model for real variables can then be written as

$$y_{jt}^{real} = \alpha_j^{real} + \lambda_j^{real\prime} f_t + \varepsilon_{jt}^{real},$$

with

$$\alpha_j^{real} = \alpha_j - \alpha_{\text{infl}}, \qquad \lambda_j^{real} = \lambda_j - \lambda_{\text{infl}},$$

and

$$\text{var}(\varepsilon_{jt}^{real}) = r_j^2 + r_{\text{infl}}^2,$$

where α_{infl}, λ_{infl}, and r_{infl}^2 correspond to the parameters of the inflation variable. We let the matrix Λ^{real} be equal to Λ, except that for all rows corresponding with nominal variables λ_j is replaced by λ_j^{real}. Likewise, we let y_t^{real} and α^{real} correspond with y_t and α, but with all nominal entries replaced by real ones. We use core inflation (excluding food and energy) as our measure of inflation. Thus, our nominal model can be easily adjusted to make statements about real growth rates.

The second data issue concerns the fact that seven of our variables are growth rates that are only observed at a quarterly frequency, while our model is based on the assumption that all growth rate variables in y_t are monthly growth rates. An empirical implementation of the model thus requires some assumptions on the relationship between the observed quarterly and the unobserved monthly growth rates. These assumptions are discussed in Appendix A.

Finally, a few words about the estimation method. Models of this type are typically estimated using an Expectation-Maximization (EM) algorithm.

Instead we employ a numerical optimizer to find the maximum likelihood estimate, which may also be viewed (from a Bayesian perspective) as the maximum a posteriori probability estimate under an improper prior distribution. To assist numerical optimization we derive the gradient of the log-likelihood function. For our numerical optimizer we use the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm, as implemented in the `optim` function in `R`.

3 Construction of the cycle

Our primary variable of interest is real per capita GDP growth. We seek to isolate the main forces driving its dynamics and examine how other variables tend to behave as real GDP traces out a typical cycle. We consider a 'main driving force' to be something with a *persistent* impact. In this context, persistent means that an important driving factor should still have a significant impact on GDP after τ months have passed.

Our identification scheme concentrates into a single factor all relevant information for forecasting GDP growth τ periods into the future. The factors are of course latent (unknown). Conceptually we thus assume that, at any point in time, sufficient information is available to infer the present value of the factors with high precision. If we would observe the factors f_t, then expected future values of $y_{t+\tau}$ would be given by

$$\mathrm{E}\left(y_{t+\tau}^{real} \mid f_t\right) = \alpha^{real} + \Lambda_\tau^{real} f_t, \qquad \Lambda_\tau^{real} = \Lambda^{real} A^\tau.$$

Given a choice of τ, our transformation matrix S makes Λ_τ^{real} block lower triangular with positive diagonal entries. As a result, the first factor contains all information relevant for forecasting $y_{1,t+\tau}^{real}$ (real per capita GDP growth); factors 1 and 2 together contain all information for $y_{2,t+\tau}^{real}$, and so on.

To find the transformation matrix S, we first select a number of 'focus' variables, equal to the number of factors, which we set to $K = 22$. The first focus variable is real GDP growth. We select the rows of Λ_τ^{real} corresponding to the focus variables and we denote this square matrix by $\Lambda_{\tau *}$. Then we apply the QR decomposition to its transpose:

$$\Lambda_{\tau *}' = QL',$$

where Q is orthogonal and L is lower triangular. This transformation of $\Lambda_{\tau *}$ requires post-multiplication by $S^{-1} = S'$. Thus we achieve our desired result by setting $S = Q'$.

The cycle shows the dynamic effects of a one-time shock, i.e. an impulse response function (IRF). In our case we study a particular realization of the vector u_t. Our cycle is an IRF with initial (at time 0) shock u:

$$\mathrm{E}\left(y_{t+h}^{real} \mid f_t = u\right) = \alpha^{real} + \Lambda^{real} A^h u \qquad (h = 0, 1, \dots).$$

We set u at a typical value, as follows. The first element u_1 equals the standard deviation of a typical 'real GDP shock', which is 1. We set the remaining elements to be their expected values conditional on $u_1 = 1$. Given that $V = I$, we thus have $u' = (1, 0, \dots, 0)$. An IRF represents a typical evolution of GDP growth. On average (typically) the factors and variables are at their long-run means, respectively 0 and α. Then, a typical event happens to the first factor: a one-standard-deviation shock. The IRF traces out the expected (again, typical) evolution of the economy following this shock.

Although the selection of focus variables determines the parameter identification, it does *not* affect the cycle. This is due to the fact that the initial shock vector is not affected by the particular focus variables chosen.

4 Empirical characterization of the business cycle

Given the theoretical framework developed in the previous two sections we can now provide an empirical characterization of the business cycle. We shall present four figures to illustrate this characterization. In addition, we shall discuss some implications of our findings and investigate the sensitivity of our results to the choice of the time horizon τ which we set at $\tau = 30$, i.e. 2.5 years.

The data used are accounted for in Appendix C, which also provides an overview of all variables. Most data come from the Federal Reserve's FRED database from January 1948 onwards. In total, 39% of data points are missing.

4.1 Graphical illustration

How does our shock dynamically drive the variables? Let us first consider our main focus variable, real per capita GDP growth.

FIGURE 1 HERE

In Figure 1 we see that annualized real GDP spikes up by up to 1.2% at the beginning of the cycle. The growth rate then drops, and after one year it becomes slightly negative for a two-year period. The estimated long-run mean for nominal GDP growth is 5.4%. If we subtract average core inflation, 3.9%, we arrive at an average real growth rate of 1.5%.

FIGURE 2 HERE

Figure 2 presents the IRFs for real (per capita) growth in consumption, labor productivity, investment, and corporate profits. Each of these receives a strong initial boost, and all exhibit dynamic patterns similar to real GDP growth. The magnitudes of the shocks differ however: consumption and productivity are on the same scale as GDP, whereas investment and profits are much larger, up to 20%. Several other variables mimic the pattern of consumption and profits, in particular industrial production, (durable) orders growth, (retail) sales growth, and oil price growth, but again with varying magnitudes.

FIGURE 3 HERE

The cyclical behavior of labor market indicators and capacity utilization is illustrated in Figure 3. The shock leads to lower unemployment, but the main effects are delayed by about one year. At its peak, unemployment drops by about 0.2%. We estimate unemployment to be 5.7% on average. If we include marginally attached workers, then unemployment averages 10.7%. Per capita job openings and capacity utilization show a similar (but obviously opposite) pattern as unemployment. At its peak, the shock leads to an increase in utilization of about 0.7%. We estimate capacity utilization to be 80% on average.

FIGURE 4 HERE

We now turn to six financial variables, as shown in Figure 4. The shock leads to immediate positive stock returns, which is not surprising since news of greater future growth is incorporated into stock prices. The initial return boost is 2%. As stock prices increase, the stock earnings yield mechanically decreases. Following the shock, excess stock returns are higher than average for a while, but then turn negative at about the same time that GDP growth turns positive. This suggests that investors are more willing to bear investment risk during good economic times.

The shock leads to both lower real and nominal bond yields, possibly because of higher expected consumption growth: smooth consumption is, *ceteris paribus*, preferable to uneven consumption and hence people will bid up the price of current consumption relative to future consumption. A lower price of future consumption relative to current consumption implies a higher interest rate. Another possible reason is uncertainty about consumption growth. The lower this uncertainty, the less people will want to save for precautionary reasons. As people save less, they bid down bond prices. Hartzmark (2014) finds that the latter factor is the most important determinant of interest rates. The reason that interest rates rise during a boom could therefore well be caused by the fact that booms also coincide with lower perceived consumption risk.

The slope of the yield curve decreases, i.e. we expect to see a flatter curve, while its curvature (not pictured) decreases. A steeper yield curve could indicate that short-term interest rates are expected to increase in the future. Evidence suggests however that a steeper slope predicts a higher excess return on long-term bonds (Fama and Bliss, 1987). We also see that corporate yield spreads are lower, at least during the first year; cf. Gilchrist and Zakrajšek (2012). It thus appears that during 'good times' investors bid up the prices of all risky assets, whether they be stocks, duration-risky bonds, or credit-risky bonds. Perhaps not unrelated is our finding that the shock leads to greater consumer confidence.

4.2 Interpretations

While the main goal of this paper is to empirically characterize a typical business cycle, we can also examine to what extent our findings are supportive of various theoretical interpretations. Are business cycles mainly driven by shocks to productivity, oil prices, housing markets, or perhaps some other variable?

We investigate whether observed patterns satisfy necessary conditions in support of various causal interpretations. We do so by examining the degree in which information contained in the first factor is predictive of variables other than real GDP. Specifically, we compute for each variable j and various horizons h:

$$\xi_{jh} = \frac{\text{var}(y_{j,t+h} \mid f_{1t})}{\text{var}(y_{j,t+h} \mid f_t)}.$$

If, for example, labor productivity is a key driver of the cycle, then we would expect that information contained in the factor driving GDP growth is also key information for future productivity. Vice versa, if the factor that dynamically drives GDP is barely predictive of productivity, then it would not be credible that productivity is driving GDP.

FIGURE 5 HERE

We present two graphs of ξ_{jh} in Figure 5. In the top panel we show ξ_{jh} for real per capita GDP growth itself. For $h = 30$ the ratio is 100% (by construction), but we see that the shock contains most predictive information for most time horizons: ξ_{jh} is over 88%, except for $h \in \{1, 2\}$. In the long run all ratios converge to 100% as the relevance of current (time-t) conditioning information decreases.

In the bottom panel of Figure 5 we examine ξ_{jh} for labor productivity. We see that the first factor contains a substantial amount of predictive information for productivity: ξ_{jh} is over 93% for all horizons (except $h \in \{1, 2\}$). This suggests that productivity (or technology) could be a key driver of business cycles.

We next examine ξ_{jh} for a selection of other variables, focusing on its values for $3 \leq h \leq 30$ (months). For oil prices and inflation the effects are less than for GDP and labor productivity. For oil prices the effects are between 85 and 92%, for inflation between 60% and 84%. Although in both cases the effects are pronounced, they are not as strong as for productivity. Moreover, a positive cycle is characterized by higher oil prices, suggesting that oil prices tend to *respond* to overall economic conditions rather than drive them; see Hamilton (1983) and Arezki, Ramey, and Sheng (2015) on the effects of oil shocks to the macroeconomy.

Neither is housing a likely key driver of the cycle: the values of ξ_{jh} are between 7% and 53% for housing starts, and between 13% and 63% for housing supply. This finding contrasts with various studies, in particular Leamer (2015). Similarly, fluctuations in consumer preferences and sentiment are an unlikely source of business cycle fluctuations, as ξ_{jh} is less than 40% for consumer confidence.

For government spending we find a higher fraction, between 84% and 94%. Government spending may therefore be a key driver, although it is difficult to believe that productivity responds so quickly to higher government spending. It is also possible that government spending responds to economic conditions, for example more spending during times of greater growth. We would expect certain types of government spending (e.g. on basic research) to potentially boost output levels over a far longer horizon. In addition, higher government spending should theoretically reduce private consumption, and this is not what we observe during our cycle; see King and Rebelo (1999, p. 974).

The theory of the productivity-driven business cycle has several testable implications. When we compare our cycle with the IRFs of a standard real business cycle model as presented by King and Rebelo (1999, p. 968), we see that our cycle shares key features with this model: consumption and investment are both greater, with investment having a significantly stronger effect. In addition, labor hours (as proxied by unemployment) increase, in contrast with the findings by Galí (1999) and Francis et al. (2014). Interest rates are also higher, both in our cycle and the model.

4.3 Sensitivity

Finally, we examine to what extent our results are sensitive to the choice of τ. We look at cycles for alternative choices, specifically $\tau = 1, 2, \ldots, 60$ months.

FIGURE 6 HERE

If we focus on real GDP growth, it appears that there are essentially two types of cycles. Figure 6 shows the baseline, or 'type-1' cycles type in blue, and the 'type-2' cycles in black. Importantly, we have flipped the sign for all type-2 cycles, thus the two cycles appear to be mirror images of each other. However a key difference between the two cycle types is the sign of the initial shock, which is positive for type-1 cycles but negative for type-2 cycles.

For most variables the general shape of two cycle types are (after flipping signs for type-2) broadly similar, although for some (as with GDP growth) the sign of the initial shock is reversed. But there are a number of exceptions, in particular for the financial variables. To begin, for type-2 cycles the stock earnings-to-price ratio tends to increase (not decrease), corporate spreads are higher (instead of lower), and bond yields are lower (instead of higher). Stock returns tend to be negative for type-2 cycles, and consumer confidence is lower, not higher.

Taken together, these results suggest that the cumulative impact of a type-2 cycle is negative. Indeed, as we can see in Figure 6, type 2 cycles share the general shape of type-1 cycles, but feature more pronounced periods of negative GDP growth.

5 Conclusion

In this paper we presented a new approach for understanding the main sources of business cycle fluctuations. We estimated a dynamic latent factor model and identified a factor by concentrating predictive power for real per capita GDP growth. We studied the dynamic effects of a typical shock to this factor, tracing out a business cycle. The cycle lasts about three years, with many variables showing similar evolutions, including GDP, consumption, productivity, investment, orders, sales, and profits. Following a positive shock, utilization of capital and labor increases, but with a delay, and consumer confidence is higher. Interest rates rise, and valuations of risky assets increase, which suggests that people are more willing to bear investment risk during economically good times. We also considered whether our results suggest various variables as the key driver of the cycle, and we offered tentative evidence that business cycles are mainly driven by shocks to general productivity.

As always, our results are contingent on the assumptions made and all assumptions can be questioned. An important assumption is model stability; our parameters do not drift and there are no regime changes. This is particularly relevant for studying the effects of public policy. As argued by Lucas (1976), inferences based on the past will be invalid for predicting the future if an important shift in (expected future) policy has taken place; see also D'Agostino, Gambetti, and Giannone (2013).

Many variables could be added to the model, but more variables do not necessarily make a better model and the numerical optimization will become more cumbersome. (See however Jungbacker and Koopman (2015) for a method of speeding up the calculation of the likelihood function for high-dimensional dynamic latent factor models.) Other possible statistical refinements include adding more lags to the latent factor VAR, incorporating higher frequency data, and using a latent factor model outside the normal class (Creal, 2015).

Acknowledgements

I am grateful to Patrick Caldon, Andy Tang, and Harald Uhlig for useful feedback.

Appendix A: Quarterly growth rates

We noted in Section 2 that some variables are growth rates that are only observed at a quarterly frequency. This concerns the following seven variables:

> GDP growth, investment growth, corporate profits, residential investment, net government spending growth, government spending growth, and labor productivity growth.

This causes a problem because our model is based on the assumption that all growth rate variables in y_t are monthly growth rates. An approach to deal with this issue was developed by Mariano and Murasawa (2003); see also Bańbura, Giannone, and Reichlin (2010, Section 3.2). Our method is slightly different and less computationally demanding.

For each quarterly observed variable j, let Y_{jt}^m denote the monthly annualized quantity, e.g. the amount of GDP produced in month t, multiplied by twelve. Likewise, let Y_{jt}^q denote the quarterly annualized quantity, e.g. the amount of GDP produced in months t, $t-1$, and $t-2$, multiplied by four. Letting y_{jt}^m and y_{jt}^q denote the monthly and quarterly annualized log-growth rates, we then have

$$y_{jt}^m = 12\left[\log(Y_{jt}^m) - \log(Y_{j,t-1}^m)\right],$$
$$y_{jt}^q = 4\left[\log(Y_{jt}^q) - \log(Y_{j,t-3}^q)\right].$$

We assume that

$$Y_{jt}^q = Y_{jt}^m \qquad (t = 1, 4, 7, \dots).$$

In other words, the annualized amount of GDP produced in a quarter equals the annualized amount of GDP produced in the third month of that quarter. This is an approximating assumption; it will, of course, not be exactly true. From this assumption it follows that

$$y_{jt}^q = \frac{1}{3}(y_{jt}^m + y_{j,t-1}^m + y_{j,t-2}^m).$$

Thus, observed quarterly growth rates are the average of three unobserved monthly growth rates. The (unobserved) monthly growth rates are elements of the vector y_t, so that

$$y_{jt}^m = \alpha_j + \lambda_j' f_t + \varepsilon_{jt}.$$

Combining the last two equations then gives

$$y_{jt}^q = \alpha_j + \frac{1}{3}\lambda_j'(f_t + f_{t-1} + f_{t-2}) + \frac{1}{3}(\varepsilon_{jt} + \varepsilon_{j,t-1} + \varepsilon_{j,t-2}).$$

Appendix B: Estimation

Setup and likelihood

Based on Appendix A we reformulate the process driving y_t. For the quarterly growth rate variables we have

$$y_t^q = \alpha^q + \frac{1}{3}\Lambda_q(f_t + f_{t-1} + f_{t-2}) + \frac{1}{3}(\varepsilon_t^q + \varepsilon_{t-1}^q + \varepsilon_{t-2}^q), \qquad \varepsilon_t^q \sim \mathrm{N}(0, R^q),$$

while for the remaining variables we have the same as before, namely

$$y_t^m = \alpha^m + \Lambda_m f_t + \varepsilon_t^m, \qquad \varepsilon_t^m \sim \mathrm{N}(0, R^m),$$

with

$$f_t = Af_{t-1} + u_t, \qquad u_t \sim \mathrm{N}(0, V).$$

In order to combine these equations into one system we define

$$\widetilde{y}_t = \begin{pmatrix} y_t^q \\ y_t^m \end{pmatrix}, \qquad \widetilde{\alpha} = \begin{pmatrix} \alpha^q \\ \alpha^m \end{pmatrix}, \qquad \widetilde{\varepsilon}_t = \begin{pmatrix} 0 \\ \varepsilon_t^m \end{pmatrix},$$

$$\widetilde{f}_t = \begin{pmatrix} f_t' & f_{t-1}' & f_{t-2}' & \varepsilon_t^{q\prime} & \varepsilon_{t-1}^{q}{}' & \varepsilon_{t-2}^{q}{}' \end{pmatrix}',$$

$$\widetilde{\Lambda} = \begin{pmatrix} \frac{1}{3}\Lambda_q & \frac{1}{3}\Lambda_q & \frac{1}{3}\Lambda_q & \frac{1}{3}I & \frac{1}{3}I & \frac{1}{3}I \\ \Lambda_m & 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \qquad \widetilde{R} = \begin{pmatrix} 0 & 0 \\ 0 & R^m \end{pmatrix},$$

and

$$\widetilde{A} = \begin{pmatrix} A & 0 & 0 & 0 & 0 & 0 \\ I & 0 & 0 & 0 & 0 & 0 \\ 0 & I & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & I & 0 & 0 \\ 0 & 0 & 0 & 0 & I & 0 \end{pmatrix}, \qquad \widetilde{V} = \begin{pmatrix} V & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & R^q & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

The full model can then be written in state space form as

$$\widetilde{y}_t = \widetilde{\alpha} + \widetilde{\Lambda}\widetilde{f}_t + \widetilde{\varepsilon}_t, \qquad \widetilde{\varepsilon}_t \sim \mathrm{N}(0, \widetilde{R}),$$

with

$$\widetilde{f}_t = \widetilde{A}\widetilde{f}_{t-1} + \widetilde{u}_t, \qquad \widetilde{u}_t \sim \mathrm{N}(0, \widetilde{V}).$$

The gradient

We wish to obtain the gradient, i.e. the derivative of the logarithm of the likelihood:

$$q(\theta \mid Y) = \frac{\partial \log(p(Y \mid \theta))}{\partial \theta}.$$

We shall obtain this derivative indirectly through the likelihood of Y and the latent factors F jointly, $p(Y, F \mid \theta)$. The *augmented* likelihood is, apart from irrelevant constants,

$$
\begin{aligned}
\log p(Y, F \mid \theta) = &-\frac{1}{2}\sum_{k=1}^{K}\sum_{t=1}^{T}(f_{kt} - a_k' f_{t-1})^2 \\
&-\frac{1}{2}\sum_{j\in\mathcal{J}^q}\sum_{t\in\mathcal{T}_j^{obs}}\left[\log(r_j^2/3) + \frac{1}{r_j^2/3}\left(y_{jt}^q - \alpha_j - \frac{1}{3}(f_t + f_{t-1} + f_{t-2})'\lambda_j\right)^2\right] \\
&-\frac{1}{2}\sum_{j\in\mathcal{J}^m}\sum_{t\in\mathcal{T}_j^{obs}}\left[\log(r_j^2) + \frac{1}{r_j^2}\left(y_{jt} - \alpha_j - f_t'\lambda_j\right)^2\right].
\end{aligned}
$$

In the formulas above, the sets $\mathcal{J}^q$ and $\mathcal{J}^m$ contain the quarterly and monthly observed variables respectively, $\mathcal{T}_j^{obs}$ contains the time periods for which data on variable j are observed, and a_k' denotes the kth row of the matrix A.

The gradient $q(\theta \mid Y)$ and the gradient $q(\theta \mid Y, F)$ of the augmented likelihood are connected by

$$
q(\theta \mid Y) = \mathrm{E}_{F|Y,\theta}[q(\theta \mid Y, F)];
$$

see e.g. Ruud (1991). Our purpose is to first obtain the partial derivatives of the log-augmented likelihood and then take the expectation. The expectation is taken with respect to the distribution $p(F \mid Y, \theta)$, the so-called *smoothing* distribution.

We next present the various elements of the gradient of the augmented kernel. We begin with λ_j and r_j, where λ_j' is the jth row of Λ, and r_j the square root of the jth diagonal element of R. For the quarterly measured variables ($j \in \mathcal{J}^q$) the partial derivatives are given by:

$$
\begin{aligned}
\frac{\partial q}{\partial \lambda_j} &= \sum_{t\in\mathcal{T}_j^{obs}} \frac{1}{r_j^2}(f_t + f_{t-1} + f_{t-2})(y_{jt} - \alpha_j - \frac{1}{3}(f_t + f_{t-1} + f_{t-2})'\lambda_j), \\
\frac{\partial q}{\partial r_j} &= \sum_{t\in\mathcal{T}_j^{obs}}\left[-\frac{1}{r_j} + \frac{1}{r_j^3/3}\left(y_{jt} - \alpha_j - \frac{1}{3}(f_t + f_{t-1} + f_{t-2})'\lambda_j\right)^2\right].
\end{aligned}
$$

For the remaining variables ($j \in \mathcal{J}^m$) the gradient elements are:

$$
\begin{aligned}
\frac{\partial q}{\partial \lambda_j} &= \sum_{t\in\mathcal{T}_j^{obs}} \frac{1}{r_j^2} f_t(y_{jt} - \alpha_j - f_t'\lambda_j), \\
\frac{\partial q}{\partial r_j} &= \sum_{t\in\mathcal{T}_j^{obs}}\left[-\frac{1}{r_j} + \frac{1}{r_j^3}\left(y_{jt} - \alpha_j - f_t'\lambda_j\right)^2\right].
\end{aligned}
$$

We turn now to the components α_j of the long-run mean vector α. Recall that a subset of these elements are equal to one other. Let α^{equal} denote this common long-run mean, and let $\mathcal{J}^{equal}$ collect the indices of the variables that share this mean. The gradient elements are then given by:

$$\frac{\partial q}{\partial \alpha_j} = \sum_{t \in \mathcal{T}_j^{obs}} \frac{1}{r_j^2}(y_{jt} - \alpha_j - f_t'\lambda_j) \qquad (j \in \mathcal{J}^m,\ j \notin \mathcal{J}^{equal}),$$

$$\frac{\partial q}{\partial \alpha_j} = \sum_{t \in \mathcal{T}_j^{obs}} \frac{1}{r_j^2/3}(y_{jt} - \alpha_j - \frac{1}{3}(f_t + f_{t-1} + f_{t-2})'\lambda_j) \qquad (j \in \mathcal{J}^q,\ j \notin \mathcal{J}^{equal}),$$

and

$$\frac{\partial q}{\partial \alpha^{equal}} = \sum_{j \in \mathcal{J}^m \cap \mathcal{J}^{equal}} \sum_{t \in \mathcal{T}_j^{obs}} \frac{1}{r_j^2}(y_{jt} - \alpha^{equal} - f_t'\lambda_j)$$
$$+ \sum_{j \in \mathcal{J}^q \cap \mathcal{J}^{equal}} \sum_{t \in \mathcal{T}_j^{obs}} \frac{1}{r_j^2/3}\left(y_{jt} - \alpha^{equal} - \frac{1}{3}(f_t + f_{t-1} + f_{t-2})'\lambda_j\right).$$

Finally, concerning the rows a_k' of the matrix A, we have

$$\frac{\partial q}{\partial a_k} = \sum_{t=1}^{T}(f_{kt} - a_k' f_{t-1})f_{t-1} \qquad (k = 1, \ldots, K).$$

Stationarity

As discussed in Section 2, all variables are assumed to be stationary. This implies that the eigenvalues $\nu_1, \ldots, \nu_K$ of the matrix A in the first-order vector autoregression $f_t = Af_{t-1} + u_t$ must all have a modulus smaller than one. To enforce the condition that $|\nu_k| < 1$ for $k = 1, \ldots, K$, we estimate the model and thus obtain an estimate of the A matrix. We then perform a Jordan decomposition on this A matrix: $T^{-1}AT = J$, where T is non-singular and J is a Jordan matrix with the eigenvalues of A on its diagonal. If the modulus of all eigenvalues is smaller than 0.995, then we work with the estimated A matrix. If, however, the modulus of some eigenvalue, say ν_k, is larger than or equal to 0.995, then we replace it by

$$\nu_k^{new} = 0.995\nu_k/|\nu_k|,$$

so that $|\nu_k^{new}| = 0.995$. We then compute J^{new} by replacing the the 'non-stationary' eigenvalues ν_k by their 'stationary' counterparts ν_k^{new} and obtain a new estimate for A:

$$A^{new} = TJ^{new}T^{-1}.$$

All eigenvalues of the new A matrix then have modulus smaller that 0.995, so that stationarity is guaranteed.

Appendix C: Data

We use data from various sources. In Table 1 we provide a list of the variables used.

TABLE 1 HERE

The left column in Table 1 gives the abbreviated name and the right column gives a description and provides the source.

The Federal Reserve's database can be found at

http://research.stlouisfed.org/fred2,

Robert Shiller's website at

www.econ.yale.edu/shiller/data.htm,

and the Fama-French data can be found at

mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

We measure all growth rates on an annualized logarithmic scale. Also, in order to make profit growth comparable over time we multiply it by corporate leverage (FRED: NCBCMDPNWMV).

We do not inflation-adjust any of the variables. Some variables, however, are only provided in inflation-adjusted terms; in those cases we inflate the data, thus giving nominal variables. Throughout we use core inflation (excluding food and energy) as our inflation measure.

We divide all variables that we believe should scale up with the size of the U.S. population (such as GDP and consumption) by population size. This gives us per capita quantities. We use the variable POP from FRED for this purpose. (Annualized population growth has varied from 0.69% to 1.81%, current growth is around 0.73%.) We use seasonally-adjusted variables where applicable. We assume that GDP, consumption, corporate profits, investment, residential investment, industrial production, retail sales, business sales, inventory, new orders, durables orders, oil prices, and government spending all have the same long-run average (nominal) growth rate.

We follow a standard approach from the yield curve literature and focus on a small number of variables which together capture the main characteristics of the entire yield curve; see e.g. Piazzesi (2010, Section 7.2). For this purpose we define three *yield factors*:

$$\text{level} = \frac{\text{3-month yield} + \text{3-year yield} + \text{10-year yield}}{3},$$
$$\text{slope} = \text{10-year yield} - \text{3-month yield},$$

and

$$\text{curve} = \text{3-year yield} - \frac{\text{3-month yield} + \text{10-year yield}}{2}.$$

Most yield data come from Gürkaynak, Sack, and Wright (2007), which can be found at

www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html.

Some yield data come from FRED, namely, the three-month and six-month, as well as the ten-year (before July 1972) yields , and the one-, three-, and five-year (before May 1961) yields. Finally, the one-month yield comes from Ibbotson Associates.

All variables are defined at the end of the period. For example, yields for January are recorded at the end of January. For variables that measure a change, for example a growth rate or return, we measure the change during the reference month.

References

Arezki, R., Ramey, V. A., and Sheng, L. (2015). News shocks in open economies: Evidence from giant oil discoveries. Working Paper 20857, National Bureau of Economic Research.

Aruoba, S. B., Diebold, F. X., and Scotti, C. (2009). Real-time measurement of business conditions. *Journal of Business & Economic Statistics*, 27(4), 417–427.

Bai, J. (2003). Inferential theory for factor models of large dimensions. *Econometrica*, 71(1), 135–171.

Bańbura, M., Giannone, D., Modugno, M., and Reichlin, L. (2013). Nowcasting and the real-time data flow. Working Paper Series 1564, European Central Bank.

Bańbura, M., Giannone, D., and Reichlin, L. (2010). Nowcasting. Working Paper Series 1275, European Central Bank.

Bańbura, M. and Modugno, M. (2014). Maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data. *Journal of Applied Econometrics*, 29(1) 133–160

Barsky, R. B. and Sims, E. R. (2011). News shocks and business cycles. *Journal of Monetary Economics*, 58(3), 273–289.

Beaudry, P. and Portier, F. (2006). Stock prices, news, and economic fluctuations. *American Economic Review*, 96(4), 1293–1307.

Beaudry, P. and Portier, F. (2014). News-driven business cycles: Insights and challenges. *Journal of Economic Literature*, 52(4), 993–1074.

Bernanke, B. S. and Boivin, J. (2003). Monetary policy in a data-rich environment. *Journal of Monetary Economics*, 50(3), 525–546.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1), 1–45.

Cochrane, J. H. (1994). Shocks. *Carnegie-Rochester Conference Series on Public Policy*, 41(1), 295–364.

Cooley, T. F. and Prescott, E. C. (1995). Economic growth and business cycles. In: *Frontiers of Business Cycle Research*, pp. 1–38. Princeton, New Jersey: Princeton University Press.

Creal, D. (2015). A class of non-Gaussian state space models with exact likelihood inference. *Journal of Business and Economic Statistics*, forthcoming.

D'Agostino, A., Gambetti, L., and Giannone, D. (2013). Macroeconomic forecasting and structural change. *Journal of Applied Econometrics*, 28(1), 82–101.

Doz, C., Giannone, D., and Reichlin, L. (2012). A quasi maximum likelihood approach for large, approximate dynamic factor models. *The Review of Economics and Statistics*, 94(4), 1014–1024.

Evans, M. D. D. (2005). Where are we now? Real-time estimates of the macroeconomy. *International Journal of Central Banking*, 1(2), 127–175.

Fama, E. F. and Bliss, R. R. (1987). The information in long-maturity forward rates. *American Economic Review*, 77(4), 680–692.

Faust, J. and Wright, J. H. (2009). Comparing Greenbook and reduced form forecasts using a large realtime dataset. *Journal of Business & Economic Statistics*, 27(4), 468–479.

Fisher, J. (2004). Technology shocks matter. Econometric Society 2004 North American Winter Meetings 14, Econometric Society.

Forni, M., Gambetti, L., Lippi, M., and Sala, L. (2013). Noisy news in business cycles. CEPR Discussion Papers 9601, Center for Economic Policy Research.

Forni, M., Gambetti, L., and Sala, L. (2014). No news in business cycles. *The Economic Journal*, 124(581), 1168–1191.

Forni, M., Giannone, D., Lippi, M., and Reichlin, L. (2009). Opening the black box: Structural factor models with large cross sections. *Econometric Theory*, 25(5), 1319–1347.

Forni, M., Hallin, M., Lippi, M., and Reichlin, L. (2000). The generalized dynamic-factor model: Identification and estimation. *The Review of Economics and Statistics*, 82(4), 540–554.

Francis, N., Owyang, M. T., Roush, J. E., and DiCecio, R. (2014). A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks. *The Review of Economics and Statistics*, 96(3), 638–647.

Francis, N. and Ramey, V. A. (2005). Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited. *Journal of Monetary Economics*, 52(8), 1379–1399.

Francis, N. and Ramey, V. A. (2006). The source of historical economic fluctuations: An analysis using long-run restrictions. NBER International Seminar on Macroeconomics 2004, NBER, pp. 17–73. Cambridge, MA: The MIT Press.

Galí, J. (1999). Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations? *American Economic Review*, 89(1), 249–271.

Ghysels, E., Santa-Clara, P., and Valkanov, R. (2004). The MIDAS touch: Mixed data sampling regression models. CIRANO Working Papers 2004S-20, CIRANO.

Giannone, D., Reichlin, L., and Small, D. (2008). Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics*, 55(4), 665–676.

Gilchrist, S. and Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4), 1692–1720.

Gürkaynak, R. S., Sack, B., and Wright, J. H. (2007). The U.S. treasury yield curve: 1961 to the present. *Journal of Monetary Economics*, 54(8), 2291–2304.

Hamilton, J. D. (1983). Oil and the macroeconomy since World War II. *Journal of Political Economy*, 91(2), 228–248.

Hamilton, J. D. (1993). *Time-Series Analysis*. Princeton, New Jersey: Princeton University Press.

Hartzmark, S. M. (2014). Economic uncertainty and interest rates. Working paper, The University of Chicago.

Hindrayanto, I., Koopman, S. J., and de Winter, J. (2014). Nowcasting and forecasting economic growth in the euro area using principal components. DNB Working Papers 415, Netherlands Central Bank, Research Department.

Jungbacker, B. and Koopman, S. J. (2015). Likelihood-based dynamic factor analysis for measurement and forecasting. *The Econometrics Journal*, 18(2), C1–C21.

King, R. G. and Rebelo, S. T. (1999). Resuscitating real business cycles. In: J. B. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*, Vol. 1, pp. 927–1007. Amsterdam: Elsevier.

Kydland, F. E. and Prescott, E. C. (1982). Time to build and aggregate fluctuations. *Econometrica*, 50(6), 1345–1370.

Leamer, E. E. (2015). Housing really is the business cycle: What survives the lessons of 2008–09? *Journal of Money, Credit and Banking*, 47(S1), 43–50.

Long, J. B. and Plosser, C. I. (1983). Real business cycles. *Journal of Political Economy*, 91(1), 39–69.

Lorenzoni, G. (2009). A theory of demand shocks. *American Economic Review*, 99(5), 2050–2084.

Lucas, R. J. (1976). Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, 1(1), 19–46.

Mariano, R. S. and Murasawa, Y. (2003). A new coincident index of business cycles based on monthly and quarterly series. *Journal of Applied Econometrics*, 18(4), 427–443.

Michelacci, C. and Lopez-Salido, D. (2007). Technology shocks and job flows. *Review of Economic Studies*, 74(4), 1195–1227.

Piazzesi, M. (2010). Affine term structure models. In: Y. Aït-Sahalia and L. P. Hansen (Eds.), *Handbook of Financial Econometrics: Tools and Techniques*, Chapter 12, pp. 691–766. San Diego: North-Holland.

Ramey, V. A. (2015). Macroeconomic shocks and their propagation. In: J. B. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Vol. 2. Amsterdam: Elsevier, forthcoming.

Ruud, P. A. (1991). Extensions of estimation methods using the EM algorithm. *Journal of Econometrics*, 49(3), 305–341.

Schorfheide, F. and Song, D. (2012). Real-time forecasting with a mixed-frequency VAR. Working Paper 701, Federal Reserve Bank of Minneapolis.

Uhlig, H. (2004a). Do technology shocks lead to a fall in total hours worked? *Journal of the European Economic Association*, 2(2-3), 361–371.

Uhlig, H. (2004b). What moves GNP? Econometric Society 2004 North American Winter Meetings 636, Econometric Society.

Name	Description
Dgdp	Nominal per capita GDP growth (FRED: GDP)
Dconsumption	Nominal per capita consumption growth (FRED: PCE)
Dprofits	Nominal per capita corporate profits (FRED: NFCPATAX)
Dinvestment	Nominal per capita investment growth (FRED: GPDI)
Dindpro	Nominal per capita industrial production growth (FRED: INDPRO, inflated by CPIAUCSL)
Dretail	Nominal per capita retail sales growth (FRED: RSAFS)
Dbussales	Nominal per capita business sales growth (FRED: TOTBUSSMSA)
Dinventory	Nominal per capita inventory growth (FRED: BUSINV)
Dneworder	Nominal per capita new orders growth (FRED: NEWORDER)
Ddurorder	Nominal per capita durable orders growth (FRED: DGORDER)
Dlaborprod	Nominal labor productivity growth (FRED: OPHNFB, inflated by CPIAUCSL)
Dresidentinv	Nominal per capita private residential fixed investment (FRED: PRFI)
Dgovspendnet	Nominal per capita government spending growth, net of transfers and interest (FRED: GCE)
Dgovspend	Nominal per capita total government spending growth (FRED: W068RCQ027SBEA)
Doilprice	Nominal oil price growth (FRED: MCOILWTICO)
PMI	Purchasing Managers Index (FRED: NAPM)
caputil	Capacity utilization (FRED: TCU)
U	Civilian unemployment rate (FRED: U)
U6	Total unemployed, plus all marginally attached workers plus total employed part-time for economic reasons (FRED: U6)
jobopenings	Per capita job openings: Total nonfarm (FRED: JTSJOL)
housestart	Per capita housing starts: Total — new privately owned housing units started (FRED: HOUST)
housesupply	Monthly supply of homes in the U.S. (FRED: MSACSR)
infl1month	Inflation (FRED: CPIAUCSL)
infl1monthcore	Inflation, excluding food & energy (FRED: CPILFESL)
EP	Robert Shiller's cyclically adjusted price-earnings
corpspread	Corporate yield spread (FRED: BAA - AAA)
yieldfac1	Yield curve 'level'
yieldfac2	Yield curve 'slope'
yieldfac3	Yield curve 'curvature'
realbondreturn	Inflation-adjusted return on one-month T-Bills
MktminRF	Equity excess return
realyield5yr	Five-year treasury inflation-indexed security, constant maturity (FRED: FII5)
confidence	Consumer confidence, Conference Board

Table 1: List of variables

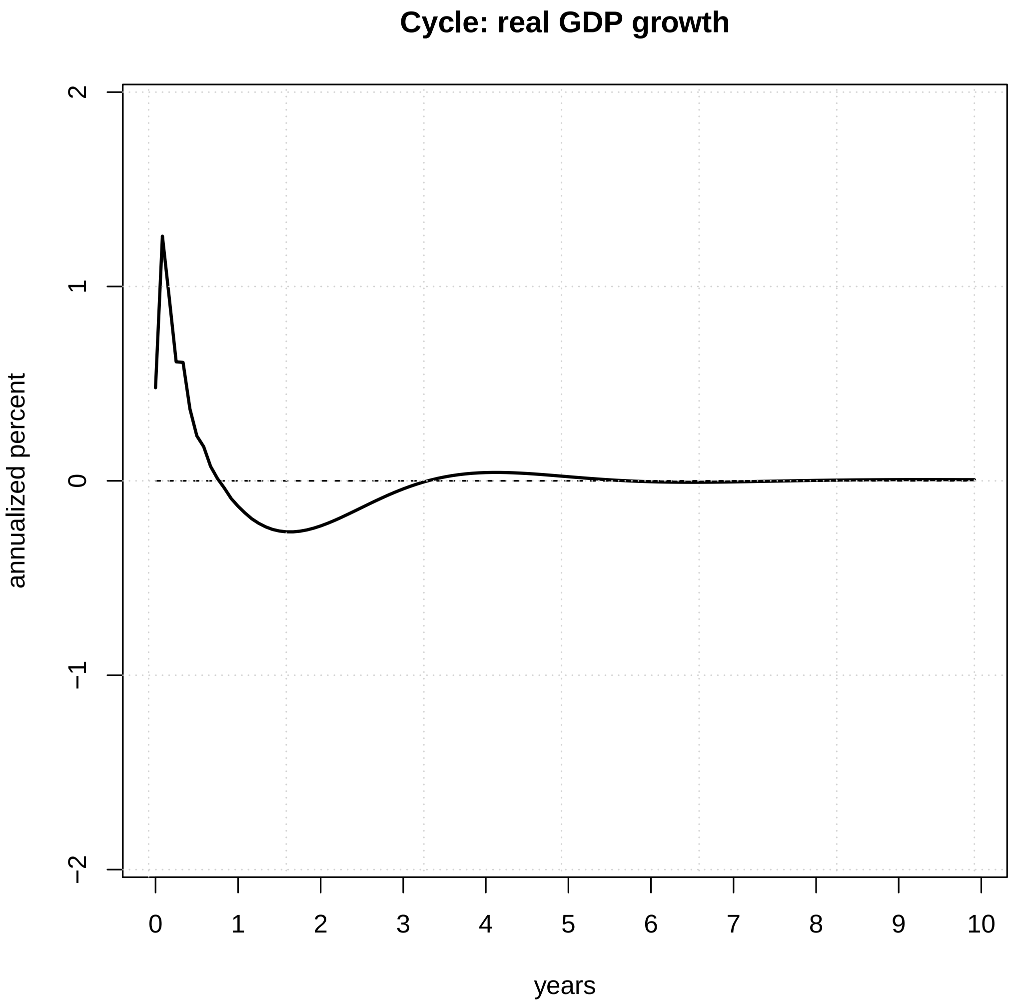

Figure 1: Cycle for real per capita GDP growth

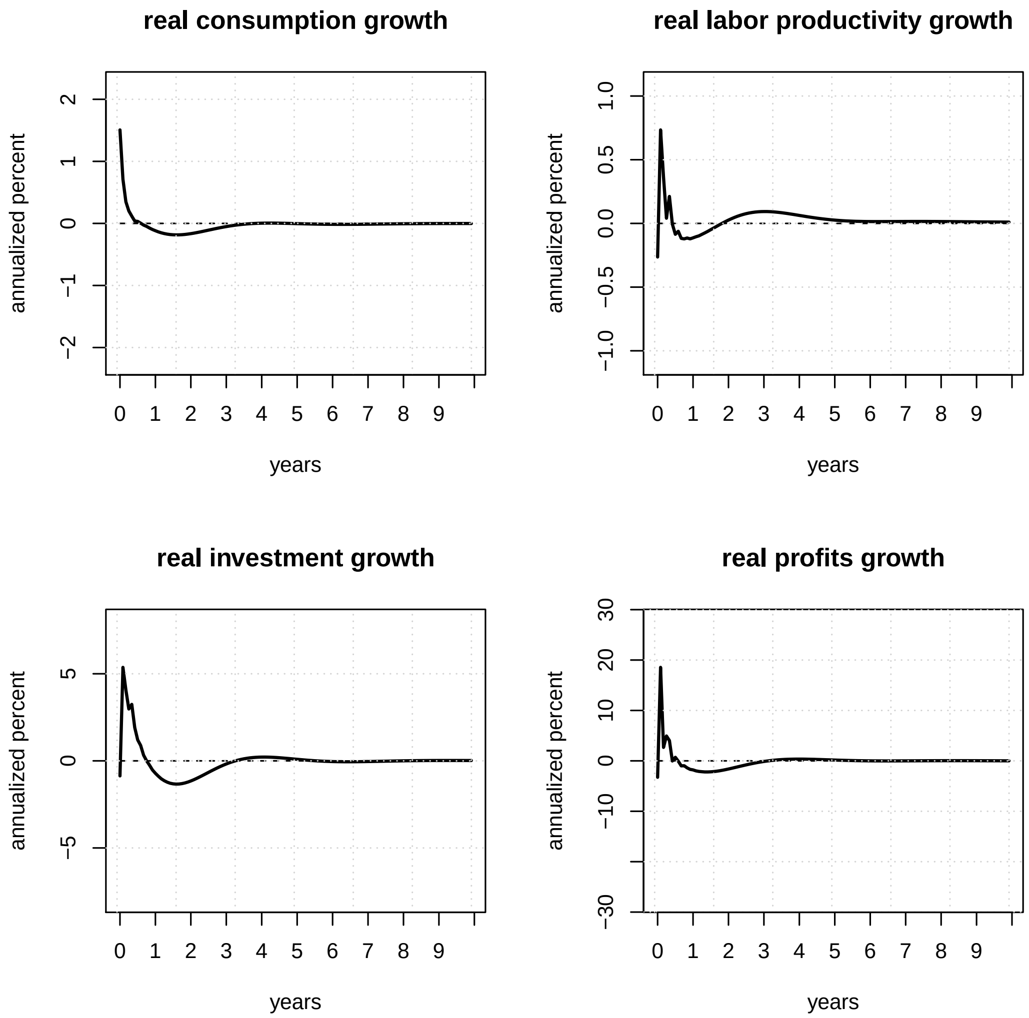

Figure 2: Cycles for real per capita consumption, labor productivity, investment, and profits

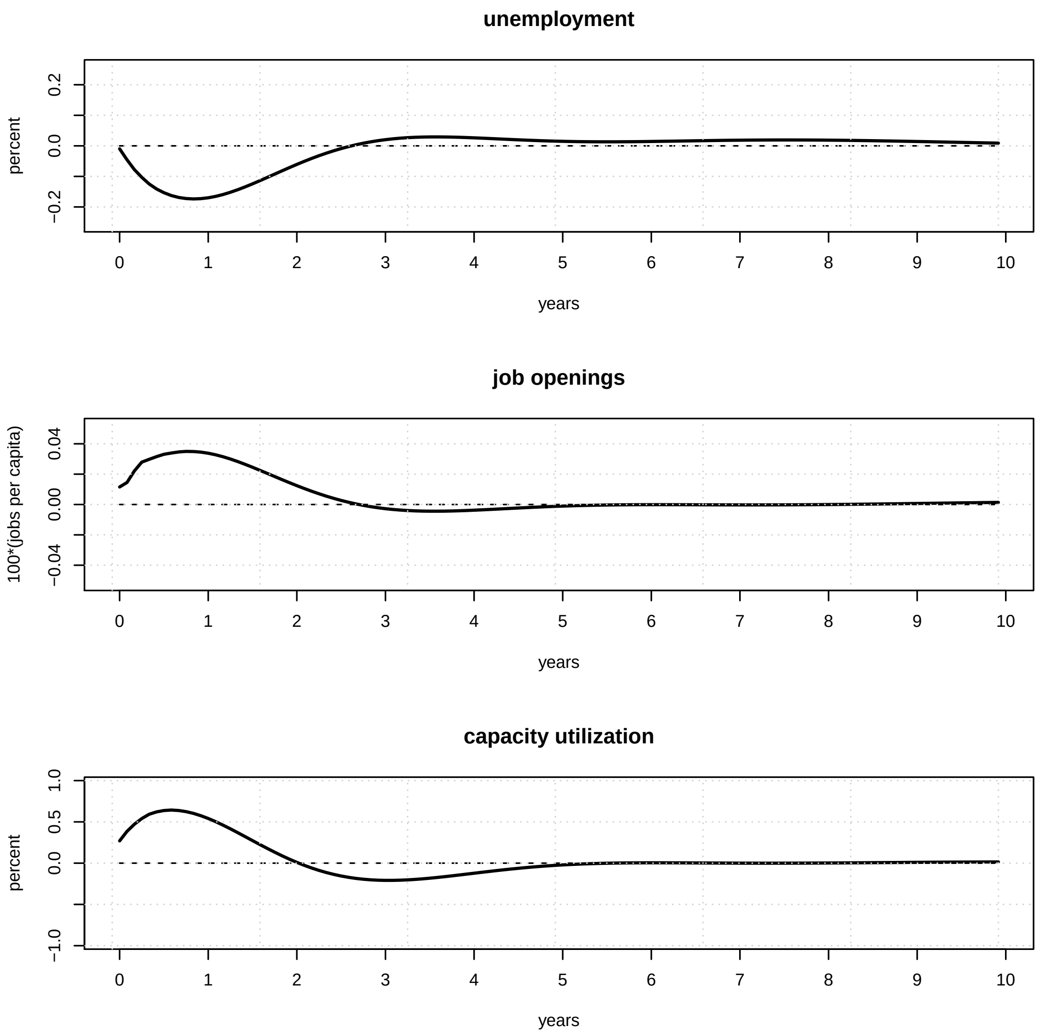

Figure 3: Cycles for unemployment, job openings, and capacity utilization

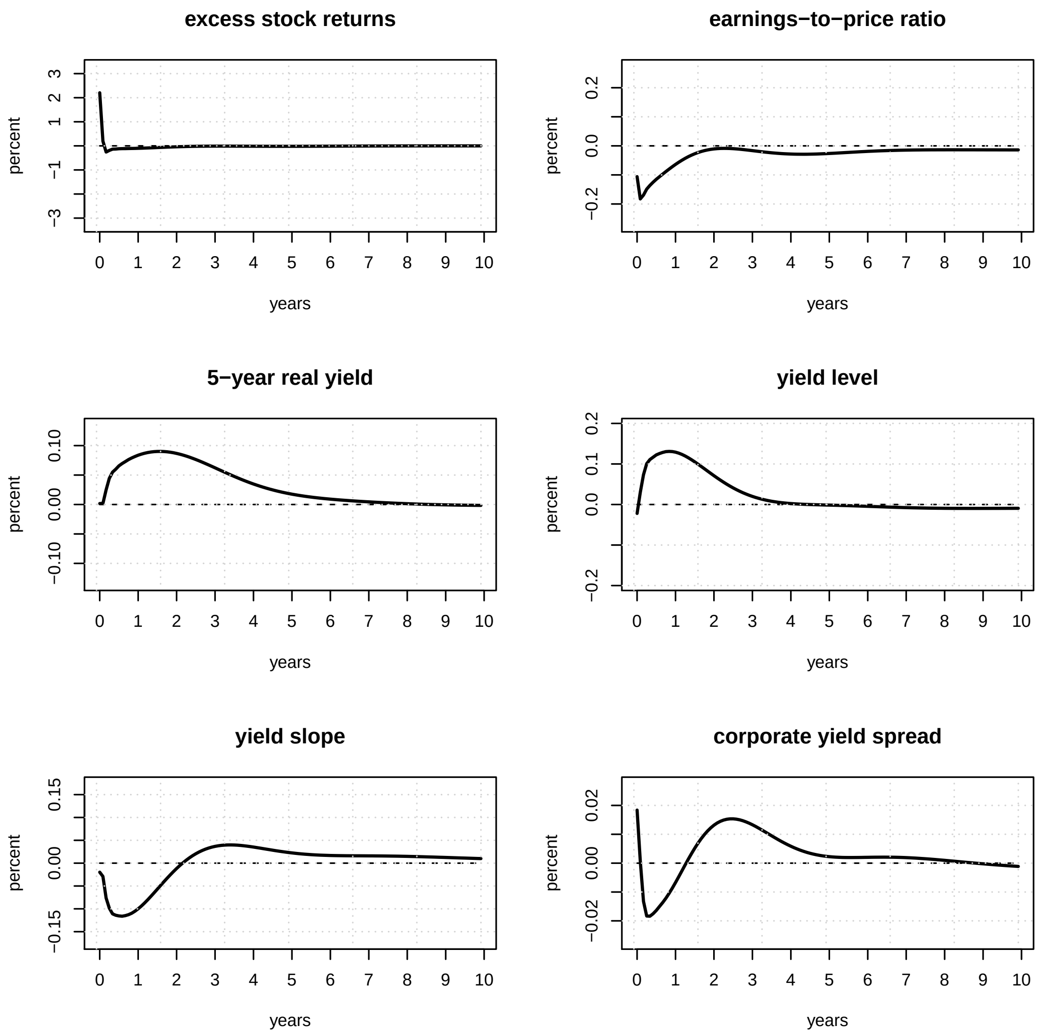

Figure 4: Cycles for various financial variables

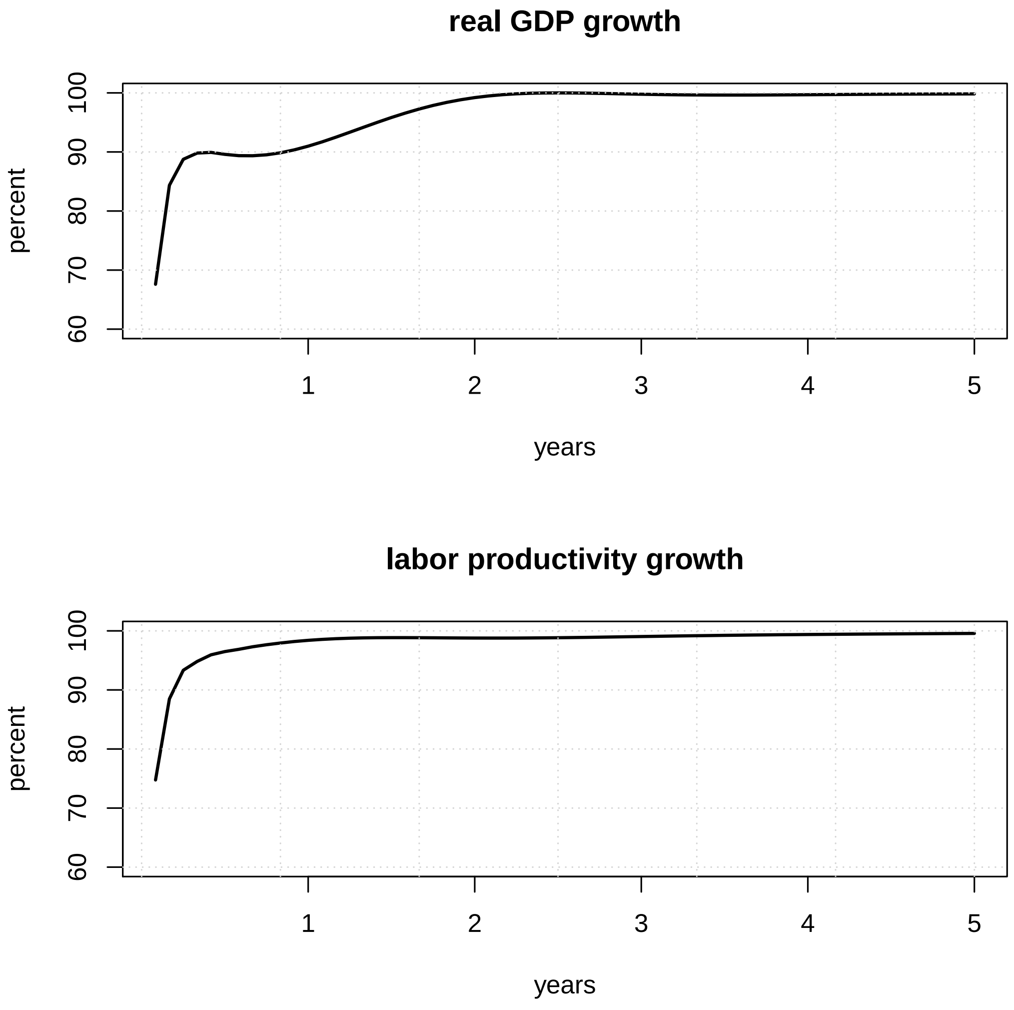

Figure 5: Fraction of conditional variance captured by the first factor: real per capita GDP growth and labor productivity growth

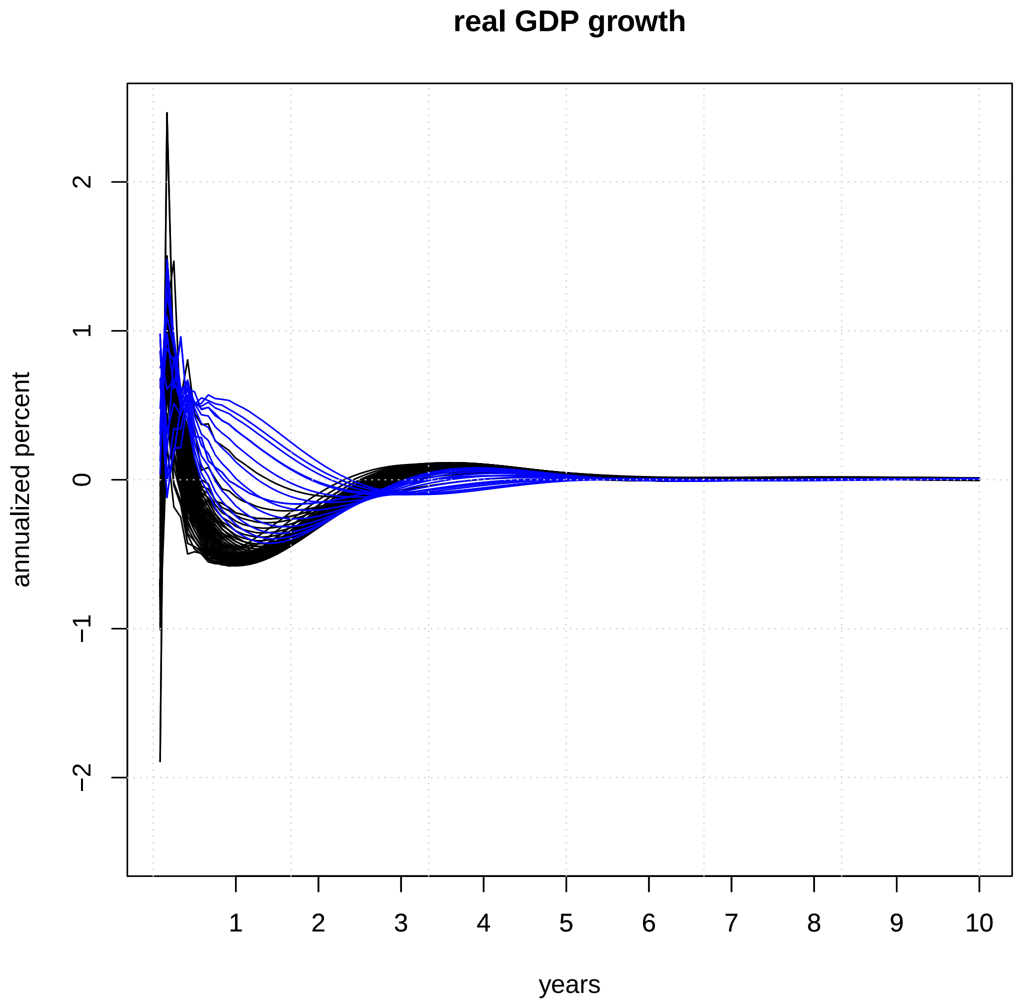

Figure 6: Cycles for $\tau = 1, 2, \ldots, 60$: real per capita GDP growth. Type-1 cycles in blue; type-2 cycles in black.