Australian National University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Financial Factors and the Business Cycle

CAMA Working Paper 44/2020
May 2020

Tino Berger
Georg-August University Göttingen
Centre for Applied Macroeconomic Analysis, ANU

Julia Richter
Georg-August University Göttingen

Benjamin Wong
Monash University, Australia
Centre for Applied Macroeconomic Analysis, ANU

Abstract

We study how financial factors shape and interact with the U.S. business cycle through a unified empirical approach where we jointly estimate financial and business cycles as well as identify their underlying drivers using a medium-scale Bayesian Vector Autoregression. First, we show, both in reduced form and when we identify a structural financial shock, that variation in financial factors had a larger role in driving the output gap post-2000 and a more modest role pre-2000. Our results suggest that the financial sector did play a role in overheating the business cycle pre-Great Recession. Second, while we document a positive unconditional correlation between the credit cycle and the output gap, the correlation of the lagged credit cycle and the contemporaneous output gap turns negative when we condition on a financial shock. The sign-switch suggests that the nature of the underlying shocks may be important for understanding the relationship between the business and financial cycles.

Keywords

Business Cycle, Financial Cycle, Financial shocks

JEL Classification

C18, E51, E32

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

Financial Factors and the Business Cycle *†

Tino Berger[1,3], Julia Richter[1], and Benjamin Wong[2,3]

[1]Georg-August University Göttingen
[2]Monash University, Australia
[3]Centre for Applied Macroeconomic Analysis, The Australian National University

Abstract

We study how financial factors shape and interact with the U.S. business cycle through a unified empirical approach where we jointly estimate financial and business cycles as well as identify their underlying drivers using a medium-scale Bayesian Vector Autoregression. First, we show, both in reduced form and when we identify a structural financial shock, that variation in financial factors had a larger role in driving the output gap post-2000 and a more modest role pre-2000. Our results suggest that the financial sector did play a role in overheating the business cycle pre-Great Recession. Second, while we document a positive *unconditional* correlation between the credit cycle and the output gap, the correlation of the lagged credit cycle and the contemporaneous output gap turns negative when we *condition* on a financial shock. The sign-switch suggests that the nature of the underlying shocks may be important for understanding the relationship between the business and financial cycles.

JEL Classification: C18, E51, E32

Keywords: Business Cycle, Financial Cycle, Financial shocks

*Email: Berger: tino.berger@wiwi.uni-goettingen.de Richter: julia.richter@wiwi.uni-goettingen.de Wong: benjamin.wong@monash.edu

†We thank James Morley and Hélène Rey for comments and helpful discussions. This research was supported by the Australia-Germany Joint Research Co-operation Scheme initiative of Universities Australia and the German Academic Exchange Service (DAAD). Benjamin Wong also acknowledges funding support from the Australian Research Council (DP190100202 and DE200100693). All remaining errors are ours. The online appendix can be found on the authors' personal websites. Berger: `https://sites.google.com/site/drtinoberger` Richter: `https://www.uni-goettingen.de/en/498940.html` Wong: `https://sites.google.com/site/benjaminwongshijie`

1 Introduction

The financial crisis of 2008-09 emphasized how developments in the financial market can spillover into the real economy. The key contribution of our paper is to quantify the role of financial factors in driving the U.S. business cycle within a unified empirical approach which allows us to jointly obtain the business and financial cycle as well as understand the underlying drivers. Building on a recent contribution by Morley and Wong (2020), our empirical approach involves estimating a medium scale Bayesian Vector Autoregression (BVAR) containing 23 U.S. macroeconomic and financial variables, and subsequently applying the Beveridge and Nelson (1981) decomposition to obtain measures of the business and financial cycle from the BVAR. A key feature of our empirical approach is that measures of the business and financial cycle obtained from the BVAR can, by construction, be decomposed into either the underlying forecast errors and/or identified structural shocks that one backs out from the BVAR. In short, we use the BVAR as a detrending device, and because we use the BVAR as a detrending device, interpretation of the business and financial cycle can be obtained using standard VAR objects such as the forecast errors or identified structural shocks.

Our empirical approach thus allows us to study the role of financial factors in driving the business cycle on three key dimensions. First, we provide an account of how much of the variation in the U.S. business cycle can be attributed to variation in the financial variables by decomposing the role of the forecast errors of these financial variables in determining the output gap, our measure of the business cycle. We consider financial variables such as credit and house prices, but also measures of financial risk, such as the VIX and credit spreads. Second, we quantify the role of an identified financial shock in driving the U.S. business cycle. A key distinction between the quantification of the role of an identified financial shock and attributing variation to forecast errors of the financial variables is the capacity to attribute causality. In particular, because variation in the financial variables can be driven by the real economy and vice versa, attributing causality requires identifying exogenous variation. Our identification of a financial shock builds directly on work by Gilchrist and Zakrajšek (2012), Caldara, Fuentes-Albero, Gilchrist, and Zakrajšek (2016), and Furlanetto, Ravazzolo, and Sarferaz (2019), but in general, the financial shock which we identify in our empirical exercise represents *exogenous* variation originating from the financial sector. Third, we study the interaction of the financial and business cycle within a common empirical model that obtains measures of both the business and financial cycle. Recall that since the financial and business cycle are obtained within the same BVAR, they are necessarily linked by common sources of variation. It is this link to common sources of variation that allows us to study how the financial and business cycle interact.

While what we do is related to the broad umbrella of work on how macro and finance interact (see, e.g. Adrian and Shin, 2010, for a review), our work is more directly related to empirical work aimed at understanding how financial factors relate to the business cycles (see, *inter alia*, Claessens, Kose, and Terrones, 2012; Jordà, Schularick, and Taylor, 2013; Aikman, Haldane, and Nelson, 2015; Albuquerque, Eichenbaum, Papanikolaou, and Rebelo, 2015). Methodologically, our empirical approach differs from the aforementioned extant work on one crucial dimension. An exercise on how much financial and business cycles interact relies crucially on how one obtains the financial and business cycle. This usually relies on implementing a filter such as the bandpass filter or the Hodrick-Prescott filter to isolate cycles and then conducting the analysis on the filtered cycles.[1] While useful, widely used filtering techniques such as the bandpass filter or the Hodrick-Prescott filter may give rise to spurious cycles, because the filtering procedure may itself impose a structure that is at odds with the underlying data generating process (see Cogley and Nason, 1995; Murray, 2003; Hamilton, 2018). A crucial advantage of our approach is that we cannot obtain spurious cycles. This is because by fitting a time series model such as a BVAR, our approach is, by construction, targeted at, and therefore consistent with, the cross-correlation and autocovariance structure of the various time series.

We note that one does not necessarily need to pre-filter the data. Unobserved components (UC) models, such as Rünstler and Vlekke (2018), have successfully estimated business and financial cycles within a unified empirical framework. As shown by Morley, Nelson, and Zivot (2003), the trend one obtains in UC models is conceptually linked to the BN decomposition through the BN decomposition of the reduced form representation of a UC model. In this regard, our work shares conceptual similarities with the UC models. However, what sets our approach apart from UC models is that we can attribute fluctuations in the cycle to either forecast errors within the BVAR or identified structural shocks by using the BN decomposition together with a BVAR. Another way of rephrasing our statement is that while UC models can model the trend and cycle, they are not designed to attribute causality or link fluctuations to observables. With our empirical approach, we can quantify, for example, how much credit mattered for the output gap as well as the role of the identified financial shock in driving the output gap. We regard the ability to quantify these effects as being crucial in understanding the role of financial factors in driving the business cycle and, at least in from this perspective, view the deviation from UC models as helpful. In short, relative to extant work, our empirical framework presents a unified framework which permits both the joint estimation of business and financial cycles, as well as linking the estimated business and financial cycles to fluctuations in observed variables or identified structural shocks.

Finally, our work is also complementary to work on how financial factors alter the output gap, albeit through applying a very different set of tools. In this vein, more structural models such as Furlanetto, Gelain, and Sanjani (2020) redefine the output gap within a DSGE environment where financial frictions are a source of inefficiencies, and thus the output gap also represents inefficiencies stemming from variation in financial frictions. More reduced form approaches such as Borio, Disyatat, and Juselius (2017) embed financial sector information

[1]Because we define cycles as it relates to some notion of a trend, our notion of what defines cycles is more accurately described as growth cycles. Classical cycles rely on dating cycles based on turning points, which is the approach taken by, for example, Claessens, Kose, and Terrones (2012). We also view the growth cycle distinction to be consistent with the framing of cycles within the broader policy discussion, as well as a reflection of much of the macroeconomics literature is couched in terms of trends and cycles. A classical cycle also does not have any natural notion of a trend and cycle, but just turning points to define a cycle. Moreover, institutions such as the Federal Reserve often frame questions such as whether the real economy is overheating as relative to its trend (i.e. the output gap), a characterization that comes much closer to the idea of growth cycles.

in conjunction with the Hodrick-Prescott filter to estimate output gaps which are "finance-neutral". Relative to the more fully structural approach by Furlanetto, Gelain, and Sanjani (2020), our approach has less structure, though we can still conduct a structural identification to quantify the role of the identified financial shock in driving the output gap. Relative to the "finance-neutral" approach, our empirical approach is more flexible and broad-based as we incorporate information from not only financial, but also other macroeconomic variables.

Our key results are as follows. First, our analysis suggests that loose financial conditions did overheat the real economy in the 2000s pre-Great Recession. From our more reduced form analysis, we find that a reasonable share of the positive output gap in the 2000s can be attributed to the excess bond premium, a credit spread constructed by Gilchrist and Zakrajšek (2012) to measure credit conditions through capturing the risk-bearing capacity of the financial sector. Our identification exercise also reveals that our identified financial shock explained a large share of the large and positive output gap in the 2000s, providing further evidence that loose credit conditions did overheat the output gap in the 2000s. Second, the role of financial factors may have been different pre and post-2000s. In particular, it appears that the role of financial factors played in the output gap was much smaller pre-2000s, but its role appears to have become much larger since the 2000s. Third, we find that while the business and credit cycle are positively correlated and vice versa, conditional on an identified financial shock, the cross-correlation between lags of the credit cycle and the contemporaneous output gap is negative. Our finding suggests that one should be careful in associating all credit booms as necessarily leading to large busts in the real economy. Indeed, our work would suggest that the *average* credit boom likely leads to a boom in the business cycle and vice versa. Only when conditioned on a financial shock, a credit boom leads to a bust in real economic activity. One interpretation, consistent with Jordà, Schularick, and Taylor (2013), suggests that it is those credit booms that originate from loose credit conditions that do lead to busts in the business cycle.

The remainder of this paper is organized as follows. Section 2 introduces the empirical framework. Section 3 presents our estimates of the financial and business cycle. Section 4 investigates the role of financial factors in driving the business cycle. Section 5 explores characteristics of the cross-correlation and lead-lag relationship between the business and financial cycle. Section 6 considers the implications of introducing a slowdown in real GDP growth into our specification. Section 7 concludes.

2 Empirical Framework

We use the Beveridge and Nelson (1981) (BN) decomposition to define the trend and cycle in our setup. Beveridge and Nelson define the trend of a time series as its long-horizon conditional expectation minus any future deterministic drift. For a time series $\{y_t\}$ which follows a random walk process with a constant drift μ, the BN trend at time t, τ_t, is

$$\tau_t = \lim_{j\to\infty} \mathbb{E}_t \left[y_{t+j} - j \cdot \mu \right]. \tag{1}$$

The cycle of the series at time t, c_t, is then defined as

$$c_t = y_t - \tau_t. \tag{2}$$

The evaluation of the conditional expectation in Equation (1) requires specifying a suitable empirical model. We build on Morley and Wong (2020) by using a medium-sized 23 variable BVAR as our empirical model. Based on the estimates of the empirical model, we then obtain trends and cycles of the various variables within the BVAR. For the business cycle, we take this as the cyclical component of real GDP. Consistent with the labeling in the wider literature and policy circles, we interchangeably refer to the business cycle as the output gap.

For measures of the financial cycle, there is less agreement about the variable of interest. For example, the cyclical component of stock prices, house prices and credit have variously been used to measure the financial cycle.[2] We also note that the Basel III regulatory accord advises policy to track the cyclical, or gap, component of the credit-to-GDP ratio, suggesting credit to be an important variable when considering the financial cycle. Guided by the broader literature, we take the cyclical component of house prices and credit as estimates of the financial cycles. Our choice of variables to consider for the financial cycle is consistent with the UC model by Rünstler and Vlekke (2018), and also consistent with the emerging consensus that the cyclical component of house prices and credit embed much of the longer frequency movement that one seeks to isolate when estimating a financial cycle (e.g., see Borio, 2014; Galati, Hindrayanto, Koopman, and Vlekke, 2016). Nonetheless, for completeness, we present results for the stock market cycle in Section C of the online appendix.

We briefly reiterate two points raised in the introduction to remind the reader of our modeling choice. First, our concept of trend and cycle is equivalent to Unobserved Components models as shown by Morley, Nelson, and Zivot (2003). However, as demonstrated by Morley and Wong (2020), the key advantage of using a BVAR is that we can directly link fluctuations in the cycles to variation of different variables within the BVAR, thus allowing us to build a richer picture of which financial variables are linked to fluctuations in the output gap. Moreover, Morley and Wong (2020) and Kamber and Wong (2020) show that standard identification tools from the SVAR literature can be easily brought into the empirical framework, a step which will be crucial for considering causality. Second, our empirical approach is immune to spurious cycles, in the Cogley and Nason (1995) and Hamilton (2018) sense, relative to using approaches such as a Hodrick-Prescott or bandpass filter (see Murray, 2003, on spurious cycles in the bandpass case).[3]

[2]We note that we use the term "cyclical component" in this statement in a generic sense as one can always extract cyclical components without involving the BN decomposition.

[3]A key point emphasized by both Cogley and Nason (1995) and Hamilton (2018) is that if the underlying data generating process was a random walk, the Hodrick-Prescott filter will attribute cycles that are spurious since the underlying time series has no forecastability, and the cycles are thus meaningless or spurious. Since

2.1 Decomposition into Trends and Cycles

Suppose we are interested in detrending K time series, where we denote each of these time series as $y_{i,t}$ where $i \in \{1, 2, \ldots, K\}$. Let $\mathbf{x_t}$ be a vector of n variables where $\Delta y_{i,t} \subset \mathbf{x_t}$.[4] We assume that $\mathbf{x_t}$ has a VAR(p) representation with the following companion form:

$$(\mathbf{X_t} - \boldsymbol{\mu}) = \mathbf{F}(\mathbf{X_{t-1}} - \boldsymbol{\mu}) + \mathbf{H}\boldsymbol{e_t}, \tag{3}$$

where $\mathbf{X_t} = \{\mathbf{x'_t}, \mathbf{x'_{t-1}}, ..., \mathbf{x'_{t-p}}\}'$, $\boldsymbol{\mu}$ is the vector of n unconditional means of $\mathbf{x_t}$, $\mathbf{F}$ is the companion matrix with eigenvalues that all are inside the unit circle, $\mathbf{H}$ maps the VAR forecast errors to the companion form, and $\boldsymbol{e_t}$ is a vector of serially uncorrelated forecast errors with covariance matrix $\boldsymbol{\Sigma}$. Denoting $\tau_{i,t}$ and $c_{i,t}$ as respectively the BN trend and cycle of the series $y_{i,t}$,

$$y_{i,t} = \tau_{i,t} + c_{i,t}. \tag{4}$$

Let $\mathbf{s_q}$ be a selector row vector with 1 at its q^{th} element, and zero otherwise. Further, let $\Delta y_{i,t}$ be in the k^{th} position of $\mathbf{x_t}$. Applying the definition of the BN decomposition, the cycle, $c_{i,t}$, can be calculated as (see Morley, 2002)

$$c_{i,t} = -\mathbf{s_k}\mathbf{F}(\mathbf{I} - \mathbf{F})^{-\mathbf{1}}(\mathbf{X_t} - \boldsymbol{\mu}). \tag{5}$$

Morley and Wong (2020) show that we can further decompose the obtained BN trends and cycles as a function of either the VAR forecast errors or structural shocks. Let $c_{ij,t}$ represent the share of the forecast error of the j^{th} variable in $\mathbf{x_t}$ on the cycle $c_{i,t}$. Similarly, let $\Delta y_{i,t}$ once again occupy the k^{th} position in $\mathbf{x_t}$. Morley and Wong (2020) show that we can write $c_{ij,t}$[5] as

$$c_{ij,t} = -\sum_{l=0}^{t-1} \mathbf{s_k}\mathbf{F^{l+1}}(\mathbf{I} - \mathbf{F})^{-\mathbf{1}}\mathbf{H}\mathbf{s'_j}\mathbf{s_j}\boldsymbol{e_{t-l}}. \tag{6}$$

Equation (6) decomposes the K cycles which we obtain through our VAR into shares of forecast errors of all the n variables contained in $\mathbf{x_t}$. We refer to Equation (6) as the informational decomposition, as it associates fluctuations in the cycles with the information contained within the other variables. At the same time, note that

$$c_{i,t} = \sum_{j=1}^{n} c_{ij,t}, \tag{7}$$

which implies that the obtained cycle from our VAR fully decomposes into the forecast

our specification nests a random walk for any differenced variable, our approach will consistently estimate the random walk process for these variables/equations, and so our approach will not fall afoul with the issue of spurious cycles.

[4] $\mathbf{x_t}$ can contain variables that are differenced or in levels. The mix of I(1) and I(0) variables does not matter as long as together, $\mathbf{x_t}$ implies a stationary VAR. We only require the variables which we are interested in detrending to be differenced, as we require variables to be I(1) in the levels to apply the BN decomposition.

[5] Morley and Wong (2020) also derive analogous expressions for the trends, but as our focus is on the business and financial cycles, we omit discussion about the trends.

errors of all the n variables contained in $\mathbf{x_t}$. Within our empirical framework, $c_{i,t}$ will represent objects of interest such as the output gap, which will be our measure of the business cycle, and the cyclical component of housing prices and credit, which represents our measure of the financial cycle. Accordingly, we will use the expression in Equation (6) to understand the role of financial variables in driving the output gap by associating fluctuations in the output gap with the forecast errors of the financial variables such as credit, house prices, stock prices, credit spreads, etc.

The decomposition in Equation (6), while informative, does not attach any causal interpretation. Attaching a casual interpretation will require identifying structural shocks. Let $\boldsymbol{\epsilon}_t$ represent a $n \times 1$ vector of orthogonal structural shocks, with the variance normalized to unity, or $\mathbb{E}\boldsymbol{\epsilon}_t'\boldsymbol{\epsilon}_t = \mathbf{I}$. The structural VAR literature shows that identifying a structural shock requires specifying a mapping

$$\boldsymbol{e}_t = \mathbf{A}\boldsymbol{\epsilon_t}, \text{ where } \mathbf{A}\mathbf{A}' = \mathbf{\Sigma}. \tag{8}$$

Let $c^S_{ij,t}$ be the share of the j^{th} structural shock on $c_{i,t}$. Using the mapping defined by Equation (8), we can substitute in Equation (6) to obtain

$$c^S_{ij,t} = -\sum_{l=0}^{t-1} \mathbf{s_k}\mathbf{F^{l+1}}(\mathbf{I}-\mathbf{F})^{-\mathbf{1}}\mathbf{H}\mathbf{A}\mathbf{s_j'}\mathbf{s_j}\boldsymbol{\epsilon_{t-l}}. \tag{9}$$

Equation (9) now allows us to interpret the business and financial cycle as a function of orthogonalized shocks, and so allows for a structural or causal interpretation. For our structural analysis, we will identify a financial shock with guidance from the wider empirical literature to understand how financial shocks drive both the business and financial cycle.

2.2 Estimation and Data

We estimate a 23 variable BVAR of U.S. macroeconomic and financial variables. The set of variables in our BVAR are real GDP, the CPI, employment, real private consumption, industrial production, capacity utilization, the unemployment rate, housing starts, the producer price index for all commodities, hours worked, nonfarm real output per hour, personal income, real gross domestic investment, the fed funds rate, the 10-year government bond yield, real M1, real M2, total credit to non-financial institutions, the S&P 500 index, real energy prices, the VIX index, real house prices, and the excess bond premium introduced by Gilchrist and Zakrajšek (2012). Most of the data is sourced from the FRED database over the sample period 1973Q1-2019Q2. Data for the excess bond premium is taken from Gilchrist and Zakrajšek (2012) its subsequent updates by the Boston Fed.[6] Most of the variables are standard, motivated in part by the specification of Banbura, Giannone, and Reichlin (2010) and Morley and Wong (2020). We provide details of the precise data source, description, and transformation in Section A of the online appendix.

[6]See `https://www.federalreserve.gov/econresdata/notes/feds-notes/2016/recession-risk-and-the-excess-bond-premium-20160408.html`.

We briefly note that our choice to work with a 23 variable BVAR is because we require a variable set that spans all the relevant information for both the business and financial cycles. More precisely, Morley and Wong (2020) show that a condition of estimating the true BN cycle is the inclusion of all the relevant forecasting information for the variables from which we are obtaining the BN cycle. At the same time, because we are making inference on the effect of a structural financial shock as part of our analysis, Forni and Gambetti (2014) show that one should include all the information that spans the SVAR shocks. The choice of the 23 variable medium-sized BVAR, as opposed to a more standard smaller six to eight variable VAR, should act as a sufficient guard against omitting relevant information.

Given the rest of the variables are standard, we only comment on the excess bond premium, which was introduced by Gilchrist and Zakrajšek (2012). The excess bond premium is a credit spread which measures the risk-bearing capacity of financial intermediaries. Faust, Gilchrist, Wright, and Zakrajšek (2013) show that the inclusion of credit spreads can help with the prediction of real economic activity. This suggests from at least the perspective of both Morley and Wong (2020) and Forni and Gambetti (2014), the inclusion of the excess bond premium, as a credit spread, is necessary as this is relevant information for aiding with the estimation of the output gap, as well as the identification of structural financial shocks. We also note that variation in the excess bond premium also plays a key role in the literature on identifying structural financial shocks (e.g. Gilchrist, Yankov, and Zakrajšek, 2009; Caldara, Fuentes-Albero, Gilchrist, and Zakrajšek, 2016), and so its inclusion within our context would also aid in the identification of structural financial shocks.

Some variables exhibit a break in the mean. If so, this implies $\boldsymbol{\mu}$ in Equation (3) and thereafter has to be adjusted. As shown by Morley and Wong (2020), these breaks in the mean can compromise the BN decomposition, as stationarity requires a variable to be mean-reverting. We thus proceed as follows. We first apply conventional transformations to the variables. To adjust for possible breaks in means, we slightly vary the treatment for the variables for which we are deriving a business or financial cycle, and the other variables.

Drift Adjustment - Business and Financial Cycle Variables For variables that we use to make inference on the business and financial cycle, a break in the mean implies a break in the drift since these variables are differenced before estimation. Given that the definition of the BN decomposition from Equation (1) depends on the drift, Kamber, Morley, and Wong (2018) show that a break in the drift can play a crucial role in obtaining reliable measures of trend and cycle. We therefore tested the variables associated with the respective financial and business cycles to ensure that the assumption of a constant drift cannot be rejected by a standard Bai and Perron (2003) test.[7] These variables under consideration are real GDP for the business cycle, and credit and house prices for the financial cycles. We found a break in the drift for credit in 2008Q1. This is not entirely surprising as the financial crisis of 2008/09 resulted in not only a stall in credit during the recession, but also a continued flattening of the drift due

[7] We tested for the break in the drift by allowing for heteroskedasticity and autocorrelation consistent (i.e. Newey and West, 1987) (HAC) standard errors.

to financial regulation post-2008 in aftermath of the crisis, notably resulting from initiatives such as the Basel Accords (notably Basel III). We therefore adjusted for a break in the drift of credit in 2008Q1.

Mean Adjustment - Other Variables For the other variables, our concern is mainly to guard against possible breaks in the mean in compromising our inference of the business and financial cycle. In particular, if there is a break in the mean in the other variables, this may imply excessive persistence instead of a quicker revision to the new (post-break) mean, and this can impart excessive persistence to our estimate of the business and financial cycle.[8] While Morley and Wong (2020) opted to difference variables if there was some evidence of a break in the mean, such an approach might be overly conservative in throwing out useful information in the level. For example, capacity utilization is a variable that exhibits a break in the mean. However, the level of capacity utilization provides a lot of information about the state of the business cycle. By differencing such a variable, we throw out a lot of useful information in the level. Kamber and Wong (2020) thus opted to adjust for breaks in the mean if there was compelling evidence to suggest so, an approach that we adapt to our setting. More precisely, we first test for a difference in the mean between the first and second half of the sample using a two-sample t-test, similar to Morley and Wong (2020). If the test rejects the null hypothesis of equal means at the 10% significance level, we follow the procedure by Kamber and Wong (2020) and use a sup-F statistic (see Andrews, 1993) to locate a break in the mean at an unknown breakpoint and use this unknown breakpoint to adjust for a break in the mean.[9] Details on the breaks are provided in Section A of the online appendix.

The estimation of the BVAR is standard. We utilize the natural-conjugate Normal-Wishart prior which draws on elements of the Minnesota Prior (e.g., see Litterman, 1986; Robertson and Tallman, 1999). Consider the VAR(p) for the vector of variables $\mathbf{x_t}$ which are demeaned before estimation:[10]

[8]The idea that excessive persistence can result from a break in the mean is not new, and has been explored and shown by Perron (1990), amongst other contributions.

[9]We tested for a break at the midpoint as a first pass as we wanted to also strike a balance against adjusting for too many breaks. If one cannot find a break in the mean using the midpoint of the sample, then we view any possible breaks in the mean as probably not sufficiently large to warrant attention. Only if we find a statistically significant difference in the mean between the first and the second half of the sample do we use the sup-F statistic to be more precise about the dating of the break.

[10]If we find a break in the mean, we adjust the $\mathbf{x_t}$ vector before estimation. This approach will be equivalent to placing a flat prior on the mean and makes the estimation of the VAR and BN decomposition straightforward. As our estimation procedure optimizes on the degree of shrinkage, the analytical properties from using the natural-conjugate prior, as opposed to Monte Carlo sampling, is a key ingredient in making our estimation procedure feasible. As noted by Morley and Wong (2020), one could model the break explicitly, though this will result in a more involved estimation procedure as we lose the analytical properties of the natural-conjugate prior and potentially makes estimation less feasible.

$$
\begin{aligned}
\mathbf{x_t} &= \mathbf{\Phi_1 x_{t-1}} + \ldots + \mathbf{\Phi}_p \mathbf{x_{t-p}} + \boldsymbol{e_t} \\
&= \begin{bmatrix} \phi_1^{11} & \ldots & \phi_1^{1n} & \phi_2^{11} & \ldots & \phi_2^{1n} & \ldots & \ldots & \phi_p^{1n} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \ddots & \vdots \\ \phi_1^{n1} & \ldots & \phi_1^{nn} & \phi_2^{n1} & \ldots & \phi_2^{nn} & \ldots & \ldots & \phi_p^{nn} \end{bmatrix} \begin{bmatrix} \mathbf{x_{t-1}} \\ \mathbf{x_{t-2}} \\ \vdots \\ \mathbf{x_{t-p}} \end{bmatrix} + \begin{bmatrix} e_{1,t} \\ \vdots \\ e_{n,t} \end{bmatrix},
\end{aligned} \tag{10}
$$

where $\mathbb{E}(\boldsymbol{e_t' e_t}) = \mathbf{\Sigma}$ and $\mathbb{E}(\boldsymbol{e_t' e_{t-i}}) = \mathbf{0}\ \forall i > 0$. We then apply shrinkage to the VAR slope coefficients using a Minnesota-type prior specification for the prior means and prior variances as follows:

$$
\mathbb{E}[\phi_i^{jk}] = 0 \tag{11}
$$

$$
Var[\phi_i^{jk}] = \begin{cases} \frac{\lambda^2}{i^2}, & j = k \\ \frac{\lambda^2}{i^2}\frac{\sigma_j^2}{\sigma_k^2}, & \text{otherwise,} \end{cases} \tag{12}
$$

where the degree of shrinkage is governed by the hyperparameter λ, with $\lambda \to 0$ shrinking to the assumption that the variables in the VAR are independent white noise processes or, equivalently for all of the differenced variables in the VAR, independent random walk processes in levels.

We obtain σ_l^2 by taking the residual variances after fitting an AR(4) on the l^{th} variable using least squares, which is a common practice (e.g., Banbura, Giannone, and Reichlin, 2010; Koop, 2013). The term $1/i^2$ governs the basic structure of the Minnesota Prior to down-weight more distant lags and the factor σ_j^2/σ_k^2 adjusts for the different scale of the data.

We follow Morley and Wong (2020) and choose λ by minimizing the one-step-ahead out-of-sample forecast error of output growth. The choice to optimize based on output growth is largely driven by our central focus on understanding how financial factors affect the output gap.

The natural conjugate Normal-Inverse-Wishart prior implies posterior moments that can be calculated either analytically or through the use of dummy observations. We will use dummy observations to estimate the BVAR (e.g., Banbura, Giannone, and Reichlin, 2010; Del Negro and Schorfheide, 2011; Woźniak, 2016). For brevity, we relegate these details to Section B of the online appendix.

3 Estimates of Business and Financial Cycles

Figure 1 presents the estimated U.S. output gap, with the bottom subplot presenting our estimate together with the associated 90% credible interval. The estimated output gap lines up with the NBER reference cycles, with turning points coinciding with NBER-dated recessions. We also note that our estimated output gap appears to be large and positive just before the

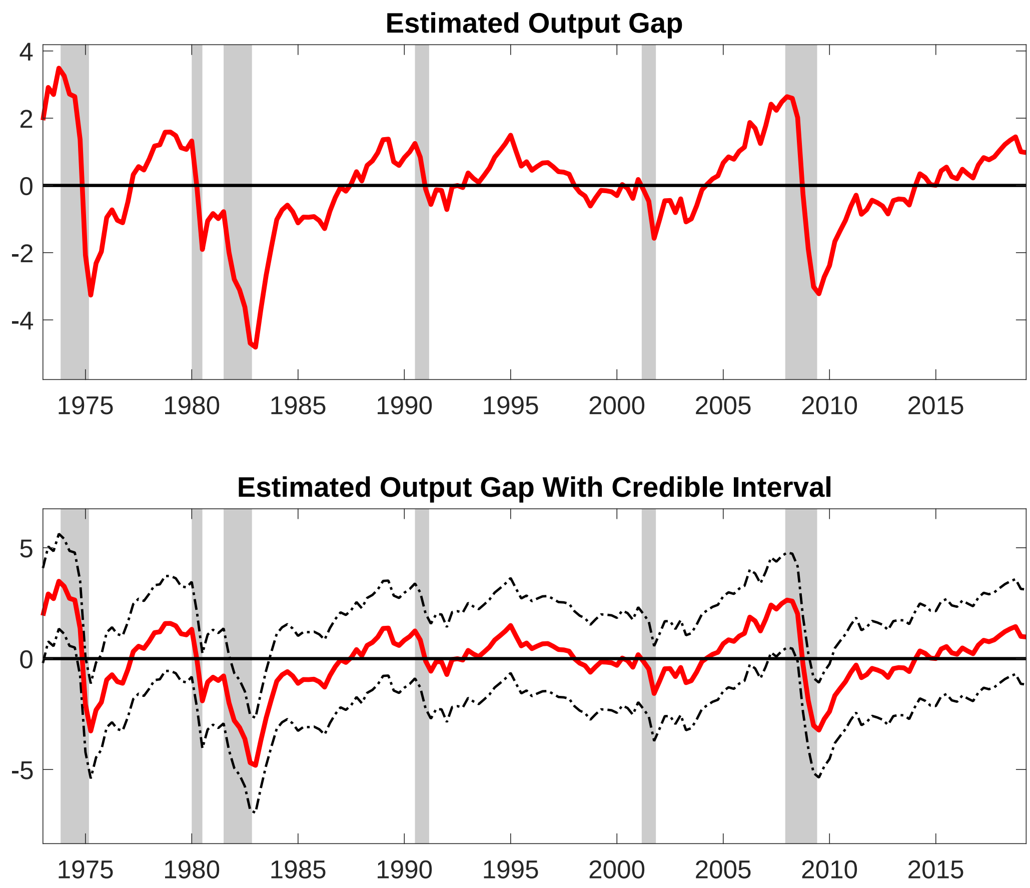

Figure 1: Estimated U.S. output gap in terms of percent deviation from the trend in real GDP. Grey shaded areas indicate NBER recessions. Bottom panel plots estimated output gap with 90% credible interval calculated as per Kamber, Morley, and Wong (2018).

Great Recession, lining up with accounts that the real economy was overheating in the 2000s (e.g., see Taylor and Wieland, 2016; Borio, Disyatat, and Juselius, 2017).

Figure 2 presents the estimated credit and house price cycle from our BVAR, which we take as estimates of the financial cycle. Our estimates are consistent with the general narratives. In particular, whether one looks at the credit or house price cycle, our estimates imply a large boom of the financial cycle in the 2000s and a bust during the Great Recession.

Recall that our estimates of the business and financial cycles do not rely on an *a priori* view on the length of these cycles, but instead only relies on an underlying BVAR and the definition of the long-horizon forecast to define the trend and cycle. As Cagliarini and Price (2017) point out, a widely held view that the financial cycle has a much longer duration than the business cycle may be partly driven by assumptions on which frequencies to isolate, potentially obscuring the distinction between assumptions and conclusions. For example, users of the bandpass filter take frequencies of $2\frac{1}{2}$ to 8 years as coinciding with the business cycle (e.g., see Baxter and King, 1999; Christiano and Fitzgerald, 2003). For the financial cycle, extant work such as Drehmann, Borio, and Tsatsaronis (2012) and Aikman, Haldane, and Nelson (2015) choose 8 to 20 or 30 years as frequencies to isolate for characterizing the financial cycle.

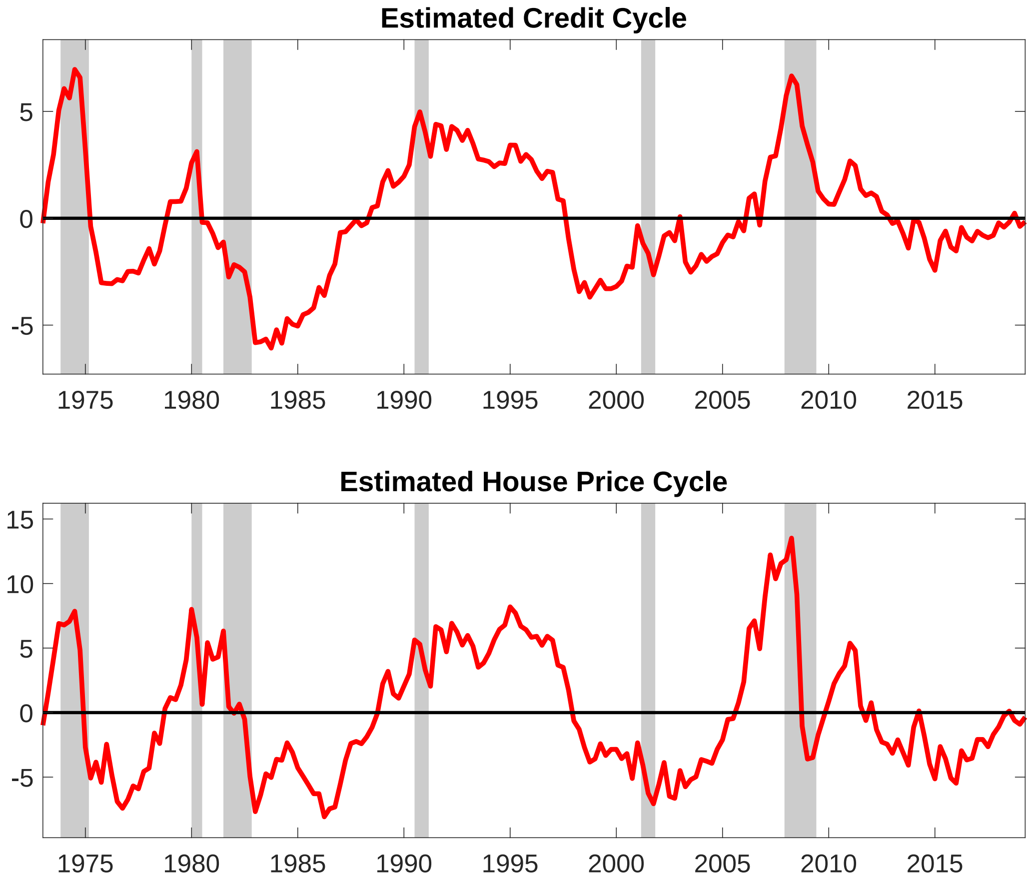

Figure 2: Estimated financial cycles from the BVAR. Units are in percent deviation from trend. Grey shaded areas indicate NBER recessions.

Because our estimates of the business and financial cycle do not rely on an *a priori* view of the length of financial and business cycles, we are in a position to reassess the view on the relative duration of the business and financial cycle through the lens of our model. Figure 3 presents the periodograms of the estimated output gap, housing cycle, and credit cycle. We highlight the frequencies between $2\frac{1}{2}$ to 8 years, 8 to 10 years, and, 10 to 20 years. Recall that $2\frac{1}{2}$ to 8 years correspond with the frequencies regularly isolated by a bandpass filter as being consistent with "business cycle frequencies" (e.g., see Baxter and King, 1999; Christiano and Fitzgerald, 2003). We find that our estimated output gap has a significant degree of fluctuations between $2\frac{1}{2}$ to 8 years. Contrast this against the housing and credit cycle, where variation across these frequencies appear to be quite limited. The peak of the spectral density of our estimated output gap occurs between 8 to 10 years with also a non-trivial component at even lower frequencies. We interpret this observation that, while fluctuations at "business cycle frequencies" are important, there is a non-trivial proportion of business cycle variation found at medium-term frequencies, consistent with Comin and Gertler (2006). The peak of the spectral density of the output gap is at about 9 years, which, while slightly outside the $2\frac{1}{2}$ to 8 years assumed by users of the bandpass filter, is not too widely at odds with the consensus view. Second, the peak of the spectral densities of the estimated housing and credit cycle both

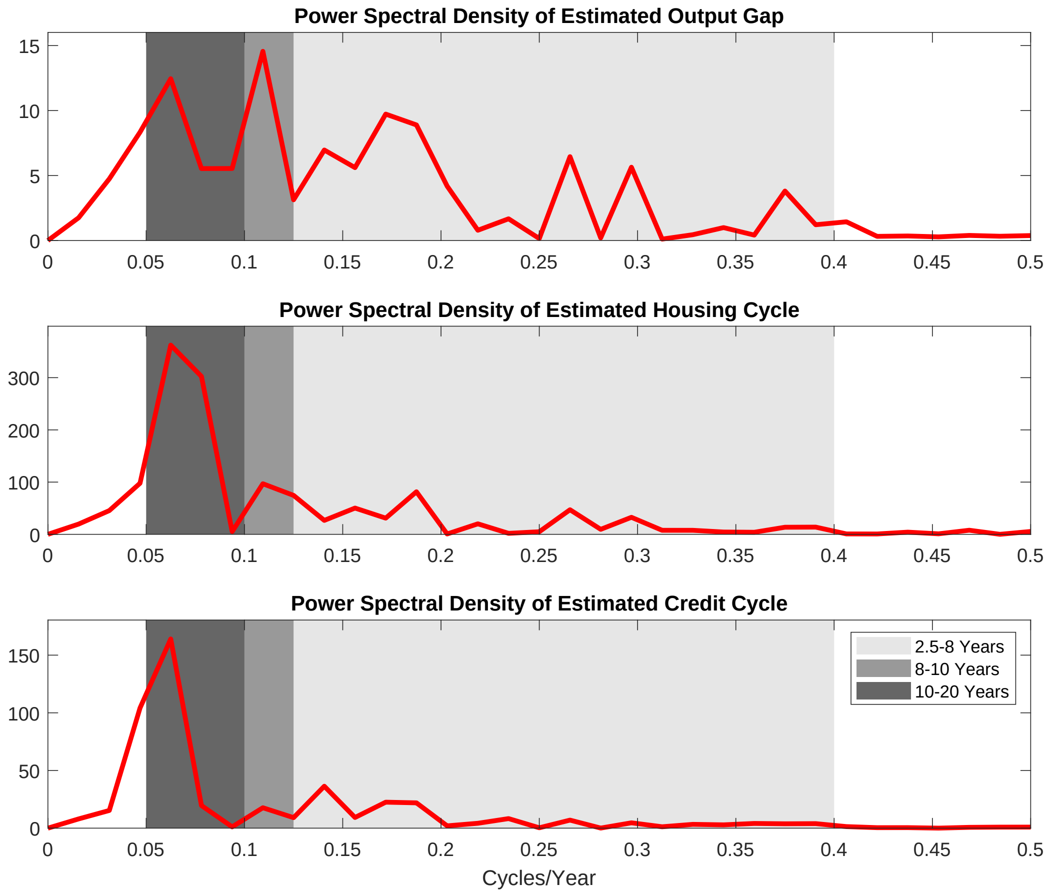

Figure 3: Power spectral density of estimated cycles. The x-axis normalizes the frequencies to number of cycles per year. The frequencies associated with $2\frac{1}{2}$ to 8 years, 8 to 10 years, and 10 to 20 years are highlighted.

lie within the 10 to 20 year window. More precisely, they peak at frequencies coinciding with 16 and 19 years respectively, very similar to extant estimates (e.g. Aikman, Haldane, and Nelson, 2015; Rünstler and Vlekke, 2018).

Overall, we find that our estimated financial cycle has a much lower frequency than our estimated business cycle, consistent with the conventional view. This conclusion still holds even when one considers that our estimated business cycle features non-trivial low-frequency fluctuations outside the traditionally assumed $2\frac{1}{2}$ to 8 years.

4 The Role for Financial Factors in Driving the Business Cycle

We now turn to the central question of the paper: what is the role of financial factors in driving the business cycle. We address this question mainly with two tools that we introduced in Section 2; the informational decomposition and structural analysis where we explicitly identify a structural financial shock.

4.1 Informational decomposition of the output gap

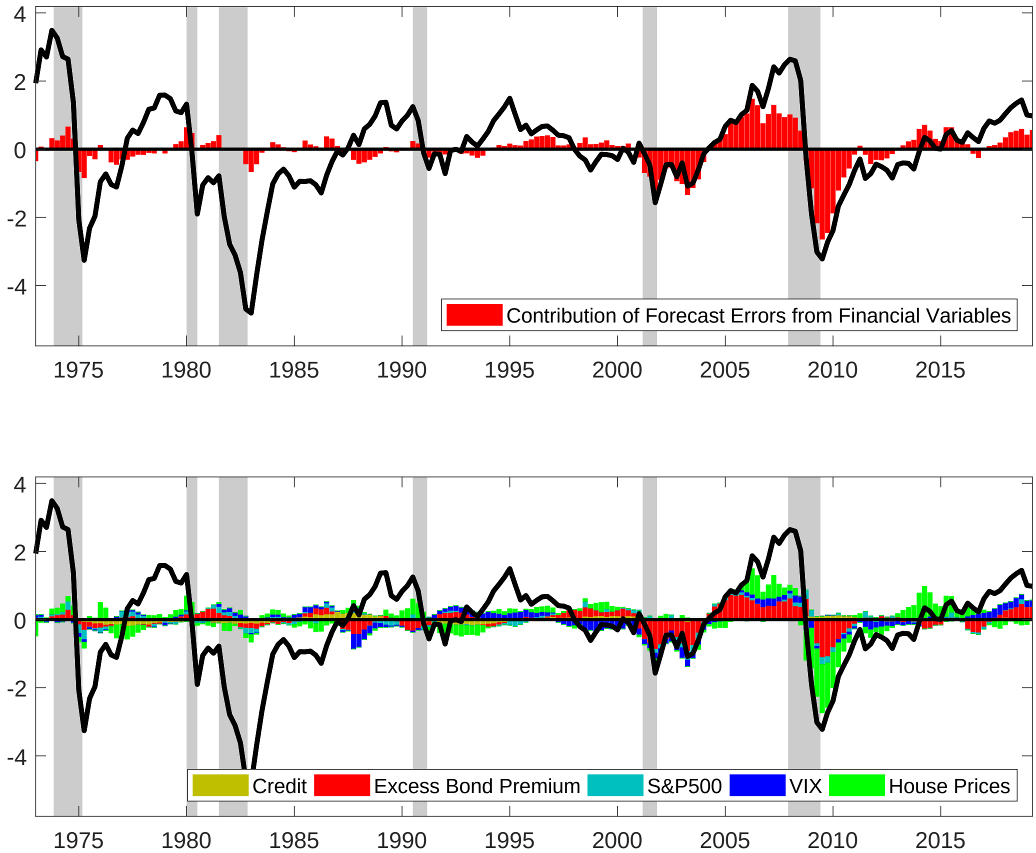

Figure 4: Informational decomposition of the output gap. Solid line plots the estimated output gap. Output gap is measured in percent deviation from the trend. Grey shaded areas indicate NBER recessions. In the top panel, the bar represents the total contribution from the BVAR forecast errors from five financial variables (credit, the excess bond premium, the S&P 500, the VIX index, and the house price) for the output gap. The bars in the bottom panel presents the contribution from each of the aforementioned variables.

Figure 4 presents the informational decomposition for the estimated output gap calculated using Equation (6). The contributions are calculated from the forecast errors of five financial variables in our BVAR system; credit, the excess bond premium, stock prices, the VIX, and house prices. The top panel sums the contribution, while the bottom panel reports the individual shares.

We document two key observations. First, the role of financial variables seems to have increased over time. While their impact is rather negligible before the year 2000, it has increased afterward. It is a more open question whether, at the end of the sample, the role of the financial variables associated with the output gap has returned to the more negligible role pre-2000. Second, financial variables have been particularly important during times where one would *a priori* attach a role for financial factors as having been important for the business cycle. For example, we find an important role for financial variables on the output gap in periods of

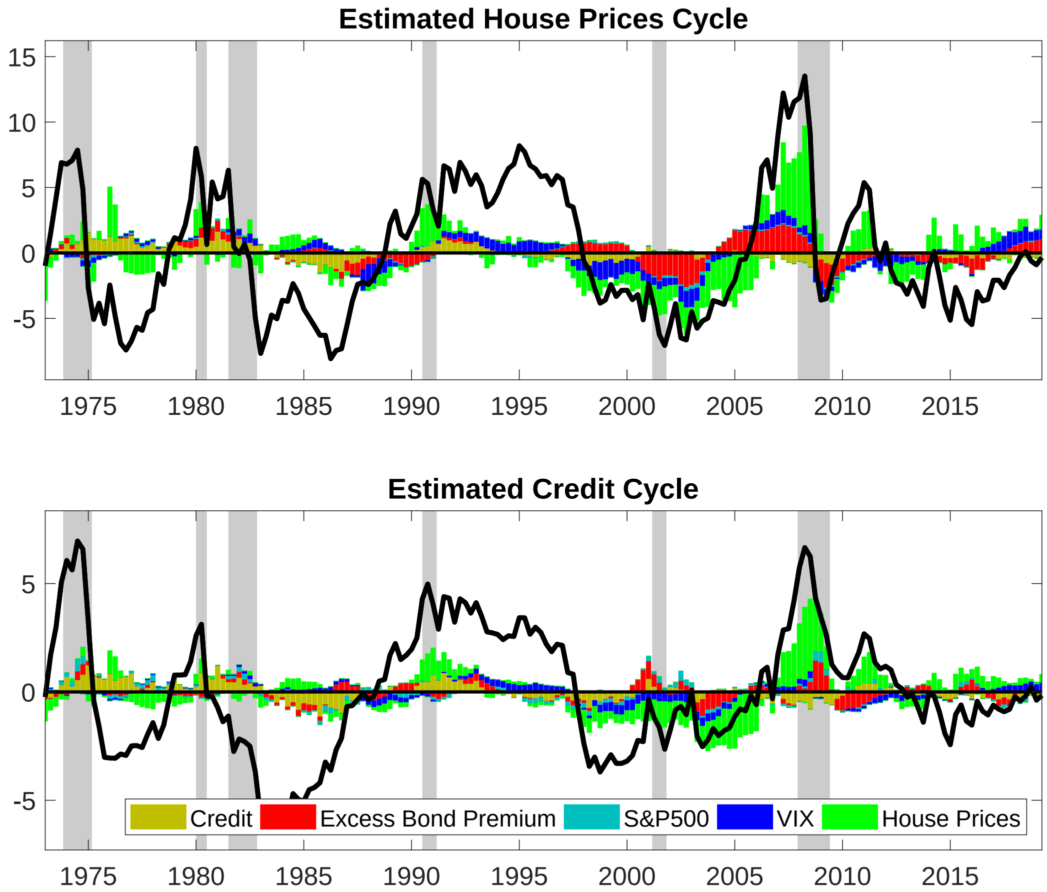

Figure 5: Informational decomposition of the estimated financial cycles. Solid line plots the estimated financial cycle, measured in percent deviation from trend. Grey shaded areas indicate NBER recessions. The bar represents the total contribution from the BVAR forecast errors from five financial variables (credit, the excess bond premium, stock prices, house prices, and the VIX) on the financial cycles for the output gap.

financial stress, such as the burst of the dot-com bubble and the outbreak of the financial crisis as well as during the built-up of large financial imbalances as seen during the 2000s.

Turning to the individual financial variables in the bottom panel of Figure 4, we find that of all the financial variables, the forecast errors from the excess bond premium and house prices contribute sizably to the output gap. As described previously, the excess bond premium reflects the risk-bearing capacity of financial intermediaries, and thus can be seen as a measure of excess credit (see Gilchrist and Zakrajšek, 2012). That we find a prominent role for information contained in the excess bond premium despite the inclusion of several other financial variables suggests that the link of how financial factors affected the output gap in the 2000s is likely linked to excess credit. Our evidence is consistent with an interpretation that excess credit contributed substantially to the overheating of the U.S. economy before the financial crisis. House prices have also been shown to play an important role in providing information about the output gap, which is consistent with Leamer's (2007) observation that "housing is the business cycle". In particular, house prices contribute to the positive output gap in the 2000s, and also explain a large share of the negative output gap in the period during and just after the 2008/09 recession. The latter is a finding that is perhaps less surprising given it is well known that the housing bust played a big role in the 2008/09 recession.

Figure 5 presents the informational decomposition of the financial variables for both the credit and house price cycles. We once again do find a larger role for the financial variables during the 2000s, relative to the pre-2000s. We also document during the 2000s, a prominent role for the excess bond premium during the booming housing cycle and a prominent role for housing in the booming credit cycle. Summing up, our results from the informational decomposition are consistent with the interpretation by Justiniano, Primiceri, and Tambalotti (2019) that excess credit supply was linked to the booming housing market. These factors were at the heart of the overheating business cycle before the Great Recession.

We stress though that the interpretation from the informational decomposition is not causal as the information contained within the financial variables could originate from shocks outside the financial sector. Nonetheless, that house prices and the excess bond premium contain information for the output gap is as a useful starting point for our causal analysis in the next subsection.

4.2 Structural Analysis

As stressed in the previous subsection, while useful, the informational decomposition cannot attribute causality. While the informational decomposition only requires fitting a standard BVAR on a set of financial and macroeconomic variables, quantifying causal effects requires explicit identifying assumptions. Moreover, while the broader literature does provide guidance on how to disentangle structural financial shocks, it is unfortunate that even if several contributions share the label “financial shock”, what each seeks to identify and isolate may not conceptually line up with one another perfectly.

While we are more agnostic as to the precise definition of a financial shock, a broad element of what we seek to isolate is exogenous variation emanating from the financial sector. Our approach is thus to draw guidance from three existing identification schemes to identify financial shocks, so that our conclusions are less sensitive to any particular identification scheme. While we will elaborate on the details of each subsequently, the three identification schemes we will employ are a Cholesky decomposition, a penalty function approach that we take guidance from Caldara, Fuentes-Albero, Gilchrist, and Zakrajšek (2016), and a sign restriction approach inspired by Furlanetto, Ravazzolo, and Sarferaz (2019). The first two identification schemes largely rely on exploiting variation in the excess bond premium for identification. As the excess bond premium is an indicator of the risk-bearing capacity of financial intermediaries, the identified financial shock in these settings is conceptually closer to exogenous variation in the financial sector’s ability to provide credit. This is also consistent with the loosening and tightening of the credit constraint, a mechanism that is very much at the heart of the financial friction/financial accelerator literature (e.g. Bernanke and Gertler, 1989; Bernanke, Gertler, and Gilchrist, 1999). The sign restriction approach by Furlanetto, Ravazzolo, and Sarferaz (2019) on the other hand, is more agnostic as to what is a financial shock. Furlanetto, Ravazzolo, and Sarferaz (2019) define and identify a financial shock as a boom in investment and stock prices. Despite conceptual differences, many of the empirical results by Furlanetto, Ravazzolo, and Sarferaz (2019) are consistent with the type of financial shock implied by the other two identification schemes, suggesting that the distinction implied by the approach by Furlanetto, Ravazzolo, and Sarferaz (2019) and that of the other two approaches is, at least empirically, perhaps not as sharp as it appears at first glance. From Equation (8), the identification of a financial shock amounts to finding a column of the $\mathbf{A}$ matrix.

Cholesky Decomposition We follow Gilchrist and Zakrajšek (2012), by identifying shocks to the excess bond premium through ordering fast-moving variables after slow-moving variables. Specifying a slow-moving and fast-moving block is a reasonably common strategy for using the Cholesky decomposition for identification within a system that features both financial and macroeconomic variables (e.g., see Christiano, Eichenbaum, and Evans, 1999; Bernanke, Boivin, and Eliasz, 2005). Within our setup, this amounts to ordering the excess bond premium after slow-moving variables such as GDP, investment, etc, and before fast-moving, which are often the financial market variables, such as stock prices. This assumes that slow-moving variables do not react contemporaneously to the financial shock and shocks in the fast-moving block. At the same time, shocks in the fast-moving block do not have a contemporaneous effect on the excess bond premium. As Gilchrist and Zakrajšek (2012) point out, the identified shock is a shock to the excess bond premium which is orthogonal to other shocks in the economy. We will interpret this shock as a structural financial shock within this setting. This identification strategy relies on utilizing (orthogonal) variation in the excess bond premium to identify financial shocks. While we are aware of possible misgivings against the zero restrictions implied by the Cholesky decomposition, we add that this is a standard identification strategy used in the wider literature (e.g., Gilchrist, Yankov, and Zakrajšek, 2009; Walentin, 2014), which at least provides a first pass at identifying a financial shock before moving on to other identification strategies.

Penalty Function Drawing inspiration from Caldara, Fuentes-Albero, Gilchrist, and Zakrajšek (2016), we consider a penalty function approach in order to identify financial shocks. This entails using a penalty function to identify the financial shock by solving for the shock to maximize the variance of the excess bond premium over the first 4 quarters.[11] Like the Cholesky decomposition, the penalty function approach also relies on orthogonal variation in the excess bond premium to identify financial shocks. The penalty function approach though, relaxes many of the zero restrictions one utilizes in the Cholesky decomposition, which one may view as being more tenable.

[11] We note that our approach differs subtly from Caldara, Fuentes-Albero, Gilchrist, and Zakrajšek (2016), who use the penalty function approach to distinguish financial shocks from uncertainty shocks. Unlike them, we do not attempt to identify an uncertainty shock. In their approach, the choice of whether to first identify the financial or uncertainty shock may matter. Therefore, strictly speaking, our approach is only identical to Caldara, Fuentes-Albero, Gilchrist, and Zakrajšek (2016) in the setting where a financial shock is identified *before* the uncertainty shock.

Real GDP	-	Housing starts	-	M1	NA
CPI Inflation	-	PPI Inflation	NA	M2	NA
Federal Funds Rate	-	Hours Worked	NA	Credit	-
Employment	-	Personal Income	-	S&P 500	-
Consumption	-	10-year rate	NA	Real Energy Prices	NA
Industrial Production	-	Productivity	NA	VIX	+
Capacity Utilization	-	Investment	-	Property prices	-
Unemployment	-	Excess Bond Premium	+	Investment/GDP ratio	-

Table 1: The table describes the sign restriction on each variable in order to identify the financial shock. NA indicates that the response of the variable to a financial shock is left unrestricted. The sign restriction is restricted to only hold upon impact.

Sign Restrictions We also consider sign restrictions to identify the financial shock. The identification of the financial shock closely mimics Furlanetto, Ravazzolo, and Sarferaz (2019). Furlanetto, Ravazzolo, and Sarferaz (2019) derive their sign restrictions by characterizing a financial shock as a shock which induces an investment and stock market boom/slump. Guided by Furlanetto, Ravazzolo, and Sarferaz (2019), Table 1 summarizes the sign restrictions to identify a financial shock. The signs are normalized where a positive financial shock leads to investment and stock market slumps. The identification strategy also imposes investment to fall more than GDP in response to a positive financial shock as an investment boom/slump forms part of their identification strategy.[12] All the sign restrictions hold only upon impact. While we do not identify more than a single financial shock, guided by Antolín-Díaz and Rubio-Ramírez (2018), we impose a narrative in addition to the sign restrictions. In smaller systems, identifying more shocks, as Furlanetto, Ravazzolo, and Sarferaz (2019) do, can add more information and yield sharper inference. Because the size of our system makes identifying more shocks a more challenging endeavor, we use the narrative sign restriction, as opposed to identifying more shocks, as a means of introducing additional information to help with the identification of the financial shock. The events we have in mind are the collapse of Lehman in September 2008 and credit freezing all round in 2008Q4. We therefore implement a narrative sign restriction that the sign of the financial shock is positive in both 2008Q3 and 2008Q4. In addition, the financial shock is the overwhelming driver of the increase in the excess bond premium between 2008Q3 to 2008Q4. This is akin to what Antolín-Díaz and Rubio-Ramírez (2018) refer to as Type B restrictions. The goal of the narrative sign restriction is to introduce more information on the identification problem.

[12]Note that as Furlanetto, Ravazzolo, and Sarferaz (2019) normalize the financial shock to induce an investment boom, so positive financial shocks in their identification causes the investment to GDP ratio to rise. However, since we normalize a positive financial shock to induce an investment slump, to make the sign of the financial shock consistent with our other two identification strategy, investment will fall more than GDP in response to a positive financial shock, so the investment to GDP ratio falls.

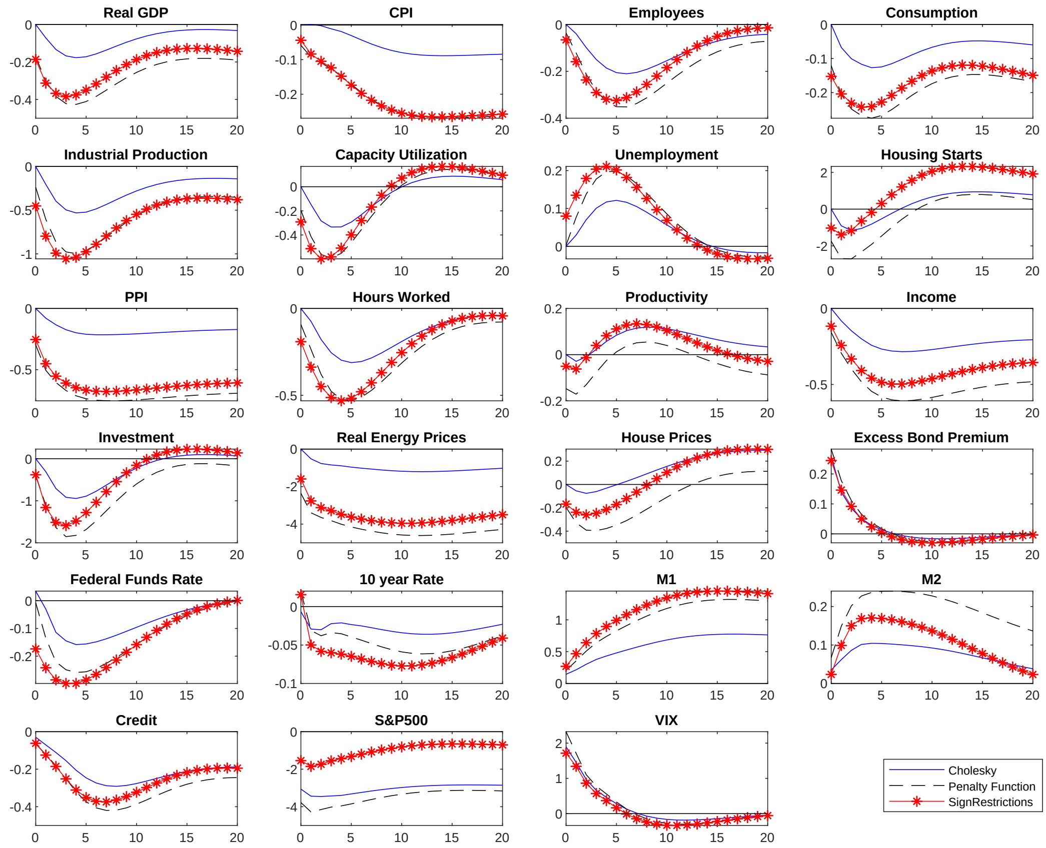

Figure 6: Estimated impulse response function to a one standard deviation financial shock. All impulse response functions are calculated from the posterior mean parameters of the estimated BVAR. The sign restriction impulse response function is chosen from a representative model using the Fry-Pagan median target from 1000 rotations of the posterior mean parameters which satisfy the sign and narrative restrictions. The x-axis is in terms of quarters after the shock. Capacity utilization, unemployment, Federal funds rate, and 10 years rate are in terms of percentage point deviation. VIX and excess bond premium are in their natural units. All other variables are in terms of percentage change.

Results from the Structural Analysis

We first present the impulse response functions to our identified financial shock. Figure 6 presents the impulse response functions to a one standard deviation structural financial shock for all three identification strategies. All the impulse response functions are conditional on the posterior mean of the BVAR estimates and we report the impulse response functions of the sign restrictions using the Fry and Pagan (2011) median target approach.[13] We explicitly account for parameter uncertainty by presenting posterior credible sets of all three identification

[13]There is a known issue of representativeness of the impulse response function as sign restrictions only identify a set and do not provide a unique solution (see Fry and Pagan, 2011). We report the Fry and Pagan (2011) median target approach here from 1000 admissible solutions conditional on the posterior mean parameters. This is just for illustrative purposes as we wish to just compare the sign of the responses of all three identification schemes.

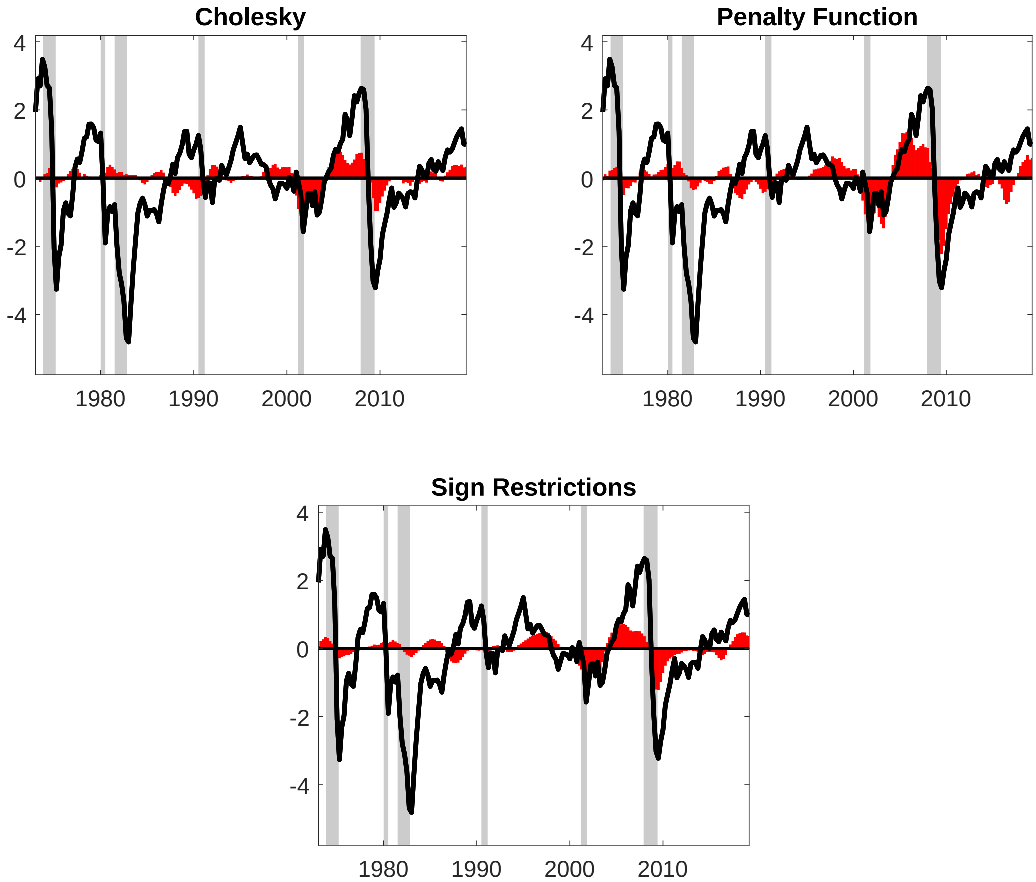

Figure 7: Contribution of the financial shock to the estimated output gap. The solid line represents the estimated output gap. The output gap is measured in percent deviation from trend. Grey shaded areas indicate NBER recessions. The title refers to the different identification schemes. The bars represent the contribution of financial shocks to the estimated output gap. The contribution from the sign restriction approach is averaged across draws which satisfy the sign and narrative restrictions.

strategies in Section D of the online appendix. On the estimated impulse response functions, we draw attention to two key points. First, while the sign restrictions do impose the responses of particular variables to financial shocks as per Table 1, the responses of all variables to a financial shock identified using all three identification strategies have the same sign. The effect of a financial shock is therefore qualitatively similar across all three identification strategies, and the difference are largely confined to the extent of the magnitude of the responses. Therefore, our results should provide at least some confidence that all three identification strategies are providing reasonable estimates of the effect of financial shocks. Second, while we restrict the sign of prices and GDP to fall in the sign restrictions, we obtain similar results with the other two identification strategies where the sign is left unrestricted. Therefore, there is at least consistent evidence that the identified financial shock in all three settings is an aggregate demand shock, or at least one where the effect on the aggregate demand side of the economy dominates.

Figure 7 presents the contribution of the financial shock to the output gap for all three different identification schemes. These shares are calculated conditional on the posterior mean of the BVAR parameters in and equivalent to reporting Equation (9) across different $\mathbf{A}$'s.[14] While the share of financial shocks on the output gap differs between the three identification strategies, we highlight two key similarities across the three different strategies. First, the share of financial shocks tend to be much smaller pre-2000s, but appear to be much larger since the 2000s. Second, financial shocks appear to contribute positively to the output gap in the 2000s before the Great Recession, and then played a large role in the negative output gap during the Great Recession. We also note that financial shocks also played a sizable negative role in the 2000/01 recession, which was associated with the bust of the dot-com bubble.

It is reassuring that even without a consensus on how to identify financial shocks, three different identification strategies provide a consistent account of how financial shocks drive the business cycle. Given the prominent role of financial shocks in contributing sizably to a large and positive output gap before the 2008/09 recession, our interpretation is that loose credit conditions originating from financial shocks in the 2000s likely fueled a boom in the business cycle which later led to the bust.

4.3 Discussion

The results of how financial shocks affect the business cycle are consistent with the more reduced form informational decomposition. In particular, the financial variables contributed more since the 2000s and played a large role in the overheating of the business cycle which is consistent with the prominence of the contribution of the more structural identified financial shock. That the excess bond premium played a prominent role in the more reduced form informational decomposition during the 2000s pre-Great Recession also suggests that the 2000s boom was driven by loose credit conditions fueled by the financial sector's appetite to undertake greater risk.

We also relate our work to contributions in the wider literature to construct both "finance-neutral" output gaps (e.g. Borio, Disyatat, and Juselius, 2017), or considering the output gap as the difference between actual output and a counterfactual in the absence of financial frictions (e.g. Furlanetto, Gelain, and Sanjani, 2020). While our work has a flavor of both, we discuss more broadly the differences and similarities to this body of work. When considering "finance-neutral" output gaps, Borio, Disyatat, and Juselius (2017) state that traditional output gap estimates are inflation centric, and thus they consider information from financial variables to

[14] For the sign restriction results, we averaged over the 1000 rotations which satisfies the sign and narrative restrictions conditional on the posterior mean parameters. Our approach to averaging across the admissible rotations is similar to Forbes, Hjortsoe, and Nenova (2018), who averaged across the different solutions when calculating their historical decomposition. Here, we average across solutions as the average of the contribution from all the, identified or unidentified, shocks across the admissible sign restriction solutions sums up to the output gap. This is because the effect of the shock and the variance of the shock automatically adjusts for each of the solutions or each $\mathbf{A}$ in Equation (9). We note this is a subtlety separate issue when we report impulse response functions in Figure 6, and why we did not average over the impulse response function. In this case, this would entail averaging over financial shocks that have different estimated variance.

estimate the transitory component of real GDP. Within our framework, our output gap has no notion of being inflation or finance centric. Instead, following on from the discussion by Evans and Reichlin (1994) and Morley and Wong (2020), when conducting a multivariate BN decomposition, any variable that is relevant for forecasting output growth is relevant for the output gap. However, the extant evidence that financial variables have proven to be relevant for forecasting output growth (see Faust, Gilchrist, Wright, and Zakrajšek, 2013) suggests the inclusion of financial variables for the estimation of the output gap within our framework. From this perspective, one can subtract the role of financial variables from the output gap in Figure 4 and regard this as the output gap if one did not incorporate information from financial variables, though we caution that this alternative output gap is not "inflation"-centric in any sense, and could best be described as the "non-financial" output gap. Moreover, this "non-financial" output gap also does not account for the fact that financial and macro variables are correlated, and so omitting financial information would merely shift some of the role played by the financial variables to macroeconomic variables, because an informational decomposition is not in any sense structural or causal. Despite these conceptual differences and the obvious caveats, our account with the "non-financial" output gap corresponds with what has been found in the "finance-neutral" output gap literature (e.g. Borio, Disyatat, and Juselius, 2017) during the 2000s. In particular, the findings from the finance-neutral output gap work suggest that the finance-neutral output gap is much larger than the inflation-centric output gap in the 2000s, consistent with the sizable role of the financial variables for the output gap in our informational decomposition. Even so, we stress while our results are consistent with regards to the view that the 2000s coincides with the perspective of the finance-neutral work, we do find a very small contribution of financial variables to the output gap pre-2000s, which suggests that any distinction of our output gap and a non-financial output gap pre-2000s is probably less relevant.

The more structural approach taken by Furlanetto, Gelain, and Sanjani (2020) views the output gap as reflecting inefficiencies arising from frictions, in the tradition of New-Keynesian DSGE models. Trend output is the counterfactual level of output in the absence of these frictions and the output gap is the difference between actual and the counterfactual output. Conceptually, the frictions in their setup are propagation mechanisms and relevant for *all* shocks. A direct comparison relative to the more structural approach of Furlanetto, Gelain, and Sanjani (2020) is naturally challenging, as a fully-specified DSGE model requires one to be explicit about the different frictions in the model. Even so, we note that a key result in their paper is that the inclusion of financial frictions implies a more positive output gap in the 2000s and before the Great Recession, consistent with our key result that the financial sector played an important role in overheating the business cycle pre-Great Recession.

5 Does the Financial Cycle Lead the Business Cycle or Vice Versa?

So far, the analysis has been focused on estimating the business and financial cycle, as well as quantifying how important financial factors have been in driving the U.S. business cycle. In this section, we focus on the links between the financial and business cycle. In particular, there is an active body of work which has been particularly interested in characterizing features on the comovement between the financial and business cycle to understand the links between them (e.g. Claessens, Kose, and Terrones, 2012; Aikman, Haldane, and Nelson, 2015; Rünstler and Vlekke, 2018; Oman, 2019).

As cross-correlations have traditionally played an important role in understanding the links between the cyclical components of different macroeconomic variables, we now adapt our empirical framework to understand cross-correlations. In particular, we are interested in shedding light on issues such as whether the financial cycle leads the business cycle or vice versa. From Equations (3) and (5), we know from Morley (2002) that $\mathbf{F}(\mathbf{I}-\mathbf{F})^{-\mathbf{1}}(\mathbf{X_t}-\boldsymbol{\mu})$ contains the estimated BN cycles. Following Kamber, Morley, and Wong (2018), the following can be used to calculate the variances of the estimated BN cycles

$$\boldsymbol{\Psi} = \mathbf{F}(\mathbf{I}-\mathbf{F})^{-\mathbf{1}}\boldsymbol{\Omega}[(\mathbf{I}-\mathbf{F})^{-\mathbf{1}}]'\mathbf{F}', \tag{13}$$

where $\boldsymbol{\Omega}$ is the variance of $\mathbf{X_t}$ and $vec(\boldsymbol{\Omega}) = [\mathbf{I}-\mathbf{F}\otimes\mathbf{F}]^{-\mathbf{1}}vec(\boldsymbol{Q})$, where

$$\boldsymbol{Q} = \begin{bmatrix} \boldsymbol{\Sigma} & \mathbf{0} & \dots \\ \mathbf{0} & \mathbf{0} & \ddots \\ \vdots & \ddots & \ddots \end{bmatrix}. \tag{14}$$

It follows that elements of $\boldsymbol{\Psi}$ will contain the cross-covariance between any pair of $c_{i,t}$ and $c_{j,t-m}$ where $i,j \in \{1,2,\dots,K\}$ and $m \in \{0,1,2,\dots\}$.[15] It is then straightforward to normalize $\boldsymbol{\Psi}$ into a correlation matrix to obtain the cross-correlation of $c_{i,t}$ and $c_{j,t-m}$, where $\Delta y_{i,t}$ and $\Delta y_{j,t}$ are respectively in the k^{th} and l^{th} position in $\mathbf{x_t}$, and

$$corr(c_{i,t}, c_{j,t-m}) = \mathbf{s_k}\boldsymbol{\psi}\mathbf{s}'_{\mathbf{nm+l}}, \tag{15}$$

where $\boldsymbol{\psi}$ is the correlation matrix associated with $\boldsymbol{\Psi}$. More precisely, Equation (15) can be used to quantify objects such as the correlation of the output gap with the credit cycle four quarters ago and vice versa, providing a richer framework to understand the interaction between the financial and business cycle. $\boldsymbol{\psi}$, though, only contains the unconditional cross-correlations between measures of the business and financial cycle. It is straightforward to modify this cross-correlation conditional on a financial shock. Let $\boldsymbol{\alpha}$ be the column of the matrix $\mathbf{A}$ which

[15] If one fitted a VAR(p) and cast it into the form implied by Equation (3), we can obtain cross-covariances up to $p-1$. To calculate the cross-covariances for cycles where $m \geq p$, one will still estimate the same VAR(p), but subsequently just augment the state vector $(\mathbf{X_t}-\boldsymbol{\mu})$ in Equation (3) with longer lags, as well as input appropriate entries in $\mathbf{F}$ to calculate $\boldsymbol{\Psi}$.

identifies the financial shock in our exercise. If we modify Equation (14) such that

$$\tilde{\boldsymbol{Q}} = \begin{bmatrix} \boldsymbol{\alpha}\boldsymbol{\alpha}' & \mathbf{0} & \dots \\ 0 & 0 & \ddots \\ \vdots & \ddots & \ddots \end{bmatrix} \tag{16}$$

and substitute $\tilde{\boldsymbol{Q}}$ for $\boldsymbol{Q}$ at every step of the calculation of $\boldsymbol{\Psi}$, we can now obtain the cross-correlations of the business and financial cycle *conditional* on a financial shock. Unconditional correlations are the outcome of various shocks, and within our framework, the financial and business cycle are just outcomes of the various, identified and unidentified, shocks. The characterization of conditional cross-correlations adds a further dimension to the analysis. In particular, while the unconditional cross-correlations are important to characterize, these may have little to do with financial shocks. Unconditional cross-correlations, like our informational decomposition exercise, also do not allow us to make causal statements. Characterizing conditional cross-correlation allows our framework to make a causal link to how financial shocks can drive particular lead-lag relationships between the business and financial cycle.

Table 2 presents unconditional correlations, as well as the unconditional and conditional 4-quarter cross-correlations between our estimates of the output gap, credit cycle and house price cycle, which we take as measures of the business and financial cycle. We also present the contemporaneous correlations between the different estimated cycles. We first focus on the top panel, which presents the unconditional cross-correlations. All entries are positive, which suggests that unconditionally, we expect booms in the financial cycle to be followed by booms in the business cycle and vice versa. While it should not be surprising that booms in the financial cycle lead to booms in the business cycle, unconditionally, this provides very little rationale for any form of regulation or macroprudential regulation to restrain credit or even house prices. A boom in the credit cycle is followed by a boom in the house price cycle and vice versa, which is consistent with the reinforcing dynamics of credit and housing booms, as documented by Jordà, Schularick, and Taylor (2015).

However, the picture changes somewhat once we condition these correlations on a financial shock, as per Equation (16). We first condition on a financial shock identified through our Cholesky and penalty function identification since these identification techniques provide a unique solution to the identification of the financial shock. For both the Cholesky and penalty function identification, we observe that once we condition on a financial shock, the credit cycle lagged 4 quarters is now strongly negatively correlated with the output gap and the house price cycle. Because sign restrictions do not point identify the financial shock, but instead produce a set of admissible solutions (see Fry and Pagan, 2011), to check for whether our sign restriction identification produces conditional correlations in line with our other two identification strategies, we count the proportion of conditional correlations from the various sign restriction solutions which are negative, and thus switch sign from the unconditional correlation.[16] This

[16]Note that the unconditional correlation is the same across all the sign restricted solutions as this quantity is derived from the same reduced form.

Unconditional Cross-Correlations

		Correlations		
		Output Gap	**House Price Cycle**	**Credit Cycle**
	Output Gap	1		
	House Price Cycle	0.57	1	
	Credit Cycle	0.54	0.79	1

		Contemporaneous (t)		
		Output Gap	**House Price Cycle**	**Credit Cycle**
Lagged 4	**Output Gap**	0.35	0.45	0.71
Quarters	**House Price Cycle**	0.28	0.87	0.73
($t-4$)	**Credit Cycle**	0.10	0.62	0.84

Conditional Cross-Correlations

Cholesky

		Contemporaneous (t)		
		Output Gap	**House Price Cycle**	**Credit Cycle**
Lagged 4	**Output Gap**	0.24	0.08	0.77
Quarters	**House Price Cycle**	0.70	0.91	-0.06
($t-4$)	**Credit Cycle**	-0.58	-0.83	0.66

Penalty Function

		Contemporaneous (t)		
		Output Gap	**House Price Cycle**	**Credit Cycle**
Lagged 4	**Output Gap**	0.02	-0.09	0.92
Quarters	**House Price Cycle**	0.58	0.69	0.54
($t-4$)	**Credit Cycle**	-0.60	-0.69	0.57

Sign Restrictions, percentage of negative correlations

		Contemporaneous (t)		
		Output Gap	**House Price Cycle**	**Credit Cycle**
Lagged 4	**Output Gap**	0	0	0.1
Quarters	**House Price Cycle**	0	0	3.3
($t-4$)	**Credit Cycle**	71.5	55.1	4.3

Table 2: Unconditional and conditional cross-correlations.

is presented in the bottom panel of Table 2. We observe a sign switch in the majority of our sign restricted solutions for the conditional correlation of the lagged credit cycle on the house price cycle, and more importantly, for the output gap. Therefore, we conclude that a majority of our sign restricted identified solutions are in line with the sign switch that we document for the Cholesky and penalty function approach.

Summing up, a key conclusion is that *unconditionally*, the credit cycle and output gap are strongly positively correlated, with the contemporaneous correlation much stronger (0.54) relative to the correlation between the lagged credit cycle and the contemporaneous output gap. While we find evidence that the credit cycle lagged 4 quarters is unconditionally positively correlated with the contemporaneous output gap, this correlation comes out rather weak (0.10). However, when we condition on a financial shock, the lagged credit cycle is strongly negatively correlated with the output gap (about -0.6). This conclusion is robust to how we identify the financial shock with either the Cholesky identification or the penalty function approach, and also a wide variety of the set of sign-restricted solutions that we obtain.

These results suggest a more nuanced view of how the business cycle interacts with the financial cycle, or more specifically the credit cycle. In general, one should not expect all credit booms to eventuate in busts. The unconditional correlations would suggest that the *average* credit boom does lead to a business cycle boom. Instead, we uncover that it is booms in the credit cycle which are driven by financial shocks that do eventuate in busts in the business cycle, given the negative conditional cross-correlation. Recall that the identified financial shock is linked to notions of excess credit emanating from the financial sector. Our results are consistent with Jordà, Schularick, and Taylor (2013), who also show that credit booms characterized by their measure of excess credit are the ones that lead to busts, or in their language, "when credit bites back". We also note that the conditionality of whether credit booms do end in business cycle busts has also been shown elsewhere (e.g., see Krishnamurthy and Muir, 2017; Coimbra and Rey, 2017; Gorton and Ordoñez, 2019), where different notions of conditionality, such as the level of the interest rate or financial fragility have been proposed. While one should exercise caution in making direct comparisons, especially with regards to whether these conditions coincide with our notion of financial shocks, our work at least provides empirical results consistent with the notion which attach a high degree of conditionality on whether credit booms end in business cycle bust. In our case, financial shocks that drive credit cycle booms appear to be the ones that lead to business cycle busts.

Our results suggest that macroprudential policy targeted at crude measures of credit growth may be too blunt of an instrument. In particular, we find that it is booms in the financial cycle induced by financial shocks, perhaps through excess credit, that eventuate into future busts in the business cycle. Our results would suggest that macroprudential policy should not curb all credit booms, especially if the goal of such policy is to prevent future busts in the business cycle, but perhaps those that originate from financial shocks. Given the way the financial shock is identified in our three settings, our work suggest that it is credit booms that originate from the financial sector undertaking excessive risk which result in excess credit that macroprudential policy should target.

6 Addressing the Possibility of a Slowdown in Real GDP Growth

We briefly address the possibility of a break in the drift of real GDP as this has a first-order implication for the measurement of the business cycle, and so may affect our results. When we

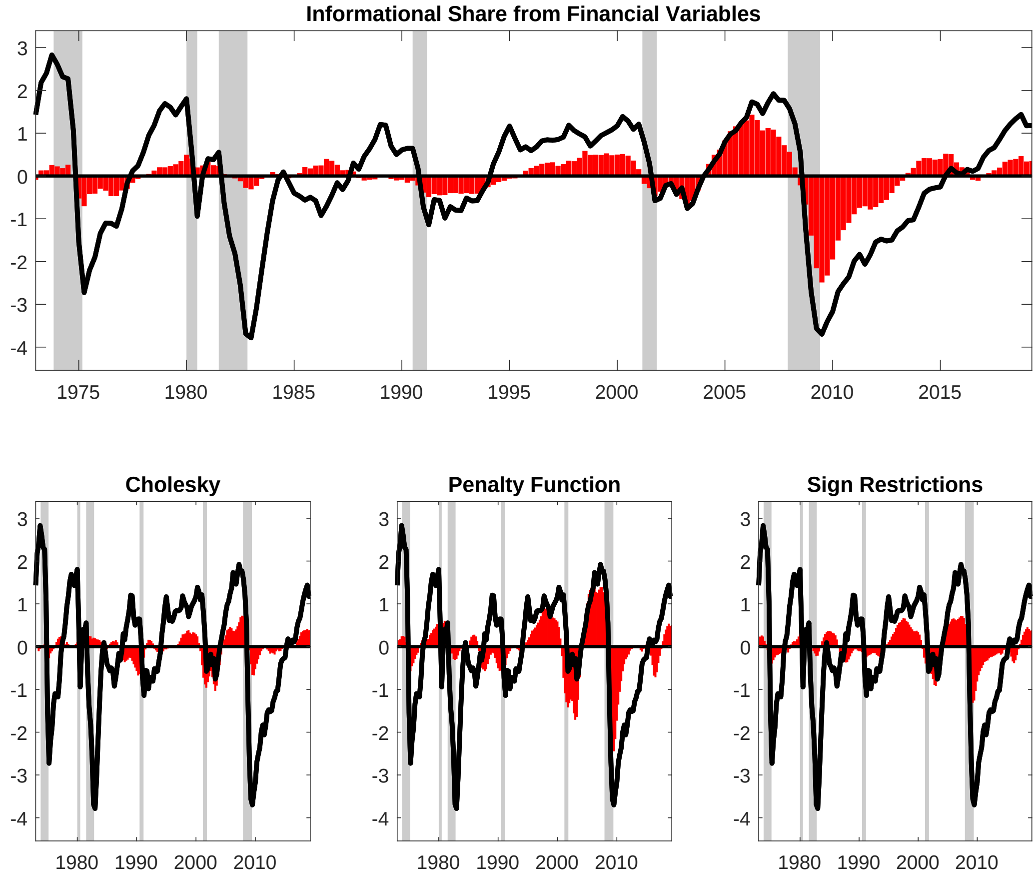

Figure 8: Allowing for a slowdown in real GDP growth in 2006Q2. The solid line is the estimated output gap. Output gap is measured in percent deviation from trend. Grey shaded areas indicate NBER recessions. The top panel presents the contribution of the forecast errors from the financial variables to the estimated output gap. The bottom three panels presents the contribution of financial shocks to the estimated output gap. The title refers to the different identification scheme. The contribution from the sign restriction approach is averaged across draws which satisfy the sign and narrative restrictions.

set up our baseline specification, we could not reject the possibility of a break in the drift of real GDP with a Bai and Perron (2003) test. However, this result of being unable to date a break is sensitive to how we adjusted for the standard errors when testing for the break. An alternative specification can date a break in 2006Q2, which is consistent with wider work (e.g. Berger, Everaert, and Vierke, 2016; Eo and Morley, 2017; Kamber, Morley, and Wong, 2018) dating a slowdown in GDP growth just before the Great Recession.[17] We note that the inherent uncertainty of whether, and if so when, a break in the drift in U.S. real GDP has occurred is not entirely surprising given the mixed evidence on the issue (see, e.g. Check and Piger, 2018).

We therefore allowed for a break in the drift of real GDP in 2006Q2 to explore the robustness of our results. Figure 8 presents the key results when considering a break in the drift in 2006Q2.

[17] To be precise, our baseline specification for the Bai and Perron (2003) test allows for heteroskedasticity and autocorrelation consistent (HAC) standard errors, which cannot date a break with the usual degree of statistical significance. If we do not allow for HAC standard errors, we will date a break in 2006Q2.

The top panel of Figure 8 presents the informational decomposition of the estimated output gap, and the role of the forecast errors of the financial variables, similar to the analysis in the top panel of Figure 4. The bottom panel presents the shock decomposition, similar to Figure 7, under our three identification schemes. We note that the estimated output gap now features a deeper output gap during the Great Recession, suggesting that allowing for a break in the drift of real GDP leads our model to interpret a greater proportion of the fall of real GDP during the Great Recession as being transitory. Nonetheless, even allowing for a break in the drift, both the informational decomposition and structural analysis are still consistent with our key results that financial factors appear more important post-2000, and played an important role in the overheating of the business cycle pre-Great Recession.

We therefore conclude that even if a slowdown in real GDP growth was a relevant concern, all the key results of our paper go through.

7 Conclusions

Utilizing new econometric tools, we study how financial factors affect the business cycle. Using a standard BVAR, our analysis can explicitly model whether the business cycle was overheating in the 2000s to understand the role of financial factors in driving the underlying fluctuations in the business cycle, as well as to characterize how the business cycle interacts with the financial cycle within a unified framework.

We find that the role of financial factors in driving the business cycle appears to be much smaller before the 2000s, but the role of financial factors appear to be much larger since the 2000s. Much of these fluctuations appear to be linked to the credit market. We make this conclusion by observing that in reduced form, the excess bond premium introduced by Gilchrist and Zakrajšek (2012) can explain the overheating business cycle in the 2000s before the Great Recession. When we explicitly identify a financial shock, we similarly find that financial shocks positively contributed to the output gap in the 2000s before the Great Recession. Also consistent with our reduced-form analysis, we find that the role of financial shock is much smaller before the 2000s.

We also uncover evidence that credit cycles are linked to the output gap, though the relationship between them would appear to be a more nuanced one. In particular, we show that credit cycle booms do not always lead to business cycle bust; only those linked to financial shocks and some notion of excess credit do. One implication of our findings is that macroprudential policy would need to distinguish between the underlying drivers of credit cycle booms rather than relying on simple rules of thumbs which prescribe unconditionally curbing all credit cycle booms.

References

Adrian, T., and H. S. Shin (2010): "Financial Intermediaries and Monetary Economics," in *Handbook of Monetary Economics*, ed. by B. M. Friedman, and M. Woodford, vol. 3 of *Handbook of Monetary Economics*, chap. 12, pp. 601–650. Elsevier.

Aikman, D., A. G. Haldane, and B. D. Nelson (2015): "Curbing the Credit Cycle," *Economic Journal*, 125(585), 1072–1109.

Albuquerque, R., M. Eichenbaum, D. Papanikolaou, and S. Rebelo (2015): "Long-run bulls and bears," *Journal of Monetary Economics*, 76(S), 21–36.

Andrews, D. W. K. (1993): "Tests for Parameter Instability and Structural Change With Unknown Change Point," *Econometrica*, 61(4), 821–856.

Antolín-Díaz, J., and J. F. Rubio-Ramírez (2018): "Narrative Sign Restrictions for SVARs," *American Economic Review*, 108(10), 2802–2829.

Bai, J., and P. Perron (2003): "Computation and Analysis of Multiple Structural Change Models," *Journal of Applied Econometrics*, 18(1), 1–22.

Banbura, M., D. Giannone, and L. Reichlin (2010): "Large Bayesian Vector Auto Regressions," *Journal of Applied Econometrics*, 25(1), 71–92.

Baxter, M., and R. G. King (1999): "Measuring Business Cycles: Approximate Band-Pass Filters For Economic Time Series," *The Review of Economics and Statistics*, 81(4), 575–593.

Berger, T., G. Everaert, and H. Vierke (2016): "Testing for time variation in an unobserved components model for the U.S. economy," *Journal of Economic Dynamics and Control*, 69(C), 179–208.

Bernanke, B., and M. Gertler (1989): "Agency Costs, Net Worth, and Business Fluctuations," *American Economic Review*, 79(1), 14–31.

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

Bernanke, B. S., M. Gertler, and S. Gilchrist (1999): "Chapter 21 The financial accelerator in a quantitative business cycle framework," vol. 1 of *Handbook of Macroeconomics*, pp. 1341 – 1393. Elsevier.

Beveridge, S., and C. R. Nelson (1981): "A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the 'Business Cycle'," *Journal of Monetary Economics*, 7(2), 151–174.

Borio, C. (2014): "The financial cycle and macroeconomics: What have we learnt?," *Journal of Banking & Finance*, 45(C), 182–198.

Borio, C., P. Disyatat, and M. Juselius (2017): "Rethinking potential output: embedding information about the financial cycle," *Oxford Economic Papers*, 69(3), 655–677.

Cagliarini, A., and F. Price (2017): "Exploring the Link between the Macroeconomic and Financial Cycles," in *Monetary Policy and Financial Stability in a World of Low Interest Rates*, ed. by J. Hambur, and J. Simon, RBA Annual Conference Volume (Discontinued). Reserve Bank of Australia.

Caldara, D., C. Fuentes-Albero, S. Gilchrist, and E. Zakrajšek (2016): "The macroeconomic impact of financial and uncertainty shocks," *European Economic Review*, 88(C), 185–207.

Check, A., and J. Piger (2018): "Structural Breaks in U.S. Macroeconomic Time Series: A Bayesian Model Averaging Approach," *mimeo*.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999): "Monetary Policy Shocks: What Have We Learned and to What End?," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 2, pp. 65–148. Elsevier.

Christiano, L. J., and T. J. Fitzgerald (2003): "The Band Pass Filter," *International Economic Review*, 44(2), 435–465.

Claessens, S., M. A. Kose, and M. E. Terrones (2012): "How do business and financial cycles interact?," *Journal of International Economics*, 87(1), 178–190.

Cogley, T., and J. M. Nason (1995): "Effects of the Hodrick-Prescott filter on trend and difference stationary time series Implications for business cycle research," *Journal of Economic Dynamics and Control*, 19(1–2), 253–278.

Coimbra, N., and H. Rey (2017): "Financial Cycles with Heterogeneous Intermediaries," Working Paper 23245, National Bureau of Economic Research.

Comin, D., and M. Gertler (2006): "Medium-Term Business Cycles," *American Economic Review*, 96(3), 523–551.

Del Negro, M., and F. Schorfheide (2011): "Bayesian Macroeconometrics," in *The Oxford Handbook of Bayesian Econometrics*, ed. by J. Geweke, G. Koop, and H. Van Dijk, pp. 293–389. Oxford University Press.

Drehmann, M., C. Borio, and K. Tsatsaronis (2012): "Characterising the financial cycle: don't lose sight of the medium term!," BIS Working Papers 380, Bank for International Settlements.

Eo, Y., and J. Morley (2017): "Why Has the U.S. Economy Stagnated since the Great Recession?," Working Paper 2017-14, University of Sydney, School of Economics.

Evans, G., and L. Reichlin (1994): "Information, Forecasts, and Measurement of the Business Cycle," *Journal of Monetary Economics*, 33(2), 233–254.

Faust, J., S. Gilchrist, J. H. Wright, and E. Zakrajšek (2013): "Credit Spreads as Predictors of Real-Time Economic Activity: A Bayesian Model-Averaging Approach," *The Review of Economics and Statistics*, 95(5), 1501–1519.

Forbes, K., I. Hjortsoe, and T. Nenova (2018): "The shocks matter: Improving our estimates of exchange rate pass-through," *Journal of International Economics*, 114(C), 255–275.

Forni, M., and L. Gambetti (2014): "Sufficient information in structural VARs," *Journal of Monetary Economics*, 66(C), 124–136.

Fry, R., and A. Pagan (2011): "Sign Restrictions in Structural Vector Autoregressions: A Critical Review," *Journal of Economic Literature*, 49(4), 938–60.

Furlanetto, F., P. Gelain, and M. Sanjani (2020): "Output Gap, Monetary Policy Trade-offs, and Financial Frictions," Working Papers 20200500, Federal Reserve Bank of Cleveland.

Furlanetto, F., F. Ravazzolo, and S. Sarferaz (2019): "Identification of Financial Factors in Economic Fluctuations," *The Economic Journal*, 129(617), 311–337.

Galati, G., I. Hindrayanto, S. J. Koopman, and M. Vlekke (2016): "Measuring financial cycles in a model-based analysis: Empirical evidence for the United States and the euro area," *Economics Letters*, 145, 83 – 87.

Gilchrist, S., V. Yankov, and E. Zakrajšek (2009): "Credit market shocks and economic fluctuations: Evidence from corporate bond and stock markets," *Journal of Monetary Economics*, 56(4), 471 – 493.

Gilchrist, S., and E. Zakrajšek (2012): "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 102(4), 1692–1720.

Gorton, G., and G. Ordoñez (2019): "Good Booms, Bad Booms," *Journal of the European Economic Association*.

Hamilton, J. D. (2018): "Why You Should Never Use the Hodrick-Prescott Filter," *The Review of Economics and Statistics*, 100(5), 831–843.

Jordà, O., M. Schularick, and A. M. Taylor (2013): "When Credit Bites Back," *Journal of Money, Credit and Banking*, 45(s2), 3–28.

Jordà, O., M. Schularick, and A. M. Taylor (2015): "Betting the house," *Journal of International Economics*, 96(S1), 2–18.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2019): "Credit Supply and the Housing Boom," *Journal of Political Economy*, 127(3), 1317–1350.

Kamber, G., J. Morley, and B. Wong (2018): "Intuitive and Reliable Estimates of the Output Gap from a Beveridge-Nelson Filter," *The Review of Economics and Statistics*, 100(3), 550–566.

Kamber, G., and B. Wong (2020): "Global Factors and Trend Inflation," *Journal of International Economics*, 122, 103265.

Koop, G. M. (2013): "Forecasting with Medium and Large Bayesian VARs," *Journal of Applied Econometrics*, 28(2), 177–203.

Krishnamurthy, A., and T. Muir (2017): "How Credit Cycles across a Financial Crisis," Working Paper 23850, National Bureau of Economic Research.

Leamer, E. E. (2007): "Housing is the business cycle," *Proceedings - Economic Policy Symposium - Jackson Hole*, pp. 149–233.

Litterman, R. B. (1986): "Forecasting with Bayesian Vector Autoregressions: Five Years of Experience," *Journal of Business & Economic Statistics*, 4(1), 25–38.

Morley, J., and B. Wong (2020): "Estimating and Accounting for the Output Gap with Large Bayesian Vector Autoregressions," *Journal of Applied Econometrics*, 35(1), 1–18.

Morley, J. C. (2002): "A State-Space Approach to Calculating the Beveridge-Nelson Decomposition," *Economics Letters*, 75(1), 123–127.

Morley, J. C., C. R. Nelson, and E. Zivot (2003): "Why Are the Beveridge-Nelson and Unobserved-Components Decompositions of GDP So Different?," *The Review of Economics and Statistics*, 85(2), 235–243.

Murray, C. J. (2003): "Cyclical Properties of Baxter-King Filtered Time Series," *The Review of Economics and Statistics*, 85(2), 472–476.

Newey, W. K., and K. D. West (1987): "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55(3), 703–708.

Oman, W. (2019): "The Synchronization of Business Cycles and Financial Cycles in the Euro Area," *International Journal of Central Banking*, 15(1), 327–362.

Perron, P. (1990): "Testing for a Unit Root in a Time Series with a Changing Mean," *Journal of Business & Economic Statistics*, 8(2), 153–162.

Robertson, J. C., and E. W. Tallman (1999): "Vector Autoregressions: Forecasting and Reality," *Economic Review*, 84(Q1), 4–18.

Rünstler, G., and M. Vlekke (2018): "Business, housing, and credit cycles," *Journal of Applied Econometrics*, 33(2), 212–226.

Taylor, J. B., and V. Wieland (2016): "Finding the Equilibrium Real Interest Rate in a Fog of Policy Deviations," *Business Economics*, 51(3), 147–154.

Walentin, K. (2014): "Business cycle implications of mortgage spreads," *Journal of Monetary Economics*, 67, 62 – 77.

Woźniak, T. (2016): "Bayesian Vector Autoregressions," *Australian Economic Review*, 49(3), 365–380.