

cege Discussion Paper

No. 415

March 2021

A Unified Approach for Jointly Estimating the Business and Financial Cycle, and the Role of Financial Factors

Tino Berger, Julia Richter, Benjamin Wong

A Unified Approach for Jointly Estimating the Business and Financial Cycle, and the Role of Financial Factors ^{*†}

Tino Berger¹, Julia Richter^{1,2}, and Benjamin Wong³

¹Georg-August University Göttingen

²German Council of Economic Experts

³Monash University, Australia

March 13, 2021

Abstract

We jointly estimate the U.S. business and financial cycle through a unified empirical approach while simultaneously accounting for the role of financial factors. Our approach uses the Beveridge-Nelson decomposition within 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 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@svr-wirtschaft.de Wong: benjamin.wong@monash.edu

†This paper reflects the personal views of the authors and not necessarily those of the German Council of Economic Experts. A previous version of this paper was circulated titled “Financial Factors and the Business Cycle”. We thank the editor, Ken West, two anonymous referees, Lutz Kilian, Mariano Kulish, James Morley, Anh D.M. Nguyen, Adrian Pagan, H el ene Rey, and James Yetman, and participants of seminars and conferences for comments and helpful discussions. We also thank Abby Nguyen for providing excellent research assistance. 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.

1 Introduction

The financial crisis of 2008-09 emphasized how developments in the financial market can spillover into the real economy, highlighting the importance to model and understand the role of the financial sector and how the financial sector of the economy interacts with the macroeconomy (see, e.g. Adrian and Shin, 2010, for a review). Within the policy sphere, it is important to understand the business and financial cycle because each is respectively used to understand imbalances in the real economy and financial sector.

The key contribution of our paper is to jointly model the business and financial cycle within a unified empirical approach. Our unified approach goes beyond just estimating both the business and financial cycle within a common empirical framework. Because we allow many variables to simultaneously evolve endogenously within a medium-sized VAR, we are also able to account for how much of the variation in the business and financial cycle can be attributed to financial variables and financial shocks. We also show how our approach unifies the estimation of business and financial cycle with SVAR work which seeks to identify financial shocks. In our work, we take the output gap, or the cyclical component of real GDP, as the business cycle, and both the housing and credit cycle, or the cyclical component of house prices and credit, as the financial cycle.

Our focus on modeling and quantifying the relationship of the business and financial cycle with financial factors is deliberate for at least two reasons. First, policy is often framed through the cyclical component of real activity and financial variables, which are the business and financial cycle respectively. For example, the output gap, or cyclical component of real GDP, is commonly used in policy settings, such as central banks, as being a summary measure of the business cycle, as well as capacity pressures. Similarly, macroprudential policy is also often framed in terms of the cyclical component of financial variables.¹ In such settings, the cyclical component of financial variables is taken to be a signal of financial imbalances and risk (for example, see Drehmann and Yetman, 2020). Our focus on the cyclical components is thus natural as this is precisely how macroeconomic stabilization and macroprudential policy is formulated. Second, we note that our approach is not unusual given broad segments of the extant literature. For example, an existing strand of the literature shares a similar focus of aiming to understand how financial factors shape the output gap, likely due to the reasons we outlined.² We also note that the practice of taking the cyclical component of house prices and credit as the financial cycle is not unusual relative to extant work (e.g., see Aikman, Haldane, and Nelson, 2015; Borio, Disyatat, and Juselius, 2017; Rünstler and Vlekke, 2018).

Briefly, our empirical approach builds off Morley and Wong (2020) and involves estimating a medium scale Bayesian Vector Autoregression (BVAR) containing both U.S. macroeconomic and financial variables, and subsequently applying the Beveridge-Nelson (BN) (1981) decom-

¹For example, macroprudential regulatory frameworks such as Basel III, treats the cyclical component of the credit-to-GDP ratio as the financial cycle.

²For example, see Aikman, Haldane, and Nelson (2015), Borio, Disyatat, and Juselius (2017), Cagliarini and Price (2017), Rünstler and Vlekke (2018), Furlanetto, Gelain, and Sanjani (2020), Constantinescu and Nguyen (2020), de Winter, Koopman, and Hindrayanto (2020) etc.

position to obtain both the output gap and measures of the financial cycle. We emphasize that our approach is unified and internally consistent to the extent that the output gap and financial cycle are obtained from the *same* time series model, namely our BVAR. We stress this is a non-trivial distinction relative to extant methods that first separately obtain the output gap and financial cycle before conducting subsequent analysis (e.g. Claessens, Kose, and Terrones, 2012; Aikman, Haldane, and Nelson, 2015; Albuquerque, Eichenbaum, Papanikolaou, and Rebelo, 2015) as it is well known how such analysis may be distorted by how one first obtains these cycles (e.g., see Canova, 1998, within the context of business cycle facts). Moreover, a key aspect of our empirical approach is that, because the output gap and financial cycle are obtained from the same BVAR, interpretation of the output gap and financial cycle are possible through standard VAR objects such as the forecast errors or identified structural shocks. It is the latter feature which will enable us to quantify the role of financial shocks for the output gap by appealing to the broader structural VAR literature (see Gilchrist and Zakrajšek, 2012; Caldara, Fuentes-Albero, Gilchrist, and Zakrajšek, 2016; Furlanetto, Ravazzolo, and Sarferaz, 2019).

Our key results are as follows. First, it appears that the role of financial factors played for both the output gap and financial cycles were much smaller pre-2000s, its role appears to have been much larger after the 2000s. In particular, 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 added somewhere between 2 to 4% to the 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. Second, we find that while the output gap and credit cycle are positively correlated, 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 an increase in the financial cycle to bust in the business cycle. Indeed, our work would suggest that the *average* credit boom is likely associated with a boom in the business cycle and vice versa. Only when conditioned on a financial shock, does the correlation switch to being negative. One interpretation, consistent with Jordà, Schularick, and Taylor (2013), suggests that it is excess credit that originates from loose credit conditions that do lead to busts in the business cycle.

We contrast our empirical approach to Borio, Disyatat, and Juselius (2017), Rünstler and Vlekke (2018) and de Winter, Koopman, and Hindrayanto (2020), which we regard as the closest in spirit to our work with regards to how one might model the relationship of financial factors to the output gap or jointly modeling the business and financial cycle. Borio, Disyatat, and Juselius (2017) use the Hodrick-Prescott (HP) filter as a starting point, and subsequently use credit growth as an exogenous variable after casting the HP filter into state-space form. As it is

well known, the HP filter may induce spurious cycles (see Cogley and Nason, 1995; Hamilton, 2018). In contrast, our approach, because it is based upon an explicitly specified time series, cannot, by construction, produce spurious cycles. Moreover, our approach does not treat credit as an exogenous variable in determining the output gap but instead allows real GDP growth, credit growth, and various macroeconomic and financial variables to evolve endogenously. This point is important because to the extent that decisions about granting or seeking credit are a function of how one views the macroeconomy, credit should be an endogenous variable. Work such as Rünstler and Vlekke (2018) and de Winter, Koopman, and Hindrayanto (2020) use Unobserved Components (UC) models to decompose real GDP, credit, and house prices into trend and cyclical components, and characterize the relationship between the subsequently extracted cyclical components. While UC models arguably are immune to spurious cycles, and thus at least from that perspective can be viewed as an improvement on the approach by Borio, Disyatat, and Juselius (2017), our approach has the advantage of linking variation from the business and financial cycles through the VAR forecast errors and/or identified financial shocks. It should also be noted, given we use the Beveridge-Nelson (BN) decomposition from a BVAR to obtain the output gap and the financial cycle, the trend and cycle from a BN decomposition and UC models are conceptually linked and identical through the reduced form of the UC model (see Morley, Nelson, and Zivot, 2003). In this regard, our empirical approach is thus conceptually akin to the UC model, except that the use of a BVAR enables us to explicitly identify the role of financial shocks, an option that is unavailable to standard UC models.

Finally, we note that part of our work also relates to broader 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. The aforementioned work by Borio, Disyatat, and Juselius (2017) embed financial sector information in conjunction with the Hodrick-Prescott filter to estimate output gaps that 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.³

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 both the business and financial cycle. Section 5 explores characteristics of the cross-correlation and lead-lag relationship between the business and financial cycle. Section 6 considers some robustness issues. Section 7 concludes.

³Though we focus on largely understand the cyclical component of real GDP and financial cycles, we also note alternative approaches such as Billio, Donadelli, Livieri, and Paradiso (2020), which seeks to understand the role of financial cycles in driving output growth by modifying a neoclassical growth model.

2 Empirical Framework

We use the Beveridge and Nelson (1981) (BN) decomposition to define the trend and cycle. 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 has a trend that follows a random walk process with a constant drift μ , the BN trend at time t , τ_t , is

$$\tau_t = \lim_{j \rightarrow \infty} \mathbb{E}_t [y_{t+j} - j \cdot \mu]. \quad (1)$$

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

$$c_t = y_t - \tau_t. \quad (2)$$

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.

Guided by the broader literature, we take the cyclical component of house prices and credit as estimates of the financial cycles, noting our choice of variables to consider for the financial cycle is also consistent with the UC model by Rünstler and Vlekke (2018). While there is less agreement about the variable of interest when measuring the financial cycle, there appears to be an 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).⁴

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 $i \in \{1, 2, \dots, K\}$. Let \mathbf{x}_t be a vector of n variables where $\Delta y_{i,t} \subset \mathbf{x}_t$.⁵ We assume that \mathbf{x}_t has a VAR(p) representation with the following companion form:

$$(\mathbf{X}_t - \boldsymbol{\mu}) = \mathbf{F}(\mathbf{X}_{t-1} - \boldsymbol{\mu}) + \mathbf{H}\mathbf{e}_t, \quad (3)$$

where $\mathbf{X}_t = \{\mathbf{x}'_t, \mathbf{x}'_{t-1}, \dots, \mathbf{x}'_{t-p}\}'$, $\boldsymbol{\mu}$ 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

⁴Drehmann, Borio, and Tsatsaronis (2012) argue that the cyclical component of house prices and credit are suitable variables to measure the financial cycle given share prices appear to have cyclical characteristics that do not accord with what one thinks of a financial cycle. The subsequent adoption by wider work to consider both credit and house price also suggests that their view has been influential in this emerging consensus. Nonetheless, for completeness, we present results for the stock market cycle in Section C of the online appendix.

⁵ \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.

errors to the companion form, and \mathbf{e}_t is a vector of serially uncorrelated forecast errors with covariance matrix Σ . Denoting $\tau_{i,t}$ and $c_{i,t}$ as respectively the BN trend and cycle of the series $y_{i,t}$,

$$y_{i,t} = \tau_{i,t} + c_{i,t}. \quad (4)$$

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)

$$c_{i,t} = -\mathbf{s}_k \mathbf{F} (\mathbf{I} - \mathbf{F})^{-1} (\mathbf{X}_t - \boldsymbol{\mu}). \quad (5)$$

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 $c_{ij,t}$ ⁶ as

$$c_{ij,t} = - \sum_{l=0}^{t-1} \mathbf{s}_k \mathbf{F}^{l+1} (\mathbf{I} - \mathbf{F})^{-1} \mathbf{H} \mathbf{s}'_j \mathbf{s}_j \mathbf{e}_{t-l}. \quad (6)$$

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

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

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

⁶Morley 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.

specifying a mapping

$$\mathbf{e}_t = \mathbf{A}\boldsymbol{\epsilon}_t, \text{ where } \mathbf{A}\mathbf{A}' = \boldsymbol{\Sigma}. \quad (8)$$

Let $c_{ij,t}^S$ 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

$$c_{ij,t}^S = - \sum_{l=0}^{t-1} \mathbf{s}_k \mathbf{F}^{l+1} (\mathbf{I} - \mathbf{F})^{-1} \mathbf{H} \mathbf{A} \mathbf{s}'_j \boldsymbol{\epsilon}_{t-l}. \quad (9)$$

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.

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), and also Berger, Morley, and Wong (2020) in a nowcasting setting, 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).⁷

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 1973Q1-2020Q1. Data for the excess bond premium is taken from Gilchrist and Zakrajšek (2012) its

⁷A 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 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.

subsequent updates by the Board of Governors.⁸ 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.

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.⁹

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 that 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 μ 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 inferences 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)

⁸See <https://www.federalreserve.gov/econresdata/notes/feds-notes/2016/recession-risk-and-the-excess-bond-premium-20160408.html>.

⁹Preliminary analysis suggests that a 15 variable BVAR may be informationally sufficient for the output gap, though it is a bit more mixed whether the 15 variable suffices for the financial cycles. Given our Bayesian shrinkage does not impose a large cost of including the additional 8 variables, we work with the 23 variable BVAR.

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.¹⁰ 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 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.¹¹ 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.¹² 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:¹³

¹⁰We tested for the break in the drift by allowing for heteroskedasticity and autocorrelation consistent (i.e. Newey and West, 1987) (HAC) standard errors.

¹¹The 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.

¹²We 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.

¹³If 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.

$$\begin{aligned}
\mathbf{x}_t &= \Phi_1 \mathbf{x}_{t-1} + \dots + \Phi_p \mathbf{x}_{t-p} + \mathbf{e}_t \\
&= \begin{bmatrix} \phi_1^{11} & \dots & \phi_1^{1n} & \phi_2^{11} & \dots & \phi_2^{1n} & \dots & \dots & \phi_p^{1n} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \ddots & \vdots \\ \phi_1^{n1} & \dots & \phi_1^{nn} & \phi_2^{n1} & \dots & \phi_2^{nn} & \dots & \dots & \phi_p^{nn} \end{bmatrix} \begin{bmatrix} \mathbf{x}_{t-1} \\ \mathbf{x}_{t-2} \\ \vdots \\ \mathbf{x}_{t-p} \end{bmatrix} + \begin{bmatrix} e_{1,t} \\ \vdots \\ e_{n,t} \end{bmatrix}, \quad (10)
\end{aligned}$$

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

$$\mathbb{E}[\phi_i^{jk}] = 0 \quad (11)$$

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

where the degree of shrinkage is governed by the hyperparameter λ , with $\lambda \rightarrow 0$ 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_i^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 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 our measure of the U.S. business cycle, the estimated U.S. output gap, together with our measure of the U.S. financial cycle, the estimated U.S. housing and credit cycle, alongside with their associated 90% credible interval. Our point estimate is based on the

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.

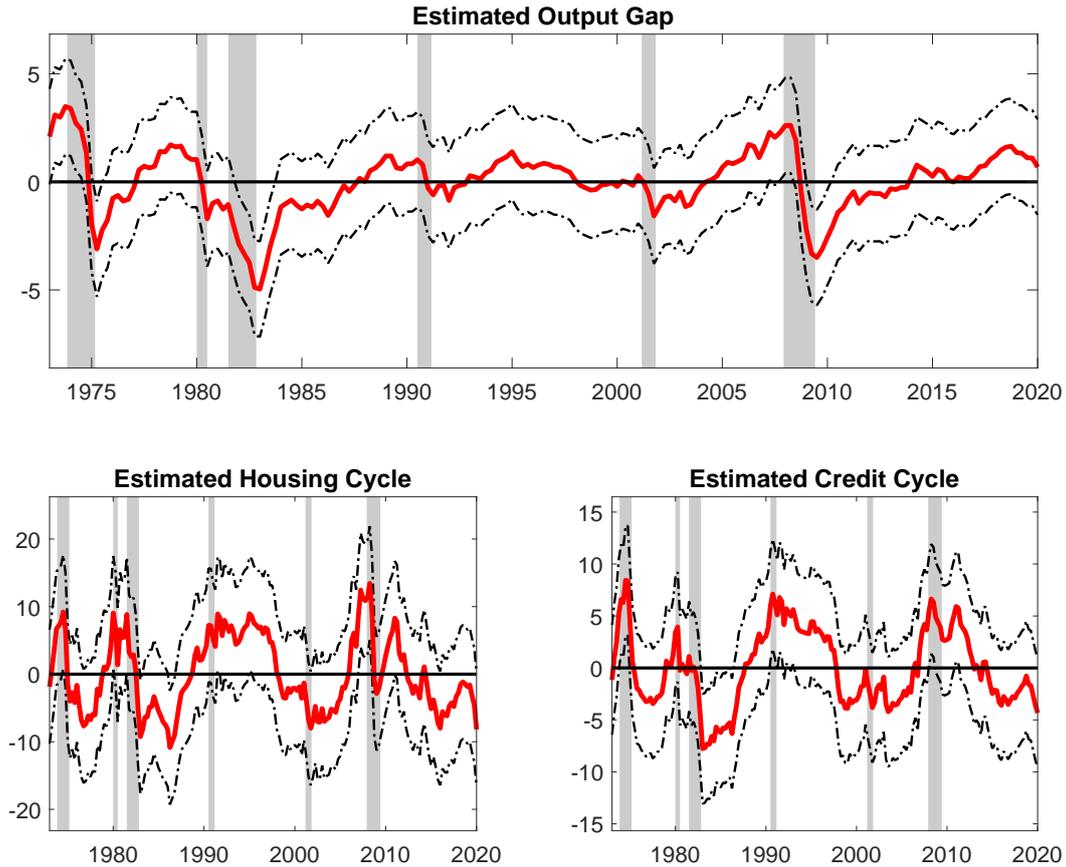


Figure 1: Estimated cycles from the BVAR. Units are in percent deviation from trend. Grey shaded areas indicate NBER recessions. 90% credible interval calculated as per Kamber, Morley, and Wong (2018)

BVAR posterior mode (i.e. we take the posterior mode of the BVAR parameters and thereafter construct the cycles by applying the BN decomposition to those BVAR parameters). 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 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). Turning to the estimates of the financial cycle, namely estimated the housing and credit 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 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 only rely on an underlying BVAR and the definition of the long-horizon forecast to define the trend and cycle. 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 can reassess the view on the relative duration of the business and financial cycle through the lens of our model. As Cagliarini and Price (2017) point out, a widely held view that the financial cycle has a much longer duration than the business cycle

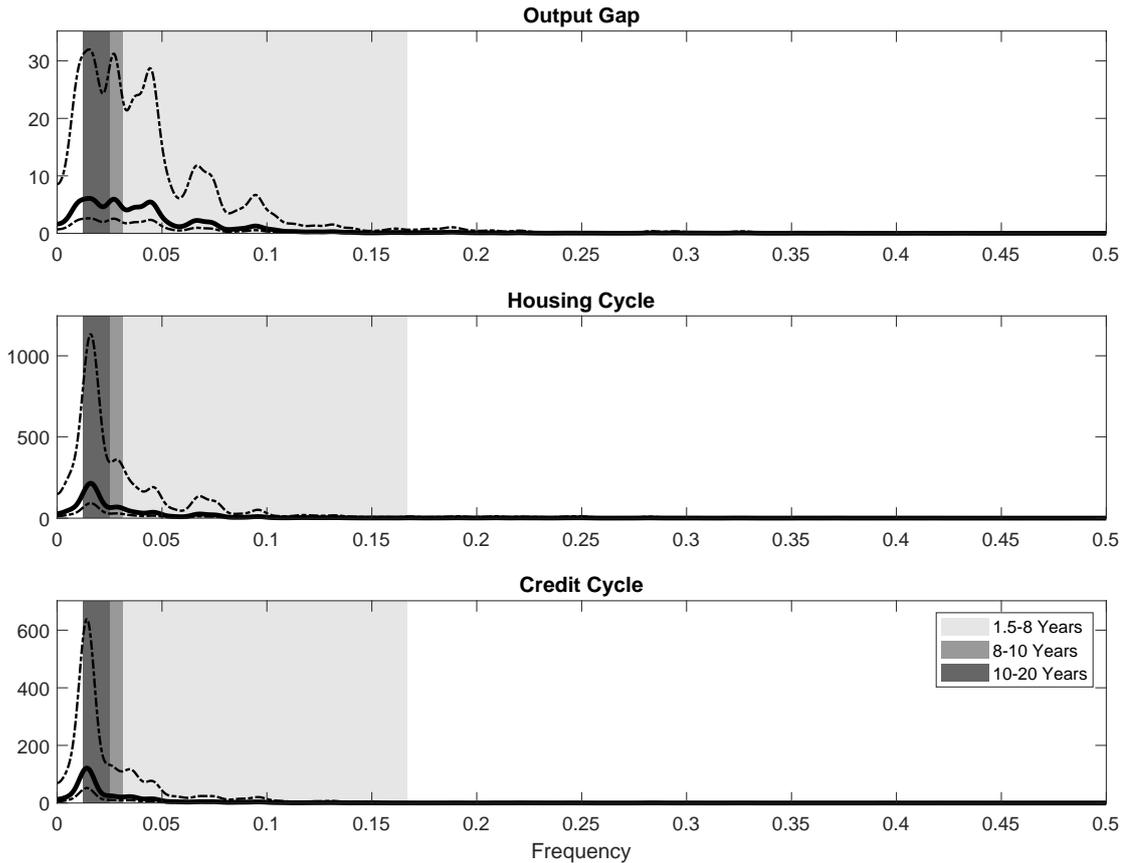


Figure 2: Estimated spectral density of the estimated cycles with 90% credible interval. The frequencies associated with $1\frac{1}{2}$ to 8 years, 8 to 10 years, and 10 to 20 years are highlighted.

may be partly driven by assumptions on which frequencies to isolate, potentially obscuring the distinction between assumptions and conclusions.¹⁴ Figure 2 presents the estimated spectral density of the estimated output gap, housing cycle, and credit cycle.¹⁵ We highlight the frequencies between $1\frac{1}{2}$ to 8 years, 8 to 10 years, and, 10 to 20 years. Recall that $1\frac{1}{2}$ 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). Our point estimate for the spectral density is similarly based on the posterior mode as per the point estimate in Figure 1. The 90% credible interval is taken from reconstructing the output gap and financial cycle from draws of the posterior distribution of the VAR coefficients and subsequently estimating the spectral density of the reconstructed output gap and financial cycle.

We find that our estimated output gap is the only cycle that features a non-trivial degree of fluctuations between $1\frac{1}{2}$ to 8 years. That is, we find very little of the variation of either

¹⁴For example, users of the bandpass filter take frequencies of $1\frac{1}{2}$ 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.

¹⁵In estimating the spectral density, we follow Schüller (2020) and use a Parzen window of $12\sqrt{T}+1$ to smooth the periodogram.

the housing or credit cycle is within the frequencies associated with $1\frac{1}{2}$ to 8 years. Instead, it appears that much of the variation of the housing and credit cycle occur at the 10-20 year frequency, with both featuring a dominant peak of the spectral densities within the 10 to 20 year window. More precisely, the dominant peak in the spectral density of the housing and credit cycle occurs 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). We note that from the posterior distribution, the dominant peak of the spectral density in the financial cycle appears fairly precisely estimated. While the output gap does feature fluctuation between the traditional business cycle frequencies of $1\frac{1}{2}$ to 8 years, we also find that a non-trivial degree of fluctuation outside the traditional business cycle frequencies. Indeed, while we note that the traditional frequencies associated with the business cycle are $1\frac{1}{2}$ to 8 years and noting the caveat that the broader literature uses different methods which may compromise comparability, Comin and Gertler (2006) emphasize non-trivial business cycle frequencies in the 2-50 year window, while Rünstler and Vlekke (2018) also find the dominant cycle to be just outside the 8 years range.

Overall, we find mixed evidence of whether the financial cycle to be substantially longer than the business cycle. A key reason for our finding is that while the peaks of the spectral density for both the housing and credit cycle appear to be very sharply identified within the 10-20 year window, the peak of the spectral density for the output gap is fraught with a large degree of uncertainty. For example, while the posterior mean difference of the implied dominant frequency of the business cycle is 10 quarters shorter than that of the credit cycle, our estimated posterior probability that the dominant frequency of the financial cycle implies a longer cycle than that implied by the dominant frequency of the business cycle is 60%, which while larger than a 50-50 probability, does on balance constitutes mixed and perhaps weak evidence.¹⁶ We also note, once again with the caveat of being in a different model setting, Kulish and Pagan (2019) tested the Rünstler and Vlekke (2018) model and are unable to reject the null hypothesis that the financial cycle in their model is longer in duration relative to the business cycle, a similar conclusion also arrived by Cagliarini and Price (2017). Through constructing the posterior distribution of the estimated spectral density, our results would suggest that imprecision involved in estimating the dominant frequency of the business cycle may reconcile the mixed evidence in the wider literature.

4 The Role of Financial Factors in Driving the Business and Financial Cycles

We now turn to the role of financial factors in driving the business and financial 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

¹⁶Note that we can make probability statements as these quantities are obtained via a Bayesian posterior distribution.

through guidance from the broader literature.

4.1 Informational decomposition of the output gap

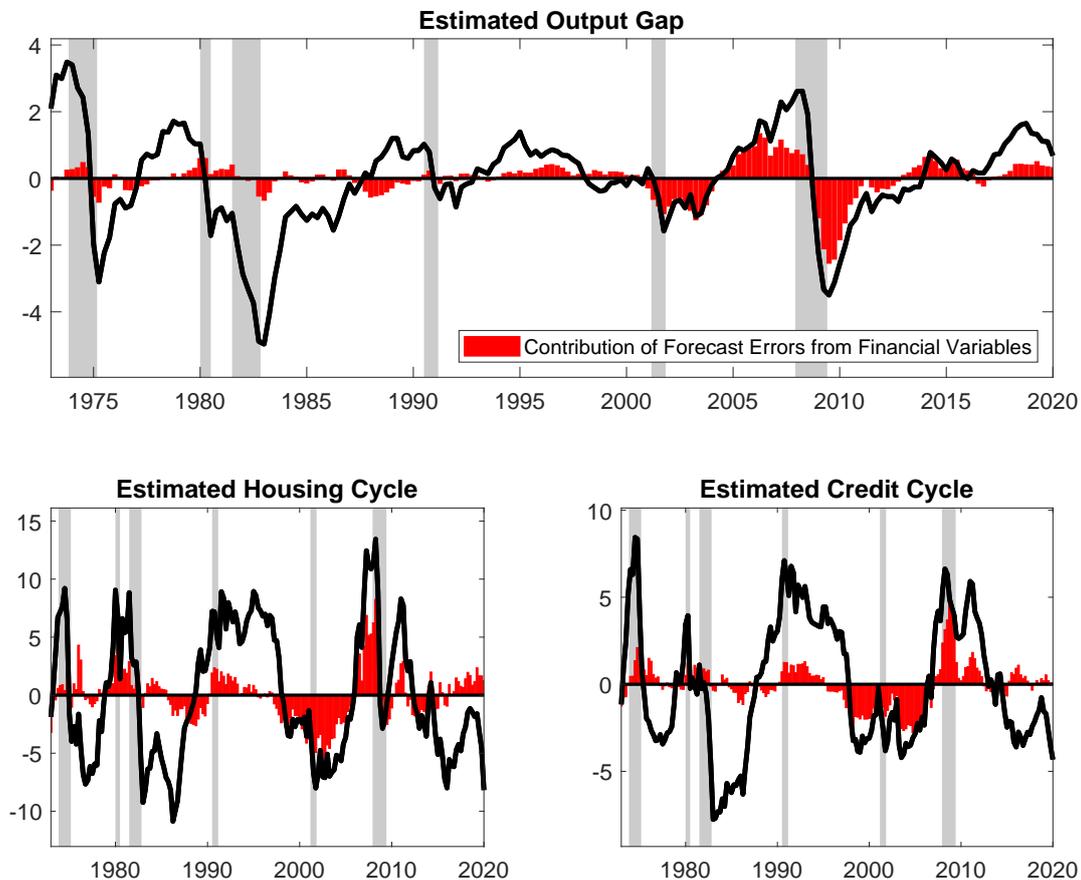


Figure 3: Informational decomposition of the estimated cycles. Solid line denotes the estimated cycle. Cycles are measured in percent deviation from the trend. Grey shaded areas indicate NBER recessions. The bars represent the total contribution of the contribution from the BVAR forecast errors from five financial variables (credit, the excess bond premium, the S&P 500, the VIX index, and the house price) The individual contributions are presented in Figure 4.

Figures 3 and 4 present the informational decomposition for the estimated output gap and financial cycles 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. Figure 4 reports the individual shares of the forecast errors of the five chosen financial variables, while Figure 3 sums up these contribution. We emphasize that the informational decomposition is not causal, so any conclusions about causal mechanisms from the information decomposition should only be viewed as suggestive. In particular, the information contained within the forecast errors of financial variables could originate from shocks outside the financial sector and/or forecast errors that have little or a

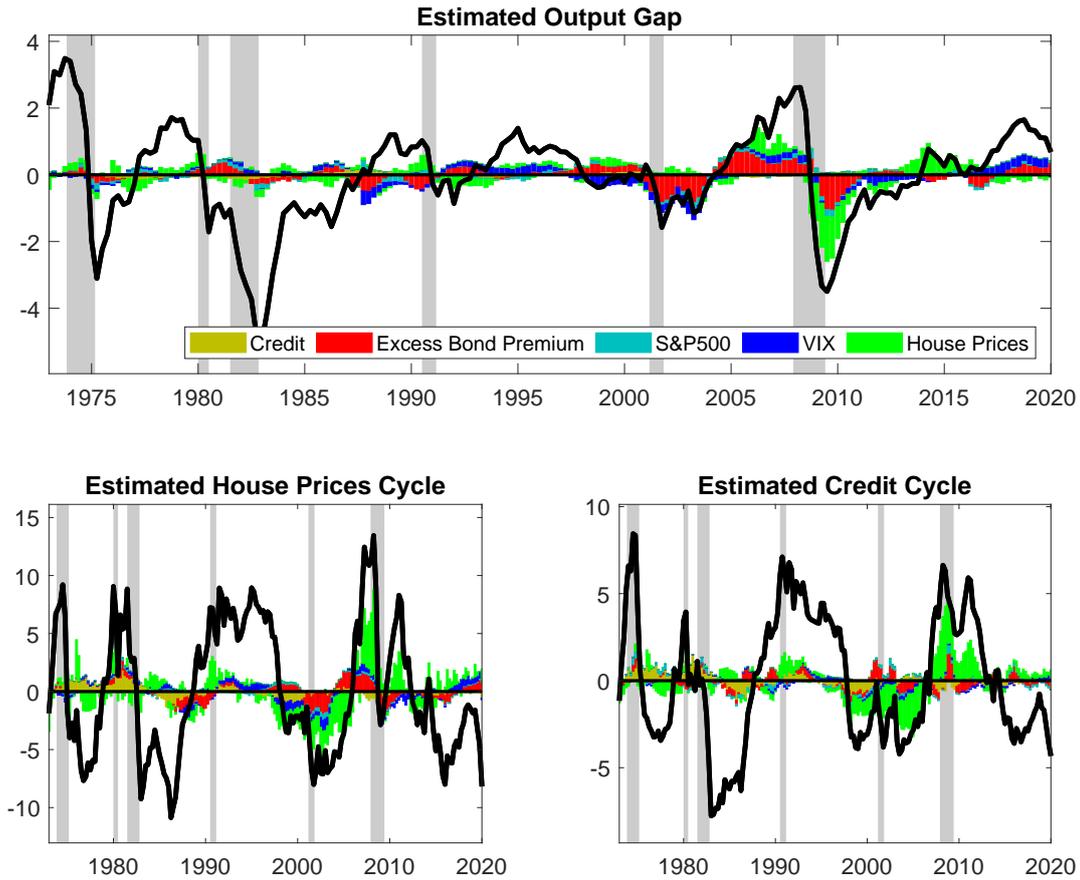


Figure 4: Informational decomposition of the estimated cycles. Solid line denotes the estimated cycle. Cycles are measured in percent deviation from the trend. Grey shaded areas indicate NBER recessions. The bars represent the individual contribution from the BVAR forecast errors from five financial variables (credit, the excess bond premium, the S&P 500, the VIX index, and the house price).

negligible role do not necessarily indicate their respective variables have no role.¹⁷

We document two general key observations from Figure 3. First, the role of financial variables seems to have been important during the 2000s, but its impact is rather negligible before the 2000s, and especially so before the mid 1990s. It is a more open question whether, towards 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

¹⁷The latter point is worth elaborating on with a stylized example. Suppose variable A Granger causes variable B, and variable B Granger causes variable C, but variable A does not Granger causes variable C. Clearly in this case, variable B matters for the estimation of the BN cycle of variable C (see Evans and Reichlin, 1994). However, the forecast errors of variable A will matter for the informational decomposition of the cycle of variable C through variable B. Therefore, even if the forecast errors of variable B do not show up in the informational decomposition of the cycle of variable C, variable B is still important, because, without the role of variable B, the forecast errors of variable A would never show up in the informational decomposition of the cycle of variable C.

variables on the output gap in periods of financial stress, such as the burst of the dot-com bubble and the outbreak of the financial crisis as well as during the build-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 sizeably to both the output gap and financial cycles. 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 the 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.

While we once again stress that the interpretation from the informational decomposition is not causal, it represents a useful starting point. That the forecast errors of house prices and the excess bond premium contain information for both the output gap and measures of the financial cycle suggest that they would have probably played a role in linking and understanding the business and financial cycle during the 2000s.

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, it is unfortunate that even if several contributions share the label “financial shock”, what each seeks to identify and isolate may not conceptually align 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 of credit availability 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 utilizing (orthogonal) variation in the excess bond premium to identify financial shocks. Mechanically, the implementation amounts to ordering the excess bond premium after slow-moving variables such as GDP, investment, etc, and before fast-moving variables, 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. Specifying a slow-moving and fast-moving block is a reasonably common strategy for using the Cholesky decomposition within a system that features both financial and macroeconomic variables (e.g., see Christiano, Eichenbaum, and Evans, 1999; Bernanke, Boivin, and Elias, 2005). 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. 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.¹⁸ Like the

¹⁸We note that our objective differs 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

Cholesky decomposition, the penalty function approach also relies on orthogonalized 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.

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.¹⁹ 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) and Ben Zeev (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 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 in 2008Q4. We therefore implement a narrative sign restriction that the

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. We will return to the issue if one wanted to also identify an uncertainty shock in the robustness section.

¹⁹Note 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 strategies, investment will fall more than GDP in response to a positive financial shock, so the investment to GDP ratio falls.

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.

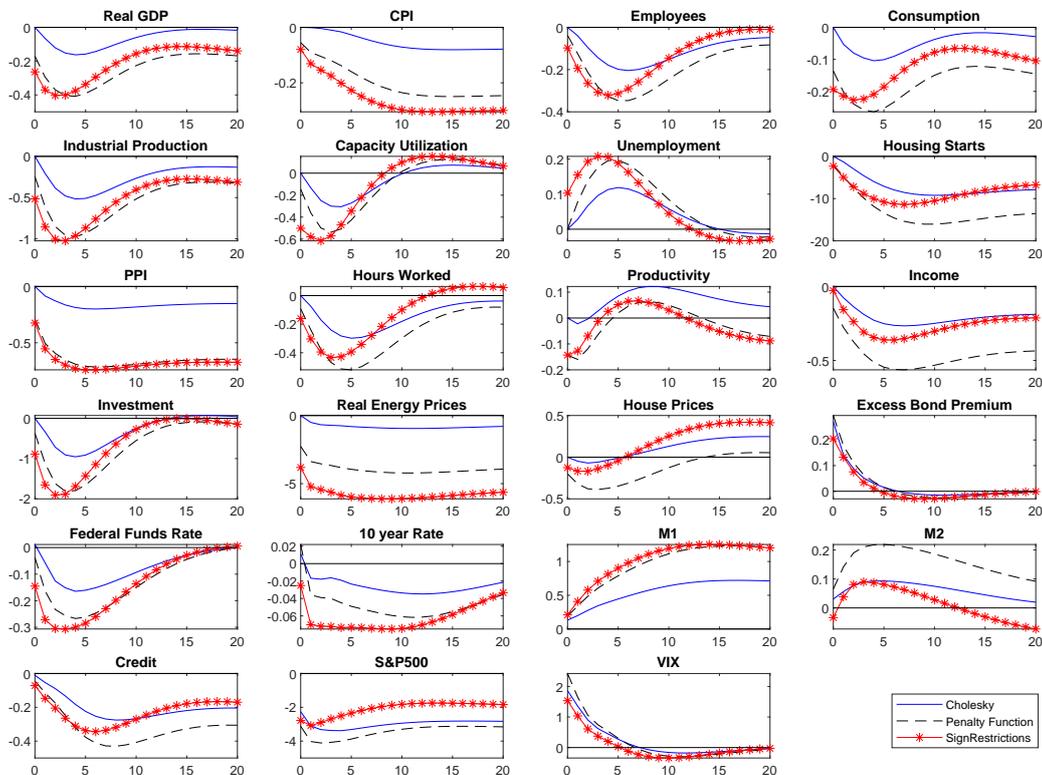


Figure 5: Estimated impulse response function to a one standard deviation financial shock. All impulse response functions are calculated from the posterior mode of the estimated BVAR. The sign restriction impulse response function is chosen from a representative model using the Fry-Pagan median target from 1000 rotations of the posterior mode 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 5 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 mode of the BVAR estimates and we report the impulse response functions of the sign restrictions using the Fry and Pagan (2011) median target approach.²⁰ We explicitly

²⁰There 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)

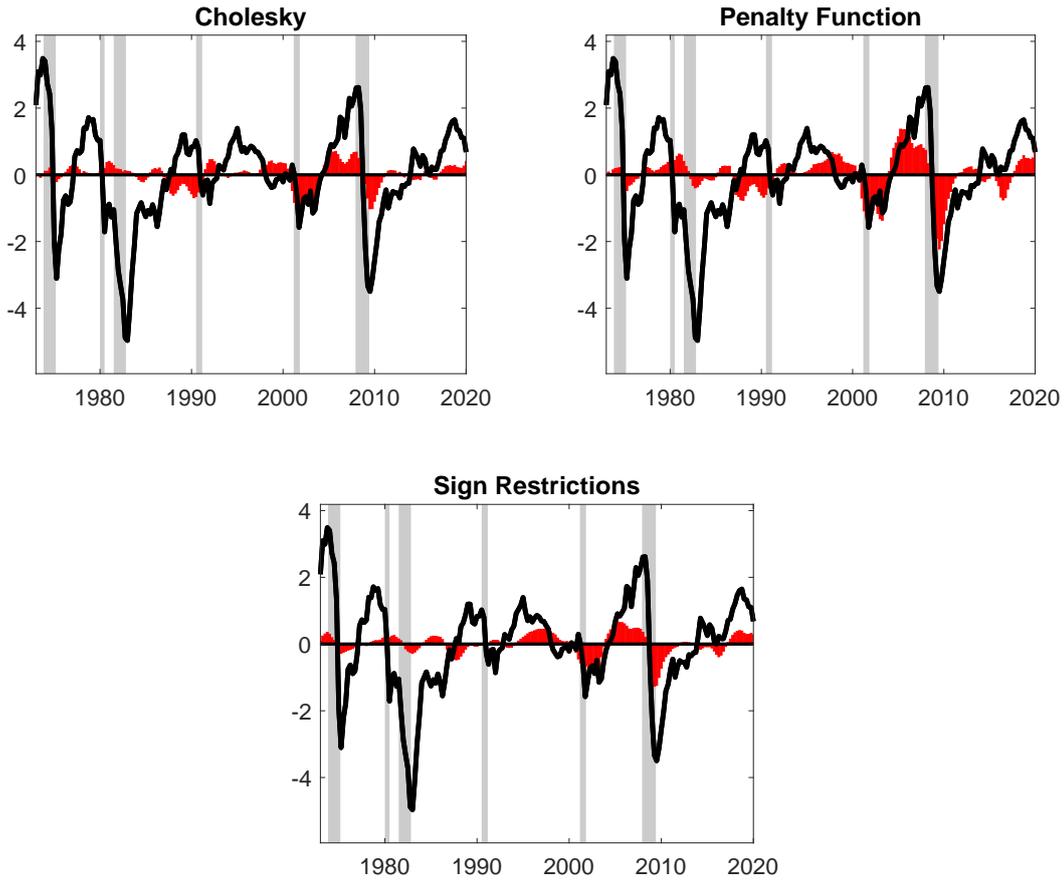


Figure 6: Contribution of the financial shock to the estimated output gap. The solid line represents the estimated output gap. The output gap is measured in percent deviation from trend. Grey shaded areas indicate NBER recessions. The title refers to the different identification schemes. The bars represent the contribution of financial shocks to the estimated output gap. The contribution from the sign restriction approach is averaged across draws that satisfy the sign and narrative restrictions.

account for parameter uncertainty by presenting posterior credible sets of all three identification 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 is 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

median target approach here from 1000 admissible solutions conditional on the posterior mode parameters. This is just for illustrative purposes as we wish to just compare the sign of the responses of all three identification schemes.

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 6 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 mode of the BVAR parameters in and equivalent to reporting Equation (9) across different \mathbf{A} 's.²¹ 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 tends 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.

To more precisely quantify how much financial shocks contributed to the overheating of the U.S. output gap in the 2000s, Figure 7 presents our estimate of how much financial shocks contributed to the U.S. output gap between 2002Q1 and 2005Q4 along with the associated credible sets and credible intervals. We choose this time period as 2002Q1 marked the first quarter after the 2000-01 recession. We choose 2005Q4 as end 2005 was the height of the asset bubble. To construct these credible sets and intervals, for each draw of the posterior distribution, we construct the implied output gap sequence of identified financial shocks, then calculate the role of financial shocks on the output gap for the time period in question.²² Because the financial shock is an identified (orthogonal) structural shock, the interpretation from Figure 7 would be our estimated counterfactual reduction in the output gap from 2002Q1 to 2005Q4 in the absence of the identified financial shock. The bounds of the 68% credible interval are taken from the 16th and 84th quantiles of the posterior distribution. Because the quantiles may obscure information about the dynamics as the role of financial shocks is derived from a path rather than a point on a distribution (see Inoue and Kilian, 2020, for the analogous argument from the perspective of an impulse response function), we also present the associated credible sets calculated via the absolute loss function as described by Inoue and Kilian (2020).²³

²¹For the sign restriction results, we averaged over the 1000 rotations which satisfies the sign and narrative restrictions conditional on the posterior mode 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 5, 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.

²²Note that this would entail subtracting the contribution of financial shocks on the output gap in 2002Q1 from the contribution of financial shocks on the output gap in 2005Q4 for each draw of the posterior distribution. For the Cholesky and penalty function identification, this effectively requires us to just take a draw from the reduced form and then construct all these associated quantities. For the sign and narrative restrictions, we have to construct membership of the posterior distribution by allowing for satisfying both the sign and narrative restriction as described by Antolín-Díaz and Rubio-Ramírez (2018), then construct the associated quantities for each draw of the posterior distribution.

²³It is a more unresolved issue whether using impulse response function, as Inoue and Kilian (2020) do, is

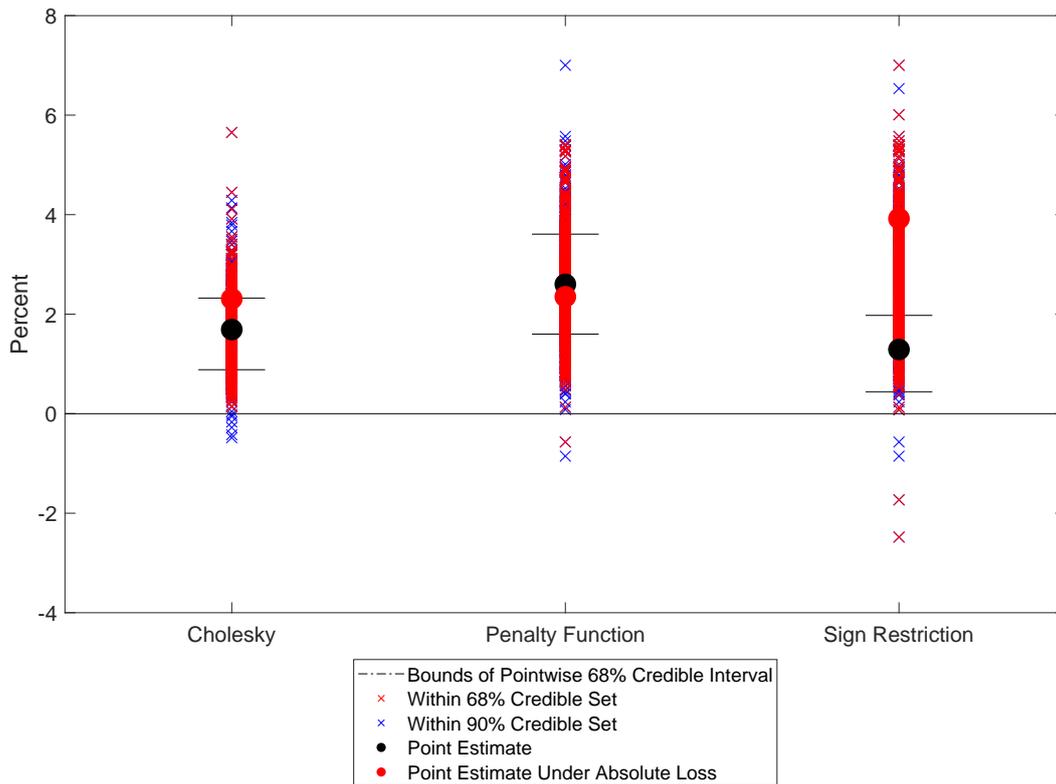


Figure 7: Contribution of the identified financial shock to the estimated output gap (in percent) for the period 2002Q1-2005Q4 under the three identification schemes. The solid lines represent the pointwise bound of the 68% credible interval. The x represent membership in either 68% or 90% credible set obtained under absolute loss function described by Inoue and Kilian (2020). The point estimates for both Cholesky and penalty function identification are obtained conditional on the mode of the VAR posterior distribution.

We take the posterior mode as our point estimate for both the Cholesky and penalty function identification, and for the sign restriction, the mean across 1000 rotations which satisfy the sign and narrative restriction but conditional on the posterior mode of the reduced form, to just retain comparability to Figure 6. We also consider an optimal point estimate under absolute loss, for the posterior draw which evaluates the minimum loss. All the point estimates, under our preferred approach conditioning on the posterior mode and under absolute loss, imply the identified financial shocks added somewhere between 2 to 4% to the output gap. In other words, in a counterfactual without the identified financial shock, the increase in the output gap between 2002Q1 to 2005Q4 would have been 2 to 4 percentage points lower, which is reasonably large, considering the historical magnitude of the estimated output gap in Figure 1. Given the

the most appropriate approach to evaluate the loss function given impulse response functions are not the focus of our analysis. We choose to evaluate the loss function based on the impulse response function to a financial shock to mostly maintain comparability with the description found in Inoue and Kilian (2020), as well as the credible sets we present in section D of the online appendix. Note that our approach would tantamount to treating the impulse response function as the primary object of interest from the BVARs, which one may argue is not necessarily true in our setting, but an appropriate compromise given the issue is still not entirely resolved. We thank Lutz Kilian for the many discussions on this issue with us.

lower bound of the 68% credible set is greater than zero under all three identifications, it implies that at least 84% of the posterior draws estimate a role of where identified financial shocks led to an increase in the output gap between 2002Q1 to 2005Q4. Turning to the credible sets, we first focus on the posterior draws within 68% credible set. Apart from 1 draw for the penalty function, and 2 draws for the sign restrictions, all elements of the credible set estimate a role for the financial shocks leading to an increase in the output gap. Note that once one moves to the credible set setting, the estimates implied by these sets are not be continuous, in the sense that we are just reporting elements associated with draws from the posterior distribution which one evaluates a smaller loss from the associated loss function. It is noteworthy while there is a greater dispersion relative to the bounds of the credible interval, almost all elements of the credible set across all three identification schemes are still bunched up between our 2 to 4% estimate. Finally, we show that even if we considered a 90% credible set, our conclusion is almost identical to using a 68% credible set.

Therefore, based on the overall evidence presented, our results point to a prominent role of financial shocks in contributing sizably to a large and positive output gap before the 2008/09 recession. Our interpretation is consistent with the notion 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. While there is some uncertainty around the estimates of how much financial shocks matter, our estimates suggest financial shocks led to between a 2 to 4% increase in the estimated output gap with the credible interval and credible sets suggesting a very high probability that financial shocks led to some degree of overheating of the business cycle between 2002Q1 to 2005Q4. 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.

To round out our analysis, we also quantify the role of the estimated financial shocks on the financial cycle. Figure 8 present these results. In general, the role of financial shocks across all three identification schemes is fairly similar. However, the role of financial shocks on both the house price cycles and credit cycles appear to be a bit different. We note that the role of financial shocks with the estimated house price cycles appearing more like the role of financial shocks with the output gap. While the role of financial shocks in the credit cycle, at least with the Cholesky and penalty function identification, appears more muted, we still find a role for financial shocks in driving the credit cycle in the 2000s. Nonetheless, we note that relative to raw fluctuations of the estimated financial cycle, the role of identified financial shocks when account for the variation in the financial cycle is still much smaller than the role that the identified financial shock has in accounting for variation in the output gap.

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 forecast errors of the financial variables contributed more since the 2000s and played a large role in the overheating of the business

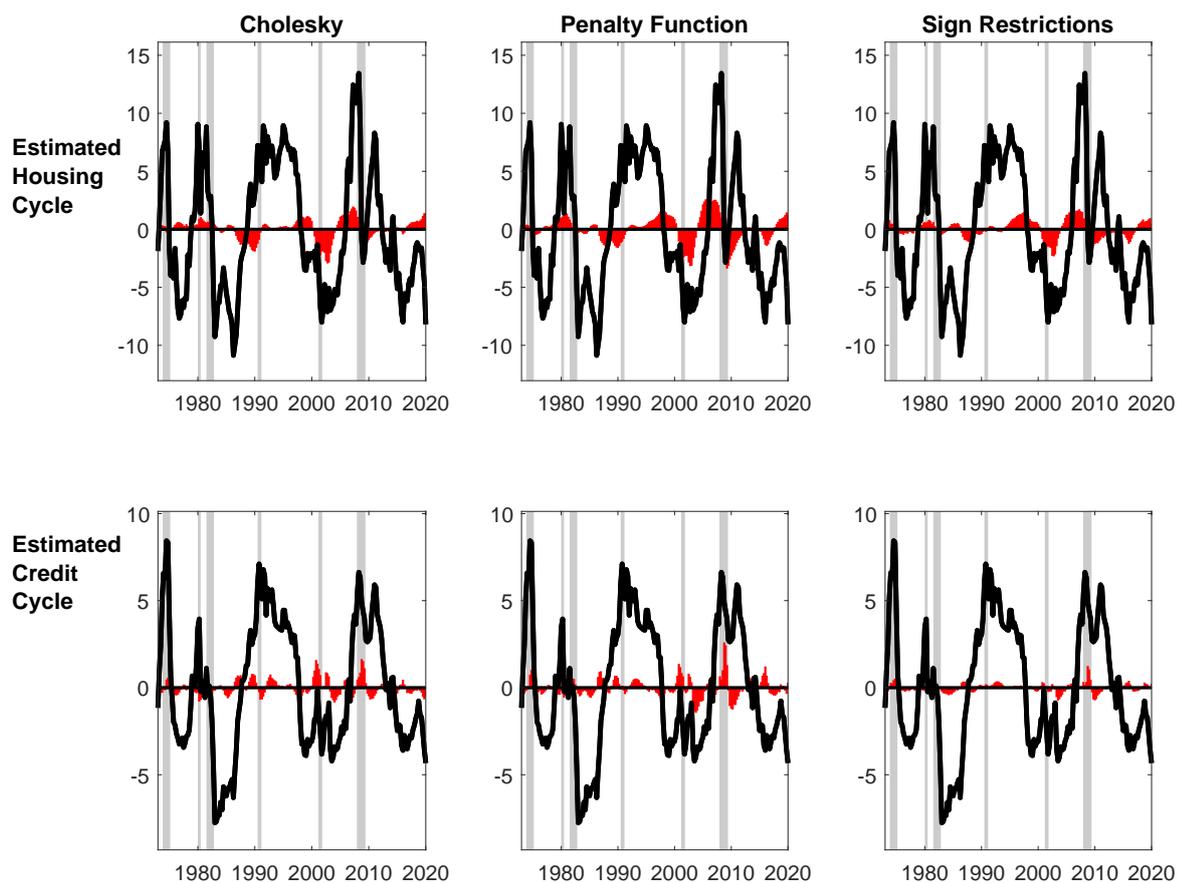


Figure 8: Contribution of the financial shock to the estimated financial cycles. The solid line represents the estimated housing cycle (top panels) and estimated credit cycle (bottom panels). The cycles are measured in percent deviation from trend. Grey shaded areas indicate NBER recessions. The headers refer to the different identification schemes. The bars represent the contribution of financial shocks to the estimated financial cycle. The contribution from the sign restriction approach is averaged across draws that satisfy the sign and narrative restrictions.

cycle, a result which is also consistent with the role of the structural identified financial shock.

From our results, it would appear that the role of the financial variables is much larger than that of financial shocks. For example, if we zoom in on the output gap during the 2000s, the role of the identified financial shocks as shown in Figure 6 is about half that of the role of financial variables, as shown by Figures 3 and 4, depending on the precise identification scheme used. While we again stress that the informational decomposition is in reduced form, and so the role of these forecast errors should not be interpreted as causal, we briefly reconcile the differences we observe between the informational decomposition in Figures 3 and 4 with the structural decomposition in Figures 6 and 8 during the 2000s boom, given a key narrative is that financial factors appear to play a role in overheating the real economy, as indicated in our decomposition of the output gap and financial cycles.

To begin, it should not be entirely surprising that the role associated with the forecast errors of the financial variables is larger than that ascribed to financial shocks. After all, the

forecast errors reflect variation from all the identified and unidentified shocks. Given we are only identifying one shock, one would expect the role of the financial shock to be much smaller than that reflected by the forecast errors of the financial variables since we expect shocks from the real economy, which we do not identify in our exercise, should also drive a non-trivial proportion of this variation in the forecast errors of the financial variables. From Figure 4, during the period from 2000 to 2008, the key financial variables whose forecast errors are driving the output gap and the financial cycles are the excess bond premium and house prices. At first glance, the role of the financial shocks driving the output gap in the 2000s is approximately the same as being ascribed to the excess bond premium.²⁴ Therefore, it would appear during the 2000s boom, the forecast error of house prices is approximately the difference between the role attributed to forecast errors of the financial variables and the role attributed to financial shocks. Note that the preceding statement does not necessarily mean house prices did not have a role in the 2000s boom. Our analysis almost certainly suggested that house prices had a role given the excess bond premium in the informational decomposition and financial shocks had a non-trivial role in the house price cycle in the 2000s. Because the excess bond premium (and financial shocks) had a non-trivial role in the house price cycle, it is almost certainly true that whatever the role the financial shocks had on the output gap in the 2000s, it had a similar role in the housing cycle.

A key insight from comparing both the informational decomposition and the structural decomposition is that it reveals that one needs to largely explain the house price forecast errors within the model to provide a fuller account of the business and financial cycle in the 2000s. Put differently, while some of the current SVAR approaches to identifying financial shocks which we explore in our structural analysis go a long way in understanding the business and financial cycle in the 2000s, one would need to find a set of, or a single, shocks which can explain the forecast errors of house prices to fully reconcile the business and financial cycle in the 2000s.

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 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

²⁴We confirm this when we looked at the role of financial shocks on the role of the excess bond premium forecast errors in the informational decomposition. While there were slight differences across the identification schemes, financial shocks accounted for most of the share of the role of the forecast errors of the excess bond premium in the informational decomposition.

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 3 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, an active body of work is 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; de Winter, Koopman, and Hindrayanto, 2020).

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 $\mathbf{F}(\mathbf{I} - \mathbf{F})^{-1}(\mathbf{X}_t - \boldsymbol{\mu})$ 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

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

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

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

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 $m \in \{0, 1, 2, \dots\}$.²⁵ 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

$$\text{corr}(c_{i,t}, c_{j,t-m}) = \mathbf{s}_k \boldsymbol{\psi} \mathbf{s}'_{\mathbf{nm}+1}, \quad (15)$$

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

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

and substitute $\tilde{\mathbf{Q}}$ for \mathbf{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

²⁵If 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 $\boldsymbol{\Psi}$.

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.

Unconditional Cross-Correlations

		Correlations		
		Output Gap	House Price Cycle	Credit Cycle
	Output Gap	1		
	House Price Cycle	0.39	1	
	Credit Cycle	0.32	0.85	1
		Contemporaneous (t)		
		Output Gap	House Price Cycle	Credit Cycle
Lagged 4	Output Gap	0.36	0.37	0.57
Quarters	House Price Cycle	0.17	0.90	0.82
($t - 4$)	Credit Cycle	0.01	0.74	0.91

Conditional Cross-Correlations

		Cholesky		
		Contemporaneous (t)		
		Output Gap	House Price Cycle	Credit Cycle
Lagged 4	Output Gap	0.31	0.20	0.50
Quarters	House Price Cycle	0.74	0.90	-0.35
($t - 4$)	Credit Cycle	-0.69	-0.89	0.78

		Penalty Function		
		Contemporaneous (t)		
		Output Gap	House Price Cycle	Credit Cycle
Lagged 4	Output Gap	0.04	-0.02	0.81
Quarters	House Price Cycle	0.53	0.61	0.29
($t - 4$)	Credit Cycle	-0.72	-0.80	0.74

Sign Restrictions, percentage of negative correlations

		Contemporaneous (t)		
		Output Gap	House Price Cycle	Credit Cycle
Lagged 4	Output Gap	0	0	10.6
Quarters	House Price Cycle	0	0	13.8
($t - 4$)	Credit Cycle	82.3	69.3	7.1

Table 2: Unconditional and conditional cross-correlations.

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).

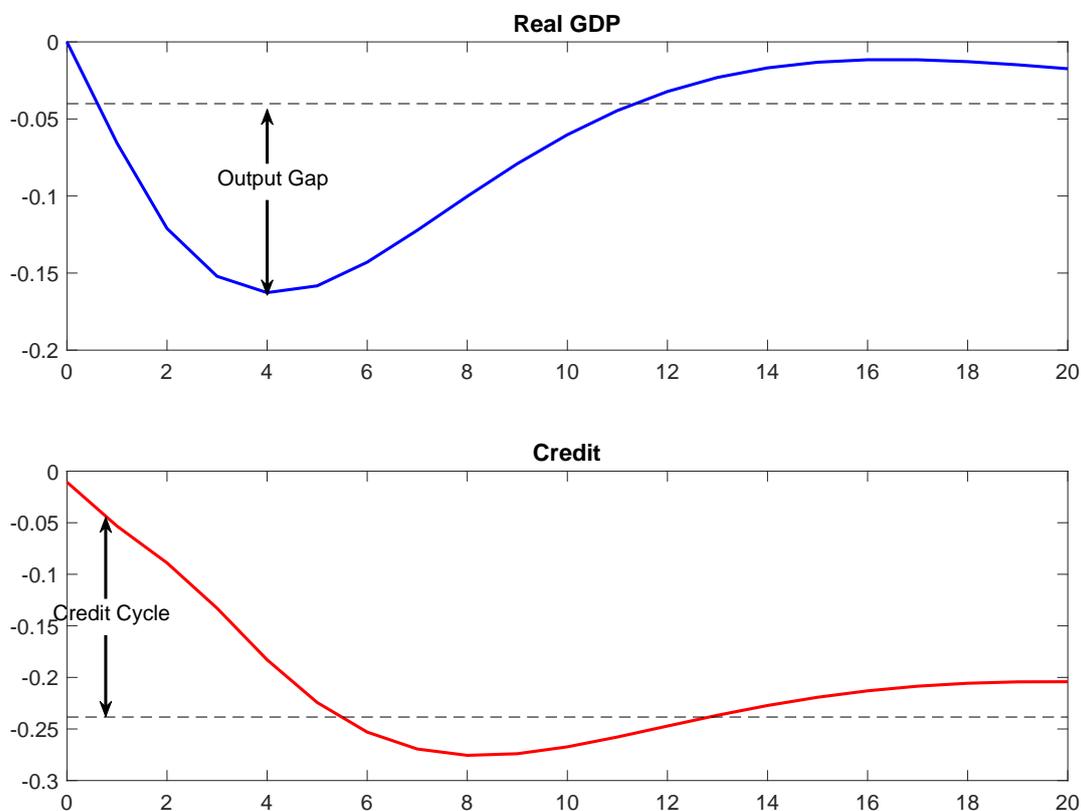


Figure 9: Impulse response function to a one standard deviation financial shock identified using Cholesky Decomposition

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.²⁶ This 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.

Figure 9 provides some intuition on why we observe a sign switch conditional on a financial shock. Presented are the impulse response function of real GDP and credit to a one standard deviation financial shock identified using the Cholesky decomposition, though the precise identification matters less given all three identification schemes show similar patterns. The impulse response functions of the level of real GDP and credit are based on cumulating the impulse response functions of real GDP growth and credit growth since both variables enter the BVAR in first differences. The definition of trend in the BN decomposition is the forecast of the long-horizon forecast. Given, by definition, the impulse response function is the response to only a financial shock being introduced into the system at time 0, the trend becomes where the level of real GDP and credit settle in the long-run. This is denoted by the dotted line in Figure 9. The difference between the impulse response function and the long horizon forecast, denoted by the dotted line, thus becomes the output gap and credit cycle which we obtain from via the BN decomposition.

The dynamics of real GDP are such that while it falls quickly in response to the financial shock, there is a hump-shaped response where the level of real GDP starts to recover 4 to 5 quarters after the final shock. This also means that the eventual fall in real GDP relative to before the financial shock is more marginal as the level eventually largely recovers. Given the level of impulse response function of real GDP is below this long-horizon level, a negative output gap opens up for up to 10 quarters after the financial shock. On the other hand, the dynamics of credit are quite different. In response to a financial shock, credit falls slowly towards its long-horizon forecast. Because credit is above its long horizon level for up to 6 quarters after the shock, a positive credit cycle opens up initially as the level of credit adjusts towards its long-horizon level. Figure 9 also clarifies the source of the negative correlation of the output gap and credit cycle conditional on a financial shock. Because the level of credit adjusts slowly, but the long-horizon level fall by more, the credit cycle and the output gap thus become negatively correlated conditional on the financial shock. The impulse response function should make clear that credit and real GDP are still positively correlated conditional on a financial shock since they both move in the same direction in response to the shock, it is only the conditional correlation of their cycles that become negatively correlated.

²⁶Note that the unconditional correlation is the same across all the sign restricted solutions as this quantity is derived from the same reduced form.

In the policy sphere, discussion of a positive credit gap is taken as an indication of excess credit. The dynamics response of the credit cycle and the output gap to an identified financial shock provides some illustration of how this occurs in the setting of our identified financial shock.²⁷ In such settings, the trend of credit is conceptually related to some notion of a sustainable or normal level of credit. In response to a financial shock, the trend of credit falls, indicating that in response to the shock, the level of sustainable, or what one would not regard as excessive, credit falls. Because credit is slow to adjust, a positive credit gap opens up. As output falls faster, and its trend remains relatively stable in response to the shock, a negative output gap quickly opens up.

Summing up, a key conclusion is that *unconditionally*, the credit cycle and output gap are strongly positively correlated, with the contemporaneous correlation much stronger (0.32) relative to the correlation between the lagged credit cycle and the contemporaneous output gap (0.01). However, when we condition on a financial shock, the lagged credit cycle is strongly negatively correlated with the output gap (about -0.7). 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. We uncover that there is a degree of conditionality when considering how the business and financial cycle may be negatively correlated. The conditionality of our statement does feature a flavor of Jordà, Schularick, and Taylor (2013), which suggests that it is excessive credit in response to loose financial conditions which do really spillover to the real economy. Based on our analysis, the average boom in the business cycle will be associated with a boom in the financial cycle and vice versa. More broadly, our results would at least suggest that macroprudential policy targeted at crude measures of credit cycles, may be too blunt of an instrument since one should not *a priori* expect all positive deviations of the financial cycle relative to trend to be associated with business cycle busts.

6 Robustness

We briefly discuss some of the following robustness issues, though relegate the presentation of these results to the online appendix.

Shifts in mean We explore two possibilities for a shift in the mean, or μ in Equation (5), as this may affect the estimation of the cycle. First, we explore the possibility of a sharp break in the drift of real GDP as this has a first-order implication for the measurement of the business cycle. When we 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 is sensitive to how we adjusted for the standard errors when testing for the break. An alternative specification dates a break in 2006Q2, which is consistent with wider work (e.g. Berger, Everaert,

²⁷We are careful to note that this discussion only pertains to our identified financial shock as these dynamics do not necessarily carry over to other types of shocks.

and Vierke, 2016; Eo and Morley, 2020; Kamber, Morley, and Wong, 2018) dating a slowdown in GDP growth just before the Great Recession.²⁸ 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 allowed for a break in the drift of U.S. GDP in 2006Q2, and present these results in Section E of the online appendix, but note that our main results are robust. Second, the breaks in may not be discrete. We therefore allowed for the possibility of a smooth change in the mean of all variables. Taking guidance from Stock and Watson (2012), we demeaned all variables before estimation by using a biweight kernel with a bandwidth of 100 quarters before estimation. We also present these results in Section E of the online appendix, but note that our results are also robust to this choice of demeaning.

Informational Sufficiency We checked if our model is informational sufficient. Taking guidance from Forni and Gambetti (2014), we constructed a factor by extracting the first principal component from FRED-QD, and tested whether this extracted factor Granger cause any of the 23 BVAR equations in an out-of-sample forecasting exercise. Using the procedure described by Clark and West (2006) to test for predictability in nested models, we did not find evidence that the extracted factor from the FRED-QD dataset Granger causes any of our VAR variables in an out-of-sample forecasting exercise, suggesting our 23 variable BVAR system is informational sufficient.

Disentangling Uncertainty from Financial Shocks It is known that it is challenging to disentangle the role of uncertainty shocks from financial shocks. Similar to Caldara, Fuentes-Albero, Gilchrist, and Zakrajšek (2016) and Furlanetto, Ravazzolo, and Sarferaz (2019), we also attempted to disentangle the role of uncertainty shocks from financial shocks. In the penalty function identification, this is a similar exercise to Caldara, Fuentes-Albero, Gilchrist, and Zakrajšek (2016) when they reverse the order of identifying uncertainty shocks first before financial shocks.²⁹ In the sign restriction setting, we identify an uncertainty shock using the same sign restriction as the financial shock, except that for the uncertainty shock, the ratio of the increase in the VIX relative to the excess bond premium is larger than the financial shock. This is effectively the same exercise to Furlanetto, Ravazzolo, and Sarferaz (2019) who attempt this disentanglement by imposing a sign restriction on the ratio of the VIX to excess bond premium. We present these results in Section F of the online appendix, but just briefly comment on the results. In the penalty function setting, it is not entirely surprising that the results are sensitive to reversing the order since Caldara, Fuentes-Albero, Gilchrist, and Zakrajšek (2016) already document sensitivity when using the VIX to identify uncertainty

²⁸To 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.

²⁹Note that because we only identify a single shock in our penalty function exercise, namely the financial shock, the role of financial shocks will be identical to a setting where one first identifies the financial shocks, then uncertainty shock using the penalty function identification.

shocks. Nonetheless, the *sum* of the effect of financial and uncertainty shocks appear to be quite similar either when one first identifies the financial shock then uncertainty shock, or vice versa. Given the sum of the shares is quite insensitive, it suggests that if one was inclined to take guidance from the penalty function identification, while pinning down the role of financial or uncertainty shocks might be tricky, the general conclusions hold if we are prepared to group the two shocks together. In the sign restriction setting, identifying a second uncertainty shock does not affect our main conclusion. In fact, the role of the identified financial shock on the output gap is almost identical between our baseline results identifying a single financial shock, or the alternative of jointly identifying both uncertainty and financial shocks.

Choice of Particular Financial and Housing Indicators We also explored using loans, rather than credit, and house prices from the OECD and Federal Housing Finance Agency, rather than the BIS in our baseline analysis. Note that some of these alternative data sources may cause mismatches with our baseline sample. These results are also presented in Section G of the online appendix. Our main results are also robust to the change in the choice of the particular financial and housing indicators we use for the empirical analysis.

7 Conclusion

Building off a standard BVAR in conjunction with the Beveridge-Nelson decomposition, we jointly estimate the U.S. business and financial cycle within a unified approach which also allows us to interpret the estimated business and financial cycles through the lens of the forecast errors or structural shocks. First, we find that the role of financial factors in driving the business cycle appears to be much larger since the 2000s. We find this result regardless of whether in the more reduced form informational decomposition, or when we identify a structural financial shock. In particular, we find evidence that the financial sector did overheat the business cycle in the 2000s before the Great Recession, with our structural analysis pointing towards financial shocks adding as much as between 2 to 4% to the output gap during the 2000s. We also uncover evidence of a more nuanced relationship between the credit cycle and the output gap. In particular, we show that the unconditional correlation between the credit cycle and the output gap is positive, but negative when conditioned on a financial shock. One implication of our findings is that macroprudential policy may need to distinguish between the underlying causes of the credit cycle rather than relying on simple rules of thumbs that prescribe unconditionally curbing all positive fluctuations in the credit cycle.

Our framework provides several interesting avenues for future work given the ability to interpret multiple cycles and also linking these fluctuations to identified shocks. In particular, while we have restricted our analysis to the U.S., one could extend our framework to understanding financial and business cycles across multiple economies. In particular, extending work such as Miranda-Agrippino and Rey (2020) to jointly modeling financial cycles across multiple economies and also teasing out whether financial cycles comove across multiple economies, and

if so what causes such comovement, would be an interesting avenue to pursue.

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.
- BEN ZEEV, N. (2018): “What can we learn about news shocks from the late 1990s and early 2000s boom-bust period?,” *Journal of Economic Dynamics and Control*, 87(C), 94–105.
- 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.
- BERGER, T., J. MORLEY, AND B. WONG (2020): “Nowcasting the Output Gap,” *Journal of Econometrics*, forthcoming.
- 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.
- BILLIO, M., M. DONADELLI, G. LIVIERI, AND A. PARADISO (2020): “On the role of domestic and international financial cyclical factors in driving economic growth,” *Applied Economics*, 52(11), 1272–1297.
- 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.
- CANOVA, F. (1998): “Detrending and business cycle facts,” *Journal of Monetary Economics*, 41(3), 475–512.
- 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.
- CLARK, T. E., AND K. D. WEST (2006): “Using out-of-sample mean squared prediction errors to test the martingale difference hypothesis,” *Journal of Econometrics*, 135(1-2), 155–186.
- 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.
- COMIN, D., AND M. GERTLER (2006): “Medium-Term Business Cycles,” *American Economic Review*, 96(3), 523–551.
- CONSTANTINESCU, M., AND A. D. M. NGUYEN (2020): “A century of gaps,” *Economic Modelling*, forthcoming.
- DE WINTER, J., S. J. KOOPMAN, AND I. HINDRAYANTO (2020): “Joint decomposition of business and financial cycles:evidence from eight advanced economies,” *mimeo*.
- 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.
- DREHMANN, M., AND J. YETMAN (2020): “Which credit gap is better at predicting financial crises? A comparison of univariate filters,” *International Journal of Central Banking*, forthcoming.
- EO, Y., AND J. MORLEY (2020): “Why Has the U.S. Economy Stagnated since the Great Recession?,” *The Review of Economics and Statistics*, forthcoming.
- 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,” *Review of Economic Dynamics*, forthcoming.

- 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.
- HAMILTON, J. D. (2018): “Why You Should Never Use the Hodrick-Prescott Filter,” *The Review of Economics and Statistics*, 100(5), 831–843.
- INOUE, A., AND L. KILIAN (2020): “Joint Bayesian Inference about Impulse Responses in VAR Models,” Working Papers 2022, Federal Reserve Bank of Dallas.
- 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.
- 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.
- KULISH, M., AND A. PAGAN (2019): “Turning point and oscillatory cycles,” CAMA Working Papers 2019-74, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, The Australian National University.
- 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.
- MIRANDA-AGRIPPINO, S., AND H. REY (2020): “U.S. Monetary Policy and the Global Financial Cycle,” *The Review of Economic Studies*, 87(6), 2754–2776.
- 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.
- NEWKEY, 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.
- SCHÜLER, Y. S. (2020): “On the credit-to-GDP gap and spurious medium-term cycles,” *Economics Letters*, 192, 109245.
- STOCK, J. H., AND M. WATSON (2012): “Disentangling the Channels of the 2007-09 Recession,” *Brookings Papers on Economic Activity*, 43(1 (Spring)), 81–156.
- 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.