

FEEDBACKS: FINANCIAL MARKETS AND ECONOMIC ACTIVITY

MARKUS BRUNNERMEIER, DARIUS PALIA, KARTHIK A. SASTRY, AND CHRISTOPHER A. SIMS

ABSTRACT. We examine the relation among measures of credit expansion, measures of financial market stress, and standard macroeconomic aggregates. We use a form of structural VAR with monthly data on 10 variables. The model explains observed variation as driven by 10 mutually independent disturbances. We identify the shocks from variation across time in their relative variability. One of them emerges as representing monetary policy. We find two distinct financial stress shocks, suggesting that attempts to create a one-dimensional index of financial stress may be misguided. Our results are consistent with the finding by others of a negative reduced form relation between household credit expansion and future output growth at certain frequencies and lags. Most variation in credit growth is nonetheless positively associated with output growth. We find the output growth decline following household credit innovations to be explained by the monetary policy response to the inflation and initial positive output growth that accompanies the credit expansion. In pseudo-out-of-sample forecasting tests, neither credit aggregates nor rate spreads provide much advance warning of the 2008-9 crisis, though they do help in recognizing the severity of the recession once it has begun.

I. INTRODUCTION

In the long run, credit aggregates tend to expand with GDP, and indeed expand faster than GDP, so that the ratio of credit to GDP is larger in rich countries and tends to grow over time. In studies of economic development, the ratio of credit to GDP is sometimes used as a measure of “financial depth,” which is thought to contribute positively to economic growth.¹ On the other hand a number of recent studies, among them Mian, Sufi and Verner (2017), Schularick and Taylor (2012), and Jordà, Schularick and Taylor (2014), claim to have demonstrated a predictive relation between rapid growth of credit and future low GDP growth or higher likelihood of crisis.

Monetary policy has strong effects on GDP growth and also, unsurprisingly, strong effects on credit growth and on spread variables that measure financial stress. Monetary

Date: December 6, 2017.

This document is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License. <http://creativecommons.org/licenses/by-nc-sa/3.0/>.

¹Some early perspectives on this topic are found, for example, in Shaw (1973), McKinnon (1973), and Goldsmith (1969). A summary of the related literature is available in World Bank (2012), p.23-25.

policy also plausibly responds to rapid credit growth or contraction and to changes in spreads. To understand the policy implications of correlations or predictive regressions relating financial variables to GDP growth, it is essential that we understand the extent to which these correlations are generated by, or mediated by, monetary policy itself.

The structural VAR literature on monetary policy effects succeeded in separating two channels of relation between inflation and interest rates — policy-generated changes in interest rates tend to reduce inflation, while interest rates on average endogenously rise with inflation to compensate investors for inflation-generated losses. For the reasons we have listed here, it seems likely that there are multiple causal channels connecting spreads, credit aggregates and business activity, and that some of these channels operate with opposite signs. It therefore seems appropriate to estimate a multiple equation model connecting these variables and to imitate if possible the structural VAR literature's approach to unraveling feedbacks in the data.

Much of the existing empirical literature in this area has used short lists of variables and has not attempted to distinguish several channels of interaction between financial variables and the macroeconomy, including the one modulated by monetary policy. Studies of the predictive power of credit growth have primarily used single-equation projection methods (e.g., Mian, Sufi and Verner (2017), Jordà, Schularick and Taylor (2014), and Jordà, Schularick and Taylor (2015)) or binary outcome (i.e., crisis or no crisis) predictive models (e.g., Schularick and Taylor (2012) and Drehmann and Juselius (2014)).² Studies focused on the information in credit spreads have looked extensively at single-equation models (e.g., Lopez-Salido, Stein and Zakrajsek (2015), and Krishnamurthy and Muir (2016)) and reduced form multi-equation models (Gilchrist, Yankov and Zakrajšek (2009) and Gilchrist and Zakrajšek (2012)). Gertler and Karadi (2015) and Caldara and Herbst (2016) introduce credit spread variables into structurally identified, multiple-equation

²Mian, Sufi and Verner (2017) is unique among these for using data outside of identified "crisis episodes." It also contains a small-scale multivariate example, with three variables (real GDP, household credit to GDP ratio, and business credit to GDP ratio), but does not endogenize interest rate dynamics or separately identify monetary policy.

frameworks with monetary policy. But these authors have a narrower focus on identifying and interpreting monetary policy shocks, relative to the rest of the system, and do not discuss the role of credit aggregates.³

There have been other studies in this area based on fully interpreted structural dynamic stochastic general equilibrium models, which of course have included estimated effects of monetary policy.⁴ These DSGE models, though, have not considered as many financial variables jointly as we consider here and have imposed more, and more arguable, identifying restrictions than we impose here.

Our model uses monthly data on industrial production (IP), the personal consumption expenditure deflator (P) household credit (HHC) business loan credit (BC), money supply (M1), the federal funds rate (R), a commodity price index (PCM), the 10 year over 3-month Treasury term spread (TS), the Gilchrist and Zakrajšek (2012) corporate bond spread (GZ) and the 3-month Eurodollar over Treasury spread (ES).⁵ The sample period runs from January 1973 to June 2015.

We use the identification-through-heteroskedasticity approach pioneered in economics by Rigobon (2003). This approach assumes that the pattern by which disturbances feed through the economy is stable across time, but that the relative sizes of the independent sources of structural disturbance in the system vary across historical periods. We began modeling time variation in disturbance variance because it is so clearly needed to accurately describe financial variables and some macroeconomic variables. We discovered as we proceeded that we obtained stable, interpretable results from this assumption alone, without the need for the short or long run restrictions on dynamics usually applied to structural VAR's. Details of the model specification are in Section II below.

³Krishnamurthy and Muir (2016) does look at both aggregates and spreads in the same framework. But their main specifications, single-equation models which can include interactions (non-linear transformations) of credit growth and credit spreads, do not solve the endogeneity problem.

⁴Christiano, Motto and Rostagno (2014), for instance, estimate a monetary DSGE model based on the contract enforcement friction of Bernanke, Gertler and Gilchrist (1999) and find that "risk shocks" which can be measured in observed credit spreads drive a significant portion of U.S. business cycle dynamics. The model uses data on credit spreads (BAA-AAA) and firm credit in addition to "standard" macro aggregates. Del Negro and Schorfheide (2013) provide a detailed comparison of the forecasting performance of this model, a standard Smets and Wouters (2007) DSGE model, and various reduced-form models.

⁵Details of the data and their sources are laid out in Section III below.

Here is a qualitative summary of our results. Details and quantitative results are in Section III below.

One of our 10 shocks, indeed the one with the most widespread effects across variables in the system, we interpret as a monetary policy shock. Responses to it match what is usually assumed about the responses to a monetary policy shock in SVAR models.⁶ This shock produces a sustained and non-trivial increase in the interbank spread (ES) variable.

There are shocks, distinct from the monetary policy shock, that move the GZ corporate bond and Eurodollar spread variables, and then later move IP in the opposite direction. These fit the idea that disturbances that originate in financial markets can have macroeconomic effects.

Several shocks generate substantial movement in household and business credit, and all but two of them move IP, if at all, in the same direction as the credit aggregates. This fits the idea that most movements in credit aggregates accompany expansion of activity and do not predict future slowdowns. There is a disturbance that moves HHC up, and then with a delay moves IP down, and another that moves BC up, and then with a delay moves IP down. But the downward movement in IP is small and barely statistically significant in both cases. There may be periods where these shocks are important, so that the credit expansion does predict future contraction in business activity, but a quantitative model that can identify such “bad” credit expansions and thereby allow a policy response would have to be multivariate to separate this component of credit growth. These disturbances increase inflation as well as credit, and initially increase IP, so it is natural that monetary policy responds to them by increasing interest rates. Our model implies that if the shocks were accompanied by monetary policy holding interest rates fixed, the decline in IP following the shocks would disappear, though of course at the cost of increased inflation.

⁶Though we pick the shock we label “monetary policy” by looking at the sign and shape of its impulse responses, this is not the same as the frequently applied “sign restriction” approach to SVAR identification. Sign restriction identification does not lead, even asymptotically, to point identification of responses, whereas our approach, if its assumptions are correct, does provide point identification.

We can also consider whether a monetary policy maker, who could identify a credit shock that predicted a future output decline, could use that knowledge to suppress the effect of the credit expansion. Since contractionary monetary policy is estimated to have negative effects on both credit and output growth, suppressing the credit expansion would only increase the subsequent output decline. Suppressing the predicted output decline via expansionary monetary policy would be feasible, but of course would increase, not decrease, the credit expansion.

In Section IV we conduct pseudo-out-of-sample forecasting experiments to see what predictive value arises from including the spread and credit aggregate variables in the system. We find that the model gives little advance warning of the 2008-9 crash, whether or not the financial variables are included, but that the model tracks the course of the recession considerably better when they are included⁷. Most of the improved tracking of the crisis period comes from including the spread variables, not the credit aggregates.

The limited predictive value for credit aggregates in our system may appear difficult to reconcile with the results of previous studies with smaller models that have found substantial predictive value for credit aggregates in forecasting future business activity or future crises. In Section V we show that results in the literature, where in smaller models credit expansion predicts slow or negative future output growth, are not inconsistent with our model. Our model implies that these small-model results are misleading if interpreted causally.

In Section VI we summarize variations on our model that we have investigated to check robustness of results, though the details of most of these variations appear in the online appendix.

II. MODELING FRAMEWORK

This section describes our empirical approach. The first two parts introduce our multivariate time series models, and the third part describes our Bayesian estimation method.

⁷This is similar to the conclusion of Del Negro and Schorfheide (2013), who compare a New Keynesian DSGE model with and without financial frictions of the form in Christiano, Motto and Rostagno (2014) and Bernanke, Gertler and Gilchrist (1999).

II.1. **The Basic Model.** We specify structural vector auto-regressive (SVAR) models with variances changing at exogenously specified dates. They can be described by the system of dynamic stochastic equations

$$A_0 y_t = \sum_{j=1}^p A_j y_{t-j} + C + \epsilon_t \quad (1)$$

where y_t is an $n \times 1$ vector of observed variables, A_0 is an $n \times n$ matrix which determines simultaneous relationships, the A_j are $n \times n$ matrices of coefficients at each lag j , C is an $n \times 1$ vector of constants, and ϵ_t is a vector of independent shocks. In the base model, these are Gaussian (normally distributed).

We exogenously separate the time span $\{1 \dots T\}$ into M subperiods and set

$$\mathbb{E} [\epsilon_t \epsilon_t'] = \Lambda_m \quad \text{if } t \text{ is in period } m \in \{1 \dots M\} \quad (2)$$

where Λ_m is a diagonal matrix. The variance of the structural shocks changes across periods, but the dynamic relationship among the variables, as determined by A_0 and the A_j , remain fixed. In different terms, the impulse responses to structural shocks will have the same shape across variance periods, but their scales will vary. Our choice of variance regimes in estimated models (discussed in Section III and presented in Table 2) is motivated by observed variation in the time series and outside knowledge about policy changes.

We could fairly easily have allowed for regime changes to evolve as a Markov-switching stochastic process, as in Sims and Zha (2006). However, so long as the regimes are persistent, few in number, and well-determined by the data, inference about the model's dynamics is not likely to be strongly affected by conditioning on the regime switch dates as if known. Of course it is plausible that the variance regime switches are not only random, but endogenously determined. Allowing for that would greatly complicate the model and, since the regime switches are few in the data, might leave the nature of the regime switch endogeneity ill-determined by the data. We leave this to future research.

Our set-up can also be illustrated in the reduced form,

$$y_t = \sum_{j=1}^p B_j y_{t-j} + D + u_t \quad (3)$$

with

$$\mathbb{E} [u_t u_t'] = A_0^{-1} \Lambda_m \left(A_0^{-1} \right)' \quad \text{if } t \text{ is in period } m \in \{1 \dots M\} \quad (4)$$

Some normalization is required, as we could multiply the rows of A_0 and Λ by scale factors without changing the implied behavior of the data. We impose the restriction

$$\frac{1}{M} \sum_{m=1}^M \lambda_{m,i} = 1 \quad \forall i \in \{1 \dots n\}$$

where $\lambda_{m,i}$ is the i th diagonal element of Λ_m . This makes the cross-period average structural variance 1 in each equation. It can be shown that, given such a normalization and the technical condition that each pair of equations differs in variance in at least one period, we can uniquely identify all n^2 parameters of A_0 (up to flipping the sign of an entire row, or permuting the order of rows).⁸ Thus the variance switching eliminates the need for arguable linear restrictions on the $(A_i)_{i=0}^p$, such as “short-run” restrictions on contemporaneous responses in A_0 or “long-run” restrictions on the sums of coefficients in the $(A_i)_{i=1}^p$. Avoiding such restrictions on A_0 , while still maintaining full identification, is a particularly appealing feature for a model including financial time series which would plausibly react to all shocks immediately at the monthly or lower frequency.

While under the model’s assumptions the impulse responses of the system will be consistently estimated in large enough samples, the model does not give names to the shocks

⁸The intuition is that if Σ_j is the reduced form residual covariance matrix for period j , the expression

$$\Sigma_i^{-1} \Sigma_j = A_0' \Lambda_i^{-1} \Lambda_j \left(A_0^{-1} \right)'$$

has the form of an eigenvalue decomposition, with the columns of A_0' the eigenvectors. As long as the eigenvalues, the diagonal elements of $\Lambda_i^{-1} \Lambda_j$, are unique (i.e., there is no k, l such that $\lambda_{j,k} / \lambda_{i,k} = \lambda_{j,l} / \lambda_{i,l}$), the rows of A_0 are therefore uniquely determined up to scale once we know Σ_i and Σ_j .

A more formal proof of this can be found, for instance, in Lanne, Lütkepohl and Maciejowska (2010).

that drive it. In our later analysis, we pick out one of our 10 estimated shocks as a monetary policy shock, and two others as reflecting disturbances originating in financial markets. Our choices of names for these shocks reflect *a priori* assumptions about what shocks with these names should look like, in terms of the responses they generate. This is similar in spirit to the “sign restriction” approach to SVAR identification. However sign restrictions on impulse responses by themselves do not provide point identification, whereas our approach does allow point identification.

Aside from flexible identification of impulse responses, another benefit of our method is efficiency. For the same reasons as standard GLS (Generalized Least Squares), our method prevents periods of large shocks from inefficiently dominating the likelihood.

II.2. Non-normality in the error distribution. The previous section’s correction for heteroskedasticity will work best if volatilities mainly change between persistent episodes or regimes. But it does not allow for the possibility of a few isolated large disturbances or outliers. For instance, the bankruptcy of Lehman Brothers in October 2008 and the 650 basis point drop in the Federal Funds rate from April to May 1980 generate outliers of around 6 standard deviations that do not disappear when we allow variance-regime switches. To guard against such large shocks distorting inference, we consider a specification in which structural errors ϵ_t have Student t distributions.⁹

In the model notation, we can introduce random parameters $\xi_{i,t}$ such that

$$\epsilon_{i,t} \sim \text{Normal}(0, \lambda_{i,t} \xi_{i,t}) \quad (5)$$

⁹We also estimated a version of the model assuming a finite scale-mixture of normals as the distribution of the residuals. That model produced very similar results to those we display and discuss from the t -distributed errors model, and had lower likelihood. (Both the t and mixture-of-normals models had much higher likelihood than models that assumed Gaussian errors.) The mixture-of-normals assumption has been used in the time series literature to better model large movements in macro variables. Lanne and Lütkepohl (2010) introduce a maximum likelihood approach to estimating a discrete normal mixture SVAR model, and Chiu, Mumtaz and Pinter (2015) describe a Bayesian Gibbs sampling algorithm with an application to a model with stochastic volatility for U.S. data. Chib and Ramamurthy (2014) present a Gibbs sampling method for estimating a DSGE model with t -distributed shocks and Cúrdia, Del Negro and Greenwald (2014) find that the assumption improves the fit of a New Keynesian DSGE model that already includes low-frequency volatility changes.

We can also think of these objects as shocks which capture, in a simple way, a high-frequency component of volatility that has no persistence across time or correlation across equations.

Giving the ζ_{it} an inverse-gamma distribution, i.e.

$$\zeta_{i,t} \sim \text{Inverse Gamma}(\text{shape} = \alpha/2, \text{rate} = 2/\alpha). \quad (6)$$

implies that each $\epsilon_{i,t}$ has an independent Student- t distribution with α degrees of freedom and unit scale.¹⁰ We chose the degrees of freedom of the t distribution to be 5.7, by fitting the sample distribution of residuals for the Gaussian-errors model.¹¹

II.3. Econometric Methodology. Equations (1) and (2), combined with the normalization of variances, describe a model with n^2 free parameters in A_0 , $(M-1)n$ free parameters in the Λ_m , and n^2p free parameters in the A_j . The model with t -distributed disturbances has another nT parameters.¹² We use Bayesian methods to update beliefs about the parameters conditional on observed data $\{y_1 \dots y_T\}$ and initial conditions $\{y_{-p-1} \dots y_0\}$.

For A_0 we specify independent Gaussian priors on all elements, centered around 100 times the identity matrix, with standard deviation 200. For $\lambda_{\cdot,i} = \{\lambda_{1,i} \dots \lambda_{M,i}\}$, the vector of variances in each equation i , we put a Dirichlet prior (with $\alpha = 2$) on $\lambda_{\cdot,i}/M$. This restricts each of the relative variances to lie in $[0, M]$ (where $M = 10$ in our main model), centers the prior on equal variances, and enforces our normalization that for each structural shock the relative variances average to one across periods. We use a variation of the ‘‘Minnesota prior’’ described in Sims and Zha (1996) on the reduced form parameters in the matrices B_j and D of equation (3). These priors, described in more detail in Appendix

¹⁰The appendix reports results from an alternative case with $\zeta_{i,t}$ as independent k -multinomial, so that the distribution is a finite mixture of normals.

¹¹All that matters to the likelihood is the shape of these distributions, not their scale, since A_0 can absorb differences in scale. However because our prior on A_0 is not scale invariant, results might have been slightly different if we had used the fitted scale for the t -distributed shocks (.78), instead of the unit scale.

¹²As mentioned previously, the ζ_{it} can equally well be called parameters or shocks. The fact that there are so many of them does not mean they cause difficulties in estimation, because each has a specified distribution. This is a special case of the general point that Bayesian inference treats parameters and shocks symmetrically. They are all unknown objects with distributions.

1, center belief loosely around independent random walks in each variable. They also imply that constant terms should not interact with near unit roots to imply rapid trend growth and that, if the dynamics are stationary, initial conditions should not be too far from the model's implied unconditional means. Conditional on A_0 , these priors imply Gaussian priors on the A_j matrices.

In the case with normal structural shocks, we sample from the posterior distribution for these parameters in a two-step process that exploits the fact that, conditional on knowing A_0 and $\{\Lambda_1 \dots \Lambda_M\}$, we can treat (1) as a system of n independent linear regressions which can be estimated with weighted least squares. In the first step of the sampling process, we use a Random Walk Metropolis algorithm to sample the elements of A_0 and $\{\Lambda_1 \dots \Lambda_M\}$ using the likelihood integrated over the A_j (which is available analytically). Then, after drawing a large MCMC sample from the marginal posterior distribution of A_0 and $\{\Lambda_1 \dots \Lambda_M\}$ we can, for each of the A_0, Λ draws, sample from the coefficients in the A_j which have a known conditional normal distribution.

The model with t -distributed shocks requires a more involved method, a (non-standard) Metropolis-in-Gibbs algorithm. The first part of the algorithm, a Monte Carlo update of A_0 and $\{\Lambda_1 \dots \Lambda_M\}$ conditional on the $\tilde{\zeta}_{i,t}$ and integrated over the $(A_j)_{j=1}^p$, is just as in the normal shocks model. But now we need to draw $(A_j)_{j=1}^p$ after each draw of (A_0, Λ) to form implied normalized residuals $\epsilon_{i,t}$. The $\tilde{\zeta}_{i,t}$, conditional on the $\epsilon_{i,t}$ are distributed independently across i and t , allowing us to draw directly from their exact conditional posterior distribution. The process is repeated recursively. Appendix 2 describes the mechanics of the both algorithms in more detail, and Appendix 3 reports diagnostic evidence of its convergence to randomly sampling the model posterior.

For a given set of model parameters, we could change the sign of the coefficients in an equation (a row of (A_0, A_1, \dots)) or change the order of the equations (permute the rows of (A_0, A_1, \dots)), without changing the implied distribution of the data. The likelihood maximum therefore recurs through the parameter space at every permutation or

| | |
|-----|---|
| IP | Industrial production |
| P | Personal consumption expenditures price index |
| HHC | Sum of commercial bank real estate and consumer loans |
| BC | Commercial bank commercial & industrial loans |
| M1 | M1 money supply |
| R | Federal funds rate |
| PCM | CRB/BLS spot (commodity) price index |
| TS | Term spread of 10 year over 3 month Treasuries |
| GZ | Gilchrist and Zakrajsek (2012) bond spread |
| ES | “TED spread” of 3-month Eurodollars over 3 month Treasuries |

TABLE 1. Data series used in model estimation.

sign change of the parameters. This means that a complete MCMC sampling of the posterior distribution would show identical impulse response distributions for all shocks, all centered at zero response — but only if the prior itself were invariant to permutations of the orderings or signs of the equations.

Our prior, because it puts positive prior means on the diagonal elements of A_0 , is not invariant to permutations and scale changes of equation coefficients. As a result, we find no indication that our posterior sampling scheme is distorting results by not eliminating draws that are permutations or sign-switches of each other. Nonetheless these methods, if applied on data for which identification did not emerge as strongly, might need to test for and eliminate permuted or sign-switched models.¹³

This probability model implies prior and posterior distributions for all (potentially non-linear) transformations of the coefficients, including the reduced form coefficients B_j and the impulse response functions for variable i to each shock j . In all reported results, following Sims and Zha (1999), we report horizon-by-horizon 68% and 90% posterior density regions as “error bands.”

III. DATA AND RESULTS

Our main specification uses monthly data on 10 time series (listed in Table 1) from January 1973 to June 2015. We include data from the 1970s because they are a valuable source of variation in the time series and because correction for time-varying variance can account for what otherwise might be interpreted as regime change in monetary policy (e.g., as in Sims and Zha (2006)). The lag length p in our model is set to 10.

Our measures of “household” and “business” credit are based on the Federal Reserve’s weekly surveys of U.S. commercial banks.¹⁴ These data are different from the quarterly and annual series, based on a more comprehensive survey of lenders and categorized based on the borrower type (including “households and non-profits,” “nonfinancial noncorporate business,” and “nonfinancial corporate business”), used in some other research.¹⁵ Appendix 7.2 includes a more detailed discussion of the differences. Although our “household credit” series includes commercial real estate loans (which cannot be separately identified for the entire sample in the data) and our “business credit” data seems to have more high-frequency variation than the corresponding quarterly series, we believe these data capture the majority of the low-frequency behaviors that are critical for existing empirical evidence of their forecasting power.¹⁶

The inclusion of three credit spreads (of interest rates over short-term Treasuries) is meant to capture several possible dimensions of credit market stress: the term spread captures inflation expectations and uncertainty about future movements in fundamentals, the bond spread captures tightness in business financing, and the TED spread captures tightness in bank financing. The first was also expected to, along with the Federal Funds rate, M1, and commodity prices, provide a sharper identification of a monetary policy

¹³This is a special case of the kind of normalization issue discussed by Hamilton, Waggoner and Zha (2007).

¹⁴These are published in the H.8 “Assets and Liabilities of Commercial Banks in the United States” release.

¹⁵In particular, the cross-country database, assembled by the Bank of International Settlements uses these quarterly data.

¹⁶One practical complication is dealing with breaks in the credit series introduced by changes in accounting standards or major entrances to or exits from the commercial bank industry. Our specific calculations for eliminating these breaks, which are particularly large in the real estate credit series, are detailed in the Appendix 7.1.

| | Start | End | Description |
|---|----------|----------|--|
| 1 | Jan 1973 | Sep 1979 | Oil crisis, stagflation, and Burns Federal Reserve |
| 2 | Oct 1979 | Dec 1982 | Volcker disinflation |
| 3 | Jan 1983 | Dec 1989 | Recovery from early 1980s recession |
| 4 | Jan 1990 | Dec 2007 | Great Moderation and Greenspan Federal Reserve |
| 5 | Jan 2008 | Dec 2010 | Great Recession |
| 6 | Jan 2011 | Jun 2015 | Zero Lower Bound, Recovery from Great Recession |

TABLE 2. Dates for variance regimes in full model specification.

| Model | MDD |
|---------------------|----------|
| Gaussian, full data | 51768.38 |
| t , full data | 52490.13 |

TABLE 3. Marginal data densities (marginal likelihoods) for two models of the full data sample (1973:1 to 2015:6). Differences between values are log Bayes factors, or log posterior odds with equal prior weights on each model.

| Shock | Jan 1973 – Sep 1979 | Oct 1979 – Dec 1982 | Jan 1983 – Dec 1989 | Jan 1990 – Dec 2007 | Jan 2008 – Dec 2010 | Jan 2011 – Jun 2015 |
|-------|------------------------|------------------------|------------------------|------------------------|------------------------|------------------------|
| 1 | 1.255 | 1.336 | 0.767 | 0.807 | 1.306 | 0.610 |
| 2 | 0.856 | 0.946 | 0.902 | 0.875 | 1.190 | 1.119 |
| 3 | 0.676 | 0.533 | 0.413 | 1.700 | 1.675 | 0.934 |
| 4 | 0.814 | 0.950 | 1.179 | 1.165 | 1.241 | 0.622 |
| 5 | 0.319 | 0.645 | 0.482 | 0.595 | 2.190 | 1.721 |
| 6 | 1.031 | 4.147 | 0.585 | 0.109 | 0.087 | 0.004 |
| 7 | 1.275 | 0.832 | 0.760 | 0.722 | 1.729 | 0.575 |
| 8 | 0.940 | 2.553 | 0.864 | 0.521 | 0.781 | 0.375 |
| 9 | 0.764 | 0.484 | 0.533 | 0.717 | 2.823 | 0.632 |
| 10 | 1.565 | 1.917 | 0.543 | 0.286 | 1.707 | 0.009 |

 TABLE 4. Posterior median relative variances for each of ten shocks in six periods, from a model with t -distributed innovations.

shock, as policy-generated rises in the short rate might be expected to have little effect on, or even lower, long rates, if the monetary tightening does succeed in lowering inflation.

We separate the full sample into six variance regimes described in Table 2. From the standpoint of estimation efficiency, we expect the separate treatment of the Volcker disinflation and Great Recession to discourage overfitting of high monetary policy and financial stress variations respectively by allowing the model to “down-weight” these periods’ residuals.

Table 3 displays estimates of the marginal data densities of the model with Gaussian distributed shocks and the model with Student's t distributed shocks. These are reported in a log points scale, so a difference of over 100 is extremely strong evidence (i.e., more than an e^{100} odds ratio) in favor of the better model. The t model clearly fits much better than the model with Gaussian errors.¹⁷

Figures 1 to 4, in four 5 by 5 "blocks," show the impulse response over five years of all 10 variables to the model's orthogonal structural shocks, scaled to draws from a unit-scale t distribution with 5.7 degrees of freedom. Since the diagonal of Λ_i is normalized to sum to one across regimes, these responses are a kind of average across regimes.

The model, despite the lack of any identifying zero restrictions on coefficients, fits recognizable monetary policy (number 6) and credit spread (9-10) shocks with significant long-term real consequences. Shock 6 is the only one that has an immediate positive R response, a delayed negative IP response, a negative (though ill-determined) long run P response, negative responses of M1 and the two credit aggregates, and a negative response of the term spread (as would be expected if the shock raises current interest rates and lowers expectations of future inflation).

The ninth and tenth shocks are the most important sources of variation in the GZ spread and the ES spread, respectively. The two spreads do not tend to move together in response to these shocks, and the two have different patterns of effects on other variables. Both depress IP. Both depress P, though in the case of shock 10 this effect is statistically weak. Shock 9, which immediately impacts the GZ spread, has a strong delayed effect in depressing BC, but modest and indeterminate-signed effect on HHC, while shock 10, which immediately impacts ES, strongly depresses HHC with ill-determined effect on BC. Shock 10 produces an expansionary movement in R, while shock 9 does not. These patterns seem to fit an interpretation that distinguishes a banking credit shock (10) from a non-bank financial disturbance (shock 9). All the effects of these shocks on other variables

¹⁷See Appendix section 3.4 for more details.

are delayed, while their effects on the spread variables are immediate. This fits an interpretation that they reflect disturbances originating in financial markets, with monetary policy at most (with shock 10) trying to partially offset their effects.

The third and fourth shocks start with an impulse to household credit and to business credit (net of inflation), respectively. They lead to persistent long-term declines in output and seem to match the “excessive credit growth” story demonstrated empirically by Mian, Sufi and Verner (2017), Schularick and Taylor (2012), and others. Perhaps surprisingly, since they used annual panel data for many countries, the responses we find to shocks 3 and 4 are quantitatively quite similar to the responses to household and firm credit found by Mian, Sufi and Verner. We find the output response with 68% confidence bands is in both cases barely significantly less than zero at the five-year or smaller horizon, but, as in the Mian-Sufi-Verner estimates, they are substantively non-trivial. However, several other shocks to which the credit variables respond substantially (shocks 1, 6, 9 and 10) move output and credit in the same direction. Over medium and long horizons, observed credit growth is as likely to result from these shocks as from the model’s third and fourth shocks. Distinguishing “good” from “bad” credit growth, then, requires observing all ten variables in this model.

Note that shocks 3 and 4 also elicit a positive movement in R , which is not surprising given that they are associated with rising inflation and, in the case of the third shock, initially rising output. If we accompany this shock with a sequence of shock-6 values that suppress this R response, we find that the inflation and credit expansion are somewhat greater and there is no output decline. In other words, the model implies that the decline in output that follows this type of credit expansion can be accounted for by the systematic monetary policy contraction that the credit shock elicits. Figure 5 shows the responses to the third and fourth shocks, and to those shocks with policy keeping the interest rate constant. (The figure omits responses that were nearly the same for the two patterns of shock.)

Though shocks 3 and 4 are followed by future declines in output growth, they are not followed by substantial movements in the spread variables GZ and ES. The spread shocks, on the other hand, are followed by substantial declines in credit aggregates. Our interpretation is that the credit expansions generated by shocks 3 and 4 are followed with a delay by slow growth due to monetary tightening, not financial market distress.

The seventh shock accounts for a substantial component of variation in P. Its immediate effect is to increase commodity prices, and to some degree to increase the GZ spread. The effect on commodity prices is persistent. With some delay, P (the PCE deflator) moves up and IP moves down by a non-trivial, but statistically marginal, amount. Neither BC nor HHC moves much. This looks like a “supply shock” originating in commodity markets.

These core impulse response results seem largely robust to the alternative error specifications. We provide full impulse response plots for versions of the model with Gaussian errors and with errors modeled as a finite scale mixture of Gaussian errors in the appendix section 4.1.

The picture is also very similar if we estimate with data only up to December 2007. Full impulse responses from this sample period are plotted in the supplemental Appendix, section 4.2. In particular, the identification of monetary policy and spread effects is very stable. There is weak (within 68%, but not 90% bands) evidence of an output response to household and business credit expansion shocks. These effects are of comparable magnitude to the estimated credit effects in the model estimated on the full dataset.

Our results suggest that the variances of these shocks change substantially among periods. Table 4 reports the variances of each of the ten structural shocks in the posterior mode *t*-errors model, over the full sample. 90% probability bands for these relative variances are quite tight, mostly within 0.8 to 1.2 times the posterior median estimate. In general, there is strong evidence of time-varying variance. Several of the shocks spike in variance during the financial crisis (period 5). The sixth shock, which we identify as a monetary policy shock, has a considerably inflated variance in the Volcker disinflation period and almost zero variance in the most recent period (near the zero lower bound).

| Month | ϵ_{it} | $dy_{i,t}$ | Month | ϵ_{it} | $dy_{i,t}$ |
|---------|-----------------|------------|----------|-----------------|------------|
| Shock 1 | | | Shock 6 | | |
| 9/2008 | -9.875 | -0.043 | 5/1980 | -22.627 | -0.067 |
| 12/1974 | -6.259 | -0.027 | 3/1980 | 10.409 | 0.031 |
| 1/1983 | 5.160 | 0.022 | 2/1981 | -10.135 | -0.030 |
| 4/1980 | -5.015 | -0.022 | 5/1981 | 9.312 | 0.028 |
| Shock 2 | | | Shock 7 | | |
| 9/2005 | 5.492 | 0.007 | 8/1976 | -4.222 | -0.066 |
| 1/1990 | 3.820 | 0.005 | 7/1975 | 4.096 | 0.064 |
| 11/2005 | -3.748 | -0.005 | 7/1974 | 3.901 | 0.060 |
| 11/2008 | -3.453 | -0.005 | 2/1974 | 3.777 | 0.059 |
| Shock 3 | | | Shock 8 | | |
| 12/1999 | 7.068 | 0.015 | 8/1974 | -7.157 | -0.014 |
| 9/2008 | 5.922 | 0.012 | 6/1981 | 6.261 | 0.012 |
| 10/2003 | -5.861 | -0.012 | 11/2008 | 5.871 | 0.011 |
| 10/2002 | 4.933 | 0.010 | 4/1980 | 4.411 | 0.008 |
| Shock 4 | | | Shock 9 | | |
| 10/2008 | 5.247 | 0.019 | 10/2008 | 18.275 | 0.023 |
| 12/1986 | 5.216 | 0.019 | 1/2009 | -7.844 | -0.010 |
| 9/2007 | 4.454 | 0.016 | 9/2008 | 7.183 | 0.009 |
| 11/1999 | 4.283 | 0.016 | 7/2002 | 6.771 | 0.009 |
| Shock 5 | | | Shock 10 | | |
| 10/2001 | -9.766 | -0.047 | 10/2008 | 10.840 | 0.019 |
| 8/2011 | 8.507 | 0.041 | 7/1974 | 8.950 | 0.016 |
| 9/2001 | 8.431 | 0.041 | 9/2008 | 7.739 | 0.014 |
| 9/2008 | 6.773 | 0.033 | 12/1973 | 7.148 | 0.013 |

TABLE 5. Four largest residuals for each shock in the main model. Point estimates are posterior medians.

Table 5 lists the 4 largest posterior median shocks for each equation. These are in standard deviation units and not scaled by the corresponding λ values. They thus show the biggest shocks, not the biggest “surprises” for the model.

The biggest of these, in shock 6, reflects a sudden easing of monetary policy in May 1980, during the recession of that year. The Federal Funds rate fell from 18 per cent to 11 percent in that month, but soon started rising again. There were also large values for shock 6 in March 1980, February 1981, and May 1981. These were all during the period of Volcker unborrowed reserve targeting, which has a high value of λ for shock 6. That is, though large, these shocks occur in a period the model has identified as a high-variance period.

The second and third largest shocks were in the two financial stress indicators, shocks 9 and 10, at the time of the Lehman collapse in October 2008.

Shock 1, which accounts for much of the variance in output and looks like a “demand” shock, was sharply negative in September 2008, reflecting the large decline in industrial production as the crisis took hold.

The September 2001 attack on the US shows up in shock 5, which looks like an accommodated money demand shock. The shock is sharply negative in October 2001 as the Fed withdrew its temporary liquidity accommodation. It was nearly as large and positive in the previous month.

Shock 7, the PCM shock, takes on large values in 1974-5.

The largest of these shocks all correspond to events that were recognizably unusual as they occurred. But they tend to come from periods with large values of λ_{it} , so they are not necessarily the biggest surprises. Looking at the largest surprises — the $\epsilon_{it} / \sqrt{\lambda_{it}}$ values — is useful because it is these residuals that have the biggest impact on model fit. Also the size and distribution across variables of these large surprises casts some doubt on our assumption of i.i.d. t -distributed scaled shocks. We discuss the surprises in the appendix section 5.

The remainder of this section reviews the model dynamics of key shocks in greater detail.

III.1. The Credit Channel of Monetary Policy. The sixth ordered shock of the model, the impulse responses of which are collected in Figure 6, satisfies the description of a monetary policy shock in the initial impulse to the Federal Funds rate, initial decrease in the 10 year over 3 month Treasury term spread, and persistent negative impact on output. At the 68% level, there is still substantial uncertainty about the responses of both consumer and commodity prices, but point estimates from the posterior mode model show persistent declines. This identification comes despite the lack of any identifying restrictions on contemporaneous responses (or the monetary policy reaction function).

The model provides substantial evidence that the effect of short rate movements is amplified by corresponding movements in interest rate spreads. In particular, the 3-month Eurodollar spread over Treasuries increases about 1 basis point per 4 in the Federal Funds rate and decays over a similar horizon. There is relatively little evidence, in contrast, that the bond premium moves in the short run (i.e., within the first 6 months) and no evidence that the effect persists longer. The general finding of financial amplification of monetary policy shocks is consistent with the empirical results of Gertler and Karadi (2015) and a variety of theoretical models which suggest that risk premia should move in response to monetary policy (e.g., Drechsler, Savov and Schnabl (2016) and Brunnermeier and Sannikov (2016)). Our empirical result is focused, however, on *inter-bank* credit conditions separate from *firm-level* credit conditions. This is concurrent with our broader empirical point that the information content of credit spreads in the multi-variate model is multi-dimensional — movements in different spreads forecast different macro-financial dynamics and, potentially, relate to different economic mechanisms.

III.2. Spread Spikes and Early Warning. Three independent shocks, ordered eight to ten, can be identified by sharp increases in spreads at $t = 0$, but only the latter two (the responses to which are plotted in Figure 7) have significant output effects. The output effect is larger and more significant for the ninth shock, associated with an initial surge in the corporate bond spread and a long-term contraction in business credit. The tenth shock, in contrast, begins with a shock to the inter-bank lending rate (of comparable magnitude to the impulse following a monetary policy shock), a significant long-term contraction in household credit, and a modestly significant short-term output contraction.

The fitting of two independent stress shocks suggests the importance of a multidimensional approach to measuring financial stress. In their long-run macro implications, the two shocks can be distinguished by sharply different implications for credit aggregates and prices. The “bond spread shock” is associated with a persistent reduction in the consumer and commodity price level and a significant decline in loans to businesses. The “inter-bank shock,” in contrast, has no long-term price effect and predicts a significant

long-term contraction in household, not business, credit. It also seems to pick up some delayed monetary loosening, potentially in response to the generated recession.

In the historical record, inter-bank shocks have almost as high a variance in the early sample (1973 to 1982) as they do in the financial crisis (Table 4). With post-2009 interbank rates very close to short rates at zero, this channel almost completely shuts down in the final variance period. The corporate bond spread shock, in contrast, is by some margin highest variance during the 2008 financial crisis.

Taken together, the impulse response and the estimated variances suggest that the macro importance of spread shocks—closely related to the forecasting value of the spread variables—is concentrated during certain high variance episodes and largest at *short* horizons. Interest rates react to shocks 9 and 10 by falling, as would be expected from a monetary policy easing. Because the response of IP to monetary policy is slower and more persistent than its response to shocks 9 and 10, it is not clear that monetary policy could do more to mitigate the output decline without creating more instability.

III.3. Credit Growth and Recessions. Our main model offers some support, within 68% error bands, of the hypothesis that excessive growth in household credit can forecast negative long-term real output growth (Figure 8). The shape of our estimated output responses to shocks three and four, with a short-term output boost and long-term contraction, are similar to those found in a small-system (household credit to GDP, business credit to GDP, and real GDP) VAR by Mian, Sufi and Verner (2017). In our model the response to shock, 3 which raises HHC by about 1 per cent over five years, is a small increase in IP, on the order of 0.1 per cent, that lasts less than a year, followed by a decline that reaches -0.1% after five years. Mian, Sufi and Verner (2017) find a larger initial rise of about 0.2%, over two years, but also a subsequent decline that reaches about -0.1% after five years. Our estimates show a five-year negative response to the BC shock (4) that is similar in size and statistical significance to the estimated response to HHC, though without the initial positive response. This also is consistent with the Mian, Sufi and Verner (2017) business credit estimated responses. All these results have fairly wide

error bands, so there is no statistically sharp difference between the models. This might not have been expected, since their results use an international panel of annual data and a smaller model, while we use monthly US data and a big model.

However, as mentioned earlier in the qualitative discussion of results, our model implies that the decline in output growth following this shock can be entirely accounted for by the rise in interest rates it elicits. The response of the system to shock 3 or shock 4, combined with a sequence of shock 6 values that keep the interest rate constant, eliminates the decline in output (Figure 5).

A different way to assess the economic significance of credit shocks in our model is to calculate the share of forecast errors that are explained by each shock (Figure 9).¹⁸ The importance of third and fourth shocks for explaining credit variation starts very high (as it has by far the largest contemporaneous impact on credit) but decays over time. Over five years, these shocks explain 20% and 24% respectively of forecast error variance household and business credit. The remainder of credit variation over these horizons is explained by the other shocks in the model, which are all associated with credit and output moving in the same direction, or with output scarcely moving at all. Over five years, the same “bad credit shocks” explain 1.4% and 1.6%, respectively, of forecast error variance in industrial production.

In the period 1990-2007, when the variance of the credit shocks is relatively high and the variance of others relatively low, the two credit shocks explain a higher fraction of five-year-ahead variance in the two aggregates (49.5% for household credit and shock 3, and 37.6% for business credit and shock 4) and a higher percentage of output variation at the same horizon (4.1% and 3.2% respectively). These numbers are still economically quite small.

¹⁸These are the squared impulse responses scaled to sum to one for each response variable in each period. Precisely, the variance decomposition of variable i is, for each j and each time horizon s , the proportion of s -step ahead forecast error variance in variable i attributable to shock j .

IV. CREDIT CONDITIONS AND FORECASTING

So far we have demonstrated that credit variables have an interesting interpretation within the model. But are they practically helpful to include, and could this have been realized before the 2008 financial crisis? We find that information in spreads can be useful for short-term forecasting at the onset of a crisis. The model with spreads does not, however, provide much advanced warning of a crisis or any clear advantage in “normal” times outside of recessions.

IV.1. Forecasting in the Recent Financial Crisis. We first focus on the 2007-08 financial crisis and its immediate aftermath. At each month between January 2007 and December 2010, we estimate (posterior modes of) models with and without credit variables using data only up to that point and then calculate 12-month forecasts. This “pseudo-out-of-sample forecasting” exercise offers a dimension in which to compare models with different data lists and gives a sense of how much changing the emphasized data in macro models would have helped in real time. We focus on the Gaussian errors specification, despite its fitting more poorly than the t model, because it seems to capture the main model dynamics and is much easier to do recursive computations with.

Figures 10, 11, and 12 plot posterior mode forecasts from our (Gaussian error) model with 10 variables, a version without the credit aggregates, and a version without the spreads, respectively, at 3-month intervals from January 2007 to October 2010. The model without spreads (Figure 12) never fully “accepts” the crisis, predicting a return to near pre-crisis growth rates at each point during the deepest contraction. The models with spreads (with or without credit aggregates) give slightly less optimistic forecasts in early 2008, at which point the bond and inter-bank spreads have elevated slightly over mid-2000s levels. But the most obvious improvement is the models’ ability to grasp the severity of the crisis during the deepest fall from mid 2008 to mid 2009. This observation is consistent with the previous section’s analysis of impulse responses, which suggested that the model could identify spread shocks which have macro effects within the first few

months. The spreads provide little advance warning of severe recession but do enhance recognition of the severe recession, and its likely persistence, once it is underway.

The addition of credit aggregates seems considerably less important. With or without credit aggregates, the model is quicker to recognize a persistent downturn. While forecasts of IP are little affected by excluding credit aggregates, the model without them consistently predicts that interest rates will start reverting to positive values from the zero lower bound, though this seems to have limited effects on forecasted output or consumer prices.

IV.2. Forecasting Power in the Entire Sample. We generalize the exercise of the previous section by calculating forecasts with versions of the main model, the no spreads model, and the no credit model estimated up to each month from October 1979 to June 2015.¹⁹ We focus on root mean squared error (RMSE) for forecasts of all variables common to the models.

Figures 13 and 14 display the evolution of these RMSE for the base model with all variables (blue), a model without spreads (red), and a model without credit aggregates (green). As suspected from the previous section, the models with spreads does a significantly better job predicting output just before and during the 2007-2009 financial crisis and recession. The model's internal projections for the Federal Funds rate are quite a bit better *at* the zero lower bound, though this comes at the cost of one set of very poor forecasts right around the final major reduction in the rate in late 2008. Any advantages in forecasting the price level and credit aggregates in the crisis are less obvious.

Outside the recent financial crisis, and potentially the early 1980s and early 2000s recessions, the no spread model seems to perform just as well if not better. We might suspect that a formal or informal comparison of models before 2008 would not clearly support the inclusion of the financial variables, even if the estimated dynamics from such a model

¹⁹The truncation at the beginning of the sample comes from the requirement of having two variance regimes to identify the parameters. Unfortunately, this cuts out some interesting macroeconomic turbulence in the 1970s.

Additionally, for the period October 1979 to December 1982, we use models with six lags because of the smaller availability of data.

look like they have “economically interesting” transmissions from spreads to macro variables. Alternatively, the model with spreads might only be better for forecasting when economic conditions worsen.²⁰

The model without credit aggregates, but with credit spreads, seems to match the full model quite closely throughout the sample. One exception seems to be the early part of the 1981-1982 recession and the subsequent uptick in growth around 1984. In several periods, including post-recession growth in the early 90s and 2010s, the no credit model is significantly *better* at predicting output. In general there is no clear pattern of the model with credit aggregates, after including spreads, doing a better job of forecasting the timing or severity of U.S. recessions.

V. ACCOUNTING FOR THE PREDICTIVE POWER OF CREDIT IN SMALL MODELS

There is a recent literature that claims to show that credit expansion predicts negative growth and/or financial crises. Our model does not suggest that credit expansion predicts financial crises and shows the negative association of credit growth with future output growth as a small component of the overall relation between credit and output. But our model does not contradict the earlier results with smaller models. It is compatible with them.

As we have already observed, the responses to shocks three and four in our model lie within error bands of the responses to credit-to-gdp-ratio shocks in Mian, Sufi and Verner (2017). Those authors also estimate single-equation regressions of growth in real output over the next three years on growth in credit in the past three years, finding that a one-standard-deviation increase in 3-year household credit growth predicts a 2.1% lower growth rate of output in the subsequent three years. they find no such effect with business credit. When we replicate those regressions using our data, we find that a one standard deviation increase in the three-year growth rate of BC predicts a 1.7 percentage point

²⁰These nuances could be captured formally by taking posterior forecasts averaged across an “ensemble” of models, the weights on which change over time (for instance, with some approximation of posterior odds). To capture them within the model might require some more complex (and possibly endogenous) modeling of regime switching.

| | Proportion of negative effects | |
|---|--------------------------------|-----------------|
| | Household Credit | Business Credit |
| without lagged IP | 0.55 | 0.56 |
| with lagged IP | 0.53 | 0.57 |
| Proportion of economically significant negative effects | | |
| | Household Credit | Business Credit |
| without lagged IP | 0.29 | 0.31 |
| with lagged IP | 0.30 | 0.33 |

TABLE 6. Probability, in simulated draws of coefficients and data from the posterior distribution under the t -distributed errors model, of a negative or “economically significant” effect of previous 3-year credit growth on subsequent three-year output growth. An economically significant effect is one that makes a one-standard-deviation change in credit growth produce a 2 percentage point or larger change in output growth rate.

decline in the subsequent three year growth rate of output. We find no such predicted effect from increased household credit growth.²¹

So our data show a similarly strong negative effect of three-year credit expansion on output growth, but shifts the strong effect from household to business credit. It could be, though, that even though the data are in this way consistent with the MSV results, the large monthly model misses this aspect of the data.

To check this, we can replicate those regressions on data simulated from our model, and we find that with substantial probability our results from simulated data are of similar size and sign to those found in the actual data. The results are in Table 6. To form the table we used three-year forward differences of the log of real output (IP) as the dependent variable, three-year backward differences in real credit (credit over price level) over output as the main independent variable, and lags of first differences of log output as an extra independent variable. In equation form, with y_t denoting log IP, hc_t the ratio of household credit deflated by PCEPI to IP, and bc_t the ratio of business credit deflated

²¹When we substitute flow of funds data on business credit and household credit for the our HHC and BC, the effects of HHC and BC increases by one standard deviation are both negative, with HHC producing a 1.6 percentage point decline and business credit a 0.6 percentage point decline.

by PCEPI to IP,

$$y_{t+3} - y_t = \alpha + \beta_h(hc_t - hc_{t-3}) + \beta_b(bc_t - bc_{t-3}) + \sum_{i=1}^k \gamma_i(y_{t-i+1} - y_{t-i}).$$

Table 6 reports probabilities for β_h and β_b . The results “with lagged IP” set $k = 3$; otherwise we set all $\gamma_i \equiv 0$.

In these simulated draws from our model’s data distribution we find the probability of credit growth coefficients smaller than 0 to be over .5 and negative enough that a one standard deviation increase in credit growth reduces output growth by 2 percentage points or more with probability about .3. Note that in the simulated data, it is about equally likely that the BC or the HHC predictive effect will be large and negative. These suggests that, according to our model, this one-equation model could suggest strong predictive power when the fully specified structural model attributes a limited causal role to credit innovations.

VI. ROBUSTNESS

We have tried a number of variants on our model to check robustness of our results. These are described in detail in the appendix. Our checks include estimating quarterly models, estimating models with triangular cholesky identification, looking for nonlinearity via various nonlinear transformations of the data, and varying the assumption on error distributions. Details of these experiments are in the appendix. Of course the model we present as our main model is itself the result of experiments like this, where we have adopted specifications when we found them fitting better. As a result, the robustness checks in the appendix do not cast doubt on our main specification.

VII. CONCLUSION

Credit conditions, monetary policy, and real activity interact dynamically through multiple channels. To study these interactions, we construct and estimate structural multiple-equation models that are identified without strong *a priori* assumptions. Our analysis

distinguishes impulses and feedbacks that focused study of individual channels might miss.

Our main model includes ten independent shocks that are identified by substantially changing volatilities across exogenously specified regimes. The data strongly favor additional corrections for fat tails in the distributions of the structural innovations, though the main qualitative conclusions are the same without them. Further refining (and possibly endogenizing) a model specification for volatility remains a task for future research, but addressing the issue in some way greatly improves model fit and affects implied dynamics.

Monetary policy is identified without any timing restrictions and seems to be amplified through inter-bank credit spreads. Two other model shocks look like “stress shocks” which originate in the financial sector and propagate to the real economy after several months of delay. The distinction between these shocks, which start with impulses to corporate bond spreads and interbank rate spreads respectively, is potentially very important for emerging research on the role of lending frictions and risk premia in the macroeconomy. A related takeaway for forecasters is that one-dimensional metrics of financial conditions may be insufficient for capturing risks for the real economy.

While these credit spread shocks do have strong real effects, they do not provide more than a few months of “advance warning” of an output contraction. In recursive-out-of-sample forecasts around the 2008 financial crisis, including additional credit spread variables only improves forecasts in a narrow window at the beginning of the downturn. Across the entire data sample, there is no clear evidence that including credit variables improves forecasting performance.

Credit aggregates in this model mainly move “passively” in the same direction as output. Two shocks generate opposite movements in household (real estate plus consumer) or business credit and output, but in all periods the magnitude of these effects is relatively small. They are accompanied by rising interest rates, and if monetary policy offsets that rise, the effect of the credit shocks on output would disappear. To the extent that this

effect is quantitatively important, a multivariate model is necessary to properly separate it from other effects. Our results demonstrate how multiple time-series analysis, without strong *a priori* restrictions, can shed light on complex interactions among policy, financial markets, and the real economy.

REFERENCES

- Bernanke, Ben S., Mark Gertler, and Simon Gilchrist.** 1999. "The financial accelerator in a quantitative business cycle framework." In *Handbook of Macroeconomics*. Vol. 1 of *Handbook of Macroeconomics*, , ed. J. B. Taylor and M. Woodford, Chapter 21, 1341–1393. Elsevier.
- Brunnermeier, Markus K., and Yuliy Sannikov.** 2016. "The I Theory of Money." Princeton University Working Paper.
- Caldara, Dario, and Edward Herbst.** 2016. "Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVARs." Board of Governors of the Federal Reserve System Finance and Economics Discussion Series Paper 2016-049.
- Chib, Siddhartha, and Srikanth Ramamurthy.** 2014. "DSGE Models with Student-t Errors." *Econometric Reviews*, 33(1-4): 152–171.
- Chiu, Ching-Wai, Haroon Mumtaz, and Gabor Pinter.** 2015. "Forecasting with VAR models: fat tails and stochastic volatility." Bank of England Bank of England working papers 528.
- Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno.** 2014. "Risk Shocks." *American Economic Review*, 104(1): 27–65.
- Cúrdia, Vasco, Marco Del Negro, and Daniel L. Greenwald.** 2014. "Rare Shocks, Great Recessions." *Journal of Applied Econometrics*, 29(7): 1031–1052.
- Del Negro, Marco, and Frank Schorfheide.** 2013. "DSGE Model-Based Forecasting." *Handbook of Economic Forecasting*, , ed. Graham Elliott and Allan Timmermann Vol. 2, Part A of *Handbook of Economic Forecasting*, Chapter 2, 57 – 140. Elsevier.
- Drechsler, Itamar, Alexi Savov, and Philipp Schnabl.** 2016. "A Model of Monetary Policy and Risk Premia." *Journal of Finance*, NA((forthcoming)): NA.
- Drehmann, Mathias, and Mikael Juselius.** 2014. "Evaluating early warning indicators of banking crises: Satisfying policy requirements." *International Journal of Forecasting*, 30(3): 759–780.

- Gertler, Mark, and Peter Karadi.** 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics*, 7(1): 44–76.
- Gilchrist, Simon, and Egon Zakrajšek.** 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review*, 102(4): 1692–1720.
- Gilchrist, Simon, Vladimir Yankov, and Egon Zakrajšek.** 2009. "Credit market shocks and economic fluctuations: Evidence from corporate bond and stock markets." *Journal of Monetary Economics*, 56(4): 471 – 493.
- Goldsmith, Raymond.** 1969. *Financial Structure and Development*. Yale University Press.
- Hamilton, James D., Daniel F. Waggoner, and Tao Zha.** 2007. "Normalization in Econometrics." *Econometric Reviews*, 26(2-4): 221–252.
- Jordà, Òscar, Moritz Schularick, and Alan M. Taylor.** 2014. "The Great Mortgaging: Housing Finance, Crises, and Business Cycles." National Bureau of Economic Research Working Paper 20501.
- Jordà, Òscar, Moritz Schularick, and Alan M. Taylor.** 2015. "Betting the house." *Journal of International Economics*, 96, Supplement 1: S2 – S18. 37th Annual NBER International Seminar on Macroeconomics.
- Krishnamurthy, Arvind, and Tyler Muir.** 2016. "How Credit Cycles across a Financial Crisis." Stanford GSB Working Paper.
- Lanne, Markku, and Helmut Lütkepohl.** 2010. "Structural Vector Autoregressions With Nonnormal Residuals." *Journal of Business & Economic Statistics*, 28(1): 159–168.
- Lanne, Markku, Helmut Lütkepohl, and Katarzyna Maciejowska.** 2010. "Structural vector autoregressions with Markov switching." *Journal of Economic Dynamics and Control*, 34(2): 121 – 131.
- Lopez-Salido, David, Jeremy C. Stein, and Egon Zakrajšek.** 2015. "Credit-Market Sentiment and the Business Cycle." Board of Governors of the Federal Reserve System Working Paper.
- McKinnon, Ronald.** 1973. *Money and Capital in Economic Development*. Washington, D.C.:Brookings Institution.

- Mian, Atif R., Amir Sufi, and Emil Verner.** 2017. "Household Debt and Business Cycles Worldwide." *Quarterly Journal of Economics*, Forthcoming.
- Rigobon, Roberto.** 2003. "Identification through Heteroskedasticity." *The Review of Economics and Statistics*, 85(4): 777–792.
- Schularick, Moritz, and Alan M. Taylor.** 2012. "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008." *American Economic Review*, 102(2): 1029–61.
- Shaw, Edward Stone.** 1973. *Financial Deepning in Economic Development*. New York:Oxford Univ. Press.
- Sims, Christopher A., and Tao Zha.** 1996. "Bayesian methods for dynamic multivariate models." Federal Reserve Bank of Atlanta FRB Atlanta Working Paper 96-13.
- Sims, Christopher A., and Tao Zha.** 1999. "Error Bands for Impulse Responses." *Econometrica*, 67(5): 1113–1155.
- Sims, Christopher A., and Tao Zha.** 2006. "Were There Regime Switches in U.S. Monetary Policy?" *American Economic Review*, 96(1): 54–81.
- Smets, Frank, and Rafael Wouters.** 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review*, 97(3): 586–606.
- World Bank.** 2012. *Global Financial Development Report 2013: Rethinking the Role of the State in Finance*. Washington, DC:World Bank.

DEPARTMENT OF ECONOMICS, PRINCETON UNIVERSITY

E-mail address: markus@princeton.edu

DEPARTMENT OF FINANCE AND ECONOMICS, RUTGERS BUSINESS SCHOOL

E-mail address: dpalia@business.rutgers.edu

DEPARTMENT OF ECONOMICS, MIT

E-mail address: ksastroy@mit.edu

DEPARTMENT OF ECONOMICS, PRINCETON UNIVERSITY

E-mail address: sims@princeton.edu

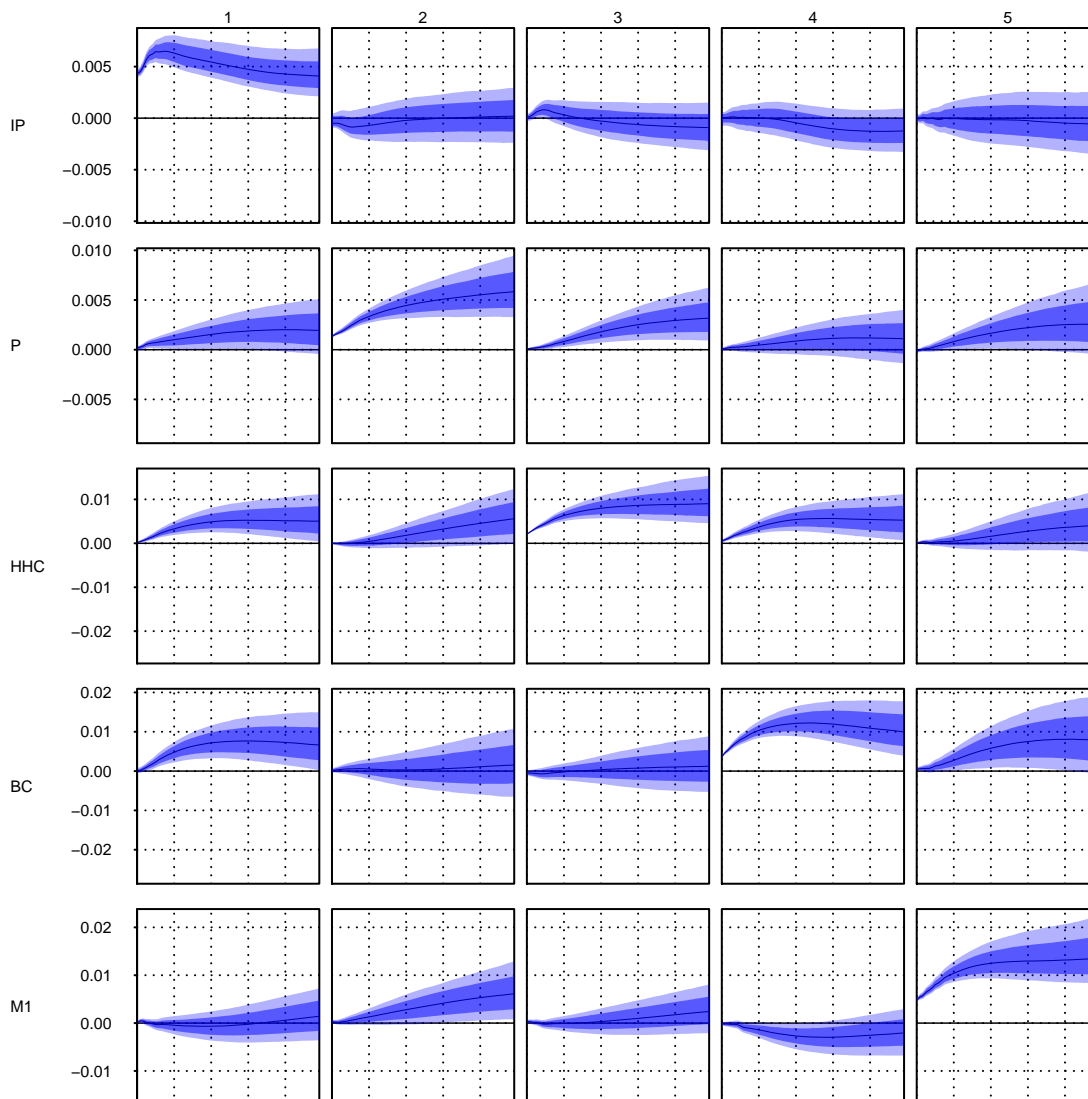


FIGURE 1. (1/4) Impulse responses to the ten orthogonal structural shocks in the model with t distributed errors over 60 months, with 68% (dark blue) and 90% (light blue) posterior uncertainty regions. Scaled to an “average” period with unit scale. Variables are in the order listed in Table 1.

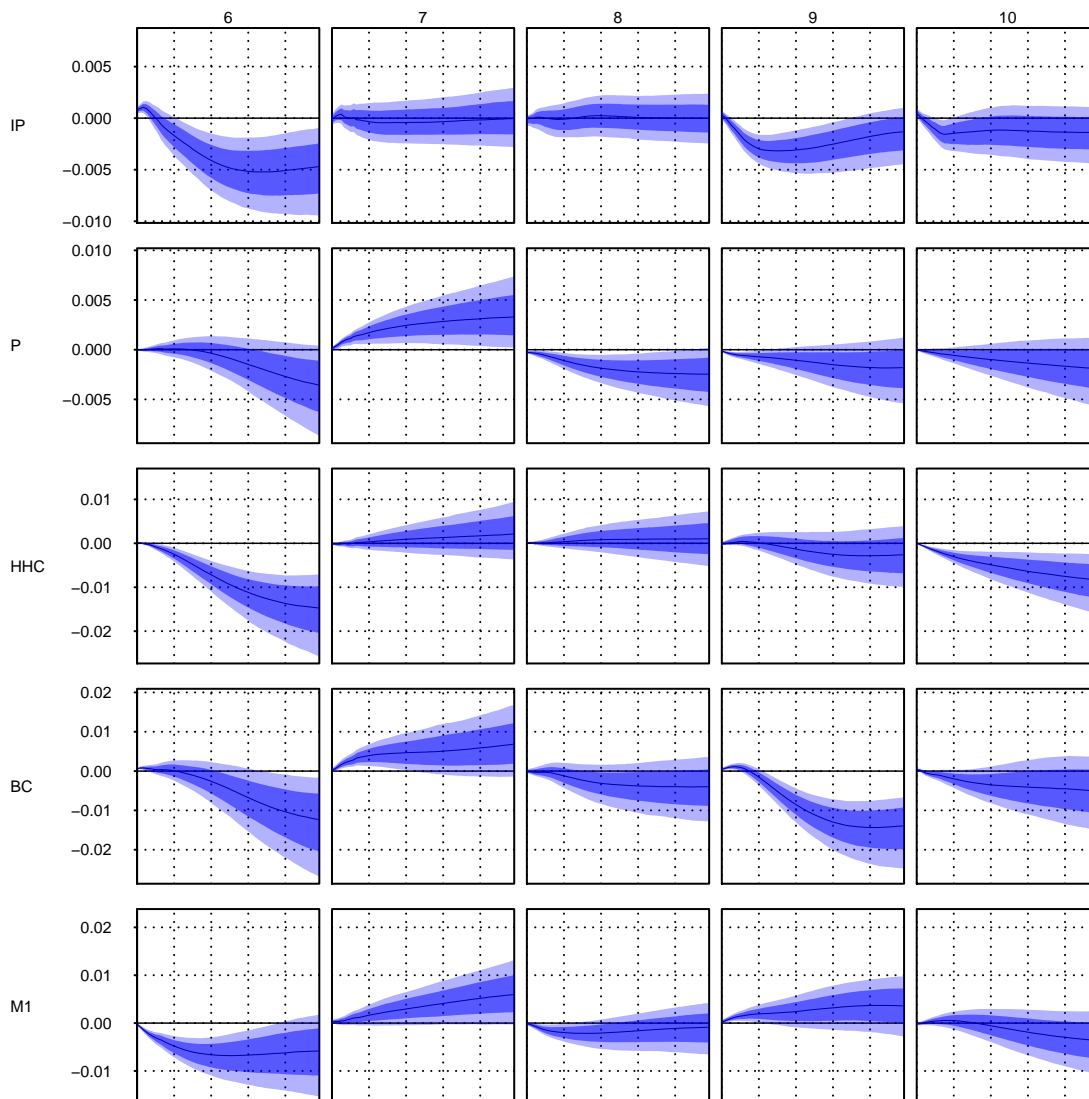


FIGURE 2. (2/4) Impulse responses to the ten orthogonal structural shocks in the model with t distributed errors over 60 months, with 68% (dark blue) and 90% (light blue) posterior uncertainty regions. Scaled to an “average” period with unit scale. Variables are in the order listed in Table 1.

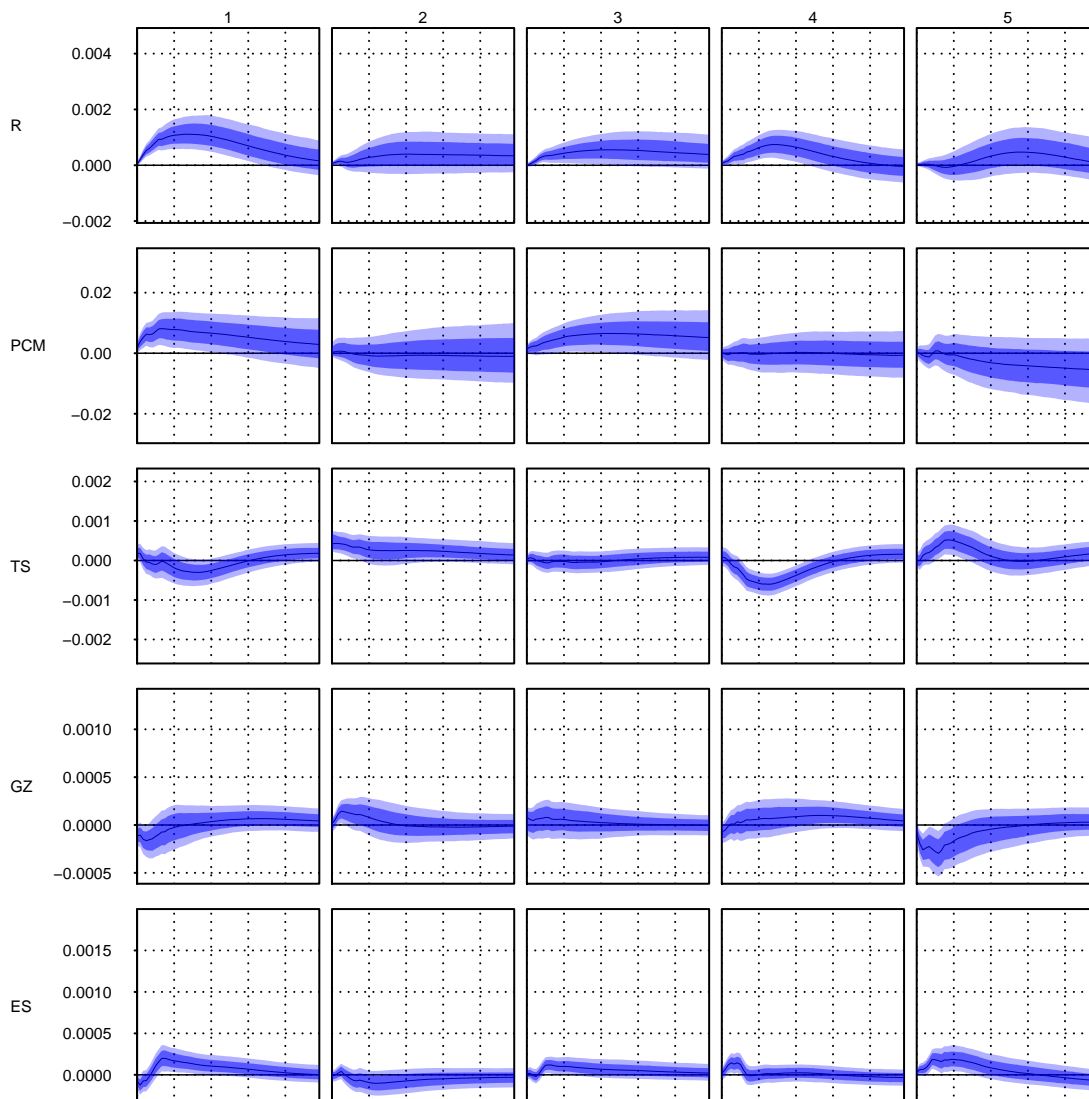


FIGURE 3. (3/4) Impulse responses to the ten orthogonal structural shocks in the model with t distributed errors over 60 months, with 68% (dark blue) and 90% (light blue) posterior uncertainty regions. Scaled to an “average” period with unit scale. Variables are in the order listed in Table 1.

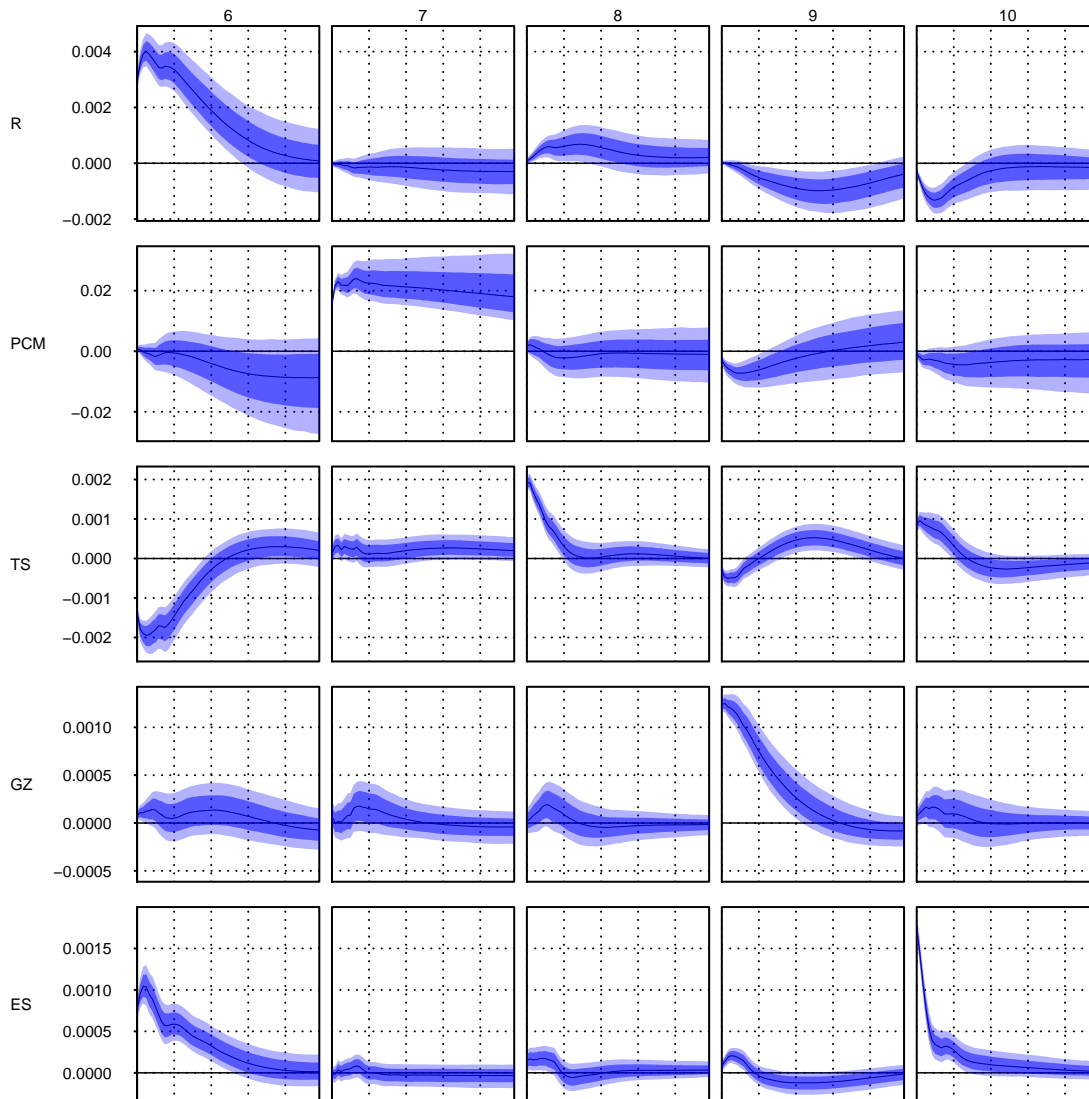


FIGURE 4. (4/4) Impulse responses to the ten orthogonal structural shocks in the model with t distributed errors over 60 months, with 68% (dark blue) and 90% (light blue) posterior uncertainty regions. Scaled to an “average” period with unit scale. Variables are in the order listed in Table 1.




FIGURE 5. The first two columns reproduce the responses to shocks 3 and 4 shown in Figures 1 and 3. The third and fourth columns combine each shock with a series of shocks to shock 6, the monetary policy shock, that keep R constant.

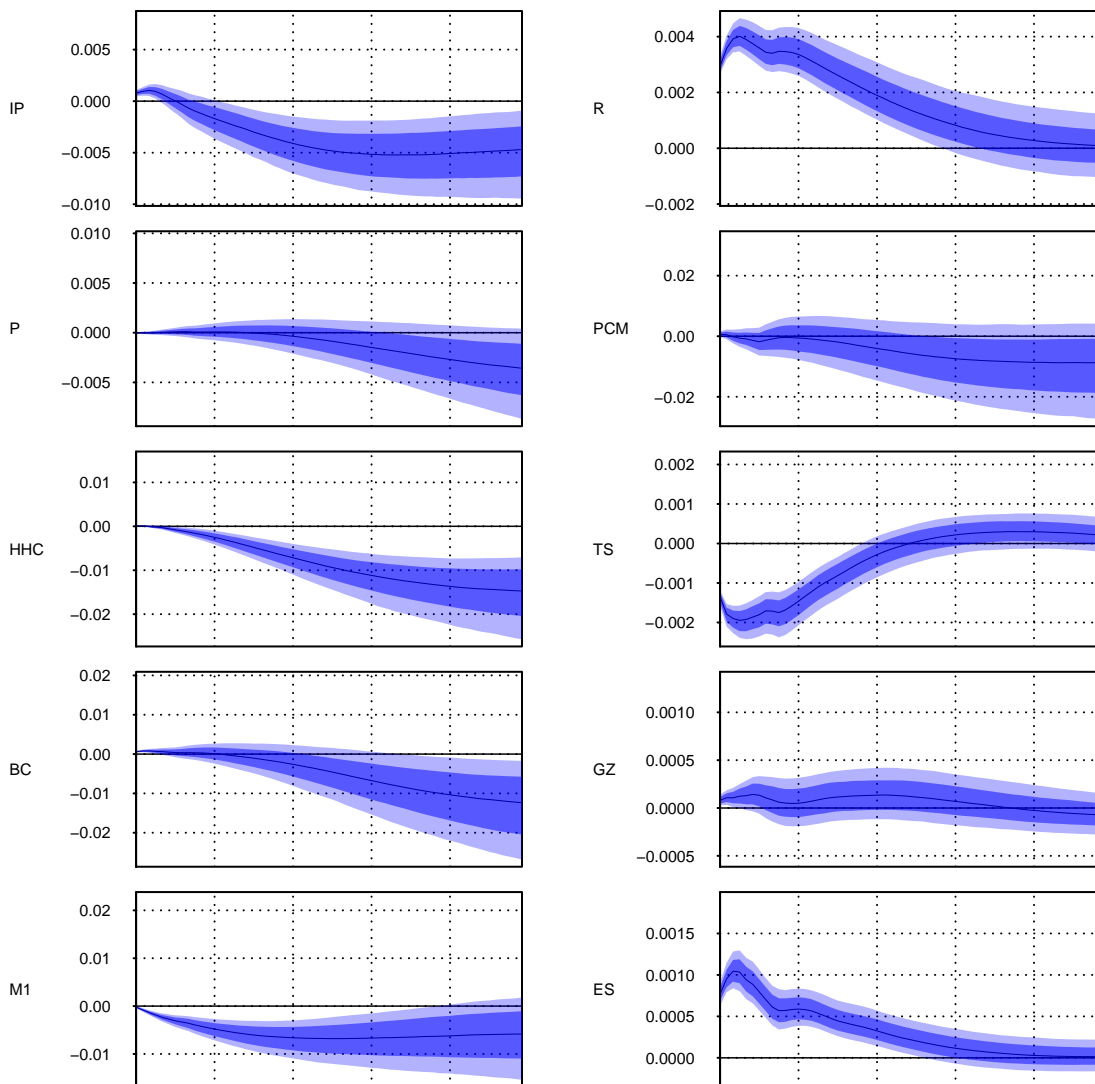


FIGURE 6. Impulse responses to the 6th ordered (monetary policy) shock, with 68% (dark blue) and 90% (light blue) posterior uncertainty regions. Scaled to an “average” period with unit scale. Variables are in the order listed in Table 1.

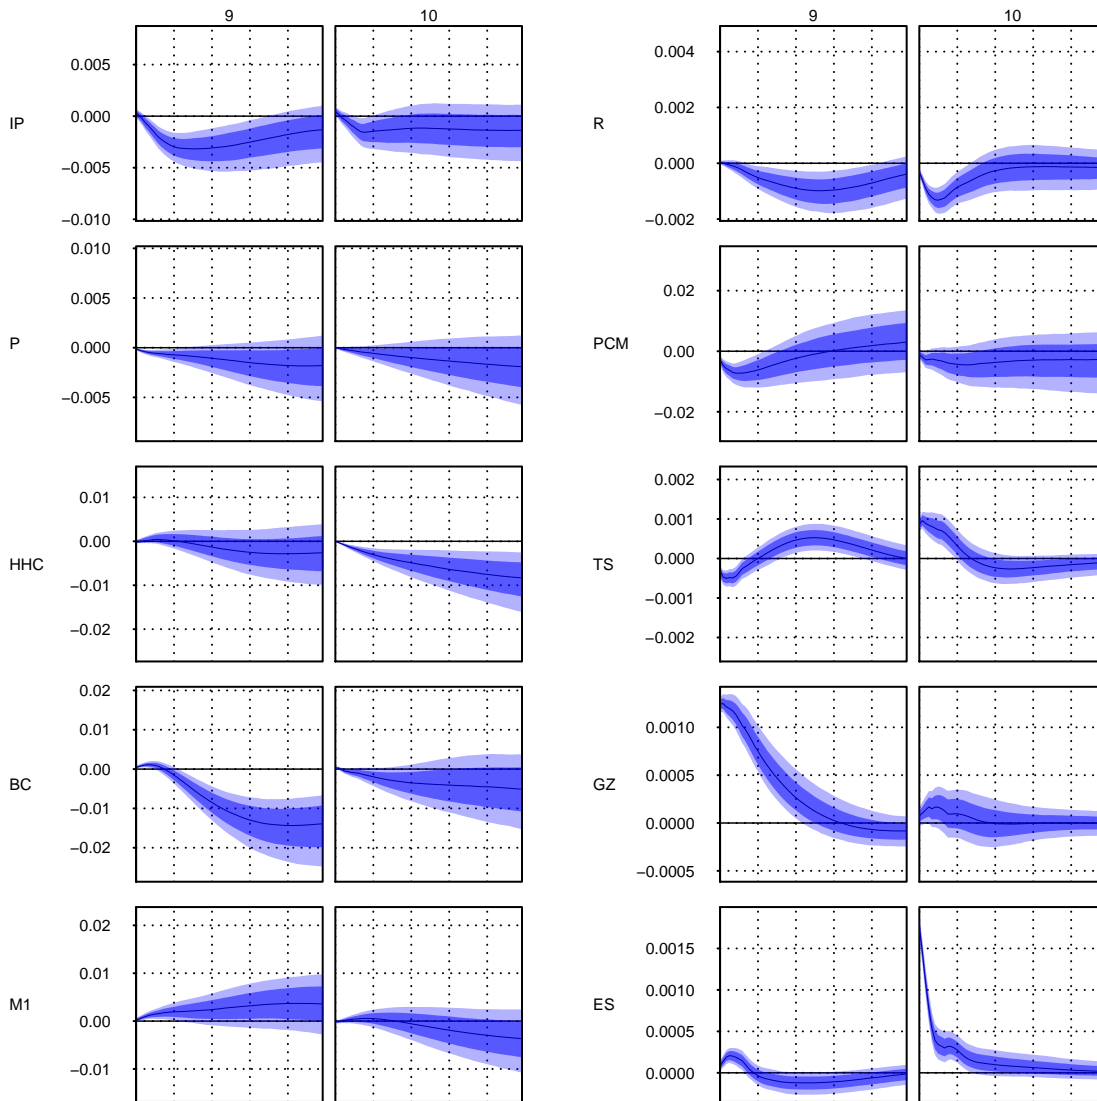


FIGURE 7. Impulse responses to the 9th (bond spread) and 10th (inter-bank spread) ordered shocks, with 68% (dark blue) and 90% (light blue) posterior uncertainty regions. Scaled to an “average” period with unit scale. Variables are in the order listed in Table 1.

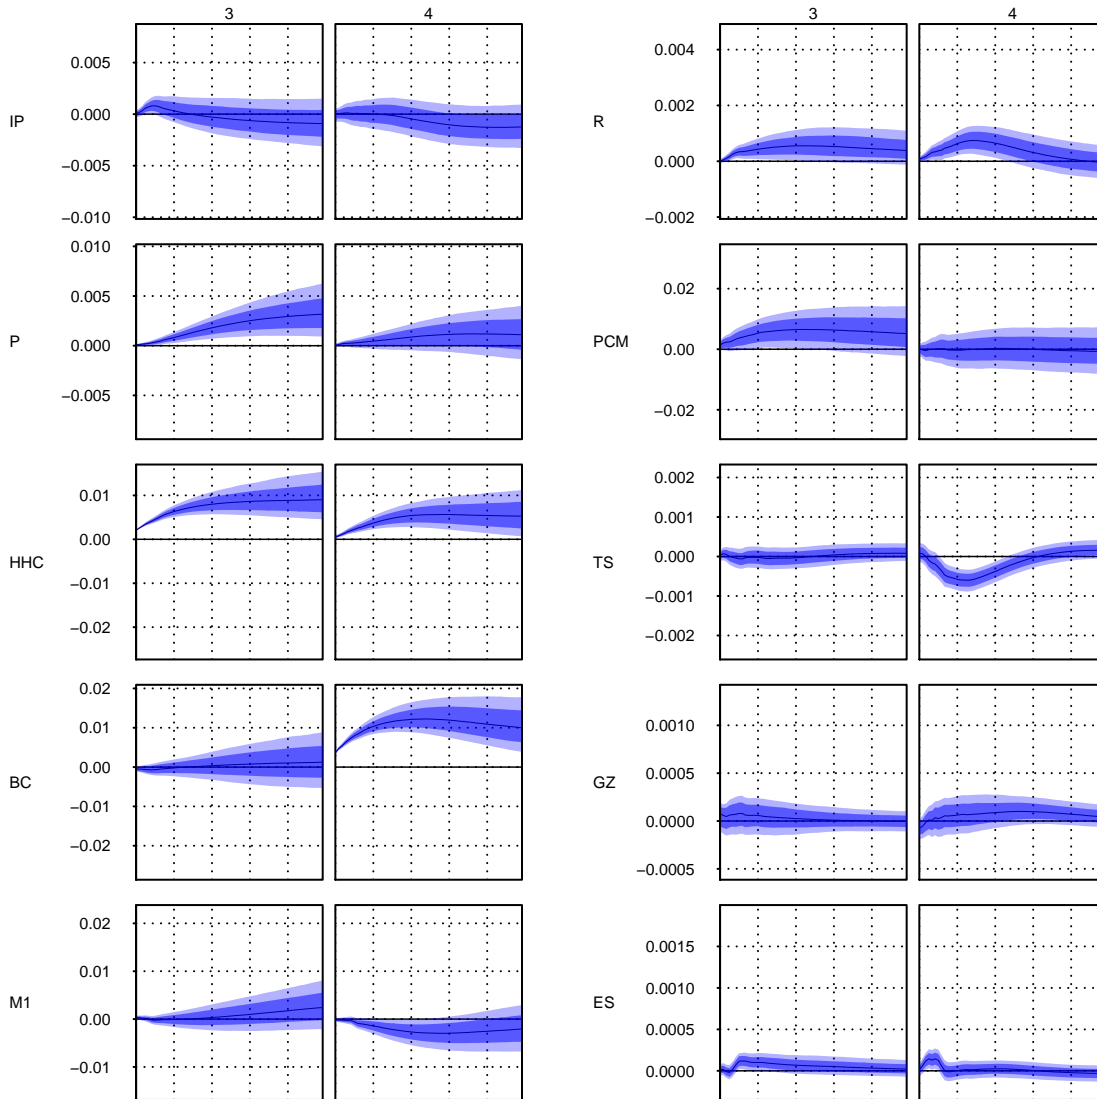


FIGURE 8. Impulse responses to the 3rd (household credit) and 4th (firm credit) ordered shocks, with 68% (dark blue) and 90% (light blue) posterior uncertainty regions. Scaled to an “average” period with unit scale. Variables are in the order listed in Table 1.

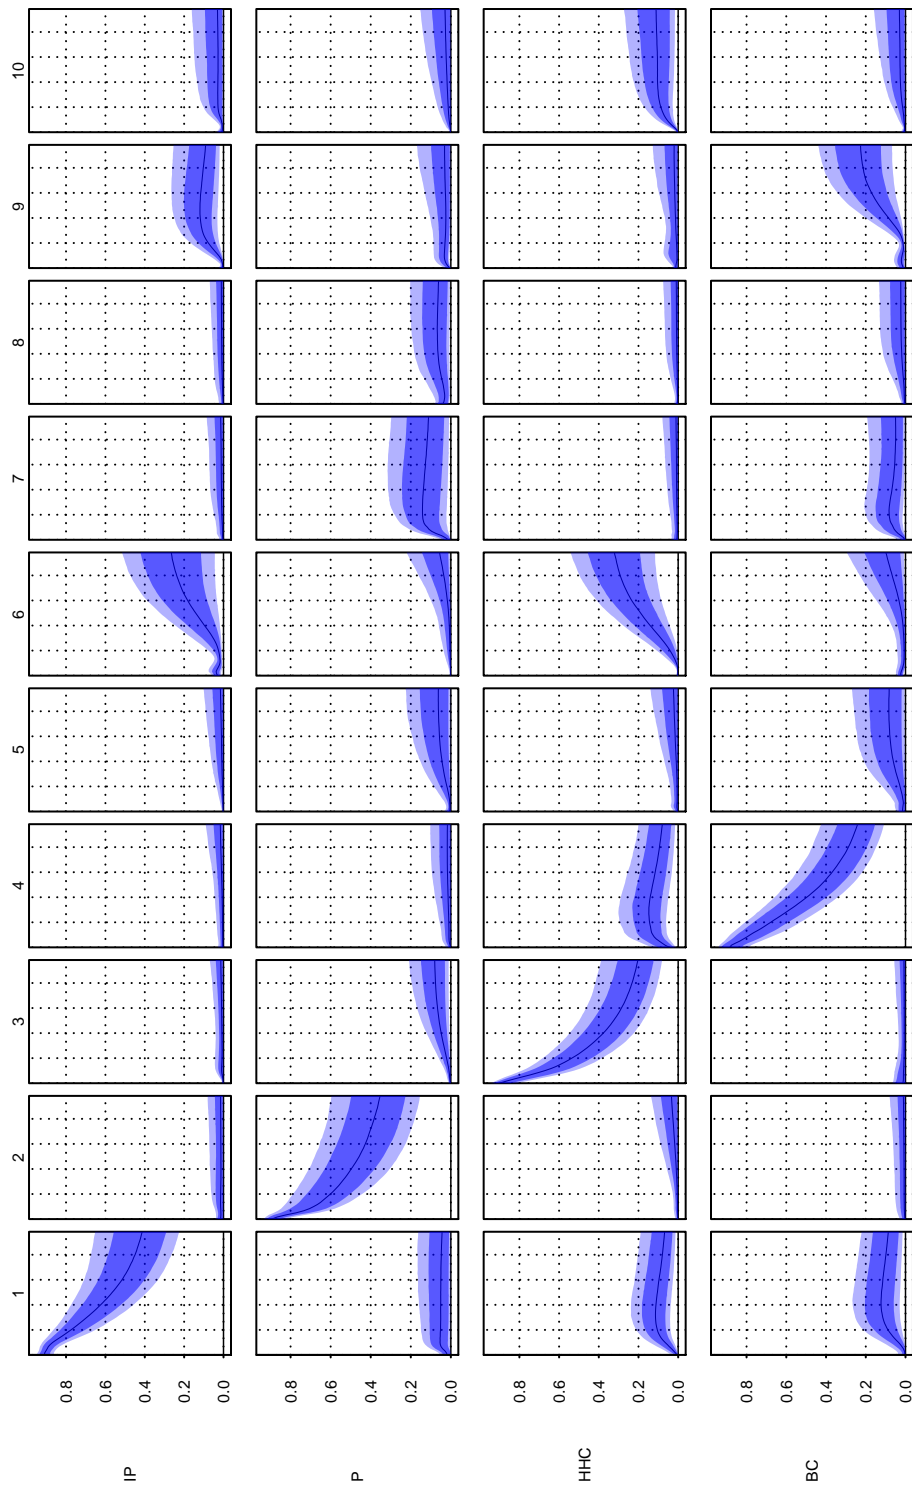


FIGURE 9. Forecast variance decompositions in the t -distributed errors model for IP, prices, household credit, and business credit over 60 months, with 68% (dark blue) and 90% (light blue) posterior uncertainty regions. Scaled to an “average” period with unit scale. Variables are in the order listed in Table 1.

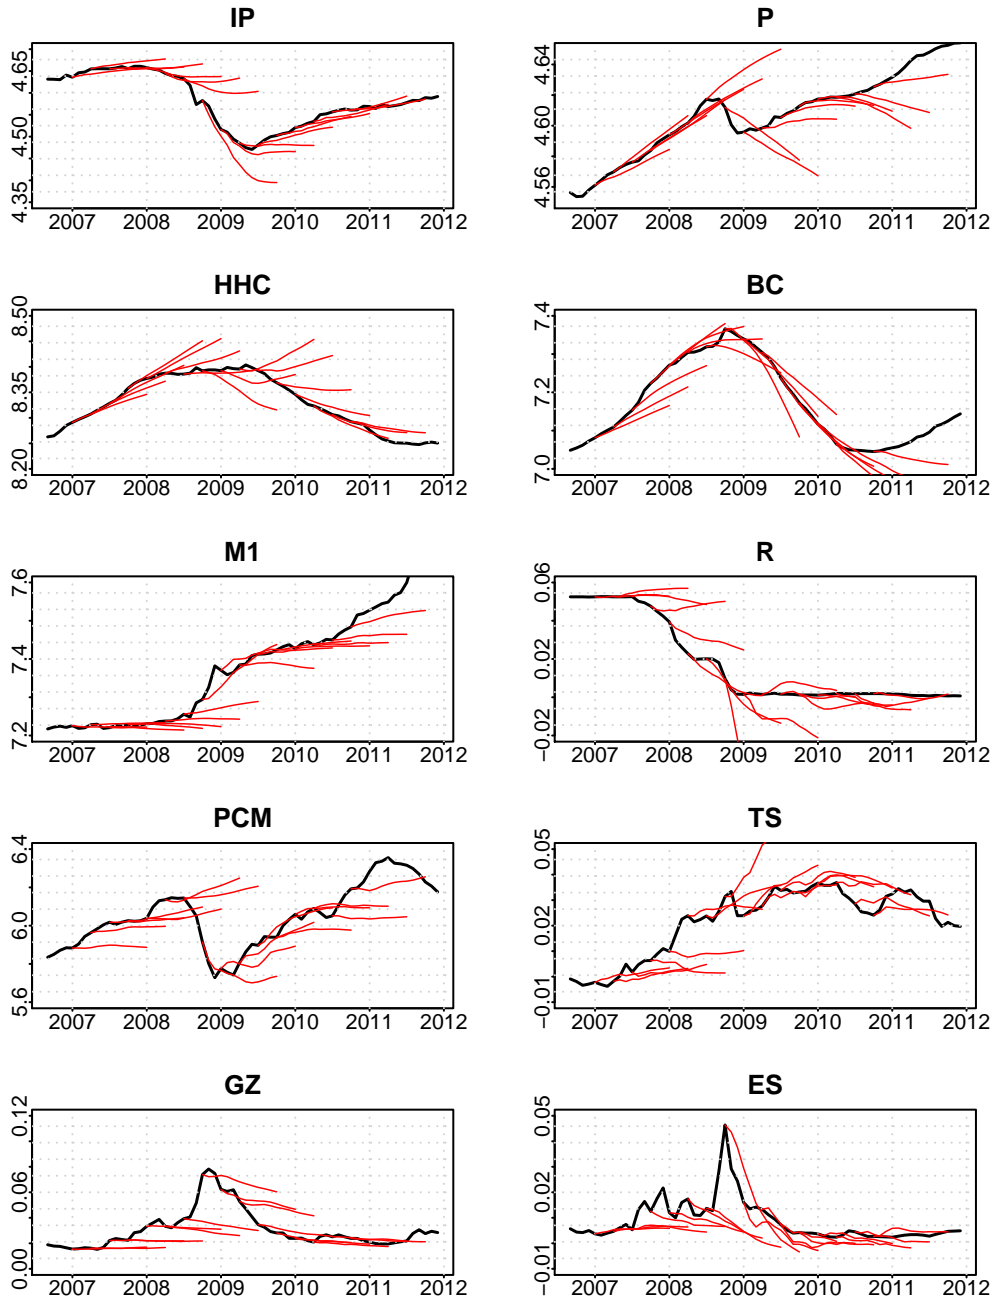


FIGURE 10. Posterior mode forecasts, from a model with all ten variables, estimated up to points in and around the Great Recession.




FIGURE 11. Posterior mode forecasts, from a model without the two credit aggregates, estimated up to points in and around the Great Recession.




FIGURE 12. Posterior mode forecasts, from a model without the three credit spreads, estimated up to points in and around the Great Recession.




FIGURE 13. Root mean squared error (RMSE) for 6-month forecasts from rolling estimations of the Gaussian errors model with all variables (blue), no credit spreads (red), and no credit aggregates (green). NBER recessions are shaded.

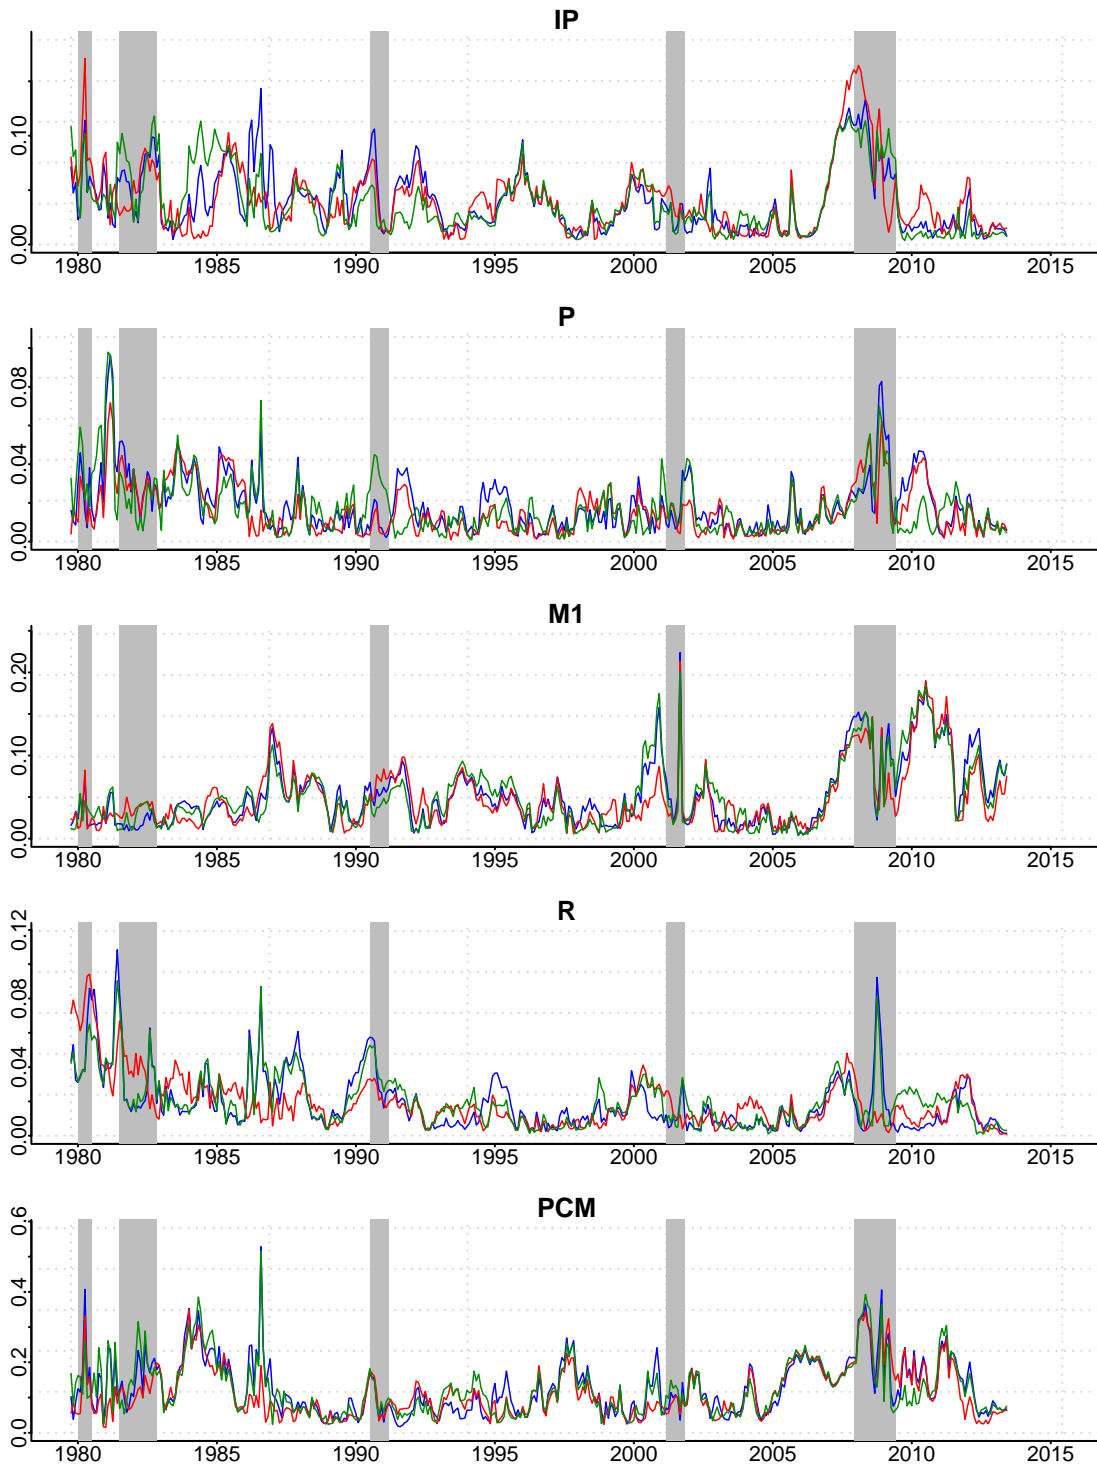


FIGURE 14. Root mean squared error (RMSE) for 24-month forecasts from rolling estimations of the main model with all variables (blue), no credit spreads (red), and no credit aggregates (green). NBER recessions are shaded.