Federal Reserve Bank of New York
Staff Reports

Multimodality in Macro-Financial Dynamics

Tobias Adrian
Nina Boyarchenko
Domenico Giannone

Staff Report No. 903
November 2019
Revised December 2020

Multimodality in Macro-Financial Dynamics
Tobias Adrian, Nina Boyarchenko, and Domenico Giannone
Federal Reserve Bank of New York Staff Reports, no. 903
November 2019; revised December 2020
JEL classification: C14, E17, E37, G01

Abstract

We estimate the evolution of the conditional joint distribution of economic and financial conditions. While the joint distribution is approximately Gaussian during normal periods, sharp tightenings of financial conditions lead to the emergence of additional modes. The U.S. economy has historically resolved quickly to the "good" mode, but we conjecture that poor policy choices could lead to prolonged periods of multimodality. We argue that multimodality arises naturally in a macro-financial intermediary model with occasionally binding intermediary constraints.

Key words: density impulse response, multimodality, nonparametric density estimator

---

Boyarchenko: Federal Reserve Bank of New York (email: nina.boyarchenko@ny.frb.org). Adrian: International Monetary Fund (email: tadrian@imf.org). Giannone: Amazon.com, Inc. (email: dgiannon2@gmail.com). The authors thank Patrick Adams and Eric Qiang for providing excellent research assistance. They also thank seminar and conference participants at the Federal Reserve Bank of New York, the Federal Reserve Board, the Bundesbank, the European Central Bank, the International Monetary Fund, Chicago Booth School of Business, Princeton University, Warwick Business School, and London Business School for helpful feedback.

1 Introduction

> "Tout le monde y croit cependant [que les erreurs sont Gaussien], me disait un jour M. Lippmann, car les expérimentateurs s'imaginent que c'est un théorème de mathématiques, et les mathématiciens que c'est un fait expérimental"
>
> (Henri Poincaré, *Calcul des probabilités*, 2nd ed., 1912, p. 171)

The theoretical literature on macro-financial dynamics has long postulated that the relationship between financial conditions and real activity is nonlinear: while financial conditions are relatively loose, the economy evolves as normal; when financial conditions tighten, the economy can slump into pronounced, lasting, macro-financial feedback loops with (extremely) adverse macroeconomic outcomes. The difficulty in testing and documenting empirically such non-linearities is that the specific nature of the non-linearity is model-specific, with no consensus in the literature on the "right" non-linear specification.

In this paper, we propose a flexible and robust non-parametric approach to characterizing non-linear system dynamics that combines kernel estimation of the one-step-ahead joint distribution (see e.g. Li and Racine, 2007, Chapter 6.2) with an efficient Monte Carlo procedure to generate term structures of joint distributions. Using non-parametric methods allows us to remain agnostic about the nature of dynamic interactions between variables of interest, allowing the data to inform us instead. If we further use independent kernels across variables, the kernel estimation can be easily imbedded into a Gibbs' sampler, allowing for estimation of joint dynamics of large number of variables of interest, as well as estimation with missing observations of variables of interest.

We apply this approach to estimate the joint dynamics of economic and financial conditions in the United States (U.S.). In our previous work (see Adrian, Boyarchenko, and Giannone, 2019), we documented that the left tail of the conditional distribution of real GDP growth moves together with the tightness of financial conditions. Examining the *joint* distribution of economic and financial conditions allows us to document a novel empirical fact. While the one-quarter-ahead joint distribution of economic and financial is (roughly) Gaussian and unimodal during "good" times, the shape of the distribution changes dramatically during periods of tight financial conditions, with additional modes emerging.

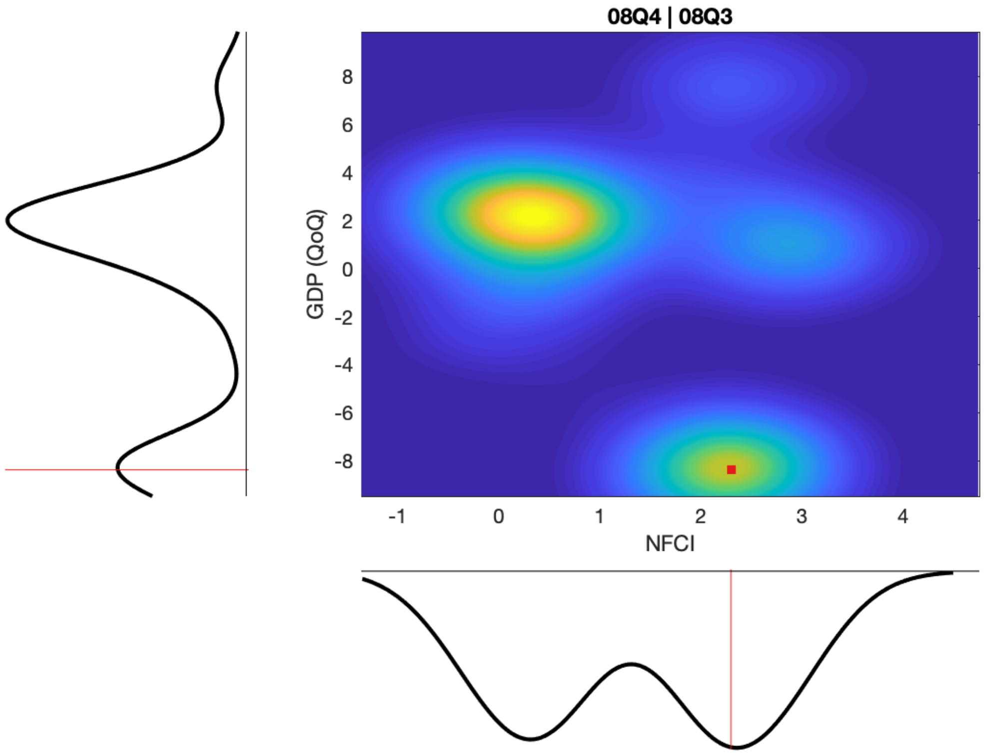

Figure 1. The figure shows the one quarter ahead marginal distributions of economic and financial conditions (real GDP growth and NFCI), together with the contour plot of the joint distribution, for 2008Q4 conditional on data as of 2008Q3. The red square indicates the ex post realization in 2008Q4.

We illustrate this finding in Figure 1, which plots the marginals and the joint one-quarter-ahead distribution of real GDP growth and financial conditions as of Q3 2008. The Figure clearly shows that, in the depth of the financial crisis, both marginal distributions exhibited multimodality: There was a possibility that the economy would resolve within a quarter to a "good" equilibrium with positive economic growth and easier financial conditions, as well as a possibility that the economy would resolve within a quarter to a "bad" equilibrium with negative economic growth and tight financial conditions.

We show that, more generally, deviations from Gaussianity only emerge during times of unusually tight financial conditions. When the tightness in financial conditions is relatively mild, such as during the 1987 stock market crash, multimodality only appears in short horizon forecasts and resolves within a couple of quarters. When the tightness in financial conditions is more widespread, such as during the 2007-2009 financial crisis, multimodality is evident in longer horizon forecasts and can take up to a year to resolve. Although in the long-run the economy always reverts back to a "good" equilibrium, with positive GDP growth and looser financial conditions, prolonged periods of multimodality or even selection of the "bad" equilibrium can lead to large losses in the level of GDP. We speculate that appropriate policy accommodation is crucial in reverting quicker to the good equilibrium but leave a formal exploration of this hypothesis for future work.

Our methodology allows us to estimate the term structure of joint distributions which, in turn, allow us to estimate density impulse response functions (DIRs) in the spirit of Gallant, Rossi, and Tauchen (1993). DIRs compute the change in the full distribution over horizons of interest, conditional on a perturbation in the one-step-ahead distribution. Analogously to impulse response functions in the linear vector autoregression setting, DIRs provide (time-varying) target moments against which models should be evaluated.

We show that our estimator performs well out-of-sample. In particular, the estimated in-sample distributions for real GDP growth and financial conditions coincide with those estimated out-of-sample. Comparing the overall out-of-sample performance of our non-parametric estimator to two linear vector auto-regressive (VARs) models – one with Gaussian standard errors and one with non-parametric standard errors – our estimator out-performs both alternatives in terms of standard performance metrics, such as probability integral transforms (PITs) and log-scores. Focusing on the forecasts during the 2007–2009 financial crisis, our out-of-sample forecast also exhibits multimodality, though the "bad" equilibrium is not as severe as that predicted in-sample. In contrast, even the VAR with non-parametric standard errors predicts a unimodal distribution, suggesting the multimodality arises from non-linear propagation of shocks, rather than non-Gaussian innovations.

The remainder of the paper is organized as follows. Section 2 discusses the theoretical foundations of multimodality. Section 3 presents our algorithm to compute the term structures of joint distributions, and discusses our contribution from an econometric perspective. Section 4 shows the empirical results for the joint distribution of economic and financial conditions in the U.S. Section 5 presents the density impulse responses. Section 6 undertakes out-of-sample forecast evaluation. Section 7 concludes.

2 The theoretical case for multimodality

Theoretical literature on economic dynamics has long postulated the possibility of multiple equilibria. In a seminal contribution, Diamond (1982) points out that thick market externalities in the labor market can generate a market failure (the so-called "coconut model"). Bryant (1983) develops a model where technological complementarities can similarly give rise to a persistent market failure via bad equilibrium selection. In the fundamental contribution of Diamond and Dybvig (1983), the possibility of bank runs give rise to multiple equilibria. And in Murphy, Shleifer, and Vishny (1989), demand spillovers generate multiple equilibria.

More recently, the macro-finance literature has focused on the wholesale banking sector as a source of equilibrium multiplicity. Gertler and Kiyotaki (2015) and Gertler, Kiyotaki, and Prestipino (2016) study models that account for the buildup and collapse of wholesale banking, and sketch out the transmission of the crises to the real sector by characterizing the sudden and discrete nature of the banking panics as well as the circumstances that makes an economy vulnerable to such panics. Sunspot runs can arise that are harmful to the economy. Whether a run equilibrium exists depends on fundamentals. The probability of a sunspot run is the outcome of a rational forecast based on fundamentals. The Gertler et al. (2016) model thus captures the movement from slow to fast runs that was a feature of the Great Recession: A weakening of banks' balance sheets increases the probability of a run, leading depositors to withdraw funds from banks.

Models with multiple equilibria naturally translate to multimodal predictive densities, with the "good" equilibrium assigned most of the probability distribution most of the time, but the "bad" equilibrium also receiving positive probability mass during times of crisis. Multimodality on its own, however, is not evidence of multiple equilibria. Instead, multimodality in the conditional distribution can arise due to sharp amplification mechanisms. Indeed, in two fundamental contributions, Morris and Shin (2000, 2002) propose unique equilibrium via imperfect knowledge while preserving amplification in models of strategic complementarity, thereby allowing for amplification mechanisms without multiplicity. In the aftermath of the financial crisis, the macro-finance literature has used this insight to model nonlinear amplification mechanisms through financial constraints, generating strong macro-financial linkages and extreme negative skewness in real outcomes in "bad" times but linear Gaussian dynamics during normal periods. Examples of such models include He and Krishnamurthy

(2013), Brunnermeier and Sannikov (2014), Adrian and Boyarchenko (2012), and Adrian and Duarte (2016). In those, and many related studies, all shocks are conditionally Gaussian, but volatility and drift depend on state variables in highly nonlinear ways due to occasionally binding constraints.

Finally, most closely related to our results, Fernández-Villaverde, Hurtado, and Nuño (2019) study a heterogeneous households model with a representative intermediary where idiosyncratic risks to household income lead to a multiplicity of stochastic steady states. In the high intermediary leverage stochastic steady state, erosion of intermediary wealth leads to prolonged periods of low wages and low capital, while shocks to intermediary wealth in the low leverage steady state resolve quickly.

3 Methodology

We are interested in constructing multi-period-ahead conditional joint distributions of economic and financial conditions. In this section, we describe how to construct the one-period-ahead conditional joint distribution. We then use efficient Markov Chain Monte Carlo (MCMC) to construct multi-period-ahead distributions from the one-period-ahead distribution.

3.1 One-period-ahead distribution

Consider a time series dataset of n_y endogenous variables $y_{i,t}$, $i = 1, \ldots, n_y$ and denote by $y_t = \left(y_{1,t}, \ldots, y_{n_y,t}\right)'$ the vector of date t realizations of the n variables. In addition to the endogenous variables y_t, suppose that we have n_x exogenous predictors x_t. In this paper, we are interested in the case when the exogenous predictors are p lags of y, so that $n_x = p \times n_y$ and

$$x_t = \left(y'_{t-1}, \ldots, y'_{t-p}\right)'.$$

That is, we are interested in estimating a distributional equivalent to vector autoregressions.

We are now ready to write down our kernel estimator (see Li and Racine, 2007, chapter 6.2). Let T be the number of observations of y_t that we have available. Then the joint distribution of y conditional on x can be estimated as

$$\hat{p}\left(y|\,x\right) = \frac{\frac{1}{T-p}\sum_{t=p+1}^{T}\mathcal{K}_{\omega_y}^{y}\left(y-y_t\right)\mathcal{K}_{\omega_x}^{x}\left(x-x_t\right)}{\frac{1}{T-p}\sum_{t=p+1}^{T}\mathcal{K}_{\omega_x}^{x}\left(x-x_t\right)}, \tag{1}$$

where $\mathcal{K}_{\omega_y}^{y}$ and $\mathcal{K}_{\omega_x}^{x}$ are independent kernels for y and x, given by

$$\mathcal{K}_{\omega_y}^{y}\left(y-y_t\right) = \prod_{i=1}^{n_y}\frac{1}{\omega_{y_i}}\varphi\left(\frac{y_i-y_{i,t}}{\omega_{y_i}}\right) \equiv \prod_{i=1}^{n_y}\mathcal{K}_{\omega_{y_i}}^{y_i}\left(y_i-y_{i,t}\right) \tag{2}$$

$$\mathcal{K}_{\omega_x}^{x}\left(x-x_t\right) = \prod_{i=1}^{n_x}\frac{1}{\omega_{x_i}}\varphi\left(\frac{x_i-x_{i,t}}{\omega_{x_i}}\right) \equiv \prod_{i=1}^{n_x}\mathcal{K}_{\omega_{x_i}}^{x_i}\left(x_i-x_{i,t}\right). \tag{3}$$

For our baseline results, we use multivariate normal kernels, so that $\varphi\left(\cdot\right)$ is the normal probability distribution function, but the kernel estimation can easily be used with alternative kernels (such as multivariate Student kernels) or be modified to accommodate dependent kernels for the endogenous and exogenous variables.

We parameterize the bandwidths as being proportional to the in-sample unconditional standard deviation of the corresponding variable: $\omega_{y_i} = c_{y_i}\sigma_{y_i}$, $\omega_{x_i} = c_{x_i}\sigma_{x_i}$, and choose a single proportionally constant $c = c_{y_i} = c_{x_j}$ for all $i = 1,\ldots,n_y$, $j = 1,\ldots,n_x$ to maximize the predictive accuracy of the resultant one-period-ahead conditional joint distribution. We plot the relationship between the bandwidth proportionality constant c and the out-of-sample (log) predictive score for the one-quarter-ahead and one-year-ahead distribution in Figure 2. The figure shows that the bandwidth selection is in-line with values that would be predicted by asymptotic criteria,[1] that there is an interior solution for the optimal constant of proportionality, and that the performance is maximized for $c \approx 0.5$.

3.2 Efficient Monte Carlo

Given an estimated one-period-ahead distribution $\hat{p}\left(y|\,x\right)$, we can use Monte Carlo simulations to estimate h-period-ahead distributions by sequentially drawing paths of y. In principle, these draws can be made directly from the inverse CDF implied by $\hat{p}\left(y|\,x\right)$ by drawing u from a (multinomial) uniform distribution and finding y that solves $y = \hat{P}^{-1}\left(u|\,x\right)$. We

[1]Li and Racine (2007) apply a version of the Silverman (1986) "rule-of-thumb" for joint unconditional density estimation where $w_j = 1.06\sigma_j T^{-1/(4+M+N)}$ for variable j, bandwidth w_j, standard deviation σ_j, sample size T, number of independent variables M, and number of dependent variables N. The rule is derived from minimizing asymptotic mean integrated square error for a Gaussian reference distribution.

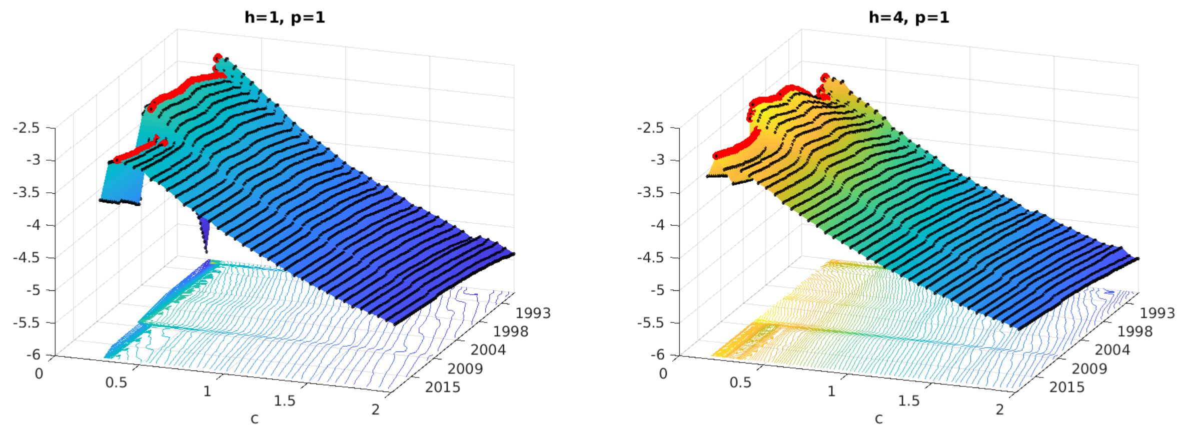

(a) One quarter ahead **(b)** One year ahead

Figure 2. Optimal bandwidth selection. This figure plots the out-of-sample log predictive scores for the one-period-ahead conditional joint distribution of real activity and financial conditions, as a function of the bandwidth proportionality constant c. Predictor variables: one lag of GDP growth and financial conditions; time period: Q1 1973 – Q1 2019.

increase the efficiency of this procedure by discretizing the state space as follows.

Algorithm 1. ***Simulating paths of*** y.

To estimate the H-period-ahead distribution of y, generate n_{sim} paths of y as follows.

1. *Discretize the state-space. Set $\kappa = \kappa_0$. For each variable j, loop through:*
   (a) *Initialize grid with bound $[\min(y_j) - \kappa\sigma_{y,j},\ \max(y_j) + \kappa\sigma_{y,j}]$ and grid point increments of $\Delta \times \sigma_{y,j}$.*
   (b) *For each grid point $y_{j,i}$, compute the kernel CDF $\Phi\left(\frac{y_j - y_{j,i}}{\omega_{y_j}}\right)$.*
   (c) *Verify that the kernel PDF integrates to one. Verify that the kernel CDF has a maximum of 1 within a tolerance of ε. If not, reset $\kappa = (1+\delta) \times \kappa$ and repeat.*
2. *For each simulated path $k = 1, \ldots, n_{sim}$, loop through each horizon $h = 1, \ldots, H$ by drawing $y^k_{t+h}|y^k_{t+h-1}$ from the grid established in Step 1 and verify the normalization condition.*

The choices κ_0, ε, Δ, and δ control the speed and accuracy of the estimator, with κ_0 controlling the size of the initial state-space, Δ controlling the fineness of the discretization grid, ε the tolerance for integration error, and δ the speed with which the size of the state space grows while the estimated density is "missing" probability mass. We set $\kappa_0 = 0.1$, $\varepsilon = 10^{-20}$, $\Delta = 1/20$, and $\delta = 0.05$.

3.3 Alternative estimators and other considerations

Multi-period conditional distributional forecasts provide a natural non-parametric counterpart to traditional VARs. Though "standard" VARs are linear, they have proven extraordinarily useful in forecasting and scenario design, helping establish many stylized facts to guide and validate economic modeling. More recent literature has extended this traditional VAR literature to nonlinear but parametric settings, as reviewed extensively in Kilian and Lutkepohl (2018). Terasvirta and Anderson (1992), van Dijk, Terasvirta, and Franses (2002), and Kilian and Taylor (2003) propose smooth transition models. Altissimo and Violante (2001) propose thresholded VARs, where recent applications include Ascari and Haber (2019), Auerbach and Gorodnichenko (2013) , Aikman, Lehnert, Liang, and Modugno (2016), and Caggiano, Castelnuovo, and Figueres (2019). Hamilton (1989), Sims and Zha (2006), Chang, Choi, and Park (2017), and Hubrich and Tetlow (2015) offer Markov switching VARs. Aruoba, Bocola, and Schorfheide (2017) use an approach of quadratic autoregressions with pruning. Relative to the approach proposed in our paper, this literature has the potential draw back of relying on particular parametrizations of the non-linearity, which may not correspond to any particular theoretical model of non-linearity.

The model that is closest to our methodology is the varying parameters VARs with stochastic volatility developed by Cogley and Sargent (2005), Primiceri (2005), Cogley and Sbordone (2008), and Del Negro and Primiceri (2015). Indeed, any non-linear model can be approximated by a time-varying parameter linear mode, (see Granger, 2008). D'Agostino, Gambetti, and Giannone (2013) show that these models generate accurate predictions in real time. The limitation of these approaches is the difficulty to assess and inspect the mechanism that generates the non-linear dynamics as the time variation of the system is exogenous and not linked to the state of the economy. Models with more interpretable mechanisms include Carriero, Clark, and Marcellino (2018) and Caldara, Scotti, and Zhong (2019) who show that VARs with stochastic volatility can generate downside growth risk similar to (Adrian et al., 2019) when shocks are allowed to affect both the conditional mean and the conditional variance.

The objective of this paper is to provide a nonparametric alternative to nonlinear system dynamics. Our approach is very general and flexible, and is focused on the in-sample fit to provide guidance to economic theory. However, we do also provide some out-of-sample evaluation of our approach in Section 6 to document that our non-parametric approach produces reasonable in real time predictions. In fact, in comparison to a linear VAR with non-Gaussian errors our forecast performance fares well. We leave more comprehensive out-of-sample evaluations for future work.

The first step in our methodology – the estimation of the one-period-ahead conditional joint distribution – can be implemented using alternative estimation methods. For example, Gallant et al. (1993) use splines. Norets and Pati (2017) and Sims (2000) deploy Bayesian non-parametric methods. Koenker, Leorato, and Peracchi (2013), and Adrian et al. (2019) use quantile or distributional regressions. We focus on kernel density estimation as it provides a relatively straightforward and easily implementable method for characterizing the distribution without imposing a parametric shape. An additional advantage is that the kernel density estimator can be easily modified to accommodate a Gibbs' sampling; that is, we can easily use a combination of the kernel density estimator and Gibbs' sampling to estimate joint distributions of large numbers of variables.

More specifically, construct the distribution of y_i conditional on x and all the other endogenous variables y_{-i} as

$$\hat{p}_i\left(y_i \middle| x, y_{-i}\right) = \frac{\frac{1}{T-p}\sum_{t=p+1}^{T}\mathcal{K}_{\omega_y}^{y}\left(y-y_t\right)\mathcal{K}_{\omega_x}^{x}\left(x-x_t\right)}{\frac{1}{T-p}\sum_{t=p+1}^{T}\mathcal{K}_{\omega_x}^{x}\left(x-x_t\right)\prod_{k\neq i}\mathcal{K}_{\omega_{y_k}}^{y_k}\left(y_k-y_{k,t}\right)}. \tag{4}$$

By construction, each $\hat{p}_i\left(y_i \middle| x, y_{-i}\right)$ corresponds to the same joint distribution of $y|x$: the numerator in (4) is the same for all $i = 1, \ldots, n_y$, and is identical to the numerator in (1). Thus, a Gibbs' sampler that draws realizations of y_t by cycling through the conditional draws from (4) converges.[2]

To make multiperiod draws with the Gibbs sampler, we modify the procedure in Algorithm 1 as follows.

Algorithm 2. ***Simulating paths of*** y ***using the Gibbs sampler.*** *To estimate the H-period-ahead distribution of y, generate n_{sim} paths of y as follows.*

1. *Discretize the state-space. Set $\kappa = \kappa_0$. For each variable j, loop through:*
   (a) *Initialize grid with bound $\left[\min\left(y_j\right)-\kappa\sigma_{y,j},\ \max\left(y_j\right)+\kappa\sigma_{y,j}\right]$ and grid point increments of $\sigma_{y,j}\times\Delta$.*

[2]See e.g. Arnold, Castillo, and Sarabia (1999), Ch. 1.

(b) For each grid point $y_{j,i}$, compute the kernel CDF $\Phi\left(\frac{y_j - y_{j,i}}{\omega_{y_j}}\right)$.

(c) Verify that the kernel PDF integrates to one. Verify that the kernel CDF has a maximum of 1 within a tolerance of ε. If not, set $\kappa = (1+\delta) \times \kappa$ and repeat.

2. Initialize $y^1_{j,t+1}$ at the conditional mean of variable j.

3. For each simulated path $k = 1, \ldots, n_{sim}$, loop through each horizon $h = 1, \ldots, H$:

(a) Draw $y^k_{1,t+h} | y^{k-1}_{2,t+h}, \ldots, y^{k-1}_{N,t+h}, y^k_{t+h-1}$ from the grid established in Step 1 and verify the normalization condition.

(b) Draw $y^k_{2,t+h} | y^k_{1,t+h}, y^{k-1}_{3,t+h}, \ldots, y^{k-1}_{N,t+h}, y^k_{t+h-1}$ from the grid established in Step 1 and verify the normalization condition.

$\vdots$

(c) Draw $y^k_{N,t+h} | y^k_{1,t+h}, \ldots, y^k_{N-1,t+h}, y^k_{t+h-1}$ from the grid established in Step 1 and verify the normalization condition.

4. Discard first DISCARD *draws.*

Finally, the kernel density approach combined with a Gibbs sampler can easily accommodate missing observations of the endogenous variables: given a candidate distribution of $y_i | y_{-i}, x$, we can draw the missing observations of y_i, and then re-estimate the conditional distribution (4).

3.4 Data

To gauge financial conditions, we use the National Financial Conditions Index (NFCI).[3] The NFCI provides a weekly estimate of U.S. financial conditions in money markets, debt and equity markets, and the traditional and shadow banking systems. The index is a weighted average of 105 measures of financial activity, each expressed relative to their sample averages and scaled by their sample standard deviations.[4] The methodology for the NFCI is described in Brave and Butters (2012) and is based on the quasi maximum likelihood estimators for large dynamic factor models developed by Doz, Giannone, and Reichlin (2012). The data for the NFCI starts in January 1973, which we use as starting point for our empirical

[3]The NFCI is computed by the Federal Reserve Bank of Chicago, available here.

[4]The list of indicators is provided here.

investigation. We average the weekly NFCI data within the quarter to obtain a quarterly NFCI series.

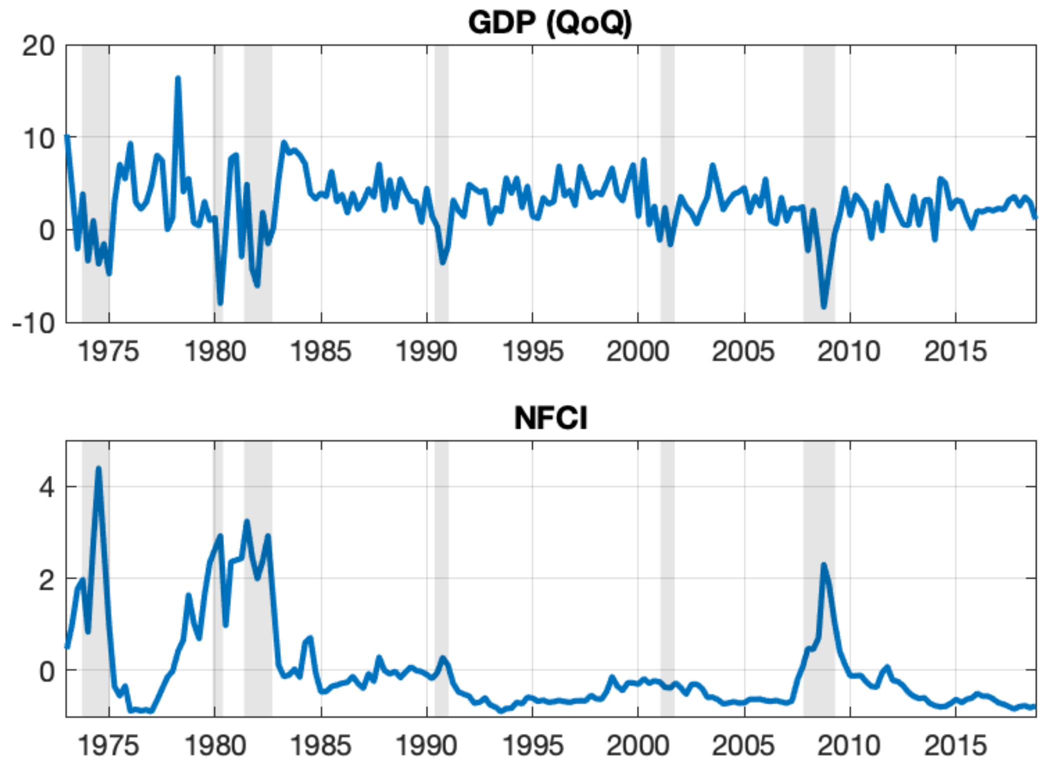

Figure 3. Economic and financial conditions. The figure shows the time series of QoQ real GDP growth and the financial conditions index, NFCI, together with NBER recession shadings.

Figure 3 shows the time series of QoQ real GDP growth, and the financial conditions index, NFCI, together with recession shadings. Real GDP growth is lower and financial conditions tighten during recessions, so that economic and financial conditions are (negatively) correlated. Adrian et al. (2019) document that, although the correlation is stronger during recessions, economic and financial conditions are correlated during normal times as well.

4 Joint distribution of economic and financial activity

We begin by examining the joint distribution of economic and financial conditions. We show that tight financial conditions lead to the emergence of multimodality in the joint distribution: the economy can either continue in a good GDP growth, loose financial conditions state or transition to an adverse GDP growth, tight financial conditions state.

4.1 Time series evolution of the joint distribution

We estimate the joint distribution of economic and financial conditions following the procedure in Section 3, using one lag of real GDP growth and NFCI as the conditioning variables. Figure 4 plots the estimated joint distribution at one, two, three, and four quarters ahead over time. Consider first the one-quarter-ahead distribution, plotted in Figure 4a. More specifically, for each quarter in the sample, Figure 4a presents the contour plot of the one-quarter-ahead joint distribution of real GDP growth and financial conditions. Contour plots shaped like symmetric disks correspond to (joint) Gaussian distributions: the conditional distribution is equally uncertain about both improvements and deteriorations to both economic and financial conditions. Figure 4a shows that this is the "normal" mode of the one-quarter-ahead joint distribution of economic and financial conditions. Though the mode of the distribution moves around over time, during "good" times, the joint distribution looks close to being symmetric around that mode.

Instead, if we focus on periods of tight financial conditions, the shape of the joint distribution changes radically and additional modes of the distribution emerge. At the one quarter horizon, multimodality emerges in the early 1980s during the Volcker disinflation episode, in 1987 around the stock market crash, in 1998 around the LTCM crisis, and starting in late 2007 for the great financial crisis. That is, when current financial conditions are tighter than normal, the estimated conditional distribution elongates and additional modes emerge. In the "good" mode, the economy returns to a good GDP growth, loose financial conditions state. In the "bad" mode(s), the economy transitions to a bad GDP growth, tight financial conditions state. When tight financial conditions are widespread in the economy, such as during the Volcker disinflation episode and during the great financial crisis, the additional modes are particularly adverse, and periods of multimodality last for multiple quarters. When tight financial conditions are more contained, such as during the 1987 stock market crash and during the LTCM crisis, the secondary modes are more moderate and the economy resolves quicker to the good mode.

In sum, Figure 4a shows that the joint distribution of economic and financial conditions exhibits distinctly non-Gaussian behavior during periods of tight financial conditions. As we showed in Adrian et al. (2019), a univariate model for the evolution of real GDP growth, even

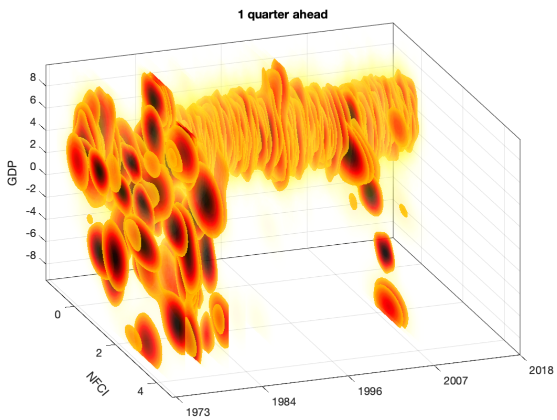

(a) One quarter ahead

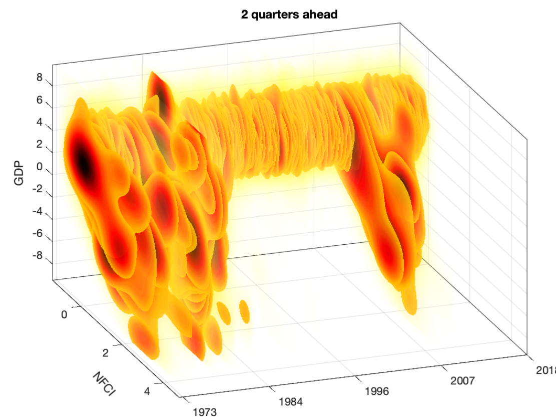

(b) Two quarters ahead

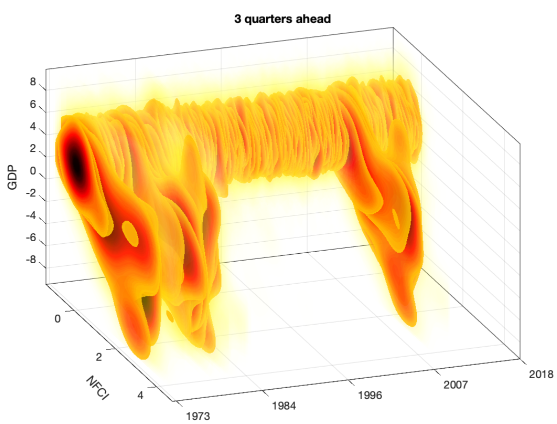

(c) Three quarters ahead

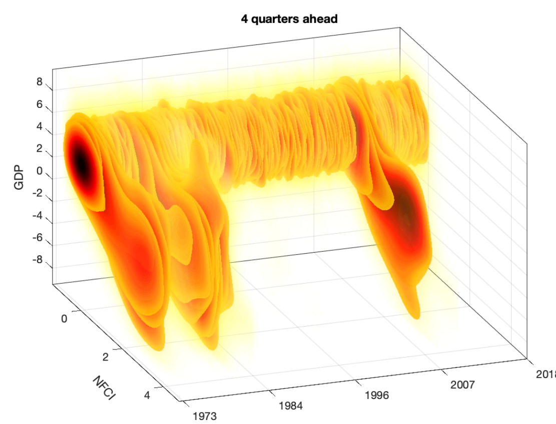

(d) Four quarters ahead

Figure 4. Joint distributions across horizons. The figure shows the joint distribution for one, two, three and four quarter ahead forecasts, as contour plots of the joint distribution in every quarter in the sample, with darker shades of red correspond to higher probability densities.

a non-parametric one, does not capture the multimodality in the distribution of real GDP growth that arises during periods of tight financial conditions. Conditioning the evolution of real GDP growth on the tightness of financial conditions is thus crucial in uncovering the time-varying multimodality we document in this paper. That is, financial conditions Granger-cause real GDP growth in a distributional sense.

Figures 4b – 4d demonstrate that this non-Gaussianity at the one-quarter horizon persists at longer forecast horizons. The tightening of financial conditions during the Volcker disinflation episode – the tightest financial conditions in our sample – generates pronounced non-Gaussianity in the joint distribution of economic and financial conditions up to four quarters out. During the great financial crisis, the forecasting distribution is non-Gaussian up to three quarters out; corresponding to the intuition that more pronounced and widerspread tightenings of financial conditions generate more persistent periods of multimodality, the non-Gaussianity after the 1987 stock market crash and the LTCM crisis get resolved within two quarters.

Figure 5 illustrates this mechanism by plotting the joint distribution across horizons up to two years ahead (eight quarters ahead), conditional on real GDP growth realization at the bottom fifth percentile, the median, and the top fifth percentile (moving across rows) and on financial conditions realization at the bottom fifth percentile, the median, and the top fifth percentile (moving across columns). The figure shows that, when financial conditions are tight (top fifth percentile), multimodality emerges and persists at even longer horizons regardless of the realization of real GDP growth. Even if real GDP growth is in the top fifth percentile of historical realizations, if financial conditions are in the tightest fifth percentile, the economy is fragile and additional modes are possible even two years out. In contrast, if real GDP growth is in the bottom fifth percentile of historical realizations, the distribution is non-Gaussian but unimodal but the non-Gaussianity disappears at horizons longer than two quarters, suggesting that the transmission mechanism between financial and GDP growth is asymmetric. If exceptionally tight financial conditions coincide with exceptionally adverse GDP growth, so that both real GDP growth and NFCI are in their respective worst fifth quantile, the one-quarter-ahead distribution is unimodal but concentrated in the bad mode, and multimodality is present from two quarters out.

Figure 6 shows that small non-Gaussianities in the distribution of QoQ real GDP growth cumulate into large non-Gaussianities for real GDP growth over multiple quarters. That

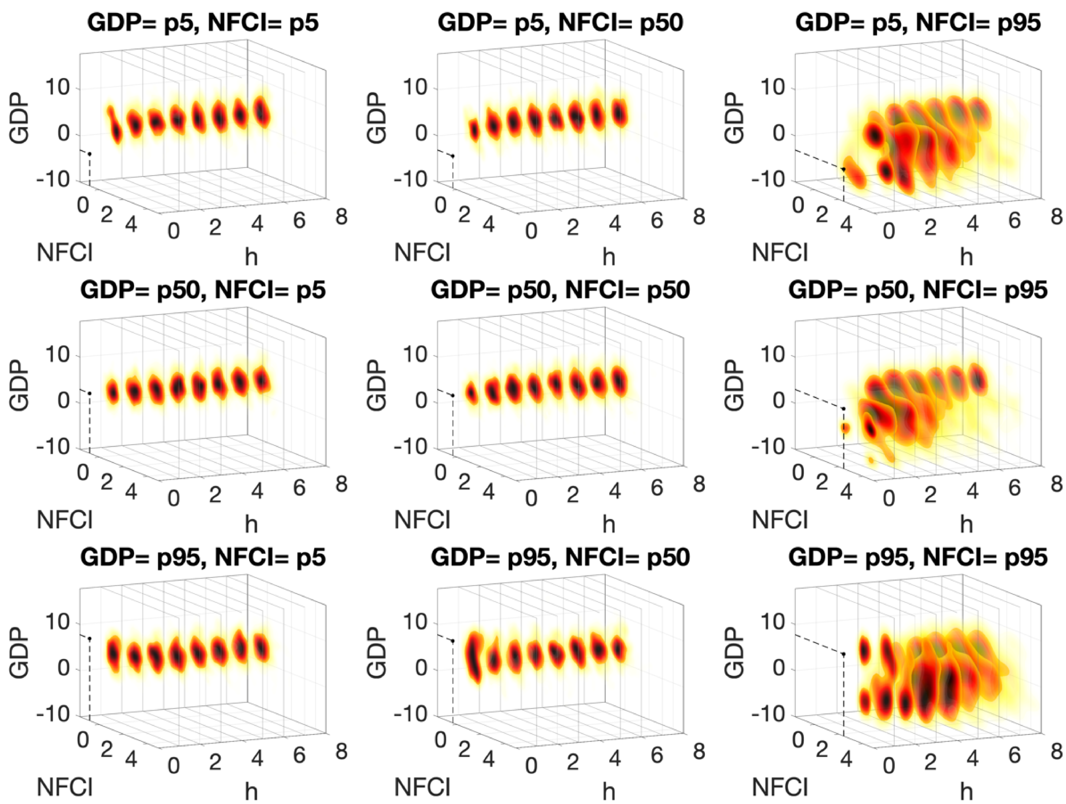

Figure 5. Inspecting the mechanism. The figure displays the evolution of joint QoQ real GDP growth and NFCI forecasts by varying the initial conditions (5 percentile, 50 percentile, and 95 percentile) for up to eight quarters ahead. Darker shades of red highlight regions of higher probability densities.

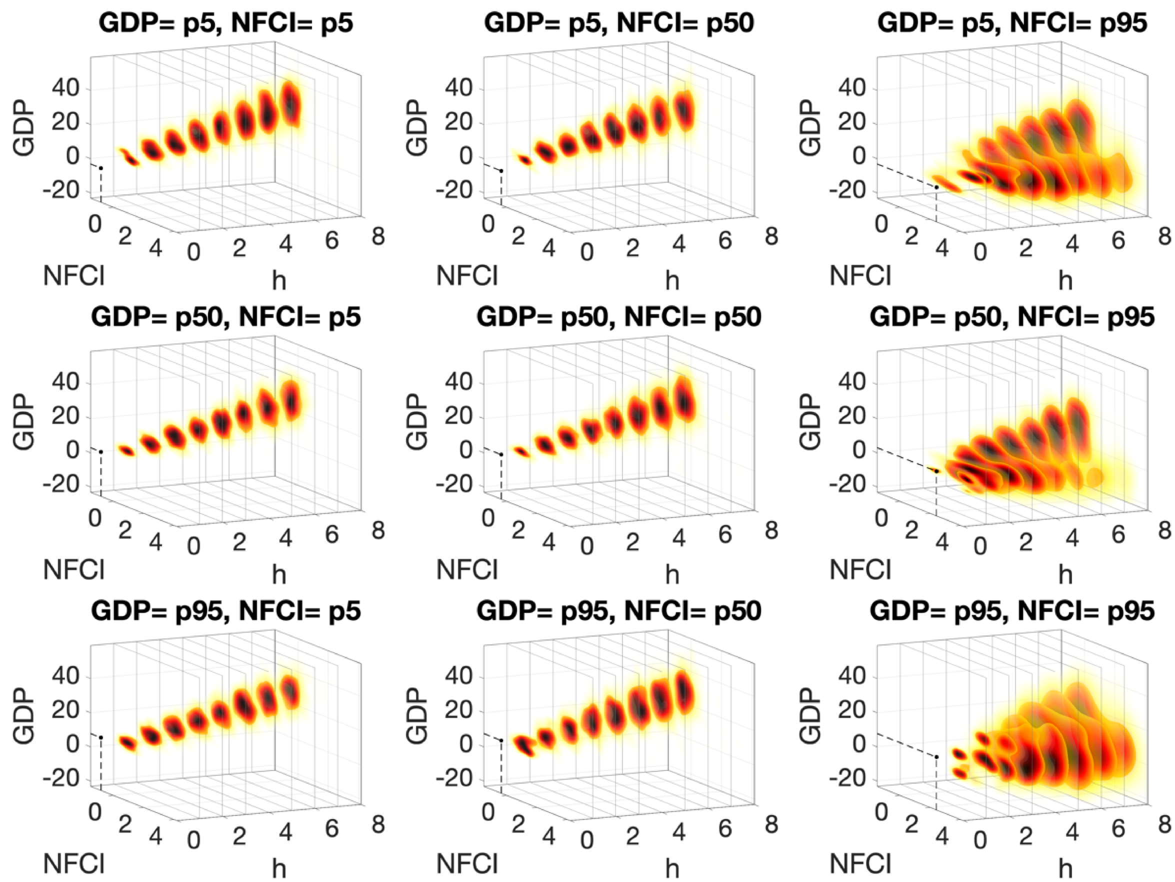

Figure 6. Distribution of cumulative growth. The figure displays the evolution of joint cumulative real GDP growth and NFCI forecasts by varying the initial conditions (5 percentile, 50 percentile, and 95 percentile) for up to eight quarters ahead. Darker shades of red highlight regions of higher probability densities.

is, even if the h-quarter ahead joint distribution of QoQ real GDP growth and financial conditions is roughly Gaussian, the joint distribution of h-quarter real GDP growth and financial conditions can still have multiple modes if the one-quarter-ahead is multimodal: although the distribution resolves eventually to the "good" state, how long the resolution takes can have profound implications on the level of GDP growth. Thus, macrofinancial linkages are at the heart of multimodality, which can manifest in two ways. If current financial conditions are tight but GDP growth are not particularly adverse, the predicted short- and medium-horizon distributions assign non-negligible probability both to the possibility of financial conditions loosening and GDP growth improving, and to the possibility of financial conditions tightening further and GDP growth declining. If, on the other hand, current financial conditions are tight and GDP growth are adverse, the short-run forecast keeps the conditional distribution concentrated at the bad mode but allows for the possibility of returning to the looser financial conditions, improved GDP growth state in the longer run.

4.2 Case study: Great Recession

> Instead of catalyzing recovery, the financial system is working against recovery.[. . .] This is a dangerous dynamic, and we need to arrest it. [. . .] We believe that action has to be sustained until recovery is firmly established. In the United States in the 30s, Japan in the 90s, and in other cases around the world, previous crises lasted longer and caused greater damage because governments applied the brakes too early. We cannot make that mistake.

Remarks introducing the Capital Assistance Program for the U.S. financial system, Treasury Secretary Geithner, February 10, 2009.[5]

In his remarks introducing the Capital Assistance Program, then-Secretary Geithner pointed to the possibility that absence of decisive policy action could lead to a negative outcome for financial conditions with persistent adverse economic consequences. Figure 7 illustrates how these dynamics played out during the Great Recession in the joint distribution of real GDP growth and financial conditions, with different columns corresponding to forecast horizons from one (leftmost column) to four (rightmost column) quarters, and different rows corre-

[5]See https://www.treasury.gov/press-center/press-releases/Pages/tg18.aspx

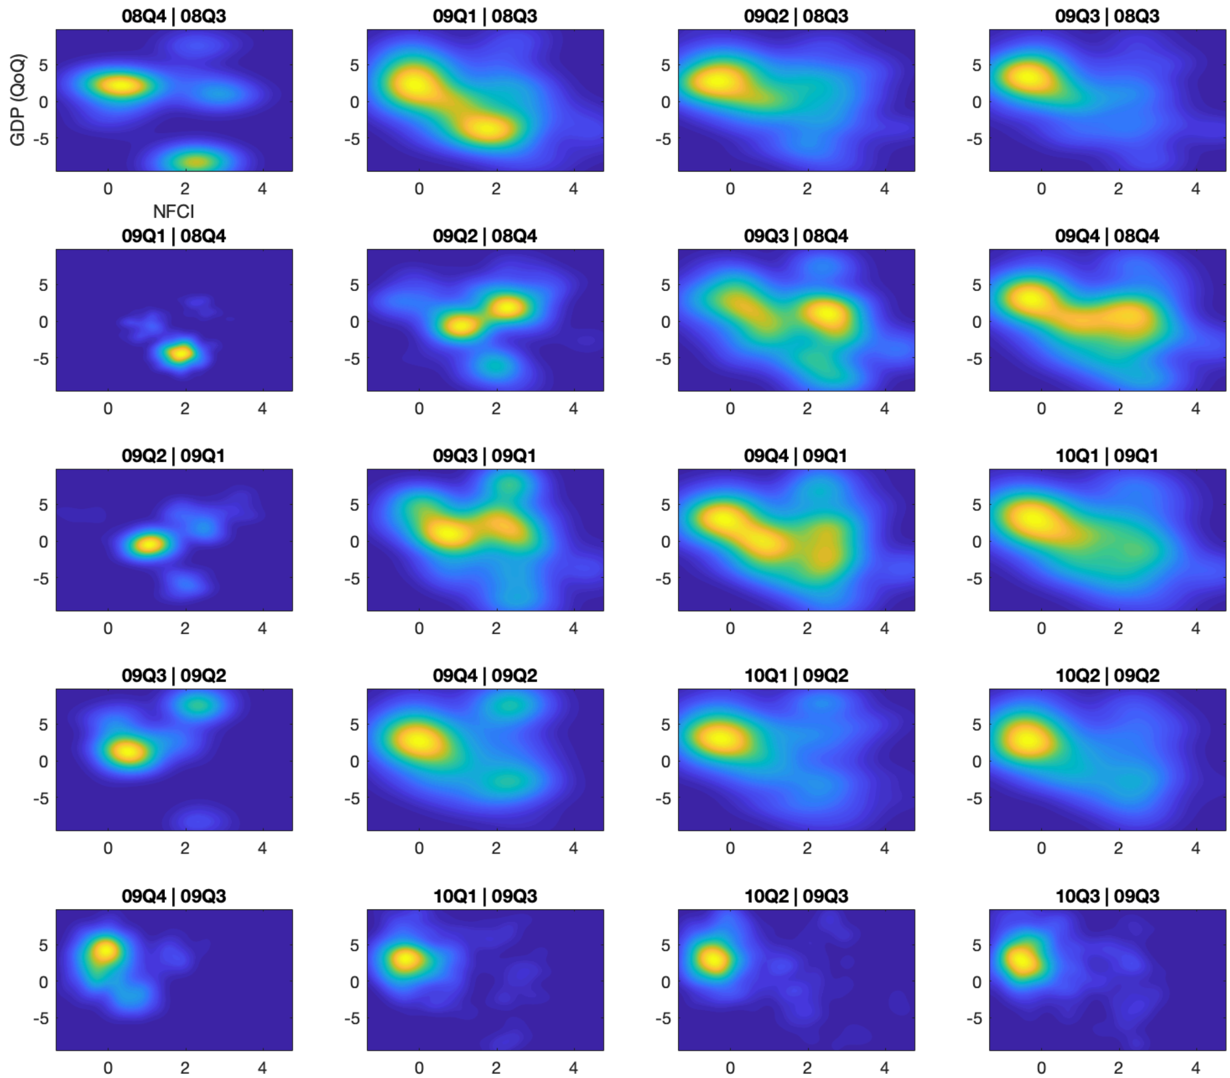

Figure 7. Joint distribution of economic and financial conditions during the Great Recession. Contour plots of 1–4 quarter-ahead density forecasts of real GDP growth and NFCI during the financial crisis. Brighter colors indicate great probability.

sponding to different conditioning information from Q3 2008 (top row) to Q3 2009 (bottom row).

Consider first the four distributions predicted using information as of Q3 2008, plotted in the top row of Figure 7. At the one quarter horizon, the predicted distribution has two easily distinguishable modes: a "good" mode, with higher growth and looser financial conditions, and a "bad" mode, with lower growth and tighter financial conditions. Importantly, the predicted distribution as of Q3 2008 assigns greater probability mass to the bad mode, and the good mode has close to zero real GDP growth and roughly average tightness of financial conditions (NFCI around 0). As we extend the horizon of the forecast, the predicted distribution assigns progressively more weight to the good mode, so that the multimodality clearly visible at the one and two quarter horizons becomes more muddled at the three and four quarter horizon, and the good outcome becomes the prevalent mode.

So how do the predicted distributions change as we move through the crisis? Focusing on the one-quarter-ahead distributions plotted in the left most column, we see that the outlook gets worse before it gets better: while the distributions forecasted as of Q4 2008 and Q1 2009 still exhibit multimodality, the good mode all but disappears. The resolution of the joint dynamics of economic and financial conditions to the good mode starts in Q2 2009 with the re-emergence of the good secondary mode to the distribution, and the full normalization to a unimodal joint distribution is completed in Q3 2009. The longer-horizon predicted distributions follow the same pattern. The two- and three-quarter-ahead distributions concentrate less mass on the good mode in Q4 2008 and Q1 2009 but start to resolve to the unimodal distribution in Q2 2009. The four-quarter-ahead distribution shows the possibility for additional modes emerging through Q2 2009 but also resolves to the unimodal norm in Q3 2009.

In sum, Figure 7 suggests that although the forecast distribution always resolves back to the good outcome at longer forecast horizons, the transition to that state can be long and arduous. The longer the economy takes to resolve to the good mode, the greater the potential losses to the level of economic activity. We leave the examination of cross-country differences in how quickly the economy resolves to the good mode to future research.

5 Density impulse response (DIR)

We turn now to the construction of counterfactual predictive densities in the spirit of Gallant et al. (1993). Comparing the (counterfactual) evolution of the conditional density after a perturbation of the one-period-ahead predictive density to the baseline evolution allows us to compute the *density* impulse response function (DIR), tracking how the entire joint distribution of economic and financial conditions responds dynamically to an initial shock. DIRs can be used in a variety of settings, from evaluating the potential policy effects to constructing dynamically consistent stress testing scenarios. In the present paper, we focus on the former application, comparing the effectiveness of improving financial and GDP growth during the financial crisis.

5.1 Constructing DIRs

Recall that the basic building block of our term structure of conditional joint distributions is the one-step-ahead conditional joint distribution, $\hat{p}\left(y_t|\, x_t\right)$. Generically, DIRs are constructed by perturbing this initial distribution, creating, say, $\hat{\rho}\left(y_t|\, x_t\right)$, and then following the procedure in Section 3.2 under the baseline model to propagate this initial disturbance across horizons to generate $\{\hat{\rho}\left(y_{t+h}|\, x_t\right)\}_{h=0}^{H}$:

Algorithm 3. ***Generating density impulse responses.***
The counterfactual distributional term structure $\{\hat{\rho}\left(y_{t+h}|\, x_t\right)\}_{h=0}^{H}$ is generated by drawing Monte Carlo paths as:

1. *Draw υ_t from the desired perturbed one-period-ahead distribution $\hat{\rho}\left(y_t|\, x_t\right)$.*

2. *Conditional on the draw υ_t, draw the two-period-ahead realization υ_{t+1} from the baseline distribution $\hat{p}\left(y_{t+1}|\, \upsilon_t\right)$.*

3. *Repeat Step 2 to horizon H.*

The density impulse response (DIR) of distribution $\hat{p}\left(y_t|\, x_t\right)$ to the initial disturbance $\hat{\rho}\left(y_t|\, x_t\right)$ is then the difference between $\{\hat{p}\left(y_{t+h}|\, x_t\right)\}_{h=0}^{H}$ and $\{\hat{\rho}\left(y_{t+h}|\, x_t\right)\}_{h=0}^{H}$ for every $h = 0, \ldots, H$.

That is, conditional on draws one-period-ahead, both the baseline and the counterfactual are performed in the same manner; the only difference is the one-period-ahead joint distribution.

Importantly, since DIRs are constructed by perturbing the *conditional* joint distributions, the resultant DIRs are state-dependent: a DIR conditional on a normal state of the economy will look different from a stress-period DIR.

As a warm-up exercise before introducing the DIRs we study in this paper, consider a standard linear VAR for economic and financial conditions. Denote by e_t the GDP growth in quarter t and by f_t the financial conditions in quarter t, so that $y_t = (e_t, f_t)'$, and consider a VAR of the form

$$y_t = A(L)\, y_{t-1} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \Sigma),$$

where $A(L)$ is a lag polynomial of order L. Using a Cholesky decomposition to identify shocks, with financial conditions ordered last, we can represent the one-step-ahead joint conditional distribution as

$$\begin{aligned} p\left(y_t \middle| y^{(t-1)}\right) = & \underbrace{\left(2\pi\sigma_{\epsilon,1}^2\right)^{-\frac{1}{2}} \exp\left(-\frac{\left(e_t - A_e(L)\, y_{t-1}\right)^2}{2\sigma_{\epsilon,1}^2}\right)}_{\text{Marginal over } e} \\ & \times \underbrace{\left(2\pi\left(\sigma_{\epsilon,12} + \sigma_{\epsilon,2}^2\right)\right)^{-\frac{1}{2}} \exp\left(-\frac{\left(f_t - A_f(L)\, y_{t-1}\right)^2}{2\left(\sigma_{\epsilon,12} + \sigma_{\epsilon,2}^2\right)}\right)}_{\text{Conditional over } f}. \end{aligned}$$

A traditional impulse response function to, for example, a financial conditions shock S_f thus propagates as a change to the mean of the conditional joint distribution, so that the counterfactual one-step-ahead distribution is given by

$$\begin{aligned} \rho_f\left(y_t \middle| y^{(t-1)}\right) = & \underbrace{\left(2\pi\sigma_{\epsilon,1}^2\right)^{-\frac{1}{2}} \exp\left(-\frac{\left(e_t - A_e(L)\, y_{t-1}\right)^2}{2\sigma_{\epsilon,1}^2}\right)}_{\text{Marginal over } e} \\ & \times \underbrace{\left(2\pi\left(\sigma_{\epsilon,12} + \sigma_{\epsilon,2}^2\right)\right)^{-\frac{1}{2}} \exp\left(-\frac{\left(f_t - A_f(L)\, y_{t-1} - S_f\right)^2}{2\left(\sigma_{\epsilon,12} + \sigma_{\epsilon,2}^2\right)}\right)}_{\text{Counterfactual conditional over } f}. \end{aligned}$$

To keep the analogy with traditional VARs, in this paper, we examine DIRs constructed by perturbing the univariate conditional distributions. Denote by e_t the GDP growth in quarter t and by f_t the financial conditions in quarter t, so that $y_t = (e_t, f_t)'$ and $x_t =$

$(e_{t-1}, f_{t-1})'$. Then we can factor the conditional joint distribution function $\hat{p}(y_t|\, x_t)$ as

$$\hat{p}(y_t|\, x_t) = \underbrace{\hat{p}_{m,e}(e_t|\, x_t)}_{\text{Marginal over } e} \underbrace{\hat{p}_{c,f}(f_t|\, e_t, x_t)}_{\text{Conditional over } f} = \underbrace{\hat{p}_{c,e}(e_t|\, f_t, x_t)}_{\text{Conditional over } e} \underbrace{\hat{p}_{m,f}(f_t|\, x_t)}_{\text{Marginal over } f}.$$

We define the DIR to financial conditions as the difference between the baseline estimated distribution and the counterfactual distribution constructed by perturbing the conditional distribution over financial conditions while keeping the marginal distribution over GDP growth fixed:

$$\hat{\rho}_f(y_t|\, x_t) \equiv \hat{p}_{m,e}(e_t|\, x_t)\, \hat{\rho}_{c,f}(f_t|\, e_t, x_t).$$

The DIR to financial conditions thus examines how the joint distribution of economic and financial conditions evolves in response to a shock that does not have a contemporaneous effect on the marginal distribution of GDP growth. Comparing to the impulse response function in a linear VAR, we can see three main differences. First, while impulse response functions in a linear VAR only propagate changes to the mean of the distribution, keeping the other moments on the distribution unchanged, while we can consider DIRs to changes in other moments of the distribution, such as a shock to the skew of the distribution. This implies a second difference: generically, DIRs change the entire shape of the conditional distribution, not just the mean as in the case of the linear VAR. Finally, since the estimated dynamics is highly non-linear, DIRs depend on the initial conditions.

Similarly, the DIR to GDP growth is the difference between the baseline estimated distribution and the counterfactual distribution constructed by perturbing the conditional distribution over GDP growth while keeping the marginal distribution over financial conditions fixed:

$$\hat{\rho}_e(y_t|\, x_t) \equiv \hat{\rho}_{c,e}(e_t|\, f_t, x_t)\, \hat{p}_{m,f}(f_t|\, x_t).$$

This is the distributional equivalent to an impulse response function to a shock to GDP growth in a VAR with Cholesky-identified shocks, where the GDP growth are ordered second.

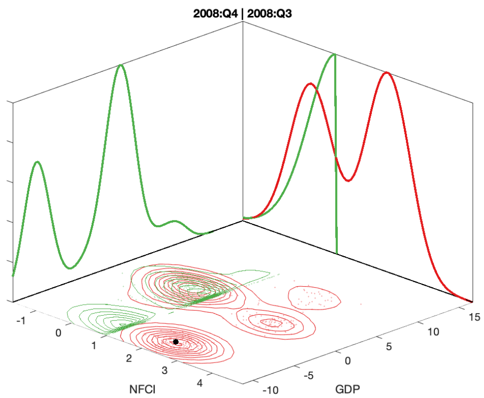

(a) One quarter ahead

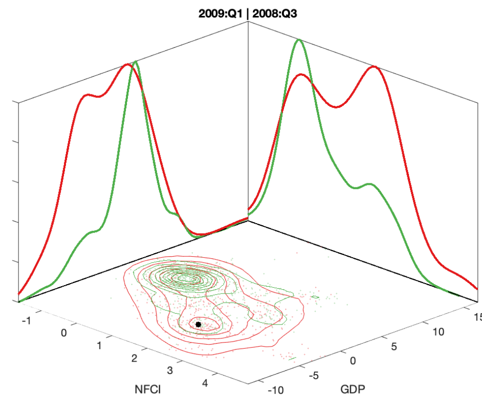

(b) Two quarters ahead

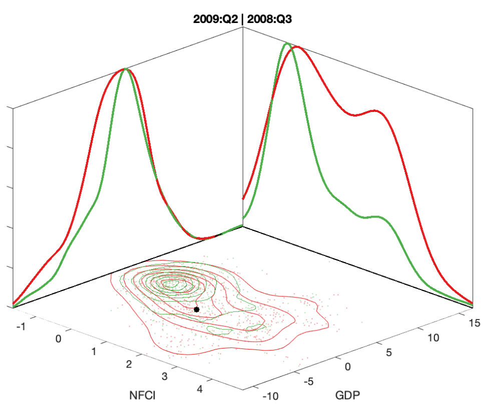

(c) Three quarters ahead

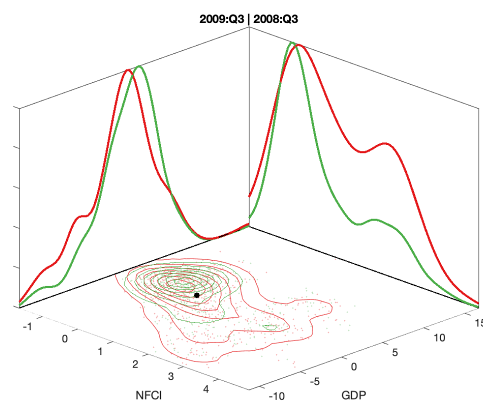

(d) Four quarters ahead

Figure 8. Baseline and counterfactual distributions for a financial conditions DIR. The figure plots how the joint distribution of economic and financial conditions changes in response to a shock to the conditional distribution of financial conditions in Q4 2008| Q3 2008 without a contemporaneous shock to the marginal distribution of GDP growth. The subfigures show the counterfactual impact of the perturbation $1-4$ quarters ahead. Baseline distribution plotted in red; counterfactual in green, with the marginal distributions for real GDP growth and NFCI shown on the sides and the joint density as contour plots in the center of each panel.

5.2 DIR to financial conditions

We begin by examining the in-sample DIR to financial conditions during the financial crisis. More specifically, we perturb the conditional distribution of financial conditions as of Q3 2008, moving some of the mass from the upper tail of the distribution (corresponding to tight financial conditions) to the center of the distribution, thus making "bad" financial conditions outcomes less likely. This DIR thus answers the question: If, in Q3 2008, policy were able to limit the possibility of extreme tightening of financial conditions during Q4 2008 without affecting the distribution of possible economic outcomes in the same quarter, what would the predicted joint distribution for the remaining quarters of the financial crisis have looked like?

Figure 8 plots the baseline (in red) and counterfactual (in green) distributions one through four quarters ahead.[6] Figure 8a shows that, at the one quarter ahead horizon, removing the upper tail of the financial conditions distribution is insufficient to remove the multimodality: the counterfactual distribution still allows for the possibility of an adverse GDP growth, tight financial conditions outcome in Q4 2008, though, by construction, the tightness of financial conditions in that mode is more mild than in the bad mode under the baseline distribution. The multimodality is resolved (to the good mode) at two quarters ahead, as shown in Figure 8b. Thus, removing the possibility of extremely tight realizations of financial conditions in Q4 2008 is sufficient to eliminate the possibility of the low economic activity and tight financial conditions in Q1 2009. Propagating this further, we see from Figures 8c – 8d that the counterfactual distributions three and four quarters ahead also exhibit thinner "bad" tails for both economic and financial conditions, where the "bad" outcomes are in the left tail for GDP growth but in the right tail for financial conditions.

5.3 DIR to GDP growth

Consider now the in-sample DIR to GDP growth during the financial crisis. Analogously to the DIR to financial conditions, we perturb the conditional distribution of GDP growth as of Q3 2008, moving some of the mass from the lower tail of the distribution (corresponding to adverse GDP growth) to the center of the distribution, thus making "bad" GDP growth

[6] Note that, by construction, the one quarter ahead baseline and counterfactual marginal distributions over economic conditions in Figure 8a coincide.

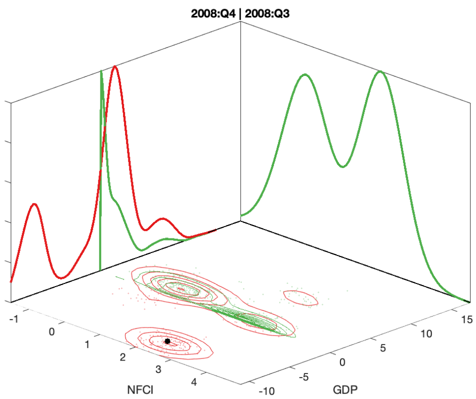

(a) One quarter ahead

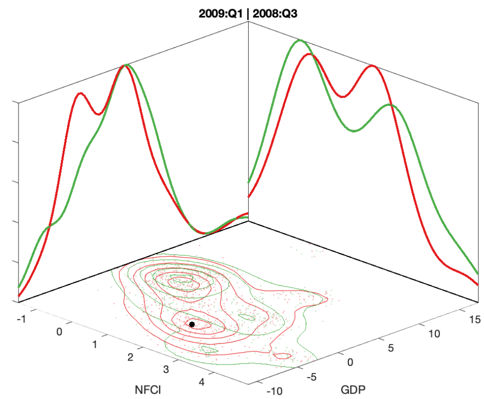

(b) Two quarters ahead

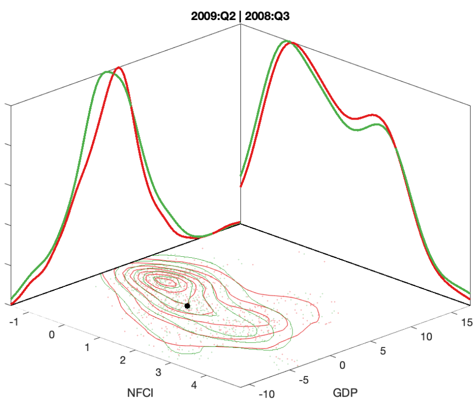

(c) Three quarters ahead

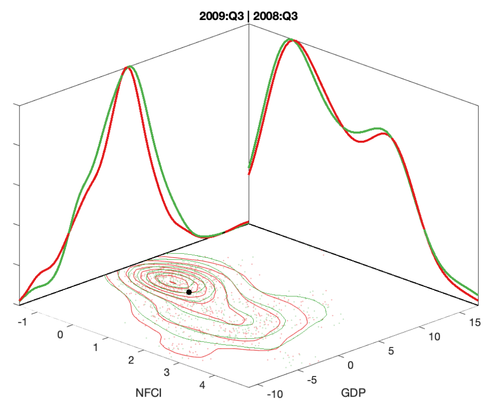

(d) Four quarters ahead

Figure 9. Baseline and counterfactual distributions for a GDP growth DIR. The figure plots how the joint distribution of GDP growth and financial conditions changes in response to a shock to the conditional distribution of GDP growth in Q4 2008| Q3 2008 without a contemporaneous shock to the marginal distribution of financial conditions. The subfigures show the counterfactual impact of the perturbation $1-4$ quarters ahead. Baseline distribution plotted in red; counterfactual in green, with the marginal distributions for real GDP growth and NFCI shown on the sides and the joint density as contour plots in the center of each panel.

outcomes less likely. This DIR thus answers the question: If, in Q3 2008, policy were able to limit the possibility of extremely adverse GDP growth during Q4 2008 without affecting the distribution of possible financial outcomes in the same quarter, what would the predicted joint distribution for the remaining quarters of the financial crisis have looked like?

Figure 9 plots the baseline (in red) and counterfactual (in green) distributions one through four quarters ahead. As with the DIR to financial conditions, Figure 9a shows that, at the one quarter ahead horizon, removing the lower tail of the real GDP growth distribution is insufficient to remove the multimodality: the counterfactual distribution still allows for the possibility of low GDP growth, tight financial conditions outcome in Q4 2008, though, by construction, real GDP growth in that mode is more positive than in the bad mode under the baseline distribution. The multimodality persists at the remaining horizons, as shown in Figures 9b–9d. Moreover, the two-quarter-ahead distribution of financial conditions under the alternative has somewhat fatter tails than under the baseline distribution; at three and four quarters ahead, the baseline and the counterfactual joint distributions are nearly identical. Thus, removing the possibility of negative real GDP growth in Q4 2008 without a corresponding improvement in the distribution of financial conditions improves the distribution of real GDP growth in the short run, but at the cost of potentially tighter financial conditions and without large long-run effects.

6 Out-of-sample evidence and alternative estimators

In this section, we evaluate the out-of-sample performance of the non-parametric estimate of the joint distribution of real GDP growth and financial conditions. We backtest the model by replicating the analysis that an economist would have done by using the proposed methodology in "real time", with the caveat that we use final revised data only as real-time data for NFCI are only available for the recent past. We produce predictive distributions using an expanding window, starting with the estimation sample that ranges from Q1 1973 to Q3 1982. We perform two types of out-of-sample analysis. First, we show that the out-of-sample distributions are similar to the in-sample distributions, in general and during the financial crisis in particular. Second, we evaluate the out-of-sample accuracy and calibration of the density forecasts relative to two VAR alternatives, analyzing the predictive score and the probability integral transform (PIT).

6.1 Out-of-sample evidence

We begin by comparing the in-sample and out-of-sample predicted distributions, plotted in Figure 10. The figure shows that the in-sample and out-of-sample estimates of the one- and four-quarter-ahead quantiles are virtually indistinguishable for both real GDP growth and financial conditions. The similarities are more striking as the Great Recession is a significant tail event that is not in the sample when estimating the out-of-sample joint distribution. When the out-of-sample quantiles do deviate noticeably from the in-sample estimates, as is the case for the four-quarters-ahead forecast of financial conditions in the 1980s, the out-of-sample forecast frequently assigns greater probability to more adverse outcomes. That is, when the recursive forecast is "wrong" relative to the full-sample one, the recursive forecast is more conservative.

To understand the differences between the in-sample and out-of-sample forecasts better, consider the one-year ahead forecast computed in real-time (out-of-sample) during the Great Recession, plotted in the right column of Figure 11. For comparison, the in-sample forecasts are reported in the left column of the Figure. Much like the in-sample forecast, the out-of-sample distribution during the Great Recession exhibits multimodality. At the beginning of the recession (forecasts as of Q3 2008 and Q4 2008), the out-of-sample multimodality is relatively mild, in particularly missing how tight financial conditions could and would get during the financial crisis. Multimodality emerges one quarter later in the out-of-sample predictive distribution. Thus, even though the out-of-sample distribution as of Q3 2008 was somewhat more benign than the in-sample distribution, the possibility for multimodality with more adverse outcomes was detected in real time, suggesting that the real GDP losses during the Great Recession could have been even larger than observed in practice.

A natural question to ask is what informs the multimodality detected by the out-of-sample estimator during the Great Recession. As we discussed in Section 4.1, prior to the Great Recession, multimodality in the one-quarter-ahead distribution arises during the Volcker disinflation episode in the early 1980s, following the stock market crash in 1987, and following the LTCM crisis. Thus, from the perspective of the out-of-sample estimator during the Great Recession, there were already prior periods where multimodality was present and persistent – for example, the multimodality during the Volcker disinflation episode lasted more than 4 quarters – and the estimator uses those experiences to shape the out-of-sample

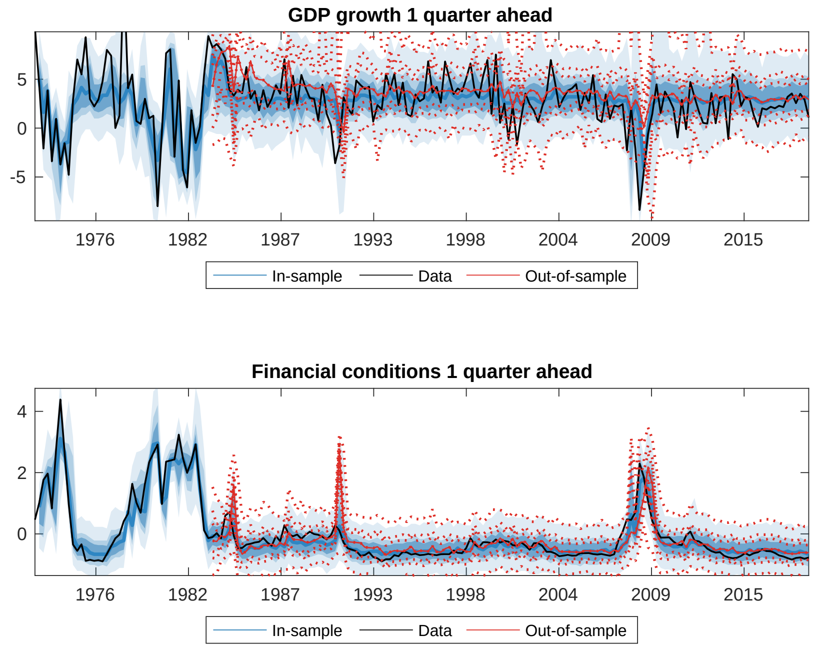

(a) One quarter ahead

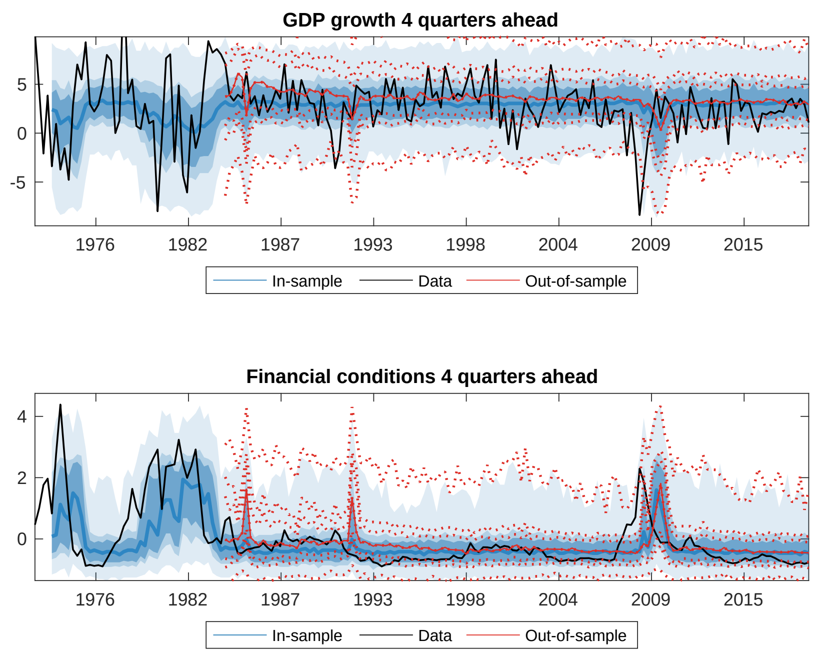

(b) Four quarters ahead

Figure 10. Out-of-sample quantiles. The figure plots the marginal distributions of real GDP growth and financial conditions one- and four-quarters-ahead in-sample (blue shaded area) and out-of-sample (red dashed lines), together with the data realization. First out-of-sample forecast as of Q2 1982.

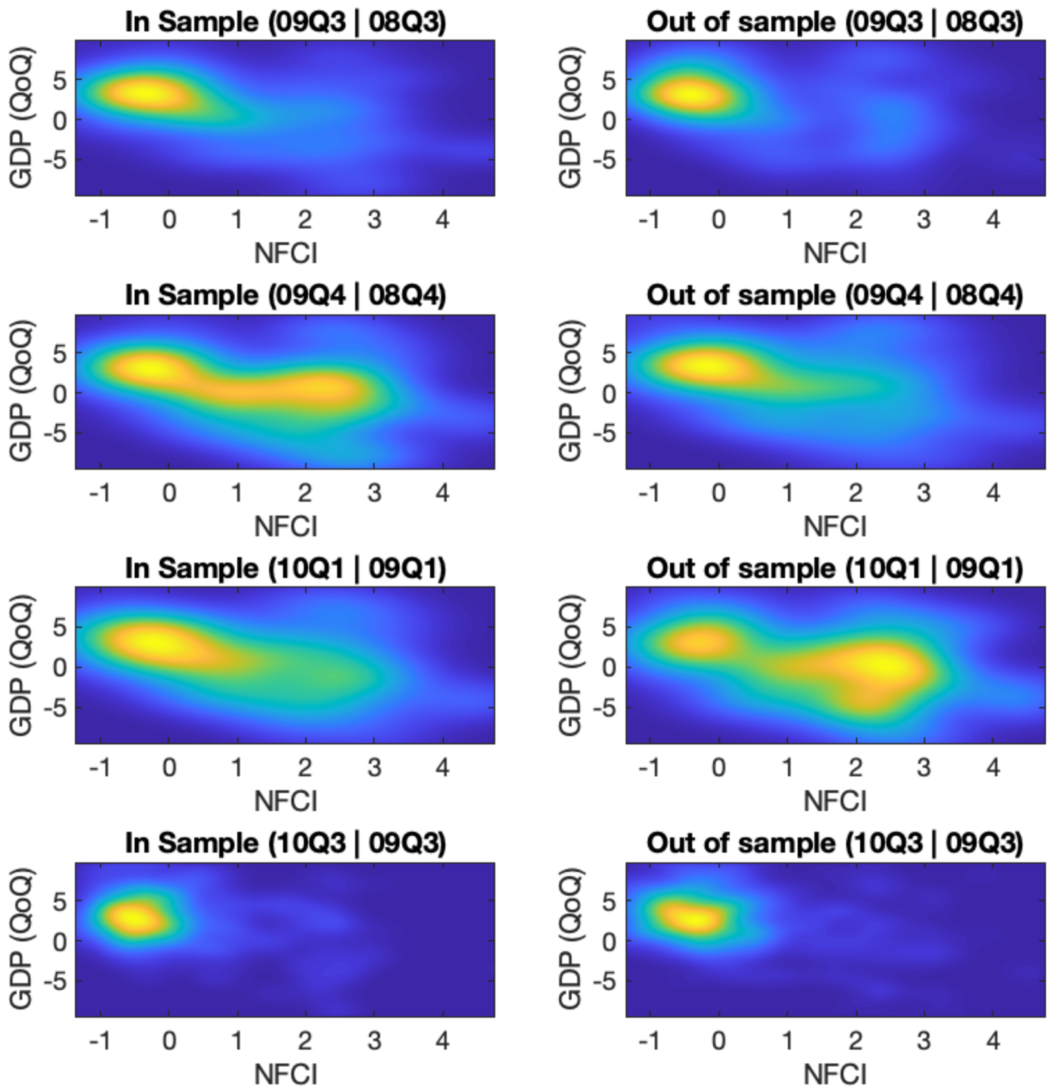

Figure 11. Out-of-sample forecast of the Great Recession. The figure reports in-sample joint densities (left column) and the out-of-sample results (right column) during the Great Recession.

predicted distribution during the Great Recession. More fundamentally, in order to posit the possibility of multiple modes, the estimator needs to have observed data corresponding to both types of modes in the past. For example, there were periods during the 1970s when tightening in financial conditions resolved to yet tighter financial conditions and sharp economic downturns but also periods when tightening in financial conditions resolved to a retracing in financial conditions and only mild economic downturns. These historical episodes are what inform the estimator as to the possibility of both types of resolution to tight financial conditions.

6.2 Alternative models

We now turn to evaluating the out-of-sample performance of our non-parametric estimator relative to two alternatives: a vector autoregression (VAR) with Gaussian errors and a VAR with non-parametric errors. More specifically, our non-parametric estimation method allows us to model the joint evolution of real GDP growth and NFCI as

$$y_{t+1} = f\left(y_t, \epsilon_{t+1}\right),$$

with implied innovations ϵ_{t+1} distributed according to $f\left(y_t, \epsilon_{t+1}\right) \sim \hat{p}\left(y_t\right)$. In Figures 12 – 13, we denote this model "NL-VAR(1)", a non-linear first order vector autoregression. Our two alternatives both take the form

$$y_{t+1} = \rho y_t + \epsilon_{t+1}.$$

In the first alternative, a linear first order VAR with Gaussian errors, $\epsilon_{t+1} \sim \mathcal{N}\left(0, \Sigma\right)$; we denote this model by "VAR(1)". The second alternative maintains the assumption of linearity but does not impose a distributional assumption on the innovations ϵ_{t+1}; we denote this model by "VAR(1)-NP". Comparing the predictions from the fully non-parametric model with those from the linear model with non-parametric errors allows us evaluate whether the multimodality we observe during periods of tightened financial conditions arises due to non-Gaussian innovations or due to non-linear transmission.

We begin by comparing the sharpness of the out-of-sample forecasts under the three models, evaluated as the log-predictive score. Figure 12 plots the log-predictive score over

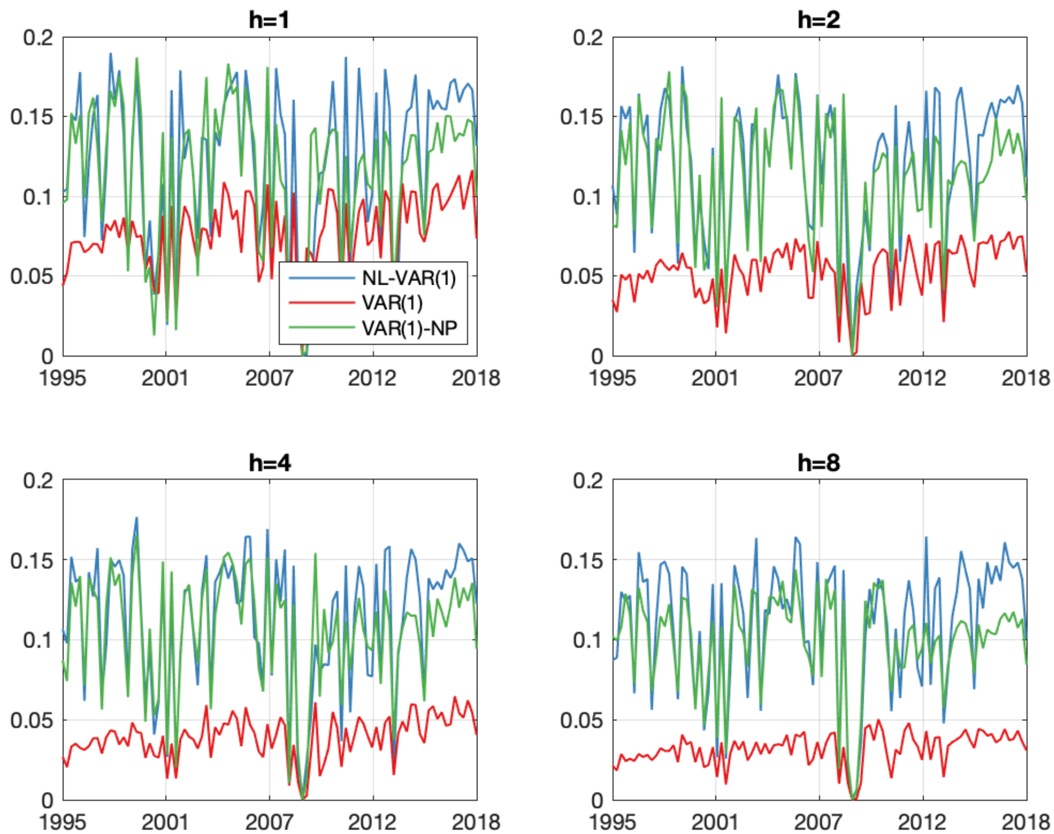

Figure 12. Forecast sharpness. The figure reports scores for the joint predictive distribution of real GDP growth and NFCI at the one-quarter-, two-quarter-, four-quarter-, and eight-quarter-ahead horizon. Log-predictive-scores for the fully non-parametric model plotted in blue ("NL-VAR(1)"); for a linear VAR with Gaussian innovations in red ("VAR(1)"); for a linear VAR with non-parametric errors in green ("VAR(1)-NL").

time at the one-quarter-, two-quarter-, four-quarter-, and eight-quarter-ahead horizon. During good times, the fully non-parametric model out-performs the two linear alternatives at all forecast horizons, and the linear VAR with non-parametric errors outperforms the linear VAR with Gaussian errors. At the one-quarter-ahead horizon, the non-parametric model underperforms the alternatives during the financial crisis. Intuitively, the fit of the linear VAR, and especially the linear VAR with Gaussian errors, is driven by outliers. Thus, at the one-quarter-ahead horizon, the linear VAR fits the crisis relatively well. The non-parametric model, instead, recognizes that the financial crisis is an outlier in the data and thus assigns a relatively low probability to the outcomes that were observed during the crisis, lowering the log-predictive score. At longer horizons, the non-parametric model out-performs the alternatives even during crisis periods. Finally, Figure A.9 in the Appendix shows that the log predictive scores for the non-parametric model estimated using the Gibbs sampler in Section 3.3 almost completely coincides with the log predictive scores when the model is estimated directly.

Consider next the calibration of the predictive distribution, evaluated through the probability integral transform (PIT). For each model and horizon, we compute the empirical cumulative distribution of the PITs, which measures the percentage of observations that are below any given quantile. In a perfectly calibrated model, the cumulative distribution of the PITs is a 45-degree line, so that the fraction of realizations below any given quantile of the predictive distribution is exactly equal to quantile probability; the close the empirical cumulative distribution of the PITs is to the 45-degree line, the better calibrated the model is. Figure 13 plots the PITs for real GDP growth and NFCI, one and four quarters ahead for the three alternative models, together with Rossi and Sekhposyan (2017) 5% confidence bands.[7] For both variables and both forecast horizons, the empirical distribution of the PITs for the non-parametric model is well within the confidence bounds and generally closer to the 45-degree line than the alternatives. Thus, the non-parametric model is better calibrated than both linear VARs. Finally, Figure A.10 in the Appendix shows that the PITs for the non-parametric model estimated using the Gibbs sampler in Section 3.3 almost completely

[7] The confidence bands should be taken as general guidance since they are derived for forecasts computed using a rolling scheme (with a constant size of the estimation sample) while we use an expanding estimation window. For the one-quarter-ahead forecasts, the bands are based on critical values derived under the null of uniformity and independence of the PIT. For the PITs of the four-quarter-ahead predictive distributions, bands are computed by bootstrapping under the assumption of uniformity only.

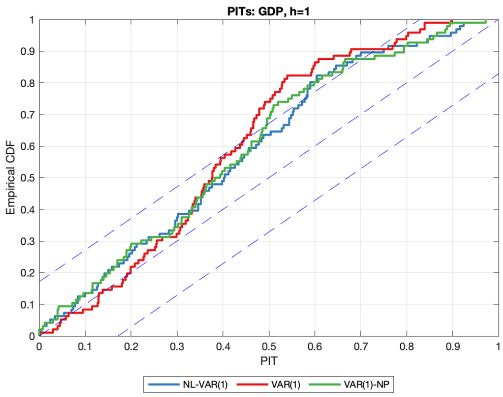

(a) Real GDP growth, one quarter ahead

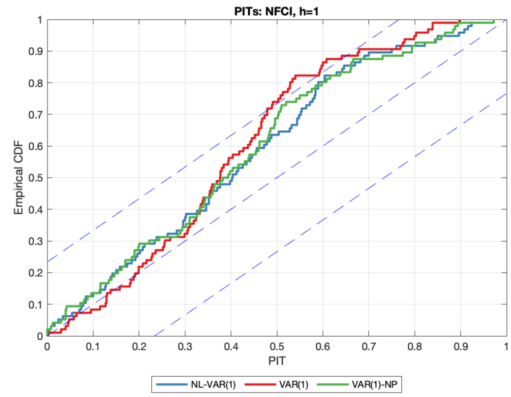

(b) NFCI, one quarter ahead

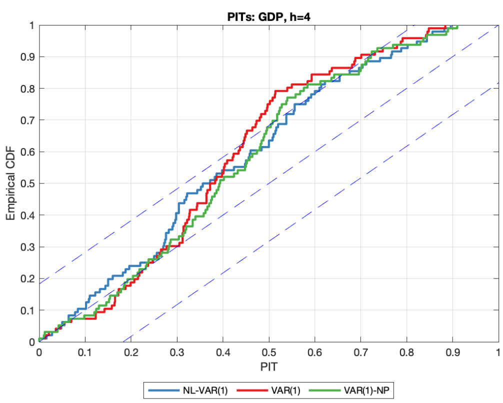

(c) Real GDP growth, four quarters ahead

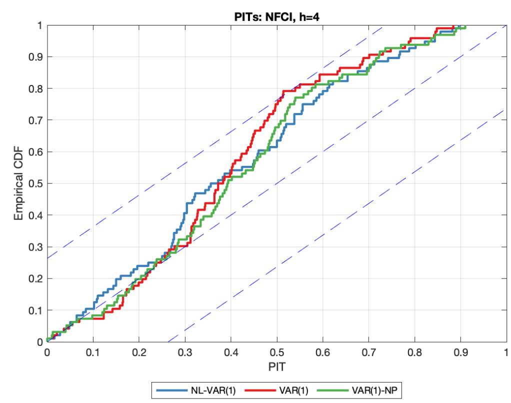

(d) NFCI, four quarters ahead

Figure 13. Calibration of the forecast. The figure reports the empirical cumulative distribution of the marginal probability integral transform (PITs) from real GDP growth and NFCI one quarter and one year ahead. PITs for the fully non-parametric model plotted in blue ("NL-VAR(1)"); for a linear VAR with Gaussian innovations in red ("VAR(1)"); for a linear VAR with non-parametric errors in green ("VAR(1)-NL"); the 5% critical values as in Rossi and Sekhposyan (2017) plotted as dashed-blue lines.

coincides with the PITs when the model is estimated directly.

We conclude our exploration of alternative models by examining the joint distribution of real GDP growth and financial conditions during the Great Recession that is predicted by the linear VAR with the density of the innovations estimated non parametrically, plotted in the last column on Figure 14. Unlike the in-sample (first column) and the out-of-sample (second column) forecasts of the non-linear model, the joint distribution predicted by the linear VAR does not exhibit multimodality. Thus, the multimodality predicted by the non-parametric model cannot be explained by non-Gaussian innovations and non-linear amplification plays a crucial role in macro-financial dynamics.

7 Conclusion

Linear vector autoregressions (VARs) have long been a workhorse of empirical macroeconomic analysis, providing empirical targets for theoretical models and allowing the evaluation of counterfactual outcomes. In this paper, we describe a distributional equivalent to VARs, creating term structures of joint distributions of variables of interest by combining a flexible non-parametric approach to estimating one-step-ahead joint distribution with an efficient Monte Carlo approach. Our method is computationally straightforward to implement, quick to estimate, and appears well behaved both in- and out-of-sample, providing a fully nonparametric alternative to linear vector autoregressions. Instead of the impulse response functions of the VAR literature, we propose density impulse responses, building on the seminal work of Gallant et al. (1993).

We apply this methodology to estimate the joint distribution of economic and financial conditions in the U.S. Estimating the joint distribution allows us to uncover a novel feature of the data: though the conditional joint distribution is unimodal and approximately Gaussian during normal times, multimodality emerges when financial conditions are tight, regardless of how benevolent GDP growth is. When the tightness of financial conditions is not extreme, the multimodality is resolved quickly to the good mode of the distribution. During periods of more extreme tightening of financial conditions, such as the 2007 – 2009 financial crisis and the Volcker disinflation of the early 1980s, the multimodality is present in the predictive distribution up to four quarters ahead. While the theoretical literature has long conjectured the possibility of multimodality in the joint distribution of economic and

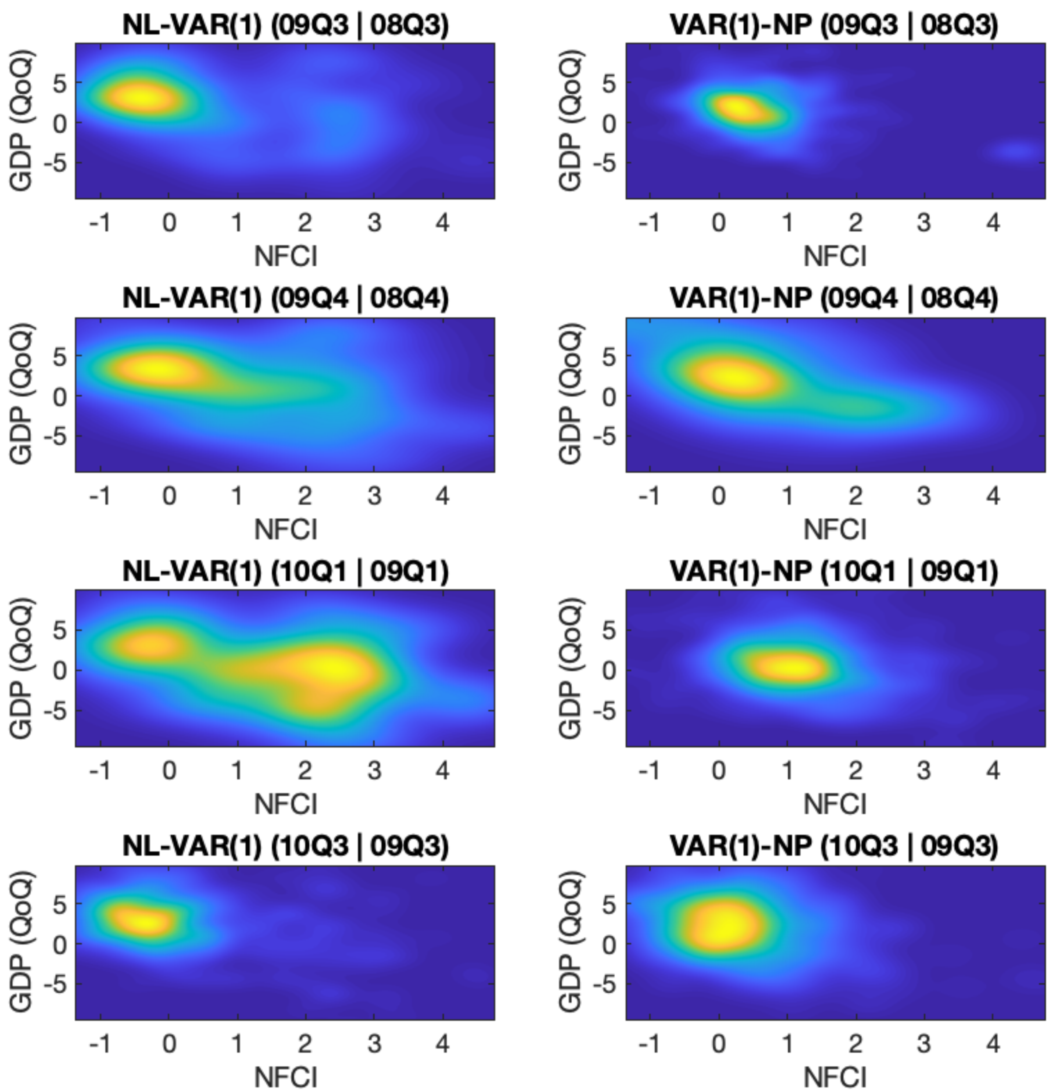

Figure 14. Forecasting the Great Recession in real time with a VAR. The figure reports the joint densities produced during the Great Recession by the baseline fully non-parametric model (left column) and the VAR with density of the innovations estimated non parametrically (right column).

financial conditions, due to either multiple equilibria or non-linear amplification mechanisms, to the best of our knowledge, our paper provides the first empirical evidence of periods of multimodality. Importantly, when the conditional distribution is multimodal, point forecasts are particularly misleading: by averaging between the two (or more) modes of the distribution, a point forecast underestimates both how good the good mode is and how bad the bad mode is.

Over our sample period, the joint distribution always converges to a unimodal distribution concentrated at the good mode in the long run. We conjecture that this is due to the fact that public policy in the U.S. over this period has aggressively aimed to counteract negative economic outcomes. But the Great Depression of 1929 – 1933 might have been an example of insufficiently aggressive policy leading to a resolution of multimodality to a unimodal distribution concentrated at the "bad" mode. That is, we conjecture that multimodality of the distribution also appeared in 1929-33, and that "bad" policy led to equilibrium selection associated with persistently bad outcomes. While we leave a careful examination of this conjecture for future research, it is worth noting that the joint distribution during the financial crisis exhibits some of these features. The distribution as of Q4 2008 showed some evidence of resolving to the bad mode, with the good mode all but disappearing from the one-quarter-ahead distribution. Moreover, even if the joint distribution always resolves to the good mode, prolonged periods of multimodality can lead to substantial losses in the level of economic output along the resolution path.

Periodic multimodality arises organically in our non-parametric approach: rather than postulating that multimodality (or non-Gaussianity more generally) exists, we allow the data to inform us about the shape of the joint distribution over time. Out-of-sample evaluation demonstrates that the periodic multimodality we detect are a genuine and robust feature of macro-financial dynamics, a feature that is overlooked with commonly used linear methods or Gaussian models. Indeed, periodic multimodality is a feature that would be overlooked by most non-linear extensions to VARs. One would need, for example, a Markov-switching VAR parametrized to have the probability of switching to the bad mode only sometimes be positive to generate the same features of the conditional joint distribution.

Our paper provides rich empirical evidence against which predictions of macro-financial models can be tested. Our results point to the importance of understanding the dynamics of the entire conditional distribution, rather than of the conditional point forecasts alone.

Thus, instead of evaluating model fit relative to linear impulse response functions, models should be evaluated relative to *density* impulse response functions. The empirical evidence in this paper suggests that nonlinear shock propagation when financial conditions are tight, either due to occasionally binding financial constraints or due to the emergence of multiple equilibria, are an important of macro-financial dynamics.

References

Adrian, T. and N. Boyarchenko (2012): "Intermediary Leverage Cycles and Financial Stability," Staff Report N. 567, Federal Reserve Bank of New York.

Adrian, T., N. Boyarchenko, and D. Giannone (2019): "Vulnerable Growth," *American Economic Review*, 109, 1263–1289.

Adrian, T. and F. Duarte (2016): "Financial Vulnerability and Monetary Policy," Staff Report N. 804, Federal Reserve Bank of New York.

Aikman, D., A. Lehnert, J. N. Liang, and M. Modugno (2016): "Financial Vulnerabilities, Macroeconomic Dynamics, and Monetary Policy," Finance and Economics Discussion Series 2016–055, Board of Governors of the Federal Reserve System (US).

Altissimo, F. and G. L. Violante (2001): "The non–linear dynamics of output and unemployment in the U.S," *Journal of Applied Econometrics*, 16, 461–486.

Arnold, B. C., E. Castillo, and J. M. Sarabia (1999): *Conditional Specification of Statistical Models*, Springer.

Aruoba, S. B., L. Bocola, and F. Schorfheide (2017): "Assessing DSGE model nonlinearities," *Journal of Economic Dynamics and Control*, 83, 34–54.

Ascari, G. and T. Haber (2019): "Sticky Prices and the Transmission Mechanism of Monetary Policy: A Minimal Test of New Keynesian Models," Discussion Paper Series 869, University of Oxford Department of Economics.

Auerbach, A. J. and Y. Gorodnichenko (2013): "Fiscal Multipliers in Recession and Expansion," in *Fiscal Policy after the Financial Crisis*, University of Chicago Press.

Brave, S. and A. Butters (2012): "Diagnosing the Financial System: Financial Conditions and Financial Stress," *International Journal of Central Banking*, 8, 191–239.

Brunnermeier, M. K. and Y. Sannikov (2014): "A Macroeconomic Model with a Financial Sector," *American Economic Review*, 104, 379–421.

Bryant, J. (1983): "A simple rational expectations Keynes–type model," *The Quarterly Journal of Economics*, 98, 525–528.

Caggiano, G., E. Castelnuovo, and J. M. Figueres (2019): "Economic Policy Uncertainty Spillovers in Booms and Busts," *Oxford Bulletin of Economics and Statistics*, 0.

Caldara, D., C. Scotti, and M. Zhong (2019): "Macroeconomic and Financial Risks: A Tale of Volatility," Mimeo, Federal Reserve Board.

Carriero, A., T. Clark, and M. Marcellino (2018): "Capturing Macroeconomic Tail Risks with Bayesian Vector Autoregressions," Mimeo, Queen Mary University of London, Bocconi University, and Federal Reserve Bank of Cleveland.

Chang, Y., Y. Choi, and J. Y. Park (2017): "A new approach to model regime switching," *Journal of Econometrics*, 196, 127–143.

Cogley, T. and T. J. Sargent (2005): "Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S," *Review of Economic Dynamics*, 8, 262–302.

Cogley, T. and A. M. Sbordone (2008): "Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve," *American Economic Review*, 98, 2101–2126.

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

Del Negro, M. and G. E. Primiceri (2015): "Time Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum," *Review of Economic Studies*, 82, 1342–1345.

Diamond, D. W. and P. H. Dybvig (1983): "Bank Runs, Deposit Insurance and Liquidity," *Journal of Political Economy*, 93, 401–419.

Diamond, P. A. (1982): "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy*, 90, 881–894.

Doz, C., D. Giannone, and L. Reichlin (2012): "A Quasi–Maximum Likelihood Approach for Large, Approximate Dynamic Factor Models," *Review of Economics and Statistics*, 94, 1014–1024.

Fernández-Villaverde, J., S. Hurtado, and G. Nuño (2019): "Financial frictions and the wealth distribution," Tech. rep., National Bureau of Economic Research, Working paper N. 26302.

Gallant, A. R., P. E. Rossi, and G. Tauchen (1993): "Nonlinear Dynamic Structures," *Econometrica*, 61, 871–907.

Gertler, M. and N. Kiyotaki (2015): "Banking, liquidity, and bank runs in an infinite horizon economy," *American Economic Review*, 105, 2011–43.

Gertler, M., N. Kiyotaki, and A. Prestipino (2016): "Wholesale banking and bank runs in macroeconomic modeling of financial crises," in *Handbook of Macroeconomics*, Elsevier, vol. 2, 1345–1425.

Granger, C. (2008): "Non-Linear Models: Where Do We Go Next - Time Varying Parameter Models?" *Studies in Nonlinear Dynamics & Econometrics*, 12, 1–11.

Hamilton, J. (1989): "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57, 357–384.

He, Z. and A. Krishnamurthy (2013): "Intermediary Asset Pricing," *American Economic Review*, 103, 732–770.

Hollo, D., M. Kremer, and M. Lo Duca (2012): "CISS-a composite indicator of systemic stress in the financial system," Working Paper Series, No 1426, ECB.

Hubrich, K. and R. J. Tetlow (2015): "Financial stress and economic dynamics: The transmission of crises," *Journal of Monetary Economics*, 70, 100–115.

Kilian, L. and H. Lutkepohl (2018): *Structural Vector Autoregressive Analysis*, no. 9781107196575 in Cambridge Books, Cambridge University Press.

Kilian, L. and M. P. Taylor (2003): "Why is it so difficult to beat the random walk forecast of exchange rates?" *Journal of International Economics*, 60, 85–107.

Koenker, R., S. Leorato, and F. Peracchi (2013): "Distributional vs. Quantile Regression," EIEF Working Papers Series 1329, Einaudi Institute for Economics and Finance (EIEF).

Li, Q. and J. Racine (2007): *Nonparametric Econometrics: Theory and Practice*, Princeton University Press.

Morris, S. and H. S. Shin (2000): "Rethinking Multiple Equilibria in Macroeconomic Modeling," *NBER Macroeconomics Annual*, 15, 139–161.

——— (2002): "Social Value of Public Information," *American Economic Review*, 92, 1521–1534.

Murphy, K. M., A. Shleifer, and R. W. Vishny (1989): "Industrialization and the Big Push," *Journal of Political Economy*, 97, 1003–1026.

Norets, A. and D. Pati (2017): "Adaptive Bayesian estimation of conditional densities," *Econometric Theory*, 33, 980–1012.

Primiceri, G. (2005): "Time Varying Structural Vector Autoregressions and Monetary Policy," *Review of Economic Studies*, 72, 821–852.

Rossi, B. and T. Sekhposyan (2017): "Alternative tests for correct specification of conditional predictive densities," Working Paper, Texas A&M University.

Silverman, B. (1986): *Density Estimation for Statistics and Data Analysis*, no. 9780412246203 in Monographs on Statistics and Applied Probability, Chapman and Hall.

Sims, C. (2000): "Using a likelihood perspective to sharpen econometric discourse: Three examples," *Journal of Econometrics*, 95, 443–462.

Sims, C. A. and T. Zha (2006): "Were There Regime Switches in U.S. Monetary Policy?" *American Economic Review*, 96, 54–81.

Stock, J. and M. Watson (1999): "Forecasting Inflation," *Journal of Monetary Economics*, 44, 293–335.

Terasvirta, T. and H. M. Anderson (1992): "Characterizing Nonlinearities in Business Cycles Using Smooth Transition Autoregressive Models," *Journal of Applied Econometrics*, 7, 119–136.

van Dijk, D., T. Terasvirta, and P. H. Franses (2002): "Smooth Transition Autoregressive Models – A Survey Of Recent Developments," *Econometric Reviews*, 21, 1–47.

Online Appendix

A Results using alternative measures

A.1 CFNAI instead of real GDP growth

In this section, we present results using he Chicago Fed National Activity Index (CFNAI) as a measure of economic conditions, instead of real GDP growth.[8] The CFNAI is a monthly index designed to measure overall economic activity. It is a weighted average of 85 existing monthly indicators of national economic activity, normalized to have an average value of zero and a standard deviation of one, with positive realizations corresponding to growth above trend.[9] The economic indicators that are included in the CFNAI are drawn from four broad categories of data: production and income; employment, unemployment, and hours; personal consumption and housing; and sales, orders, and inventories. The derived index provides a single, summary measure of a factor common to these national economic data. The methodology for the CFNAI is based on the principal component estimator for large dynamic factor models developed by Stock and Watson (1999). Figure A.1 plots the time series of CFNAI together with the time series of NFCI.

Figure A.2 shows the marginal and joint distribution of CFNAI and NFCI one quarter ahead. The marginal distribution of economic conditions has more volatile lower quantiles than upper quantiles, reflecting a negative correlation between the conditional mean and the conditional variance. The marginal distribution of financial conditions is more symmetric. As with the results in the main body of the paper, the joint distribution is roughly Gaussian during normal periods but, during periods of tight financial conditions, the shape of the joint distribution changes radically and additional modes of the distribution emerge. Figure A.3 shows that this multilmodality during periods of tighter financial conditions survives up to four quarters ahead. Finally, Figure A.4 shows that the features of the Great Recession – the emergence of a secondary negative mode in Q3-Q4 2008 and a slow resolution to the good mode – that we documented in the main body of the paper remain even when we measure economic conditions using CFNAI instead of real GDP growth.

A.2 CISS instead of NFCI

In this section, we present results using Composite Indicator of Systemic Stress (CISS) as a measure of financial conditions, instead of NFCI.[10] Figure A.5 plots the time series of the U. S. CISS together with real GDP growth.

Figure A.6 shows the marginal and joint distribution one quarter ahead of real GDP growth and CISS one quarter ahead. The marginal distribution of economic conditions has more volatile lower quantiles than upper quantiles, reflecting a negative correlation between the conditional mean and the conditional variance. The marginal distribution of financial conditions is more symmetric. As with the results in the main body of the paper, the joint distribution is roughly Gaussian during normal periods but, during periods of tight financial conditions, the shape of the joint distribution changes radically and additional modes of the distribution emerge. Figure A.7 shows that this multilmodality during periods of tighter financial conditions survives up to four quarters ahead. Finally, Figure A.8 shows that the features of the Great Recession – the emergence of a secondary negative mode in Q3-Q4 2008 and a slow resolution to the good mode – that we documented in the main body of the paper remain even when we measure financial conditions using U. S. CISS instead of NFCI.

[8] The CFNAI is computed by the Federal Reserve Bank of Chicago, available here.

[9] The list of 85 indicators is provided here.

[10] The CISS is computed by the European Central Bank. CISS combines 15 raw, mainly market-based financial stress measures that are split equally into five categories, namely the financial intermediaries sector, money markets, equity markets, bond markets and foreign exchange markets into a single index, following the methodology in Hollo, Kremer, and Lo Duca (2012).

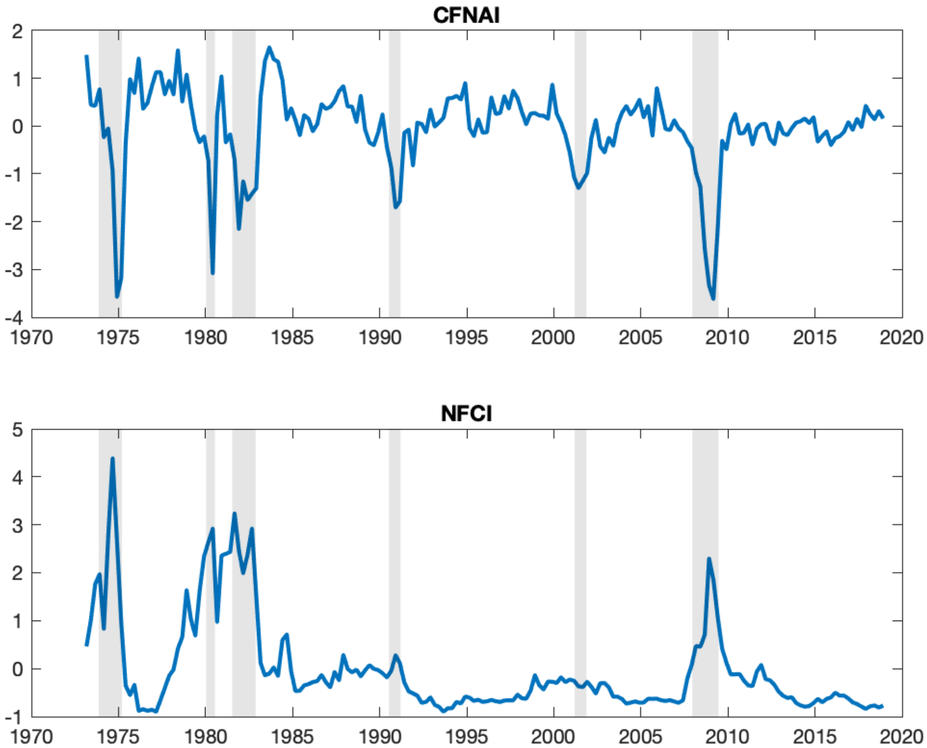

Figure A.1. Economic and financial conditions. The figure shows the time series of the economic conditions index CFNAI and the financial conditions index NFCI, together with NBER recession shadings.

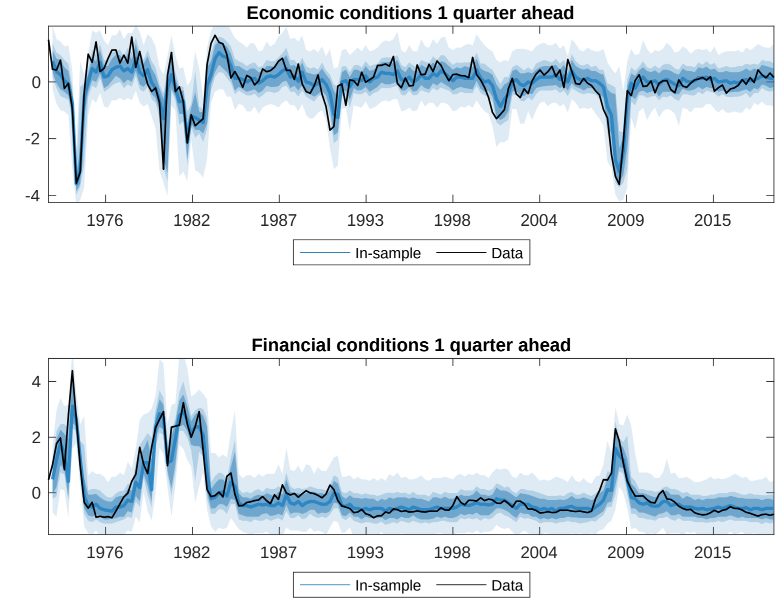

(a) Marginal: One quarter ahead

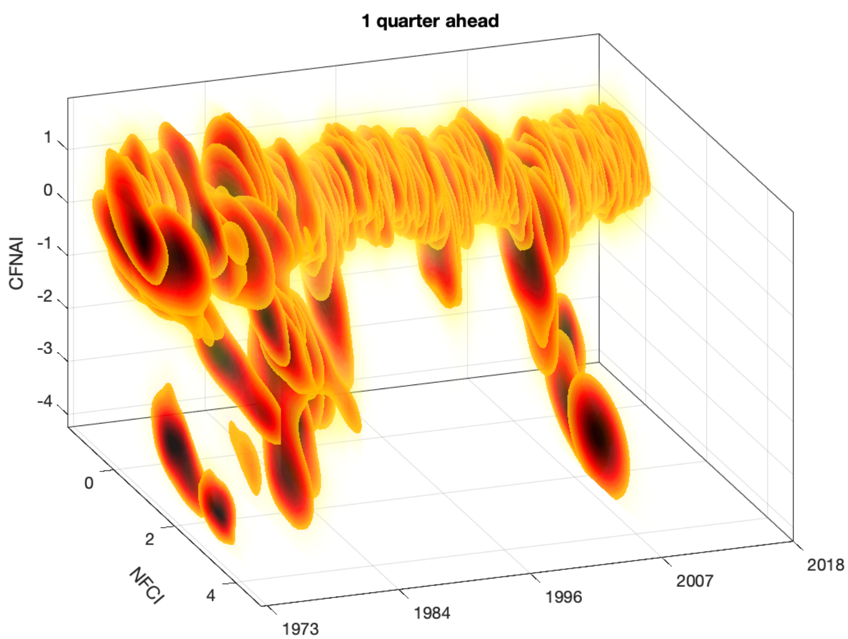

(b) Joint: One quarter ahead

Figure A.2. Density forecasts with CFNAI. The figure shows the marginal (top) and joint distributions (bottom) for one quarter ahead forecasts. In Figure A.2a, the blue shades give the 68%, and 95% quantile bands, while the solid blue line gives the median. Figure A.2b presents the contour plots of the joint distribution in every quarter in the sample, with darker shades of red correspond to higher probability densities.

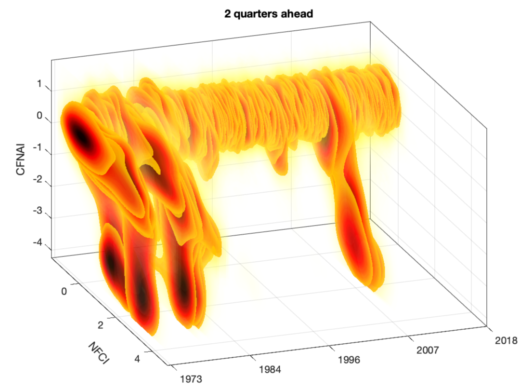

(a) Two quarters ahead

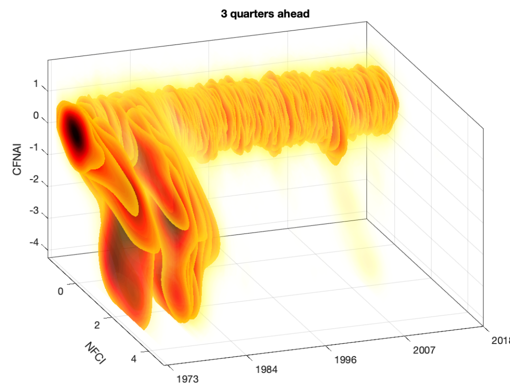

(b) Three quarters ahead

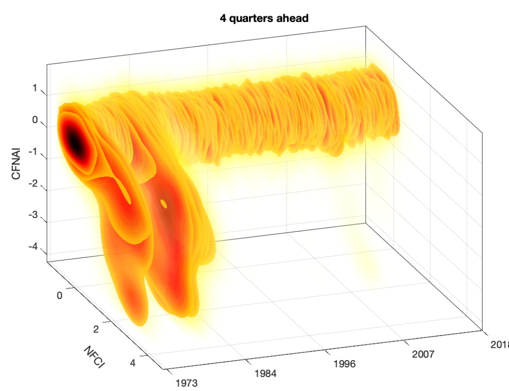

(c) Four quarters ahead

Figure A.3. Joint distributions of CFNAI and NFCI across horizons. The figure shows the joint distribution of CFNAI and NFCI for two, three and four quarter ahead forecasts, as contour plots of the joint distribution in every quarter in the sample, with darker shades of red correspond to higher probability densities.

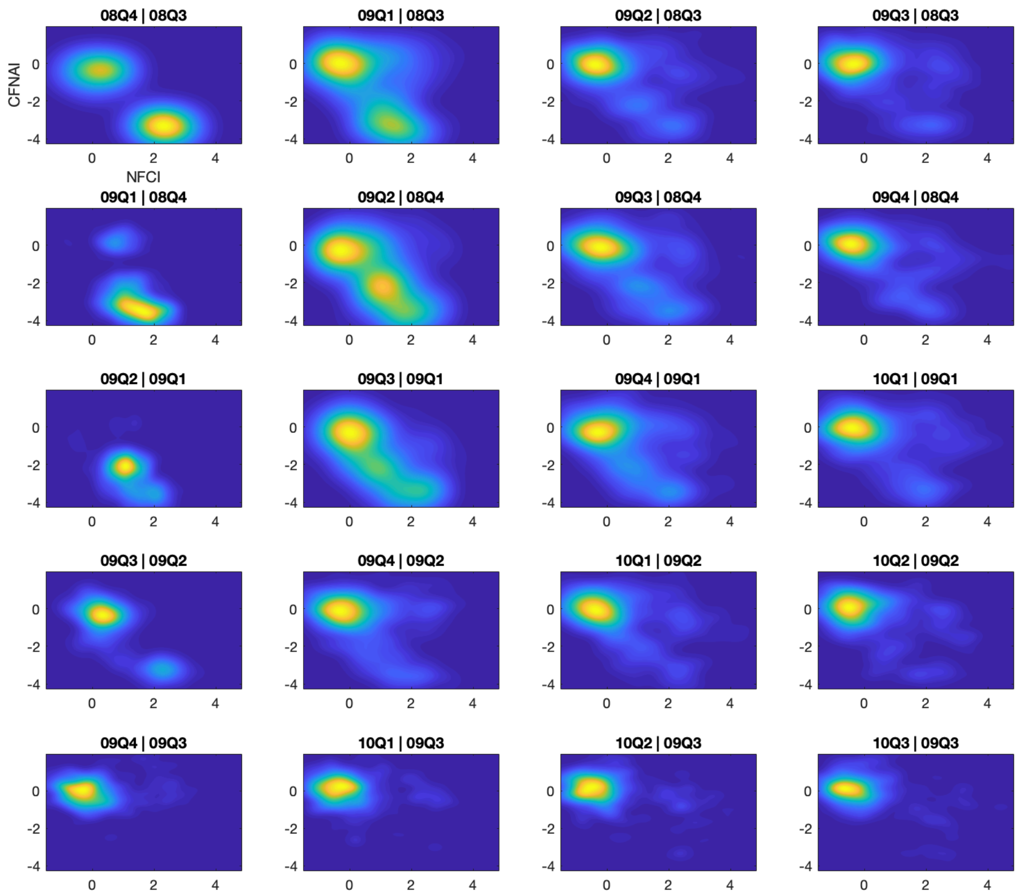

Figure A.4. Joint distribution of economic and financial conditions during the Great Recession. Contour plots of 1–4 quarter-ahead density forecasts of CFNAI and NFCI during the financial crisis. Brighter colors indicate great probability.

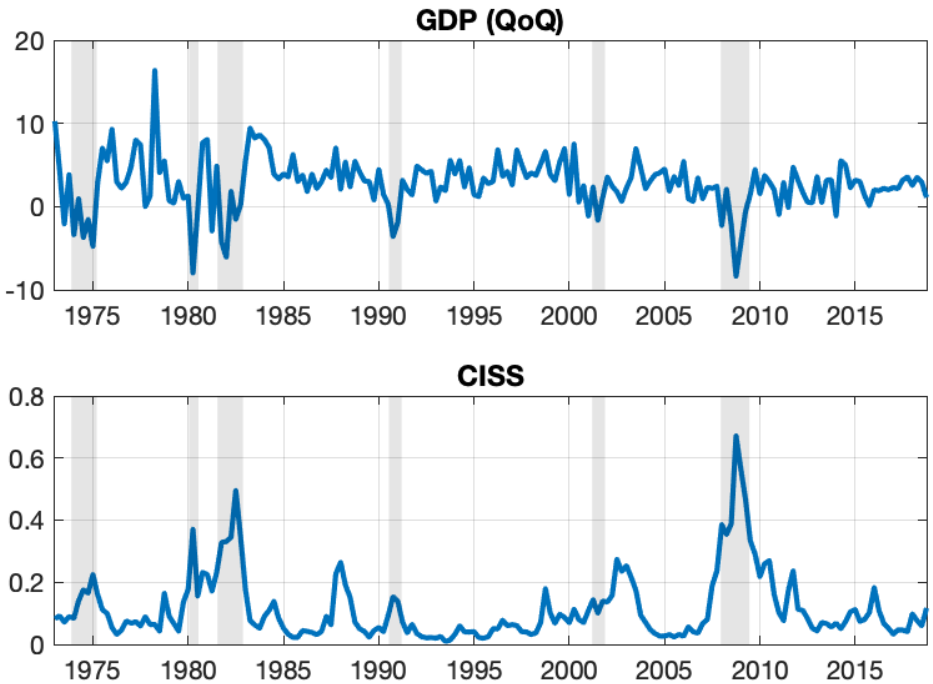

Figure A.5. Economic and financial conditions. The figure shows the time series of quarter-over-quarter real GDP growth and the financial conditions index CISS, together with NBER recession shadings.

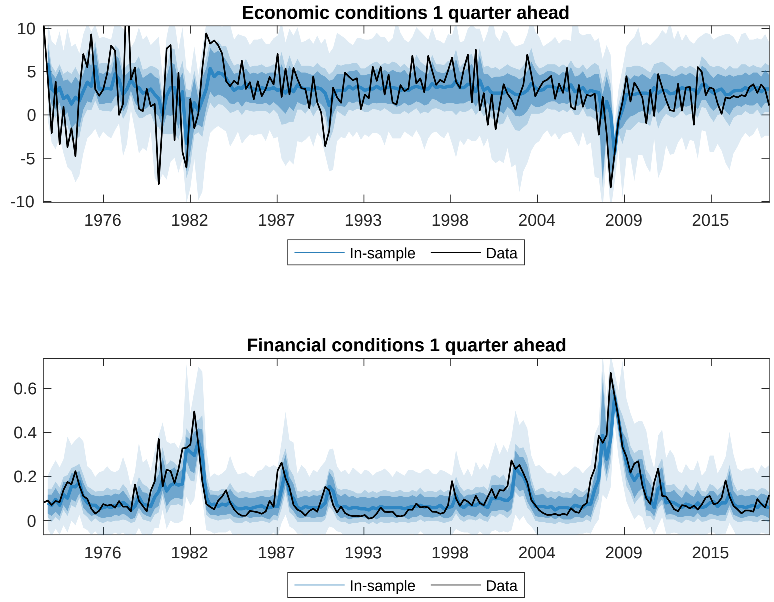

(a) Marginal: One quarter ahead

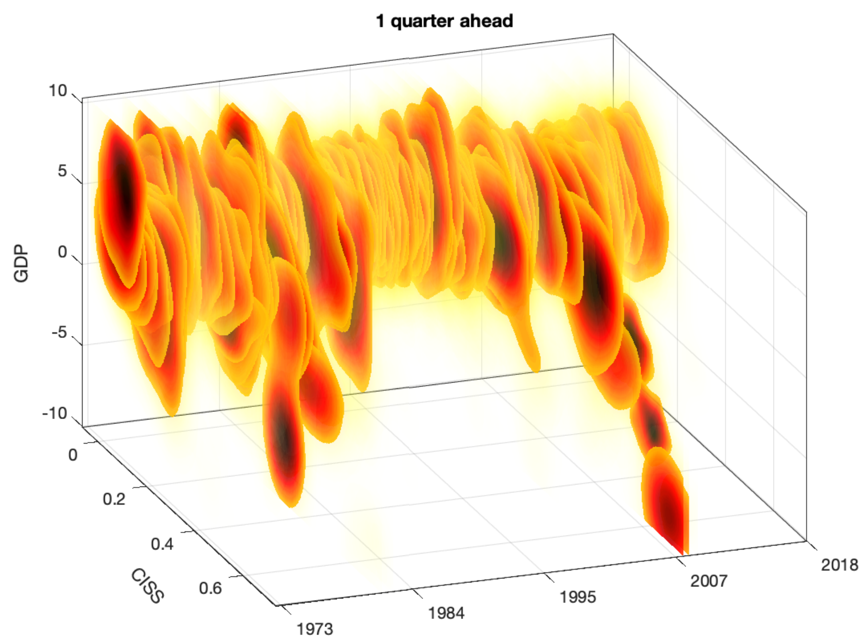

(b) Joint: One quarter ahead

Figure A.6. Density forecasts with CISS. The figure shows the marginal (top) and joint distributions (bottom) for one quarter ahead forecasts. In Figure A.6a, the blue shades give the 68%, and 95% quantile bands, while the solid blue line gives the median. Figure A.6b presents the contour plots of the joint distribution in every quarter in the sample, with darker shades of red correspond to higher probability densities.

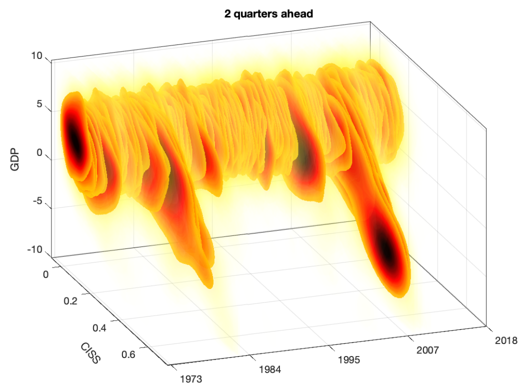

(a) Two quarters ahead

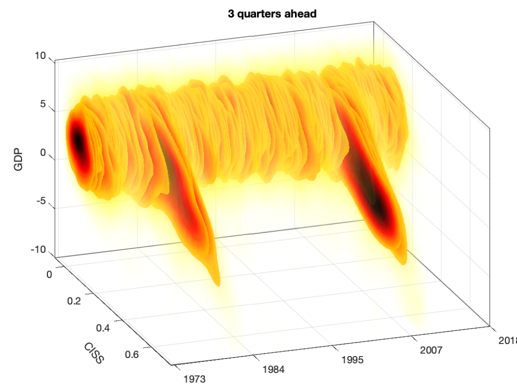

(b) Three quarters ahead

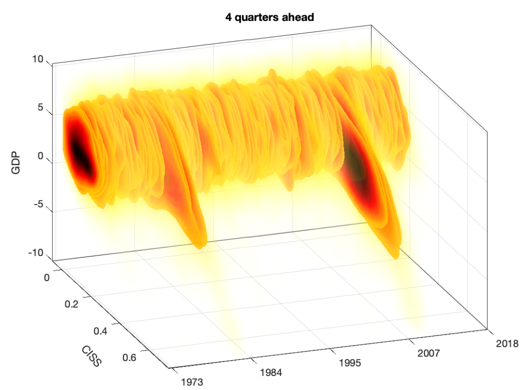

(c) Four quarters ahead

Figure A.7. Joint distributions of real GDP growth and CISS across horizons. The figure shows the joint distribution of real GDP growth and CISS for two, three and four quarter ahead forecasts, as contour plots of the joint distribution in every quarter in the sample, with darker shades of red correspond to higher probability densities.

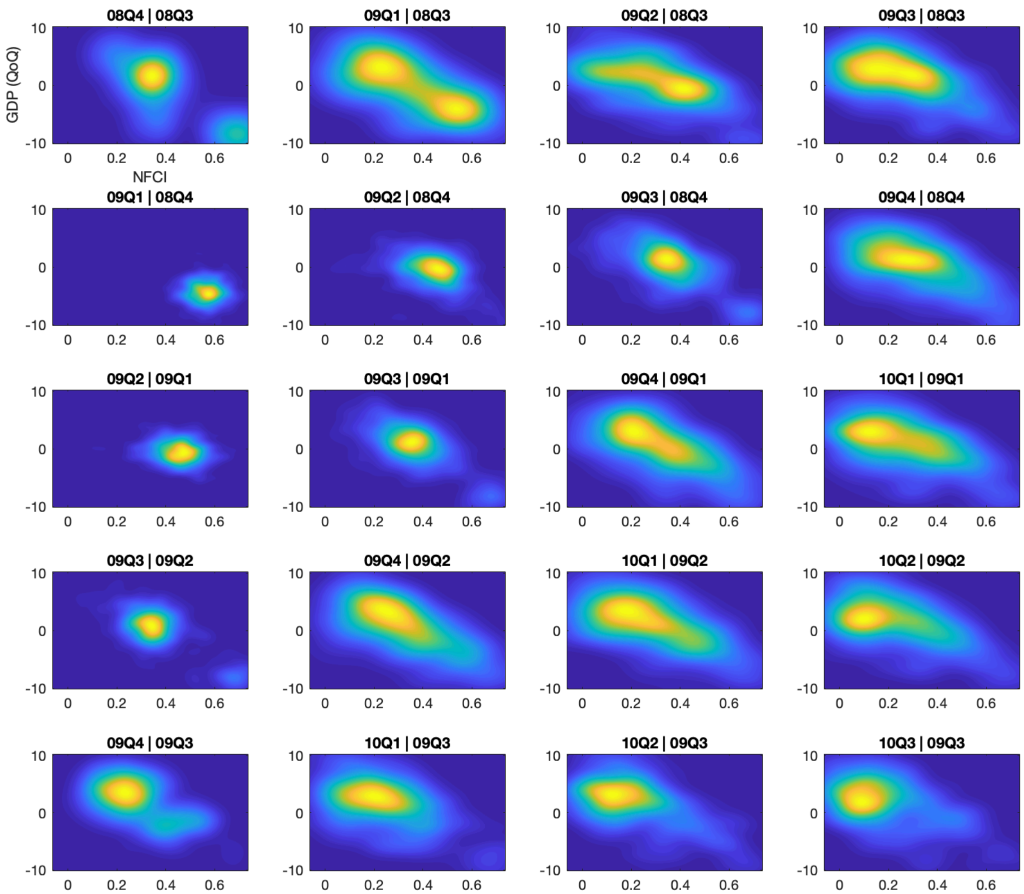

Figure A.8. Joint distribution of economic and financial conditions during the Great Recession. Contour plots of 1–4 quarter-ahead density forecasts of real GDP growth and CISS during the financial crisis. Brighter colors indicate great probability.

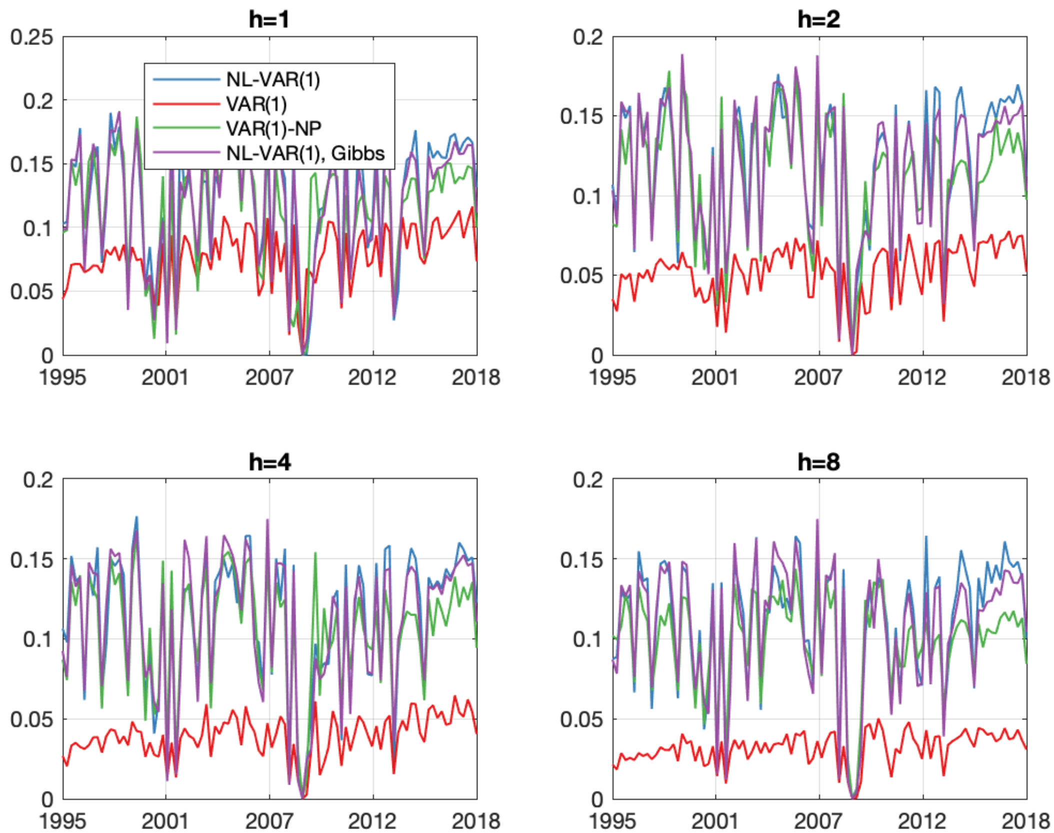

Figure A.9. Forecast sharpness of the Gibbs sampler. The figure reports scores for the joint predictive distribution of real GDP growth and NFCI at the one-quarter-, two-quarter-, four-quarter-, and eight-quarter-ahead horizon. Log-predictive-scores for the fully non-parametric model plotted in blue ("NL-VAR(1)"); for a linear VAR with Gaussian innovations in red ("VAR(1)"); for a linear VAR with non-parametric errors in green ("VAR(1)-NL"); for the fully non-parametric model estimated via a Gibbs' sample in purple ("NL-VAR(1), Gibbs").

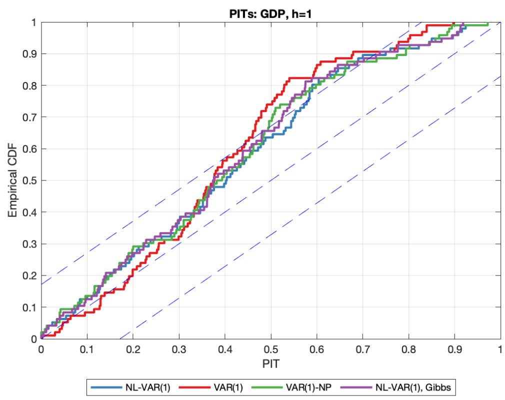

(a) Real GDP growth, one quarter ahead

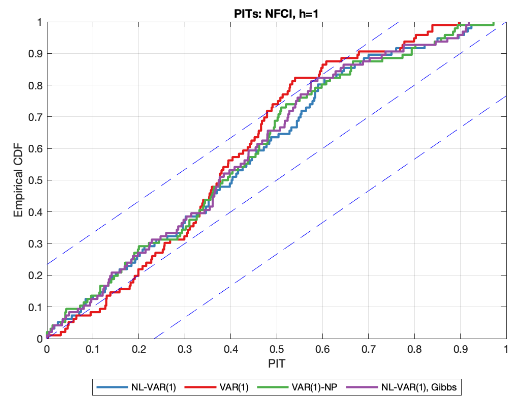

(b) NFCI, one quarter ahead

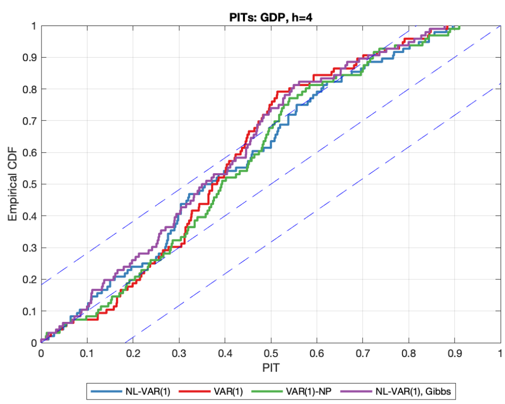

(c) Real GDP growth, one quarter ahead

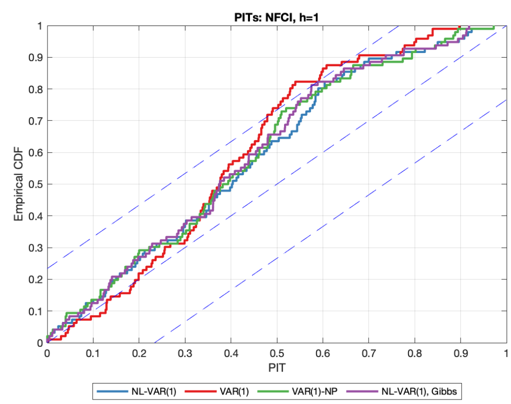

(d) NFCI, four quarters ahead

Figure A.10. Calibration of the forecast. The figure reports the empirical cumulative distribution of the marginal probability integral transform (PITs) from real GDP growth and NFCI one quarter and one year ahead. PITs for the fully non-parametric model plotted in blue ("NL-VAR(1)"); for a linear VAR with Gaussian innovations in red ("VAR(1)"); for a linear VAR with non-parametric errors in green ("VAR(1)-NL"); for the fully non-parametric model estimated via a Gibbs' sample in purple ("NL-VAR(1), Gibbs"); the 5% critical values as in Rossi and Sekhposyan (2017) plotted as dashed-blue lines.