



The
University
Of
Sheffield.

Department
Of
Economics

Tackling Large Outliers in Macroeconomic Data with Vector Artificial Neural Network Autoregression

Vito Polito, Yunyi Zhang

Sheffield Economic Research Paper Series

SERPS no. 2022004

ISSN 1749-8368

21 March 2022

Tackling Large Outliers in Macroeconomic Data with Vector Artificial Neural Network Autoregression

Vito Polito[‡] and Yunyi Zhang[‡]

March 21, 2022

Abstract

We develop a regime switching vector autoregression where artificial neural networks drive time variation in the coefficients of the conditional mean of the endogenous variables and the variance covariance matrix of the disturbances. The model is equipped with a stability constraint to ensure non-explosive dynamics. As such, it is employable to account for nonlinearity in macroeconomic dynamics not only during typical business cycles but also in a wide range of extreme events, like deep recessions and strong expansions. The methodology is put to the test using aggregate data for the United States that include the abnormal realizations during the recent Covid-19 pandemic. The model delivers plausible and stable structural inference, and accurate out-of-sample forecasts. This performance compares favourably against a number of alternative methodologies recently proposed to deal with large outliers in macroeconomic data caused by the pandemic.

Keywords: Nonlinear time series, Regime switching models, Extreme events, Covid-19, Macroeconomic forecasting.

JEL: C45, C5, E37

*The authors gratefully acknowledge comments received by participants at the 11th RCEA Money-Macro-Finance Conference; feedback and suggestions from Danilo Cascaldi-Garcia, Ana Galvao, Benny Hartwig, Gary Koop, Juan Paez-Farrell, Ivan Petrella, Jeremy Smith, Christoph Thoenissen and Shaun Vahey.

[†]University of Sheffield; CESifo; email: v.polito@sheffield.ac.uk.

[‡]Xiamen University; email: yunyi.zhang@xmu.edu.cn.

1 Introduction

Inference and dynamic analysis with the most popular model of time series macroeconomics, the vector autoregression (VAR), have become incredibly challenging since March 2020. This is because macroeconomic and financial variables have begun to record values that are several standard deviations away from their historical averages due to the effect of the Covid-19 pandemic. Unrestricted VARs interpret these large observations as permanent. Consequently the estimated transmission mechanism of the economy becomes highly unstable, if not explosive, preventing VARs from accomplishing two of the tasks for which they are most used, structural inference and forecasting.

The overarching view among economists dealing with this issue is that these extreme realizations can only be the outcome of a temporary jump in volatility, without long lasting structural effects on aggregate time series, see Schorfheide and Song (2020), Primiceri and Tambalotti (2020), Lenza and Primiceri (2020), Carriero et al. (2021). As a result, two main approaches have emerged to tackle this issue in unrestricted VARs. The first ignores the extreme observations and estimates the VAR on data pertinent to the pre-pandemic, on the ground that the recent observations are outliers generated by large shocks and should therefore be excluded from the estimation. The second retains the extreme observations within the estimation sample but employs an ad-hoc modification of the variance covariance matrix of the disturbances with the effect of downweighting data from the pandemic period.

In this paper we describe a data-driven methodology that allows an unrestricted VAR to deliver sensible (non-explosive) macroeconomic responses to shocks and forecasts in the presence of large swings in the data, including those caused by the pandemic. This is based on a regime switching VAR where artificial neural networks (ANNs) connect the model coefficients to the dynamics of an observable variable that captures a broad range of possible states of the economy, such as for example those observed over a typical business cycle, or during extreme events that cause deep recessions or strong expansions. We term this model vector artificial neural network autoregression (VANNAR).

ANNs have long been used to detect nonlinearity in time series models. The VANNAR differs from existing ANN models in two important dimensions. First, it accounts for nonlinearity in both the mean of the endogenous variables and the variance covariance

matrix of the disturbances. Second, it ensures stability via a stationarity constraint on the coefficients describing the linear dynamic of the endogenous variables. The first feature allows the model to tackle both structural changes in the economy and heteroscedastic disturbances, the second to retain stability - and therefore deliver meaningful structural inference and forecasts - even in the presence of outliers in the data, including abnormal realisations such as those caused by the pandemic.

The proposed VANNAR is in the spirit of the current literature dealing with the estimation of time series models at the time of Covid-19, in that it conditions variation in the model coefficients to the dynamic evolution of the state of the economy and it ensures stability even when an unrestricted VAR would deliver explosive dynamics. However, it does so without relying on a specific view about the effects of extreme events on the first and second moments of the relevant time series, or needing to commit to a particular modification of the variance covariance matrix of the VAR disturbances. For this reason, the proposed VANNAR provides a very flexible framework, suitable for both ex-ante and ex-post quantitative analysis. Further, by conditioning volatility on observed data dynamics, the VANNAR implicitly adjusts for both permanent and temporary changes in the data, as well as those caused by abnormal realizations.

The idea of ensuring stability in a VAR using a stationarity constraint is not new. It is routinely used in the estimation of VAR models with time varying coefficients driven by stochastic volatility, pioneered by Cogley and Sargent (2001) and Primiceri (2005). In these models, the constraint is applied to select from the posterior draws of the VAR coefficients only those that satisfy stability. There is however no guarantee that a set of stable draws can be found, particularly if the VAR is confronted with extreme observations. The way we implement the stability constraint in the VANNAR is different, as we explicitly include it in the maximization of the log-likelihood function. Thus if the model is estimated with frequentist methods, such as maximum likelihood, parameter estimates will always deliver stable dynamics. If it is estimated with Bayesian methods, the posterior estimates will also retain stability, being based on the constrained log-likelihood.

To test the proposed methodology we employ the macroeconomic data for the United States used in a recent study on VAR estimation after March 2020 by Lenza and Primiceri

(2020). This includes monthly observations of seven indicators covering the labour market, the real and the nominal activity. The model is estimated with a Bayesian method over three different samples. The first terminates in February 2020, thus excluding the large outliers caused by the pandemic in the United States, which start from March 2020. The other two samples terminate in June and September 2020, respectively. We find that the VANNAR specification with only one ANN in the variance covariance matrix gives the best statistical fit in each of these three samples. This data-driven result is of course consistent with the predominant conjecture in the existing literature that the instability of unrestricted VARs estimated on the latest data is due to a temporary jump in volatility rather than a change in the structure of the economy.

We undertake two types of empirical analyses. First, as in Lenza and Primiceri (2020), we compute (i) the responses of the macroeconomy to an unanticipated increase in the rate of unemployment identified from the Cholesky decomposition of the variance covariance matrix and (ii) out-of-sample forecasts conditional on external projections (and actual data depending on the estimation sample) of the rate of unemployment. The results from this analysis are summarized as follows. The VANNAR consistently delivers plausible structural inference, since both the real and nominal activity deteriorate in response to an unanticipated increase in the rate of unemployment regardless of the estimation sample (**R1**). Macroeconomic responses remain stable during the pandemic (**R2**). The magnitude of the macroeconomic responses and their uncertainty display a temporary increase during the pandemic periods (**R3**). Out-of-sample forecasts capture the disruption on the economy caused by the pandemic, as well as the slow recovery from June 2020 onward (**R4**). The forecast accuracy is higher once the model is tasked to predict the recovery phase of the economy rather than the sharp disruption in the early months of the pandemic (**R5**).

We then compare the results from the estimated VANNAR against those obtained from four alternative approaches recently proposed to tackle the abnormal realizations caused by the pandemic when using an unrestricted VAR. All these employ a linear VAR, but differ in the specification of the variance covariance matrix. The first uses constant volatility but estimates the VAR with data until February 2020, following Schorfheide and Song (2020) and Primiceri and Tambalotti (2020). The second, proposed by Lenza and Prim-

iceri (2020), uses information from observed data post-pandemic to modify the variance covariance matrix such that volatility increases in the first three months of the pandemic and gradually reduces from the fourth month onward. The third and fourth alternatives follow the proposal of Carriero et al. (2021) of using stochastic volatility adjusted to include either fat-tailed errors, as in Jacquier et al. (2004), or the outlier-augmented specification of Stock and Watson (2016). The results from this comparative analysis are summarized as follows. Structural inference undertaken with any alternative model is generally consistent with that from the VANNAR, since the economy’s response to an unanticipated increase in the rate of unemployment is similar across the four alternatives to that inferred from the VANNAR, displaying a temporary increase in magnitude and uncertainty when estimated on the latest data (**R6**). All alternative models yield out-of-sample forecasts qualitatively similar to those from the VANNAR (**R7**). However, none of them can produce more accurate forecasts relatively to the VANNAR, regardless of the estimation sample and for the majority of forecast horizons considered (**R8**).

Related Literature. The paper contributes methodologically and empirically to the rapidly growing literature on the estimation of unrestricted VARs in the presence of large outliers caused by the Covid-19 pandemic, typified by the works of Schorfheide and Song (2020), Primiceri and Tambalotti (2020), Lenza and Primiceri (2020), Carriero et al. (2021). The methodological contribution consists of describing a flexible data-driven approach, which does not rely on ad-hoc adjustment of either the data sample (to exclude the outliers) or the variance covariance matrix (to downweight the outliers). As such, it is employable to account for a wide range of nonlinear macroeconomic dynamics, such as those of a typical business cycle, extreme events, or abnormal outliers. The empirical contribution consists of providing a synthesis and comparison of structural inference and forecast analysis from (the VANNAR and) the main approaches so far proposed in this literature.

The paper is related to the large recent literature dealing with the problem of forecasting the aggregate macroeconomic effects of the Covid-19 shock. Broadly speaking, existing approaches use either structural (epidemiological and) economic models applied to recent data, see for Baqaee et al. (2020), Baqaee and Farhi (2020), Eichenbaum et al. (2021) and Guerrieri et al. (2020), or time series models to infer the effects of the Covid-19 pandemic

on the economy from past history, see Barro (2020), Ludvigson et al. (2020) and Jordà et al. (2020). The paper provides a time series approach to infer the effects of the Covid-19 pandemic on the economy without discarding information from the latest data.

The VANNAR belongs to the family of regime switching VARs widely used to model nonlinearity in macroeconomic data, namely the vector threshold, smooth transtion and Markov-switching autoregression - VTAR, VSTAR and VMSAR - originally developed by Granger and Terasvirta (1993) and Hamilton (1994). In these models, the economy is assumed to either jump (VTAR and VMSAR) or move gradually (VSTAR) between a finite number (often two) of predefined binding linear regimes. The switch is deterministic in the VTAR and VSTAR, stochastic in the VMSAR. The VANNAR is similar to the VTAR and VSTAR in that it conditions regime changes on observable variables. The VANNAR however does not require the a-priory specification of binding regimes, since these are implicitly set by the dynamics of the same variables that determine transition across states. Thus, the proposed methodology has two main advantages compared to existing regime switching VARs. It is more parsimonious and the determination of the regimes is data driven.

ANNs have long been used to model nonlinearity in macroeconomic and financial time series. Most of the studies in this literature focus only on nonlinearity in the mean assuming volatility constant, see for example Swanson and White (1995, 1997), Stock and Watson (1996, 1998), Teräsvirta et al. (2005) and Bredahl Kock and Teräsvirta (2016). The VANNAR extends these models as it employs ANNs to account also for time variation in the variance covariance matrix of the disturbances. ANN models with heteroscedastic errors are however not new. Medeiros et al. (2008) develop a multivariate ANN model with GARCH volatility, while McAleer et al. (2008) consider a time series model with time variation in the variance covariance matrix driven by ANNs in the main diagonal. The VANNAR differs from these two specifications because it allows time variation in the entire variance covariance matrix of the disturbances to be driven by the ANN. Further, the stability constraint included in the VANNAR ensures the model's ability to deliver sensible structural inference and forecasts when dealing with any type of nonlinearity in time series analysis, including that in the latest macroeconomic data.

The paper proceeds as follows. Section 2 describes the proposed VANNAR methodology, covering issues surrounding the model specification, the definition and identification of the log-likelihood, the estimation of the parameters, structural inference and forecasting. Section 3 presents the empirical results from the estimated VANNAR. Section 4 compares these with the alternative approaches. Section 5 concludes with a summary. Supplementary material is provided in a separate Appendix.

2 Model

Specification. Let \mathbf{y}_t be a vector of endogenous variables and \mathbf{z}_t a vector of explanatory variables, either exogenous or related to some of the variables in \mathbf{y}_t . The VANNAR is specified as follows:

$$\mathbf{y}_t = \mathbf{\Lambda}\mathbf{x}_t + \mathbf{\Gamma}\mathbf{w}(\boldsymbol{\alpha}_y, \boldsymbol{\beta}_y; \mathbf{z}_t) + \mathbf{u}_t, \quad (1)$$

$$\mathbf{u}_t \sim i.i.d. (\mathbf{0}, \boldsymbol{\Sigma}_t), \quad (2)$$

$$\boldsymbol{\Sigma}_t = \boldsymbol{\Sigma}_0 + \mathbf{\Omega}\mathbf{w}(\boldsymbol{\alpha}_\Sigma, \boldsymbol{\beta}_\Sigma; \mathbf{z}_t), \quad (3)$$

$$E(\mathbf{z}_t) = \mathbf{0}, \text{ var}(\mathbf{z}_t) = \mathbf{1}. \quad (4)$$

The conditional mean of \mathbf{y}_t on the right side of equation (1) includes three terms: the linear component $\mathbf{\Lambda}\mathbf{x}_t$, where $\mathbf{\Lambda} = [\mathbf{\Lambda}_1, \dots, \mathbf{\Lambda}_p, \boldsymbol{\lambda}]$ is a matrix of coefficients and $\mathbf{x}_t = [\mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-p}, \mathbf{1}']'$; the nonlinear component $\mathbf{\Gamma}\mathbf{w}(\boldsymbol{\alpha}_y, \boldsymbol{\beta}_y; \mathbf{z}_t)$, where $\mathbf{\Gamma} = [\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_q]$ is a matrix of ANN output weights and $\mathbf{w}(\boldsymbol{\alpha}_y, \boldsymbol{\beta}_y; \mathbf{z}_t) = [g(\boldsymbol{\alpha}'_{y,1}\mathbf{z}_{t-1} - \beta_{y,1}), \dots, g(\boldsymbol{\alpha}'_{y,q}\mathbf{z}_{t-q} - \beta_{y,q})]'$ a vector of ANNs driven by lagged values of \mathbf{z}_t ; the stochastic term \mathbf{u}_t , a vector of independent and identically distributed reduced form disturbances with zero mean and time varying variance covariance matrix $\boldsymbol{\Sigma}_t$, as per (2). The variance covariance matrix in equation (3) includes the constant term $\boldsymbol{\Sigma}_0$ and the time varying component $\mathbf{\Omega}\mathbf{w}(\boldsymbol{\alpha}_\Sigma, \boldsymbol{\beta}_\Sigma; \mathbf{z}_t)$, where $\mathbf{\Omega} = [\boldsymbol{\Sigma}_1, \dots, \boldsymbol{\Sigma}_s]$ is a matrix of output weights and $\mathbf{w}(\boldsymbol{\alpha}_\Sigma, \boldsymbol{\beta}_\Sigma; \mathbf{z}_t) = [g(\boldsymbol{\alpha}'_{\Sigma,1}\mathbf{z}_{t-1} - \beta_{\Sigma,1}), \dots, g(\boldsymbol{\alpha}'_{\Sigma,s}\mathbf{z}_{t-s} - \beta_{\Sigma,s})]'$ is a vector of ANNs. \mathbf{z}_t is standardized to have zero mean and unit variance, as per (4).

Each ANN $g(\boldsymbol{\alpha}'_{i,k}\mathbf{z}_{t-k} - \beta_{i,k})$ in the conditional mean, $i = \mathbf{y}$ and $k = 1, \dots, q$, and in the variance covariance matrix, $i = \boldsymbol{\Sigma}$ and $k = 1, \dots, s$, takes the form of an activation

function that squashes into single coefficients hidden layers formed from lagged values of the explanatory vector \mathbf{z}_t . In the empirical analysis we employ two popular activation functions in modern neural networks, the logistic and the rectified linear unit (ReLU). Following the ANN terminology, the polynomials $\boldsymbol{\alpha}'_{i,k}\mathbf{z}_{t-k} - \beta_{i,k}$ are referred to as hidden layers, the coefficients $\boldsymbol{\alpha}_{i,k}$ as input weights and $\beta_{i,k}$ as biases. Each hidden layer in the conditional mean provides a linkage between \mathbf{y}_t and a specific lagged value of \mathbf{z}_t that is nested/hidden within the ANNs $g(\boldsymbol{\alpha}'_{y,k}\mathbf{z}_{t-k} - \beta_{y,k})$, $k = 1, \dots, q$. Similarly, each hidden layer in equation (3), provides a linkage between the current period variance covariance matrix of the disturbances and a given lagged value of \mathbf{z}_t that is nested/hidden within the ANNs $g(\boldsymbol{\alpha}'_{\Sigma,k}\mathbf{z}_{t-k} - \beta_{\Sigma,k})$, $k = 1, \dots, s$. The impact of these nonlinear terms on the conditional mean of the endogenous variables and on the variance covariance matrix of the disturbances is measured by the output weight matrices $\boldsymbol{\Gamma}$ and $\boldsymbol{\Omega}$, respectively.

The VANNAR in (1)-(4) describes a flexible framework that encompasses two popular approaches in time series analysis. Setting either \mathbf{z}_t or all $g(\cdot)$ equal to zero yields the linear VAR model with homoscedastic disturbances. Setting instead all $g(\boldsymbol{\alpha}'_{\Sigma,k}\mathbf{z}_{t-k} - \beta_{\Sigma,k}) = 0$ while retaining $g(\boldsymbol{\alpha}'_{i,y}\mathbf{z}_{t-k} - \beta_{y,k}) \neq 0$ yields the typical multivariate ANN model with nonlinear mean and constant variance used by Swanson and White (1995, 1997), Stock and Watson (1996, 1998), Teräsvirta et al. (2005) and Bredahl Kock and Teräsvirta (2016). There is however an important difference between the ANN models used in these studies and the specified VANNAR. In these models, each ANN depends on the same vector of explanatory variables. Consequently, gradient algorithms used to compute maximum likelihood estimates often breakdown as a result of multicollinearity causing the information matrix to be singular.¹ The proposed specification of the VANNAR mitigates this problem because it conditions each ANN on a different lagged value of the explanatory variables.²

The VANNAR in (1)-(4) is specified in its simplest form as a single layer model. This can be easily extended to include additional layers in both the conditional mean and variance covariance matrix, for example by repeatedly replacing the vector of explanatory variables \mathbf{z}_{t-k} in any given ANN with additional ANNs that also depend on \mathbf{z}_{t-k} . The proposed

¹If the ANNs were to depend on the same vector of explanatory variables, issues of multicollinearity would be compounded in the VANNAR, given the inclusion of ANNs in the variance covariance matrix.

²We evaluate the significance of this different conditioning of the ANNs in Section 3.

specification has the main advantage of being parsimonious and avoiding the issues of multicollinearity highlighted above.

The Constrained Log-Likelihood. When \mathbf{z}_t is observable, the likelihood function of the VANNAR model in (1)-(4) can be formulated analytically and its parameters estimated with either frequentist or Bayesian methods. There are, however, three main issues with the evaluation of the likelihood function. The first two are typical of the econometric analysis of ANN models, concerning the unidentifiability of the log-likelihood and the indeterminacy of the ANN structure required to model a time series. The third issue, related to the presence of large outliers in the data, concerns the stability of the VANNAR dynamics returned by the likelihood-based estimation. Before outlining the estimation algorithm, we state the log-likelihood function and tackle each of these three issues separately.

The log-likelihood. Given a sample of $t = 1, \dots, T$ observations and the assumption on the disturbances in (2), the conditional log-likelihood function of the VANNAR in (1)-(4) can be formed as:

$$\ln L_T(\boldsymbol{\Theta}) = \frac{1}{T - p^*} \sum_{t=p^*+1}^T \left(-\frac{n}{2} \ln 2\pi - \frac{1}{2} \ln |\boldsymbol{\Sigma}_t| - \frac{1}{2} \mathbf{u}_t' \boldsymbol{\Sigma}_t^{-1} \mathbf{u}_t \right), \quad (5)$$

where $\boldsymbol{\Theta}$ is a vector collecting all VANNAR coefficients, $p^* = \max(p, q, s)$ and n denotes the dimension of \mathbf{y}_t . To ensure the positive definiteness on the variance covariance matrix (3) and its determinant, each matrix of coefficients $\boldsymbol{\Sigma}_k$ is decomposed as

$$\boldsymbol{\Sigma}_k = \mathbf{T}_k \mathbf{T}_k', \quad (6)$$

where \mathbf{T}_k is a lower triangular matrix, $k = 0, \dots, s$.

For the subsequent analysis, it is convenient to employ the partition $\boldsymbol{\Theta} = [\boldsymbol{\Theta}_1, \boldsymbol{\Theta}_2]$, where $\boldsymbol{\Theta}_1 = \boldsymbol{\Lambda}$ includes the linear coefficients of the conditional mean in (1) whereas $\boldsymbol{\Theta}_2 = [\boldsymbol{\Gamma}, \mathbf{T}_0, \dots, \mathbf{T}_s, \boldsymbol{\alpha}_{\mathbf{y},1}, \dots, \boldsymbol{\alpha}_{\mathbf{y},q}, \beta_{\mathbf{y},1}, \dots, \beta_{\mathbf{y},q}, \boldsymbol{\alpha}_{\boldsymbol{\Sigma},1}, \dots, \boldsymbol{\alpha}_{\boldsymbol{\Sigma},s}, \beta_{\boldsymbol{\Sigma},1}, \dots, \beta_{\boldsymbol{\Sigma},s}]$ includes the coefficients pertinent to the constant part of the variance covariance matrix and to the ANNs.

Unidentifiability. Unidentifiability of the log-likelihood is a well known issue in most ANN models, see Medeiros et al. (2006) for homoscedastic ANN models, McAleer et al.

(2008) and Medeiros et al. (2008) for heteroscedastic ANN models. We extend here this discussion to the VANNAR. Three features of the model cause lack of identifiability. The first refers to the possibility of including irrelevant ANNs. This is because in the mean equation (1) the input weights $\alpha_{\mathbf{y},k}$ and bias coefficient $\beta_{\mathbf{y},k}$ are unidentified if the corresponding output weights $\mu_k = \mathbf{0}$, for any $k = 1, \dots, q$. At the same time, μ_k and $\beta_{\mathbf{y},k}$ are unidentified if $\alpha_{\mathbf{y},k} = \mathbf{0}$. Similarly, in equation (3) the input weights $\alpha_{\Sigma,k}$ and bias coefficient $\beta_{\Sigma,k}$ are unidentified if the output weight matrix $\Sigma_k = \mathbf{0}$, for any $k = 1, \dots, s$. Further, Σ_k and $\beta_{\Sigma,k}$ are unidentified if $\alpha_{\Sigma,k} = 0$. The second refers to the interchangeability of the ANNs. This means that it is possible to permute the ANNs and the output weights in both equations (1) and (3) and obtain the same log-likelihood. The third feature causing unidentifiability refers to the re-scaling of the ANNs modelled through logistic functions, since $g(x) = 1 - g(-x)$.

To deal with unidentifiability, Stock and Watson (1998) proceed by conditioning the estimation of the log-likelihood function (5) on multiple sets of initial parameter values. Medeiros et al. (2006), Medeiros et al. (2008) and McAleer et al. (2008) proceed using restrictions on the ANNs that ensure identification of a unique value for the log-likelihood function. Following their approach, a possible set of sufficient restrictions for exact identification of the parameters in equations (1)-(4) is:

$$\alpha_{\mathbf{y},k} > 0, k = 1, \dots, q, \text{ and } \alpha_{\Sigma,k} > 0, k = 1, \dots, s, \quad (7)$$

$$\beta_{\mathbf{y},1} \leq \dots \leq \beta_{\mathbf{y},q} \text{ and } \beta_{\Sigma,1} \leq \dots \leq \beta_{\Sigma,s}, \quad (8)$$

$$\mu_{\mathbf{y},k} \neq 0, k = 1, \dots, q, \text{ and } \sigma_{k,r,c} \neq 0, r, c = 1, \dots, q, k = 1, \dots, s, \quad (9)$$

where $\sigma_{k,r,c}$ denotes the element in row r and column c of Σ_k , $k = 1, \dots, s$. The restrictions in (7) ensure that no irrelevant ANN is included in either the mean equation or the variance covariance matrix of the VANNAR. Restrictions (8) ensure that the value of the VANNAR log-likelihood is not invariant to the ordering of the ANNs in the mean and variance covariance matrix equations. Restrictions (9) imply that input weights and bias coefficients in both the mean and variance covariance matrix equations are always identified.

Indeterminacy. ANN modelling is often motivated by advocating the so-called universal approximator property. This states that under mild regularity conditions, a nonlinear pro-

cess can be approximated arbitrarily accurately by a linear combination of a large enough number of ANNs, see for example Hornik et al. (1989). However, this offers no guidance on the number/type of explanatory variables, ANNs and layers required to achieve a desired degree of accuracy. Avoiding over parametrization is therefore paramount when using ANNs. One possibility is to adapt the simple-to-general approach used by Swanson and White (1995, 1997) and Medeiros et al. (2006) to the VANNAR in (1)-(4). This starts by selecting the optimal lag length for the homoscedastic VAR representation of \mathbf{y}_t on the basis of one or more information criteria. This is then extended including ANNs in the mean and variance covariance matrix starting with the smallest possible VANNAR structure (one ANN in either the mean or the variance covariance matrix), and then gradually adding further ANNs. The optimal VANNAR specification is deemed to be reached once no reduction in the information criteria is found by adding further ANNs to a given specification.³

Stability. Necessary and sufficient conditions to ensure the stability of ANN models do not exist in the literature, since nonlinearity prevents the analytical formulation of these conditions.⁴ To ensure stability of the VANNAR dynamics we proceed as follows. Suppose the coefficients in Θ_2 are known. Then the VANNAR is linear in the remaining parameters and the conditional mean (1) can be equivalently written in the companion form $\mathbf{Y}_t = \mathbf{A}\mathbf{Y}_{t-1} + \mathbf{U}_t$, where $\mathbf{Y}_t = [\mathbf{y}'_t, \dots, \mathbf{y}'_{t-p+1}, \mathbf{y}'_{t-p}]'$, $\mathbf{U}_t = [\mathbf{v}'_t, \dots, \mathbf{0}', \mathbf{0}']'$, $\mathbf{v}_t = \boldsymbol{\lambda} + \Gamma\mathbf{w}(\boldsymbol{\alpha}_y, \boldsymbol{\beta}_y; \mathbf{z}_t) + \mathbf{u}_t$ and \mathbf{A} is a matrix in which the remaining coefficients in $\boldsymbol{\Lambda}$ are re-arranged in a form compatible with \mathbf{Y}_t . The companion system can be employed as a tool to monitor stability. According to this, the VANNAR in (1)-(4) is stable if:

$$\det(\mathbf{I} - \mathbf{A}\zeta) \neq 0, \forall \zeta \in \mathbb{C}, |\zeta| \leq 1, \quad (10)$$

which states that, conditional on Θ_2 , the VANNAR is stable if the eigenvalues of \mathbf{A} have modulus less than one. Whether (10) can effectively ensure the stability of the VANNAR

³Of course it would be possible to determine the optimal VANNAR structure using instead a general-to-specific approach. This is often found to be less accurate but computationally more efficient than simple-to-general method, see for example MacKay (1992). For the empirical analysis, we adopt the simple-to-general search approach also because existing applications in macroeconomics often find models with only a few ANNs to yield the best statistical fit.

⁴In principle, stability could be evaluated numerically by simulating trajectories of \mathbf{y}_t for all possible sample histories and checking that these are stationary. In practice, aside from being computationally demanding, this offers no guarantee that stable paths can be found.

in (1)-(4) depends on \mathbf{z}_t . If this includes only exogenous variables, then (10) provides a necessary and sufficient condition for the stability of the VANNAR. If \mathbf{z}_t includes at least one of the variables in the endogenous vector \mathbf{y}_t , then (10) is neither a necessary nor a sufficient condition for stability. However, in this case it would at least ensure the stability of a first-order approximation of the VANNAR, thus still providing a useful device to monitor stability in practice. We find this is actually the case in the empirical analysis in the next section, where the stability constraint (10) ensures that the estimation procedure delivers stable dynamics even in the presence of the extreme observations caused by pandemic.

Estimation. There are two popular algorithms for the estimation of ANN models. In the absence of the stability constraint (10), the first is a gradient-based method tackling the estimation of all parameters at once. The second, often referred to as the neural method, is more appealing as it exploits the possibility of writing the model in a linear form conditional on a subset of the parameters, thereby estimating the remaining parameters analytically.⁵ We describe here how to adapt both algorithms to the specified VANNAR. We find that, once accounting for the stability constraint (10), the option of solving for some of the parameters analytically is no longer feasible and the two algorithms become equivalent.

Gradient Method. The full information maximum likelihood estimator of the parameters in the VANNAR can be calculated by maximizing the log-likelihood function in (5) subject to the restrictions in (6)-(9) and the stability constraint in (10). Thus, the estimation is formulated as a constrained nonlinear optimization problem. This can be solved with sequential quadratic programming (SQP), which represents the state of the art in nonlinear programming methods. The basic idea is to formulate a programming sub-problem based on a quadratic approximation of the Lagrangian of the original problem. This can then be solved with standard iterative algorithms, such as Newton’s method.⁶

Neural Method. This involves repeating until convergence two sequential steps. The first consists of determining a draw of Θ_2 . As this includes all the parameters in the variance covariance matrix and in the ANNs, restrictions (6)-(9) can be enforced at this

⁵For a textbook description, see Chapter 19 of Martin et al. (2013).

⁶In the special case of maximizing the unconstrained log-likelihood (5), then the SQP method reduces to the Newton’s method. If the log-likelihood maximization is subject only to equality constraints, then the SQP method is equivalent to applying Newton’s method to the first-order optimality conditions. An overview of SQP is found in Nocedal and Wright (2006).

stage. Conditional on Θ_2 , the VANNAR is linear in the remaining parameters. Thus, in the second step $\mathbf{v}_t = \mathbf{y}_t - \Gamma \mathbf{w}(\boldsymbol{\alpha}_y, \boldsymbol{\beta}_y; \mathbf{z}_t)$ is computed and the estimation problem reduces to finding the vector Θ_1 that maximizes $\sum_{t=p^*+1}^T \frac{1}{2}(\mathbf{v}_t - \Theta_1 \mathbf{x}_t)' \boldsymbol{\Sigma}_t^{-1} (\mathbf{v}_t - \Theta_1 \mathbf{x}_t)$ subject to (10). Once Θ_1 is obtained, the VANNAR log-likelihood can be evaluated from (5).

If stability is not a concern, the neural method is preferable because Θ_1 can be determined analytically via generalized least squares.⁷ This simplifies the estimation of the model parameters by side-stepping many of the numerical issues associated with the convergence of the gradient algorithm. However, the analytical solution is no longer an option for the neural method once estimation is constrained by (10), as also estimation of Θ_1 requires the solution of a highly nonlinear constrained optimization problem. Thus, when accounting for (10) the neural method is effectively equivalent to the gradient method.

Both the gradient-based and the neural methods described above account for the stability constraint (10) in the formulation of the Lagrangian function. As noted above, this is different from maximizing (5) subject to the restrictions in (6)-(9) and verifying (10) ex-post, which would not guarantee stability. It is this formulation of the maximum likelihood optimization problem that helps finding stable dynamics from the VANNAR even when there are abnormal observations in the data. We use a Bayesian method to estimate the VANNAR parameters. A detailed description of the estimation algorithm is provided in Appendix A. In essence, we proceed as follows. First, we draw a large number of possible sets of parameter values Θ from a random distribution and identify out of these the twelve that yield the highest unrestricted log-likelihood in (5). We then use each of these plus the OLS estimates from the VAR as separate initial values for the maximization of the unrestricted log-likelihood in (5). The resulting thirteen sets of optimal values are then used for a subsequent round of maximization of the log-likelihood function in (5) this time including the restrictions in (6)-(9) and the stability constraint in (10). We then select out of these the parameters vector that yields the highest log-likelihood as starting value for the final round of estimation, which is based on the Markov chain Monte Carlo algorithm of Chernozhukov and Hong (2003). This is supported by a Minnesota prior to cope with the high dimensionality of the parameter space and a Student- t prior on Θ_2 to ensure the

⁷Using $\text{vec}(\Theta_1') = (\sum_{t=p^*+1}^T \boldsymbol{\Sigma}_t^{-1} \otimes \mathbf{x}_t' \mathbf{x}_t)^{-1} \text{vec}(\sum_{t=p^*+1}^T (\mathbf{x}_t' \mathbf{v}_t \boldsymbol{\Sigma}_t^{-1}))$.

restrictions in (6)-(9) and regularize the ANN parameter values. Only posterior draws that satisfy (10) are retained at this final stage.

Impulse Response Function and Forecasts. The transmission mechanism of shocks in the VANNAR depends on the state of the economy, the size and sign of shocks. The algorithm we use to compute responses to shocks adapts to the VANNAR the methodology of generalized impulse response functions (IRFs) of Koop et al. (1996) extended to account for sign restrictions. In essence, conditional on a given period t we identify structural shocks through the Cholesky factorization of the variance covariance matrix and generate trajectories of the endogenous variables by iterating forward the VANNAR, accounting for uncertainty in the parameter estimates and in the shocks realization. We then measure from these trajectories the median IRF and confidence bands. When \mathbf{z}_t includes one or more variables in \mathbf{y}_t , the IRF accounts for the interaction between the evolution of the endogenous variables and the VANNAR parameters through the ANNs included in the conditional mean equation in (1) and the variance covariance matrix in (3). No direct feedback enters the IRF if \mathbf{z}_t is exogenous, though the responses to shocks still vary according to the state of the economy, the size and sign of shocks.

Forecasts of the endogenous variables in the VANNAR can be computed using a similar algorithm to that of the IRF, by simulating trajectories of the variables in \mathbf{y}_t based on different draws from the posterior of the parameters and sequences of the reduced form disturbances over the chosen forecast horizon. The simulated trajectories can then be used to obtain numerical approximations of the forecast density and the desired moments. As for the IRF, the simulated trajectories include the uncertainty in the parameter estimates, the uncertainty in the shocks realization and the feedback from \mathbf{z}_t on the VANNAR coefficients. An important special case occurs when \mathbf{z}_t is exogenous and its trajectory is known over the forecast horizon. As such, the dynamic evolution of the ANNs can be computed in advance and the VANNAR can be cast in an equivalent state space form with time dependent coefficients. The Kalman smoother can then be employed to generate forecasts of the variables in the endogenous vector. When \mathbf{z}_t includes one or more endogenous variables of the VANNAR but their path is pre-specified, the Kalman smoother algorithm can still be used to forecast remaining endogenous variables.

The algorithms for computing IRFs and forecasts from the VANNAR are described in detail in Appendix B. As outlined above, all these methods for dynamic analysis rely on trajectories simulated from the model. Dijk et al. (2002) observe that the main issue of this approach is that the simulated densities can be distorted if some of the draws from the posterior of the parameters lead to unstable trajectories and thus to ‘wild’ IRFs or forecasts. Of course, the constraint (10) ensures stability of all the draws from the posterior of the parameters, thereby preventing such a distortion in simulated densities from the VANNAR.

3 Empirical Results: The VANNAR

Data. The data are those employed by Lenza and Primiceri (2020) in their study on VAR estimation after March 2020. The vector \mathbf{y}_t includes monthly observations of the seven variables: the rate of *unemployment*; total *employment* in the nonfarm sector; total real personal consumption, *PCE*; real personal consumption of services, *PCE services*; the price index of the two consumption aggregates, *PCE (price)* and *PCE services (price)*, respectively; and the price index of core inflation, *core PCE (price)*.⁸ Lenza and Primiceri (2020) use observations from December 1988 to September 2020 to avoid the larger data variability during the pre-1990 period and focus on the instability caused by the pandemic from March 2020 onward. Of course, since the VANNAR allows time varying parameters in the mean and variance covariance matrix, we could use a longer sample and extend beyond September 2020. We retain the same sample period for the sake of comparison.

To appraise the significance of instability in these data, we have evaluated when the stability constraint (10) binds while recursively estimating a linear VAR with homoscedastic disturbances starting with a sample including observations up to January 2000 until the most recent available data at the time of writing, April 2021.⁹ We have carried out this exercise using a VAR with two lags, since this is the optimal lag length according to the Schwarz Information Criterion (SIC) in the majority of the estimation samples, and with

⁸All data are taken from the online database of the Federal Reserve Bank of St. Louis accessed on June 2021. See Lenza and Primiceri (2020) for a more detailed description.

⁹The recursive estimation is carried out as follows. We first estimate the VAR using all observations until January 2000 and evaluate the stability constraint. Next we add the observations of February 2000, re-estimate the VAR and re-evaluate the stability constraint. We proceed in this way until April 2021.

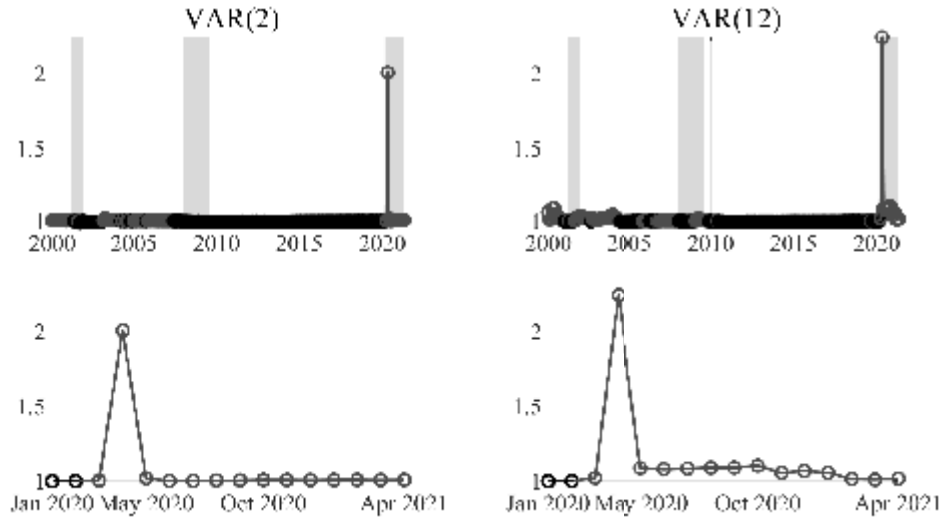


Figure 1: Maximum absolute value of the eigenvalue from each VAR recursively estimated between January 2000 and April 2021. Instances in red indicate unstable VAR. NBER recessions in grey.

twelve lags, as is often used in VAR with monthly data.

Figure 1 plots the maximum (absolute value of the) eigenvalue from each VAR recursively estimated during January 2000 and April 2021 in the top panels, zooming on January 2020 onward in the bottom panels. Instances in which the stability constraint (10) binds (maximum eigenvalue greater than one) are highlighted in red, meaning that the VAR estimated up to that period is unstable. Three main results are clearly visible. Firstly, the issue of VAR instability is not peculiar to the pandemic period, but occurs many times before: the stability constraint binds in about 29 and 35 percent of the estimated VARs when using two and twelve lags, respectively. Secondly, the extent of instability since March 2020 is unprecedented, as shown by the largest eigenvalue being above two at the start of the pandemic period. Third, the issues of VAR instability and its abnormal increase during the pandemic hold under any conventional specification of the VAR lag length.

The above results are due neither to the specific data frequency nor to the particular sample choice. To check this, we have considered an earlier starting date, January 1959, and converted the monthly observations of the seven indicators into quarterly averages. We have then re-estimated the VAR recursively over the 169 samples ending from 1979Q1 to

2021Q1. The stability constraint binds in about 43 and 49 percent of these samples when using a VAR with two and four lags, respectively. The VAR with two lags is unstable in the last three quarters of 2020, whereas the VAR with four lags is unstable in the the last three quarters of 2020 and in the first quarter of 2021.¹⁰

The Estimated Model. To appraise the effect of the abnormal variation in the data caused by the pandemic, the VANNAR is estimated using observations first prior to the pandemic, i.e. up to February 2020, and then including the first four and seven months of the pandemic, i.e. samples ending in June and September 2020, respectively.

The estimation of the VANNAR requires prior specification of the activation functions, of the explanatory variables, and of the number of lags and ANNs. Regarding the activation functions, the VANNAR is estimated using for all ANN terms in equations (1) and (3) either logistic, $g(\cdot) = 1/[1 + \exp(\boldsymbol{\alpha}'_{i,k}\mathbf{z}_{t-k} - \beta_{i,k})]$, or ReLU, $g(\cdot) = \max\{0, \boldsymbol{\alpha}'_{i,k}\mathbf{z}_{t-k} - \beta_{i,k}\}$, functions, the two most popular choices in modern neural networks. The choice of the explanatory variables is not trivial since there is no clear-cut theoretical prescription for what these should be. We use $\mathbf{z}_t = u_{t-1}$, where u_{t-1} is the rate of unemployment standardized according to (4). This is justified for three main reasons. Firstly, the rate of unemployment is often used as an indicator of the state of the economy in the macroeconomic literature using regime switching models, on the ground that it provides a proxy for the amount of slack in the economy, see for example Rothman (1998), Dijk et al. (2002), Skalin and Teräsvirta (2002), Ramey and Zubairy (2018).¹¹ Secondly, because the rate of unemployment is also an endogenous variable, IRFs and forecasts can account for direct feedback from the state of the economy to the conditional mean of the endogenous variables and the variance covariance matrix of the disturbances. Thirdly, the speed of change in the state of the economy, as measured by the month-to-month variation in the rate of unemployment, determines the speed of variation in the VANNAR coefficients. This is important because it ensures that the model can capture both gradual and more sudden changes in the coeffi-

¹⁰Further, as shown by Schorfheide and Song (2020) and Carriero et al. (2021), VAR models estimated using data of the United States from March 2020 onward are unstable when including several different types of macroeconomic indicators, and not just the specific ones used in this empirical analysis.

¹¹Ramey and Zubairy (2018) subtract the natural rate from the actual unemployment rate and it is the differential between these two that they employ as proxy for economic slack. In the VANNAR, this differential is implicitly accounted for by the estimated bias coefficient β .

Table 1: Simple-to-general search for VANNAR specification (logistic activation functions).

q	s	February 2020				June 2020				September 2020			
		LL	AIC	HIQ	SIC	LL	AIC	HIQ	SIC	LL	AIC	HIQ	SIC
0	0	4.61	-8.52	-7.97	-7.14	1.75	-2.79	-2.25	-1.43	1.68	-2.67	-2.13	-1.31
1	0	-18.40	37.55	38.14	39.03	-10.76	22.26	22.85	23.73	-47.20	95.15	95.72	96.60
0	1	4.84	-8.81	-8.14	-7.11	2.09	-3.32	-2.66	-1.64	1.83	-2.81	-2.14	-1.14
1	1	4.54	-8.18	-7.46	-6.38	1.91	-2.91	-2.21	-1.14	1.61	-2.32	-1.62	-0.55
2	1	4.45	-7.93	-7.19	-6.05	1.74	-2.54	-1.80	-0.67	1.80	-2.66	-1.92	-0.80
1	2	4.53	-7.99	-7.15	-5.88	1.69	-2.32	-1.49	-0.23	1.72	-2.39	-1.57	-0.31
2	2	4.51	-7.90	-7.03	-5.70	0.88	-0.65	0.22	1.54	1.80	-2.49	-1.63	-0.32

Note: Results for $q = s = 0$ in the June and September 2020 samples are disregarded for VANNAR specification since the model is unstable in these two samples.

cients, including those due to the pandemic, as illustrated below. The choice of the number of lags and ANNs in the VANNAR is entirely data driven and determined according to the simple-to-general search procedure described in Section 2.

Table 1 reports the log-likelihood and information criteria obtained from the simple-to-general search of the VANNAR specification in the three samples when the activation function is chosen to be logistic. According to the SIC, the best VAR lag length is equal to two in each sample. This is equivalent to the VANNAR with $q = s = 0$ reported in the first row of the table. As we do not impose the stability constraint on the VAR, the estimates for the June and September 2020 samples are unstable and for this reason the specification $q = s = 0$ is disregarded for the purpose of model selection in these two samples. According to most of the diagnostics in Table 1, the VANNAR with only one ANN in the variance covariance matrix gives the best statistical fit in all three samples.¹² None of the VANNAR estimated using ReLU activation functions improves on the outcomes in Table 1.¹³ Thus the estimated VANNAR that we use for the subsequent empirical analysis employs logistic functions and is based on $q = 0$ and $s = 1$.¹⁴ While being entirely data driven, this result is

¹²As a robustness check we re-estimated all models in Table 1 using \mathbf{z}_{t-1} in the additional ANNs for $q, s > 1$ rather than \mathbf{z}_{t-k} . We still find that the VANNAR with $q = 0$ and $s = 1$ yields the best fit. We did not test for specifications based on $q, s > 2$ or additional layers.

¹³These results are not included here for reason of space, but are reported in Appendix C.

¹⁴A clear pattern visible from Table 1 is that the log-likelihood estimated for each model is almost halved as the sample changes from February to either June or September 2020. This reduction is the necessary price to pay to fit each VANNAR on the recent observations that include much larger variability compared to historical standards.

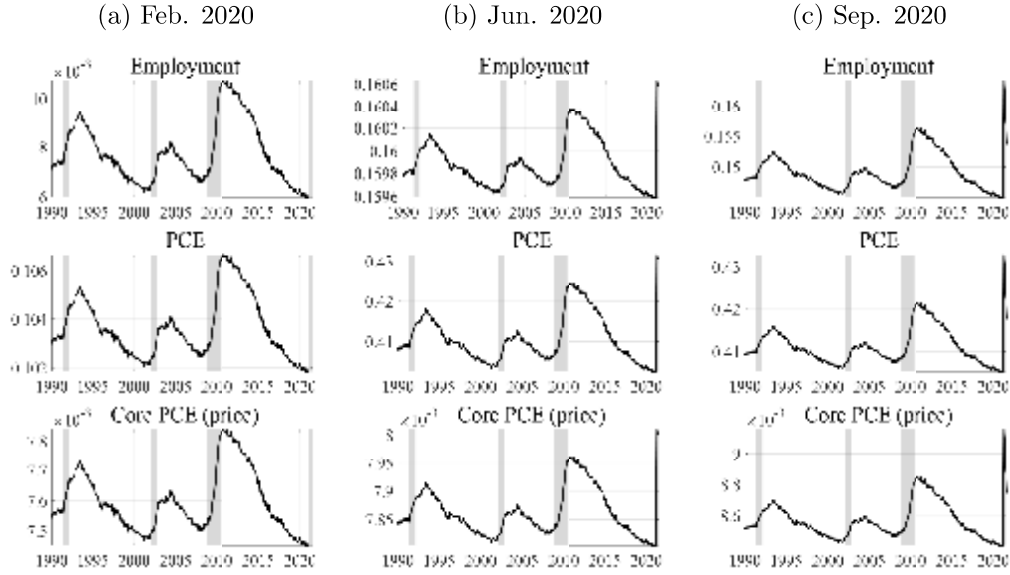


Figure 2: Time-varying volatility in the VANNAR residuals as measured by the diagonal elements of Σ_t . NBER recessions in grey.

evidently consistent with the view underpinning the current literature on VAR estimation since March 2020, namely that the extreme realizations caused by the pandemic are due to temporary jumps in volatility, rather than structural changes in macroeconomic data.

Figure 2 displays the the time-varying volatility in the VANNAR residuals, measured by the diagonal elements of Σ_t , across the three samples. Since volatility is affected by the ANN terms, the figure provides information on the type of nonlinearity captured by the VANNAR in each estimation sample. From here on, we present the results for employment, aggregate consumption, PCE, and the general price level, core PCE(price), since these are representative of the main compartments of the aggregate economy covered in the dataset, i.e. the labour market, the real and nominal sector.¹⁵ The vertical axes are not harmonized across the estimation samples to facilitate visualization. The following three results are worth observing. Firstly, volatility displays patterns similar to those estimated from a regime switching VAR where the transition variable is the rate of unemployment, in that it increases during periods of high unemployment around recessions, see for example Ramey

¹⁵The results for the other macroeconomic indicators are reported in Appendix C. The overall picture does not changes significantly once considering the nonlinear effects stemming from the ANNs on the covariances in (3). We omit reporting these for reason of space.




Figure 3: Median, 68 and 95 percent confidence bands of responses to unanticipated increase in the rate of unemployment from the VANNAR. All variables are in logs $\times 100$.

and Zubairy (2018).¹⁶ Secondly, the VANNAR accounts for the extreme observations as estimation goes beyond February 2020 through a sudden upswing in volatility. Thirdly, the latest increase in volatility is larger but relatively short-lived compared to those observed during previous recessions, as visible from the fact that by September 2020 volatility has already reduced to levels close to those reached in the aftermath of the Great Recession.¹⁷

Impulse Responses. Figure 3 shows the responses of employment, consumption and the aggregate price level to an unanticipated increase in the rate of unemployment over 60 months, with their 68 and 95 percent confidence bands. As in Lenza and Primiceri (2020), the shock is identified from the Cholesky decomposition of the variance covariance matrix and sign restrictions ensure that the rate of unemployment is positive for the entire simulation horizon. The responses are measured from the VANNARs estimated over the three samples, and the shock occurs at the end of each of sample. Two results are

¹⁶Similar patterns are obtained from regime switching VARs that condition volatility on cyclical variables such as the growth rate of GDP or indicators of financial risk.

¹⁷In Appendix D, we show that many of these features are also observable from the estimates of macroeconomic volatility obtained under several of the alternative methodologies recently proposed to deal with large outliers caused by pandemic.

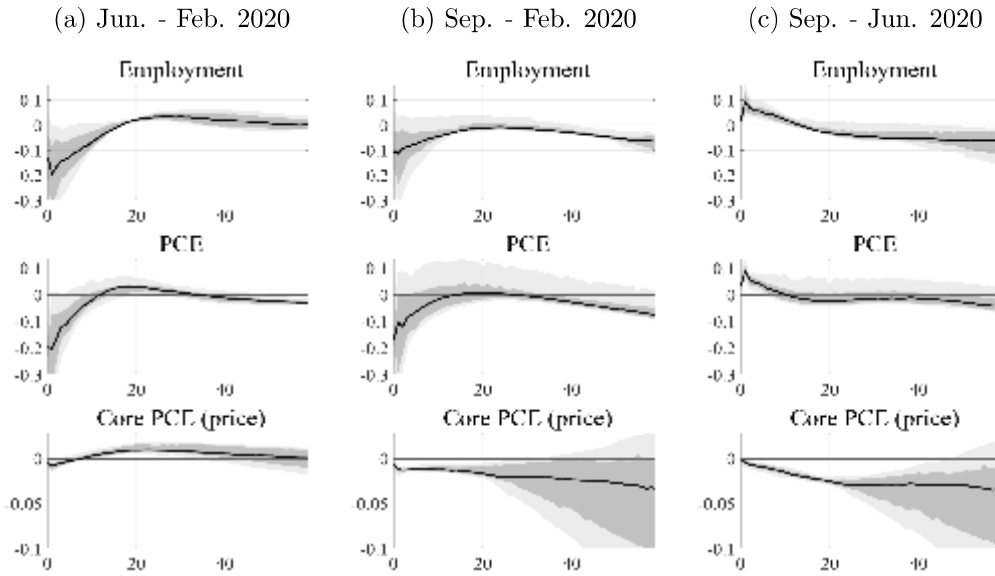


Figure 4: Median, 68 and 95 percent confidence bands of response differentials to unanticipated increase in the rate of unemployment from the VANNAR. All variables are in logs $\times 100$.

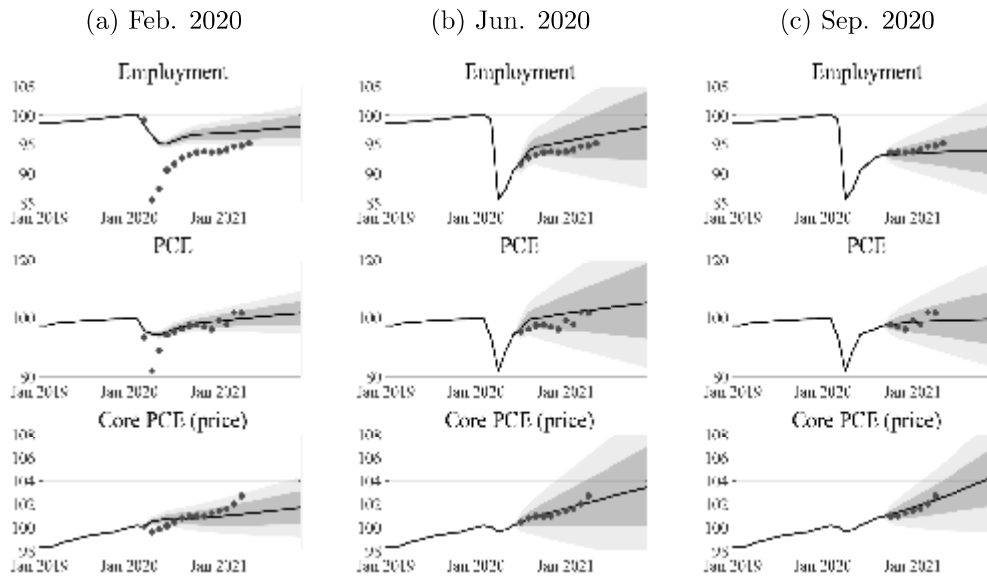


Figure 5: Median, 68 and 95 percent confidence bands of out-of-sample forecasts from the VANNAR conditional on the unemployment prediction from the Blue Chip Forecast and actual data (Feb. and Jun. 2020). All variables are in logs $\times 100$.

immediately visible. First, the estimated VANNAR yields plausible inference about the structure of the economy in all three samples (**R1**). The unemployment shock causes a fall in employment, aggregate consumption and prices over the five years of the simulation horizon. This is consistent with the typical response of the economy to a persistent aggregate demand shock and with the macroeconomic responses estimated by Lenza and Primiceri (2020). It is also consistent with the evidence from survey data suggesting that households in the United States expect high unemployment and low inflation in the aftermath of the Covid-19 shock, see Coibion et al. (2020).¹⁸ Second, structural inference remains stable once the VANNAR is estimated in the two samples that include the abnormal observations caused by the pandemic, June and September 2020 (**R2**). Third, the magnitude of the macroeconomic responses and the uncertainty surrounding them over the simulation horizon increases during the pandemic periods (**R3**). This is particularly visible from the responses of employment and consumption, on impact and over the first twenty months of the simulation horizon.

To formally evaluate how the macroeconomic responses change due to the pandemic we constructed a test based on the empirical distribution of the differentials (for each variable and time horizon) between 1000 posteriors of the IRF estimated across two different samples. Since all responses in Figure 3 are below zero, a negative (positive) differential indicates increase (decrease) in the response magnitude at the observed horizon.

Figure 4 displays the median, 68 and 95 percent confidence bands of the response differentials between the estimation samples June and February, September and February, and September and June 2020. We highlight the following results. For the real variables the increase in the responses magnitude from February to June 2020 is statistically significant on impact and up to about the first twelve months with 68 percent of confidence. These differentials reduce over the pandemic period and are found to be not different from zero from a statistical stance. The differentials for the nominal activity responses are smaller than those of the real activity, though still significant between three to six months.

¹⁸Using data from a smaller survey conducted at the beginning of March 2020, Binder (2020) finds that United States' consumers expect high unemployment but also high inflation as a result of the pandemic. The contrasting evidence about inflation expectations compared to the survey of Coibion et al. (2020), which includes responses gathered in April 2020, may be because respondents had time to update their information about price dynamics since March 2020.

Table 2: Root Mean Square Forecast Errors from VANNAR.

Sample: Variable/Horizon	February 2020						June 2020				September 2020		
	1	3	7	10	12	14	1	3	7	10	1	3	7
Employment	0.83	7.79	3.24	3.21	2.82	2.40	0.69	1.74	2.75	3.11	0.69	1.09	2.12
PCE	2.06	5.31	1.15	2.84	2.19	2.49	1.71	3.61	4.90	6.04	1.33	3.08	4.83
Core PCE (price)	0.10	0.79	0.45	0.61	0.83	1.56	0.32	0.63	1.21	1.70	0.33	0.57	1.06

That change of macroeconomic volatility is associated with change in the magnitude of the economy response to shocks is not a new finding. It is for example a recurrent result from the Great Moderation macroeconomic literature, which gives evidence of permanent structural breaks in leading macroeconomic indicators around the mid-1980s, see for example Boivin et al. (2010). In contrast, the above results suggest that the increase in the macroeconomic responses magnitude and uncertainty caused by the recent pandemic are only temporary.¹⁹

Conditional Forecasts. Figure 5 shows the median, 68 and 95 percent confidence bands of the out-of-sample forecasts of employment, aggregate consumption and prices until December 2021 estimated from the VANNARs against the actual data (red dots) available up to April 2021. Table 2 reports the corresponding root mean square forecast errors (RMSFEs) at selected time horizons.²⁰ As in Lenza and Primiceri (2020), for the sample ending in September 2020 the forecast is conditioned on the consensus unemployment prediction from October 2020 to December 2021 taken from the October 2020 release of the Blue Chip Forecast. For the other two samples the forecast is further conditioned on the actual unemployment rate data until September 2020. We highlight two main results. First, the out-of-sample forecasts capture both the disruption on the wider economy caused by the pandemic from February 2020, as well as the slow recovery from June 2020 onward (**R4**). Second, according to the visual evidence and the computed RMSFEs, the forecast accuracy is higher once the model is tasked to predict the recovery phase of the economy

¹⁹To further verify this statement, we have re-estimated the VANNAR using data up to April 2021. The IRFs and differentials are similar to those shown in Figures 3 and 4 for September 2020. These are reported in Appendix C.

²⁰The RMSFEs are computed from the forecast errors measured out of 1000 posteriors of the conditional forecasts for any starting period and horizon.

(samples terminating in June and September 2020) rather than the sharp disruption in the early months of the pandemic (sample ending in February 2020) (**R5**).²¹

4 Empirical Results: Alternative Approaches

Specification. We compare the responses to shocks and forecasts presented above with those obtained from four alternative approaches recently proposed to deal with the VAR instability caused by the Covid-19 pandemic. These are all based on a VAR with constant coefficients in the equations for the mean but differ for the specification of the variance covariance matrix.²² The first uses a constant variance covariance matrix and estimates the VAR on data up until February 2020, in the same spirit of Schorfheide and Song (2020) and Primiceri and Tambalotti (2020). We refer to this as the VAR-C.²³ The second is the VAR with the ad-hoc modification of the variance covariance matrix proposed by Lenza and Primiceri (2020), hereafter VAR-LP. Using information from observed data, this assumes that volatility increases in the first three months of the pandemic, between March and May 2020, and gradually decreases from June onward. Thus, when estimation is based on the pre-pandemic period, no adjustment to the variance covariance matrix is required and the VAR-LP is effectively equivalent to the VAR-C.²⁴ The third and fourth alternatives

²¹As for the IRF, forecasts uncertainty increases significantly during the first four months of the pandemic, reducing only marginally afterwards. The effect of this change in uncertainty on the forecasts is of course accounted for in the computed RMSFEs.

²²For reason of space, we only describe in the paper the main features of each approach. The interested reader can refer to the original sources for specific details. Each alternative model is estimated using Bayesian methods with the support of a standard Minnesota prior and on the same data employed for the VANNAR. Further details on estimation and additional results are given in Appendix D.

²³Schorfheide and Song (2020) consider a mixed-frequency VAR including eleven macroeconomic indicators of the United States. They measure only out-of-sample forecasts, using first the VAR estimated with data up to the end of 2019, and then a sequence of VARs recursively estimated up to June 2020. They find that the VAR estimated until the end of 2019 gives more stable and reasonable forecasts than those estimated on samples that include the early pandemic periods. Primiceri and Tambalotti (2020) consider a VAR including monthly observations up to February 2020 of six indicators of the aggregate United States economy. They forecast the impact of an ad-hoc shock that mimics the observed effects of Covid-19 on the aggregate economy, under alternative scenarios regarding the duration of the pandemic. Our VAR-C is in the spirit of the VAR proposed in these two studies in that we curtail the estimation sample to the pre-pandemic period. As the objective is the comparison with the VANNAR, we neither evaluate the effects of using mixed-frequency data nor consider the construction of an ad-hoc shock based on observed aggregates.

²⁴Lenza and Primiceri (2020) compute both for IRFs and out-of-sample forecasts. Our results for the VAR-LP are based on the replication code made available by the authors on their webpage.

follow the proposal of Carriero et al. (2021) of enhancing a VAR with stochastic volatility allowing for temporary jumps in the variances of the disturbances through either fat-tailed errors drawn from a Student- t distribution, as pioneered by Jacquier et al. (2004), or the outlier-augmented stochastic volatility specification of Stock and Watson (2016). We refer to these last two alternatives as VAR-SV t and VAR-SVO, respectively.²⁵

It is useful to compare the variance covariance matrix of the last three approaches with that of the VANNAR. In these three models volatility is formulated as $\Sigma_t = \mathbf{B}^{-1}\mathbf{H}_t^m(\mathbf{B}^{-1})'$; where $m = LP, SVt, SVO$, \mathbf{B} is a lower triangular matrix; and \mathbf{H}_t^m is a diagonal matrix given by $\mathbf{H}_t^{LP} = s_t^2\mathbf{D}$, $\mathbf{H}_t^{SVt} = \mathbf{O}_t^1\mathbf{D}_t(\mathbf{O}_t^1)'$ and $\mathbf{H}_t^{SVO} = \mathbf{O}_t^2\mathbf{D}_t(\mathbf{O}_t^2)'$ in the VAR-LP, VAR-SV t , and VAR-SVO, respectively. In the \mathbf{H}_t^m matrices, \mathbf{D} is diagonal, with coefficients being either constant, in the VAR-LP, or stochastic, in the VAR-SV t and VAR-SVO. The remaining elements capture drifts in volatility and their specification depends on whether the sample includes abnormal outliers. In the estimation sample terminating in February 2020 s_t is set to be equal to one, while \mathbf{O}_t^1 and \mathbf{O}_t^2 are set as identity matrices. Thus the VAR-LP is equivalent to the VAR-C, while both the VAR-SV t and VAR-SVO reduce to a VAR with stochastic volatility. However, s_t , \mathbf{O}_t^1 and \mathbf{O}_t^2 are appropriately modified to downweight the abnormal outliers in the June and September estimation samples. In specific, s_t takes three different values between March and May 2020, and decays following an autoregressive process from June onward. Both \mathbf{O}_t^1 and \mathbf{O}_t^2 are set as diagonal matrices of mutually and serially i.i.d. unobserved states drawn from either a Student- t distribution (with five degrees of freedom) in the VAR-SV t or a uniform distribution (with support between two and twenty) in the VAR-SVO. Aside from not requiring these ad-hoc changes, the variance covariance matrix of the VANNAR in equation (3) differs from these alternative models in two main dimensions. Firstly, it measures time variation additively through the ANNs, as in McAleer et al. (2008). Secondly, it does not restrict the matrices Σ_k , $k \geq 1$, to be diagonal. Thus, unlike the alternative specifications, the VANNAR allows time variation in the variances and covariances to be independent from each other.

²⁵Carriero et al. (2021) use monthly observations for sixteen macroeconomic and financial indicators for the United States from March 1959 to September 2020. They employ the VAR-SV t and VAR-SVO only for out-of-sample forecasting. As replication codes were unavailable at the time of writing, the results for the VAR-SV t and VAR-SVO models presented in this paper are based on our implementation of the algorithms described in their paper.




Figure 6: Median, 68 and 95 percent confidence bands of responses to unanticipated increase in the rate of unemployment from the VAR-C (first column only) and the VAR-LP (all columns). All variables are in logs $\times 100$.




Figure 7: Median, 68 and 95 percent confidence bands of responses to unanticipated increase in the rate of unemployment from the VAR-SVt. All variables are in logs $\times 100$.




Figure 8: Median, 68 and 95 percent confidence bands of responses to unanticipated increase in the rate of unemployment from the VAR-SVO. All variables are in $\log \times 100$.

Table 3: Median response differentials between alternative models and the VANNAR.

Sample: Model/Horizon	February 2020				June 2020				September 2020			
	1	12	24	60	1	12	24	60	1	12	24	60
Employment												
VAR-C	-0.002	0.007	0.032	0.009	0.128*	0.061	-0.000	0.010	0.112	0.048	0.044	0.074
VAR-LP	-0.002	0.007	0.032	0.009	-0.091	-1.577**	-1.838**	-1.642*	0.088	-0.066	-0.095	-0.075
VAR-SV t	-0.003	0.068	0.096*	0.054	0.064	-0.056	-0.095	-0.101	0.070	0.010	0.015	0.037
VAR-SVO	-0.003	0.068	0.096*	0.054	-0.166	-0.520**	-0.668**	-0.772*	0.083	0.061	0.064	0.071
PCE												
VAR-C	0.046	0.016	0.041	0.029	0.243*	0.024	0.017	0.061	0.214*	0.031	0.037	0.105
VAR-LP	0.046	0.016	0.041	0.029	-0.205	-1.213*	-0.950	-1.441*	0.172*	-0.041	-0.015	-0.019
VAR-SV t	0.044	0.040	0.024	-0.008	0.043	-0.341*	-0.565*	-0.901**	0.089	-0.188	-0.311*	-0.446*
VAR-SVO	0.044	0.040	0.024	-0.008	-0.254	-0.803**	-1.159**	-1.860**	0.177*	-0.007	-0.046	-0.020
Core PCE (price)												
VAR-C	0.005	-0.010	-0.021	-0.035	0.007	-0.012	-0.031	-0.034	0.010	0.004	-0.003	0.002
VAR-LP	0.005	-0.010	-0.021	-0.035	-0.010	-0.464**	-0.887**	-1.245**	0.007	-0.037	-0.068*	-0.090
VAR-SV t	0.009	0.020	0.032*	0.041*	0.059	0.040	0.038	0.034	0.010	0.034	0.052	0.095
VAR-SVO	0.009	0.020	0.032*	0.041*	0.064*	0.074	0.091	0.135	0.012	0.033*	0.053*	0.084*

Note: * and ** indicate significance at 68 and 95 percent confidence, respectively.

Impulse Responses. Figures 6, 7 and 8 present the median, 68 and 95 percent confidence bands of the responses of employment, aggregate consumption and the price level to an unanticipated increase in the rate of unemployment estimated from the VAR-LP, VAR-SV t and VAR-SVO, respectively. As noted above, the VAR-LP estimated on the sample ending by February 2020 is equivalent to the VAR-C. Thus, the macroeconomic responses from the VAR-C can be simply observed from the first column of Figure 6. In each model, the shock to the rate of unemployment is identified using the same protocol of Lenza and Primiceri (2020) used for the VANNAR.²⁶

The main result we highlight is that, to a great extent, the responses from the four alternative models yield structural inference similar to that obtained from the VANNAR (R6). The responses of the real economy (employment and consumption) are qualitatively similar across the four alternatives and to those obtained from the VANNAR, in that the shock generates a sharp and long-lasting reduction in the real activity. The responses of the nominal economy estimated from the VAR-C and the VAR-LP are also similar to those from the VANNAR, in that the shock leads to a persistent reduction in the price level, but differ from those estimated from the VAR-SV t and the VAR-SVO which do not show similar reductions. Further, the macroeconomic responses from the VAR-LP, the VAR-SV t and the VAR-SVO display an increase in magnitude and uncertainty as the estimation sample changes from February to June 2020 thereby including the early pandemic observations.²⁷ Uncertainty reduces only partially as the economy progresses over the pandemic but it is still higher than that measured for the pre-pandemic period, as visible when comparing the bands around the responses in June 2020 with those in September 2020.

To further evaluate these results, we construct an empirical test of the difference between the responses from the alternative models and the VANNAR.²⁸ Table 3 presents the

²⁶Vertical axes in the figures are not harmonized across the estimation samples to facilitate visualization. In particular, note how the scale increases by a factor of 10 when moving from February to June 2020.

²⁷By construction, the macroeconomic responses from the VAR-C are time invariant, being exactly the same for February, June and September 2020.

²⁸For each alternative model we compute 1000 posteriors of the IRF for any of the three samples and at any given horizon. We then subtract (from each of these) 1000 posteriors of the corresponding IRF estimated from the VANNAR to construct empirical densities of the response differentials. We extract from these densities the median differential, the 68 and 95 percent confidence bands. Evidence in favour of macroeconomic responses from an alternative model being different from those of the VANNAR is gained if zero is not included in the confidence bands at a given statistical level.

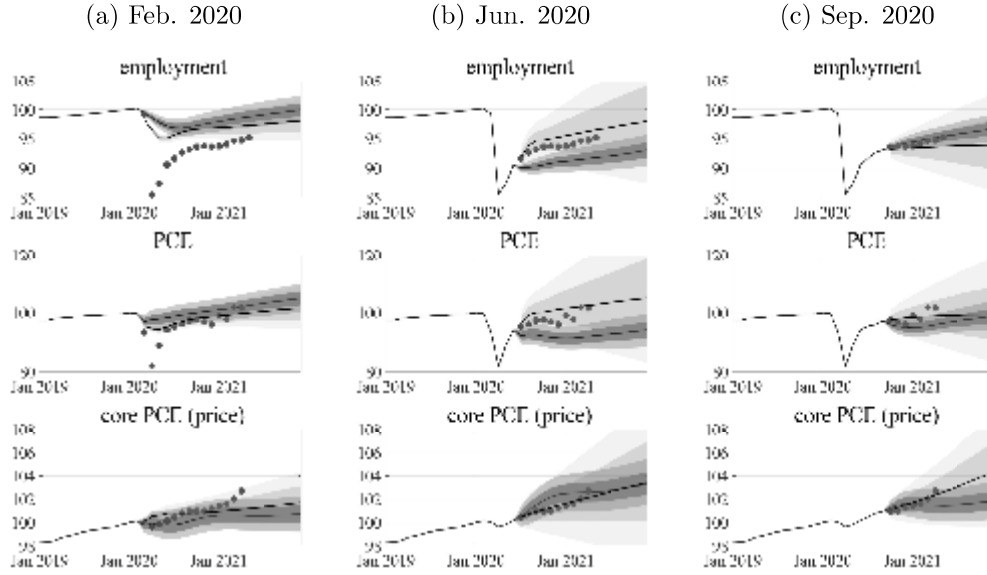


Figure 9: Median, 68 and 95 percent confidence bands of out-of-sample forecasts from VAR-C (blue) and VANNAR (grey) conditional on the unemployment prediction from the Blue Chip Forecast and actual data (Feb. and Jun. 2020). All variables are in logs $\times 100$.

results from this exercise for the responses of employment, aggregate consumption and price level evaluated at four given horizons over the three estimation samples. Each row reports the median differential of a given alternative model relative to the VANNAR, with one and two stars indicating whether this is different from zero with 68 and 95 percent confidence, respectively. According to these results, the macroeconomic responses estimated from the VANNAR are statistically not different from those obtained under any of the alternative models for the majority of samples and time horizons. The few instances where the responses are statistically different tend to be concentrated within the June sample and for the variables measuring the real activity, employment and consumption.²⁹

Conditional Forecasts. Figures 9, 10, 11 and 12 display the median, 68 and 95 percent confidence bands of the conditional forecasts of employment, aggregate consumption and price level from the VANNAR against those from the VAR-C, VAR-LP, VAR-SVO and VAR-SV t , respectively. As for the VANNAR, the forecasts are computed starting from the end of each estimation sample. In each figure, solid black lines until the start of the

²⁹As a further robustness check, we recomputed the test using the mean differentials as opposed to the mean. The results are essentially unaltered.




Figure 10: Median, 68 and 95 percent confidence bands of out-of-sample forecasts from VAR-LP (blue) and VANNAR (grey) conditional on the unemployment prediction from the Blue Chip Forecast and actual data (Feb. and Jun. 2020). All variables are in logs \times 100.




Figure 11: Median, 68 and 95 percent confidence bands of out-of-sample forecasts from VAR-SVt (blue) and VANNAR (grey) conditional on the unemployment prediction from the Blue Chip Forecast and actual data (Feb. and Jun. 2020). All variables are in logs \times 100.




Figure 12: Median, 68 and 95 percent confidence bands of out-of-sample forecasts from VAR-SVO (blue) and VANNAR (grey) conditional on the unemployment prediction from the Blue Chip Forecast and actual data (Feb. and Jun. 2020). All variables are in logs $\times 100$.

forecasts indicate actual data, red dots denote actual data over the forecasting horizon until April 2021. The main results observable from these figures is that each of the alternative models yield out-of-sample forecasts qualitatively similar to those from the VANNAR (**R7**). This is supported by three observations. Firstly, no model estimated with data up until February 2020 can accurately predict the effects of the pandemic on the wider economy. This is particularly evident for the large decline in the real economic activity more than the nominal activity. Secondly, the accuracy of the forecasts visibly improves once considering the estimates in June and September 2020, though still most of the models appear to understate the observed recovery of aggregate consumption. Thirdly, under all alternative models, macroeconomic uncertainty increases once estimation includes the latest data.

To formally evaluate how the forecasts from these alternative models compare to those of the VANNAR, we constructed an empirical test of the ratio between the RMSFEs estimated under each alternative model and the VANNAR, based on 1000 posteriors of each model conditional forecast for any starting period and horizon. The resulting empirical densities are then used to evaluate the mean RMSFE ratios, the 68 and 95 percent confidence bands.

Table 4: Root Mean Square Forecast Errors Ratios between alternative models and the VANNAR.

Sample: Model/Horizon	February 2020						June 2020				September 2020		
	1	3	7	10	12	14	1	3	7	10	1	3	7
Employment													
VAR-C	0.42**	1.34**	1.27	1.44	2.21	6.65	12.53*	5.90	4.11	3.47	3.01	4.90	1.49
VAR-LP	0.41*	1.34**	1.27	1.46	2.23	7.01	12.69	6.39	5.10	5.17	8.27	2.51	1.54
VAR-SV t	0.78	1.31*	1.99**	2.14**	3.23**	10.27**	36.89*	23.55	24.90	20.99	48.33	33.46	10.84
VAR-SVO	0.78	1.31*	1.99**	2.14**	3.23**	10.27**	3.38	2.00	2.15	1.70	2.03	2.21	1.40
PCE													
VAR-C	1.86*	1.73**	8.93	5.02	7.00	7.77	107.26	11.38	6.43	21.15	7.48	2.42	4.31
VAR-LP	1.82*	1.73**	10.13	4.91	7.38	3.91	230.85	12.52	6.79	30.49	8.93	2.87	4.52
VAR-SV t	2.26*	1.15	6.44	2.59	3.51	10.43	53.08	8.02	4.43	21.22	7.48	5.20	3.33
VAR-SVO	2.26*	1.15	6.44	2.59	3.51	10.43	62.98	6.46	3.27	13.01	2.26	2.66	2.45
Core PCE (price)													
VAR-C	15.68*	0.71	7.10	7.51	3.81	3.16	6.23	6.70	11.28	1.47	6.74	16.35	4.26
VAR-LP	16.96*	0.69	6.78	6.89	3.94	3.31	14.70	15.38	20.77	4.27	3.17	11.68	4.36
VAR-SV t	8.74	1.43	2.52	3.56	1.43	0.84*	1.55	4.00	5.27	2.43	2.83	11.29	2.72
VAR-SVO	8.74	1.43	2.52	3.56	1.43	0.84*	1.53	2.99	1.08	2.15	4.08	5.82	2.64

Note: Values above (below) one indicate that the forecast from a given alternative model does not (does) improve that from the VANNAR. * and ** indicate significance at 68 and 95 percent confidence, respectively.

Table 4 presents the results from the computed mean RMSFE ratios. Values above (below) one indicate that the forecast from a given alternative model does not (does) improve that from the VANNAR, with one and two stars indicating 68 and 95 percent confidence, respectively. According to the results in the table, none of the alternative models can produce more accurate forecasts relatively to the VANNAR, regardless of the estimation sample and for the majority of forecast horizons considered (**R8**). In particular, considering February 2020, the VANNAR forecasts for employment, over most of the forecast horizon, and consumption, up to three months ahead, are more accurate than most alternatives. Forecasts of consumption beyond three months and of the price level turn out to be not statistically different. As the starting period of the forecast moves to either June or September 2020, the mean RMSFE ratios are generally greater than one in most instances, pointing towards the greater accuracy of the VANNAR forecasts relatively to the alternatives considered. However, in almost all instances the improvement in the forecast performance turns out to be not statistically significant at the conventional levels due to the large standard errors from the alternative approaches in these two samples, as visible

from figures 9, 10, 11 and 12.³⁰

5 Conclusion

In this paper we develop a regime switching vector autoregression where artificial neural networks drive time variation in the coefficients of the conditional mean of the endogenous variables and the variance covariance matrix of the disturbances. The model, which we term vector artificial neural network autoregression (VANNAR), includes a stability constraint to ensure non-explosive dynamics in any instance where a (non)linear VAR would otherwise be unstable. The proposed VANNAR methodology is entirely data driven and designed to capture not only many of the typical characteristics observed in nonlinear time series data, such as randomness, cycles and stochastic trends, but also more extreme events, such as jumps, asymmetric adjustments and large outliers that cannot be captured adequately by linear homoscedastic VARs.

We confront the proposed methodology with one of the latest challenge in empirical macroeconomics, the estimation of VARs on aggregate data that include the large outliers caused by the Covid-19 pandemic. The empirical results show that the VANNAR delivers plausible (**R1**) and stable (**R2**) structural inference, with the magnitude and uncertainty of the macroeconomic responses to shocks temporarily increasing during the pandemic periods (**R3**). The model forecasts capture only in part the macroeconomic slump observed in the first three months of the pandemic, tracking much better the subsequent slow recovery phase (**R4** and **R5**). As a robustness exercise, we compare these results with those from four alternative methodologies recently proposed to deal with VAR estimation in the presence of the large outliers. We find that structural inference undertaken with any alternative model is generally consistent with that from the VANNAR (**R6**). All alternative models yield out-of-sample forecasts qualitatively similar to those from the VANNAR (**R7**), but none of them can produce more accurate forecasts relatively to the VANNAR (**R8**).

³⁰For example, the estimated RMSFE ratio between the VAR-LP and the VANNAR for PCE at the one month horizon is very large (230.87) because the density of the forecasts from the VAR-LP is very diffuse while that from the VANNAR is much tighter at the one month horizon. Thus the RMSFEs sampled from the VAR-LP are large, while those from the VANNAR are much smaller, hence the large ratio. The very diffuse density of the VAR-LP forecasts also determines the statistical insignificance of the RMSFE ratio.

References

- Baqae, D. and E. Farhi (2020). Supply and demand in disaggregated keynesian economies with an application to the covid-19 crisis. Technical report, National Bureau of Economic Research.
- Baqae, D., E. Farhi, M. J. Mina, and J. H. Stock (2020). Reopening scenarios. Technical report, National Bureau of Economic Research.
- Barro, R. J. (2020). Non-pharmaceutical interventions and mortality in us cities during the great influenza pandemic, 1918-1919. Technical report, National Bureau of Economic Research.
- Binder, C. (2020). Coronavirus fears and macroeconomic expectations. *The Review of Economics and Statistics* 102(4), 721–730.
- Boivin, J., M. T. Kiley, and F. S. Mishkin (2010). How has the monetary transmission mechanism evolved over time? In *Handbook of Monetary Economics*, Volume 3, pp. 369–422. Elsevier.
- Bredahl Kock, A. and T. Teräsvirta (2016). Forecasting macroeconomic variables using neural network models and three automated model selection techniques. *Econometric Reviews* 35(8-10), 1753–1779.
- Carriero, A., T. E. Clark, M. G. Marcellino, and E. Mertens (2021). Addressing COVID-19 outliers in BVARs with stochastic volatility. Technical report, CEPR DP No. DP15964.
- Chernozhukov, V. and H. Hong (2003). An MCMC approach to classical estimation. *Journal of Econometrics* 115(2), 293–346.
- Cogley, T. and T. J. Sargent (2001). Evolving post-world war II US inflation dynamics. *NBER Macroeconomics Annual* 16, 331–373.
- Coibion, O., Y. Gorodnichenko, and M. Weber (2020). The cost of the covid-19 crisis: Lockdowns, macroeconomic expectations, and consumer spending. Technical report, National Bureau of Economic Research.

- Dijk, D. v., T. Teräsvirta, and P. H. Franses (2002). Smooth transition autoregressive models—a survey of recent developments. *Econometric Reviews* 21(1), 1–47.
- Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2021). The macroeconomics of epidemics. *The Review of Financial Studies* 34(11), 5149–5187.
- Granger, C. W. and T. Terasvirta (1993). *Modelling non-linear economic relationships*. Oxford University Press.
- Guerrieri, V., G. Lorenzoni, L. Straub, and I. Werning (2020). Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages? Technical report, National Bureau of Economic Research.
- Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press.
- Hornik, K., M. Stinchcombe, H. White, et al. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks* 2(5), 359–366.
- Jacquier, E., N. G. Polson, and P. E. Rossi (2004). Bayesian analysis of stochastic volatility models with fat-tails and correlated errors. *Journal of Econometrics* 122(1), 185–212.
- Jordà, Ò., S. R. Singh, and A. M. Taylor (2020). Longer-run economic consequences of pandemics? *The Review of Economics and Statistics*, 1–29.
- Koop, G., M. H. Pesaran, and S. M. Potter (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics* 74(1), 119–147.
- Lenza, M. and G. E. Primiceri (2020). How to estimate a VAR after March 2020. Technical report, National Bureau of Economic Research.
- Ludvigson, S. C., S. Ma, and S. Ng (2020). Covid19 and the macroeconomic effects of costly disasters. *NBER Working Paper* (w26987).
- MacKay, D. J. (1992). A practical Bayesian framework for backpropagation networks. *Neural Computation* 4(3), 448–472.
- Martin, V., S. Hurn, and D. Harris (2013). *Econometric modelling with time series: specification, estimation and testing*. Cambridge University Press.

- McAleer, M., M. C. Medeiros, and D. Slottje (2008). A neural network demand system with heteroskedastic errors. *Journal of Econometrics* 147(2), 359–371.
- Medeiros, M. C., M. McAleer, D. Slottje, V. Ramos, and J. Rey-Maqueira (2008). An alternative approach to estimating demand: Neural network regression with conditional volatility for high frequency air passenger arrivals. *Journal of Econometrics* 147(2), 372–383.
- Medeiros, M. C., T. Teräsvirta, and G. Rech (2006). Building neural network models for time series: a statistical approach. *Journal of Forecasting* 25(1), 49–75.
- Nocedal, J. and S. J. Wright (2006). Sequential quadratic programming. *Numerical optimization*, 529–562.
- Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies* 72(3), 821–852.
- Primiceri, G. E. and A. Tambalotti (2020). Macroeconomic forecasting in the time of covid-19. *Manuscript, Northwestern University*.
- Ramey, V. A. and S. Zubairy (2018). Government spending multipliers in good times and in bad: evidence from us historical data. *Journal of Political Economy* 126(2), 850–901.
- Rothman, P. (1998). Forecasting asymmetric unemployment rates. *The Review of Economics and Statistics* 80(1), 164–168.
- Schorfheide, F. and D. Song (2020). Real-time forecasting with a (standard) mixed-frequency VAR during a pandemic. *PIER Working Paper*.
- Skalin, J. and T. Teräsvirta (2002). Modeling asymmetries and moving equilibria in unemployment rates. *Macroeconomic Dynamics* 6, 202–241.
- Stock, J. H. and M. W. Watson (1996). Evidence on structural instability in macroeconomic time series relations. *Journal of Business & Economic Statistics* 14(1), 11–30.

- Stock, J. H. and M. W. Watson (1998). A comparison of linear and nonlinear univariate models for forecasting macroeconomic time series. Technical report, National Bureau of Economic Research.
- Stock, J. H. and M. W. Watson (2016). Core inflation and trend inflation. *The Review of Economics and Statistics* 98(4), 770–784.
- Swanson, N. R. and H. White (1995). A model-selection approach to assessing the information in the term structure using linear models and artificial neural networks. *Journal of Business & Economic Statistics* 13(3), 265–275.
- Swanson, N. R. and H. White (1997). A model selection approach to real-time macroeconomic forecasting using linear models and artificial neural networks. *The Review of Economics and Statistics* 79(4), 540–550.
- Teräsvirta, T., D. Van Dijk, and M. C. Medeiros (2005). Linear models, smooth transition autoregressions, and neural networks for forecasting macroeconomic time series: A re-examination. *International Journal of Forecasting* 21(4), 755–774.