

Impulse Response Analysis of Structural Nonlinear Time Series Models

Giovanni Ballarin*
University of Mannheim

November 6, 2023

First version: May 30, 2023

[Link to most recent version](#)

Abstract: Linear time series models are the workhorse of structural macroeconomic analysis. Yet, economic theory as well as data suggest that nonlinear and asymmetric effects might be key to understanding the potential effects of sudden economic changes. This paper proposes a new semi-nonparametric sieve approach to estimate impulse response functions of nonlinear time series within a general class of structural models. Using physical dependence conditions, I prove that a two-step procedure can flexibly accommodate nonlinear specifications, avoiding the choice of fixed parametric forms. Sieve impulse responses are proven to be consistent by deriving uniform estimation guarantees, while an iterative algorithm makes it straightforward to compute them in practice. Simulations show that the proposed semi-nonparametric approach provides insurance against misspecification at minor efficiency costs. In a US monetary policy application, I find that the sieve GDP response associated with a rate hike is, at its peak effects, 16% larger than that of a linear model. Finally, when studying interest rate uncertainty shocks, sieve responses imply up to 54% and 71% stronger contractionary effects on production and inflation, respectively.

*E-mail: giovanni.ballarin@gess.uni-mannheim.de. I thank Otilia Boldea, Timo Dimitriadis, Juan Carlos Escanciano, Lyudmila Grigoryeva, Klodiana Istrefi, Marina Khismatullina, So Jin Lee, Yuiching Li, Sarah Mouabbi, Andrey Ramirez, Christoph Rothe, Carsten Trenkler and Mengshan Xu, as well as the participants of the Econometrics Seminar at the University of Mannheim, the 2023 ENTER Jamboree, the 10th HKMETrics Workshop, the GSS Weekly Seminar at Tilburg University and the Internal Econometrics Seminar at Vrije Universiteit Amsterdam for their comments, suggestions and feedback.

1 Introduction

This paper presents a semi-nonparametric method to study the structural dynamic effects of unpredictable shocks in a class of nonlinear time series models.

Linear models are the foundation of economic structural time series modeling. The nature of linear models makes them especially tractable and apt at describing fundamental interactions and processes. For example, large classes of macroeconomic models in modern New Keynesian theory can be reduced to linear VARMA form via linearization techniques. This often justifies the application of the linear time series toolbox from a theoretical point of view. Concurrently, the work of [Sims \(1980\)](#) on VARs reinvigorated the strain of macroeconometric literature that seeks to study dynamic economic relationships. [Brockwell and Davis \(1991\)](#), [Hamilton \(1994b\)](#) and [Lütkepohl \(2005\)](#) provide detailed overviews of linear time series modeling and its developments. When the objects of interest are solely dynamic effects, the local projection (LP) approach of [Jordà \(2005\)](#) has also gained popularity as an alternative thanks to its flexibility and ease of implementation. LPs do not directly impose a linear model on the conditional distribution of the time series, but rather consist of linear lag regressions. Throughout this paper, the key dynamic effect under discussion will be the impulse response function (IRF), which is the common inference object of both linear VARMA and LP analyses.

Nonlinear methods seek to flexibly study the dependence structure between variables of interest by accommodating a potentially complex model structure. In recent years, research in nonlinear and asymmetric effects has grown, partly due to the increasing availability of data, making it feasible to estimate more elaborate models ([Fuleky, 2020](#)). From a macroeconomic perspective, one can imagine at least three broad categories of nonlinearities that may be important to study. Sign-dependence of impulse responses is a potential key factor in the evaluation of monetary policy, as the specific effects of an interest rate change might be mitigated if the central bank implements a rate drop rather than a rate hike, while some others might be enhanced ([Debortoli et al., 2020](#)). If impulse responses are size-sensitive, large shocks and small shocks can have vastly different economic impacts, meaning that the policymaker must account for nonlinear scaling in the intensity of an intervention ([Tenreyro and Thwaites, 2016](#)). Finally, if the researcher’s objective lies in studying exogenous changes impacting a variable that is nonlinear by definition, such as volatility indexes, any valid structural model should account for this feature.

The main contribution of this work is the development of an approach that allows estimating structural IRFs which can account for general nonlinear effects. This goal entails solving two related issues: first, structural identification of shocks, so that it is possible to give a valid economic interpretation to impulse responses; second, estimation of nonlinear functions in the setting of dependent data. In a linear setup, identification and estimation can be considered as distinct problems, but when working with nonlinear models these questions become intertwined. Without specific assumptions, nonlinear model classes are much too vast in terms of complexity: there are too many channels for any variable to affect any other. Disentangling such channels thus becomes impossible, and one cannot structurally interpret IRFs and dynamic effects such

as multipliers. This problem can be solved by being more precise about the classes of models one is willing to entertain. I consider the structural nonlinear framework originally proposed by [Gonçalves et al. \(2021\)](#), which involves selecting one variable to identify the structural shocks of interest, X_t , and treating it separately from all other series, a vector Y_t , included in the model. By imposing a few additional assumptions on the dependence structure of innovations, one is able to include general nonlinear effects of X_t and its lags onto Y_t . By further allowing the lags of Y_t to influence X_t , this setup permits nonlinear dynamics to propagate to all variables over time. The significant upside of this paradigm is that structural identification is built-in, instead of being treated as a separate step. The latter path is most often taken in the literature by implementing the generalized impulse response function (GIRF) proposed by [Koop et al. \(1996\)](#). [Kilian and Lütkepohl \(2017\)](#) have, however, highlighted that common linear identification strategies such as long-term and sign restrictions are generally impossible to impose in general nonlinear models, since closed-form expressions are not available but in a handful of special cases.

A weakness of the framework in [Gonçalves et al. \(2021\)](#) is that it requires choosing a specific functional form for the nonlinear components of the model, such as the negative-censoring map or a cubic map. These are used to tease out the sign and size effects of shocks.¹ Yet, correct prior knowledge of such terms is often unreasonable, especially in multivariate, multi-lag models. The natural way to avoid selecting a parametric nonlinear specification is to resort to semi-nonparametric techniques. Nonparametric time series methods have a long history in econometrics ([Härdle et al., 1997](#)), but until recently not much progress has been made in applying them to studying dynamic effects. Impulse response functions are objects that depend on the global properties of the model and, to be more precise, defining an IRF requires iterating shock perturbations over time. In a nonlinear model, the perturbation depends on the variables' state, so that one must consider the shock's effects across possible states. That is, different features of the nonlinear model such as level, slope, curvature must be evaluated over a range of values. Therefore, in this setting, an econometrician must provide error guarantees that are uniform over the variables' domain. In this work, I combine the uniform inference framework of [Chen and Christensen \(2015\)](#) with the structural nonlinear time series scheme discussed above. The general idea is to resort to semi-nonparametric series estimation and work in a physical dependence setup ([Wu, 2005](#)). On the one hand, I argue that physical dependence is a natural way of imposing assumptions that lead to estimable models, being more transparent than standard mixing conditions. On the other hand, the series approach makes it easy to estimate models with linear and nonlinear components of the type considered in this paper. It also provides well-developed theoretical results to study uncertainty. Under appropriate regularity assumptions, I show that a two-step semi-nonparametric series estimation procedure is able to consistently recover the structural model in a uniform sense. This result encompasses the generated regressors' problem, which arises in the second step due to the structural identification strategy. Lastly, I prove that the nonlinear impulse response function estimates obtained are themselves asymptotically consistent and, thanks to an iterative algorithm, straightforward to

¹The negative-censoring map applied to variable a is $a \mathbb{N} \max\{a, 0\}$.

compute in practice.

To validate the proposed methodology, I provide simulation evidence. The first set of results shows that, with realistic sample sizes, the efficiency costs of the semi-nonparametric procedure are small compared to correctly-specified parametric estimates. A second set of simulations demonstrates that whenever the nonlinear parametric model is mildly misspecified the large-sample bias is large, while for semi-nonparametric estimates it is negligible. Finally, I study how the IRFs computed with the new method compare with the ones from two previous empirical exercises. In a small, quarterly model of the US macroeconomy, I find that the parametric nonlinear and nonlinear appear to underestimate by intensity the GDP responses by 13% and 16%, respectively, after a large exogenous monetary policy shock. Moreover, sieve responses achieve maximum impact a year before their linear counterparts. Then, I evaluate the effects of interest rate uncertainty on US output, prices, and unemployment following [Istrefi and Mouabbi \(2018\)](#). In this exercise, the impact on industrial production of a one-deviation increase in uncertainty is approximately 54% stronger according to semi-nonparametric IRFs than the comparable linear specification. These findings suggest that structural impulse responses predicated on linear specifications might be appreciably underestimating shock effects.

Related Literature. Nonlinear models for dependent data have been extensively developed with the aim of analyzing diverse types of series, see e.g. the monographs of [Tong \(1990\)](#), [Fan and Yao \(2003\)](#), [Gao \(2007\)](#), [Tsay and Chen \(2018\)](#). [Teräsvirta et al. \(2010\)](#) provide a thorough discussion of nonlinear economic time series modeling, but, by only presenting the generalized IRF (GIRF) approach proposed by [Koop et al. \(1996\)](#), [Potter \(2000\)](#) and [Gourieroux and Jasiak \(2005\)](#), they do not explicitly address *structural* analysis.

Parametric nonlinear specifications are common prescriptions, for example, in time-varying models ([Auerbach and Gorodnichenko, 2012](#), [Caggiano et al., 2015](#)) and state-depend models ([Ramey and Zubairy, 2018](#)). They have been and are commonly used in time-homogeneous models. [Kilian and Vega \(2011\)](#) provide a structural analysis of the effects of GDP on oil price shocks and, in contrast to previous literature, find that asymmetries play a negligible role: they do this by including a negative-censoring transformation of the structural variable and testing for significance. [Caggiano et al. \(2017\)](#), [Pellegrino \(2021\)](#) and [Caggiano et al. \(2021\)](#) use interacted VAR models to estimate effects of uncertainty and monetary policy shocks. From a finance perspective, [Forni et al. \(2023a,b\)](#) study the economic effects of financial shocks. Their generalized VMA specification, which is based on that of [Debortoli et al. \(2020\)](#), sets that innovations be transformed with the quadratic map.² [Gambetti et al. \(2022\)](#) study news shocks asymmetries by imposing that news changes enter their autoregressive model with a pre-specified threshold function.

Extension of nonparametric methods to nonlinear time series have already been discussed in the recent literature. For example, [Kanazawa \(2020\)](#) proposed to use radial basis function neural networks to estimate a nonlinear time series model of the US macroeconomy. This work focuses

²I will discuss how their nonlinear model setup compares to the one I consider below.

on estimating the GIRF of [Koop et al. \(1996\)](#), with its structural limitations: productivity is assumed to be a fully exogenous variable. [Gourieroux and Lee \(2023\)](#) provide a framework for nonparametric kernel estimation and inference of IRFs via local projections. Yet, they primarily work in the one-dimensional case and only mention economic identification in multivariate setups from the perspective of linear VARs. The work possibly closest to the present paper seems to be that of [Lanne and Nyberg \(2023\)](#), who develop a nearest-neighbor approach to impulse responses estimation that builds on the local projection idea and the GIRF concept. These papers, save for [Gourieroux and Lee \(2023\)](#), do not fully develop an asymptotic theory for their estimators, which makes it hard to judge the econometric assumptions under which they are applicable.

Outline. The remainder of this paper is organized as follows. Section 2 provides the general framework for the structural model. Section 3 describes the two-step semi-nonparametric estimation strategy, provides a thorough treatment of physical dependence assumptions and reports the key uniform consistency guarantees. Section 4 is devoted to the discussion of nonlinear impulse response function computation, validity and consistency. In Section 5, I report simulation results that show the performance of the proposed method, while in Section 6 I discuss empirical applications. Finally, Section 7 concludes. All proofs and additional technical results, as well as secondary plots, can be found in Appendices B and C, respectively.

Notation. A (vector) random variable will be denoted in capital or Greek letters, e.g. Y_t or \mathbf{y}_t , while its realization will be in lowercase Latin letters, that is y_t . For a process $\{Y_t\}_{t \in \mathbb{Z}}$, we write $Y_{t:S} = (Y_t, Y_{t-1}, \dots, Y_{S-1}, Y_S)$, as well as $Y_{:t} = (\dots, Y_{t-2}, Y_{t-1}, Y_t)$ for the left-infinite history and $Y_{t:\infty} = (Y_t, Y_{t+1}, Y_{t+2}, \dots)$ for its right-infinite history. The same notation is also used for random variable realizations. For a matrix $A \in \mathbb{R}^{d \times d}$ where $d \geq 1$, $\|A\|$ is the spectral norm, $\|A\|_{\infty}$ is the supremum norm and $\|A\|_r$ for $0 < r < \infty$ is the r -operator norm. For a random vector or matrix, I will use $\| \cdot \|_{L^r}$ to denote the associated L^r norm.

2 Model Framework

In this section, I introduce the nonlinear time series model that will be considered throughout the paper. This model setup will be a generalization of the one developed in [Gonçalves et al. \(2021\)](#) by letting the form of nonlinear components to remain unspecified until estimation. The idea behind the partial structural identification scheme is simple: if Z_t is the full vector of time series of interest, one must choose one series, call it X_t , as the structural variable, and add specific assumption on its dynamic effects on the remaining series, vector Y_t . The central goal will be the estimation of the impulse responses of Y_t due to a shock in X_t .

2.1 Model

Let $Z_t = (X_t, Y_t)'$ where $X_t \in \mathbb{R}^d$ and $Y_t \in \mathbb{R}^{d_Y}$, and let $d = 1 + d_Y$ be the dimension of Z_t . The structural nonlinear data generating process has form

$$B_0 Z_t = b + B_1 Z_{t-1} + F_1 X_t, \quad (1)$$

where $b = (b_1, b_2)'$ and $Z_{t-1} = (Z_{t-1}^1, Z_{t-1}^2)'$ are partitioned accordingly. I assume that model (1) imposes a *linear* dependence of observables on Y_t and its lags, while series X_t can enter *nonlinearly*. That is, $B_1 = B_1 L + \dots + B_p L^{p-1}$ and $F_1 = F_0 + F_1 L + \dots + F_p L^p$ are linear and functional lag polynomials, respectively.³

Matrices F_0, \dots, F_p are functional in the sense that their entries consist of real univariate functions. The product between $F_p L^p$ and X_t is to be interpreted as functional evaluation, i.e.

$$F_p L^p X_t = \begin{pmatrix} f_{p,1}(X_t) \\ \vdots \\ f_{p,d}(X_t) \end{pmatrix},$$

where $f_{j,l} \in \Lambda$ for $j = 0, \dots, p$, $l = 1, \dots, d$, and Λ is a sufficiently regular function class.⁴ The modeling choice to remain within the autoregressive time series class with additive lag structure has two core advantages. First, it yields a straightforward generalization to classical linear models (Lütkepohl, 2005, Kilian and Lütkepohl, 2017). Second, it keeps semi-nonparametric estimation of nonlinear components feasible. Additivity in variables and lags means that the curse of dimensionality involved with multivariate nonparametric estimation is effectively mitigated (Fan and Yao, 2003).

Let the lag polynomials be given by

$$B_1 = \begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix}, \quad F_1 = \begin{pmatrix} 0 \\ F_{21} \end{pmatrix}.$$

This structural formulation means that the model equation for X_t is restricted to be linear in all regressors. It also implies that X_t does not depend contemporaneously on itself. Note that as long as $B_{12} = 0$, X_t still depends upon nonlinear functions of its own lags, which enter via lags of Y_t . Next, I impose that $B_0 \in \mathbb{R}^{d \times d}$ has the form

$$B_0 = \begin{pmatrix} 1 & 0 \\ B_{0,12} & B_{0,22} \end{pmatrix},$$

where $B_{0,22}$ is non-singular and normalized to have unit diagonal. The structural model is thus

³This is a minor abuse of notation compared to e.g. Lütkepohl (2005). The choice to use a matrix notation is due to the ease and clarity of writing a (multivariate) additive nonlinear model such as (1) in a manner consistent with standard linear VAR models. In cases where a real matrix $A \in \mathbb{R}^{d \times d}$ is multiplied with a conformable functional matrix F , I simply assume the natural product of a scalar times a function, e.g. $A_{ij} F_k$, where F_k is a function, returning a new real function.

⁴To fix ideas, one may think of Λ^q , the Hölder function class of smoothness $q \geq 0$ and domain $M \subset \mathbb{R}$. We shall make more precise assumptions regarding Λ in Section 3 when discussing model estimation.

given by

$$\begin{aligned} X_t &= b_1 + B_{12}\rho Lq Y_{t-1} + B_{11}\rho Lq X_{t-1} + \varepsilon_{1t}, \\ B_{0,22} Y_t &= b_2 + B_{22}\rho Lq Y_{t-1} + B_{21}\rho Lq X_{t-1} + B_{0,12} X_t + F_{21}\rho Lq X_{t-2}. \end{aligned}$$

Moreover, it follows that B_0^{-1} exists and has form

$$B_0^{-1} = \begin{pmatrix} 1 & 0 \\ B_0^{21} & B_0^{22} \end{pmatrix}.$$

The constraints on B_0 yield a structural identification assumption and require that X_t be pre-determined with respect to Y_t (Gonçalves et al., 2021). By introducing

$$\mu: B_0^{-1}b, \quad A\rho Lq: B_0^{-1}B\rho Lq \quad \text{and} \quad G\rho Lq: B_0^{-1}F\rho Lq,$$

one thus obtains

$$\begin{aligned} X_t &= \mu_1 + A_{12}\rho Lq Y_{t-1} + A_{11}\rho Lq X_{t-1} + \varepsilon_{1t}, \\ Y_t &= \mu_2 + A_{22}\rho Lq Y_{t-1} + A_{21}\rho Lq X_{t-1} + G_{21}\rho Lq X_t + B_0^{21} \varepsilon_{1t} + B_0^{22} \varepsilon_{2t}, \end{aligned} \quad (2)$$

or, equivalently,

$$Z_t = \mu + A\rho Lq Y_{t-1} + G\rho Lq X_t + u_t, \quad (3)$$

where $u_t = r u_{1t}, u_{2t}^s, u_{1t} + \varepsilon_{1t}$ and $u_{2t} = B_0^{21} \varepsilon_{1t} + B_0^{22} \varepsilon_{2t}$. Given the structure of B_0^{-1} , one can see that $A_{12}\rho Lq = B_{12}\rho Lq$, $A_{11}\rho Lq = B_{11}\rho Lq$ and $G_{11}\rho Lq = 0$. Importantly, one must also notice that $A_{12}\rho Lq$ and $G_{21}\rho Lq = B_0^{22}F_{21}\rho Lq$ might now be not properly identified without further assumptions. Since $A_{21}\rho Lq$ is not necessarily zero, linear effects of lags of X_t on Y_t can enter by means of both lag polynomials. To resolve this issue, I therefore assume that the functional polynomial $G_{21}\rho Lq$ contains, at lags greater than zero, only *nonlinear* components.⁵

Example 2.1. (A Simple Bivariate Model). To give a concrete example of (2), assume that one wants to model the effects of monetary policy shocks on U.S. GDP growth following Romer and Romer (2004). Then, let

$$\begin{aligned} X_t &= \varepsilon_{1t}, \\ Y_t &= \mu_2 + A_2 Y_{t-1} + G\rho X_{tq} + B_0^{21} \varepsilon_{1t} + \varepsilon_{2t}, \end{aligned}$$

where X_t are the policy shocks, which are assumed to be i.i.d., while Y_t is a macroeconomic variable whose responses the researcher is interested in, e.g. GDP growth or PCE inflation. This setup is very minimal, and I assume here, for the sake of simplicity, that endogeneity of ε_{2t} does not pose a problem. Then, the term $G\rho X_{tq} + B_0^{21} \varepsilon_{1t} = G\rho_{1tq} + B_0^{21} \varepsilon_{1t} = H\rho_{1tq}$ fully captures any contemporaneous effect of monetary policy shocks on Y_t . When $G\rho_{1tq} = 0$, $H\rho_{1tq}$

⁵When using a semi-nonparametric estimation strategy with B-splines, this will be feasible to implement numerically. When using wavelets, this also is a natural approach. In practice, however, some care must be taken to avoid constructing collinear regression matrices.

and the model are linear. If $G(p_{1t}) = \max\{0, \alpha p_{1t}\}$ for some $\alpha > 0$, function H is piece-wise linear: contractionary and expansionary shocks have, in general, different effects on Y_t , but shocks with the same sign have proportional impact. As a final example, if $G(p_{1t}) = \alpha p_{1t}^3$ then $H(p_{1t})$ is a third-degree polynomial, so that both sign and size of monetary policy shocks are fundamental determinants of Y_t 's impulse response function. In principle, to correctly quantify the repercussions of a specific monetary intervention a researcher must model all of these effects, unless they have a strong prior belief that either or both can be safely ignored. More complex nonlinear and asymmetric relations are also possible. A more robust strategy - as proposed in the present work - is to avoid choosing G (or H) as part of the model's specification, but rather to empirically estimate it jointly with all other coefficients.

Remark 2.1. (Constrained Models). The general approach of leaving $F(p_{1t})$ unconstrained is appealing when no precise economic intuition or information is available. However, there might be cases where the functional form of the nonlinear component is either partially known, or can be restricted. A simple restriction is that of a uniform functional over lags,

$$F(p_{1t}) = F + FL + FL^2 + \dots + FL^p.$$

This is a constraint effectively imposed by e.g. [Gonçalves et al. \(2021\)](#), [Kilian and Vega \(2011\)](#) and other references. They do this by fully specifying F , but nonparametric constraints may be desired, e.g. monotonicity. Constrained estimation of $F(p_{1t})$ is addressed in [Remark 3.2](#) below.

The system of equations in (2) provides the so-called *pseudo-reduced form* model. By design, one does not need to identify the model fully, meaning that fewer assumptions on Z_t and ϵ_t are needed to estimate the structural effects of p_{1t} on Y_t . This comes at the cost of not being able to simultaneously study structural effects with respect to p_{2t} . An associated problem is that, in general, $G_{21}(p_{1t})X_t$ is correlated with innovation u_{2t} through $B_0^{21} p_{1t}$. The main challenge to structural shock identification of p_{1t} thus lies in the fact that if $B_0^{21} > 0$ and $G_{21}(p_{1t}) > 0$, there is endogeneity in the equations for Y_t since X_t depends linearly on p_{1t} . [Gonçalves et al. \(2021\)](#) address the issue by proposing a two-step estimation procedure wherein one explicitly controls for p_{1t} by using regression residuals \hat{p}_t . In [Section 3](#) below, I show that this approach also allows for consistent semi-nonparametric estimation of structural impulse responses.

Remark 2.2. (Identification Schemes). [Forni et al. \(2023a,b\)](#) provide an alternative nonlinear structural identification framework to that of [Gonçalves et al. \(2021\)](#). Their approach was originally introduced in [Debortoli et al. \(2020\)](#) and is based on the VMA form of the time series. Using the current notation, suppose that the structural representation of Z_t is given by

$$Z_t = b + Q(p_{1t})F(p_{1t}) + B(p_{1t})\epsilon_t$$

where ϵ_t are independent structural shocks with zero mean and identity covariance, while p_{1t} identifies, e.g., financial innovations and shocks. $Q(p_{1t})$ and $B(p_{1t})$ are both linear lag polynomials and F is a nonlinear function to be specified by the researcher. Imposing some additional

assumptions, the reduced-form assumed by [Forni et al. \(2023a\)](#) is

$$Z_t = \mu + A(p)LqZ_t + Q_0 F(p^{-1}t) + B_0 \varepsilon_t, \quad (4)$$

where $F(p) = \chi^2$ in their baseline specification. [Forni et al. \(2023b\)](#) use an analogous model, while [Debortoli et al. \(2020\)](#) also consider more general setups where Q_0 is replaced by a general lag polynomial $D(p)Lq$. These kinds of structural assumptions are similar but not identical to the ones imposed in [Gonçalves et al. \(2021\)](#) and this paper. For (4) to overlap with (2), one must assume that X_t is exogenous and independently distributed, so that its level does not affect the mapping of ε_t through F . That is, (4) requires that only the *shocks* have nonlinear effects, not the structural variable itself. The upside of this approach is that one can directly and explicitly model asymmetry in the innovation process. The drawbacks are that, without a clear identification of a structural variable, one must fully identify B_0 . Moreover, function F remains to be specified a priori. Note, however, that if innovation sequence ε_t is observable, a generalization of the semi-nonparametric estimation results of this paper to the framework of [Debortoli et al. \(2020\)](#) would be straightforward.

I now state some preliminary assumptions for the model.

Assumption 1. ε_{1t} and ε_{2t} are mutually independent time series such that

$$\begin{matrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{matrix} \text{ i.i.d. } \begin{matrix} 0, & \begin{matrix} \frac{2}{1} & 0 \\ 0 & \Sigma_2 \end{matrix} \end{matrix}$$

where Σ_2 is a diagonal positive definite matrix.

Assumption 2. ε_t is strictly stationary, ergodic and such that $\sup_t E|\varepsilon_t| < \infty$.

Assumption 3. The roots of equation $\det(pI_d - A(p)Lq) = 0$ are outside the complex unit circle.

Assumption 1 follows [Gonçalves et al. \(2021\)](#). Assumption 2 is a high-level assumption on the properties of process ε_t and is common in the analysis of structural time series. Assumption 3 ensures that it is possible to invert lag polynomial $pI_d - A(p)Lq$ in order to define impulse responses, as done below. However, Assumption 2 and 3 will not be sufficient to make sure that (2) is estimable from data, and in Section 3 additional constraints on $A(p)Lq$ and $G(p)Lq$ will be required in order to apply semi-nonparametric estimation. Moreover, Assumption 2 is not easily interpretable: functional lag polynomial $G(p)Lq$ makes it impossible to reduce semi-structural equations (2) to an explicit infinite moving average form.

I will resolve both the former (sufficiency) and latter (interpretability) issue by using the nonlinear dynamic model framework outlined by [Pötscher and Prucha \(1997\)](#). It will allow introducing regularity assumptions on the dependence of Z_t which enable the derivation of consistency of impulse response estimates.

2.2 Structural Nonlinear Impulse Responses

Starting from pseudo-reduced equations (2), by letting $\Psi pLq = p/d = ApLqLq^{-1}$ one can further derive that

$$Z_t = \Theta pLq_t + \Gamma pLq X_t, \quad (5)$$

where

$$\mu: \Psi p1q \begin{matrix} \mu_1 \\ \mu_2 \end{matrix}, \quad \Theta pLq: \Psi pLq B_0^{-1}, \quad \text{and} \quad \Gamma pLq: \Psi pLq \begin{matrix} 0 \\ G_{21} pLq \end{matrix}.$$

To formally define impulse responses, it is useful to partition the polynomial ΘpLq according to

$$\Theta pLq: \Theta_{1pLq} \Theta_{2pLq},$$

where Θ_{1pLq} represents the first column of matrices in ΘpLq , and Θ_{2pLq} the remaining d_Y columns.

Given impulse $p \in \mathbb{R}$ at time t , define the shocked innovation process as $1s p q = s$ for $s \leq t$ and $1t p q = t$, as well as the shocked structural variable as $X_{sp} q = X_t$ for $s \leq t$ and $X_{sp} q = X_{sp} Z_{t-1, t-2, \dots, t-1, \dots, s} q$ for $s \neq t$. Further, let

$$\begin{aligned} Z_{t-h}: & \quad \Theta_{1pLq} 1t-h + \Theta_{2pLq} 2t-h + \Gamma pLq X_t, \\ Z_{t-hp} q: & \quad \Theta_{1pLq} 1t-hp q + \Theta_{2pLq} 2t-h + \Gamma pLq X_{tp} q, \end{aligned}$$

be the time- t baseline and shocked series, respectively. The unconditional impulse response is given by

$$\text{IRF}_{hp} q = \mathbb{E} r Z_{t-hp} q - Z_{t-hs}. \quad (6)$$

The difference between shock and baseline is clearly

$$\begin{aligned} Z_{t-hp} q - Z_{t-h} &= \Theta_{h,1} \Gamma pLq X_{tp} q - \Gamma pLq X_t \\ &= \Theta_{h,1} p \Gamma_0 X_{t-hp} q - \Gamma_0 X_{t-h} q - \dots - p \Gamma_h X_{tp} q - \Gamma_h X_{ts} q, \end{aligned}$$

therefore the unconditional IRF reduces to

$$\text{IRF}_{hp} q = \Theta_{h,1} \mathbb{E} r \Gamma_0 X_{t-hp} q - \Gamma_0 X_{t-hs} - \dots - \mathbb{E} r \Gamma_h X_{tp} q - \Gamma_h X_{ts}. \quad (7)$$

Notice that, in (7), while one can linearly separate expectations in the impulse response formula, terms $\mathbb{E} r \Gamma_j X_{t-jp} q - \Gamma_j X_{t-js}$ for $0 \leq j \leq h$ cannot be meaningfully simplified. Coefficients Γ_j are functional, therefore it is not possible to collect them across $X_{t-jp} q$ and X_{t-j} . Moreover, these expectations involve nonlinear functions of lags of X_t and cannot be computed explicitly. To address this issue, Section 4 provides an iterative procedure that makes computation of nonlinear impulse responses in (7) straightforward.

Remark 2.3. (Local Projection Approaches). As mentioned in the introduction, in recent years there has been growing interest in nonlinear IRF estimation procedures, and, accordingly, ways to generalize the LP framework. [Jordà \(2005\)](#) already suggested that nonlinear impulse responses

can, in principle, be directly estimated with local projections via the so-called *flexible local projection* approach. The flexible LP method relies on the Volterra expansion of time series to account for nonlinearities. There are multiple issues with this method. First, [Jordà \(2005\)](#) does not directly state how the validity of Volterra series implies the autoregressive form used in the LP regression. Second, the flexible LP proposal is fundamentally equivalent to adding polynomial factors to the linear regression specification. Thus, it is effectively a semi-nonparametric method, yet [Jordà \(2005\)](#) does not provide a theoretical analysis from this viewpoint. Moreover, no criterion or empirical rule-of-thumb for selecting the truncation order of the Volterra expansion are suggested, which becomes a key issue in practice. Due to these concerns, application of flexible LPs seems hard to justify from an econometric perspective.⁶ [Lanne and Nyberg \(2023\)](#) propose to nonparametrically recover the conditional mean function with a nearest-neighbor (k -NN) regression estimator. Their method is very flexible, but requires appropriately choosing the neighborhood size k and a distance measure for histories of realizations, and the authors do not theoretically address these issues. Very recently, [Gourieroux and Lee \(2023\)](#) have considered nonlinear IRF estimation with kernel-based methods by means of a novel conditional quantile representation of the process. They prove kernel LP estimators based on such representation are consistent, and that the direct estimator is asymptotically normal. The theory is developed only for the univariate case, with an autoregressive structure of lag order one, limiting the applicability of their procedure.

3 Estimation

Pseudo-reduced form model (2) can be compactly rewritten as

$$\begin{aligned} X_t &= \Pi_1^1 W_{1t} + u_{1t}, \\ Y_t &= \Pi_2^1 W_{2t} + u_{2t}, \end{aligned} \tag{8}$$

where

$$\begin{aligned} \Pi_1 &: \quad 1, A_{1,11}, \dots, A_{p,11}, A_{1,12}, \dots, A_{p,12}^1 \in \mathbb{R}^{1 \times pd}, \\ \Pi_2 &: \quad 2 \quad G_{1,21} \quad G_{p,21} \quad A_{1,22} \quad A_{p,22} \quad B_0^{21 \times 1}, \\ Z_{t \ 1:t \ p} &: \quad X_{t \ 1}, \dots, X_{t \ p}, Y_{t \ 1}^1, \dots, Y_{t \ p}^1 \in \mathbb{R}^{pd}, \\ W_{1t} &: \quad 1, Z_{t \ 1:t \ p}^1 \in \mathbb{R}^{1 \times pd}, \\ W_{2t} &: \quad 1, X_t, Z_{t \ 1:t \ p}^1 \in \mathbb{R}^{3 \times pd}. \end{aligned}$$

Additionally, let $W_1 = [W_{11}, \dots, W_{1n}]^1$ and $W_2 = [W_{21}, \dots, W_{2n}]^1$ be the design matrices for X_t and Y_t , respectively.

⁶Moreover, the complexity of estimating Volterra kernels grows exponentially with the kernel order, and thus more sophisticated approaches have been proposed to make estimation feasible, see e.g. [Sirotko-Sibirskaya et al. \(2020\)](#) and [Movahedifar and Dickhaus \(2023\)](#).

Two-step Estimation Procedure. Since W_{2t} is an infeasible vector of regressors, to estimate Π_2 one can use $\mathbb{W}_{2t} = \rho_1, X_t, Z_t^1, \dots, Z_t^p, \rho_1 q^1$, which now contains generated regressors in the form of residual ρ_{1t} . This approach is an adaptation of the two-step procedure put forth by [Gonçalves et al. \(2021\)](#), where I allow for semi-nonparametric estimation:

1. Regress X_t onto W_{1t} to get estimate $\hat{\Pi}_1$ and compute residuals $\rho_{1t} = X_t - \hat{\Pi}_1' W_{1t}$.
2. Fit Y_t using \mathbb{W}_{2t} to get estimate $\hat{\Pi}_2$. Since $G_{1,21}, \dots, G_{p,21}$ contain functional parameters, a semi-nonparametric estimation method is required.
3. Compute coefficients in Θ_{pLq} and \mathbb{P}_{pLq} from $\hat{\Pi}_1$ and $\hat{\Pi}_2$.
4. Consider the two paths with time t shocks ϵ_t versus ϵ_t to construct the unconditional IRF, average over histories as well as future shocks by using the algorithm detailed in [Proposition 4.1](#) or [Proposition 4.2](#).

[Gonçalves et al. \(2021\)](#) only allow for pre-determined nonlinear transforms of X_t . The core contribution of this paper is allowing $G_{1,21}, \dots, G_{p,21}$ to be estimated in a nonparametric way. I focus on series estimation in order to build on the extensive theory available in the setting of dependent data ([Chen, 2013](#), [Chen and Christensen, 2015](#)). This further adds to the framework of [Gonçalves et al. \(2021\)](#), as their regularity assumptions are stated only as preconditions for a uniform LLN to hold and are not easy to interpret.

Remark 3.1. (Alternative Estimation Approaches). One does not need to limit estimation of the nonlinear functional parameters $G_{1,21}, \dots, G_{p,21}$ to series-type estimators. The literature on nonparametric regression is mature, and thus kernel ([Tsybakov, 2009](#)), nearest-neighbor ([Li and Racine, 2009](#)), partitioning ([Cattaneo et al., 2020](#)) and deep neural network ([Farrell et al., 2021](#)) estimators are all potentially valid alternatives. For example, [Huang et al. \(2014\)](#) use kernel regression to perform density estimation and regression under physical dependence. However, thanks to both availability of uniform inference results (see also [Belloni et al. 2015](#)) and ease of implementation, series methods stand out as a choice for semi-nonparametric time series estimation and nonlinear impulse response computation.

In the remainder of this section, I first introduce the semi-nonparametric series estimation strategy in detail. Then, I outline the core assumptions of the sieve setup. Special focus is put on the dependence structure of the data: rather than directly assuming α -mixing as in [Chen and Christensen \(2015\)](#), I shall consider physical dependence assumptions ([Wu, 2005](#)) to provide transparent conditions on the model itself that, if satisfied, ensure consistency. I prove that the proposed two-step semi-nonparametric procedure is uniformly consistent under physical dependence assumptions. These assumptions can be imposed directly on the model, and, as such, may be empirically checked, if necessary. The uniform asymptotic guarantees are first stated for the infeasible estimator involving true innovations ϵ_{1t} and later extended to encompass feasible estimator $\hat{\Pi}_2$.

3.1 Semi-nonparametric Series Estimation

Starting from (8), one can introduce the i th-row coefficient matrices

$$\begin{aligned} G_i^{21} &= rG_{1,21} & G_{\rho,21} s_j, \\ A_i^{22} &= rA_{1,22} & A_{\rho,22} s_j, \end{aligned}$$

and B_{0i}^{21} accordingly. Consider now the regression problem for each individual component of Y_t ,

$$Y_{it} = G_i^{21} X_{t:t-\rho} + A_i^{22} Y_{t-1:t-\rho} + B_{0i}^{21} + u_{2it},$$

where $X_{t:t-\rho} = (X_t, \dots, X_{t-\rho})'$ and $i = 1, \dots, d_Y$. For simplicity of notation, I suppress intercept B_{0i}^{21} , but this is without loss of generality. Since G_i^{21} consists of $1-\rho$ functional coefficients and A_i^{22} can be segmented into ρ row vectors of length d_Y , it is possible to rewrite the above as

$$Y_{it} = \sum_{j=0}^{\rho} g_{ij}^{21} X_{t-j} + \sum_{j=1}^{\rho} A_{ij}^{22} Y_{t-j} + B_{0i}^{21} + u_{2it}. \quad (9)$$

I will use $\beta_{2,i} = (G_i^{21}, A_i^{22}, B_{0i}^{21})'$ to identify the vector of coefficients in the equation for the i th component of Y_t . From (9), Π_2^1 can be decomposed in d_Y rows of coefficients, i.e.

$$\begin{aligned} Y_{1t} &= \beta_{2,1} \\ \vdots & \\ Y_{d_Y t} &= \beta_{2,d_Y} \end{aligned} \quad W_{2t} + u_{2t}$$

and one can treat each equation separately.

A semi-nonparametric series estimator for (9) is built on the idea that, if functions g_{ij}^{21} belong to an appropriate functional space, one can construct a growing collection of sets of basis functions – called a *sieve* – which, linearly combined, progressively approximate g_{ij}^{21} . That is, one can reduce the infinite dimensional problem of estimating the functional coefficients in $\beta_{2,i}$ to a linear regression problem. Although (9) features a sum of possibly nonlinear functions in X_{t-j} and Y_{t-j} , as well as linear functions of Y_{t-j} and 1_t , constructing a sieve is straightforward.⁷

Assume that $g_{ij}^{21} \in \Lambda$, where Λ is a sufficiently regular function class to be specified in the following, and let \mathbf{B}_Λ be a sieve for Λ . Let b_1, \dots, b_p be the collection of $p \geq 1$ sieve basis functions from \mathbf{B}_Λ and define

$$\begin{aligned} b(x) &= (b_1(x), \dots, b_p(x))', \\ \mathbf{B} &= (b(X_{1:1-\rho}), \dots, b(X_{n:n-\rho})). \end{aligned}$$

The sieve space for $\beta_{2,i}$ is $\mathbf{B}_\Lambda^{1-p} \in \mathbb{R}^{1-pd_Y}$, where here \mathbb{R} identifies the space of linear functions. Since the nonparametric components of Π_2^1 are linearly separable in the lag dimension, I take

⁷See [Chen \(2007\)](#) for a comprehensive exposition of sieve estimation. [Chen and Shen \(1998\)](#) and [Chen \(2013\)](#) also provide additional examples of partially linear semi-nonparametric models under dependence.

\mathbf{B}_Λ^{1-p} to be a direct product of sieve spaces.⁸ Importantly, the same sieve can be used for all components of Y_t , as I assume the specification of the model does not change across i .

Let $b_{\cdot,1K}, \dots, b_{\cdot,KK}$ be the sieve basis in $\mathbf{B}_\Lambda^{1-p} \subset \mathbb{R}^{1-pd_Y}$ which, for $\ell \neq 1$ and K $p = p_1 + \dots + p_d_Y$, is given by

$$\begin{aligned} b_{\cdot,1K} &= \mathbb{1}_{W_{2t} \leq q}, & b_{\cdot,1} &= X_{t,q}, \\ & \vdots & & \\ b_{\cdot,p_1 q} &= \mathbb{1}_{W_{2t} \leq q}, & b_{\cdot,p_1} &= X_{t,p_1 q}, \\ b_{\cdot,p_1 q} &= \mathbb{1}_{W_{2t} \leq q}, & Y_{t,1,1}, & \\ & \vdots & & \\ b_{\cdot,p_K q} &= \mathbb{1}_{W_{2t} \leq q}, & Y_{t,p,d_Y}, & \\ b_{\cdot,KK} &= \mathbb{1}_{W_{2t} \leq q}, & & \mathbb{1}_{t^c}. \end{aligned}$$

where q fixes the size of the nonparametric component of the sieve. Note that K , the overall size of the sieve, grows linearly in q , which itself controls the effective dimension of the nonparametric component of the sieve, $b_{\cdot,1}, \dots, b_{\cdot,K}$. In all theoretical results, I will focus on the growth rate of K rather than q , as asymptotically they differ at most by a constant multiplicative factor.

The regression equation for $Y_{2,i}$ is

$$Y_{2,i} = \sum_{j=1}^K b_{2,i,j} W_{2t} + u_{2i},$$

where $Y_{2,i} = (Y_{2i1}, \dots, Y_{2in})'$ and $u_{2i} = (u_{2i1}, \dots, u_{2in})'$. The estimation target is the conditional expectation $\mathbb{E} Y_{2,i} | W_{2t} = w$ under the assumption $\mathbb{E} u_{2i} | W_{2t} = 0$. By introducing

$$\begin{aligned} b^K &= (b_{\cdot,1K}, \dots, b_{\cdot,KK}), \\ B &= (b^K | W_{21}, \dots, b^K | W_{2n}). \end{aligned}$$

the *infeasible least squares series estimator* $\hat{\rho}_{2,i}$ is given by

$$\hat{\rho}_{2,i} = b^K | w | B' B^{-1} B^K Y_{2,i}.$$

Similarly, consider the feasible series regression matrices

$$\begin{aligned} \hat{b}^K &= (\hat{b}_{\cdot,1K}, \dots, \hat{b}_{\cdot,KK}), \\ \hat{B} &= (\hat{b}^K | W_{21}, \dots, \hat{b}^K | W_{2n}). \end{aligned}$$

⁸It is not necessary to consider the more general case of tensor products of 1D sieve functions, as it would be the case for a general $p_1 + \dots + p_d_Y$ -dimensional function $G_t^{21}(X_t, X_{t-1}, \dots, X_{t-p})$. As previously discussed, the additive structure avoids the curse of dimensionality which in nonlinear time series modeling is often a primary concern when working with moderate sample sizes (Fan and Yao, 2003).

Thus, the *feasible least squares series estimator* is

$$\hat{\beta}_{2,i} = (W_i' W_i)^{-1} W_i' Y_i.$$

Given that the semi-nonparametric estimation problem is the same across i , to further streamline notation, where it does not lead to confusion I will let β_2 be a generic coefficient vector belonging to \mathbb{R}^{d_Y} , as well as define β_2 , Y and W_2 accordingly.

Remark 3.2. (Constrained Sieve Estimation). The idea of constrained estimation was only briefly touched upon in Remark 2.1. In fully parametric nonlinear models, constraints are often imposed out of necessity or simplicity. If, say, $G_{1,21}$ is constituted only of the negative-censoring map, it is unclear why $G_{2,21}$ would be constituted instead of quadratic or cubic functions, for example. That is, *specific* parametric assumptions can be either unreasonable or hard to justify in practice.⁹ Yet, constrained semi-nonparametric estimation might be desirable at times.

If the shape of the regression function is to be constrained to ensure e.g. non-negativity, monotonicity or convexity, [Chen \(2007\)](#) gives examples of shape-preserving sieves, like cardinal B-spline wavelets. Constraints on a generic sieve can also be imposed at estimation time. For example, for simplicity suppose $d_Y = 1$ and $\rho = 2$, and that one wants to impose $G_{1,21} = G_{2,21}$. The constrained sieve estimator then solves

$$\min_{\beta} \sum_{t=1}^n (Y_t - \beta' W_{2t})^2 \quad \text{subject to} \quad \beta \geq 0, \beta' W_{2t} \geq 0 \quad \forall t.$$

Analysis of restricted or constrained estimators, however, is still a challenging problem in non-parametric theory, c.f. [Horowitz and Lee \(2017\)](#), [Freyberger and Reeves \(2018\)](#), [Chetverikov et al. \(2018\)](#). Misspecification in particular is complex to address. Accordingly, I will not be imposing any specific restrictions on the nonlinear functions in Π_2 outside the ones necessary to derive uniform asymptotic theory.

Spline Sieve. The B-spline sieve $\text{BSplp}(r, \mathbf{m})$ of degree $r \geq 1$ over $[0, 1]$ can be constructed using the Cox-de Boor recursion formula. Alternatively, an equivalent way of constructing the spline sieve is as follows. For simplicity, let $d_Y = 1$ and let $0 = m_1 < \dots < m_{r+1} = 1$ be a set of knots. Then

$$b_{\text{spline}}(x) = \sum_{j=1}^r \beta_j B_j(x), \quad \text{where } B_j(x) = \prod_{k=1}^j (x - m_{k-1}) \prod_{k=j+1}^r (m_k - x).$$

The resulting spline sieve is piece-wise polynomial of degree r . Moreover, notice that in practice the spline sieve already contains a linear and constant term, so care must be taken to avoid collinearity (for example, by not including an additional intercept and linear term in X_t in the series regression).

⁹For more precise examples and a more in-depth discussion, see Section 2.1 of [Chen \(2013\)](#).

3.2 Distributional and Sieve Assumptions

To develop the asymptotic uniform consistency theory, I rely on the general theoretical framework established by [Chen and Christensen \(2015\)](#). Basic distributional and sieve assumptions can be carried over from their setup mostly unchanged.

Assumption 4. (i) $\{u_{tPZ}\}$ are such that $\{u_{tPZ}\}$ i.i.d. ρ_0, Σ_0 , (ii) $\{u_{1tPZ}\}$ and $\{u_{2tPZ}\}$ are mutually independent, (iii) $\{u_{tPZ}\} \in E$ for all $t \in \mathbb{Z}$ where $E \in \mathbb{R}^{d_Y}$ is compact, convex and has nonempty interior.

Assumption 5. (i) $\{Z_t\}$ is a strictly stationary and ergodic time series, (ii) $\{X_t\} \in X$ for all $t \in \mathbb{Z}$ where $X \in \mathbb{R}$ is compact, convex and has nonempty interior, (iii) $\{Y_t\} \in Y$ for all $t \in \mathbb{Z}$ where $Y \in \mathbb{R}^{d_Y}$ is compact, convex and has nonempty interior.

Assumptions 4(i)-(ii) are a repetition of Assumption 1. As W_{2t} depends only on $X_{t:t-p}$, $Y_{t-1:t-p}$, and u_{1t} , Assumption 1 also implies that entries of u_{2t} are independent of W_{2t} , so that $\mathbb{E}u_{2t} | W_{2t} = 0$.¹⁰ Assumption 5(i) also follows from Assumption 2. However, thanks to the results derived in Section 3.3, below I will impose more primitive conditions on the model for Z_t that allow to recover 5(i). Assumption 4(iii) and Assumptions 5(ii)-(iii) imply that X_t , Y_t , as well as u_{tPZ} are bounded random variables. In (semi-)nonparametric estimation, imposing that X_t be bounded almost surely is a standard assumption. Since lags of Y_t and innovations u_{tPZ} contribute linearly to all components of Z_t , it follows that they too must be bounded. Unbounded regressors are more complex to handle when working in the nonparametric setting. Generalization from bounded to unbounded domains under dependence has already been discussed by e.g. [Fan and Yao \(2003\)](#). [Chen and Christensen \(2015\)](#) also allow for an expanding support by using weighted sieves. I leave this extension for future work.

It is, however, important to highlight that bounded support assumptions are relatively uncommon in time series econometrics. This is clear when considering the extensive literature available on linear models such as, e.g., state-space, VARIMA and dynamic factor models ([Hamilton, 1994a](#), [Lütkepohl, 2005](#), [Kilian and Lütkepohl, 2017](#), [Stock and Watson, 2016](#)). Avoiding Assumptions 4(iii) and 5(iii) can possibly be achieved with a change in the model's equations – so that, for example, lags of Y_t only effect X_t either via bounded functions or not at all – so I do not discuss this approach here. In practice, Assumptions 4(ii) and 5(ii)-(iii) are not excessively restrictive, as most credibly stationary economic series often have reasonable implicit (e.g. inflation) or explicit bounds (e.g. employment rate).¹¹

Let $\mathcal{F}_t = \sigma(\dots, u_{1t-1}, u_{2t-1}, Y_{t-1}, \dots, u_{1t}, u_{2t}, Y_t)$ be the natural filtration defined up to time t . Thanks to Assumptions 4 and 5 the following moment requirements hold trivially.

¹⁰Moreover, for any given i , the sequence $\{u_{2itPZ}\}$ is i.i.d. over time index t .

¹¹This is not true, of course, when modeling extreme events like natural disasters, wars or financial crises. To study these types of series, however, researchers often apply specialized models. Thinking in this direction, a future development could be to extend the framework presented here to allow for innovations with unbounded support.

Assumption 6. (i) $\mathbb{E}r|u_{2it}|^2|F_{t-1}$ is uniformly bounded for all $t \in \mathbb{Z}$ almost surely, (ii) $\mathbb{E}r|u_{2it}|^2 \leq \mathcal{S}$ for some $\mathcal{S} > 0$, (iii) $\mathbb{E}r|Y_{it}|^2 \leq \mathcal{S}$ is uniformly bounded for all $t \in \mathbb{Z}$ almost surely, and (iv) $\mathbb{E}r|Y_{it}^2|F_{t-1} \leq \mathcal{S}$ for any $\mathcal{S} > 0$.

Now let $W_2 \in \mathbb{R}^d$ be the domain of W_{2t} . By assumption, W_2 is compact and convex and is given by the direct product

$$W_2 = X^{1-p} \times Y^p \times E_1,$$

where E_1 is the domain of structural innovations u_{1t} i.e. $E = E_1 \times E_2$.

Assumption 7. Define $\kappa_{K,n} := \sup_{W \in W_2} b^K(W)$ and

$$\kappa_{K,n} := \inf_{W \in W_2} \mathbb{E}r|b^K(W_{2t}) - b^K(W_{2t}^*)|^2.$$

It holds:

- (i) There exist $\kappa_1, \kappa_2 \neq 0$ s.t. $\sup_{W \in W_2} b^K(W) \leq \kappa_1 n^{-1} \kappa_2$.
- (ii) There exist $\kappa_1^{-1} \neq 0, \kappa_2 > 0$ s.t. $\kappa_{K,n} \geq \kappa_1 n^{-1} \kappa_2$.
- (iii) $\inf_{W \in W_2} \mathbb{E}r|b^K(W_{2t}) - b^K(W_{2t}^*)|^2 > 0$ for all K and n .

Assumption 7 provides mild regularity conditions on the families of sieves that can be used for the series estimator. More generally, letting W_2 be compact and rectangular makes Assumption 7 hold for commonly used basis functions (Chen and Christensen, 2015).¹² In particular, Assumption 7(i) holds with $\kappa_1 = 0$ since the domain is fixed over the sample size.

In the proofs, it is useful to consider the orthonormalized sieve basis. Let

$$\begin{aligned} \mathcal{B}^K(W) &:= \mathbb{E}r|b^K(W_{2t}) - b^K(W_{2t}^*)|^2 \mathcal{B}^K(W), \\ \mathcal{B} &:= \{\mathcal{B}^K(W_{21}), \dots, \mathcal{B}^K(W_{2n})\} \end{aligned}$$

be the orthonormalized vector of basis functions and the orthonormalized regression matrix, respectively.

Assumption 8. It holds that $\mathcal{B}^1 \mathcal{B} \{n\} \rightarrow 0$ in probability.

Assumption 8 is the key assumption imposed by Chen and Christensen (2015) to derive uniform convergence rates under dependence. They prove that if $\{W_{2t}\}_{t \in \mathbb{Z}}$ is strictly stationary and α -mixing – with either geometric or algebraic decay, depending on the sieve family of interest – then Assumption 8 holds. Let $(\Omega, \mathcal{Q}, \mathbb{P})$ be the underlying probability space and define

$$\rho(A, B) := \frac{1}{2} \sup_{\mathcal{P}_I, \mathcal{Q}_J} \left| \mathbb{P}(A \cap B) - \mathbb{P}(A)\mathbb{P}(B) \right|$$

where A, B are two σ -algebras, $\mathcal{P}_I \in \mathcal{A}$, $\mathcal{Q}_J \in \mathcal{B}$ and the supremum is taken over all finite partitions of Ω . The h -th α -mixing coefficient of process $\{W_{2t}\}_{t \in \mathbb{Z}}$ is defined as

$$\rho(h) := \sup_t \rho(\mathcal{P}_{\{1, \dots, W_{2t-1}, W_{2t}\}}, \mathcal{Q}_{\{W_{2t+h}, W_{2t+h+1}, \dots\}})$$

¹²See Chen (2007), Belloni et al. (2015) for additional discussion and examples of sieve families.

and W_{2t} is said to be *geometric* or *exponential β -mixing* if $\rho(h) \propto \exp(-\beta h)$ for some $\beta > 0$ and $\beta > 0$. The main issue with mixing assumptions is that they are, in general, hard to compute and evaluate. Therefore, especially in nonlinear systems, assuming that $\rho(h)$ decays exponentially over h imposes very high-level assumptions on the model. There are, however, many setups in which it is known that β -mixing holds under primitive assumptions (see [Chen \(2013\)](#) for examples).

In the next subsection, I will argue that using a different concept of dependence - one rooted in a physical understanding of the underlying stochastic process - leads to imposing transparent assumptions on the model's structure.

3.3 Physical Dependence Conditions

Consider now a *non-structural model* of the form

$$Z_t = \Phi Z_{t-1} + \epsilon_t. \quad (10)$$

This is a generalization of semi-reduced model (3) where linear and nonlinear components are absorbed into one functional term and B_0 is the identity matrix.¹³ Indeed, note that models of the form $Z_t = \Phi Z_{t-1}, \dots, Z_{t-p} + \epsilon_t$ can be rewritten as (10) using a companion formulation. If ϵ_t is stochastic, (10) defines a causal nonlinear stochastic process. More generally, it defines a nonlinear difference equation and an associated dynamical system driven by ϵ_t . Throughout this subsection, I shall assume that $Z_t \in \mathbb{R}^{d_z}$ as well as $\epsilon_t \in \mathbb{R}^{d_z}$.

Relying on the framework of [Pötscher and Prucha \(1997\)](#), I now introduce explicit conditions that allow to control dependence in nonlinear models by using the toolbox of physical dependence measures developed by [Wu \(2005, 2011\)](#). The aim is to use a dynamical system perspective to address the question of imposing meaningful assumptions on nonlinear dynamic models. This makes it possible to give more primitive conditions under which one can actually estimate (8) in a semi-nonparametric way.

Stability. An important concept for dynamical system theory is that of stability. Stability turns out to play a key role in constructing valid asymptotic theory, as it is well understood in linear models. It is also fundamental in developing the approximation theory of nonlinear stochastic systems.

Example 3.1. (Linear System). As a motivating example, first consider the linear system

$$Z_t = BZ_{t-1} + \epsilon_t$$

where we may assume that $\epsilon_t \in \mathbb{R}^{d_z}$, $\epsilon_t \in \mathbb{R}^{d_z}$, is a sequence of i.i.d. innovations.¹⁴ It is well-known that this system is stable if and only if the largest eigenvalue of B is strictly less than

¹³In this specific subsection, shock identification does not play a role and, as such, one can safely ignore B_0 .

¹⁴One could alternatively think of the case of a deterministic input, setting $\epsilon_t = P_t \delta_t$ where $P_t \delta_t$ is a Dirac density on the deterministic sequence $\delta_t \in \mathbb{R}^{d_z}$.

one in absolute value (Lütkepohl, 2005). For a higher order linear system, $Z_t = BpLqZ_{t-1} + t$ where $BpLq = \begin{bmatrix} B_1 & B_2 & \dots & B_p \end{bmatrix}$, stability holds if and only if $|\max_i \rho(Bq)| < 1$ where

$$\mathbf{B} : \begin{bmatrix} B_1 & B_2 & & B_p \\ I_{d_Z} & 0 & & 0 \\ 0 & I_{d_Z} & & 0 \\ \vdots & \vdots & & \vdots \\ 0 & & I_{d_Z} & 0 \end{bmatrix}$$

is the companion matrix.

Extending the notion of stability from linear to nonlinear systems requires some care. Pötscher and Prucha (1997) derived generic conditions allowing to formally extend stability to nonlinear models by first analyzing *contractive* systems.

Definition 3.1 (Contractive System). Let $Z_t \in \mathbb{R}^{d_Z}$, $t \in E$, where $t \in \mathbb{N}$ is generated according to

$$Z_t = GpZ_{t-1} + tq.$$

The system is contractive if for all $pz, z^1 \in \mathbb{R}^{d_Z}$ and $pe, e^1 \in E$

$$Gpz, q - Gpz^1, q \leq C_Z \|z - z^1\| + C_e \|e - e^1\|$$

holds with Lipschitz constants $0 \leq C_Z < 1$ and $0 \leq C_e < \delta$.

Sufficient conditions to establish contractivity are

$$\sup_i \text{stack}_{i-1}^{d_Z} \left(\frac{BG}{BZ} pz^i, e^i \right) \leq 1 \quad (11)$$

and

$$\frac{BG}{B} < \delta, \quad (12)$$

where the stacking operator $\text{stack}_{i-1}^{d_Z} r$ progressively stacks the rows, indexed by i , of its argument (which can be changing with i) into a matrix. Values $pz^i, e^i \in \mathbb{R}^{d_Z} \times E$ change with index i as the above condition is derived using the mean value theorem, therefore it is necessary to consider a different set of values for each component of Z_t .

It is easy to see, as Pötscher and Prucha (1997) point out, that contractivity is often a too strong condition to be imposed. Indeed, even in the simple case of a scalar AR(2) model $Z_t = b_1 Z_{t-1} + b_2 Z_{t-2} + t$, regardless of the values of $b_1, b_2 \in \mathbb{R}$ contractivity is violated. This is due to the fact that in a linear AR(2) model studying contractivity reduces to checking $\|\mathbf{B}\| < 1$ instead of $|\max_i \rho(Bq)| < 1$, and the former is a stronger condition than the latter.¹⁵ One can weaken contractivity – which must hold for G as a map from Z_{t-1} to Z_t – to the idea of *eventual*

¹⁵See Pötscher and Prucha (1997), pp.68-69.

contractivity. That is, intuitively, one can impose conditions on the dependence of Z_{t+h} on Z_t for $h \geq 1$ sufficiently large. To do this formally, I first introduce the definition of system map iterates.

Definition 3.2 (System Map Iterates). *Let $Z_t \in \mathbb{R}^{d_z}$, $\mathbf{e}_t \in \mathbb{R}^{d_e}$ where $\{Z_t, \mathbf{e}_t\}_{t \in \mathbb{N}}$ is generated from a sequence $\{\mathbf{u}_t\}_{t \in \mathbb{N}}$ according to*

$$Z_t = G^h Z_{t-1}, \mathbf{e}_t.$$

The h -order system map iterate is defined to be

$$G^{ph} : \mathbb{R}^{d_z} \times \mathbb{R}^{d_e} \rightarrow \mathbb{R}^{d_z} \times \mathbb{R}^{d_e} \\ G^{ph}(Z_t, \mathbf{e}_{t-1}, \dots, \mathbf{e}_{t-h}) = G^p(Z_t, \mathbf{e}_{t-1}, \dots, \mathbf{e}_{t-h}), \mathbf{e}_{t-h}$$

where \circ signifies function composition and $G^{0q} Z_t = Z_t$.

To shorten notation, in place of $G^{ph}(Z_t, \mathbf{e}_{t-1}, \dots, \mathbf{e}_{t-h})$ I shall use $G^{ph}(Z_t, \mathbf{e}_{t-1:t-h})$. Additionally, for $1 \leq j \leq h$, the partial derivative

$$\frac{\partial G^{ph}}{\partial \mathbf{e}_j}$$

for some fixed h is to be intended with respect to \mathbf{e}_j , the j -th entry of the input sequence. This derivative does not depend on the time index since by assumption G is time-invariant and so is G^{ph} .

Taking again the linear autoregressive model as an example,

$$Z_{t+h} = G^{ph}(Z_t, \mathbf{e}_{t-1:t-h}) = B_1^h Z_t + \sum_{i=0}^{h-1} B_1^i \mathbf{e}_{t-h+i}$$

since $G^p Z_t = B_1 Z_t$. If B_1 determines a stable system, then $B_1^h \rightarrow 0$ as $h \rightarrow \infty$ since G^h converges to zero, and therefore $B_1^h \approx C_Z^{-1}$ for h sufficiently large. It is thus possible to use system map iterates to define stability for higher-order nonlinear systems.

Definition 3.3 (Stable System). *Let $Z_t \in \mathbb{R}^{d_z}$, $\mathbf{e}_t \in \mathbb{R}^{d_e}$, where $\{Z_t, \mathbf{e}_t\}_{t \in \mathbb{N}}$ is generated according to the system*

$$Z_t = G^h Z_{t-1}, \mathbf{e}_t.$$

The system is stable if there exists $h \geq 1$ such that for all $\mathbf{z}^1, \mathbf{z}^2 \in \mathbb{R}^{d_z}$ and $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_h \in \mathbb{R}^{d_e}$,

$$G^{ph}(\mathbf{z}^1, \mathbf{e}_{1:h}) - G^{ph}(\mathbf{z}^2, \mathbf{e}_{1:h}) \leq C_Z \|\mathbf{z}^1 - \mathbf{z}^2\| + C \|\mathbf{e}_{1:h}\|$$

holds with Lipschitz constants $0 \leq C_Z \leq 1$ and $0 \leq C < \infty$.

It is important to remember that this definition encompasses systems with an arbitrary finite autoregressive structure, i.e., $Z_t = G^p Z_{t-p}, \dots, Z_{t-1}, \mathbf{e}_t$ for $p \geq 1$, thanks to the companion formulation of the process. An explicit stability condition, similar to that discussed above for

contractivity, can be derived by means of the mean value theorem. Indeed, for a system to be stable it is sufficient that, at iterate h ,

$$\sup_{\#} \text{stack}_{j=1}^{d_Z} \frac{B G^{p,h} q}{B Z} p_{Z^j, e_{1:h}^j} q \quad Z^j P Z, e_{1:h}^j P \quad i \quad i=1 \quad E \quad + \quad 1 \quad (13)$$

and

$$\sup_{\#} \frac{B G^{p,h} q}{B_j} p_{Z, e_{1:h}} q \quad Z P Z, e_{1:h} P \quad i=1 \quad E \quad + \quad \mathfrak{B}, \quad j = 1, \dots, h. \quad (14)$$

Pötscher and Prucha (1997) have used conditions (11)-(12) and (13)-(14) as basis for uniform laws of large numbers and central limit theorems for L^r -approximable and near epoch dependent processes.

Physical Dependence. Wu (2005) first proposed alternatives to mixing concepts by proposing dependence measures rooted in a dynamical system view of a stochastic process. Much work has been done to use such measures to derive approximation results and estimator properties, see for example Wu et al. (2010), Wu (2011), Chen et al. (2016), and references within.

Definition 3.4. If for all $t \in \mathbb{Z}$, Z_t has finite r th moment, where $r \neq 1$, the functional physical dependence measure Δ_r is defined as

$$\Delta_r(p, h) : \sup_t \|Z_t - h\|_{L^r} \quad G^{p,h} p_{Z_t^1, t-1:t-h} q \quad L^r$$

where $\| \cdot \|_{L^r} = (\mathbb{E} | \cdot |^r)^{1/r}$, Z_t^1 is due to $F_t^1 = \dots, \mathcal{F}_{t-1}^1, \mathcal{F}_t^1$ and $\mathbf{t}_{t \in \mathbb{Z}}^1$ is an independent copy of $\mathbf{t}_{t \in \mathbb{Z}}$.

Chen et al. (2016), among others, show how one may replace the geometric α -mixing assumption with a physical dependence assumption.¹⁶ They show that the key sufficient condition is for $\Delta_r(p, h)$ to decay sufficiently fast as h grows.

Definition 3.5 (Geometric Moment Contracting Process). $\mathbf{t}_{Z_t \in \mathbb{Z}}$ is geometric moment contracting (GMC) in L^r norm if there exists $a_1 > 0$, $a_2 > 0$ and $\rho \in (0, 1)$ such that

$$\Delta_r(p, h) \leq a_1 \exp(-a_2 h^\rho).$$

GMC conditions can be considered more general than α -mixing, as they encompass well-known counterexamples, e.g., the known counterexample provided by $Z_t = \rho Z_{t-1} + \epsilon_t$ for $t \in \mathbb{Z}$ i.i.d. Bernoulli r.v.s (Chen et al., 2016). In the following proposition I prove that if contractivity or stability conditions as defined by Pötscher and Prucha (1997) hold for G and $\mathbf{t}_{t \in \mathbb{Z}}$ is an i.i.d. sequence, then process $\mathbf{t}_{Z_t \in \mathbb{Z}}$ is GMC under weak moment assumptions.

¹⁶I adapt the definitions of Chen et al. (2016) to work with a system of the form $Z_t = G p_{Z_{t-1}, \epsilon_t} q$.

Proposition 3.1. Assume that $\{u_{tPZ}, t \in E, \mathbb{R}^{dz}$ are i.i.d. and $\{Z_t\}_{t \in PZ}$ is generated according to

$$Z_t = G_p Z_{t-1} + \epsilon_t,$$

where $Z_t \in \mathbb{R}^{dz}$ and G is a measurable function.

- (a) If contractivity conditions (11)-(12) hold, $\sup_{t \in PZ} \|Z_t\|_{L^r} < \infty$ for $r \geq 2$ and $\|G_p\|_{L^r} < 1$ for some $p, q \in PZ \in E$, then $\{Z_t\}_{t \in PZ}$ is GMC with

$$\Delta_r(p, q) \leq a \exp(-h|q|)$$

where $h = \log C_{Zq}$ and $a = 2 \|Z_t\|_{L^r} < \infty$.

- (b) If stability conditions (13)-(14) hold, $\sup_{t \in PZ} \|Z_t\|_{L^r} < \infty$ for $r \geq 2$ and $BG\{BZ\} \leq M_Z$, then $\{Z_t\}_{t \in PZ}$ is GMC with

$$\Delta_r(p, q) \leq \bar{a} \exp(-h|q|)$$

where $h = \log C_{Zq}$ and $\bar{a} = 2 \|Z_t\|_{L^r} \max\{M_Z^{h-1}, 1\} C_Z < \infty$.

Proposition 3.1 is important in that it links the GMC property to transparent conditions on the structure of the nonlinear model. It also immediately allows handling multivariate systems, while previous work has focused on scalar systems (c.f. Wu (2011) and Chen et al. (2016)).

Finally, it is now possible to show that if $\{W_{2t} u_{tPZ}\}$ satisfies physical dependence assumptions, then Assumption 8 is fulfilled, c.f. Lemma 2.2 in Chen and Christensen (2015) for α -mixing assumptions.

Lemma 3.1. If Assumption 7(iii) holds and $\{W_{2t} u_{tPZ}\}$ is strictly stationary and GMC then one may choose an integer sequence $q = q(n) \leq n/2$ with $\rho_n(q) \leq K \Delta_r(p, q) \leq \rho_1(q)$ for $5/2 \leq p \leq 2 + 2/r$ and $r \geq 2$ such that

$$\rho_{B^1, B^2}(n) \leq I_K \cdot O_P \left(K, n \right) \leq \frac{C}{n} \frac{q \log K}{n} \leq \rho_P(p, q)$$

provided $K, n \leq \frac{a}{\rho \log K} \{n\} \leq \rho_1(q)$.

It can be seen that Lemma 3.1 holds by setting $a = \frac{a}{K \log n} \{n\} \leq \rho_1(q)$ and choosing $q(n) \leq \frac{1}{\log K} n^{r-1}$, where ρ is the GMC factor introduced in Proposition 3.1. Therefore, the rate is the same as the one derived by Chen and Christensen (2015) for exponentially α -mixing regressors. As shown in Proposition 3.1, system contractivity and stability conditions both imply geometric moment contractivity, meaning that in place of Assumption 8 one may require the following.

Assumption 9. For $r \geq 2$ it holds either:

- (i) $\{Z_t\}_{t \in PZ}$ is GMC in L^r norm,
(ii) $\{Z_t\}_{t \in PZ}$ is generated according to $Z_t = \Phi_p(Z_{t-1}, \dots, Z_{t-p}; \epsilon_t)$ where $\sup_{t \in PZ} \|Z_t\|_{L^r} < \infty$ and Φ is either contractive according to Definition 3.1 or stable according to Definition 3.3.

It is straightforward to prove that if GMC conditions are imposed on $\{Z_t\}_{t \in \mathcal{P}Z}$, this implies that $\{W_{2t}\}_{t \in \mathcal{P}Z}$ is also GMC.¹⁷ Therefore, Lemma 3.1 applies and Assumption 8 as well as Assumption 5(i) are verified.

3.4 Uniform Convergence and Consistency

Since the key asymptotic condition of Chen and Christensen (2015) is upheld under GMC assumptions, their uniform convergence bound on the approximation error of the series estimator can be applied. In order to do so, one must also impose some regularity conditions on \mathcal{G}_2 .

Without loss of generality, let $X \in [0, 1]^s$ and let $\|\cdot\|_{\mathcal{G}_2} = \sup_{W \in \mathcal{P}Y} |g_2(W)|$ be the sup-norm of the conditional mean function $g_2(W)$.

Assumption 10. The unconditional density of X_t is uniformly bounded away from zero and infinity over X .

Assumption 11. For all $1 \leq i \leq d_Y$ and $0 \leq j \leq p$, the restriction of g_{ij}^{21} to $[0, 1]^s$ belongs to the Hölder class $\Lambda^s(p, 1)$ of smoothness $s \geq 1$.

Assumptions 10 and 11 are standard in the nonparametric regression literature. One only needs to restrict the complexity of functions g_{ij}^{21} since, for any i , the remainder of $\mathcal{G}_{2,i}$ consists of linear functions. More precisely, what is really needed is that the nonparametric components of the sieve given by b_{1K}, \dots, b_{JK} are able to approximate g_{ij}^{21} well enough.

Assumption 12. Sieve B belongs to $\text{BSpl}(r, [0, 1]^{d_Y}, r)$, the B-spline sieve of degree r over $[0, 1]^{d_Y}$, or $\text{Wav}(p, [0, 1]^{d_Y}, r)$, the wavelet sieve of regularity r over $[0, 1]^{d_Y}$, with $r \geq \max\{s, 1\}$.

In the remainder of the paper, I will consider the cubic spline sieve ($r = 3$), but theoretical results are stated in the more general setting. Moreover, d will be the effective dimension of the joint estimation domain for G_i^{21} .

Theorem 3.1 (Chen and Christensen (2015)). *Let Assumptions 4, 5, 6, 7, 9, 10, 11 and 12 hold. If*

$$K = pn \{ \log pn \}^{d(p+2s)},$$

then

$$\rho_2 = \|\mathcal{G}_2\|_{\mathcal{G}_2} = O_P \{ pn \{ \log pn \}^{s(p+2s)} \}$$

provided that $s \geq 2$ (in Assumption 6) and $d \geq 2s$.

In Theorem 3.1 the sup-norm consistency rate generally depends on the dimension d and thus, in principle, the curse of dimensionality slows down convergence compared to parametric estimation. Fortunately, under the current structural model assumptions, the nonlinear functional components in \mathcal{G}_2 are linearly separable in the lag dimension, and thus one may take $d = 1$ as effective dimension. This also means that condition $d \geq 2s$ is trivially satisfied.

¹⁷A formal argument can be found in Appendix B.

Two-step Consistency. The following theorem ensures that the two-step estimation procedure produces consistent estimates. Since for impulse response functions one needs to study the iteration of the entire structural model, this result is stated in terms of the full coefficient matrices.

Theorem 3.2. *Let $\{Z_t, u_t\}_{t \in \mathbb{Z}}$ be determined by structural model (1). Under Assumptions 1, 4, 5, 6, 7, 9, 10, 11 and 12, let $\hat{\Pi}_1$ and $\hat{\Pi}_2$ be the least squares and semi-nonparametric series estimators for Π_1 and Π_2 , respectively, based on the two-step procedure. Then,*

$$\|\hat{\Pi}_1 - \Pi_1\|_{\infty} = O_P(\rho n^{-1/2})$$

and

$$\|\hat{\Pi}_2 - \Pi_2\|_{\infty} = O_P\left(\frac{K}{n}\right) + O_P\left(\frac{K}{n}\right),$$

where $\hat{\Pi}_2$ is the infeasible series estimator involving $\hat{\Pi}_1$.

Sup-norm bounds for $\|\hat{\Pi}_2 - \Pi_2\|_{\infty}$ follow immediately from Lemma 2.3 and Lemma 2.4 in Chen and Christensen (2015). In particular, choosing the optimal nonparametric rate $K = \rho n \{\log \rho n\}^{d/(p-2s)}$ for the infeasible estimator would yield

$$\|\hat{\Pi}_2 - \Pi_2\|_{\infty} = O_P\left(\rho n \{\log \rho n\}^{s/(p-2s)}\right)$$

as per Theorem 3.1. The condition for consistency in Theorem 3.2 reduces to

$$\frac{K^3}{n} = o(1),$$

since for B-spline and wavelet sieves $\frac{K}{n} \rightarrow 0$ and $\frac{K}{n} \rightarrow \frac{1}{K}$. It is simple to show that if for the feasible estimator $\hat{\Pi}_2$ the same rate $\rho n \{\log \rho n\}^{d/(p-2s)}$ is chosen for K , the consistency condition in the above display is fulfilled assuming $s \neq 1$ and $d = 1$.¹⁸

Remark 3.3. (Hyperparameter Selection). An important practical question when applying any series or kernel-type methods is the selection of hyperparameters. For the former, this entails the choice of the sieve's size K . Although theory provides only asymptotic rates, a number of methods can be used to select K , such as cross-validation, generalized cross-validation and Mallows's criterion (Li and Racine, 2009). In the case of piece-wise splines, once size is selected, knots can be chosen to be the K uniform quantiles of the data. This ensures knots are not located in regions of the domain with very few observations. In simulations and applications, for simplicity, I select sieve sizes manually and locate knots approximately following empirical quantiles. In unreported numerical experiments, I check that results are robust to moderate changes in the number and approximate locations of spline knots.

¹⁸The rate for K may be optimized by balancing the uniform (infeasible) rate with the error due to residuals. Since this paper is not concerned with finding the optimal rate, I do not perform this exercise here.

4 Impulse Response Analysis

Once the model's linear, functional and structural coefficient are consistently estimated, computation of nonlinear impulse responses must be addressed. As discussed in Section 2, nonlinear IRFs are generally hard to lay hands on, since the functional MA(\mathcal{S}) form of the process is highly non-trivial. In this section, I will provide an explicit, iterative algorithm to compute responses that is numerically straightforward and does not require the construction of moving average functional coefficients. Moreover, since to derive uniform bounds it is assumed that the data has compact support, I will introduce a novel yet familiar IRF definition, called the relaxed impulse response function, which is compatible with boundedness. Lastly, I prove that semi-nonparametric IRF estimates are consistent with respect to their population counterparts.

4.1 Computation

Recall from equation (7) in Section 2.1 that impulse responses involve two moving average lag polynomials, $\Theta_p L_q$ for the linear model component and $\Gamma_p L_q$ for the nonlinear component, respectively. As a first step, one can derive a semi-explicit recursive algorithm for computing $\text{IRF}_{h,p,q}$ in a manner that does not involve simulations of the innovations process.

Proposition 4.1 (Gonçalves et al. (2021), Proposition 3.1). *Under Assumptions 1, 2 and 3, for any $h = 0, 1, \dots, H$, let*

$$V_{j,p,q} := \mathbb{E} \Gamma_j X_{t-j,p,q} \mathbb{S} \quad \mathbb{E} \Gamma_j X_{t-j,s}.$$

To compute

$$\text{IRF}_{h,p,q} = \Theta_{h,1} \sum_{j=0}^h V_{j,p,q},$$

the following steps can be used:

(i) For $j = 0$, set $X_{t,p,q} = X_t$ and $V_{0,p,q} = \mathbb{E} \Gamma_0 X_{t,p,q} \mathbb{S} = \mathbb{E} \Gamma_0 X_{t,s}$.

(ii) For $j = 1, \dots, h$, let

$$X_{t-j,p,q} = X_{t-j} - \Theta_{j,11} \sum_{k=1}^j \rho \Gamma_{k,11} X_{t-j-k,p,q} - \Gamma_{j,11} X_{t-j-k,q} \\ - \sum_{j=p} X_{t-j;t} \mathbb{Q},$$

where ρ_j are implicitly defined and depend on $\Theta_p L_q$ and $\Gamma_p L_q$.

(iii) For $j = 1, \dots, h$, compute

$$V_{j,p,q} = \mathbb{E} \Gamma_j \sum_{j=p} X_{t-j;t} \mathbb{Q} \mathbb{S} = \mathbb{E} \Gamma_j X_{t-j,s}.$$

The proof of Proposition 4.1 is identical to that in Gonçalves et al. (2021), with the only variation being that in the current setup it is not possible to collect the nonlinear function across

$X_{t-j} \rho$ and $X_{t-j} \kappa$. Computation of $X_{t-j} \rho$ in step (ii) involves recursive evaluations of nonlinear functions, which is why the algorithm is semi-explicit. For each horizon h , one needs to evaluate $h-1$ iterations of $X_{t-j} \rho$. Importantly, however, this approach dispenses from the need to simulate innovations ε_{t-j}^h as the joint distribution of $\varepsilon_{t-h+1}, \varepsilon_{t-j+1}, \dots, \varepsilon_t$ already contains all relevant information. [Gonçalves et al. \(2021\)](#) naturally argue that the algorithm outlined in Proposition 4.1 is significantly more efficient than schemes involving Monte Carlo simulations like e.g. the one used by [Kilian and Vigfusson \(2011\)](#).

However, ε_{t-j}^h are combinations of real and functional matrices and closed-form derivation is numerically impractical. Note that, by the definition of IRFs, the following *explicit* iterative algorithm is also valid.

Proposition 4.2. *In the same setup of Proposition 4.1, to compute $\text{IRF}_{h\rho}$ the following steps can be used:*

(i¹) For $j=0$, let $X_{t\rho} = X_t$ and

$$\text{IRF}_{0\rho} = B_0^{21} \mathbb{E} \begin{matrix} 0 \\ G_{21,0} X_{t\rho} \end{matrix} \mathbb{E} \begin{matrix} 0 \\ G_{21,0} X_t \end{matrix}.$$

(ii¹) For $j=1, \dots, h$, let

$$\begin{aligned} X_{t-j\rho} &= \mu_1 A_{12\rho} Y_{t-j\rho} + A_{11\rho} X_{t-j-1\rho} + \varepsilon_{t-j}, \\ Y_{t-j\rho} &= \mu_2 A_{22\rho} Y_{t-j-1\rho} + H_{21\rho} X_{t-j\rho} + B_0^{21} \varepsilon_{t-j} + u_{2t-j}, \end{aligned}$$

where $H_{21\rho} = A_{21\rho} L$ and $u_{2t-j} = B_0^{22} \varepsilon_{t-j}$. Setting $Z_{t-j\rho} = \rho X_{t-j\rho}$, $Y_{t\rho} = \rho Y_{t\rho}$ it holds

$$\text{IRF}_{h\rho} = \mathbb{E} Z_{t-j\rho} = \mathbb{E} Z_{t-j}.$$

Proposition 4.2 follows directly from the definition of the unconditional impulse response (6) combined with explicit iteration of the semi-reduced form (2) and sidesteps the MA(8) formulation in (7). Step (i¹) is trivial in nature. Step (ii¹) may not seem useful when compared to (ii), since, in practice, innovations ε_{t-1} and u_{2t} are not available. However, let

$$\hat{\rho}, \hat{A}_{11\rho}, \hat{A}_{12\rho}, \hat{A}_{21\rho}, \hat{A}_{22\rho}, \hat{B}_0^{21}$$

be estimates of the model's coefficients derived, for example, from series estimator $\hat{\Pi}_1$ and $\hat{\Pi}_2$. In sample, one can compute residuals $\hat{\rho}_{1t}$ and $\hat{\rho}_{2t}$, and by definition it holds

$$\begin{aligned} X_t &= \hat{\rho}_1 \hat{A}_{12\rho} Y_{t-1} + \hat{A}_{11\rho} X_{t-1} + \hat{\rho}_{1t}, \\ Y_t &= \hat{\rho}_2 \hat{A}_{22\rho} Y_{t-1} + \hat{A}_{21\rho} X_t + \hat{B}_0^{21} \hat{\rho}_{1t} + \hat{\rho}_{2t}. \end{aligned}$$

This means that one can readily construct the shocked sequence recursively as

$$\begin{aligned} \hat{\mathcal{R}}_{t-j\rho} &= \hat{\rho}_1 \hat{A}_{12\rho} \hat{\mathcal{V}}_{t-j-1\rho} + \hat{A}_{11\rho} \hat{\mathcal{R}}_{t-j-1\rho} + \hat{\rho}_{1t-j}, \\ \hat{\mathcal{V}}_{t-j\rho} &= \hat{\rho}_2 \hat{A}_{22\rho} \hat{\mathcal{V}}_{t-j-1\rho} + \hat{A}_{21\rho} \hat{\mathcal{R}}_{t-j\rho} + \hat{B}_0^{21} \hat{\rho}_{1t-j} + \hat{\rho}_{2t-j}, \end{aligned}$$

for $j = 1, \dots, h$ where $\mathcal{X}_{t|p} = \mathcal{X}_t$, $\mathcal{X}_{t-s} = \mathcal{X}_{t-s}$ for all $s \neq 1$ and similarly for $\mathcal{V}_{t|p}$. To evaluate a structural IRF, over a sample of size n one can compute

$$\text{IRF}_{h|p} = \frac{1}{n} \sum_{j=1}^n \mathcal{V}_{t-j|p} \mathcal{Y}_t,$$

which is still considerably less demanding than Monte Carlo simulations. Additionally, the advantage in implementing steps (i¹)-(ii¹) over the procedure in Proposition 4.1 is that, when $\mathcal{P}_{21|L}$ is a semi-nonparametric estimate, iterating model equations is numerically much more straightforward than handling functional MA matrices $\mathcal{P}_j u_j^h$.

4.2 Nonlinear Responses with Relaxed Shocks

Following Proposition 4.1, the sample impulse response would be

$$\text{IRF}_{h|p} = \Theta_{h,1} \sum_{j=0}^h \bar{\mathcal{V}}_{j|p}, \quad (15)$$

where

$$\bar{\mathcal{V}}_{j|p} = \frac{1}{n} \sum_{t=1}^n \mathcal{P}_j \mathcal{P}_{j|p} \mathcal{X}_{t-j|p} = \mathcal{P}_j \mathcal{X}_{t-j}$$

and Θ , \mathcal{P} and \mathcal{P}_j are plug-in estimates of the respective quantities based on $\hat{\mathcal{H}}_1$ and $\hat{\mathcal{H}}_2$. However, under Assumptions 4 and 5, the construction of impulse response (15) is improper. This can be immediately seen by noticing that, at impact,

$$\mathcal{X}_{t|p} = \mathcal{P}_j \mathcal{X}_{t-j} = \mathcal{X}_t,$$

meaning that $\mathbb{P}(\mathcal{X}_{t|p} \in \mathcal{R} \mid \mathcal{X}_{t-j} = x) = 0$ since there is a translation of size j in the support of \mathcal{X}_t . The problem is rooted in the fact that the standard definition of IRF involves a translation of the distribution of time t structural innovations, which is incompatible with the assumptions imposed in Section 3 to derive semi-nonparametric consistency.

There are multiple ways to address this issue. One option, which would require substantial technical work, is to extend Theorem 3.2 to encompass regressors with unbounded or expanding domains. A potential direction could be coupling the weighted sieves of Chen and Christensen (2015) with appropriately defined shocks. Instead, I propose to take a more direct approach by changing the *type* of structural shock one studies in a way consistent with bounded domains for all variables.

Definition 4.1. A mean-shift structural shock $\mathcal{X}_{t|p}$ is a transformation of \mathcal{X}_t such that

$$\mathbb{P}(\mathcal{X}_{t|p} \in E_1) = 1 \quad \text{and} \quad \mathbb{E}(\mathcal{X}_{t|p}) = \mu.$$

A mean-shift shock is such that the distribution of time t innovations is shifted to have mean μ , while retaining support E almost surely. This definition is natural in that it makes evaluating the effect of the MA(\mathcal{B}) component of the unconditional IRF straightforward. With




Figure 1: Example of symmetric shock relaxation. Unperturbed (left, blue) versus shocked (right, orange) densities of innovations ε_t . The shock relaxation function (right, gray) and together determine the form of the relaxed shock used to compute the IRF.

a mean-shift shock, at impact it holds

$$X_{t|t} - X_t = \varepsilon_t \rho(\varepsilon_t)$$

yet $\varepsilon_t \rho(\varepsilon_t)$ is not known unless the transformation for the mean-shock is itself known. Unfortunately, the assumption that the mean of $\varepsilon_t \rho(\varepsilon_t)$ is exactly equal to 0 requires that the distribution of ε_t be known to properly choose a mean-shift transform. If instead one is willing to assume only that $\mathbb{E}[\varepsilon_t \rho(\varepsilon_t)] = 0$, it is possible to sidestep this requirement by introducing a *shock relaxation function*.

Definition 4.2 (Shock Relaxation Function). *A shock relaxation function is a map $\rho : E_1 \rightarrow \mathbb{R}$ such that $\rho(z) \geq 0$ for all $z \in E_1$, $\rho(z) \neq 0$ for all $z \in E_1$ and there exists $z_0 \in E_1$ for which $\rho(z_0) = 1$.*

In general, choosing a shock relaxation function without taking into account the shape of domain E_1 does not necessarily imply that the relaxed shocks will not push the structural variable out-of-bounds. Therefore, I also introduce the notion of *compatibility*.

Definition 4.3 (Compatible Relaxation). *Consider a shock $\varepsilon \in \mathbb{R}$ and let $E_1 = [a, b]$.*

(i) If $a < 0$, ρ is said to be right-compatible with ε if

$$\rho(z) \propto \frac{b - z}{|b - a|} \text{ for all } z \in E_1.$$

(ii) If $b > 0$, ρ is said to be left-compatible with ε if

$$\rho(z) \propto \frac{z - a}{|b - a|} \text{ for all } z \in E_1.$$

(iii) Given shock size $|\varepsilon| > 0$, ρ is said to be compatible if it is both right- and left-compatible.

By setting

$$\varepsilon_t \rho(\varepsilon_t) = \varepsilon_t \rho(\varepsilon_t)$$

where ρ is compatible with γ , it follows that $X_t \rho \gamma = X_t \rho \gamma$ and $|\mathbb{E} \rho \gamma| = |\mathbb{E} \rho \gamma|$ / since $\mathbb{E} \rho \gamma \in \mathbb{R}$, by definition of ρ . If γ is a bump function, a relaxed shock is a structural shock that has been mitigated proportionally to the density of innovations at the edges of E_1 and the squareness of ρ . For better intuition, Figure 1 provides a graphical rendition of shock relaxation of a symmetric error distribution with a bump function.

Remark 4.1. The definition of compatible relaxation function is *static*, as it considers only the impact effect of a shock. Nonetheless, the assumption that $X_t \in \mathbb{R}$ for all t must also hold for $X_t \rho \gamma$, the shocked structural variable. In theory, given ρ , one can always either expand X or strengthen ρ so that compatibility is enforced at all horizons $1 \leq h \leq H$. For simulations, where one has access to the data generating process, the choice of domains and relaxation functions can be done transparently. In practice, some care is required. When working with empirical data, unless one is willing to assume X_t is wholly exogenous – as in Section 6.1 with monetary policy shocks – or strictly autoregressive, some scenarios are more amenable to analysis with the framework presented here than other. In Section 6.2, following Istrefi and Mouabbi (2018), I will let X_t be a non-negative uncertainty measure, so that negative shocks are harder to study without producing sequences that contain *negative* uncertainty values. Thus, I will focus on positive, contractionary shocks.

For a given X_t , transformation $X_t \rho \gamma$ is not directly applicable since ρ is not observed. In practice, therefore, I will consider

$$\mathcal{X}_t \rho \gamma : X_t \rho \gamma.$$

For simplicity of notation, let $\Gamma_t : \rho \gamma$. Similarly to Step (ii) of Proposition 4.1, given a path $X_t \rho \gamma$ one finds

$$X_t \rho \gamma = X_t \rho \gamma - \Theta_{j,11} \Gamma_t \sum_{k=1}^j \rho \Gamma_{k,11} X_t \rho \gamma - \Gamma_{k,11} X_t \rho \gamma \\ \rho X_t \rho \gamma.$$

The relaxed-shock impulse response is thus given by

$$\mathbb{I}RF_{h\rho \gamma} : \mathbb{E} Z_t \rho \gamma - Z_t \rho \gamma - \Theta_{h,1} \mathbb{E} \rho \gamma \sum_{k=1}^j \mathbb{E} \Gamma_k X_t \rho \gamma - \Gamma_k X_t \rho \gamma.$$

In what follows, I show that by replacing Γ_t with $\rho \gamma$ it is possible to consistently estimate unconditional expectations involving $X_t \rho \gamma$ as well as $X_t \rho \gamma$, and thus $\mathbb{I}RF_{h\rho \gamma}$, by averaging over sample realizations.

4.3 Relaxed Impulse Response Consistency

For a given $\rho \in \mathbb{R}$ and compatible shock relaxation function ρ , vector $V_j \rho \gamma$ is the nonlinear component of impulse responses. One can focus on a specific variable's response by introducing,

for $1 \leq j \leq d$,

$$V_{j, \rho, \mathbf{q}} = \frac{1}{n} \sum_{t=1}^{n_j} \Gamma_{j, \rho} X_{t, j:t; \Gamma_{t, \mathbf{q}}} \Gamma_{j, X_{t, j}},$$

where $V_{j, \rho, \mathbf{q}}$ is the horizon j nonlinear effect on the j th variable and $\Gamma_{j, \rho}$ is the j th component of functional vector Γ_j . For the sake of notation I also define

$$v_{j, \rho} X_{t, j:t; \Gamma_{t, \mathbf{q}}} = \Gamma_{j, \rho} X_{t, j:t; \Gamma_{t, \mathbf{q}}} \Gamma_{j, X_{t, j}}.$$

Let $\hat{v}_{j, \rho} X_{t, j:t; \Gamma_{t, \mathbf{q}}}$ be its sample equivalent, so that

$$\hat{v}_{j, \rho} X_{t, j:t; \Gamma_{t, \mathbf{q}}} = \hat{\mathbf{p}}_{j, \rho} \hat{\mathbf{p}}_{j, \rho} X_{t, j:t; \Gamma_{t, \mathbf{q}}} \hat{\mathbf{p}}_{j, X_{t, j}},$$

$$\hat{\mathbf{p}}_{j, \rho, \mathbf{q}} = \frac{1}{n} \sum_{t=1}^{n_j} \hat{\mathbf{p}}_{j, X_{t, j:t; \Gamma_{t, \mathbf{q}}}}$$

and

$$\hat{\mathbf{Y}}_{\text{IRF}_{h, \rho, \mathbf{q}}} = \Theta_{h, 1} \sum_{t=1}^n \hat{\mathbf{p}}_{\rho, t, \mathbf{q}} \sum_{j=0}^h \hat{\mathbf{p}}_{j, \rho, \mathbf{q}}$$

for $1 \leq j \leq d$.

Theorem 4.1. *Let $\hat{\mathbf{Y}}_{\text{IRF}_{h, \rho, \mathbf{q}}}$ be a semi-nonparametric estimate for the horizon h relaxed shock IRF of variable ρ . Under the same assumptions as in Theorem 3.2*

$$\hat{\mathbf{Y}}_{\text{IRF}_{h, \rho, \mathbf{q}}} \stackrel{P}{\rightarrow} \mathbf{Y}_{\text{IRF}_{h, \rho, \mathbf{q}}}$$

for any fixed integers $0 \leq h \leq \mathcal{B}$ and $1 \leq j \leq d$.

5 Simulations

This section is devoted to analyzing the empirical performance of the two-step semi-nonparametric estimation strategy discussed above. I will consider the two simulation setups employed by [Gonçalves et al. \(2021\)](#), with focus on bias and MSE of the estimated relaxed shocked impulse response functions. Additionally, I provide simulations under a modified design which highlight how in larger samples the non-parametric sieve estimator consistently recovers impulse responses, while a least-squares estimator constructed with a pre-specified nonlinear transform does not. In all simulations, I use a B-spline sieve of order 1.

5.1 Benchmark Bivariate Design

The first simulation setup involves a bivariate DGP where the structural shock does not directly affect other observables. This is a simple environment to check that indeed the two-step estimator recover the nonlinear component of the model and impulse responses are consistently estimated, and that the MSE does not worsen excessively.

I consider three bivariate data generation processes. DGP 1 sets X_t to be a fully exogenous

innovation process,

$$\begin{aligned} X_t &= \varepsilon_{1t}, \\ Y_t &= 0.5Y_{t-1} + 0.5X_t + 0.3X_{t-1} + 0.4\max\{0, X_t\} - 0.3\max\{0, X_{t-1}\} + \varepsilon_{2t}. \end{aligned} \quad (16)$$

DGP 2 adds an autoregressive component to X_t , but maintains exogeneity,

$$\begin{aligned} X_t &= 0.5X_{t-1} + \varepsilon_{1t}, \\ Y_t &= 0.5Y_{t-1} + 0.5X_t + 0.3X_{t-1} + 0.4\max\{0, X_t\} - 0.3\max\{0, X_{t-1}\} + \varepsilon_{2t}. \end{aligned} \quad (17)$$

Finally, DGP 3 add an endogenous effect of Y_{t-1} on the structural variable by setting

$$\begin{aligned} X_t &= 0.5X_{t-1} + 0.2Y_{t-1} + \varepsilon_{1t}, \\ Y_t &= 0.5Y_{t-1} + 0.5X_t + 0.3X_{t-1} + 0.4\max\{0, X_t\} - 0.3\max\{0, X_{t-1}\} + \varepsilon_{2t}. \end{aligned} \quad (18)$$

Following Assumption 1, innovations are mutually independent. To accommodate Assumptions 4 and 5, both ε_{1t} and ε_{2t} are drawn from a truncated standard Gaussian distribution over $[-3, 3]$.¹⁹ All DGPs are centered to have zero intercept in population.

I evaluate bias and MSE plots using 1000 Monte Carlo simulation. For a chosen horizon H , the impact of a relaxed shock on ε_{1t} is evaluated on Y_{t+h} for $h = 1, \dots, H$. To compute the population IRF, I employ a direct simulation strategy that replicates the shock's propagation through the model and I use 10 000 replications. To evaluate the estimated IRF, the two-step procedure is implemented: a sample of length n is drawn, the linear least squares and the semi-nonparametric series estimators of the model are used to estimate the model and the relaxed IRF is computed following Proposition 4.2. For the sake of brevity, I discuss the case of $\alpha = 1$ and I set the shock relaxation function to be

$$\rho(z) = \exp\left(-1 - \frac{z^4}{3}\right) \mathbb{1}_{|z| \leq 1}$$

over interval $[-3, 3]$ and zero everywhere else.²⁰ Choices of $\alpha = 1$ and $\beta = 0.5$ yield similar results in simulations, so I do not discuss them here.

Figure 2 contains the results for sample size $n = 240$. This choice is motivated by considering the average sample sizes found in most macroeconomic settings: it is equivalent to 20 years of monthly data or 60 yearly of quarterly data (Gonçalves et al., 2021). The benchmark method is an OLS regression that relies on a priori knowledge of the underlying DGP specification. Given the moderate sample size, to construct the cubic spline sieve estimator of the nonlinear component of the model I use a single knot, located at 0. The simulations in Figure 2 show that while the MSE is slightly higher for the sieve model, the bias is comparable across methods. Note

¹⁹Let $\varepsilon_{it} \sim \mathcal{N}(0, 1)$ for $i = 1, 2$, then the truncated Gaussian innovations used in simulation are set to be $\varepsilon_{it} = \min\{\max\{-3, \varepsilon_{it}\}, 3\}$. The resulting r.v.s have a non-continuous density with two mass points at -3 and 3. However, in practice, since these masses are negligible, for the moderate sample sizes used this choice does not create issues.

²⁰It can be easily checked that this choice of ρ is compatible with shocks of size $0 \leq |\varepsilon| \leq 1$.




Figure 2: Simulations results for DGPs 1-3.

that for DGP 3, due to the dependence of the structural variable on non-structural series lags, the MSE and bias increase significantly, and there is no meaningful difference in performance between the two estimation approaches.

5.2 Structural Partial Identification Design

To showcase the validity of the proposed sieve estimator under the type of partial structural identification discussed in the paper, I again rely on the simulation design proposed by [Gonçalves et al. \(2021\)](#). All specifications are block-recursive, and require estimating the contemporaneous effects of a structural shock on non-structural variables, unlike in the previous section.

The form of the DGPs is

$$B_0 Z_t = B_1 Z_{t-1} + C_0 \tilde{f}(\chi_t) + C_1 \tilde{f}(\chi_{t-1}) + \epsilon_t,$$

where in all variations of the model

$$B_0 = \begin{pmatrix} 1 & 0 & 0 \\ 0.45 & 1 & 0.3 \\ 0.05 & 0.1 & 1 \end{pmatrix}, \quad C_0 = \begin{pmatrix} 0 \\ 0.2 \\ 0.08 \end{pmatrix}, \quad \text{and} \quad C_1 = \begin{pmatrix} 0 \\ 0.1 \\ 0.2 \end{pmatrix}.$$

I focus on the case $\tilde{f}(\chi) = \max\{0, \chi\}$, since this type of nonlinearity is simpler to study. DGP 4 treats χ_t as an exogenous shock by setting

$$B_1 = \begin{pmatrix} 0 & 0 & 0 \\ 0.15 & 0.17 & 0.18 \\ 0.08 & 0.03 & 0.6 \end{pmatrix};$$

DGP 5 add serial correlation to χ_t ,

$$B_1 = \begin{pmatrix} 0.13 & 0 & 0 \\ 0.15 & 0.17 & 0.18 \\ 0.08 & 0.03 & 0.6 \end{pmatrix};$$

and DGP 6 includes dependence on Y_{t-1} ,

$$B_1 = \begin{pmatrix} 0.13 & 0.05 & 0.01 \\ 0.15 & 0.17 & 0.18 \\ 0.08 & 0.03 & 0.6 \end{pmatrix}.$$

For these data generating processes, I employ the same setup of simulations with DGPs 1-3, including the number of replications as well as the type of relaxed shock. as well as the sieve grid. Here too I evaluate MSE and bias of both the sieve and the correct specification OLS estimators with as sample size of $n = 240$ observations. The results in Figure 3 show again that there is little difference in terms of performance between the semi-nonparametric sieve approach and a correctly-specified OLS regression.




Figure 3: Simulations results for DGPs 4-6.




Figure 4: Plot of nonlinear function $\rho(x)$ used in DGP 7.

5.3 Model Misspecification

The previous sections report results that support the use of the sieve IRF estimator in a sample of moderate size, since it performs comparably to a regression performed with a priori knowledge of the underlying DGP. I now show that the semi-nonparametric approach is also robust to model misspecification compared to simpler specifications involving fixed choices for nonlinear transformations.

To this end, I modify DGP 2 to use a smooth nonlinear transformation to define the effect of structural variable X_t on Y_t . That is, there is no compounding of linear and nonlinear effects. The autoregressive coefficient in the equation for X_t is also increased to make the shock more persistent. The new data generating process, DGP 2¹, is, thus, given by

$$\begin{aligned} X_t &= 0.8X_{t-1} + \varepsilon_{1t}, \\ Y_t &= 0.5Y_{t-1} + 0.9\rho(X_t) + 0.5\rho(X_{t-1}) + \varepsilon_{2t}. \end{aligned} \tag{19}$$



where $\rho(x) = \rho(x - 1) + 0.5 \tanh(\rho(x - 1) - 2)$, which is plotted in Figure 4.

To emphasize the difference in estimated IRFs, in this setup I focus on $\rho(x)$, which requires adapting the choice of innovations and shock relaxation function. In simulations of DGP 2¹, ε_{1t} and ε_{2t} are both drawn from a truncated standard Gaussian distribution over $r \in [-5, 5]$. The shock relaxation function of this setup is given by



$$\rho(x) = \exp\left(-1 - \frac{z}{5}\right) \mathbb{1}_{\{z > 3.9\}}$$

over interval $r \in [-5, 5]$ and zero everywhere else. This form of $\rho(x)$ is adapted to choices of $\rho(x)$ such that $0 < \rho(x) < 2$. The sieve grid now consists of 4 equidistant knots within $p \in [-5, 5]$. I use the same numbers of replications as in the previous simulations. Finally, the regression design is identical to that used for DGP 2 under correct specification.

The results obtained with sample size $n = 2400$ are collected in Figure 5. I choose this larger sample size to clearly showcase the inconsistency of impulse responses under misspecification: as it can be observed, the simple OLS estimator involving the negative-censoring transform



(a) 2



(b) 2

Figure 5: Simulations results for DGP 7.

produces IRF estimates with consistently worse MSE and bias than those of the sieve estimator at almost all horizons. Similar results are also obtained for more moderate shocks $\sigma = 1$, but the differences are less pronounced. These simulations suggest that the semi-nonparametric sieve estimator can produce substantially better IRF estimates in large samples than methods involving nonlinear transformations selected a priori.

In this setup, it is also important to highlight the fact that the poor performance of OLS IRF estimates does not come from $\rho(x)$ being “complex”, and, thus, hard to approximate by combinations of simple functions. In fact, if in DGP 2¹ function $\rho(x)$ is replaced by $\gamma \rho(x) : \rho(x) = 1 - q$, the differences between sieve and OLS impulse response estimates become minimal in simulations, with the bias of the latter decreasing by approximately an order of magnitude (see Figure 8 in Appendix C). This is simply due the fact that $\gamma \rho(x)$ is well approximated by $\max\{0, x\}$ directly. However, one then requires either prior knowledge or sheer luck when constructing the nonlinear transforms of X_t for an OLS regression. The proposed series estimator, instead, just requires an appropriate choice of sieve. Many data-driven procedures to select sieves in applications have been proposed, see for example the discussion in Kang (2021).

6 Empirical Applications

In this section, I showcase the practical utility of the proposed semi-nonparametric sieve estimator by considering two applied exercises. First, I revisit the empirical analysis of Gonçalves et al. (2021), which is itself based on the work of Tenreyro and Thwaites (2016). This provides both linear and nonlinear benchmarks for the monetary policy responses within a compact economic model. I find that, although the differences between approaches are mild, nonparametric IRFs in fact provide counter-evidence to the conclusions reported by Gonçalves et al. (2021). In the second application, I compare the linear and nonlinear impulse responses that are produced by uncertainty shocks in the setup studied by Istrefi and Mouabbi (2018). Here, sieve-estimated IRFs show differences in shape, timing and intensity, chiefly when the sign of the shock changes.

6.1 Monetary Policy Shocks

The objective of the empirical analysis in Gonçalves et al. (2021) is to analyze the effects of a monetary policy shock on a model of the US macroeconomy. Structural identification is achieved via a narrative approach, following the seminal work of Romer and Romer (2004).

The four-variable model is set up identically to the one of Gonçalves et al. (2021), Section 7. Let $Z_t = \rho X_t, \text{FFR}_t, \text{GDP}_t, \text{PCE}_t$ ¹, where X_t is the series of narrative U.S. monetary policy shocks, FFR_t is the federal funds rate, GDP_t is log real GDP and PCE_t is PCE inflation.²¹ As

²¹In Gonçalves et al. (2021) p. 122, it is mentioned that CPI inflation is included in the model, but both in the replication package made available by one the authors (<https://sites.google.com/site/lkilian2019/research/code>) from which I source the data, and Tenreyro and Thwaites (2016), PCE inflation is used instead. Moreover, the authors say that both the FFR and PCE enter the model in first differences, yet in their code these variables are kept in levels. I keep their original formulation to allow for a proper comparison between estimation methods.

a pre-processing step, GDP is transformed to log GDP and then linearly detrended. The data is available quarterly and spans from 1969:Q1 to 2007:Q4. As in [Tenreyro and Thwaites \(2016\)](#), I use a model with one lag, $\rho = 1$. Narrative shock X_t is considered to be an i.i.d. sequence, i.e. $X_t = \epsilon_t$, therefore I assume no dependence on lagged variables when implementing pseudo-reduced form (2). Like in [Gonçalves et al. \(2021\)](#), I consider positive and negative shocks of size $|z| = 1$. As such, I choose

$$p(z) = \exp(-|z|) \propto \exp(-|z|) \left(\frac{z}{4} \right)^6 + 1$$

to be the shock relaxation functions. Figure 10 in Appendix C provides a check for the validity of given the sample distribution of X_t . Knots for sieve estimation are located at $t = -1, 0, 1$. The model is block-recursive, and the structural formulation of Section 2.1 allows identifying the U.S. monetary policy shocks without the need to impose additional assumptions on the remaining shocks. [Gonçalves et al. \(2021\)](#), following [Tenreyro and Thwaites \(2016\)](#), use two nonlinear transformations, $F(x) = \max(0, x)$ and $F(x) = x^3$, to try to gauge how negative versus positive and large versus small shocks, respectively, affect the U.S. macroeconomy. For clarity, below I refer to this approach as “parametric nonlinear method”. Since the authors find that the inclusion of a cubic term does not meaningfully change impulse responses, I focus on comparing the IRFs estimated via sieve regression with the ones obtained by setting $F(x) = \max(0, x)$, as well as by not including nonlinear terms (i.e. linear IRFs).

Figure 6 shows the estimated impulse response to both a positive and negative unforeseen monetary policy shock. The impact on the federal funds rate is consistent across all three procedures, but there are important differences in GDP and inflation responses. In case of an exogenous monetary tightening change, the parametric nonlinear response for GDP, unlike in the case of linear and parametric nonlinear IRFs, is nearly zero at impact and has a monotonic decrease until around 10 quarters ahead. The change in shape is meaningful, as the procedure of [Gonçalves et al. \(2021\)](#) still yields a small short-term upward jump in GDP when a monetary tightening shock hits. Moreover, after the positive shock, the sieve GDP responses reaches its lowest value 4 and 2 quarters before the linear and parametric nonlinear responses, while its size is 13% and 16% larger, respectively.²² Finally, the sieve PCE response is positive for a shorter interval, but looks to be more persistent once it turns negative also 10 months after impact.

When the shock is expansionary, sieve IRFs show a pronounced asymmetry, even more than that of parametric nonlinear responses. One can observe that semi-nonparametric federal funds rate IRF is marginally mitigated compared to the alternative estimates. An important puzzle is due to the clearly negative impact on GDP. Indeed, both types of nonlinear responses show a drop in output in the first 5 quarters. Also note that the PCE inflation has a positive spike the first couple of quarters after impact. Such a quick change seems unrealistic, as one does not expect inflation to suddenly reverse sign, but, as [Gonçalves et al. \(2021\)](#) also remark, the overall

²²The strength of this effect changes across different shocks sizes, as Figure 12 in Appendix C proves. As shocks sizes get smaller, nonlinear IRFs, both parametric and sieve, show decreasing negative effects.

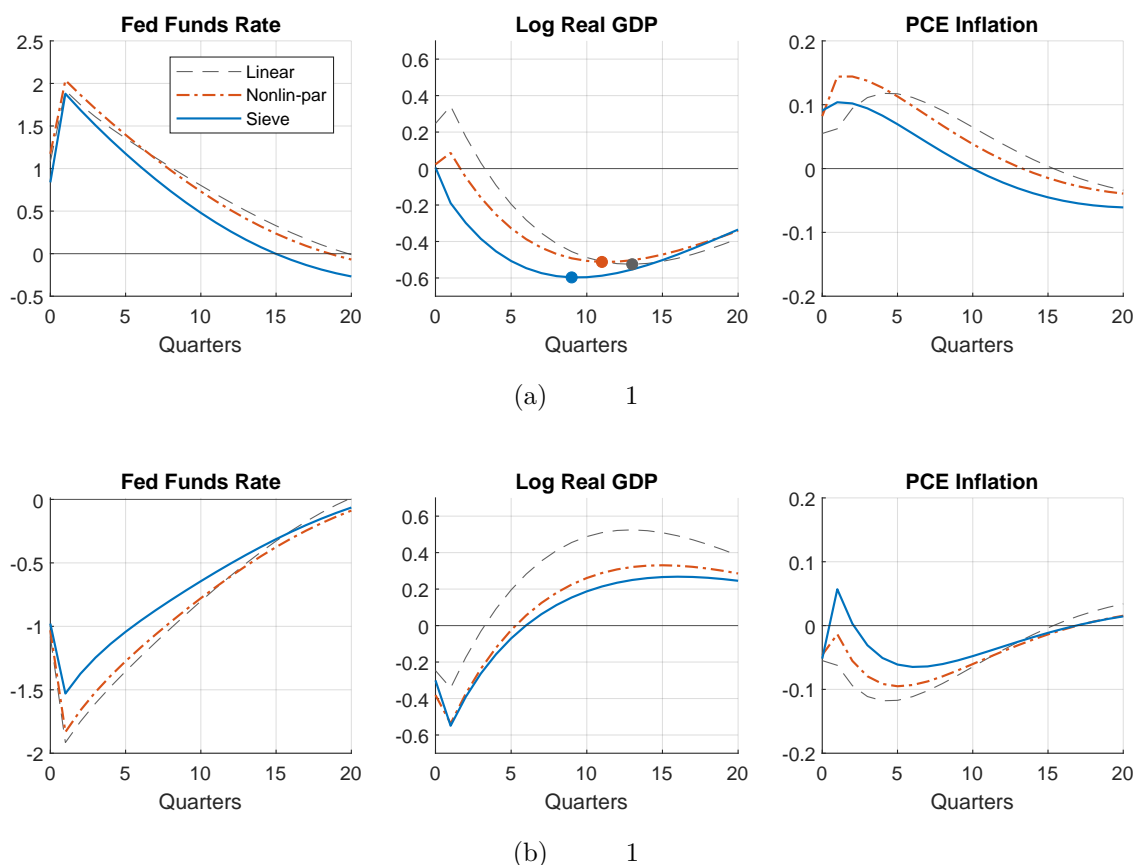


Figure 6: Effect of an unexpected U.S. monetary policy shock on federal funds rate, GDP and inflation. Linear (gray, dashed), parametric nonlinear with $F_{p,q} = \max\{0, x_q\}$ (red, point-dashed) and sieve (blue, solid) structural impulse responses. For $\sigma = 1$, the lowest point of the GDP response is marked with a dot.

impact on inflation of both shocks is small when compared to the change in federal funds rate.

This comparison between methods, and specifically the nature of nonparametric impulse responses, provides evidence that a small econometric model, such as the one studied by [Tenreiro and Thwaites \(2016\)](#), may be inadequate to fully capture the dynamic effects of monetary policy shocks. In both setups, however, impulse response interpretation is only suggestive, as confidence bands are missing and only pointwise IRFs are available. Whether the puzzles highlighted above would persist after accounting for estimation uncertainty is an important research question that I leave for future analysis.

6.2 Uncertainty Shocks

Uncertainty in interest rates appears to be a significant factor in recent economic history. Starting with the fundamental changes brought forth by the unprecedented measures of unconventional monetary policy after the 2007-2008 financial crisis, to the powerful economic stimuli during the COVID-19 pandemic, and finally the subsequent interest rate tightening and inflation phenomenon of 2022, central banks and institutional agents are often very concerned

about uncertainty. Since traditional central bank policymaking is heavily guided by the principle that the central bank *can* and *should* influence expectations, controlling the (perceived) level of ambiguity in current and future commitments is key.

Istrefi and Mouabbi (2018) provide an analysis of the impact of unforeseen changes in the level of subjective interest rate uncertainty on the macroeconomy. They derive a collection of new indices based on short- and long-term profession forecasts. Their empirical study goes in depth into studying the different components that play a role in transmitting uncertainty shocks, but here I will focus on re-evaluating their structural impulse response estimates under the light of potentially-missing nonlinear effects. For the sake of simplicity, my evaluation will focus only on the 3-months-ahead uncertainty measure for short-term interest rate maturities (3M3M) and the US economy.²³

Like in Istrefi and Mouabbi (2018), let $Z_t = (\rho X_t, IP_t, CPI_t, PPI_t, RT_t, UR_t)'$ be a vector where X_t is the chosen uncertainty measure, IP_t is the (log) industrial production index, CPI_t is the CPI inflation rate, PPI_t is the producer price inflation rate, RT_t is (log) retail sales and UR_t is the unemployment rate. The nonlinear model specification is given by

$$Z_t = \mu + A_1 Z_{t-1} + A_2 Z_{t-2} + F_1 \rho X_{t-1} + F_2 \rho X_{t-2} + DW_t + u_t,$$

where W_t includes a linear time trend and oil price OIL_t .²⁴ The data has monthly frequency and spans the period between May 1993 and July 2015.²⁵ Note here that, following the identification strategy of Gonçalves et al. (2021), nonlinear functions F_1 and F_2 are to be understood as not effecting X_t , which is the structural variable. The linear VAR specification of Istrefi and Mouabbi (2018) is recovered by simply assuming $F_1 = F_2 = 0$ prior to estimation. Since they use recursive identification and order the uncertainty measure first, this model too is block-recursive.

I consider a positive shock with intensity ρ , where σ is the standard deviation of structural innovations. In this empirical exercise, the relaxation function is given by

$$\rho \rho \quad | \quad |z| \propto \frac{1}{4} \exp \left(- \frac{|z|}{4} \right)^8 \quad 1 \quad 1$$

and I set $t_0, 1, 0.3u$ to be the cubic spline knots. As 3M3M is a non-negative measure of uncertainty, some care must be taken to make sure that the shocked paths for X_t do not reach negative values. Figure 14 in Appendix C shows that the relaxation function is compatible, and also that the shocked nonlinear paths of X_t with impulse ρ and 1 all do not cross below zero.

Figure 7 presents both the linear and nonlinear structural impulse responses obtained. Importantly, even though Istrefi and Mouabbi (2018) estimate a Bayesian VAR model and here

²³Istrefi and Mouabbi (2018) also provide comparisons with results obtained with the other uncertainty measures, which they comment are all very similar to the ones obtained with 3M3M. Their paper additionally evaluates a number of other highly developed countries.

²⁴Inclusion of linear exogenous variables in the semi-nonparametric theoretical framework detail in Section 3 is straightforward as long as one can assume that they are stationary and weakly dependent. The choice of using $\rho = 2$ is identical to that of the original authors, based on BIC.

²⁵I reuse the original data employed by the authors, who kindly shared it upon request, but rescale retail sales (RT_t) so that the level on January 2000 equals 100.

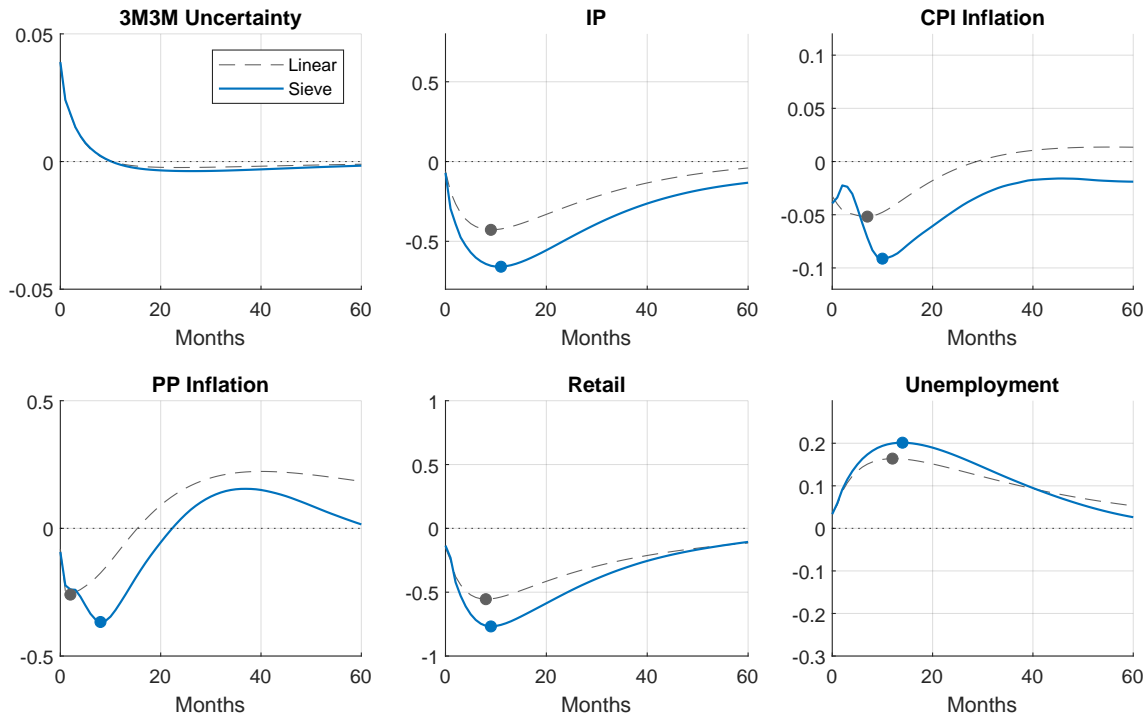


Figure 7: Effect of an unexpected, one-standard-deviation uncertainty shock to US macroeconomic variables. Linear (gray, dashed) and sieve (blue, solid) structural impulse responses. The extreme points of the responses are marked with a dot.

I consider a frequentist vector autoregressive benchmark, the shape of the IRFs is retained, c.f. the median response in the top row of their Figure 4. When uncertainty increases, industrial production drops, and the size and extent of this decrease is intensified in the nonlinear responses. In fact, the sieve IP response reaches a value that is 54% lower than that of the respective linear IRF.²⁶ A similar behavior holds true for retail sales (38% lower) and unemployment (23% higher), proving that this shock is more profoundly contractionary than suggested by the linear VAR model. Further, CPI and PP inflation both show short-term fluctuations which strengthen the short- and medium-term impact of the shock. CPI and PP nonlinear inflation responses are 76% and 41% stronger than their linear counterpart, respectively. These differences suggest that linear IRFs might be both under-estimating the short-term intensity and misrepresenting long-term persistence of inflation reactions. From another perspective, [Nowzohour and Stracca \(2020\)](#) presented evidence that consumer consumption growth, credit growth and unemployment do not co-move with the policy uncertainty index (EPU) of [Baker et al. \(2016\)](#), but are negatively correlated with financial volatility. Given the strength of nonlinear IRFs, this discrepancy may also suggest that the 3M3M uncertainty measure partially captures the financial channel, too.

The introduction of nonlinear terms in the structural VAR of [Istrefi and Mouabbi \(2018\)](#) thus provides evidence that fundamental impulse response features might otherwise be missed. Indeed, Figure 13 in Appendix C - which plots regression functions of endogenous variables

²⁶Figure 15 in Appendix C confirms that this difference is consistent over a range of shock sizes, too.

with respect to X_t - proves that high and low uncertainty levels may have significantly different effects on endogenous economic variables. In particular, at the second lag, tail effects appear to be milder, while at low levels changes in uncertainty have more pronounced impact.

7 Conclusion

This paper studies the application of semi-nonparametric series estimation to the problem of structural impulse response analysis for time series. After first discussing the partial identification model setup, I have used the conditions of system contractivity and stability to derive physical measures of the dependence for nonlinear systems. In turn, these allow to derive primitive conditions under which series estimation can be employed and structural IRFs are consistently estimated. The simulation results prove that this approach is valid in moderate samples and has the added benefit of being robust to misspecification of the nonlinear model components. Finally, two empirical applications showcase the utility in departing from both linear and parametric nonlinear specifications when estimating structural responses.

There are many possible avenues for extending the results I have presented here. A key aspect that I have not touched upon is inference in the form of confidence intervals: the theory of [Chen and Christensen \(2015\)](#) does not encompass uniform inference, and, as such, additional results have to be developed. Indeed, (uniform) confidence bands are necessary to properly quantify the uncertainty of IRF estimates. [Belloni et al. \(2015\)](#) give a uniform asymptotic inference theory, but their derivations are limited to non-dependent data. [Li and Liao \(2020\)](#) and [Cattaneo et al. \(2022\)](#) provide theoretical coupling results that could be exploited in order to handle time series data. [Chen and Christensen \(2018\)](#) give a theory of uniform inference for panel IV setups, which could possibly be generalized to handle nonlinear IRFs. In the spirit of [Kang \(2021\)](#), it would be also important to derive inference results that are uniform in the selection of series terms, as, in practice, a data-driven procedure for selecting \mathcal{K} should be used. Studying other sieve spaces, such as neural networks or shape-preserving sieves ([Chen, 2007](#)), would also be highly desirable. The latter can be especially useful in contexts where economic knowledge suggests that the nonlinear components of the model are e.g. strictly monotonic increasing or convex. Finally, sharpening of convergence rates used in the main proofs is of independent interest.

References

- Auerbach, A. J. and Gorodnichenko, Y. (2012). Measuring the Output Responses to Fiscal Policy. *American Economic Journal: Economic Policy*, 4(2):1–27.
- Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring Economic Policy Uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636.
- Belloni, A., Chernozhukov, V., Chetverikov, D., and Kato, K. (2015). Some new asymptotic theory for least squares series: Pointwise and uniform results. *Journal of Econometrics*, 186(2):345–366.
- Brockwell, P. J. and Davis, R. A. (1991). *Time Series: Theory and Methods*. Springer Science & Business Media.
- Caggiano, G., Castelnuovo, E., Colombo, V., and Nodari, G. (2015). Estimating Fiscal Multipliers: News From A Non-linear World. *The Economic Journal*, 125(584):746–776.
- Caggiano, G., Castelnuovo, E., and Figueres, J. M. (2017). Economic policy uncertainty and unemployment in the United States: A nonlinear approach. *Economics Letters*, 151:31–34.
- Caggiano, G., Castelnuovo, E., and Pellegrino, G. (2021). Uncertainty shocks and the great recession: Nonlinearities matter. *Economics Letters*, 198:109669.
- Cattaneo, M. D., Farrell, M. H., and Feng, Y. (2020). Large sample properties of partitioning-based series estimators. *The Annals of Statistics*, 48(3):1718–1741.
- Cattaneo, M. D., Masini, R. P., and Underwood, W. G. (2022). Yurinskii’s Coupling for Martingales. *Working Paper*.
- Chen, X. (2007). Chapter 76 Large Sample Sieve Estimation of Semi-Nonparametric Models. In Heckman, J. J. and Leamer, E. E., editors, *Handbook of Econometrics*, volume 6, pages 5549–5632. Elsevier.
- Chen, X. (2013). Penalized Sieve Estimation and Inference of Semiparametric Dynamic Models: A Selective Review. In Acemoglu, D., Arellano, M., and Dekel, E., editors, *Advances in Economics and Econometrics*, pages 485–544. Cambridge University Press, 1 edition.
- Chen, X. and Christensen, T. M. (2015). Optimal uniform convergence rates and asymptotic normality for series estimators under weak dependence and weak conditions. *Journal of Econometrics*, 188(2):447–465.
- Chen, X. and Christensen, T. M. (2018). Optimal sup-norm rates and uniform inference on nonlinear functionals of nonparametric IV regression. *Quantitative Economics*, 9(1):39–84.
- Chen, X., Shao, Q.-M., Wu, W. B., and Xu, L. (2016). Self-normalized Cramér-type moderate deviations under dependence. *The Annals of Statistics*, 44(4):1593–1617.

- Chen, X. and Shen, X. (1998). Sieve Extremum Estimates for Weakly Dependent Data. *Econometrica*, 66(2):289.
- Chetverikov, D., Santos, A., and Shaikh, A. M. (2018). The Econometrics of Shape Restrictions. *Annual Review of Economics*, 10(1):31–63.
- Debortoli, D., Forni, M., Gambetti, L., and Sala, L. (2020). Asymmetric Effects of Monetary Policy Easing and Tightening. *Working Paper*.
- Fan, J. and Yao, Q. (2003). *Nonlinear time series: nonparametric and parametric methods*, volume 20. Springer.
- Farrell, M. H., Liang, T., and Misra, S. (2021). Deep Neural Networks for Estimation and Inference. *Econometrica*, 89(1):181–213.
- Feng, B. Q. (2003). Equivalence constants for certain matrix norms. *Linear Algebra and its Applications*, 374:247–253.
- Forni, M., Gambetti, L., Maffei-Faccioli, N., and Sala, L. (2023a). Nonlinear transmission of financial shocks: Some new evidence. *Journal of Money, Credit and Banking*.
- Forni, M., Gambetti, L., and Sala, L. (2023b). Asymmetric effects of news through uncertainty. *Macroeconomic Dynamics*, pages 1–25.
- Freyberger, J. and Reeves, B. (2018). Inference under Shape Restrictions. *Working Paper*.
- Fuleky, P., editor (2020). *Macroeconomic Forecasting in the Era of Big Data: Theory and Practice*, volume 52 of *Advanced Studies in Theoretical and Applied Econometrics*. Springer International Publishing, Cham.
- Gambetti, L., Maffei-Faccioli, N., and Zoi, S. (2022). Bad News, Good News: Coverage and Response Asymmetries. *Working Paper*.
- Gao, J. (2007). *Nonlinear Time Series: Semiparametric and Nonparametric Methods*. Chapman and Hall/CRC.
- Gonçalves, S., Herrera, A. M., Kilian, L., and Pesavento, E. (2021). Impulse response analysis for structural dynamic models with nonlinear regressors. *Journal of Econometrics*, 225(1):107–130.
- Gourieroux, C. and Jasiak, J. (2005). Nonlinear Innovations and Impulse Responses with Application to VaR Sensitivity. *Annales d'Économie et de Statistique*, pages 1–31.
- Gourieroux, C. and Lee, Q. (2023). Nonlinear impulse response functions and local projections. *Working Paper*.
- Hamilton, J. D. (1994a). State-space models. *Handbook of Econometrics*, 4:3039–3080.
- Hamilton, J. D. (1994b). *Time Series Analysis*. Princeton University Press.

- Härdle, W., Lütkepohl, H., and Chen, R. (1997). A Review of Nonparametric Time Series Analysis. *International Statistical Review*, 65(1):49–72.
- Horn, R. A. and Johnson, C. R. (2012). *Matrix Analysis*. Cambridge University Press, second edition.
- Horowitz, J. L. and Lee, S. (2017). Nonparametric estimation and inference under shape restrictions. *Journal of Econometrics*, 201(1):108–126.
- Huang, Y., Chen, X., and Wu, W. B. (2014). Recursive Nonparametric Estimation for Time Series. *IEEE Transactions on Information Theory*, 60(2):1301–1312.
- Istrefi, K. and Mouabbi, S. (2018). Subjective interest rate uncertainty and the macroeconomy: A cross-country analysis. *Journal of International Money and Finance*, 88:296–313.
- Jordà, Ò. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182.
- Kanazawa, N. (2020). Radial basis functions neural networks for nonlinear time series analysis and time-varying effects of supply shocks. *Journal of Macroeconomics*, 64:103210.
- Kang, B. (2021). Inference In Nonparametric Series Estimation with Specification Searches for the Number of Series Terms. *Econometric Theory*, 37(2):311–345.
- Kilian, L. and Lütkepohl, H. (2017). *Structural Vector Autoregressive Analysis*. Themes in Modern Econometrics. Cambridge University Press, Cambridge.
- Kilian, L. and Vega, C. (2011). Do energy prices respond to us macroeconomic news? a test of the hypothesis of predetermined energy prices. *Review of Economics and Statistics*, 93(2):660–671.
- Kilian, L. and Vigfusson, R. J. (2011). Are the responses of the us economy asymmetric in energy price increases and decreases? *Quantitative Economics*, 2(3):419–453.
- Koop, G., Pesaran, M. H., and Potter, S. M. (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics*, 74(1):119–147.
- Lanne, M. and Nyberg, H. (2023). Nonparametric Impulse Response Analysis in Changing Macroeconomic Conditions. *Working Paper*.
- Li, J. and Liao, Z. (2020). Uniform nonparametric inference for time series. *Journal of Econometrics*, page 14.
- Li, Q. and Racine, J. S. (2009). *Nonparametric econometric methods*. Emerald Group Publishing.
- Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*. New York : Springer, Berlin.

- Movahedifar, M. and Dickhaus, T. (2023). On the closed-loop Volterra method for analyzing time series. *Working Paper*.
- Nowzohour, L. and Stracca, L. (2020). More than a feeling: Confidence, uncertainty, and macroeconomic fluctuations. *Journal of Economic Surveys*, 34(4):691–726.
- Pellegrino, G. (2021). Uncertainty and monetary policy in the US: A journey into nonlinear territory. *Economic Inquiry*, 59(3):1106–1128.
- Pötscher, B. M. and Prucha, I. (1997). *Dynamic nonlinear econometric models: Asymptotic theory*. Springer Science & Business Media.
- Potter, S. M. (2000). Nonlinear impulse response functions. *Journal of Economic Dynamics and Control*, 24(10):1425–1446.
- Ramey, V. A. and Zubairy, S. (2018). Government spending multipliers in good times and in bad: evidence from us historical data. *Journal of political economy*, 126(2):850–901.
- Romer, C. D. and Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American economic review*, 94(4):1055–1084.
- Sims, C. A. (1980). Macroeconomics and Reality. *Econometrica*, 48(1):1–48.
- Sirotko-Sibirskaya, N., Franz, M. O., and Dickhaus, T. (2020). Volterra bootstrap: Resampling higher-order statistics for strictly stationary univariate time series. *Working Paper*.
- Stock, J. H. and Watson, M. W. (2016). Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. In *Handbook of Macroeconomics*, volume 2, pages 415–525. Elsevier.
- Tenreyro, S. and Thwaites, G. (2016). Pushing on a string: US monetary policy is less powerful in recessions. *American Economic Journal: Macroeconomics*, 8(4):43–74.
- Teräsvirta, T., Tjøstheim, D., and Granger, C. W. J. (2010). *Modelling Nonlinear Economic Time Series*. Oxford University Press.
- Tong, H. (1990). *Non-linear Time Series: a Dynamical System Approach*. Oxford University Press.
- Tropp, J. A. (2012). User-Friendly Tail Bounds for Sums of Random Matrices. *Foundations of Computational Mathematics*, 12(4):389–434.
- Tsay, R. S. and Chen, R. (2018). *Nonlinear Time Series Analysis*, volume 891. John Wiley & Sons.
- Tsybakov, A. B. (2009). *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer, New York ; London.

- Wu, W. B. (2005). Nonlinear system theory: Another look at dependence. *Proceedings of the National Academy of Sciences*, 102(40):14150–14154.
- Wu, W. B. (2011). Asymptotic theory for stationary processes. *Statistics and its Interface*, 4(2):207–226.
- Wu, W. B., Huang, Y., and Huang, Y. (2010). Kernel estimation for time series: An asymptotic theory. *Stochastic Processes and their Applications*, 120(12):2412–2431.

Appendix

A Preliminaries

Matrix Norms. Let

$$\|A\|_r := \max_{\|x\|_r = 1} \|Ax\|_r$$

be the r -operator norm of matrix $A \in \mathbb{C}^{d_1 \times d_2}$. The following Theorem establishes the equivalence between different operator norms as well as the compatibility constants.

Theorem A.1 (Feng (2003)). Let $1 \leq p, q \leq \infty$. Then for all $A \in \mathbb{C}^{d_1 \times d_2}$,

$$\|A\|_p \leq \sqrt[p]{\|A\|_q} \sqrt[q]{\|A\|_p},$$

where

$$\sqrt[p]{\|A\|_q} = \begin{cases} 1 & \text{if } p \neq q, \\ \|A\|_q & \text{if } p = q. \end{cases}$$

This norm inequality is sharp.

In particular, if $p \leq q$ then it holds

$$\frac{1}{\|A\|_q^{1/p - 1/q}} \|A\|_p \leq \|A\|_q \leq \|A\|_p^{1/q} \|A\|_q^{1/p}.$$

B Proofs

B.1 GMC Conditions and Proposition 3.1

Lemma B.1. Assume that $\{u_t\}_{t \in \mathbb{P}Z}$, $\{P_t\}_{t \in E}$, \mathbb{R}^{d_z} are i.i.d., and $\{Z_t\}_{t \in \mathbb{P}Z}$ is generated according to

$$Z_t = G_p(Z_{t-1}, u_t),$$

where $Z_t \in \mathbb{R}^{d_z}$ and G is a measurable function. If either

- (a) Contractivity conditions (11)-(12) hold, $\sup_{t \in \mathbb{P}Z} \|P_t\|_{L^r} \leq \mathcal{B}$ and $\|G_p(z, u)\|_{L^r} \leq \mathcal{B}$ for some $\|z, u\| \leq E$;
- (b) Stability conditions (13)-(14) hold, $\sup_{t \in \mathbb{P}Z} \|P_t\|_{L^r} \leq \mathcal{B}$ and $\|BG\|_{L^r} \leq M_Z \leq \mathcal{B}$;

then

$$\sup_t \|Z_t\|_{L^r} \leq \mathcal{B} \text{ w.p.1.}$$

Proof.

- (a) In a first step, we show that, given event Ω , realization $\{Z_t\}_{t \in \mathbb{P}Z}$ is unique with probability one. To do this, introduce initial condition Z for $t = 1$ such that $Z \in \mathbb{R}^{d_z}$ and $\|Z\| \leq E$. Define

$$Z_t^p = G_p^p(Z_{t-1}^p, u_t), \quad t = 1, 2, \dots$$

Further, let $Z_t^{1p, q}$ be the realization with initial condition $z^1 \bullet Z$ and innovation realizations $t \geq 1: t p, q$. Note that

$$Z_t^{1p, q} - Z_t^{1p, q} \preceq C_Z \|z^1 - z^1\|,$$

which goes to zero as $\|\tilde{N}\| \rightarrow 0$. Therefore, if we set $Z_t^{1p, q} := \lim_{\|\tilde{N}\| \rightarrow 0} Z_t^{1p, q}$, $Z_t^{1p, q}$ is unique with respect to the choice of Z w.p.1. A similar recursion shows that

$$Z_t^{1p, q} \preceq C_Z \|z^1 - z^1\| + \sum_{k=0}^{t-1} C_Z^k C_{t-k}^{p, q}.$$

By norm equivalence, this implies

$$\begin{aligned} \|Z_t^{1p, q}\|_{L^r} &\preceq C_Z \|z^1 - z^1\|_r + \sum_{k=0}^{t-1} C_Z^k C_{t-k}^{p, q} \\ &\preceq C_Z \|z^1 - z^1\|_r + C_Z \sup_{t \in \mathbb{P}^Z} C_{t-1}^{p, q} \|\mathcal{B}\|, \end{aligned}$$

and taking the limit $\|\tilde{N}\| \rightarrow 0$ proves the claim.

- (b) Consider again distinct initial conditions $z^1 \bullet Z$ and innovation realizations $t \geq 1: t p, q$, yielding $Z_t^{1p, q}$ and $Z_t^{1p, q}$, respectively. We may use the contraction bound derived in the proof of Proposition 3.1 (b) below, that is,

$$\|Z_t^{1p, q} - Z_t^{1p, q}\|_r \preceq C_Z C_2 \|z^1 - z^1\|_r,$$

where $C_2 \geq 0$ is a constant. With trivial adjustments, the uniqueness and limit arguments used for (a) above apply here too. □

Proof of Proposition 3.1.

- (a) By assumption it holds that for all $p, z, z^1 \in \mathbb{P}^Z$ and $p, e, e^1 \in \mathbb{P}^E$

$$\|G_p z, q - G_p z^1, q\| \preceq C_Z \|z - z^1\| + C \|e - e^1\|$$

holds, where $0 \preceq C_Z \leq 1$ and $0 \preceq C \leq \mathcal{B}$. The equivalence of norms directly generalizes this inequality to any r -norm for $r \geq 2$. We study $\|Z_t^h - Z_t^{1h}\|_r$ where Z_t^{1h} is constructed with a time- t perturbation of the history of Z_t^h . Therefore, for any given t and $h \geq 1$ it holds that

$$\begin{aligned} \|Z_t^h - Z_t^{1h}\|_r &\preceq C_Z \|G^{p/h} z_{t-1:t-h}^1 - G^{p/h} z_{t-1:t-h}\| + \|G^{p/h} z_{t-1:t-h}^1 - G^{p/h} z_{t-1:t-h}\|_r \\ &\preceq C_Z^h \|z_t - z_t^1\|_r, \end{aligned}$$

since sequence $z_{t-1:t-h}$ is common between Z_t^h and Z_t^{1h} . Clearly then

$$\|Z_t^h - Z_t^{1h}\|_r \preceq 2 \|z_t - z_t^1\|_r \exp(-h)$$

for $\sup_t \|C_Z\| \leq \log p$. Letting $a = 2$ and shifting time index t backward by h , since $\sup_t \|Z_t\| \leq \log p$ w.p.1 from Lemma B.1 the result for L^r follows with $\epsilon = 1$.

(b) Proceed similar to (a), but notice that now we must handle cases of steps $1 \leq h \leq h$. Consider iterate $h = 1$, for which

$$\begin{aligned} \|Z_{t-h} - G^{p/h} Z_{t-1:h}^1\|_r &\leq C_Z \|G^{p/h} Z_{t-1:h}^1 - G^{p/h} Z_{t-1:h}^1\|_r \\ &\leq C_Z^h \|Z_{t-1:h}^1 - Z_{t-1:h}^1\|_r \\ &\leq C_Z^h M_Z \|Z_t - Z_{t-h}^1\|_r \end{aligned}$$

by the mean value theorem. Here we may assume that $M_Z \leq 1$ otherwise we would fall under case (a), so that $M_Z \leq M_Z^2 \leq \dots \leq M_Z^{h-1}$. More generally,

$$\|Z_{t-h} - G^{p/h} Z_{t-1:h}^1\|_r \leq C_Z^{j p/h} \max_t M_Z^{h-1} \|Z_t - Z_{t-h}^1\|_r$$

for $j p/h \leq t \leq h$. Result (b) then follows by noting that $j p/h \leq h-1$ and then proceeding as in (a) to derive GMC coefficients. □

Companion and Lagged Vectors. The assumption of GMC for a process translates naturally to vectors that are composed of stacked lags of realizations. This, for example, is important in the discussion of Section 3 when imposing Assumption 9, since one needs that series regressors $\{W_{2t}\}_{t \in \mathcal{P}_Z}$ be GMC.

Recall that $W_{2t} = (X_t, X_{t-1}, \dots, X_{t-p}, Y_t, Y_{t-1}, \dots, Y_{t-p+1})$. Here we shall reorder this vector slightly to be

$$W_{2t} = (X_t, X_{t-1}, Y_{t-1}, \dots, X_{t-p}, Y_{t-p+1}).$$

For $h \geq 0$ and $1 \leq l \leq h$, let $Z_{t-h}^1 : \Phi^{p/h} Z_{t-h}^1, \dots, Z_{t-h}^1$; $t-1:h$ be the a perturbed version of Z_t , where $Z_{t-h}^1, \dots, Z_{t-h}^1$ are taken from an independent copy of $\{Z_t\}_{t \in \mathcal{P}_Z}$. Define

$$W_{2t}^1 = (X_t^1, X_{t-1}^1, Y_{t-1}^1, \dots, X_{t-p}^1, Y_{t-p+1}^1).$$

Using Minkowski's inequality

$$\begin{aligned} \|W_{2t-h} - W_{2t-h}^1\|_{L^r} &\leq \|X_{t-h} - X_{t-h}^1\|_{L^r} + \sum_{j=1}^p \|Z_{t-h-j} - Z_{t-h-j}^1\|_{L^r} \\ &\leq \sum_{j=0}^p \|Z_{t-h-j} - Z_{t-h-j}^1\|_{L^r}, \end{aligned}$$

thus, since $p \geq 0$ is fixed finite,

$$\sup_t \|W_{2t-h} - W_{2t-h}^1\|_{L^r} \leq \sum_{j=0}^p \Delta_r(p, h-j) \leq p \log p \leq a_{1Z} \exp a_{2Z} h.$$

Above, a_{1Z} and a_{2Z} are the GMC coefficients of $\{Z_t\}_{t \in \mathcal{P}_Z}$.

B.2 Lemma 3.1 and Matrix Inequalities under Dependence

In order to prove Lemma 3.1, the idea is to modify the approach of [Chen and Christensen \(2015\)](#), which relies on Berbee's Lemma and an interlaced coupling, to handle variables with physical dependence. [Chen et al. \(2016\)](#) provide an example on how to achieve this when working with self-normalized sums. In what follows I modify their ideas to work with random dependent matrices.

First of all, I recall below a Bernstein-type inequality for independent random matrices of [Tropp \(2012\)](#).

Theorem B.1. *Let $\{\Xi_i\}_{i=1}^n$ be a finite sequence of independent random matrices with dimensions $d_1 \times d_2$. Assume $\mathbb{E}\Xi_i = 0$ for each i and $\max_{1 \leq i \leq n} \|\Xi_i\| \leq R_n$ and define*

$$S_n^2 := \max_{i=1}^n \left(\mathbb{E} \|\Xi_{i,n} \Xi_{j,n}^T\| + \mathbb{E} \|\Xi_{i,n}^T \Xi_{j,n}\| \right).$$

Then for all $z \neq 0$,

$$\mathbb{P} \left(\left\| \sum_{i=1}^n \Xi_i \right\| \geq z \right) \leq p d_1 d_2 \exp \left(- \frac{z^2}{n q S_n^2} \right).$$

The main exponential matrix inequality due to [Chen and Christensen \(2015\)](#), Theorem 4.2 is as follows.

Theorem B.2. *Let $\{X_i\}_{i=1}^n$ where $X_i \in \mathbb{R}^{d_1 \times d_2}$ be a ρ -mixing sequence and let $\Xi_{i,n} = \Xi_n(X_i)$ for each i where $\Xi_n : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^{d_1 \times d_2}$ be a sequence of measurable $d_1 \times d_2$ matrix-valued functions. Assume that $\mathbb{E}\Xi_{i,n} = 0$ and $\|\Xi_{i,n}\| \leq R_n$ for each i and define*

$$S_n^2 := \max_{i=1}^n \left(\mathbb{E} \|\Xi_{i,n} \Xi_{j,n}^T\| + \mathbb{E} \|\Xi_{i,n}^T \Xi_{j,n}\| \right).$$

Let $l \leq q \leq n/2$ be an integer and let $I = \{q, \dots, n\}$ when $q \leq n$ and $I = \emptyset$ otherwise. Then, for all $z \neq 0$,

$$\mathbb{P} \left(\left\| \sum_{i=1}^n \Xi_{i,n} \right\| \geq z \right) \leq \frac{n}{q} p q \mathbb{P} \left(\left\| \sum_{i \in I} \Xi_{i,n} \right\| \geq z \right) + 2 p d_1 d_2 \exp \left(- \frac{z^2}{n q S_n^2} \right),$$

where $\sum_{i \in I} \Xi_{i,n} = 0$ whenever $I = \emptyset$.

To fully extend Theorem B.2 to physical dependence, I will proceed in steps. First, I derive a similar matrix inequality by directly assuming that random matrices $\Xi_{i,n}$ have physical dependence coefficient $\Delta_r^\Xi \leq \rho$. In the derivations I will use that

$$\frac{1}{p d_2 q^{1/2} q^{1/r}} \|A\|_r \leq \|A\|_2 \leq p d_1 q^{1/2} q^{1/r} \|A\|_r.$$

for $r \geq 2$.

Theorem B.3. Let $\{u_{jPZ}\}$ be a sequence of i.i.d. variables and let $\{\Xi_{i,n} u_{jPZ}^n\}$,

$$\Xi_{i,n} = G_n^{\Xi}(\dots, i-1, i),$$

for each i , where $\Xi_n: X \rightarrow \mathbb{R}^{d_1 \times d_2}$, be a sequence of measurable $d_1 \times d_2$ matrix-valued functions. Assume that $\mathbb{E} \Xi_{i,n} = 0$ and $\|\Xi_{i,n}\| \leq R_n$ for each i and define

$$S_n^2 = \max_{i \neq j} \mathbb{E} \|\Xi_{i,n} \Xi_{j,n}^T\|, \mathbb{E} \|\Xi_{i,n} \Xi_{j,n}^T\|.$$

Additionally assume that $\|\Xi_{i,n}\|_{L^r} \leq B$ for $r \geq 2$ and define the matrix physical dependence measure $\Delta_r^{\Xi}(\rho, \eta)$ as

$$\Delta_r^{\Xi}(\rho, \eta) = \max_{1 \leq i \leq n} \|\Xi_{i,n} - \Xi_{i,n}^{\eta}\|_{L^r},$$

where $\Xi_{i,n}^{\eta} = G_n^{\Xi}(\dots, i-h-1, i-h, i-h-1, \dots, i-1, i)$ for independent copy $\{u_{jPZ}\}$. Let $1 \leq q \leq n/2$ be an integer and let $I = \{qn, \dots, n\}$ when $qn \leq n$ and $I = \emptyset$ otherwise. Then, for all $z \geq 0$,

$$\mathbb{P} \left(\sum_{i=1}^n \|\Xi_{i,n}\| \geq z \right) \leq \frac{n^{r-1}}{q^r p d_2 q^{r(2-1/z)}} \Delta_r^{\Xi}(\rho, \eta) \mathbb{P} \left(\sum_{i \in I} \|\Xi_{i,n}\| \geq z \right) + 2p d_1 d_2 \exp \left(-\frac{z^2/2}{nqS_n^2} \right) + qR_n z/3,$$

where $\sum_{i \in I} \|\Xi_{i,n}\| = 0$ whenever $I = \emptyset$.

Proof. To control dependence, we can adapt the interlacing block approach outlined by [Chen et al. \(2016\)](#). To interlace the sum, split it into

$$\sum_{i=1}^n \|\Xi_{i,n}\| = \sum_{j \in J_e} W_j + \sum_{j \in J_o} W_j + \sum_{i \in I} \|\Xi_{i,n}\|,$$

where $W_j = \sum_{i=qj-1}^{qj} \|\Xi_{i,n}\|$ for $j = 1, \dots, \lfloor n/q \rfloor$ are the blocks, $I = \{qn, \dots, n\}$ if $qn \leq n$ and J_e and J_o are the subsets of even and odd numbers of $\{1, \dots, \lfloor n/q \rfloor\}$, respectively. For simplicity define $J = J_e \cup J_o$ as the set of block indices and let

$$W_j^i = \mathbb{E} \|W_j\| \mid \{u_{jPZ}\}_{j \in J, j \neq i}.$$

Note that by construction $\{W_j^i\}_{j \in J_e}$ are independent and also $\{W_j^i\}_{j \in J_o}$ are independent. Using the triangle inequality we find

$$\begin{aligned} \mathbb{P} \left(\sum_{i=1}^n \|\Xi_{i,n}\| \geq z \right) &\leq \mathbb{P} \left(\sum_{j \in J} \|W_j\| + \sum_{i \in I} \|\Xi_{i,n}\| \geq z \right) \\ &\leq \mathbb{P} \left(\sum_{j \in J} \|W_j\| \geq z \right) + \mathbb{P} \left(\sum_{i \in I} \|\Xi_{i,n}\| \geq z \right) \\ &= \mathbb{P} \left(\sum_{j \in J_e} \|W_j\| \geq z \right) + \mathbb{P} \left(\sum_{j \in J_o} \|W_j\| \geq z \right) + \mathbb{P} \left(\sum_{i \in I} \|\Xi_{i,n}\| \geq z \right) \\ &= I + II + III + IV. \end{aligned}$$

We keep term IV as is. As in the proof of [Chen and Christensen \(2015\)](#), terms II and III consist of sums of independent matrices, where each W_j^i satisfies $W_j^i \preceq qR_n$ and

$$\max \left(\mathbb{E} \|W_j^i W_j^{i1}\|, \mathbb{E} \|W_j^{i1} W_j^i\| \right) \preceq qS_n^2.$$

Then, using the exponential matrix inequality of [Tropp \(2012\)](#),

$$\mathbb{P} \left(\sum_{j \in J_e} W_j^i \preceq z \right) \preceq \rho d_1 \exp \left(-\frac{z^2 \{2\}}{nqS_n^2} \right) \preceq qR_n z \{3\}.$$

The same holds for the sum over J_o . Finally, we use the physical dependence measure Δ_r^Ξ to bound I . Start with the union bound to find

$$\begin{aligned} \mathbb{P} \left(\sum_{j \in J} \rho W_j - W_j^i \preceq z \right) &\preceq \mathbb{P} \left(\sum_{j \in J} W_j - W_j^i \preceq z \right) \\ &\preceq \frac{n}{q} \mathbb{P} \left(W_j - W_j^i \preceq \frac{q}{n} z \right), \end{aligned}$$

where we have used that $\text{tr}\{q\} \preceq n\{q\}$. Since W_j and W_j^i differ only over a \mathcal{G} -algebra that is q steps in the past, by assumption

$$W_j - W_j^i \preceq_{L^r} q \Delta_r^\Xi \rho q q,$$

which implies, by means of the r th moment inequality,

$$\mathbb{P} \left(W_j - W_j^i \preceq \frac{q}{n} z \right) \preceq \mathbb{P} \left(\rho d_2 q^{1\{r-1\}2} \|W_j - W_j^i\|_r \preceq \frac{q}{n} z \right) \preceq \frac{n^r}{q^r \rho d_2 q^{r\{2-1\}Z^r}} \Delta_r^\Xi \rho q q.$$

where $\rho d_2 q^{1\{r-1\}2}$ is the operator norm equivalence constant such that $\|X\|_r \preceq \rho d_2 q^{1\{r-1\}2} \|X\|$ ([Feng, 2003](#)). Therefore,

$$\mathbb{P} \left(\sum_{j \in J} \rho W_j - W_j^i \preceq z \right) \preceq \frac{n^{r-1}}{q^r \rho d_2 q^{r\{2-1\}Z^r}} \Delta_r^\Xi \rho q q$$

as claimed. \square

Notice that the first term in the bound is weaker than that derived by [Chen and Christensen \(2015\)](#). The \mathcal{G} -mixing assumption and Berbee's Lemma give strong control over the probability $\mathbb{P} \left(\sum_{j \in J} \rho W_j - W_j^i \preceq z \right)$. In contrast, assuming physical dependence means we have to explicitly handle a moment condition. One might think of sharpening [Theorem B.3](#) by sidestepping the r th moment inequality (c.f. avoiding Chebyshev's inequality in concentration results), but I do not explore this approach here.

The second step is to map the physical dependence of a generic vector time series $\{X_i\}_{i \in \mathbb{Z}}$ to matrix functions.

Proposition B.1. *Let $\{X_i\}_{i \in \mathbb{Z}}$ where $X_i \in \mathbb{R}^p$ for $i \in \mathbb{Z}$ be a sequence*

with finite r th moment, where $r \geq 0$, and functional physical dependence coefficients

$$\Delta_r(\rho, h) = \sup_i \|X_i - G^{(h)}(X_{i-1:i-h})\|_{L^r}$$

for $h \geq 1$. Let $\Xi_{i,n} = \Xi_n(X_i)$ for each i where $\Xi_n : X \rightarrow \mathbb{R}^{d_1 \times d_2}$ be a sequence of measurable $d_1 \times d_2$ matrix-valued functions such that $\Xi_n = \rho(v_1, \dots, v_{d_2})$ for $v \in \mathbb{R}^{d_1}$. If $\Xi_{i,n} \in \mathcal{S}$ and

$$C_{\Xi} := \sup_{x \in X} \|\rho(x)\| \in C_{\Xi} \in \mathcal{S},$$

then matrices $\Xi_{i,n}$ have physical dependence coefficients

$$\Delta_r^{\Xi}(\rho, h) = \sup_i \|\Xi_{i,n} - \Xi_{i,n}^h\|_{L^r} \leq \frac{d_2}{d_1} \frac{d_2}{d_1} C_{\Xi} \Delta_r(\rho, h),$$

where $\Xi_{i,n}^h = \Xi_n(G^{(h)}(X_{i-1:i-h}))$.

Proof. To derive the bound, we use $\Xi_n(X_i)$ and $\Xi_n(X_i^h)$ in place of $\Xi_{i,n}$ and $\Xi_{i,n}^h$, respectively, where $X_i^h = G^{(h)}(X_{i-1:i-h})$. First we move from studying the operator r -norm (recall, $r \geq 2$) to the Frobenius norm,

$$\|\Xi_n(X_i) - \Xi_n(X_i^h)\|_r \leq \rho d_2^{1/2} \|\Xi_n(X_i) - \Xi_n(X_i^h)\|_F.$$

where as intermediate step we use the 2-norm. Let $\Xi_n = \rho(v_1, \dots, v_{d_2})$ for $v \in \mathbb{R}^{d_1}$ and $v \in \mathbb{R}^{d_1 \times d_2}$, so that

$$\|\Xi_n\|_F = \left(\sum_{i=1}^{d_2} \|v_i\|_2^2 \right)^{1/2}$$

where $v = \rho(v_1, \dots, v_{d_2})$. Since $v : X \rightarrow \mathbb{R}^{d_1}$ are vector functions, the mean value theorem gives that

$$\|\Xi_n(X_i) - \Xi_n(X_i^h)\|_F \leq \left(\sum_{i=1}^{d_2} C_{\Xi_i}^2 \|X_i - X_i^h\|_2^2 \right)^{1/2} \leq \frac{d_2}{d_1} C_{\Xi} \|X_i - X_i^h\|_2.$$

Combining results and moving from the vector r -norm to the 2-norm yields

$$\|\Xi_n(X_i) - \Xi_n(X_i^h)\|_r \leq \rho d_2^{1/2} \left(\rho d_1^{1/2} \right)^{1/r} C_{\Xi} \|X_i - X_i^h\|_r.$$

The claim involving the L^r norm follows immediately. \square

The following Corollary, which specifically handles matrix functions defined as outer products of vector functions, is immediate and covers the setups of series estimation.

Corollary B.1. Under the conditions of Proposition B.1, if

$$\Xi_n(X_i) = \rho(X_i) \rho(X_i^1) = Q_n$$

where $\Xi_n : X \rightarrow \mathbb{R}^d$ is a vector function and $Q_n \in \mathbb{R}^{d \times d}$ is nonrandom matrix, then

$$\Delta_{r,p} \Xi_n \approx d^{3/2} \lambda^{2/r} C \Delta_{r,p} h_n,$$

where $C := \sup_{x \in X} \|\nabla_x h(x)\|$.

Proof. Matrix Q_n cancels out since it is nonrandom and appears in both $\Xi_n \rho(X_i; q)$ and $\Xi_n \rho(X_i^n; q)$. Since $\Xi_n \rho(X_i; q)$ is square, the ratio of row to column dimensions simplifies. \square

The following Corollaries to Theorem B.3 can now be derived in a straightforward manner.

Corollary B.2. Under the conditions of Theorem B.3 and Proposition B.1, for all $z \neq 0$

$$\mathbb{P} \left(\sum_{i=1}^n \Xi_{i,n} \neq z \right) \approx \frac{n^{r-1} \rho d_2 q^{2/r} \lambda^{2/r} \rho d_1 q^{1/2} \lambda^{1/r} C_{\Xi} \Delta_{r,p} q q}{q^r z^r} \mathbb{P} \left(\sum_{i=1}^n \Xi_{i,n} \neq z \right) \\ \approx 2 \rho d_1 \rho d_2 q \exp \left(- \frac{z^2 \lambda^{2/r}}{n q S_n^2} \right) \frac{1}{q R_n z^3}.$$

where $\Delta_{r,p} q$ is the functional physical dependence coefficient of X_i .

Corollary B.3. Under the conditions of Theorem B.3 and Proposition B.1, if $q = \rho n \lambda$ is chosen such that

$$\frac{n^{r-1} \rho d_2 q^{2/r} \lambda^{2/r} \rho d_1 q^{1/2} \lambda^{1/r} C_{\Xi} \Delta_{r,p} q q}{q^r} \approx \rho \lambda$$

and $R_n \approx \frac{1}{q \log \rho d_1 \rho d_2 q} \approx \rho S_n \approx \frac{1}{n \lambda}$ then

$$\sum_{i=1}^n \Xi_{i,n} = O_P \left(S_n \frac{1}{n q \log \rho d_1 \rho d_2 q} \right).$$

This result is almost identical to Corollary 4.2 in Chen and Christensen (2015), with the only adaptation of using Theorem B.3 as a starting point. Condition $R_n \approx \frac{1}{q \log \rho d_1 \rho d_2 q} \approx \rho S_n \approx \frac{1}{n \lambda}$ is simple to verify by assuming, e.g., $q = \rho n \lambda \log \rho n \lambda$ since $\log \rho d_1 \rho d_2 q \approx \log \rho K$ and $K = \rho n \lambda$.

Note that when $d_1 = d_2 = K$, which is the case of interest in the series regression setup, the first condition in Corollary B.3 reduces to

$$K^{5/2} \lambda^{2/r} \rho d_1 q^{1/2} \lambda^{1/r} C_{\Xi} \Delta_{r,p} q q \approx \rho \lambda,$$

which also agrees with the rate of Corollary B.1. Assumption 7(i) and a compact domain further allow to explicitly bound factor C_{Ξ} by

$$C_{\Xi} \leq K^{-2},$$

so that the required rate becomes

$$K \Delta_{r,p} q q \approx \rho \lambda, \quad \text{where } \lambda = \frac{3}{2} \frac{r}{2} \approx 2.$$

Proof of Lemma 3.1. The proof follows from Corollary B.3 by the same steps of the proof of Lemma 2.2 in Chen and Christensen (2015). Simply take $\Xi_{i,n} = n^{-1} \rho \bar{b} \lambda^K \rho X_{i,q} \bar{b} \lambda^K \rho X_{i,q}^1 / \lambda^K q$ and note that $R_n \asymp n^{-1} p_1 \asymp \frac{2}{K,n} \frac{2}{K,n} q$ and $S_n \asymp n^{-2} p_1 \asymp \frac{2}{K,n} \frac{2}{K,n} q$. \square

For Lemma 3.1 to hold under GMC assumptions a valid choice for $q\rho nq$ is

$$q\rho nq = n^{-1} \log p K n^{r-1} q$$

where ρ as in Proposition 3.1. This is due to

$$\begin{aligned} \frac{n^{-r-1}}{q} q K \Delta_r p q q &\lesssim \frac{n^{r-1}}{q^r} K \exp p q q \\ &\lesssim \frac{n^{r-1} K}{\log p K n^{r-1} q^r} p K n^{r-1} q^{-1} \\ &= \frac{1}{\log p K n^{r-1} q^r} p q. \end{aligned}$$

Note then that, if $\frac{K,n}{K,n} \lesssim 1$ and $\frac{K,n}{K,n} \lesssim \frac{?}{K}$, since

$$\frac{K,n}{K,n} \frac{C}{\frac{q \log K}{n}} \lesssim \frac{C}{\frac{K \log p K n^{r-1} q \log p K q}{n}} \lesssim \frac{C}{\frac{K \log p n^{-r-2} q \log p n q}{n}} \lesssim \frac{C}{\frac{K \log p n q^2}{n}},$$

to satisfy Assumption 8 we may assume $\frac{a}{K \log p n q^2 \{n\}} \asymp p_1 q$ as in Remark 2.3 of Chen and Christensen (2015) for the case of exponential ρ -mixing regressors.

B.3 Theorem 3.2

Before delving into the proof of Theorem 3.2, note that we can decompose $\hat{\Pi}_2 = \Pi_2$ as

$$\hat{\Pi}_2 = \Pi_2 + \rho \hat{\Pi}_2 + \hat{\Pi}_2 q + \rho \hat{\Pi}_2 + \hat{\Pi}_2 q + \rho \hat{\Pi}_2 + \Pi_2 q,$$

where $\hat{\Pi}_2$ is the projection of Π_2 onto the linear space spanned by the sieve. The last two terms can be handled directly with the theory developed by Chen and Christensen (2015). Specifically, their Lemma 2.3 controls the second term (variance term), while Lemma 2.4 handles the third term (bias term). This means here we can focus on the first term, which is due to using generated regressors ρ_{1t} in the second step.

Since $\hat{\Pi}_2$ can be decomposed in d_Y rows of semi-nonparametric coefficients, i.e.,

$$\begin{array}{ccc} & 2,1 & \\ Y_t & \vdots & W_{2t} \quad \sigma_{2t}, \\ & 2,d_Y & \end{array}$$

we further reduce to the scalar case. Let ρ_2 be any row of Π_2 and, with a slight abuse of notation, Y the vector of observations of the component of Y_t of the same row, so that one may write

$$\rho_2 \rho X q = \rho_2 \rho X q = \bar{b}^K \rho X q \mathbb{B}^1 \mathbb{B} = \mathbb{B} \mathbb{B}^{-1} Y = \bar{b}^K \rho X q \mathbb{B}^1 \mathbb{B} = \mathbb{B}^1 \mathbb{B} = \mathbb{B}^1 Y$$

where $\mathcal{B}^K \rho \times q \Gamma_{B,2}^{1\{2\}} b^K \rho \times q$ is the orthonormalized sieve according to $\Gamma_{B,2} : \mathbb{E} r b^K \rho W_{2t} q b^K \rho W_{2t} q^1$'s, \mathcal{B} is the *infeasible* orthonormalized design matrix (involving ρ_{1t}) and \mathcal{B} is *feasible* orthonormalized design matrix (involving ρ_{1t}). In particular, note that

$$\mathcal{B} = B R_n, \text{ where } R_n : \begin{matrix} 0 & 0 & \rho_{11} & \dots & 11 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \rho_{1n} & \dots & 1n \end{matrix} \in \mathbb{R}^{n \times K},$$

which implies $\mathcal{B} = B R_n \Gamma_{B,2}^{1\{2\}} : R_n$.

The next Lemma provides a bound for the difference $\rho \mathcal{B}^1 \mathcal{B} \{nq\} - \rho \mathcal{B}^1 B \{nq\}$ that will be useful in the proof of Theorem 3.1 below.

Lemma B.2. *Under the setup of Theorem 3.1, it holds*

$$\rho \mathcal{B}^1 \mathcal{B} \{nq\} - \rho \mathcal{B}^1 B \{nq\} = O_{pp} \frac{a}{K} \overline{K} \{nq\}.$$

Proof. Using the expansion $\mathcal{B}^1 \mathcal{B} = B^1 B + \rho \mathcal{B}^1 R_n + R_n^1 B + q + R_n^1 R_n$, one immediately finds that

$$\rho \mathcal{B}^1 \mathcal{B} \{nq\} - \rho \mathcal{B}^1 B \{nq\} \approx 2 B^1 R_n \{n\} + R_n^1 R_n \{n\}.$$

The second right-hand side factor satisfies $R_n^1 R_n \{n\} \approx \frac{2}{K,n} R_n^1 R_n \{n\}$. Moreover,

$$\begin{aligned} R_n^1 R_n \{n\} &= \frac{1}{n} \sum_{t=1}^n \rho \rho_{1t} \rho_{1t}^2 \\ &= \frac{1}{n} \sum_{t=1}^n \rho \Pi_1 \mathbf{\hat{H}}_1 q^1 W_{1t} W_{1t}^1 \rho \Pi_1 \mathbf{\hat{H}}_1 q \\ &\approx \Pi_1 \mathbf{\hat{H}}_1^2 W_1^1 W_1 \{n\} \\ &= O_{pp} n^{-1} q, \end{aligned}$$

since $W_1^1 W_1 \{n\} = O_{pp} 1q$. Under Assumption 12, $\frac{2}{K,n} \{n\} = O_{pp} \frac{a}{K} \overline{K} \{nq\}$ since B-splines and wavelets satisfy $\frac{2}{K,n} \mathbf{\hat{A}} 1$. Consequently, $R_n^1 R_n \{n\} = O_{pp} \frac{a}{K} \overline{K} \{nq\}$.

Factor $B^1 R_n \{n\}$ is also straightforward, but depends on sieve dimension K ,

$$\begin{aligned} B^1 R_n \{n\} &\approx \frac{1}{n} \sum_{t=1}^n \mathcal{B}^K \rho W_{2t} q \rho \rho_{1t} \rho_{1t}^2 \\ &= \frac{1}{n} \sum_{t=1}^n \mathcal{B}^K \rho W_{2t} q W_{1t}^1 \rho \Pi_1 \mathbf{\hat{H}}_1 q \\ &\approx \Pi_1 \mathbf{\hat{H}}_1 B^1 W_1 \{n\} \\ &= O_{pp} \frac{a}{K} \overline{K} \{nq\}, \end{aligned}$$

since $B^1 W_1 \{n\} = O_{pp} \frac{a}{K} \overline{K} \{nq\}$ as the column dimension of W_1 is fixed. The claim then follows by

noting $O_P(\frac{a}{\sqrt{K}})$ is the dominating order of convergence. \square

Proof of Theorem 3.2. Since $\hat{\Pi}_1$ the least squares estimator of a linear equation, the rate of convergence is the parametric rate $n^{-1/2}$. The first result is therefore immediate.

For the second step, we consider

$$\hat{\Pi}_2 - \Pi_2 \leq \hat{\Pi}_2 - \Pi_2 + \Pi_2 - \Pi_2,$$

and bound explicitly the first right-hand side term. For a given component of the regression function,

$$|\rho_2(x) - \hat{\rho}_2(x)| \leq |I| + |II|.$$

We now control each term on the right side.

(1) It holds

$$\begin{aligned} |I| &\leq \sup_{x \in \mathcal{W}_2} |\rho_2(x) - \hat{\rho}_2(x)| \leq \sup_{x \in \mathcal{W}_2} |\rho_2(x) - \Pi_2(x)| + \sup_{x \in \mathcal{W}_2} |\Pi_2(x) - \hat{\Pi}_2(x)| \\ &\leq \sup_{x \in \mathcal{W}_2} |\rho_2(x) - \Pi_2(x)| + \sup_{x \in \mathcal{W}_2} |\Pi_2(x) - \hat{\Pi}_2(x)|. \end{aligned}$$

Let A_n denote the event on which $\sup_{x \in \mathcal{W}_2} |\Pi_2(x) - \hat{\Pi}_2(x)| \leq \frac{1}{K}$, so that $\sup_{x \in \mathcal{W}_2} |\Pi_2(x) - \hat{\Pi}_2(x)| \leq \frac{1}{K}$ on A_n . Notice that since $\sup_{x \in \mathcal{W}_2} |\rho_2(x) - \Pi_2(x)| \leq O_P(1)$ (Lemma B.2) and, by assumption, $\sup_{x \in \mathcal{W}_2} |\Pi_2(x) - \hat{\Pi}_2(x)| \leq O_P(1)$, then $\mathbb{P}(A_n^c) \leq o(1)$. On A_n then

$$|I| \leq \frac{1}{K} + \sup_{x \in \mathcal{W}_2} |\Pi_2(x) - \hat{\Pi}_2(x)| \leq \frac{1}{K} + \sup_{x \in \mathcal{W}_2} |\Pi_2(x) - \hat{\Pi}_2(x)|.$$

From $\hat{\Pi}_2 = \frac{1}{n} \sum_{t=1}^n \rho_2(x_t) \rho_2(x_t)^\top$ it follows that

$$\hat{\Pi}_2 - \Pi_2 = \frac{1}{n} \sum_{t=1}^n (\rho_2(x_t) \rho_2(x_t)^\top - \Pi_2)$$

on A_n , meaning

$$|I| \leq O_P\left(\frac{1}{K} + \frac{1}{\sqrt{n}}\right)$$

as $\sup_{x \in \mathcal{W}_2} |\rho_2(x) \rho_2(x)^\top - \Pi_2| \leq O_P(1)$ and $\mathbb{P}(A_n^c) \leq o(1)$.

(2) Again we proceed by uniformly bounding $|II|$ according to

$$|II| \leq \sup_{x \in \mathcal{W}_2} |\rho_2(x) \rho_2(x)^\top - \Pi_2| \leq \sup_{x \in \mathcal{W}_2} |\rho_2(x) \rho_2(x)^\top - \Pi_2| + \sup_{x \in \mathcal{W}_2} |\Pi_2(x) - \hat{\Pi}_2(x)|.$$

The last factor has order $O_P(\frac{1}{\sqrt{n}})$ since Π_2 is growing in row dimension with K . For the middle term, introduce

$$\Delta_B := \sup_{x \in \mathcal{W}_2} |\rho_2(x) \rho_2(x)^\top - \Pi_2|$$

and event

$$B_n: \left\{ \|\hat{B}^1 - B^1\|_{\infty} \leq \frac{\Delta_B}{\sqrt{n}} \right\} \times \left\{ \|\hat{K} - K\|_{\infty} \leq \frac{\Delta_B}{\sqrt{n}} \right\}.$$

On B_n , we can apply the bound (Horn and Johnson, 2012)

$$\|\hat{B}^1 - B^1\|_{\infty} \leq \frac{\|\hat{B}^1 - B^1\|_{\infty} \|\hat{K} - K\|_{\infty}}{\lambda_{\min}(\hat{B}^1 - B^1)} \leq \frac{\Delta_B}{\lambda_{\min}(B^1)} \frac{\Delta_B}{\sqrt{n}}.$$

Since $\|\hat{B}^1 - B^1\|_{\infty} \leq O_P\left(\frac{\Delta_B}{\sqrt{n}}\right)$ by Lemma B.2, we get

$$\|\hat{B}^1 - B^1\|_{\infty} = O_P\left(\frac{\Delta_B}{\sqrt{n}}\right)$$

on B_n . Finally, using $\mathbb{P}(\|A - B\|_{\infty} \leq \epsilon) \leq \mathbb{P}(\|A\|_{\infty} \leq \epsilon) + \mathbb{P}(\|B\|_{\infty} \leq \epsilon)$ we note that $\mathbb{P}(B_n^c) \rightarrow 0$ so that the bound asymptotically holds irrespective of event B_n .

Thus, we have shown that

$$\begin{aligned} \|\hat{\rho}_2 - \rho_2\|_{\infty} &= O_P\left(\frac{\Delta_B}{\sqrt{n}}\right) + O_P\left(\frac{\Delta_B}{\sqrt{n}}\right) \\ &= O_P\left(\frac{\Delta_B}{\sqrt{n}}\right) \end{aligned}$$

as clearly $\frac{\Delta_B}{\sqrt{n}} \rightarrow 0$ and, as discussed in the proof of Lemma B.2, $\frac{\Delta_B}{\sqrt{n}} \leq O_P\left(\frac{\Delta_B}{\sqrt{n}}\right)$. This bound is uniform in X and holds for each of the (finite number of) components of $\hat{\mathbf{H}}_2$, therefore the proof is complete. \square

B.4 Theorem 4.1

Before proving impulse response consistency, I show that compositions of the model's autoregressive nonlinear maps are also consistently estimated at any fixed horizon. This means that the "functional moving average" coefficient matrices Γ_j involved in Proposition 4.1 can be consistently estimated with $\hat{\mathbf{H}}_1$ and $\hat{\mathbf{H}}_2$.

Lemma B.3. *Under the assumptions of Theorem 3.2 and for any fixed integer $j \neq 0$ it holds*

$$\|\hat{\Gamma}_j - \Gamma_j\|_{\infty} = o_P(1).$$

Proof. By definition, recall that $\Gamma_j L = \Psi_j L G_j$ where $\Psi_j = \rho I_d - A_j L$. Since $\Psi_j L$ is an MA(\mathcal{B}) lag polynomial, we have that

$$\Gamma_j L = \sum_{k=0}^{\infty} \Psi_k L^k = \rho G_0 + G_1 L + \dots + G_p L^p,$$

where $\Psi_0 = I_d$, $\Psi_k \in \mathbb{R}^{d \times d}$ are purely real matrices and G_0 is a functional vector that may also contain linear components (i.e. allow linear functions of X_t). This means that Γ_j is a convolution

of real and functional matrices,

$$\Gamma_j = \sum_{k=1}^{\min(j, \rho)} \Psi_j \cdot k G_k.$$

The linear coefficients of Λ_j can be consistently estimated by $\hat{\mathbf{H}}_1$ and $\hat{\mathbf{H}}_2$, and thus plug-in estimate $\hat{\Psi}_j$ is consistent for Ψ_j (Lütkepohl, 2005). Therefore,

$$\begin{aligned} \hat{\mathbf{P}}_j - \Gamma_j &\leq \sum_{k=1}^{\min(j, \rho)} \Psi_j \cdot k G_k - \hat{\Psi}_j \cdot k \hat{G}_k \\ &\leq \sum_{k=1}^{\min(j, \rho)} \|\Psi_j - \hat{\Psi}_j\| \cdot k G_k + \sum_{k=1}^{\min(j, \rho)} \hat{\Psi}_j \cdot k |G_k - \hat{G}_k| \\ &\leq \sum_{k=1}^{\min(j, \rho)} O_p(1) C_{G,k} + O_p(1) O_p(1) \\ &= O_p(1), \end{aligned}$$

where $C_{G,k}$ is a constant and $|G_k - \hat{G}_k| = O_p(1)$ as a direct consequence of Proposition 3.2. \square

Note. Since we assume that the model respects either contractivity or stability conditions, the impulse responses must decay (eventually) exponentially fast to zero. This means that by “stitching” bounds appropriately, one should also be able to achieve convergence *uniformly* over $h = 0, 1, \dots, \mathcal{S}$.

Recall now that the sample estimate for the relaxed-shock impulse response is

$$\hat{\mathbf{Y}}_{\text{RF}, h, \rho, \mathbf{q}} = \Theta_{h,1} \sum_{t=1}^n \mathbf{p}_{1,t} \sum_{j=0}^h \hat{\mathbf{V}}_{j, \rho, \mathbf{q}}$$

where

$$\hat{\mathbf{V}}_{j, \rho, \mathbf{q}} = \frac{1}{n} \sum_{t=1}^n \hat{\mathbf{V}}_{j, \rho, \mathbf{q}} \cdot \mathbf{X}_{t-j:t} \cdot \mathbf{P}_t = \frac{1}{n} \sum_{t=1}^n \hat{\mathbf{V}}_{j, \rho, \mathbf{q}} \cdot \mathbf{X}_{t-j:t} \cdot \mathbf{P}_t.$$

Therefore, the estimated horizon h impulse response of the j th variable is

$$\hat{\mathbf{Y}}_{\text{RF}, h, \rho, \mathbf{q}} = \Theta_{h,1} \sum_{t=1}^n \mathbf{p}_{1,t} \sum_{j=0}^h \frac{1}{n} \sum_{t=1}^n \hat{\mathbf{V}}_{j, \rho, \mathbf{q}} \cdot \mathbf{X}_{t-j:t} \cdot \mathbf{P}_t.$$

Lemma B.4. Under the assumptions of Theorem 4.1, let $\mathbf{x}_{j:0} = (\mathbf{x}_j, \dots, \mathbf{x}_0) \in \mathbb{R}^J$ and $\mathbf{P} \in \mathbb{R}^{E_1}$ be nonrandom quantities. Let \mathbf{r} be the relaxed shock determined by \mathbf{r} , \mathbf{q} , and \mathbf{P} . Then

- (i) $\sup_{\mathbf{x}_{j:0}} \|\hat{\mathbf{V}}_{j, \rho, \mathbf{q}} \cdot \mathbf{x}_{j:0} \cdot \mathbf{r}\| = O_p(1)$,
- (ii) $\sup_{\mathbf{x}_{j:0}} \|\hat{\mathbf{V}}_{j, \rho, \mathbf{q}} \cdot \mathbf{x}_{j:0} \cdot \mathbf{r}\| = O_p(1)$,

for any fixed integers $j \geq 0$ and $\mathbf{P} \in \mathbb{R}^{E_1}$.

Proof.

(i) From Proposition 4.1, we have that

$$\mathbb{P}_j \mathbb{P}_{X_{j:0}; \mathbf{q}} \left(\mathbf{X}_j - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}} \prod_{k=1}^j \mathbb{P}_{\Gamma_{k,11} X_{j-k} \mathbb{P}_{\mathbf{q}}} \Gamma_{k,11} X_{j-k} \mathbf{q},$$

thus

$$\begin{aligned} \left| \mathbb{P}_j \mathbb{P}_{X_{j:0}; \mathbf{q}} \left(\mathbf{X}_j - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}} - \prod_{k=1}^j \mathbb{P}_{\Gamma_{k,11} X_{j-k} \mathbb{P}_{\mathbf{q}}} \Gamma_{k,11} X_{j-k} \mathbf{q} \right| \\ \leq \prod_{k=1}^j \mathbb{P}_{\Gamma_{k,11} X_{j-k} \mathbb{P}_{\mathbf{q}}} \Gamma_{k,11} X_{j-k} \mathbf{q} \leq \prod_{k=1}^j \mathbb{P}_{\Gamma_{k,11} X_{j-k} \mathbf{q}} \Gamma_{k,11} X_{j-k} \mathbf{q}. \end{aligned}$$

This yields

$$\sup_{X_{j:0}} \left| \mathbb{P}_j \mathbb{P}_{X_{j:0}; \mathbf{q}} \left(\mathbf{X}_j - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}} - \prod_{k=1}^j \mathbb{P}_{\Gamma_{k,11} X_{j-k} \mathbf{q}} \Gamma_{k,11} X_{j-k} \mathbf{q} \right| \leq 2^j \sup_{X_{j:0}} \mathbb{P}_{\Gamma_{k,11} X_{j-k} \mathbf{q}} \Gamma_{k,11} X_{j-k} \mathbf{q}.$$

Since j is finite and fixed and the uniform consistency bound of Lemma B.3 holds, a fortiori $\sup_{X_{j:0}} \mathbb{P}_{\Gamma_{k,11} X_{j-k} \mathbf{q}} \Gamma_{k,11} X_{j-k} \mathbf{q} = o_{\mathbb{P}}(1)$.

(ii) Similarly to above,

$$\begin{aligned} \left| \mathbb{P}_j \left(\mathbf{X}_j - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}} - \prod_{k=1}^j \mathbb{P}_{\Gamma_{k,11} X_{j-k} \mathbf{q}} \Gamma_{k,11} X_{j-k} \mathbf{q} \right| \\ \leq \mathbb{P}_j \left(\mathbf{X}_j - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}} \leq \mathbb{P}_j \left(\mathbf{X}_j - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}} \\ \leq 2 \mathbb{P}_j \left(\mathbf{X}_j - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}} \leq C_{\Gamma,j,l} \mathbb{P}_j \left(\mathbf{X}_j - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}} \leq C_{\Gamma,j,l} \mathbb{P}_j \left(\mathbf{X}_j - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}}, \end{aligned}$$

where we have used that $\mathbb{P}_j \left(\mathbf{X}_j - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}} \leq \mathbb{P}_j \left(\mathbf{X}_j - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}}$ to derive the first term in the second line. In the last line, $C_{\Gamma,j,l}$ is a constant such that

$$\mathbb{P}_j \left(\mathbf{X}_j - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}} \leq \prod_{k=1}^j \mathbb{P}_{\Gamma_{k,11} X_{j-k} \mathbf{q}} \Gamma_{k,11} X_{j-k} \mathbf{q} \leq C_{\Gamma,j,l} \mathbb{P}_j \left(\mathbf{X}_j - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}}.$$

The claim then follows thanks to Lemma B.3 and (i). □

In what follows, define $\mathbb{P}_j \left(\mathbf{X}_t - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}}$ to be a version of $\mathbb{P}_j \left(\mathbf{X}_t - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}}$ that is constructed using coefficient estimates from $\mathbf{t}\hat{\mathbf{H}}_1, \hat{\mathbf{H}}_2 \mathbf{u}$ but evaluated on the true innovations \mathbf{X}_t .

Proof of Theorem 4.1. If we introduce

$$\mathbb{R}\mathbb{F}_{h, \mathbf{p}, \mathbf{q}} : \Theta_{h,1} \mathbf{p} \mathbf{q} = \prod_{t=1}^n \mathbb{P}_{\mathbf{q}} \left(\mathbf{X}_t - \Theta_{h,1} \right) \mathbb{P}_{\mathbf{q}} \prod_{j=0}^h \frac{1}{n} \prod_{t=1}^j \mathbb{P}_j \left(\mathbf{X}_t - \Theta_{j,11} \right) \mathbb{P}_{\mathbf{q}},$$

then clearly

$$\begin{aligned} \mathbb{E} \mathbb{R} F_{h, p, q} & \mathbb{E} \mathbb{R} F_{h, p, q} \asymp \mathbb{E} \mathbb{R} F_{h, p, q} \mathbb{E} \mathbb{R} F_{h, p, q} \mathbb{E} \mathbb{R} F_{h, p, q} \mathbb{E} \mathbb{R} F_{h, p, q} \\ & I \quad II. \end{aligned}$$

To control II , we can observe

$$\begin{aligned} II & \asymp \Theta_{h, 1} n^{-1} \sum_{t=1}^n p_{1t} q \Theta_{h, 1} \mathbb{E} r p_{1t} q s \\ & \quad \sum_{j=0}^h \frac{1}{n} \sum_{t=1}^{n_j} p_j, X_{t:j;t}; r_t \mathbb{E} r v_j, X_{t:j;t}; r_s \\ & \asymp \Theta_{h, 1} \Theta_{h, 1} n^{-1} \sum_{t=1}^n p_{1t} q \Theta_{h, 1} n^{-1} \sum_{t=1}^n p_{1t} q \mathbb{E} r p_{1t} q s \\ & \quad \sum_{j=0}^h \frac{1}{n} \sum_{t=1}^{n_j} p_j, X_{t:j;t}; r_t \mathbb{E} r v_j, X_{t:j;t}; r_s \\ & \asymp \Theta_{h, 1} \Theta_{h, 1} n^{-1} \sum_{t=1}^n p_{1t} q \Theta_{h, 1} n^{-1} \sum_{t=1}^n p_{1t} q \mathbb{E} r p_{1t} q s \\ & \quad \sum_{j=0}^h \frac{1}{n} \sum_{t=1}^{n_j} p_j, X_{t:j;t}; r_t v_j, X_{t:j;t}; r_t \\ & \quad \sum_{j=0}^h \frac{1}{n} \sum_{t=1}^{n_j} v_j, X_{t:j;t}; r_t \mathbb{E} r v_j, X_{t:j;t}; r_s. \end{aligned}$$

The first two terms in the last bound are $O_P(p_1 q)$ since $\Theta_{h, 1} \Theta_{h, 1} = O_P(p_1 q)$, as discussed in Lemma B.3, and $n^{-1} \sum_{t=1}^n p_{1t} q \mathbb{E} r p_{1t} q s$ by a WLLN. For the other terms in the last sum above, we similarly note that

$$\frac{1}{n} \sum_{t=1}^{n_j} p_j, X_{t:j;t}; r_t v_j, X_{t:j;t}; r_t = O_P(p_1 q)$$

from Lemma B.4, while thanks again to a WLLN it holds

$$\frac{1}{n} \sum_{t=1}^{n_j} v_j, X_{t:j;t}; r_t \mathbb{E} r v_j, X_{t:j;t}; r_s = O_P(p_1 q).$$

Since h is fixed finite, this implies that $II = O_P(p_1 q)$.

Considering now I , we can write

$$\begin{aligned} I & \asymp \Theta_{h, 1} n^{-1} \sum_{t=1}^n p p_{1t} q p_{1t} q \sum_{j=0}^h \frac{1}{n} \sum_{t=1}^{n_j} p_j, X_{t:j;t}; p_t p_j, X_{t:j;t}; r_t \\ & I^1 \quad I^2. \end{aligned}$$

Since by assumption φ is a bump function, thus continuously differentiable over the range of t , by the mean value theorem

$$n^{-1} \int_{t-1}^t \rho \rho_{1t} q - \rho_{1t} q \asymp n^{-1} \int_{t-1}^t |q'| / \rho_{1t} - 1 t$$

for a sequence $t \in \mathcal{I}_n$ of evaluations of first-order derivative q' at values \bar{t} in the interval with endpoint t and ρ_t . One can use $|q'| \asymp C$ with a finite positive constant C , and by recalling that $\rho_{1t} - 1 t = \rho \Pi_1 - \mathbf{H}_1 q' W_{1t}$ one thus gets


$$n^{-1} \int_{t-1}^t \rho \rho_{1t} q - \rho_{1t} q \asymp C \frac{1}{n} \int_{t-1}^t \rho \Pi_1 - \mathbf{H}_1 q' W_{1t} \asymp C \Pi_1 - \mathbf{H}_1 \frac{1}{n} \int_{t-1}^t W_{1t} - 2 = o_P(\rho_{1t} q).$$

This proves that term I^1 is itself $o_P(\rho_{1t} q)$. Finally, to control I^2 , we use that by construction estimator \mathbf{H}_2 is composed of sufficiently regular functional elements i.e. B-spline estimates of order 1 or greater. Thanks again to the mean value theorem


$$\begin{aligned} \frac{1}{n} \int_{t-1}^t \rho_j \cdot X_{t-j:t} \rho_t - \rho_j \cdot X_{t-j:t} r_t &\asymp \frac{1}{n} \int_{t-1}^t \rho_j \cdot X_{t-j:t} \rho_t - \rho_j \cdot X_{t-j:t} r_t \\ &\asymp C_{\rho^j, j} \frac{1}{n} \int_{t-1}^t \rho_{1t} - 1 t \end{aligned}$$

for any fixed j and some $C_{\rho^j, j} > 0$. This holds since ρ_j is uniformly continuous by construction. Note that we have assumed that the nonlinear part of Π_2 belongs to a Hölder class with smoothness $s_j > 1$ (for simplicity, assume here that s is integer, otherwise a similar argument can be made). Then, even though $C_{\rho^j, j}$ depends on the sample, it is bounded above in probability for n sufficiently large. Following the discussion of term I^1 , we deduce that the last line in the display above is $o_P(\rho_{1t} q)$. As h is finite and independent of n , it follows that also I^2 is of order $o_P(\rho_{1t} q)$. \square

C Additional Plots



(a) 2



(b) 2

Figure 8: Simulation results for DGP 2¹ when considering τ in place of \cdot .




Figure 9: Estimated nonlinear regression functions for the narrative U.S. monetary policy variable. Contemporaneous (left side) and one-period lag (right side) effects are shown, linear and nonlinear functions. For comparison, linear VAR coefficients (dark gray) and the identity map (light gray, dashed) are shown as lines.




Figure 10: Comparison of histograms and shock relaxation function for a positive (left) and negative (right) shock in monetary policy. Original (blue) versus shocked (orange) distribution of the sample realization of π_t . The dashed vertical line is the mean of the original distribution, while the solid vertical line is the mean after the shock.

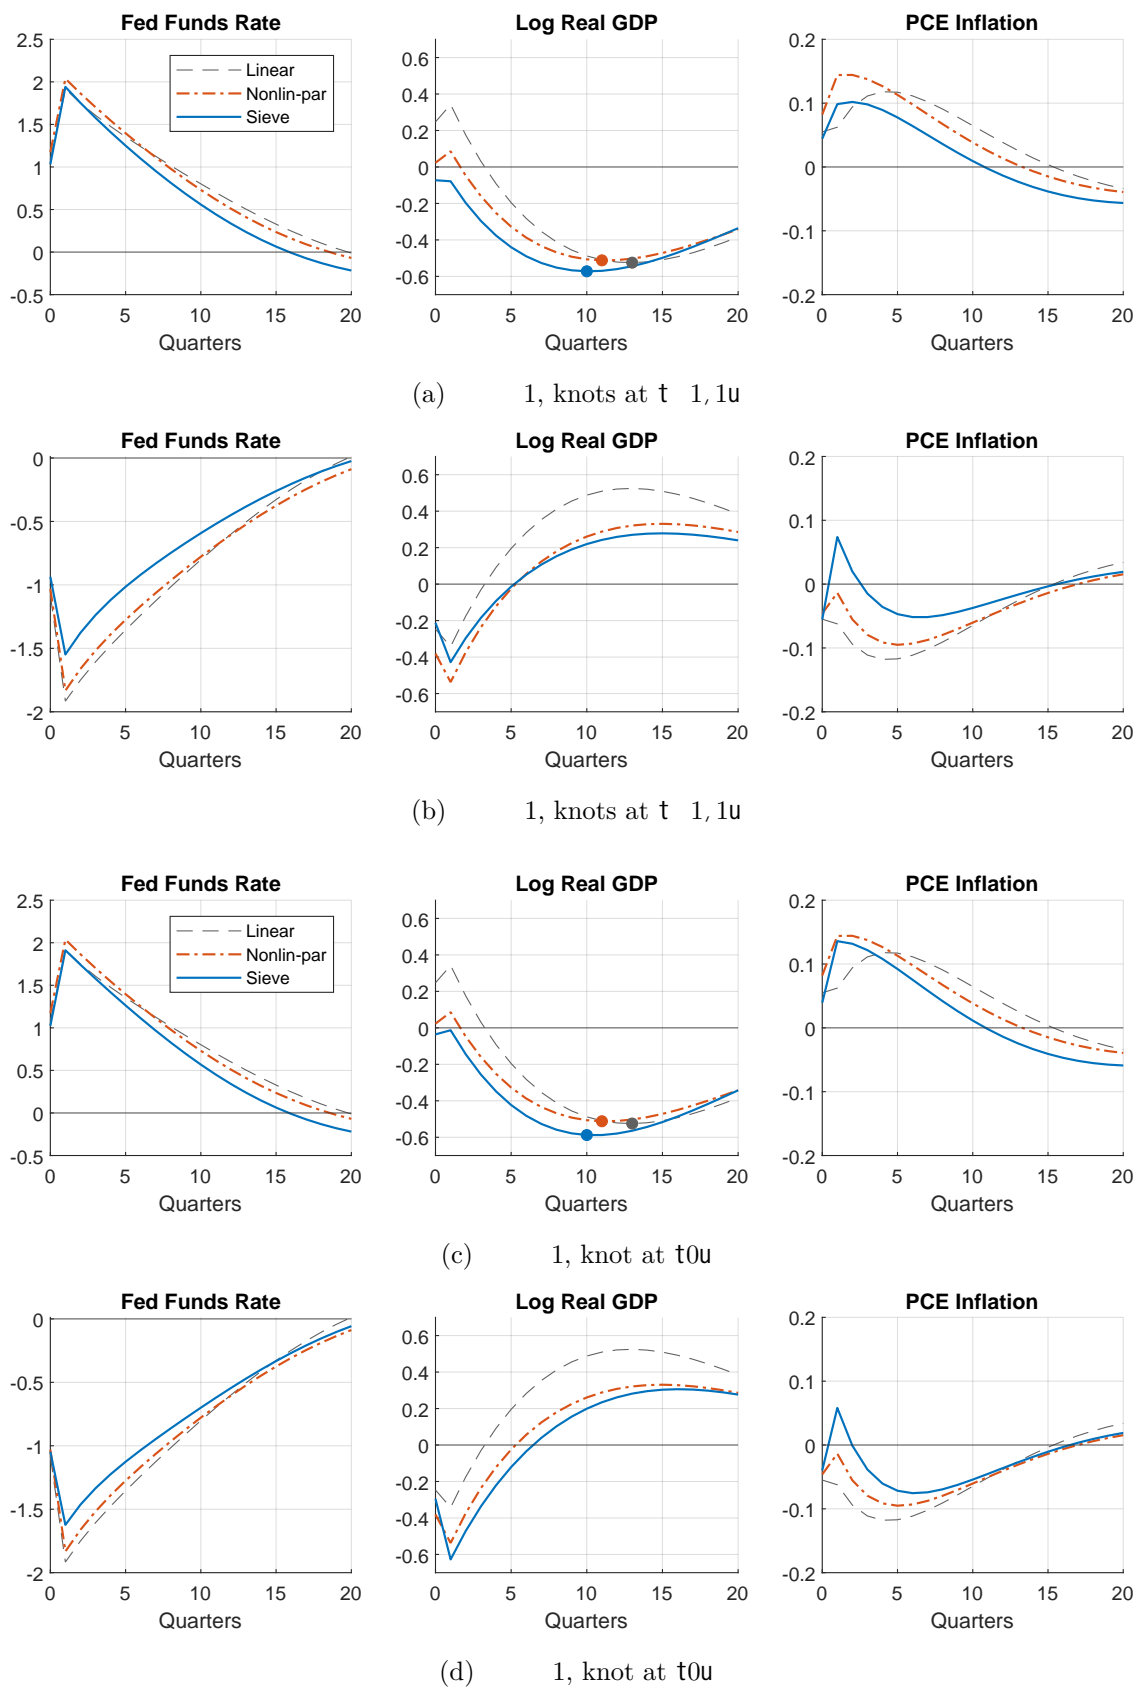


Figure 11: Robustness plots for U.S. monetary policy shock when changing knots compared to those used in Figure 6. Note that linear and parametric nonlinear responses do not change.

GDP




Figure 12: Relative changes in the GDP impulse responses function when the size of the shock is reduced from that used in Figure 6. The standard deviation of X_t is $\sigma_{\epsilon,1} = 0.5972$. Linear IRFs are re-scaled such that for all values of δ the linear response at $h = 0$ is one in absolute value. Nonlinear IRFs are re-scaled by δ times the linear response scaling factor.

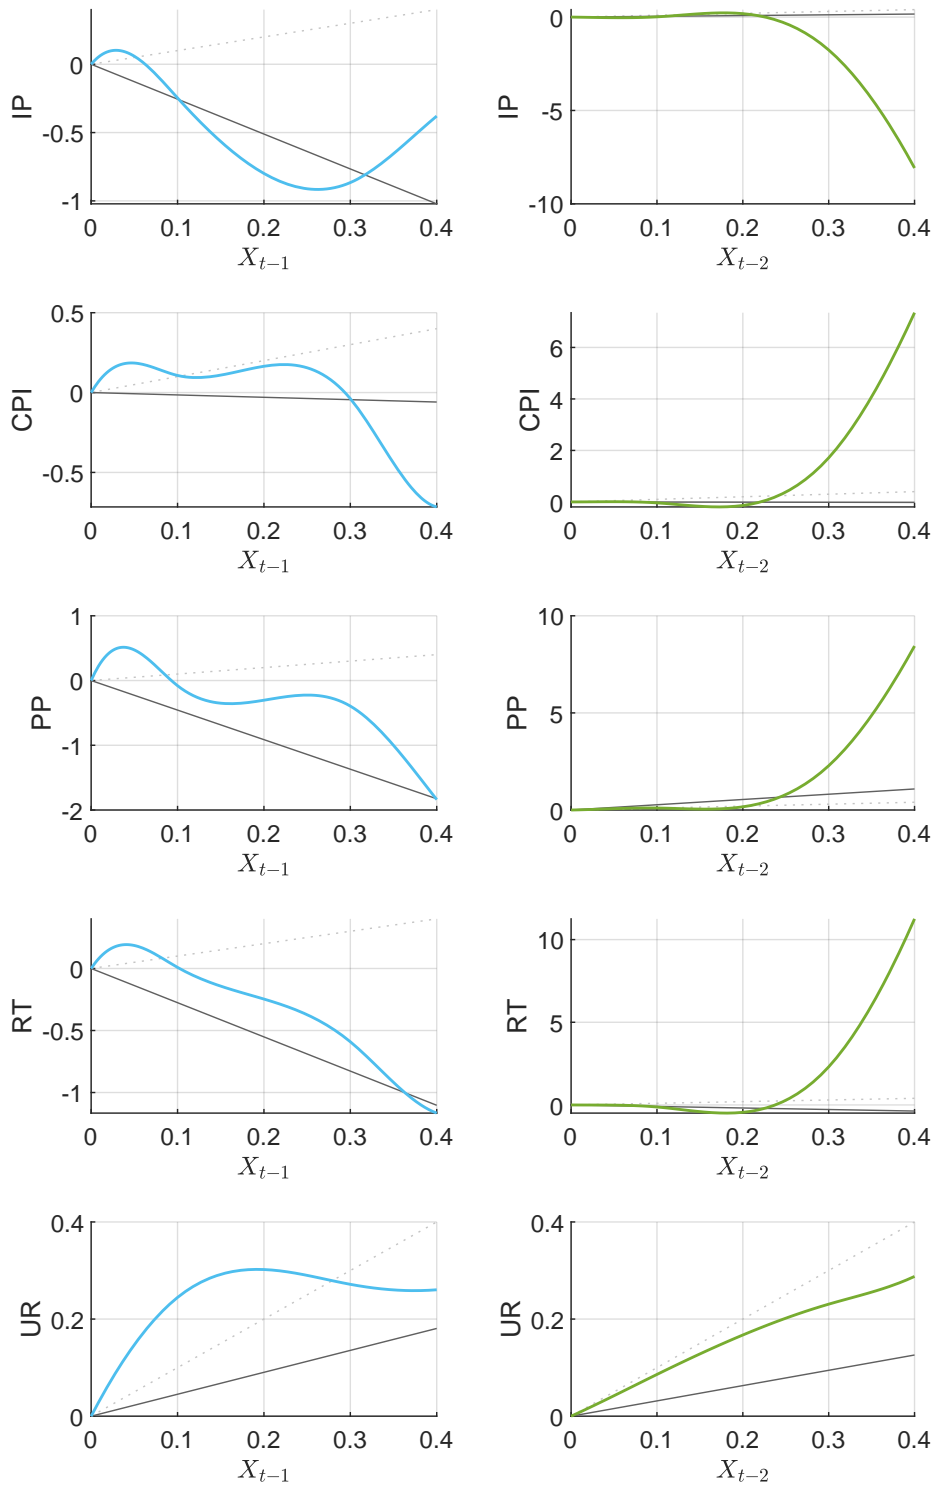


Figure 13: Estimated nonlinear regression functions for the 3M3M subjective interest rate uncertainty measure. One-period (left side) and two-period lag (right side) effects are shown, combining linear and nonlinear functions. For comparison, linear VAR coefficients (dark gray) and the identity map (light gray, dashed) are shown as lines.





Figure 14: [**Top**] Histograms and shock relaxation function for a one-standard-deviation shock in interest rate uncertainty. Original (blue) versus shocked (orange) distribution of the sample realization of x_t . The dashed vertical line is the mean of the original distribution, while the solid vertical line is the mean after the shock. [**Bottom**] Envelope (min-max) of shocked paths for one-standard-deviation impulse response.

IP

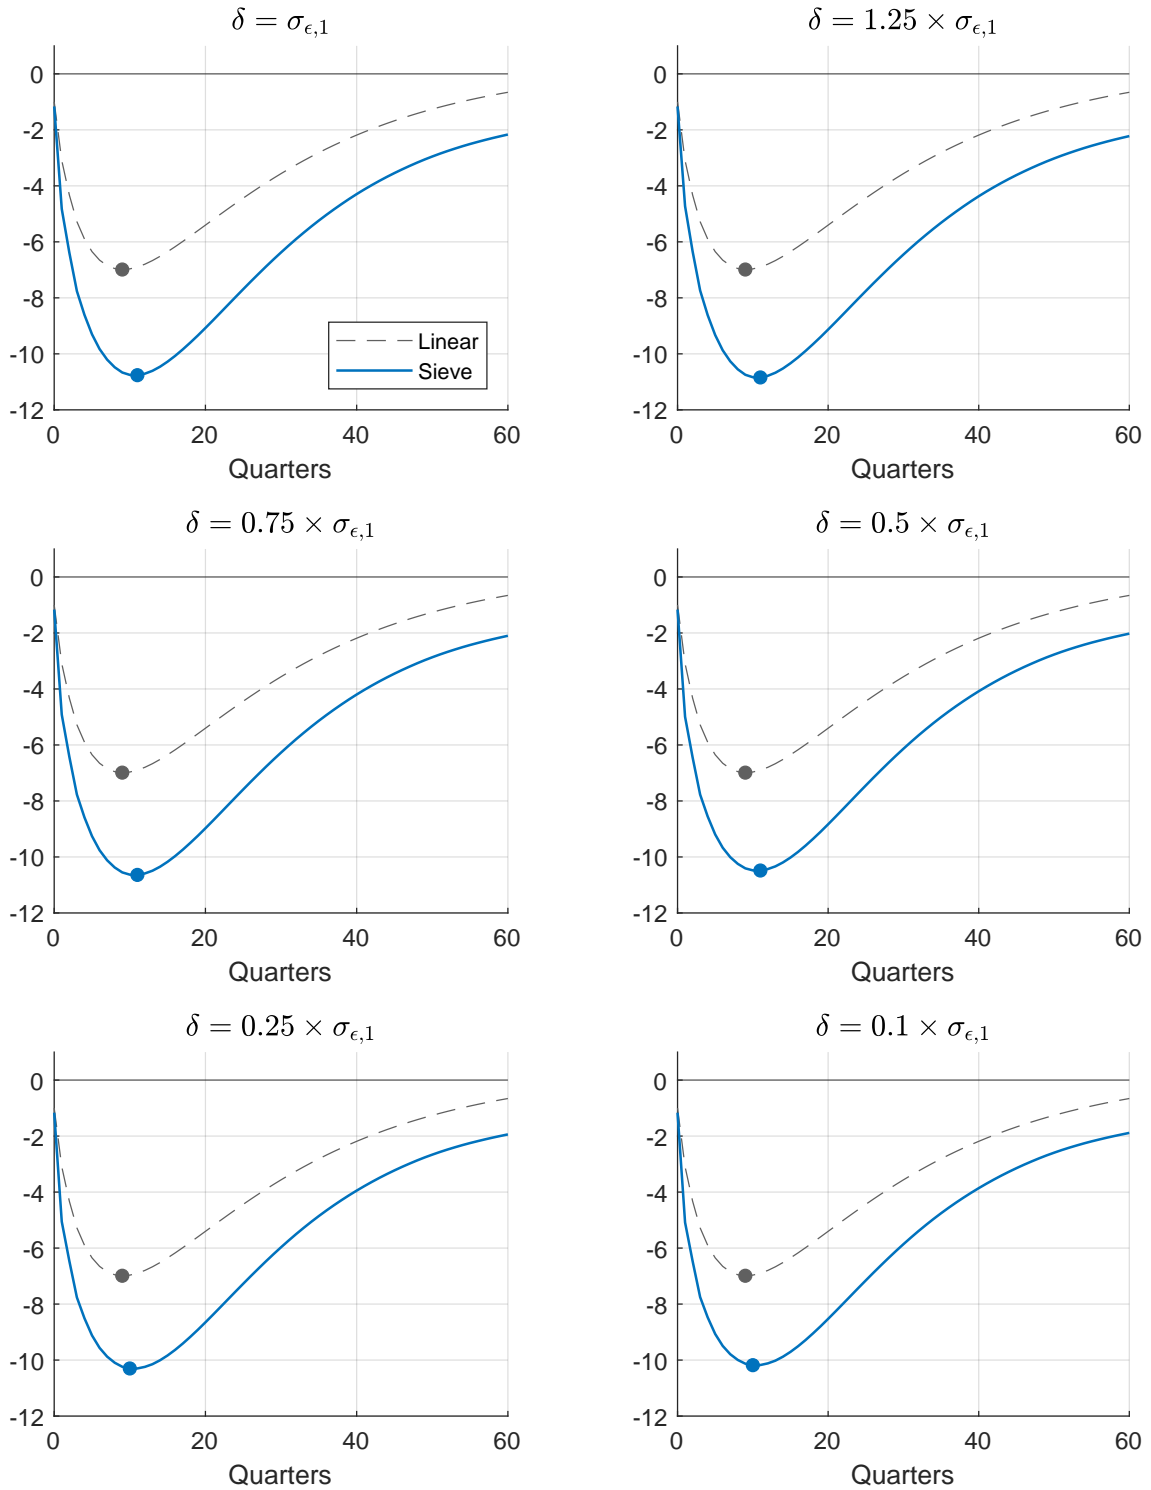


Figure 15: Relative changes in the industrial production impulse responses function when the size of the shock is reduced from that used in Figure 7. The standard deviation of ϵ_{1t} is $\sigma_{\epsilon,1} = 0.0389$. Linear IRFs are re-scaled such that for all values of δ the linear response at $h = 0$ is one in absolute value. Nonlinear IRFs are re-scaled by δ times the linear response scaling factor.

CPI

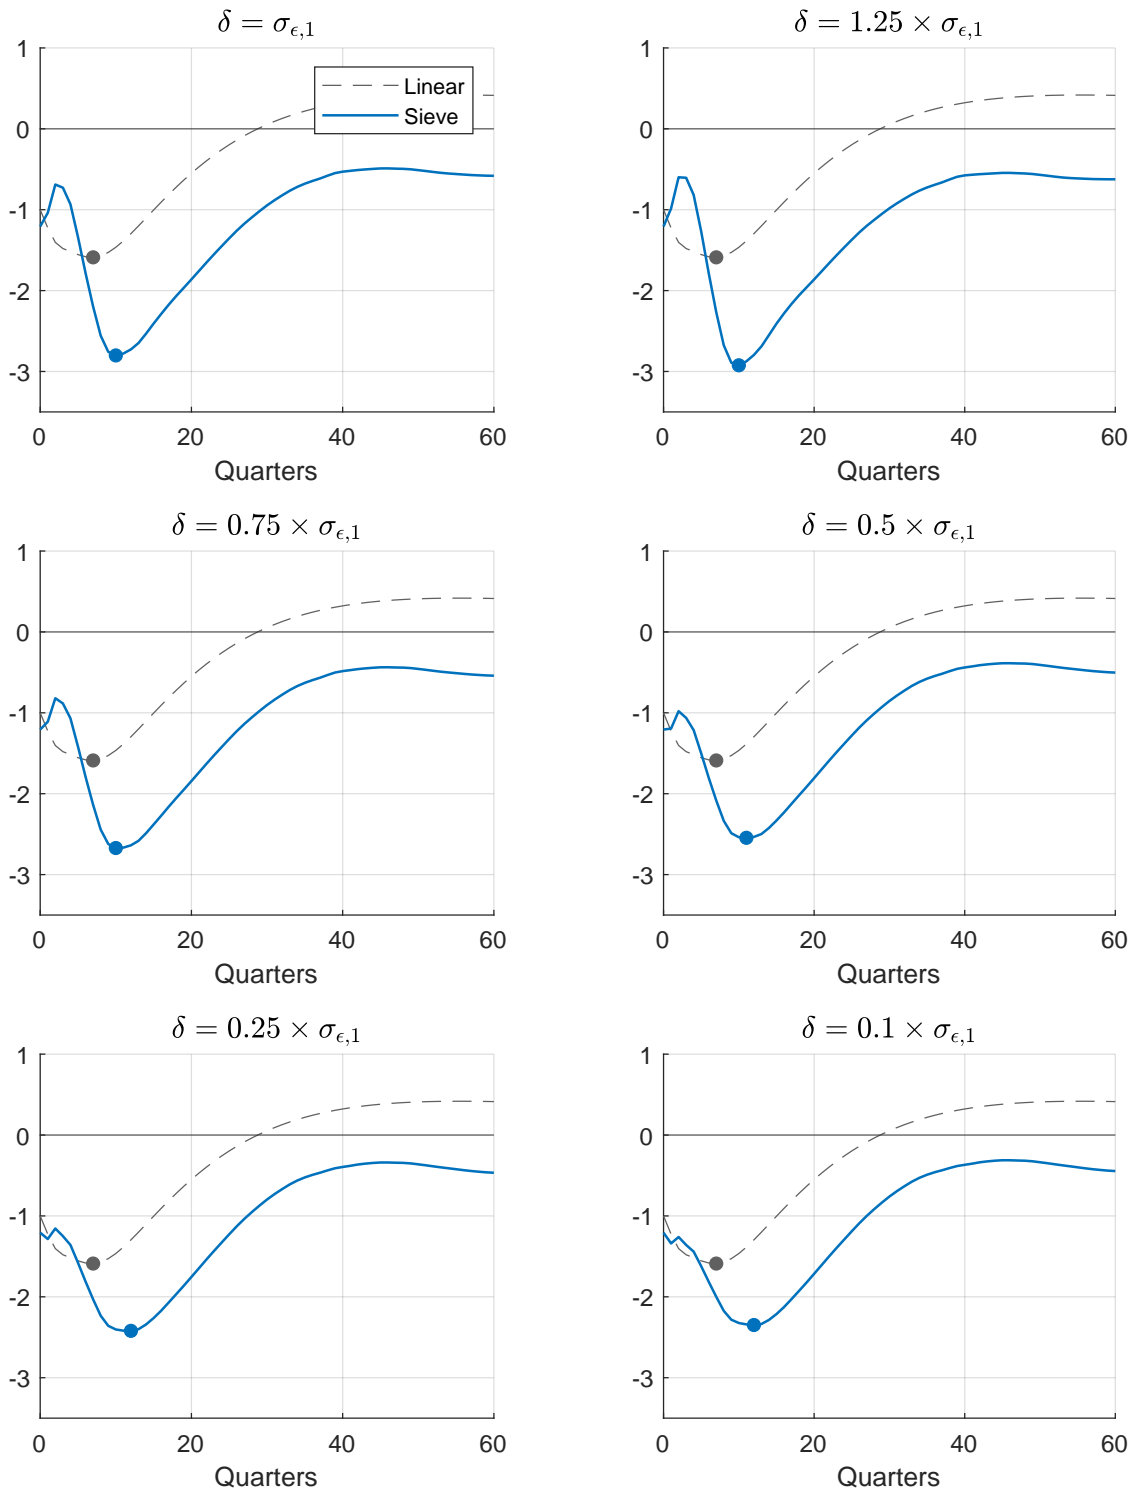


Figure 16: Relative changes in the CPI impulse responses function when the size of the shock is reduced from that used in Figure 7. The standard deviation of ϵ_{1t} is $\sigma_{\epsilon,1} = 0.0389$. Linear IRFs are re-scaled such that for all values of δ the linear response at $h = 0$ is one in absolute value. Nonlinear IRFs are re-scaled by δ times the linear response scaling factor.