

Estimating Nonlinear Business Cycle Mechanisms with Linear Vector Autoregressions: A Monte Carlo Study*

KARSTEN KOHLER[†]  and ROBERT CALVERT JUMP[‡] 

[†]*Economics Department, Leeds University Business School, Maurice Keyworth Building, Woodhouse, Leeds, LS2 9JT, United Kingdom (e-mail: k.kohler@leeds.ac.uk)*

[‡]*Institute for Political Economy, Governance, Finance and Accountability, University of Greenwich, Old Royal Naval College, Park Row, Greenwich, London, SE10 9LS, United Kingdom (e-mail: r.g.calvertjump@greenwich.ac.uk)*

Abstract

The paper investigates how well linear vector autoregressions (VARs) identify endogenous cycle mechanisms and cycle frequencies when the underlying process is a nonlinear limit cycle. We conduct Monte Carlo simulations with five nonlinear models in which cycles are driven by the interaction of two state variables. We find that while linear VARs quantitatively underestimate the strength of the interaction mechanism, they successfully identify the qualitative presence of a cycle mechanism in most cases (55%–100%). Our results further suggest that linear VARs are surprisingly successful at estimating cycle frequencies of nonlinear processes.

I. Introduction

Since the rise to popularity of real business cycle theory in the 1980s, and then New Keynesian DSGE models in the 1990s, macroeconomics has been dominated by the use of linearized difference equations to model business cycles. Edward C. Prescott (1986, p. 10) famously argued that,

‘some systems of low-order linear stochastic difference equations with a non-oscillatory deterministic part, and therefore no cycle, display key business cycle features [...] I thus do not refer to business cycles, but rather to business cycle phenomena, which are nothing more nor less than a certain set of statistical properties of a certain set of important aggregate time series.’

Recently, however, the idea of endogenous cycles has re-emerged in macroeconomic theory. Azariadis (2018) and Galí (2018), for example, call for new models with complex eigenvalues and endogenous propagation mechanisms to account for periodic fluctuations

JEL Classification numbers: C15, C32, E32.

*We thank the editor Jonathan Temple, an anonymous referee, Roberto Dieci, Giorgos Galanis, Alexander Guschanski, Achilleas Mantes, Christian Proaño, Ron Smith, Engelbert Stockhammer, Frank Westerhoff, and Rafael Wildauer for helpful comments. All errors are the authors’.

in output. An important stimulus has been empirical research on financial cycles that finds dominant frequencies of 13–16 years in aggregate financial variables such as private credit and house prices (Borio, 2014; Aikman, Haldane and Nelson, 2015; Rünstler and Vlekke, 2017; Stockhammer *et al.*, 2019; Strohsal, Proaño and Wolters, 2019). Borio (2014, p. 186) argues that these financial boom-bust episodes should be seen ‘as the result of endogenous forces that perpetuate (irregular) cycles’, rather than the result of random shocks propagated by linear mechanisms.

The re-emergence of endogenous cycles in macroeconomic theory is typified by the recent contribution of Beaudry, Galizia and Portier (2020). They show that various US macroeconomic indicators, such as hours worked per person and the unemployment rate, exhibit a pronounced cycle of around 10 years in length. They build a stochastic New Keynesian model with strategic complementarities in consumption, in which nonlinearities generate a limit cycle that matches the periodicities found in the data. Endogenous fluctuations are also a major theme in heterogeneous agents models of asset price dynamics (e.g. Brock and Hommes, 1997; Chiarella, 1992; Dieci and He, 2018). Similarly, the literature on behavioural New Keynesian models investigates endogenous cycle mechanisms based on strategic interaction (Branch and McGough, 2010; Calvert Jump and Levine, 2019).

Despite the re-emergence of endogenous cycles in macroeconomic theory and the existence of various types of nonlinear time series models, linear vector autoregressive (VAR) models still dominate macroeconometrics. Among other reasons, this is because they provide flexible approximations to a variety of time series processes. In particular, linear ARMA processes can be thought of as approximations to the Wold decomposition of regular stationary processes (Francq and Zakoian, 1998), and ARMA models can be found which approximate the spectra of any stationary process, nonlinear or otherwise (Brockwell and Davis, 2006, pp. 130–133).

However, while VAR models can be thought of as approximations to a variety of time series processes, their estimated coefficients will rarely correspond to the structural parameters of a nonlinear data-generating process (DGP). This is a simple implication of omitted variable (or omitted nonlinearity) bias. For example, Beaudry, Galizia and Portier (2017) use a Monte Carlo simulation to show that univariate linear autoregressive models fail to identify local instability when estimated on data generated by a nonlinear limit cycle process, despite the linear model providing (in principle) a good approximation to second moments. Specifically, Beaudry *et al.* (2017) find that the moduli of locally unstable processes, as defined by their linearization, are underestimated by linear time series models.

While we know that linear time series models underestimate the moduli of limit cycle processes, the extent to which they accurately estimate the endogenous cycle mechanisms and cycle lengths of nonlinear limit cycle processes is unknown. In this paper, we provide an assessment of the quantitative relevance of this type of omitted nonlinearity bias in an applied, small sample context. Using a series of Monte Carlo studies, we explore what the linear projection of a vector process $\{X_t\}$ on its own history – that is, a VAR model – can tell us about the linear part – that is, the Jacobian matrix – of canonical nonlinear business cycle models. Specifically, we study whether linear VAR models are a reliable method of identifying the cyclical interaction mechanisms posited by a

broad class of endogenous business cycle models. By interaction mechanism, we mean a causal mechanism in which two state variables interact such that the first variable is increasing in the lags of the second variable, while the second variable is decreasing in the lags of the first. In the first-order, bivariate models we focus on, such an interaction mechanism is a critical condition for complex eigenvalues, and thus genuine oscillations (Stockhammer *et al.*, 2019). Such a set-up can be found in Minskyan financial cycle models (e.g. Asada, 2001), New Keynesian financial accelerator models (e.g. Kiyotaki and Moore, 1997), or models of housing cycles (Dieci and Westerhoff, 2012; Dieci and Westerhoff, 2016). Given the parsimony and flexibility of linear VAR models, alongside the re-emergence of limit cycle models in macroeconomic theory, it is of great practical use to know how well linear VARs identify such interaction mechanisms in the presence of nonlinearities.¹

Our Monte Carlo studies consist of a sequence of linear VAR models estimated on artificial data generated from five nonlinear business cycle models. Using these data, we summarize the information contained in the Jacobian of nonlinear business cycle models in two ways. First, we compare the eigenvalues of estimated linear VAR models to the eigenvalues of the underlying nonlinear DGP. And second, we ask if linear VAR models can detect the interaction mechanisms that generate the eigenvalues of the DGP.

Our first contribution is to confirm and extend the finding of Beaudry *et al.* (2017) to the multivariate case, that linear autoregressive models generally fail to identify the presence of local instability when estimated on data generated from nonlinear limit cycle models. Our second contribution is to document what is, to the best of our knowledge, a novel finding: linear VARs successfully identify the qualitative presence of a cyclical interaction mechanism in the majority of cases. When the DGP is a first-order system with white noise errors, a linear VAR(1) accomplishes cycle detection rates of 100%. In more realistic settings with serially correlated shock processes and lag choice based on algorithms, cycle detection rates range between 55% and almost 100%. We also find that linear VARs generate relatively few false positives: they incorrectly identify cycle mechanisms in a benchmark nonlinear model in less than 1% of cases. Our final contribution is to document that VARs are surprisingly successful at estimating cycle frequencies of nonlinear processes. Although local instability and the magnitude of the interaction mechanism tend to be underestimated, the downward bias in the estimated length of the limit cycle is typically smaller. Thus, as in Fernández-Villaverde and Rubio-Ramirez (2005), we find that models with omitted nonlinearities can still provide accurate point estimates for certain parameters of nonlinear DGPs – in this case, even when the nonlinearities produce limit cycles.

Our paper is related to a number of different research programmes in applied macroeconomics. First and foremost, it is related to the extant literature on nonlinear cycle models (Brock and Hommes, 1997; Bischi *et al.*, 2001; Wegener, Westerhoff and Zaklan, 2009; Branch and McGough, 2010; Dieci and Westerhoff, 2012; Dieci and

¹ Focussing on the interaction mechanism is especially relevant when researchers are not primarily interested in the functional form of the DGP but rather in the interaction mechanism itself. For example, recent studies of financial cycles (Borio, 2014; Rünstler and Vlekke, 2017; Stockhammer *et al.*, 2019; Strohsal *et al.*, 2019; Juselius and Drehmann, 2020) focus on interactions between financial variables and the real economy, but do not present strong theoretical reasons to pin down a specific functional form for these interactions.

He, 2018; Calvert Jump and Levine, 2019; Beaudry *et al.*, 2020), as well as empirical work on business and financial cycles with linear estimation techniques (Stockhammer *et al.*, 2019; Strohsal *et al.*, 2019). Stockhammer *et al.* (2019) investigate finance-driven business cycles and find evidence consistent with the presence of cyclical interaction mechanisms between GDP and interest rates as well as corporate debt in several advanced economies. Similarly, Strohsal *et al.* (2019) study the bivariate interaction of financial and real cycles in the frequency domain using cross-spectra between GDP and financial variables, such as credit and house prices. Both Stockhammer *et al.* (2019) and Strohsal *et al.* (2019) rely on linear VARs to investigate financial–real interactions and disregard the potential issue of nonlinearity.

Second, it is related to the literature on nonlinear time series models, such as threshold or smooth transition (vector) autoregressive models (Teräsvirta, Tjøstheim and Granger, 2010; Kilian and Lütkepohl, 2017; Teräsvirta, 2018), which add flexible nonlinearities to workhorse linear models. Most closely related is the study by Beaudry *et al.* (2017) that finds that the moduli of locally unstable processes are underestimated by linear models. Finally, our paper is also related to the literature on linear versus nonlinear estimation in DSGE models. In a seminal study, Fernández-Villaverde and Rubio-Ramírez (2005) use a Monte Carlo study to compare the estimation properties of linear versus nonlinear approximations to data generated from a DSGE. They conclude that, despite nonlinear approximations providing a considerably better fit to their simulated data, models with omitted nonlinearities can still provide accurate point estimates for certain parameters of nonlinear DGPs and should not be abandoned. However, they did not consider DGPs that produce limit cycles, which is the focus of our paper.

The paper is structured as follows. Section II discusses cycle mechanisms in linear VAR models and how bias can arise from omitted nonlinearities alongside the more familiar Hurwicz and serial correlation biases. Section III introduces a linear benchmark as well as the five nonlinear business cycle models that will be considered as DGPs in our Monte Carlo studies. Section IV describes the Monte Carlo design. Section V presents the main results and various robustness tests. Section VI concludes.

II. Estimating 2D cycle mechanisms with linear VARs

Cycle mechanisms in VAR models

Following the approach proposed in Stockhammer *et al.* (2019), consider a first-order bivariate linear VAR model,

$$\begin{bmatrix} y_t \\ f_t \end{bmatrix} = \begin{bmatrix} a_1 & a_2 \\ b_1 & b_2 \end{bmatrix} \begin{bmatrix} y_{t-1} \\ f_{t-1} \end{bmatrix} + \begin{bmatrix} u_{yt} \\ u_{ft} \end{bmatrix}, \quad (1)$$

in which y_t and f_t denote two arbitrary macroeconomic variables – for example, a measure of real activity and debt as in Kiyotaki and Moore (1997), flow and stock variables as in Beaudry *et al.* (2020), or house prices and the housing stock as in (Dieci and Westerhoff, 2012; Dieci and Westerhoff, 2016) – and u_{yt} and u_{ft} are error terms. As is well known, the presence of a pair of complex conjugate eigenvalues of the coefficient matrix

in (1) generates oscillatory dynamics. A necessary condition for complex eigenvalues in this simple model is,

$$a_2b_1 < 0, \tag{2}$$

that is, the off-diagonal elements have opposite sign. To see this, recall that the eigenvalues of the coefficient matrix in (1) are the roots of the characteristic equation $\lambda^2 - \lambda Tr + Det = 0$, where $Tr = a_1 + b_2$ and $Det = a_1b_2 - a_2b_1$ are the trace and determinant of the coefficient matrix, respectively. A necessary and sufficient condition for the roots to be complex is a negative discriminant $Tr^2 - 4Det$. Simplifying $Tr^2 - 4Det = (a_1 + b_2)^2 - 4(a_1b_2 - a_2b_1)$ to $(a_1 - b_2)^2 + 4(a_2b_1)$ demonstrates that $Tr^2 - 4Det$ can only be negative if $a_2b_1 < 0$.

As argued in Stockhammer *et al.* (2019), this provides a straightforward intuition that is useful for structuring theories of endogenous business cycles: oscillations in (1) stem from an interaction mechanism between y_t and f_t in which an increase in one variable induces an increase in the second variable, which in turn drags down the first. In a theory of finance-driven business cycles, for example, an increase in output encourages firms to take on more debt, but the resulting debt burden might provoke a retrenchment of investment and economic activity, resulting in oscillatory dynamics in output and debt. Importantly, although the condition in (2) is only a necessary one, it is nonetheless of critical interest from a theoretical perspective, as it embodies the main driving force of endogenous cycles. In the limit, if $a_2 = b_1 = 0$, the only causal linkage between y_t and f_t would be via the error terms. For this reason, we summarize the information contained in our VAR models by the interaction mechanism in (2).

In addition, given the recent interest in the length of business and financial cycles (Borio, 2014; Rünstler and Vlekke, 2017; Stockhammer *et al.*, 2019; Strohsal *et al.*, 2019), we also analyse the eigenvalues λ of the coefficient matrix, from which we obtain the implied cycle length,

$$L = \frac{2\pi}{\omega} = \frac{2\pi}{\arccos\left(\frac{re(\lambda)}{|\lambda|}\right)} = \frac{2\pi}{\arctan\left(\frac{im(\lambda)}{re(\lambda)}\right)}, \tag{3}$$

where ω is the frequency, $|\lambda| = \sqrt{re^2 + im^2}$ is the modulus, and re and im are the real and imaginary part of a complex eigenvalue.

Sources of estimation bias

Suppose first that the two-dimensional DGP is linear. Two familiar sources of bias may arise in a VAR in this case Hamilton (1994, chap. 8).² If the error terms are serially uncorrelated, the OLS estimator of the coefficients would suffer from the standard Hurwicz bias in which correlation between error terms and lagged dependent variables exists in a finite sample. This bias disappears asymptotically. If the error terms are

²This paper focuses on stationary processes. See Abadir, Hadri and Tzavalis (1999), Doornik, Nielsen and Rothenberg (2003), and Hamilton (1994, chap. 18) for sources of bias in non-stationary VARs.

serially correlated, the estimator becomes inconsistent, that is, the bias does not appear asymptotically.

Suppose next that the DGP is nonlinear and the Jacobian matrix evaluated at the equilibrium is given by

$$J = \begin{bmatrix} \tilde{a}_1 & \tilde{a}_2 \\ \tilde{b}_1 & \tilde{b}_2 \end{bmatrix}. \quad (4)$$

A linear VAR model estimated on data generated from this model will be consistent for its linear projection, but the VAR coefficient matrix in (1) will not, in the general case, coincide with the Jacobian in (4). To fix intuition on this third source of bias, consider a simple nonlinear AR(1) model for some variable x_t , in which $E[x_t|x_{t-1}] = g(x_{t-1})$. Suppose for example that,

$$x_t = \alpha + \beta x_{t-1} + \delta f(x_{t-1}) + \epsilon_t, \quad (5)$$

in which $f(x_{t-1})$ is the nonlinear part of the conditional expectation function, and $f'(x_{t-1}) = 0$ when evaluated at the relevant steady state. If instead a linear AR(1) model,

$$x_t = a + bx_{t-1} + \eta_t, \quad (6)$$

is estimated, then the functional form misspecification takes the form of omitted variable bias, which is described by the standard formula,

$$\text{bias} = E[b] - \beta = \delta \frac{\text{Cov}[x_{t-1}, f(x_{t-1})]}{\text{Var}[x_{t-1}]}. \quad (7)$$

Looked at in this way, the effect of omitted nonlinearities depends on the covariance between the lagged dependent variable and the omitted nonlinear part of the DGP.

Omitting a (relevant) nonlinearity from the VAR model in (1) means that error terms u_{yt} and u_{ft} are correlated with the lagged dependent variables y_{t-1} and f_{t-1} . This correlation does not go to zero as the sample size increases. As a result, we cannot expect OLS estimators of a_1 , a_2 , b_1 and b_2 in (1) to be consistent for \tilde{a}_1 , \tilde{a}_2 , \tilde{b}_1 and \tilde{b}_2 in (4). While the main interest of this paper lies in assessing this omitted nonlinearity bias, all three sources of biases might be present in linear VAR models estimated on data generated by nonlinear processes. Therefore, our Monte Carlo design will evaluate their relative importance.

III. Endogenous business cycle models

We conduct our analysis using data generated from five separate models. To provide a benchmark environment in which we expect linear VARs to behave well, we first consider the classic linear business cycle model due to Samuelson (1939). We then use the recent nonlinear business cycle model in Beaudry *et al.* (2020) as our central test of the ability of linear VARs to estimate cycle mechanisms. Lastly, four further nonlinear limit cycle models and a nonlinear model that does not generate limit cycles are examined to ensure robustness.




Figure 1. Simulation of SM39

Notes: Simulation of the system $C_t = 0.4C_{t-1} + 0.4I_{t-1}$, $I_t = -1.2C_{t-1} + 0.8I_{t-1}$. Upper panel: deterministic version. Middle panel: stochastic version with white noise errors $\epsilon_t \sim N(0, 1)$. Lower panel: stochastic version with serially correlated error terms $u_t = 0.8u_{t-1} + \epsilon_t$

A linear benchmark model

Consider first the classic linear business cycle model due to Samuelson (1939), henceforth SM39. The structure of this model is straightforward and is outlined formally in Online Appendix S1 (section A). In SM39, interactions between the Keynesian multiplier effect on consumption and the accelerator effect on investment give rise to business cycles. An increase in investment boosts consumption via the multiplier, which in turn, given the lag structure of the model, drags down investment. The upper panel of Figure 1 plots damped oscillations in a deterministic version of SM39, while the lower two panels present simulated dynamics from a stochastic version of the model with white noise and serially correlated disturbances, respectively. By comparing the top panel with the bottom two panels, it can be seen that fluctuations eventually fade away in the deterministic model, but are sustained by the addition of exogenous shocks.

A nonlinear benchmark model

Now consider the recent model in Beaudry *et al.* (2020), who propose a nonlinear endogenous business cycle model in which fluctuations are shock-independent. Local instability is introduced through a demand complementarity: the expansionary effects of an increase in spending by some agents improve credit availability in the aggregate,

allowing more agents to increase their spending. As the system moves away from its steady state, demand complementarities become weaker. Far away from the locally unstable steady state, the system thus becomes locally stable again. Beaudry *et al.* (2020) introduce this mechanism into a New Keynesian business cycle model with intertemporally optimizing forward-looking agents. They show that limit cycles can emerge even in an environment with rational agents.

We use the simplified reduced-form system proposed in Beaudry, Galizia and Portier (2016); Beaudry *et al.* (2020), henceforth BGP20, as the baseline nonlinear business cycle model for our Monte Carlo experiments. The model is summarized by,

$$I_t = \alpha_0 - \alpha_1 X_t + \alpha_2 I_{t-1} + F(I_t), F(0) = 0, \tag{8}$$

$$X_t = (1 - \delta)X_{t-1} + I_{t-1}, \quad \alpha_1, \alpha_2, \delta, F'(I_t) \in (0, 1), \tag{9}$$

where X_t denotes a capital or durable consumption good and I_t denotes investment in that good. Investment depends negatively on the stock of accumulated assets and positively on feedback effects. The term $F(I_t)$ captures positive demand externalities.

Beaudry *et al.* (2016) show that as $F'(I^*)$ increases from zero to one, a limit cycle around a unique equilibrium emerges via a Neimark–Sacker bifurcation. This limit cycle is stable if $F'''(I^*)$ is sufficiently negative. These restrictions on the properties of $F(I_t)$ imply a sigmoid-shaped function, and Beaudry *et al.* (2016) suggest the logistic function $F(I) = \frac{1}{1+e^{-I}}$ as a candidate. Using this functional form, and eliminating the contemporaneous X_t in (8), yields,

$$I_t = \alpha_0 - \alpha_1(1 - \delta)X_{t-1} + (\alpha_2 - \alpha_1)I_{t-1} + \frac{\alpha_3}{1 + e^{-I_t}}, \quad \alpha_3 > 0 \tag{10}$$

$$X_t = (1 - \delta)X_{t-1} + I_{t-1}, \tag{11}$$

with $\alpha_0 = -\alpha_3/2$ so that $F(0) = 0$. The steady state (I^*, X^*) is derived in Online Appendix S1 (section B). The Jacobian evaluated at this steady state is given by,

$$J_{BGP20}(I^*, X^*) = \begin{bmatrix} \frac{(\alpha_2 - \alpha_1)(1 + e^{-I^*})^2}{(1 + e^{-I^*})^2 - \alpha_3 e^{-I^*}} & -\frac{\alpha_1(1 - \delta)(1 + e^{-I^*})^2}{(1 + e^{-I^*})^2 - \alpha_3 e^{-I^*}} \\ 1 & 1 - \delta \end{bmatrix}. \tag{12}$$

The coefficients on the off-diagonal of (12) exhibit opposite signs as long as α_3 is sufficiently small, so that the critical condition for a cycle mechanism (2) is satisfied. The limit cycle in this model is driven by an interaction mechanism between expenditures and the stock of assets, where strong feedback effects of spending near the steady state create local instability, which is contained by a weakening of these feedback effects for extreme values of spending. Figure 2 illustrates these cyclical dynamics for a deterministic and stochastic version of the model, respectively.

To assess the robustness of the simulation results to different limit cycle models, we consider four further endogenous business cycle models that differ in their functional




Figure 2. Simulation of BGP20

Notes: Simulation of the system (10)–(11) with $\alpha_0 = -1$, $\alpha_1 = 0.4$, $\alpha_2 = 0.9$, $\alpha_3 = 2$, $\delta = 0.4$. Upper panel: deterministic version. Middle panel: stochastic version with white noise errors $\epsilon_t \sim N(0, 1)$. Lower panel: stochastic version with serially correlated error terms $u_t = 0.8u_{t-1} + \epsilon_t$

forms – see Online Appendix S1 (section A) for a detailed discussion of these models. Tables 1 and 2 provide an overview of the DGPs, their main properties and the parameterizations used in the simulations.³ Figure 3 illustrates the dynamics of these business cycle models for the deterministic and stochastic cases. All systems display periodic endogenous fluctuations with different cycle frequencies, which are rendered less regular and more persistent by additive serially correlated noise.

IV. Monte Carlo design

For the Monte Carlo experiments, we use stochastic versions of the DGPs reported in Tables 1 and 2. We first consider the case where shock processes are serially uncorrelated and the corresponding lag order p of the VAR is correctly specified ($p = 1$). Second, in a more common scenario, shock processes follow an AR(1) process as assumed in Beaudry *et al.* (2020) and other benchmark macroeconomic models (e.g. Smets and Wouters, 2003) (see Online Appendix S1 (section C) for a discussion of how we evaluate the interaction

³With the exception of BGP20, all parameterizations are taken from the respective studies and meet the conditions for endogenous cycle dynamics. In BGP20 no parameterization is provided. We chose a parameterization that satisfies the conditions for a Neimark–Sacker bifurcation discussed in Beaudry *et al.* (2016). Figure 2 confirms the presence of a stable limit cycle for this parameterization.

TABLE 1
Overview of simulated DGPs, part 1

	Samuelson (1939) (SM39)	Baudry et al. (2016, 2020) (BGP20)	Bischi et al. (2001) (BDRS01)
Type	Keynesian multiplier-accelerator	New Keynesian w/ demand complementarities	Kaldorian multiplier-accelerator
Interaction mechanism	Consumption (C)–investment (I)	Expenditure (I)–asset stock (X)	Output (Y)–capital stock (K)
Dynamics	Damped oscillations	Limit cycle	Limit cycle
Type of nonlinearity	None	Sigmoid function: logistic	Sigmoid function: arctangent
Reduced form	$C_t = c(C_{t-1} + I_{t-1} + G)$; $I_t = \beta[c(C_{t-1} + I_{t-1} + G) - C_{t-1}]$	$I_t = \alpha_0 - \alpha_1(1 - \delta)X_{t-1} + (\alpha_2 - \alpha_1)I_{t-1} + \frac{\alpha_0}{1 + e^{-\alpha_1 I_{t-1}}}$; $X_t = (1 - \delta)X_{t-1} + I_{t-1}$	$Y_t = Y_{t-1} + \alpha[\sigma\mu + \gamma(\sigma\mu/\delta - K_{t-1}) + \arctan(Y_{t-1} - \mu) - \sigma Y_{t-1}]$; $K_t = (1 - \delta)K_{t-1} + \sigma\mu + \gamma(\sigma\mu/\delta - K_{t-1}) + \arctan(Y_{t-1} - \mu)$
Parameterization	$c = 0.4$ $\beta = 2$ $G = 0$	$\alpha_0 = -1$ $\alpha_1 = 0.4$ $\alpha_2 = 0.9$ $\alpha_3 = 2$	$\alpha = 1.2$ $\gamma = 0.6$ $\delta = 0.2$ $\mu = 10$
Jacobian (at steady state)	$\begin{bmatrix} 0.4 & 0.4 \\ -1.2 & 0.8 \end{bmatrix}$	$\delta = 0.4$ $\begin{bmatrix} 1 & -0.48 \\ 1 & 0.6 \end{bmatrix}$	$\sigma = 0.4$ $\begin{bmatrix} 1.72 & -0.72 \\ 1 & 0.2 \end{bmatrix}$
Cycle condition	$a_2 b_1 = -0.48$	$a_2 b_1 = -0.48$	$a_2 b_1 = -0.72$
Eigenvalue (λ)	$\lambda = 0.6 \pm 0.663i$	$\lambda = 0.8 \pm 0.663i$	$\lambda = 0.96 \pm 0.377i$
modulus ($ \lambda $)	$ \lambda = 0.894$	$ \lambda = 1.039$	$ \lambda = 1.032$
cycle length (L)	$L = 7.52$	$L = 9.08$	$L = 16.8$

Notes: See Online Appendix S1 (section A) for a detailed discussion of models. $L = \frac{2\pi}{\omega} = \frac{2\pi}{\arccos\left(\frac{\text{re}(\lambda)}{|\lambda|}\right)}$.

TABLE 2
Overview of simulated DGPs, part 2

	Neubert and Kot (1992) (NK92)	Wegener et al. (2009) (WWZ09)	Dieci and Westerhoff (2012) (DWF12)
Type			
Interaction mechanism	Predator-prey dynamics	Metzlerian inventory cycles with heterogeneous expectations	Asset price speculation in housing market
Dynamics	Predator (y) and prey (x) population	Inventories (S)—expected sales (U)	Housing stock (z)—house prices (p)
Type of nonlinearity	Limit cycle	Limit cycle	Limit cycle
Reduced form	Second-order polynomial	Rational function	Rational function
	$y_t = cx_{t-1}y_{t-1}$; $x_t = (r + 1)y_{t-1} - rx_{t-1}^2 - cx_{t-1}y_{t-1}$	see Online Appendix S1 (section A)	$z_t = ep_{t-1} + dz_{t-1}$; $p_t = (1 - c - e)p_{t-1} + \frac{f p_{t-1} - g h p_{t-1}^2}{1 + h p_{t-1}^2} - dz_{t-1}$
Parameterization	$c = 2.1$ $r = 3.1$	$b = 0.75$ $c = 0.3$ $d = 1$ $f = 0.1$ $k = 0.1$ $\bar{C} = 30$ $\bar{T} = 10$	$c = 0.05$ $d = 0.98$ $e = 0.5$ $f = 0.08$ $g = 1$ $h = 1$
Jacobian (at steady state)	$\begin{bmatrix} 1 & 1.624 \\ -1 & -0.476 \end{bmatrix}$	$\begin{bmatrix} 0.848 & -0.253 \\ 0.975 & 0.975 \end{bmatrix}$	$\begin{bmatrix} 0.98 & 0.5 \\ -0.98 & 1.25 \end{bmatrix}$
Cycle condition	$a_2 b_1 = -1.624$	$a_2 b_1 = -0.246$	$a_2 b_1 = -0.49$
Eigenvalue (λ), modulus ($ \lambda $)	$\lambda = 0.262 \pm 1.039i$ $ \lambda = 1.071$	$\lambda = 0.911 \pm 0.492i$ $ \lambda = 1.036$	$\lambda = 1.115 \pm 0.687i$ $ \lambda = 1.310$
cycle length (L)	$L = 4.75$	$L = 12.7$	$L = 11.4$

Notes: See Online Appendix S1 (section A) for a detailed discussion of models. $L = \frac{2\pi}{\omega} = \frac{2\pi}{\arccos\left(\frac{\operatorname{re}(\lambda)}{|\lambda|}\right)}$.

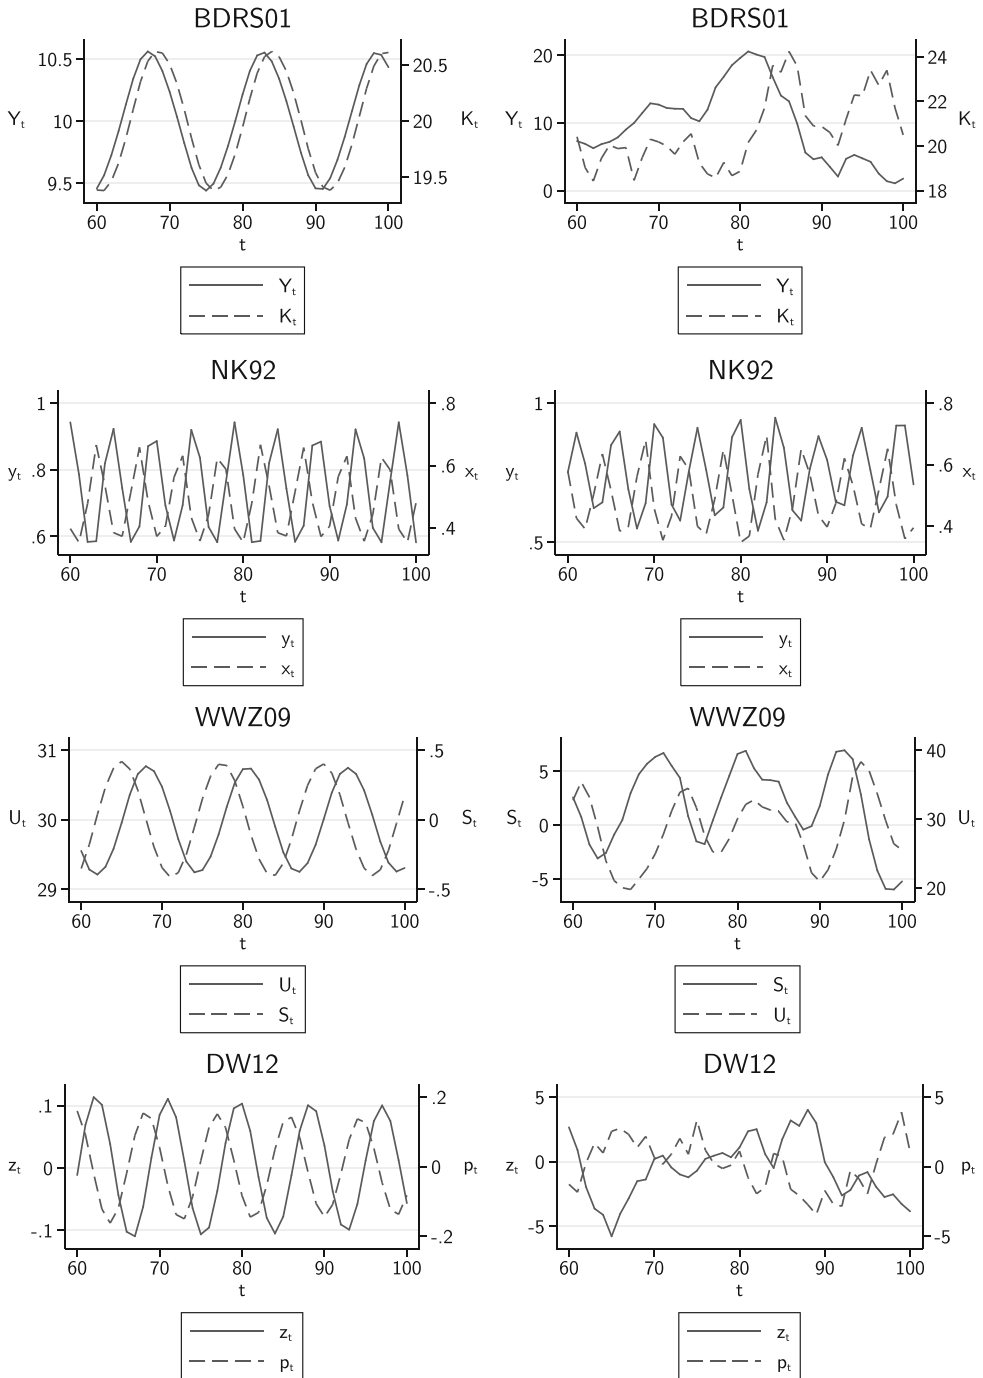


Figure 3. Simulated DGPs, deterministic (left) and stochastic (right)
 Notes: Parameterizations as in Tables 1 and 2. Stochastic versions (right panel) with serially correlated error terms $u_t = 0.8u_{t-1} + \epsilon_t$, with $\epsilon_t \sim N(0, 1)$. For NK92, the initial condition was set to $y_0 = x_0 = 0.6$ and the error variance to $\sigma_\epsilon = 0.01$ to ensure trajectories remain in the positive quadrant

mechanism in this case). As a result, the appropriate lag order of the VAR(p) has to be chosen based on diagnostic checks or information criteria as in Stockhammer *et al.* (2019). We use an AR(1) process of the form $u_t = \rho u_{t-1} + \epsilon_t$, where $\epsilon_t \sim N(0, 1)$ and $\rho = 0.8$. To determine the lag order, we start with a minimum lag length of 2. If there is serial correlation in the residuals, the lag length is successively increased. A maximum number of lags is specified and estimates are excluded from further consideration if their residuals still exhibit serial correlation. We check the robustness of this approach below by using instead the Akaike information criterion (AIC) for lag selection.

Note that in the case of the VAR(p), the eigenvalues can no longer be directly mapped to the Jacobian matrix of the deterministic component of the system. To determine a unique cycle length from the VAR(p) models with up to p complex eigenvalues, we average over those estimated cycle lengths that lie within conventional business and financial cycle frequencies of 3–20 periods (see Beneš and Vávra, 2005 for a similar approach).

The Monte Carlo design is given by the following algorithm:

1. Draw samples of size T from the DGP.
2. Estimate a linear VAR(p) on the sample. In the scenario with serially uncorrelated error terms, set $p = 1$ and proceed with step 6. In the scenario with AR(1) shocks, set $p = 2$ and proceed with step 3.
3. Run a Lagrange Multiplier test for serial correlation in the errors.
4. If the null hypothesis of no serial correlation is rejected at the 10% level, re-estimate the VAR with a further lag added.
5. Repeat steps 3 and 4 until either there is no more serial correlation or a maximum of 4 lags is reached.
6. Save the estimated coefficients and relevant information criteria.
7. Repeat the above steps $N = 1,000$ times.

We emulate a realistic small sample scenario for applied time series analyses with annual macroeconomic data (baseline: $T = 50$). A maximum of four lags is imposed as most researchers working with annual data will be reluctant to consider more than four lags. A VAR(4) that still exhibits serial correlation may raise concerns of model misspecification. We therefore exclude estimated VAR(4) models for which serial correlation does not vanish.

As noted in section II, we are interested in the interaction mechanism summarized by (2), the eigenvalues of the estimated model, and the corresponding dominant cycle length. For those n estimated models that pass the serial correlation test, we store the estimated coefficients on the off-diagonal of coefficient matrix on the first lag – that is, a_2 and b_1 in (1). Testing condition (2), that is, $a_2 b_1 < 0$, requires a one-sided statistical test on a nonlinear combination of estimators. We calculate the standard error with the delta method (see Oehlert, 1992) and run a one-sided t -test on condition (2) imposing a significance level of 10%. Estimated test statistics that pass this test are deemed statistically significant, which we denote by $(a_2 b_1 < 0)^*$. For the estimated eigenvalues, we plot sampling distributions on Argand diagrams and compare these to the true eigenvalues of the underlying limit

TABLE 3

Results for SM39 and BGP20, VAR(1) with AR(0) shocks and VAR(p) with AR(1) shocks, averages

	SM39		BGP20	
	VAR(1)	VAR(p)	VAR(1)	VAR(p)
$a_2\hat{b}_1$	-0.48	-0.49	-0.34	-0.30
bias ($a_2\hat{b}_1$)	-0.20%	-1.77%	28.9%	37.4%
NRMSE ($a_2\hat{b}_1$)	1.00	1.00	2.66	1.39
($a_2\hat{b}_1 < 0$)	100%	99.9%	100%	96.1%
($a_2\hat{b}_1 < 0$)*	100%	96.2%	99.9%	73.5%
$ \hat{\lambda} _{dom}$	0.89	0.90	0.87	0.86
bias ($ \hat{\lambda} _{dom}$)	-1.04%	0.21%	-16.5%	-16.8%
$\hat{r}e$	0.59		0.65	
bias ($\hat{r}e$)	-1.78%		-19.4%	
$\hat{i}m$	0.66		0.58	
bias ($\hat{i}m$)	-0.66%		-12.8%	
\hat{L}	7.50	7.72*	8.66	9.13*
bias (\hat{L})	-0.33%	2.63%	-4.62%	0.58%
Nonnorm (% of n)	7.9%	6.93%	5.2%	6.91%
Lags	1	2.27	1	2.26
Serial corr(% of N)		10.5%		8.8%
n	1,000	895	1,000	912

Notes: Total number of runs $N = 1,000$; n : number of VARs that did not suffer from serial correlation; sample size $T = 50$; $a_2\hat{b}_1$: estimated cycle condition; bias: % deviation from absolute true value; NRMSE: root mean squared error normalized by standard deviation; ($a_2\hat{b}_1 < 0$): relative frequency (in n) with which the cycle condition was satisfied; ($a_2\hat{b}_1 < 0$)*: relative frequency (in n) with which the condition was statistically significant at the 10% level or lower; $|\hat{\lambda}|_{dom}$: dominant modulus; re : real part of eigenvalue; im : imaginary part of eigenvalue, L : cycle length. Nonnorm: relative frequency (in n) with which estimated VARs exhibited non-normal residuals. Serial corr: relative frequency (in N) with which estimated VARs exhibited serially correlated residuals after the inclusion of 4 lags. *Average over $\hat{L}_j \in (3, 20)$, where $j = 1, \dots, p$.

cycle process to assess how frequently estimated VARs suggest local instability and what they can tell us about the underlying limit cycle.

To assess the relative importance of omitted nonlinearity bias compared with bias stemming from lagged dependent variables and serial correlation, we explore the asymptotic properties of the estimators of interest for different cases that allow us to disentangle these sources of bias. Finally, we also conduct Jarque–Bera normality tests on the residuals of the VAR to assess whether non-normality can be taken as a sign of underlying limit cycle processes.⁴

V. Simulation results

Main results

Table 3 summarizes the results from the VAR(1) with AR(0) shock processes and the VAR(p) with AR(1) shocks for the baseline models SM39 and BGP20 (sample size $T = 50$). For the linear model SM39, all estimated coefficients and eigenvalues exhibit less than 2% bias. The normalized root mean squared error (NRMSE) of the cycle mechanism $a_2\hat{b}_1$ is equal to unity. In each draw, the critical condition (2) for a cycle mechanism is satisfied and statistically significant in 96% of draws in the VAR(p). Indeed,

⁴We test the null hypothesis of normality using a 10% significance level.

the linear first-order model SM39 with AR(1) error terms is formally equivalent to a VAR(2) with white noise errors. We thus expect the estimated VAR(p) to asymptotically identify the off-diagonal coefficients of SM39 correctly.

To shed light on the different sources of bias in the linear case, Figure 4 plots the bias for different sample sizes. The left panel presents results from a VAR(1) with serially uncorrelated errors, whereas the right panel depicts the case of a VAR(1) with AR(1) errors (i.e. without lag adjustment to remove serial correlation). With serially uncorrelated errors, the autoregressive coefficient a_1 exhibits the well-known Hurwicz bias that is rapidly decreasing in the sample size. Notably, the estimated cycle mechanism a_2b_1 appears to be rather unaffected by this bias. Similarly, the real and imaginary parts of the complex eigenvalue as well as the modulus are strongly affected by the Hurwicz bias for small sample sizes, but the implied cycle length is much less so. With serially correlated errors, the estimators become inconsistent with asymptotic biases in the estimators for the cycle condition and cycle length of around 45% (but a much stronger bias in the autoregressive parameter). Online Appendix S1 (section D) reports results from a VAR(p) confirming that lag adjustment is effective in rendering the estimator for the cycle condition consistent in the presence of AR(1) shocks. For sample sizes of $T > 250$, the bias becomes practically irrelevant (less than 1%).

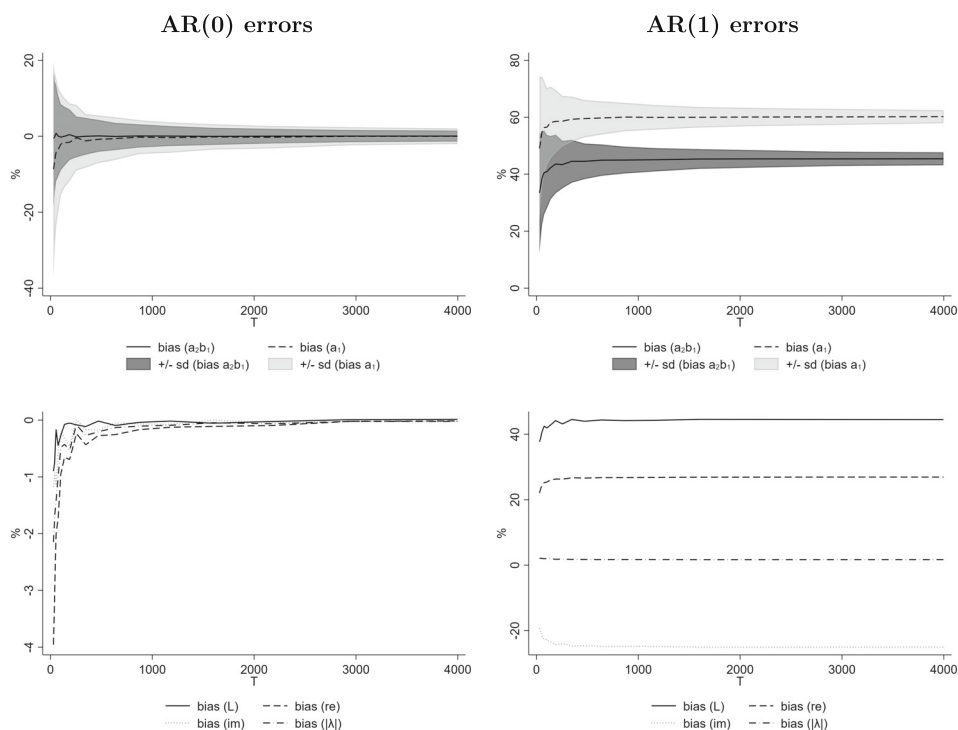


Figure 4. Asymptotic bias in VAR(1) for SM39, AR(0) and AR(1) errors

Notes: Lines represent the bias (in %) in estimated statistics from Monte Carlo simulations with $N = 1,000$ runs for increasing sample sizes $T = 30, \dots, 4,000$. Shaded areas denote one standard deviation around the lines

We can now contrast these results with the nonlinear model BGP20, where the estimated coefficients are additionally biased as a result of the unmodelled nonlinearity. We first note that in Table 3 the estimated real part of the eigenvalue and the modulus are underestimated in size. The VAR therefore fails to identify the presence of local instability in BGP20 (the estimated modulus is below unity) – in line with the findings for a univariate process in Beaudry *et al.* (2017). The main interest of this paper, however, lies in the presence of complex eigenvalues that stem from interaction mechanisms inherent to endogenous cycle models. We note that the strength of the interaction mechanism is underestimated, with $a_2\hat{b}_1 = -0.34$ being around 29% smaller in absolute terms than the true value $a_2b_1 = -0.48$. This bias becomes larger in the VAR(p).

Despite this underestimation of the strength of the interaction mechanism, the critical condition for a cycle mechanism is qualitatively satisfied and statistically significant in most cases, with cycle detection rates of 99.9% (VAR(1)) and 73.5% (VAR(p)). We further note that the real and imaginary part of the complex eigenvalue are both underestimated, but the downward bias in the imaginary part is smaller (–12.8%) compared with the downward bias in the real part (–19.4%). Interestingly, the implied cycle length exhibits a much smaller bias than the individual parts of the complex conjugate (–4.6%). Figure 5 plots the sampling distribution of the estimated complex eigenvalues for BGP20 (VAR(1)).⁵ It can be seen that while both the real and imaginary parts are underestimated, the angle $\omega = \arctan(\text{im}(\lambda)/\text{re}(\lambda))$ of the estimated eigenvalue in polar coordinates is relatively close to its true value. While a higher real part raises the cycle length, a higher imaginary part decreases it. The downward biases in both parts thus partly offset each other, which explains the comparatively small bias in the estimated cycle length. Finally, it is worth noting that the VARs do not appear to suffer from a significantly higher occurrence of serial correlation or non-normal errors compared with the linear benchmark.

To assess the relative importance of omitted nonlinearity bias, Figure 6 explores the asymptotic properties of the VAR(1) for BGP20. In the case without serial correlation (left panel), the Hurwicz bias decreases rapidly, and for sample sizes of $T > 100$, the remaining bias can effectively be exclusively attributed to the omitted nonlinearity. The biases in the cycle mechanism and cycle length approach constant values of around 30% and –4%, respectively, that are driven by the nonlinearity. In the case with serially correlated errors (right panel), the asymptotic biases in the cycle condition and length are substantially larger (by around 37%-pts and 65%-pts, respectively), as they are driven by both the omitted nonlinearity and unmodelled serial correlation.

This exercise suggests that the bias due to unmodelled nonlinearities dominates the Hurwicz bias from the inclusion of lagged dependent variables, which becomes practically irrelevant for medium sample sizes. However, bias due to unmodelled serial correlation is substantial and quantitatively similar to the bias from unmodelled nonlinearities. The use of higher-order lags substantially reduces serial correlation bias, as expected.⁶

⁵Online Appendix S1 (section E) provides further details on the sampling distribution of estimators for BGP20.

⁶Online Appendix S1 (section F) further explores the robustness to different parameterizations of BGP20. The results are very similar.




Figure 5. Argand diagram, BGP20, VAR(1) with AR(0) shocks

Notes: The grey shaded areas summarize the sampling distribution of the eigenvalues (limited to the positive quadrant of the Argand diagram) of a VAR(1) estimated on 10,000 runs of BGP20. The black point • marks the mode of the sampling distribution. The true eigenvalue of the linear part of BGP20 is $\lambda_+ = 0.8 + 0.66i$ (as in Table 1) and the dashed line connects that point with the origin

Overall, these results suggest that the linear VAR performs very well in identifying the critical condition for a cycle mechanism in the nonlinear baseline models. While it tends to underestimate the true magnitude of the interaction mechanism, it allows for an accurate qualitative identification of the presence of a cycle mechanism and a reasonable gauge of the cycle length.

Robustness tests

This section presents a number of robustness tests and extensions to our main results.

Other DGPs

First, to assess robustness to different forms of nonlinearity, Table 4 displays results from four further nonlinear limit cycle models summarized in Tables 1 and 2. We observe that for all nonlinear DGPs the VAR underestimates the true value of the cycle mechanism (a_2b_1) in absolute terms. However, the cycle condition is still correctly identified in the majority of samples. The VARs estimated on NK92 and DW12 exhibit very high cycle detection rates of almost 100% and 99%, respectively. For WWZ09 and BDRS01, the cycle mechanism is correctly identified in 80% and 55% of cases, respectively.

While NK92 exhibits the highest cycle detection rate among the nonlinear models, the VAR underestimates its cycle mechanism in absolute terms by about 44%. However,




Figure 6. Asymptotic bias in VAR(1) for BGP20, AR(0) and AR(1) errors

Notes: Lines represent the bias (in %) in estimated statistics from Monte Carlo simulations with $N = 1,000$ runs for increasing sample sizes $T = 30, \dots, 4,000$. Shaded areas denote one standard deviation around the lines

as the interaction mechanism in this model is comparatively strong, the cycle condition is still qualitatively satisfied in each sample draw. For DW12, which also exhibits a high detection rate, the VAR estimate of the cycle condition is much less biased compared with NK92 (less than 3%). The better performance compared with the other nonlinear models is likely due to the fact that for DW12 the off-diagonal elements of the Jacobian are linear. In comparison, BGP20, WWZ09, and BDRS01 exhibit nonlinear off-diagonal elements and a moderate strength of the interaction mechanism at the steady state. Here the VAR still performs well overall, but worse compared with the models with stronger interaction mechanisms (NK92) or linear off-diagonal elements (DW12).

Figure 7 examines the effect of variations in the sample size ($T \in \{50, 100, 150, 200\}$). For all nonlinear models except DW12, the estimator converges to a fixed value that is different from the true value. For DW12, the estimator appears to converge to the true value, presumably because it is the only model that is linear in the off-diagonal coefficients. Overall, the linear VAR performs reasonably well in identifying the cycle condition (2), but with differences depending on the underlying DGP.

The estimated real parts of the eigenvalue and moduli are again underestimated in size across all nonlinear DGPs. Our study thus confirms and extends to the multivariate case

TABLE 4
Results of other DGPs, VAR(1) with AR(0) shocks and VAR(p) with AR(1) shocks, averages

	BDRS01		NK92		WWZ09		DW12	
	VAR(1)	VAR(p)	VAR(1)	VAR(p)	VAR(1)	VAR(p)	VAR(1)	VAR(p)
$a_2\hat{b}_1$	-0.30	-0.13	-1.44	-0.91	-0.20	-0.18	-0.50	-0.48
bias ($a_2\hat{b}_1$)	58.3%	82.5%	11.6%	44.2%	17.2%	27.2%	-2.37%	2.95%
NRMSE ($a_2\hat{b}_1$)	5.66	5.57	7.00	3.33	1.37	1.17	1.00	1.00
($a_2\hat{b}_1 < 0$)	100%	89.7%	100%	100%	100%	97.5%	100%	100%
($a_2\hat{b}_1 < 0$)*	100%	55.3%	100%	100%	99.7%	79.7%	100%	98.9%
$ \hat{\lambda} _{dom}$	0.71	0.83	0.96	0.99	0.87	0.90	0.46	0.82
bias ($ \hat{\lambda} _{dom}$)	-31.4%	-19.3%	-10%	-7.68%	-16.4%	-13.5%	-65.2%	-37.4%
$\hat{r}e$	0.60		0.23		0.74		0.35	
bias ($\hat{r}e$)	-31.4%		-12.4%		-18.5%		-68.7%	
\hat{im}	0.36		0.94		0.44		0.23	
bias (\hat{im})	-4.75%		-9.86%		-10.1%		-66.7%	
\hat{L}	12.4	11.3*	4.72	5.62*	11.9	11.9*	9.76	10.3*
bias (\hat{L})	-26%	-32.8%	-0.49%	18.4%	-6.53%	-6.19%	-14.3%	-9.32%
Nonnorm (% of n)	5.9%	6.46%	26.4%	15.5%	6.1%	7.33%	5.4%	6.14%
Lags	1	2.29	1	3.27	1	2.28	1	2.32
Serial corr(% of N)		8.7%		63.3%		11.3%		10.4%
n	1,000	913	1,000	367	1,000	887	1,000	896

Notes: Total number of runs $N = 1,000$; n : number of VARs that did not suffer from serial correlation; sample size $T = 50$; $a_2\hat{b}_1$: estimated cycle condition; bias: % deviation from absolute true value; NRMSE: Root mean squared error normalized by standard deviation; ($a_2\hat{b}_1 < 0$): relative frequency (in n) with which the cycle condition was satisfied; ($a_2\hat{b}_1 < 0$)*: relative frequency (in n) with which the condition was statistically significant at the 10% level or lower; $|\hat{\lambda}|_{dom}$: dominant modulus; re : real part of eigenvalue; im : imaginary part of eigenvalue, L : cycle length. Nonnorm: relative frequency (in n) with which estimated VARs exhibited non-normal residuals. Serial corr: relative frequency (in N) with which estimated VARs exhibited serially correlated residuals after the inclusion of 4 lags. *Average over $\hat{L}_j \in (3, 20)$, where $j = 1, \dots, p$.

the finding from Beaudry *et al.* (2017) for a broader class of nonlinearities. Going beyond Beaudry *et al.* (2017), our robustness tests show that the bias in the estimated cycle length is generally lower compared with the bias in the estimated modulus in the VAR(1). In the VAR(p), the bias in the estimated cycle length is less than 10% for WWZ09 and DW12, but between 18% and 33% (in absolute terms) for NK92 and BDRS01. We explore below to what extent these results are related to our cycle length selection rule.

Two further observations can be made. First, NK92 is the only model in which serial correlation occurs significantly more often (63%) compared with the linear baseline. We conjecture that this is due to the high cycle frequency of the deterministic component of the DGP ($L = 4.8$). The VARs estimated on the other nonlinear models do not suffer from an elevated occurrence of serial correlation, suggesting that serial correlation is not a useful sign for underlying limit cycle processes. Second, only NK92 exhibits a somewhat increased frequency of non-normality. Thus, non-normality does not seem to be a reliable indicator of limit cycles either.

Overall, this confirms the main result that the linear VAR underestimates the magnitude of cycle mechanisms but successfully uncovers their existence. It also estimates the cycle length quite accurately, especially in the VAR(1). How well the VAR performs appears to depend on two factors. First, if the nonlinearity only affects the main diagonal but not the off-diagonal elements, the likelihood that the cycle mechanism is identified is higher.

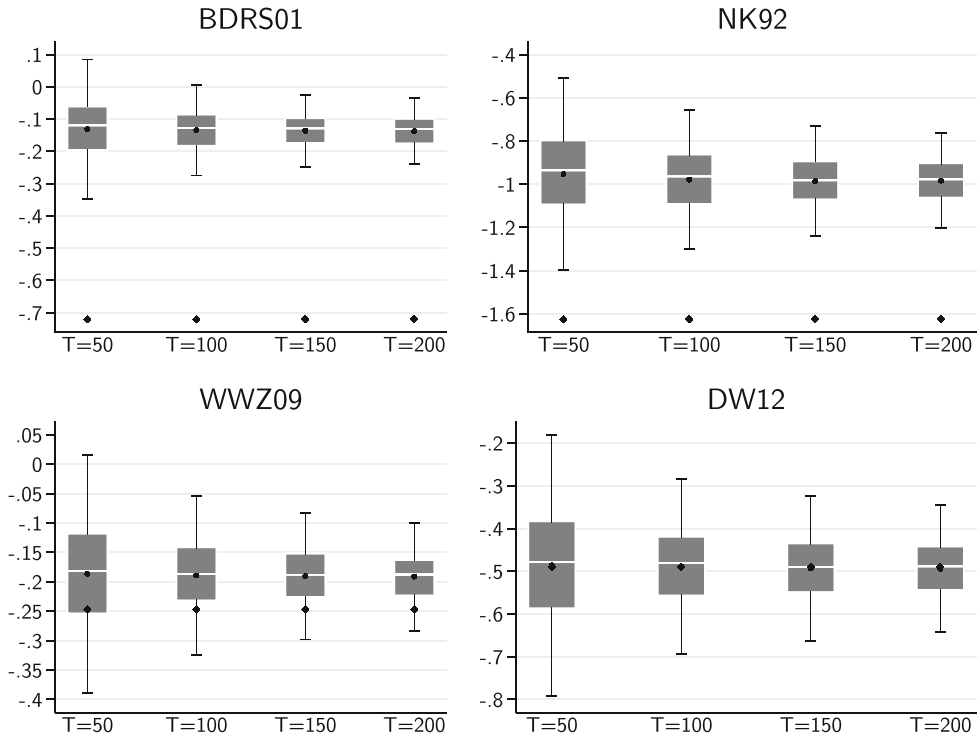


Figure 7. Box plots for $a_2\hat{b}_1$ for different sample sizes, other DGPs, AR(1) shocks

Notes: Boxes represent the first and third quartile, circles mark the mean, and whiskers represent 2 standard deviations around the mean. Diamonds mark the true values

Second, if the cyclical interaction mechanism of a nonlinear model is quantitatively strong, it will have a higher chance of being detected by a linear VAR.

Sensitivity to false positives

Next, we investigate the sensitivity of the linear VAR to false positives in the detection of cycle mechanisms. To this end, we consider the nonlinear model proposed in Brock and Hommes (1997), henceforth BH97, that generates fluctuations but does not contain a cycle mechanism of the kind discussed in section II. Instead, the model gives rise to aperiodic chaotic fluctuations resulting from saddle instability. A formal discussion of the model can be found in Online Appendix S1 (section A) and simulation results are reported in Online Appendix S1 (section G).

The estimated cycle condition is close to its true value of zero. It is statistically significant in less than 1% of cases. Despite the omitted nonlinearity, the estimator appears to converge to the true value for increasing sample sizes, which is likely due to the off-diagonal elements being equal to zero at the steady state. We also note a substantially elevated number of models that exhibit non-normal errors (72% in the case of the VAR(1)) as well as serial correlation (16%). This suggests that serial correlation and non-normality are more reliable indicators of nonlinearity when the DGPs produce chaotic fluctuations as opposed to periodic limit cycles. Overall, this exercise suggests that the linear VAR is quite robust to false positives.

Alternative approaches to obtain the cycle length

We further compare our baseline results to an alternative approach to obtain a unique cycle length when there is more than one complex eigenvalue. In the baseline approach, an average cycle length is calculated over a range of conventional cycle frequencies (3–20 periods). An alternative approach is to use only the cycle length implied by the dominant eigenvalue. Online Appendix S1 (section D) documents the asymptotic properties of the two approaches for the linear model SM39 and the nonlinear model BGP20.

The results suggest that there are trade-offs. The baseline approach exhibits a comparatively small bias for small sample sizes (less than 5% for $T < 200$). However, averaging over a conventional range appears to render the estimator asymptotically biased. By contrast, the alternative approach based on the dominant eigenvalue performs better for larger sample sizes. In the linear case, it asymptotically converges to the true value. In the nonlinear case, a negative asymptotic bias remains, but this is smaller compared with the baseline approach (–4.7% compared with –8.4%). However, the alternative approach performs poorly for small sample sizes, where the bias is large. In sum, there is no clearly preferable approach for selecting a cycle length when there are multiple complex eigenvalues. In practice, researchers might want to draw on additional knowledge to select the cycle length they suspect to be driven by the deterministic component of the DGP.

Alternative lag selection rules

Finally, we check the robustness of our results to the lag selection algorithm (see Online Appendix S1 (section H)). The results remain very similar when using the AIC instead of serial correlation tests to determine the lag length. Notably, the results are highly robust despite the fact that 19%–68% of estimated VARs exhibit serial correlation when using this lag selection method.

VI. Conclusion

Recent macroeconomic research has revived the notion of endogenous business and financial cycles (Borio, 2014; Stockhammer *et al.*, 2019; Beaudry *et al.*, 2020). This paper investigated what commonly used linear vector autoregressions would tell us about the interaction mechanisms that drive nonlinear limit cycles. We conducted Monte Carlo simulations in which linear VARs were estimated on data from five nonlinear business cycle models. Our findings suggest that despite a tendency to underestimate the moduli of complex eigenvalues and the magnitude of cycle mechanisms, linear VARs identify the presence of cycle mechanisms quite well, with cycle detection rates between 55% and almost 100%. The detection rate tends to be higher (i) when the nonlinearity only affects the autoregressive component of the model but not the interaction mechanism itself and (ii) the larger is the quantitative magnitude of the interaction mechanism. Linear VARs also appear to have low false positive error rates of less than 1% in the numerical experiment conducted in this paper. Lastly, there is only a comparatively small downward bias in estimated cycle lengths, suggesting that linear VARs successfully pick up the frequencies generated by nonlinear cycle processes.

These findings have important practical implications for research on business cycles and financial cycles with linear VARs, which is increasingly interested in cycle frequencies and interactions between real and financial variables (Stockhammer *et al.*, 2019; Strohsal *et al.*, 2019; Juselius and Drehmann, 2020). In the absence of clear-cut theoretical assumptions about the specific functional form of potentially nonlinear data-generating processes, linear VARs continue to be the default choice in this line of research. Our results lend support to this research in showing that linear VARs are capable of identifying the existence of cycle interaction mechanisms and can provide a reasonable approximation of cycle frequencies, even if these may stem from nonlinear processes. More generally, our results suggest that if key macroeconomic time series are indeed driven by nonlinear limit cycle processes as suggested in (Beaudry *et al.*, 2017; Beaudry *et al.*, 2020), these oscillations are likely to show up as damped oscillations in linear VARs with conventional macroeconomic indicators.

Our findings also point to a number of limitations of the linear approach. First, while qualitatively strong at identifying cycle mechanisms, the linear VAR may not detect them when the interaction mechanism is quantitatively small. Second, our findings demonstrate that the linear VAR fails to identify local instability inherent in limit cycle models of endogenous business cycles. This confirms previous studies with univariate processes (Beaudry *et al.*, 2017) and extends this result to the multivariate case, as well as a broader class of functional forms. For the applied researcher, this means that the estimated moduli from linear VARs do not contain useful information about possible instabilities of the underlying process. Provided there are theoretical priors about the nature of the underlying nonlinearity, nonlinear time series models may be a useful option (see Kilian and Lütkepohl, 2017; Teräsvirta *et al.*, 2010; Teräsvirta, 2018). For example, smooth transition models might adequately approximate the sigmoid nonlinearities found in most macroeconomic limit cycle models. A multivariate time series model with an adequately flexible nonlinear part might be able to identify simultaneously interaction mechanisms and instabilities: we leave this natural extension to future work.

Overall, linear VARs are a parsimonious method for the qualitative identification of endogenous cycle mechanisms and may stimulate further investigation into potential instabilities and nonlinearities in the underlying process. For researchers that are primarily interested in testing for the critical condition for a cycle mechanism or estimating cycle lengths, the linear VAR is therefore a convenient tool, despite the limitations already noted. For researchers who are primarily interested in identifying and estimating impulse response functions to structural shocks, eigenvalues of the reduced-form VAR are often plotted in Argand diagrams as a simple ‘eyeball’ test for dynamic stability, and then ignored. Our results suggest that the eigenvalues of the reduced-form VAR contain useful information about cycle frequencies and cyclical mechanisms that could complement impulse response functions when informing model choice and selection. A careful consideration of the cyclical properties of estimated linear VAR models can therefore be a useful step in applied macroeconomics research.

Final Manuscript Received: April 2021

References

- Abadir, K. M., Hadri, K. and Tzavalis, E. (1999). 'The influence of VAR dimensions on estimator biases', *Econometrica*, Vol. 67, pp. 163–181.
- Aikman, D., Haldane, A. G. and Nelson, B. D. (2015). 'Curbing the credit cycle', *Economic Journal*, Vol. 125, pp. 1072–1109.
- Asada, T. (2001). 'Nonlinear dynamics of debt and capital: a post-Keynesian analysis', in Aruka Y. (ed), *Evolutionary Controversies in Economics. A New Transdisciplinary Approach*, Springer, Japan, Tokyo, pp. 73–88.
- Azariadis, C. (2018). 'Riddles and models: a review essay on Michel De Vroey's *A History of Macroeconomics from Keynes to Lucas and Beyond*', *Journal of Economic Literature*, Vol. 56, pp. 1538–1576.
- Beaudry, P., Galizia, D. and Portier, F. (2016). *Putting the Cycle Back into Business Cycle Analysis*, NBER Working Paper No. 22825, National Bureau of Economic Research, Cambridge, MA.
- Beaudry, P., Galizia, D. and Portier, F. (2017). 'Is the macroeconomy locally unstable and why should we care?', *NBER Macroeconomics Annual*, Vol. 31, pp. 479–530.
- Beaudry, P., Galizia, D. and Portier, F. (2020). 'Putting the cycle back into business cycle analysis', *American Economic Review*, Vol. 110, pp. 1–47.
- Beneš, J. and Vávra, D. (2005). *Eigenvalue Filtering in VAR Models with Application to the Czech Business Cycle*, ECB Working Paper Series No. 549.
- Bischi, G. I., Dieci, R., Rodano, G. and Saltari, E. (2001). 'Multiple attractors and global bifurcations in a Kaldor-type business cycle model', *Journal of Evolutionary Economics*, Vol. 11, pp. 527–554.
- Borio, C. (2014). 'The financial cycle and macroeconomics: what have we learnt?', *Journal of Banking & Finance*, Vol. 45, pp. 182–198.
- Branch, W. A. and McGough, B. (2010). 'Dynamic predictor selection in a New Keynesian model with heterogeneous expectations', *Journal of Economic Dynamics and Control*, Vol. 34, pp. 1492–1508.
- Brock, W. A. and Hommes, C. H. (1997). 'A rational route to randomness', *Econometrica*, Vol. 65, pp. 1059–1095.
- Brockwell, P. and Davis, R. (2006). *Time Series: Theory and Methods*, Springer, New York, NY.
- Calvert Jump, R. and Levine, P. (2019). 'Behavioural new Keynesian models', *Journal of Macroeconomics*, Vol. 59, pp. 59–77.
- Chiarella, C. (1992). 'The dynamics of speculative behaviour', *Annals of Operations Research*, Vol. 37, pp. 101–123.
- Dieci, R. and He, X.-Z. (2018). 'Heterogeneous agent models in finance', in Hommes C. and LeBaron B. (eds), *Handbook of Computational Economics* Amsterdam, North-Holland and Oxford: Elsevier, Vol. 4, pp. 257–328.
- Dieci, R. and Westerhoff, F. (2012). 'A simple model of a speculative housing market', *Journal of Evolutionary Economics*, Vol. 22, pp. 303–329.
- Dieci, R. and Westerhoff, F. (2016). 'Heterogeneous expectations, boom-bust housing cycles, and supply conditions: a nonlinear economic dynamics approach', *Journal of Economic Dynamics and Control*, Vol. 71, pp. 21–44.
- Doornik, J. A., Nielsen, B. and Rothenberg, T. J. (2003). 'The influence of VAR dimensions on estimator biases: comment', *Econometrica*, Vol. 71, pp. 377–383.
- Fernández-Villaverde, J. and Rubio-Ramírez, J. F. (2005). 'Estimating dynamic equilibrium economies: linear versus nonlinear likelihood', *Journal of Applied Econometrics*, Vol. 20, pp. 891–910.
- Francq, C. and Zakoian, J.-M. (1998). 'Estimating linear representations of nonlinear processes', *Journal of Statistical Planning and Inference*, Vol. 68, pp. 145–165.
- Galí, J. (2018). 'The state of new Keynesian economics: a partial assessment', *Journal of Economic Perspectives*, Vol. 32, pp. 87–112.
- Hamilton, J. D. (1994). *Time Series Analysis*, Princeton University Press, Princeton, NJ.
- Juselius, M. and Drehmann, M. (2020). 'Leverage dynamics and the burden of debt', *Oxford Bulletin of Economics and Statistics*, Vol. 82, pp. 347–364.
- Kilian, L. and Lütkepohl, H. (2017). *Structural Vector Autoregressive Analysis*, Cambridge University Press, Cambridge, UK.

- Kiyotaki, N. and Moore, J. (1997). 'Credit cycles', *Journal of Political Economy*, Vol. 105, pp. 211–248.
- Neubert, M. G. and Kot, M. (1992). 'The subcritical collapse of predator populations in discrete-time predator-prey models', *Mathematical Biosciences*, Vol. 110, pp. 45–66.
- Oehlert, G. W. (1992). 'A note on the delta method', *The American Statistician*, Vol. 46, pp. 27–29.
- Prescott, E. C. (1986). 'Theory ahead of business cycle measurement', *Federal Reserve Bank of Minneapolis Quarterly Review*, Vol. 10, pp. 9–22.
- Rünstler, G. and Vlekke, M. (2017). 'Business, housing, and credit cycles', *Journal of Applied Econometrics*, Vol. 33, pp. 212–226.
- Samuelson, P. A. (1939). 'Interactions between the multiplier analysis and the principle of acceleration', *The Review of Economics and Statistics*, Vol. 21, pp. 75–78.
- Smets, F. and Wouters, R. (2003). 'An estimated dynamic stochastic general equilibrium model of the euro area', *Journal of the European Economic Association*, Vol. 1, pp. 1123–1175.
- Stockhammer, E., Calvert Jump, R., Kohler, K. and Cavallero, J. (2019). 'Short and medium term financial-real cycles: an empirical assessment', *Journal of International Money and Finance*, Vol. 94, pp. 81–96.
- Strohsal, T., Proaño, C. R. and Wolters, J. (2019). 'Characterizing the financial cycle: evidence from a frequency domain analysis', *Journal of Banking & Finance*, Vol. 106, pp. 568–591.
- Teräsvirta, T. (2018). *Nonlinear Models in Macroeconometrics* Oxford Research Encyclopedia of Economics and Finance, Oxford University Press, Oxford, UK.
- Teräsvirta, T., Tjøstheim, D. and Granger, C. W. J. (2010). *Modelling Nonlinear Economic Time Series*, Oxford University Press, Oxford, UK.
- Wegener, M., Westerhoff, F. and Zaklan, G. (2009). 'A Metzlerian business cycle model with nonlinear heterogeneous expectations', *Economic Modelling*, Vol. 26, pp. 715–720.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Appendix S1. Supporting Information