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

Karsten Kohler† Robert Calvert Jump ‡

Forthcoming in *Oxford Bulletin of Economics and Statistics*

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.

JEL Codes: C15, C32, E32

Keywords: vector autoregression, limit cycles, endogenous cycles, business and financial cycles, cycle frequency

*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'.

†Economics Department, Leeds University Business School, Maurice Keyworth Building, Woodhouse, Leeds LS2 9JT. 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. E-Mail: r.g.calvertjump@greenwich.ac.uk.

1 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 linearised 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 nonoscillatory 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 in output. An important stimulus has been empirical research on financial cycles that finds dominant frequencies of 13 to 16 years in aggregate financial variables such as private credit and house prices (Aikman et al. 2015, Borio 2014, Rünstler and Vlekke 2017, Stockhammer et al. 2019, Strohsal et al. 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 et al. (2020). They show that various U.S. 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 Zakoïan 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 et al. (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 linearisation, 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 – i.e., a VAR model – can tell us about the linear part – i.e., 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, 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 summarise the

[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, Juselius and Drehmann 2020, Rünstler and Vlekke 2017, Stockhammer et al. 2019, Strohsal et al. 2019) 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.

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-Ramírez (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 (Beaudry et al. 2020, Bischi et al. 2001, Branch and McGough 2010, Brock and Hommes 1997, Dieci and Westerhoff 2012, Dieci and He 2018, Calvert Jump and Levine 2019, Wegener et al. 2009), 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 short-term real 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 (Kilian and Lütkepohl 2017, Teräsvirta

et al. 2010, 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 2 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 3 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 4 describes the Monte Carlo design. Section 5 presents the main results and various robustness tests. Section 6 concludes.

2 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,

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

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, 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,

$$a_2 b_1 < 0, \tag{2}$$

i.e., 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 $Det = a_1 b_2 - a_2 b_1$ 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 $Tr^2 - 4Det = (a_1 + b_2)^2 - 4(a_1 b_2 - a_2 b_1)$ to $(a_1 - b_2)^2 + 4(a_2 b_1)$ demonstrates that $Tr^2 - 4Det$ can only be negative if $a_2 b_1 < 0$.

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 summarise 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,

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

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 serially correlated, the estimator becomes inconsistent, i.e., the bias does not appear asymptotically.

[2]This paper focuses on stationary processes. See Abadir et al. (1999), Doornik et al. (2003), and Hamilton (1994, chap.18) for sources of bias in non-stationary VARs.

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

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

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 $\mathsf{E}[x_t|x_{t-1}] = g(x_{t-1})$. Suppose for example that,

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

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,

$$x_t = a + b x_{t-1} + \eta_t, \tag{6}$$

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

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

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.

3 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.

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 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 improves 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 optimising 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 et al. (2016, 2020), henceforth BGP20, as the baseline nonlinear business cycle model for our Monte Carlo experiments. The model is summarised by,

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

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

where X_t denotes a capital or durable consumption good and I_t denotes investment in that

Figure 1: Simulation of SM39

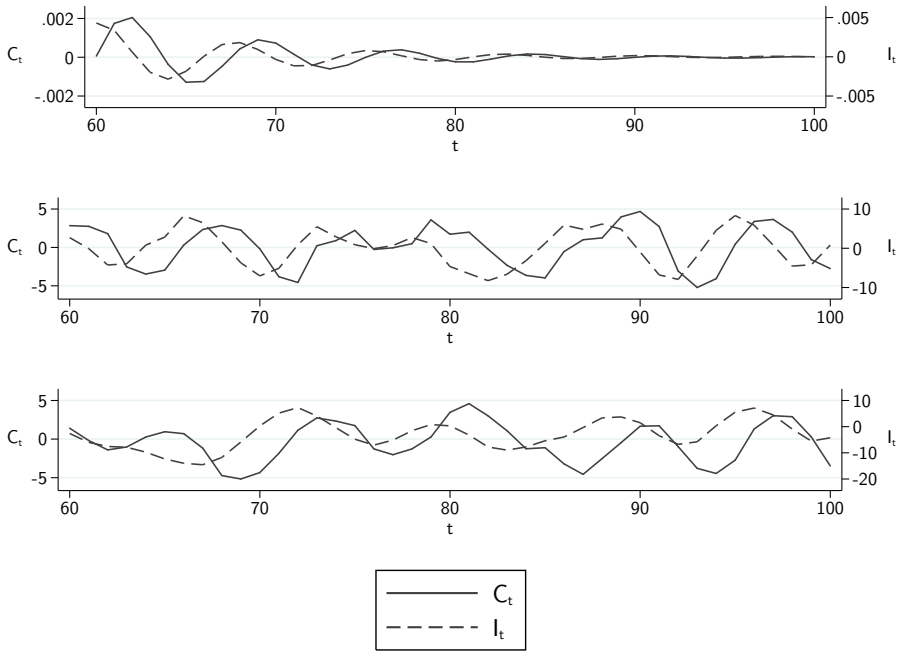

Notes: Simulation of the system $C_t = 0.4C_{t-1} + 0.4I_{t-1}$, $I_t = -1.2C_{t-1} + 0.8I_{t-1}$. Upper panel: deterministic version. Middle panel: stochastic version with white noise errors $\epsilon_t \sim N(0,1)$. Lower panel: stochastic version with serially correlated error terms $u_t = 0.8u_{t-1}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

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,

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

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

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

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

Figure 2: Simulation of BGP20

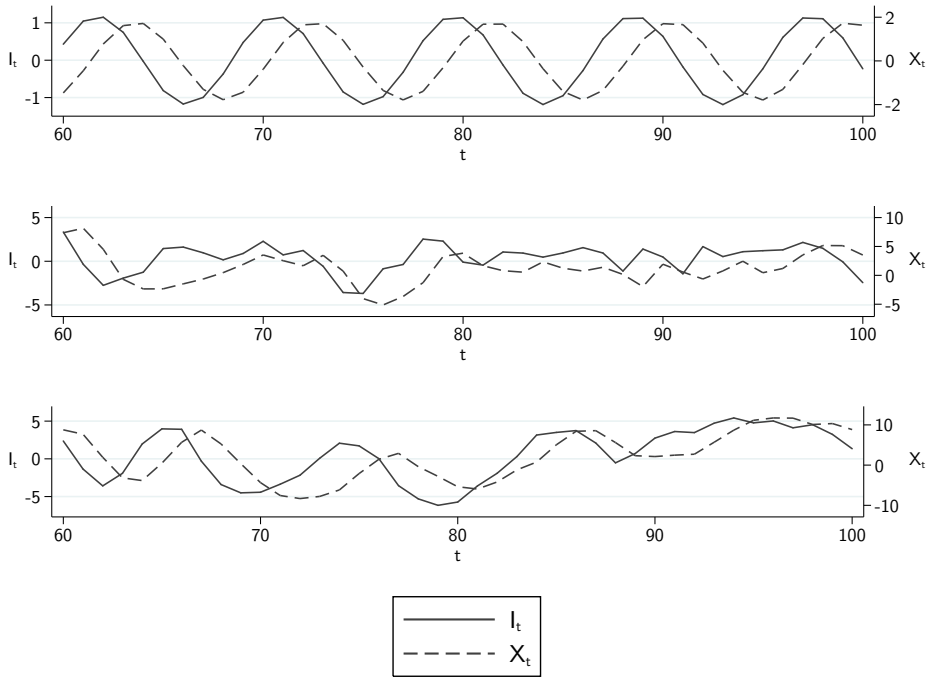

Notes: Simulation of the system (10)-(11) with $\alpha_0 = -1$, $\alpha_1 = 0.4$, $\alpha_2 = 0.9$, $\alpha_3 = 2$, $\delta = 0.4$. Upper panel: deterministic version. Middle panel: stochastic version with white noise errors $\epsilon_t \sim N(0, 1)$. Lower panel: stochastic version with serially correlated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

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 forms – see Online Appendix A for a detailed discussion of these models. Tables 1-2 provide an overview of the DGPs, their main properties and the parameterisations 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,

[3]With the exception of BGP20, all parameterisations are taken from the respective studies and meet the conditions for endogenous cycle dynamics. In BGP20 no parameterisation is provided. We chose a parameterisation 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 parameterisation.

Table 1 Overview of simulated DGPs, part 1

	Samuelson (1939) (SM39)	**Beaudry 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	$C_t = c(C_{t-1} + I_{t-1} + G)$; $I_t = \beta[c(C_{t-1} + I_{t-1} + G) - C_{t-1}]$	$I_t = \alpha_0 - \alpha_1(1-\delta)X_{t-1} + (\alpha_2 - \alpha_1)I_{t-1} + \frac{\alpha_3}{1+e^{-I_t}}$; $X_t = (1-\delta)X_{t-1} + I_{t-1}$	$Y_t = Y_{t-1} + \alpha[\sigma\mu + \gamma(\sigma\mu/\delta - K_{t-1}) + \arctan(Y_{t-1} - \mu) - \sigma Y_{t-1}]$; $K_t = (1-\delta)K_{t-1} + \sigma\mu + \gamma(\sigma\mu/\delta - K_{t-1}) + \arctan(Y_{t-1} - \mu)$
Parameterisation	$c = 0.4$ $\beta = 2$ $G = 0$	$\alpha_0 = -1$ $\alpha_1 = 0.4$ $\alpha_2 = 0.9$ $\alpha_3 = 2$ $\delta = 0.4$	$\alpha = 1.2$ $\gamma = 0.6$ $\delta = 0.2$ $\mu = 10$ $\sigma = 0.4$
Jacobian (at steady state)	$\begin{bmatrix} 0.4 & 0.4 \\ -1.2 & 0.8 \end{bmatrix}$	$\begin{bmatrix} 1 & -0.48 \\ 1 & 0.6 \end{bmatrix}$	$\begin{bmatrix} 1.72 & -0.72 \\ 1 & 0.2 \end{bmatrix}$
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 ($\lvert\lambda\rvert$),	$\lvert\lambda\rvert = 0.894$,	$\lvert\lambda\rvert = 1.039$,	$\lvert\lambda\rvert = 1.032$,
cycle length (L)	$L = 7.52$	$L = 9.08$	$L = 16.8$

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

Table 2 Overview of simulated DGPs, part 2

	Neubert and Kot (1992) (NK92)	**Wegener et al. (2009) (WWZ09)**	**Dieci and Westerhoff (2012) (DW12)**
Type	Predator-prey dynamics	Metzlerian inventory cycles with heterogeneous expectations	Asset price speculation in housing market
Interaction mechanism	Predator (y) and prey (x) population	Inventories (S)–expected sales (U)	Housing stock (z)–house prices (p)
Dynamics	Limit cycle	Limit cycle	Limit cycle
Type of nonlinearity	Second-order polynomial	Rational function	Rational function
Reduced form	$y_t = cx_{t-1}y_{t-1}$; $x_t = (r+1)x_{t-1} - rx_{t-1}^2 - cx_{t-1}y_{t-1}$	see Online Appendix A	$z_t = ep_{t-1} + dz_{t-1}$; $p_t = (1-c-e)p_{t-1} + \frac{fp_{t-1} - gh p_{t-1}^3}{1+hp_{t-1}^2} - dz_{t-1}$
Parameterisation	$c = 2.1$ $r = 3.1$	$b = 0.75$ $c = 0.3$ $d = 1$ $f = 0.1$ $k = 0.1$ $\bar{C} = 30$ $\bar{I} = 10$	$c = 0.05$ $d = 0.98$ $e = 0.5$ $f = 0.08$ $g = 1$ $h = 1$
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),	$\lambda = 0.262 \pm 1.039i$,	$\lambda = 0.911 \pm 0.492i$,	$\lambda = 1.115 \pm 0.687i$,
modulus ($\lvert\lambda\rvert$),	$\lvert\lambda\rvert = 1.071$,	$\lvert\lambda\rvert = 1.036$,	$\lvert\lambda\rvert = 1.310$,
cycle length (L)	$L = 4.75$	$L = 12.7$	$L = 11.4$

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

Figure 3: Simulated DGPs, deterministic (left) and stochastic (right)

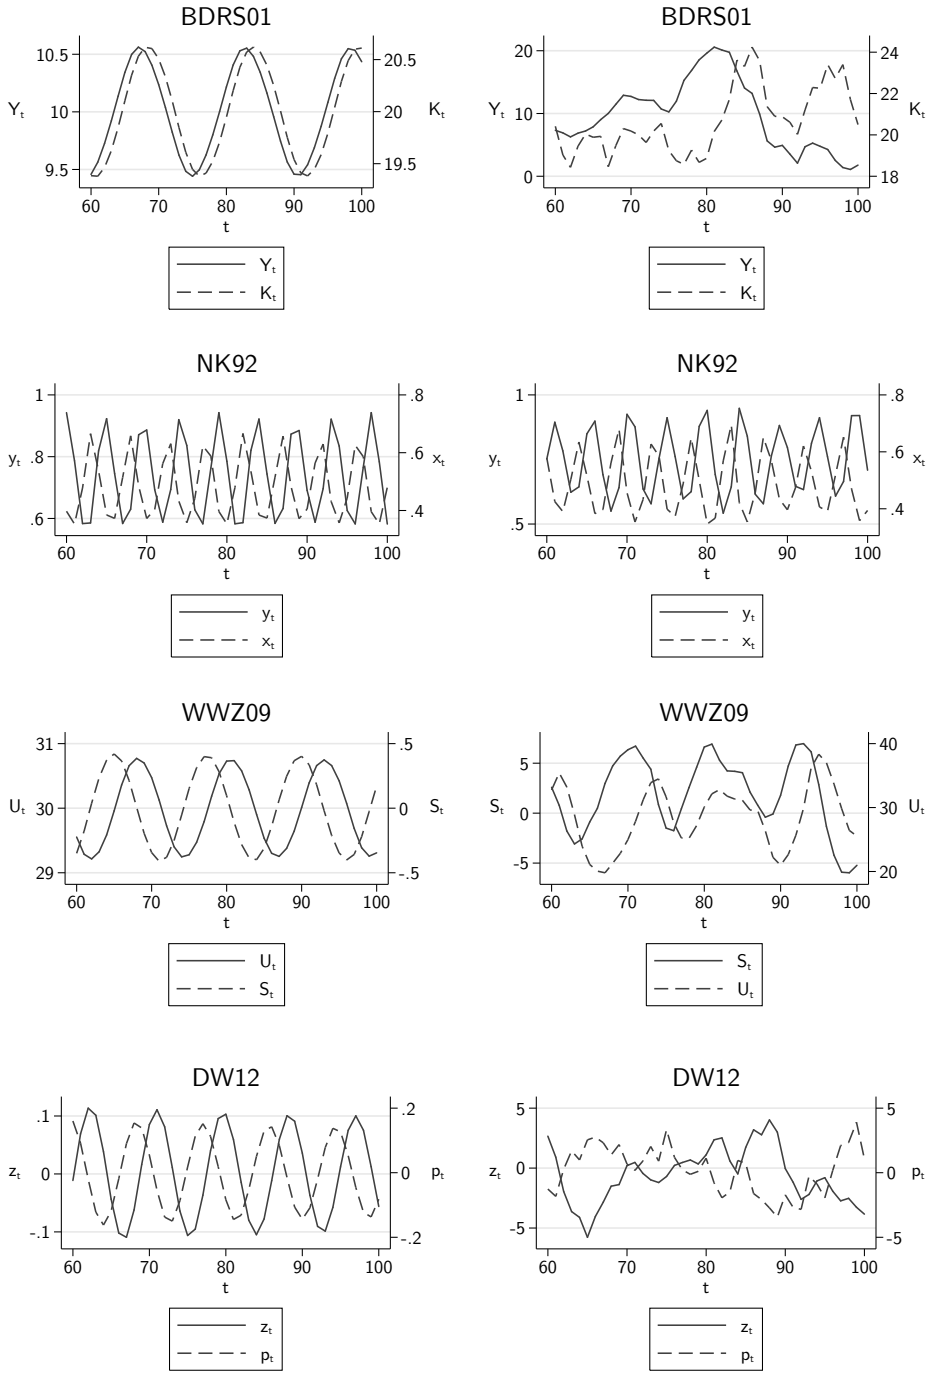

Notes: Parameterisations as in Tables 1-2. Stochastic versions (right panel) with serially correlated error terms $u_t = 0.8u_{t-1} + \epsilon_t$, with $\epsilon_t \sim N(0,1)$. For NK92, the initial condition was set to $y_0 = x_0 = 0.6$ and the error variance to $\sigma_\epsilon = 0.01$ to ensure trajectories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

which are rendered less regular and more persistent by additive serially correlated noise.

4 Monte Carlo design

For the Monte Carlo experiments we use stochastic versions of the DGPs reported in Tables 1-2. 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 C for a discussion of how we evaluate the interaction 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 to 20 periods (see Beneš and Vàvra 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 2, we are interested in the interaction mechanism summarised 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 – i.e. a_2 and b_1 in (1). Testing condition (2), i.e. $a_2 b_1 < 0$, 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 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 to 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]

5 Simulation results

Main results

Table 3 summarises 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 normalised root mean squared error (NRMSE) of the cycle mechanism $a_2 b_1$ 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, 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

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

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)
$\hat{a_2 b_1}$	-0.48	-0.49	-0.34	-0.30
bias ($\hat{a_2 b_1}$)	-0.20%	-1.77%	28.9%	37.4%
NRMSE ($\hat{a_2 b_1}$)	1.00	1.00	2.66	1.39
($\hat{a_2 b_1} < 0$)	100%	99.9%	100%	96.1%
($\hat{a_2 b_1} < 0$)*	100%	96.2%	99.9%	73.5%
$\hat{\lvert\lambda\rvert}_{dom}$	0.89	0.90	0.87	0.86
bias ($\hat{\lvert\lambda\rvert}_{dom}$)	-1.04%	0.21%	-16.5%	-16.8%
$\hat{re}$	0.59		0.65	
bias ($\hat{re}$)	-1.78%		-19.4%	
$\hat{im}$	0.66		0.58	
bias ($\hat{im}$)	-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	1000	895	1000	912

Notes: Total number of runs $N = 1,000$; n: number of VARs that did not suffer from serial correlation; sample size $T = 50$; $\hat{a_2 b_1}$: estimated cycle condition; bias: % deviation from absolute true value; NRMSE: Root mean squared error normalised by standard deviation; ($\hat{a_2 b_1} < 0$): relative frequency (in n) with which the cycle condition was satisfied; ($\hat{a_2 b_1} < 0$)*: relative frequency (in n) with which the condition was statistically significant at the 10% level; $\lvert\lambda\rvert_{dom}$: 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 $j = 1, .., p$.

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 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 E 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%).

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 $\hat{a_2b_1} = -0.34$ 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 to 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 real and imaginary part 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

[5]Online Appendix F provides further details on the sampling distribution of estimators for BGP20.

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

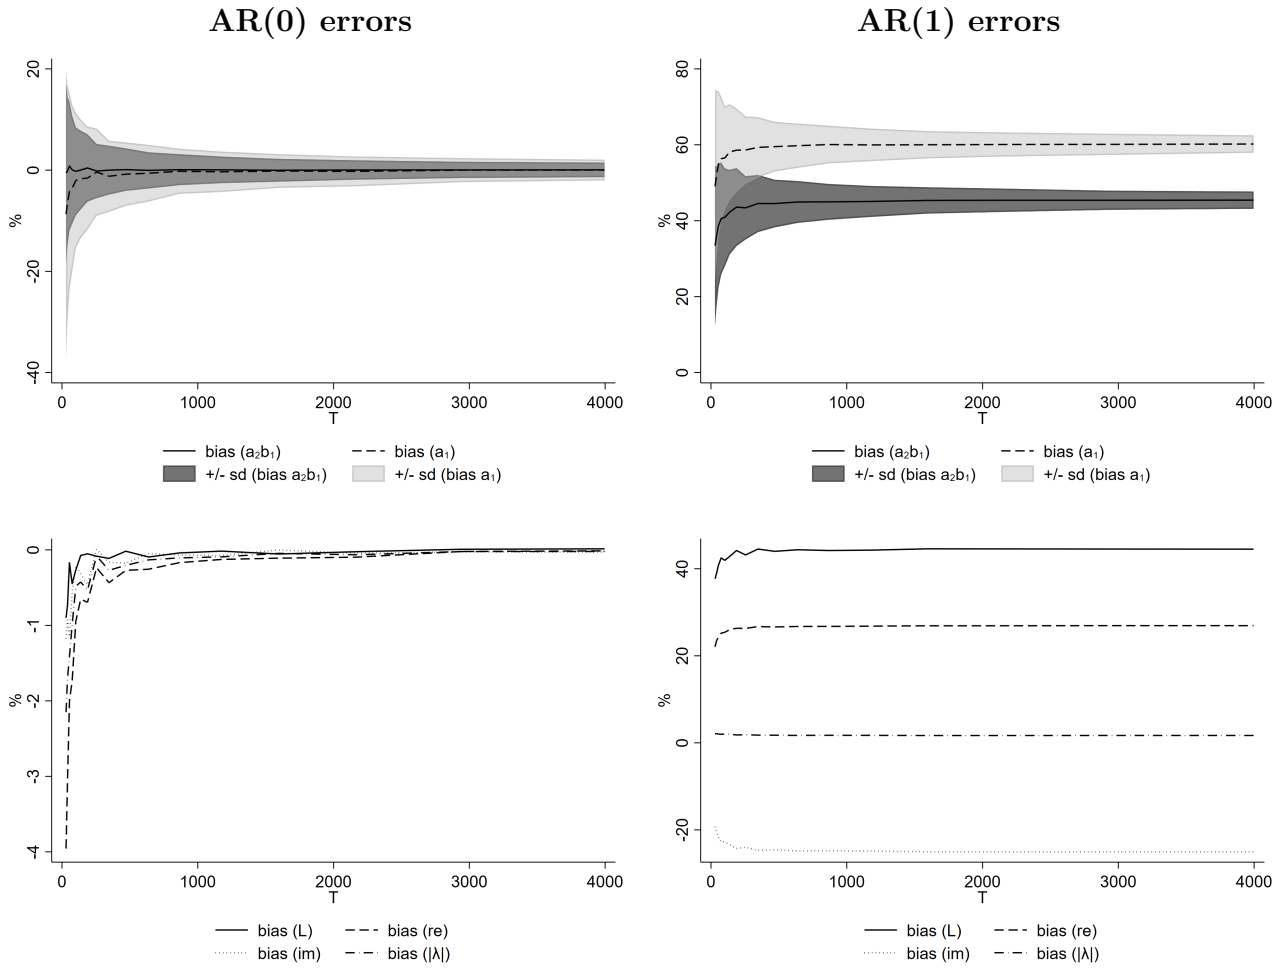

Notes: Lines represent the bias (in %) in estimated statistics from Monte Carlo simulations with $N = 1,000$ runs for increasing sample sizes $T = 30, ..., 4000$. Shaded areas denote one standard deviation around the lines.

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 to 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.

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

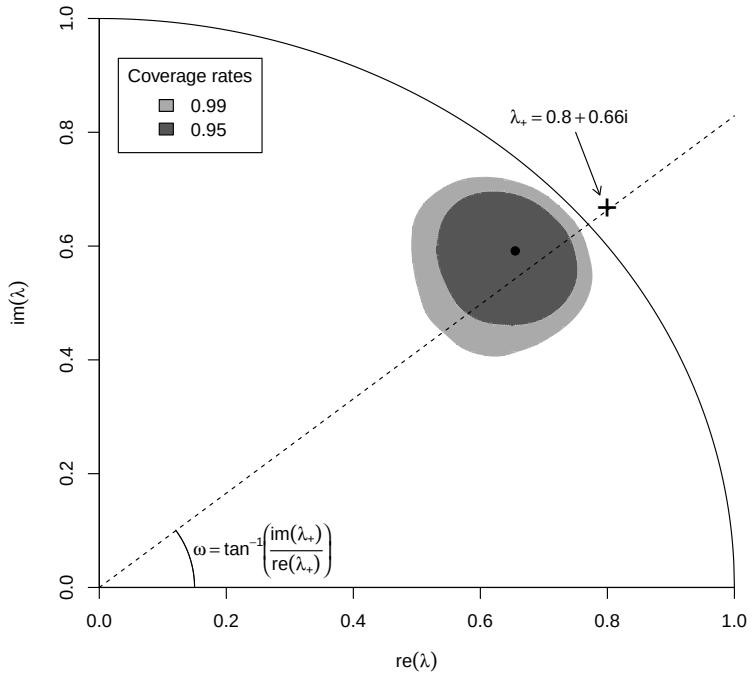

Notes: The grey shaded areas summarise the sampling distribution of the eigenvalues (limited to the positive quadrant of the Argand diagram) of a VAR(1) estimated on 10,000 runs of BGP20. The black point • marks the mode of the sampling distribution. The true eigenvalue of the linear part of BGP20 is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

This exercise suggests that the bias due to unmodelled nonlinearity 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]

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.

[6]Online Appendix G further explores the robustness to different parameterisations of BGP20. The results are very similar.

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

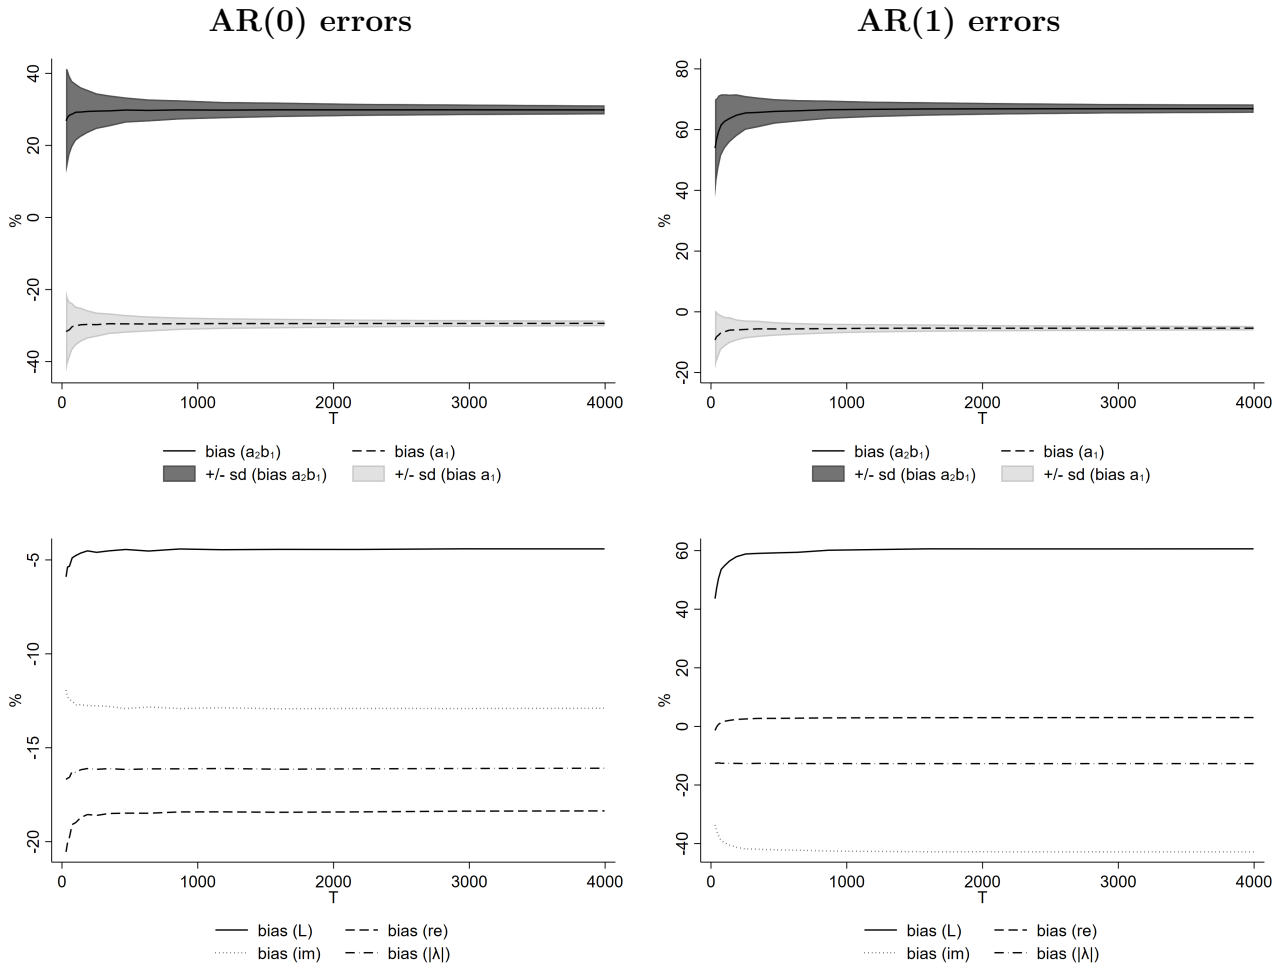

Notes: Lines represent the bias (in %) in estimated statistics from Monte Carlo simulations with $N = 1,000$ runs for increasing sample sizes $T = 30, ..., 4000$. Shaded areas denote one standard deviation around the lines.

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 summarised in Tables 1-2. 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.

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%
$\hat{\lvert\lambda\rvert}_{dom}$	0.71	0.83	0.96	0.99	0.87	0.90	0.46	0.82
bias ($\hat{\lvert\lambda\rvert}_{dom}$)	-31.4%	-19.3%	-10%	-7.68%	-16.4%	-13.5%	-65.2%	-37.4%
$\hat{re}$	0.60		0.23		0.74		0.35	
bias ($\hat{re}$)	-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	1000	913	1000	367	1000	887	1000	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 normalised 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; $\lvert\lambda\rvert_{dom}$: 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 $j = 1, .., p$.

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, 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 to NK92 (less than 3%). The better performance compared to 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 to 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

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

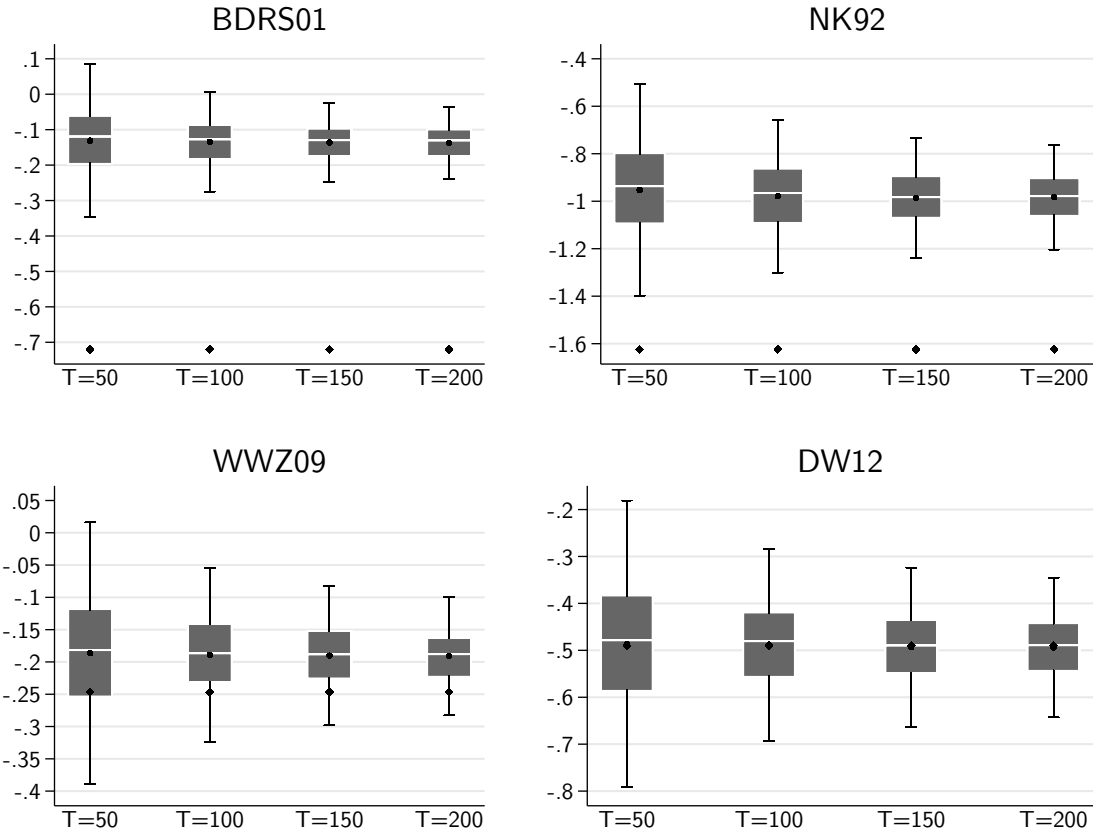

Notes: Boxes represent the first and third quartile, circles mark the mean, and whiskers represent 2 standard deviations around the mean. Diamonds mark the true values.

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 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 to 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 to 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. 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 2. 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 A and simulation results are reported in Online Appendix H.

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 produces 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 to 20 periods). An alternative approach is to use the cycle length implied by the dominant eigenvalue only. Online

Appendix E 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 to the baseline approach (-4.7% compared to -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 I). 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.

6 Conclusion

Recent macroeconomic research has revived the notion of endogenous business and financial cycles (Beaudry et al. 2020, Borio 2014, Stockhammer et al. 2019). 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 (Juselius and Drehmann 2020, Stockhammer et al. 2019, Strohsal et al. 2019). 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, 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 that 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.

References

Abadir, K. M., Hadri, K. and Tzavalis, E. (1999), 'The influence of VAR dimensions on estimator biases', *Econometrica* **67**(1), 163–181.

Aikman, D., Haldane, A. G. and Nelson, B. D. (2015), 'Curbing the Credit Cycle', *Economic Journal* **125**(585), 1072–1109.

Asada, T. (2001), Nonlinear Dynamics of Debt and Capital: A Post-Keynesian Analysis, *in* Y. Aruka, 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* **56**(4), 1538–1576.

Beaudry, P., Galizia, D. and Portier, F. (2016), Putting the Cycle Back into Business Cycle Analysis, NBER Working Paper 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* **31**(1), 479–530.

Beaudry, P., Galizia, D. and Portier, F. (2020), 'Putting the Cycle Back into Business Cycle Analysis', *American Economic Review* **110**(1), 1–47.

Beneš, J. and Vàvra, D. (2005), 'Eigenvalue filtering in VAR models with application to the Czech Business Cycle', *ECB Working Paper Series* (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* **11**(5), 527–554.

Borio, C. (2014), 'The Financial Cycle and Macroeconomics: What Have We Learnt?', *Journal of Banking & Finance* **45**, 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* **34**(8), 1492–1508.

Brock, W. A. and Hommes, C. H. (1997), 'A Rational Route to Randomness', *Econometrica* **65**(5), 1059–1095.

Brockwell, P. and Davis, R. (2006), *Time Series: Theory and Methods*, Springer.

Calvert Jump, R. and Levine, P. (2019), 'Behavioural New Keynesian models', *Journal of Macroeconomics* **59**, 59–77.

Chiarella, C. (1992), 'The Dynamics of Speculative Behaviour', *Annals of Operations Research* **37**(1), 101–123.

Dieci, R. and He, X.-Z. (2018), Heterogeneous agent models in finance, *in* C. Hommes and B. LeBaron, eds, '*Handbook of Computational Economics*', Vol. 4, Elsevier, pp. 257–328.

Dieci, R. and Westerhoff, F. (2012), 'A simple model of a speculative housing market', *Journal of Evolutionary Economics* **22**(2), 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* **71**, 21–44.

Doornik, J. A., Nielsen, B. and Rothenberg, T. J. (2003), 'The influence of VAR dimensions on estimator biases: Comment', *Econometrica* **71**(1), 377–383.

Fernández-Villaverde, J. and Rubio-Ramírez, J. F. (2005), 'Estimating dynamic equilibrium economies: linear versus nonlinear likelihood', *Journal of Applied Econometrics* **20**(7), 891–910.

Francq, C. and Zakoïan, J.-M. (1998), 'Estimating linear representations of nonlinear processes', *Journal of Statistical Planning and Inference* **68**(1), 145–165.

Galí, J. (2018), 'The State of New Keynesian Economics: A Partial Assessment', *Journal of Economic Perspectives* **32**(3), 87–112.

Goodwin, R. M. (1967), A Growth Cycle, *in* C. Feinstein, ed., 'Socialism, Capitalism, and Economic Growth', Cambridge University Press, Cambridge, UK.

Hamilton, J. D. (1994), *Time Series Analysis*, Princeton University Press, Princeton, New Jersey.

Juselius, M. and Drehmann, M. (2020), 'Leverage dynamics and the burden of debt', *Oxford Bulletin of Economics and Statistics* **82**(2), 347–364.

Kaldor, N. (1940), 'A Model of the Trade Cycle', *Economic Journal* **50**(197), 78.

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* **105**(2), 211–248.

Metzler, L. A. (1941), 'The Nature and Stability of Inventory Cycles', *The Review of Economics and Statistics* **23**(3), 113–129.

Neubert, M. G. and Kot, M. (1992), 'The subcritical collapse of predator populations in discrete-time predator-prey models', *Mathematical Biosciences* **110**, 45–66.

Oehlert, G. W. (1992), 'A Note on the Delta Method', *The American Statistician* **46**(1), 27–29.

Prescott, E. C. (1986), 'Theory Ahead of Business Cycle Measurement', *Federal Reserve Bank of Minneapolis Quarterly Review* **10**(4), 9–22.

Rünstler, G. and Vlekke, M. (2017), 'Business, housing, and credit cycles', *Journal of Applied Econometrics* **33**(2), 212–226.

Samuelson, P. A. (1939), 'Interactions between the Multiplier Analysis and the Principle of Acceleration', *The Review of Economics and Statistics* **21**(2), 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* **1**(5), 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* **94**, 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* **106**, 568–591.

Teräsvirta, T. (2018), 'Nonlinear models in macroeconometrics', *Oxford Research Encyclopedia of Economics and Finance* .

Teräsvirta, T., Tjøstheim, D. and Granger, C. W. J. (2010), *Modelling nonlinear economic time series*, Oxford University Press Oxford.

Wegener, M., Westerhoff, F. and Zaklan, G. (2009), 'A Metzlerian business cycle model with nonlinear heterogeneous expectations', *Economic Modelling* **26**(3), 715–720.

ONLINE APPENDIX

A Endogenous business cycle models

A.1 Samuelson (1939) [SM39]

The structural model in Samuelson (1939) is given by:

$$Y_t = G + C_t + I_t \tag{A.1}$$
$$I_t = \beta(C_t - C_{t-1}) \tag{A.2}$$
$$C_t = cY_{t-1}, \tag{A.3}$$

where Y_t: output; G: exogenous government expenditures; C_t: consumption; I_t: investment; β: accelerator effect of changes in consumption on investment; c: marginal propensity to consume. Shifting A.1 one period back and substitution into A.3, and substitution of A.3 in A.2 yields the reduced form:

$$C_t = c(C_{t-1} + I_{t-1} + G); \qquad c \in (0,1) \tag{A.4}$$
$$I_t = \beta[c(C_{t-1} + I_{t-1} + G) - C_{t-1}]; \qquad \beta > 0, \tag{A.5}$$

with Jacobian matrix:

$$J_{SM39}(C, I) = \begin{bmatrix} c & c \\ \beta(c-1) & \beta c \end{bmatrix}. \tag{A.6}$$

Since $c < 1$, we have $c\beta(c-1) < 0$, so that condition (2) for a cycle mechanism is satisfied: investment boosts consumption, which in turn drags down investment dynamics.

A.2 Bischi, Dieci, Rodano, and Saltari (2001) [BDRS01]

Bischi et al. (2001) propose a discrete time version of the Kaldor (1940) model. Unlike Samuelson (1939), Kaldor (1940) assumed an unstable goods market and introduced nonlinear saving and investment functions that prevent explosive oscillations. He assumed that the capital stock exerts a negative effect on investment due to a decreasing marginal efficiency of capital. Bischi et al. (2001) specify the sigmoid shaped investment function described in Kaldor (1940) as an arctangent-function.[7]

[7]The consumption function is linear as nonlinearity is not required to obtain limit cycles in the Kaldor

The structural model given by:

$$Y_t = Y_{t-1} + \alpha[I_{t-1}(Y_{t-1}, K_{t-1}) - S_{t-1}(Y_{t-1}, K_{t-1})] \tag{A.7}$$

$$K_t = K_{t-1} + I_{t-1}(Y_{t-1}, K_{t-1}), \tag{A.8}$$

where K_t: capital stock; S_t: saving, and all other variables as defined previously. Bischi et al. (2001) use the following saving and investment functions:

$$S_t = \sigma Y_t \tag{A.9}$$

$$I_t = \sigma\mu + \gamma(\frac{\sigma\mu}{\delta} - K_t) + \arctan(Y_{t-1} - \mu) \tag{A.10}$$

with $\alpha, \mu, \delta, \gamma > 0; \sigma \in (0, 1)$; where α: speed of adjustment of output to excess demand/supply; σ: propensity to save; δ: depreciation rate; μ: normal level of income (so that $\sigma\mu/\delta$: normal level of capital stock). With $\sigma(1 + \gamma/\delta) \geq 1$, there is a unique steady state $(Y^*, K^*) = (\mu, \mu\frac{\sigma}{\delta})$.

Substitution of A.9 and A.10 into A.7 and A.8 yields the reduced form:

$$Y_t = Y_{t-1} + \alpha[\sigma\mu + \gamma(\sigma\mu/\delta - K_{t-1}) + \arctan(Y_{t-1} - \mu) - \sigma Y_{t-1}]; \tag{A.11}$$

$$K_t = (1 - \delta)K_{t-1} + \sigma\mu + \gamma(\sigma\mu/\delta - K_{t-1}) + \arctan(Y_{t-1} - \mu), \tag{A.12}$$

with Jacobian matrix:

$$J_{BDRS01}(Y^*, K^*) = \begin{bmatrix} 1 + \frac{\alpha}{1+(Y^*-\mu)^2} - \alpha\sigma & -\alpha\gamma \\ \frac{1}{1+(Y^*-\mu)^2} & 1 - \delta - \gamma \end{bmatrix}. \tag{A.13}$$

It can be seen that the coefficients on the off-diagonal always exhibit opposite signs (independently of the steady state value Y^*). Higher output pushes up capital accumulation, whereas capital drags down output dynamics. The inherent nonlinearity creates local instability near the steady state but renders the system stable away from the steady state.

A.3 Neubert and Kot (1992) [NK92]

Population dynamics driven by the interaction between a predator hunting on a prey are a key subject of mathematical biology that inspired economic dynamics (e.g. Goodwin 1967, Kiyotaki and Moore 1997). Neubert and Kot (1992) provide a discrete time model that captures

model.

population dynamics with non-overlapping generations:

$$N_t = N_{t-1} + rN_{t-1}(1 + \frac{N_{t-1}}{K} - eN_{t-1}P_{t-1}) \tag{A.14}$$

$$P_t = bN_{t-1}P_{t-1} + (1-d), \tag{A.15}$$

where N_t and P_t: population densities of prey and predator population, respectively; K: carrying capacity of prey; e, b, d: foraging efficiency, birth rate, and death rate of the predator.

A rescaling of variables using $x_t = N_t/K$, $y_t = eP_t/bK$, and $c = bK$, yields:

$$y_t = cx_{t-1}y_{t-1} \qquad r > 0, c > 1 \tag{A.16}$$

$$x_t = (r+1)x_{t-1} - rx_{t-1}^2 - cx_{t-1}y_{t-1}. \tag{A.17}$$

A key feature of predator-prey models is the presence of three steady states: one at the origin which corresponds to the extinction of both species, one at $(0, 1)$ where the predator is extinct and the prey survives at its carrying capacity (here normalised to unity), and a positive one at (y^*, x^*) where both species coexist. The third steady state is given by $\left(\frac{r(c-1)}{c^2}, \frac{1}{c}\right)$. The Jacobian matrix is:

$$J_{NK92}(y^*, x^*) = \begin{bmatrix} cx^* & cy^* \\ -cx^* & (r+1) - 2rx^* - cy^* \end{bmatrix}. \tag{A.18}$$

For the third steady state, the interaction mechanism between the two species is captured by the opposite signs on the off-diagonal. In contrast to Goodwin (1967), where the coefficients on the main diagonal are zero (so that the only possible outcome is a centre), the discrete time predator-prey model allows for a richer set of dynamics. Neubert and Kot (1992) show that as c exceeds a critical value, a Neimark-Sacker bifurcation takes place giving rise to a stable limit cycle.

A.4 Wegener, Westerhoff and Zaklan (2009) [WWZ09]

Wegener et al. (2009) propose a nonlinear extension of Metzler (1941)'s linear inventory cycle model by introducing heterogeneous expectations. Firms can choose between two different strategies to predict future sales: an extrapolative predictor that extrapolates past trends in consumption and a regressive predictor that assumes consumption will return towards its equilibrium value. Firms tend to switch to a regressive rule when consumption is far away from the steady state, while extrapolative rules become more attractive for values of consumption near the steady state. The introduction of heterogeneous expectation formation and strategy

switching allows for a richer set of dynamics compared to Metzler (1941). The structural form of the model is given by:

$$Y_t = \bar{I} + S_t + U_t \tag{A.19}$$

$$S_t = kU_t - [kU_{t-1} - (bY_{t-1} - U_{t-1})] \tag{A.20}$$

$$U_t = w_t[bY_{t-1} + c(Y_{t-1} - \bar{C})] + (1 - w_t)(bY_{t-1} + f[\bar{C} - bY_{t-1}] \tag{A.21}$$

$$w_t = \frac{1}{1 + d(\bar{C} - bY_{t-1})^2}, \tag{A.22}$$

with $b \in (0, 1)$, $f \in [0, 1]$, $c \geq 1$, $k > 0$ and Y_t: output; $\bar{I}$: exogenous investment expenditures; S_t: production of consumption goods for stocks; U_t: expected sale of consumption goods; k: ratio of desired inventory stocks to expected sales; b: marginal propensity to consume; w_t: weight of extrapolative expectation formation in aggregate expected sales; c: speed of deviation from equilibrium consumption assumed by extrapolative expectation rule; $\bar{C}$: equilibrium level consumption;[8] f: adjustment speed towards equilibrium assumed by regressive expectation rule; d: popularity of regressive expectations. Substitution of A.19 into A.20 and A.21, and rearranging yields:

$$S_t = bS_{t-1} + kU_t + (b - k - 1)U_{t-1} + b\bar{I} \tag{A.23}$$

$$\begin{aligned} U_t = {} & [w_t b(1 + c) + (1 - w_t)b(1 - f)](S_{t-1} + U_{t-1} + \bar{I}) \\ & + [(1 - w - t)f - cw_t]\bar{C}. \end{aligned} \tag{A.24}$$

Elimination of the contemporaneous U_t yields a 3D representation:

$$\begin{aligned} S_t = {} & \{k[w_t b(1 + c) + (1 - w_t)b(1 - f)] + b\}(S_{t-1} + \bar{I}) + \\ & \{k[w_t b(1 + c) + (1 - w_t)b(1 - f)] + b - k - 1\}U_{t-1} + \\ & k[(1 - w_t)f - cw_t]\bar{C} \end{aligned} \tag{A.25}$$

$$\begin{aligned} U_t = {} & [w_t b(1 + c) + (1 - w_t)b(1 - f)](S_{t-1} + U_{t-1} + \bar{I}) \\ & + [(1 - w_t)f - cw_t]\bar{C} \end{aligned} \tag{A.26}$$

$$w_t = \frac{1}{1 + d[\bar{C} - b(\bar{I} + S_{t-1} + U_{t-1})]^2}. \tag{A.27}$$

A.27 can be used to eliminate w_t from A.25 and A.26, yielding a 2D representation. For the parameterisation considered in this paper (see Table 2), condition (2) is satisfied such that

[8] Note that in the steady state (S^*, U^*, w^*), $\bar{C} = S^* + U^*$.

an increase in expected sales slows down the accumulation of inventories, while an increase in inventory stocks boosts expected sales. As the parameters b, c, and k exceed critical values, the unique equilibrium loses its local stability and a supercritical limit cycle emerges.

A.5 Dieci and Westerhoff (2012) [DW12]

Dieci and Westerhoff (2012) study a model with strategy switching in speculative housing markets. House prices are driven by excess demand, while the price-elastic supply is determined by construction and depreciation of the housing stock. The demand for houses has a real and a speculative component. The real component is determined by current prices, whereas the speculative component is governed by expected future prices. As in Wegener et al. (2009), agents can choose between an extrapolative and a regressive forecasting rule. When prices are close to the steady state, the extrapolative rule becomes more attractive, but when prices deviate strongly from their fundamental values, more and more agents switch to the regressive rule. The structural form is given by:

$$P_t = P_{t-1} + a(D_{t-1} - S_{t-1}) \tag{A.28}$$

$$S_t = dS_{t-1} + eP_t \tag{A.29}$$

$$D_t = D_t^R + D_t^S \tag{A.30}$$

$$D_t^R = b - cP_t \tag{A.31}$$

$$D_t^E = f(P_t - \bar{P}) \tag{A.32}$$

$$D_t^{MR} = g(\bar{P} - P_t) \tag{A.33}$$

$$D_t^S = w_t D_t^E + (1 - w_t) D_t^{MR} \tag{A.34}$$

$$w_t = \frac{1}{1 + h(P_t - \bar{P})^2}, \tag{A.35}$$

where P_t: house price; a: sensitivity of houses prices to excess demand; D_t: total demand for houses; S_t: housing supply; $(1-d)$: depreciation rate of houses; e: sensitivity of housing construction to house prices; D_t^R: real demand for houses; D_t^S: total speculative demand for houses; b: autonomous real demand for houses; c: sensitivity of real housing demand to house prices; D_t^E: extrapolative component of speculative demand; f: sensitivity of extrapolative demand to deviations of house prices from their steady state value; $\bar{P}$: fundamental house price; D_t^{MR}: regressive component of speculative demand; g: sensitivity of regressive demand to deviations of house prices from their steady state value, h: popularity of regressive expectations.

Introducing the auxiliary variable $Z_t = S_{t-1}$, substitution of A.30-A.34 into A.28-A.29, and writing the variables in terms of deviation from their steady state ($z_t = Z_t - \bar{Z}$, $p_t = P_t - \bar{P}$)

yields:

$$z_t = ep_{t-1} + dz_{t-1}, \qquad e > 0, d \in (0,1) \tag{A.36}$$

$$p_t = (1-c-e)p_{t-1} + \frac{fp_{t-1} - ghp_{t-1}^3}{1+hp_{t-1}^2} - dz_{t-1}, \qquad c > 0. \tag{A.37}$$

If $f \leq c + e/(1-d)$, there is a unique steady given by $(z^*, p^*) = (0,0)$. The Jacobian matrix is:

$$J_{DW12}(z^*, p^*) = \begin{bmatrix} d & e \\ -d & 1-c-e+\frac{f-hp^{*2}(f+hgp^{*2}+3g)}{(1+hp^{*2})^2} \end{bmatrix}. \tag{A.38}$$

An increase in house prices leads to an increase in housing supply, which in turn feeds negatively back into house prices, so that condition (2) holds. As the extrapolative demand for houses becomes more price-sensitive (reflected in an increase in f), the steady state becomes locally unstable and a Neimark-Sacker bifurcation leads to a stable limit cycle.

A.6 BH97

Brock and Hommes (1997) pioneered a model of price formation in markets where supply depends on heterogeneous price expectations. Agents can switch between two different forecasting rules depending on their relative profitability: (i) a costly rational expectations predictor that correctly predicts the equilibrium price and (ii) a costless naive predictor that assumes future prices are equal to their present value. The structural form is given by:

$$D_t = n_{1t-1}S_{1t} + n_{2t-1}S_{2t-1} \tag{A.39}$$

$$D_t = A - Bp_t \tag{A.40}$$

$$S_{1t} = bp_{t+1} \tag{A.41}$$

$$S_{2t} = bp_t \tag{A.42}$$

$$n_{1t} = \frac{\exp(\beta(\pi_{1t} - C))}{\exp(\beta(\pi_{1t} - C)) + \exp(\beta\pi_{2t})} \tag{A.43}$$

$$n_{2t} = 1 - n_{1t} \tag{A.44}$$

$$\pi_{1t} = \frac{b}{2}p_{t+1}^2 - C \tag{A.45}$$

$$\pi_{2t} = \frac{b}{2}p_t(2p_{t+1} - p_t) \tag{A.46}$$

$$m_t = n_{1t} - n_{2t}, \tag{A.47}$$

where D_t: demand; S_t: supply; S_{1t}, S_{2t}: supply of producers that follow a rational expectations and a naive expectations rule, respectively; n_{1t}, n_{2t}: share of producers that follow a rational expectations and a naive expectations rule, respectively; A: autonomous demand; B: sensitivity of demand with respect to price; p_t: price; b: sensitivity of supply with respect to (expected) price; π_{1t}, π_{2t}: realised profits of producers that follow a rational expectations and a naive expectations rule, respectively; β: speed with which producers switch between predictors; C: information cost of the rational expectations predictor; m_t: difference between the fraction of agents that uses the rational expectations predictor.

Defining p_t in terms of deviations from its steady state by setting $A = 0$ and then substituting A.40-A.47 into the equilibrium condition A.39 and solving for p_t yields A.48. Substitution of A.43-A.46 into A.47 and elimination of the contemporaneous value for p_t through substitution yields A.49:[9]

$$p_t = \frac{-b(1-m_{t-1})}{2B+b(1+m_{t-1})}p_{t-1} \qquad b, B > 0 \qquad \text{(A.48)}$$

$$m_t = \tanh\left\{\frac{\beta}{2}\left[\frac{b}{2}\left(\frac{b(1-m_{t-1})}{2B+b(1+m_{t-1})}+1\right)^2 p_{t-1}^2 - C\right]\right\} \qquad \beta, C > 0, \qquad \text{(A.49)}$$

The model has a unique fixed point at $(p^*, m^*) = (0, \tanh(-\beta C/2))$. At this steady state, the Jacobian matrix is given by:

$$J_{BH97}(p^* = 0, m^*) = \begin{bmatrix} \frac{-b(1-m^*)}{2B+b(1+m^*)} & 0 \\ 0 & 0 \end{bmatrix}. \qquad \text{(A.50)}$$

As the parameter β exceeds a critical value, a strange attractor emerges that gives rise to irregular fluctuations.

[9] Using the fact that $\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$.

B Derivation of steady state of BGP20

Imposing $I_t = I_{t-1} = I^*$ and $X_t = X_{t-1} = X^*$ in (10)-(11) yields:

$$I^* = \alpha_0 - \alpha_1(1-\delta)X^* + (\alpha_2 - \alpha_1)I^* + \frac{\alpha_3}{1+e^{-I^*}}, \qquad \alpha_3 > 0 \tag{B.1}$$

$$X^* = (1-\delta)X^* + I^*. \tag{B.2}$$

Solving B.2 for X^* and substitution into B.1 yields:

$$I^*(1 + \frac{\alpha_1}{\delta} - \alpha_2) = \alpha_0 + \frac{\alpha_3}{1+e^{-I^*}} \tag{B.3}$$

$$X^* = \frac{I^*}{\delta}. \tag{B.4}$$

For $\alpha_0 = -\alpha_3/2$, we have $F(0) = \alpha_0 + \frac{\alpha_3}{1+e^{-0}} = 0$, so that the steady state is given by $I^* = X^* = 0$.

C Assessing the interaction mechanism when the lag length is greater than 1

The VAR(1) in (1) is consistent for both the linear projection and condition (2), and estimates of the eigenvalues of the linear projection can be obtained from the estimated coefficient matrix. In practice, however, estimated VAR(1) models will often suffer from serially correlated errors of the form:

$$\begin{bmatrix} u_{yt} \\ u_{ft} \end{bmatrix} = \sum_{i=1}^{p} A_i \begin{bmatrix} u_{yt-i} \\ u_{ft-i} \end{bmatrix} + \begin{bmatrix} \epsilon_{yt} \\ \epsilon_{ft} \end{bmatrix}. \tag{C.1}$$

In our Monte Carlo simulation, we emulate this practical problem through the use of AR(1) shock processes. As discussed in Stockhammer et al. (2019), the VAR(1) in (1) in conjunction with the serially correlated error terms (C.1) can be rewritten as an unrestricted VAR(p). The structural parameters driving any cycles are then assumed to be the parameters in (1) – i.e. the coefficient matrix on the first lag. A consequence is that the estimated eigenvalues of the VAR(p) are nonlinear functions of the coefficients in the companion matrix of the VAR(p) and therefore no longer correspond directly to the eigenvalues of the VAR(1). If we assume that A_i is diagonal then a_2 and b_1 in (1) are the only coefficients that are still identified, and thus the critical condition for the existence of a cycle mechanism in (2) can be evaluated.

D Asymptotic bias in VAR(p)

Figure A1: Asymptotic bias in VAR(p) with AR(1) errors

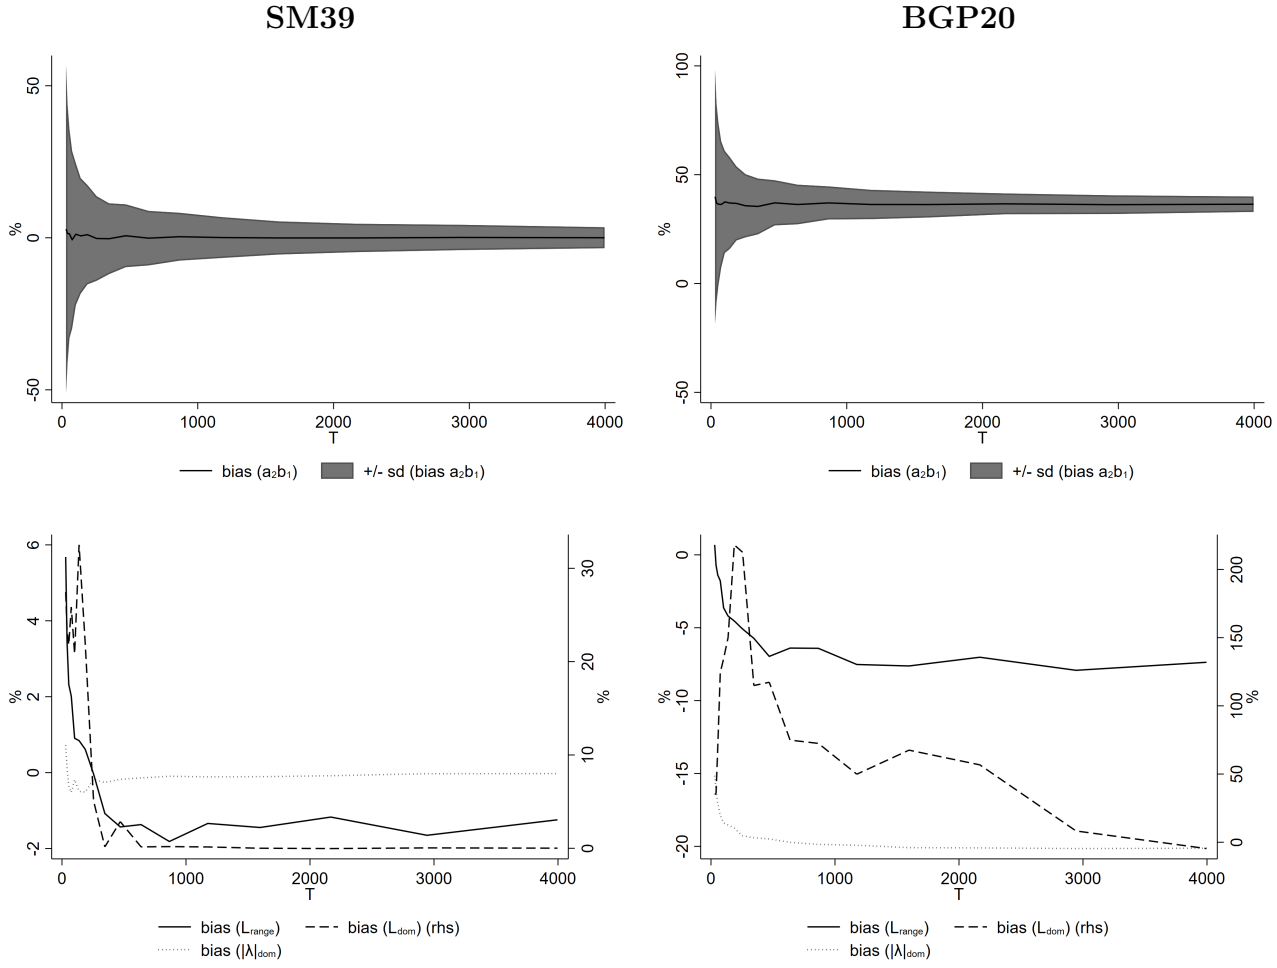

Notes: Lines represent the bias (in %) in estimated statistics from Monte Carlo simulations with $N = 1,000$ runs for increasing sample sizes $T = 30, ..., 10000$. L_{range}: average cycle length over the range $(3, 20)$; L_{dom}: cycle length associated with dominant eigenvalue. Bias of L_{dom} is displayed on the right-hand side axis.

E Distribution of statistics, BGP20

Figure A2: Histograms, BGP20, VAR(1) with AR(0) shocks

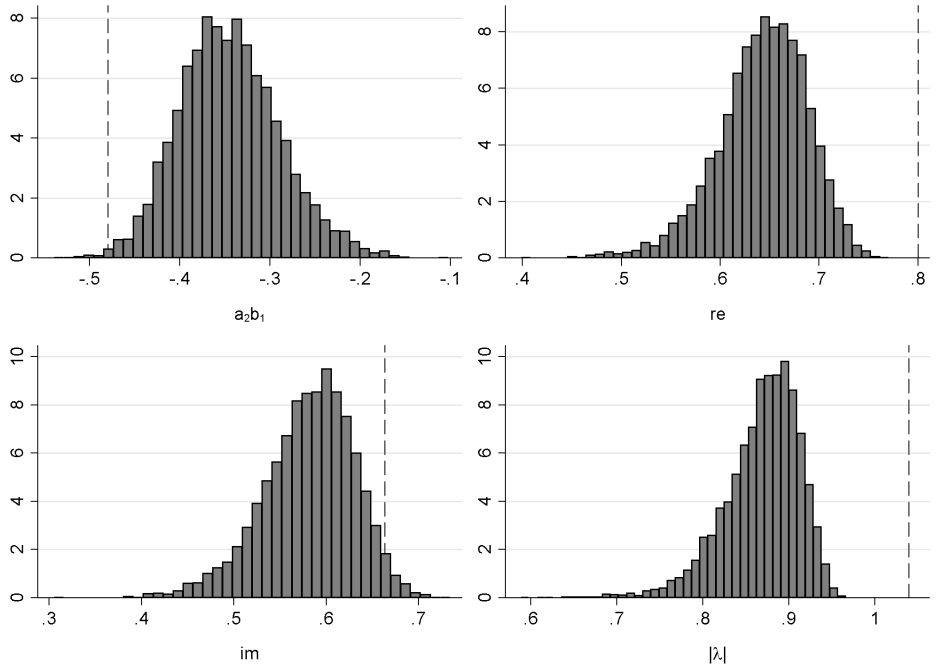

Notes: Number of Monte Carlo runs: $N = 1000$. Vertical lines mark the true value.

F Different parameterisations for BGP20

Figure A3: Different parameterisations of $\alpha 1$, BGP20, AR(0) and AR(1) shocks

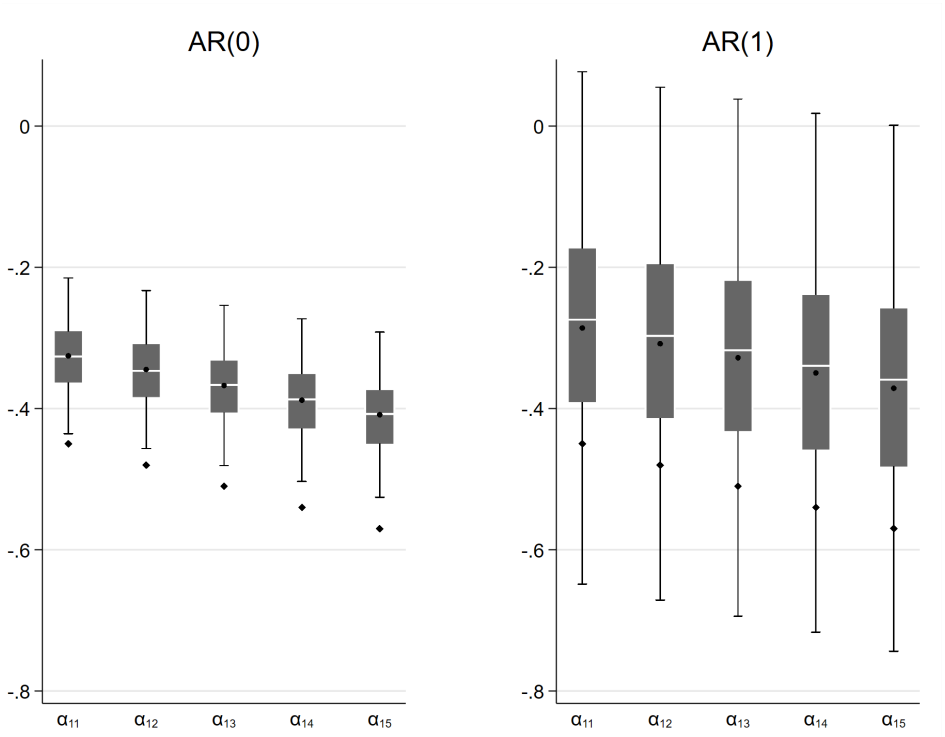

Notes: Boxes represent the first and third quartile, circles mark the mean, and whiskers represent 2 standard deviations around the mean. Diamonds mark the true value. $\alpha_{11} = 0.375$, $\alpha_{12} = 0.4$ (baseline), $\alpha_{13} = 0.425$, $\alpha_{14} = 0.45$, $\alpha_{15} = 0.475$. Sample size: $T = 50$.

The model exhibits limit cycle dynamics for all parameterisations used in Figure A3, but the strength of the interaction mechanism increases successively. It can be seen that the true magnitude of the interaction effect is underestimated, but the bias remains constant across parameterisations. The presence of serially correlated shock processes (right panel) renders the estimator less precise. A higher magnitude of the interaction effect increases the cycle detection rate.

G Sensitivity to false positives

Figure A4: Simulation of BH97, deterministic and stochastic

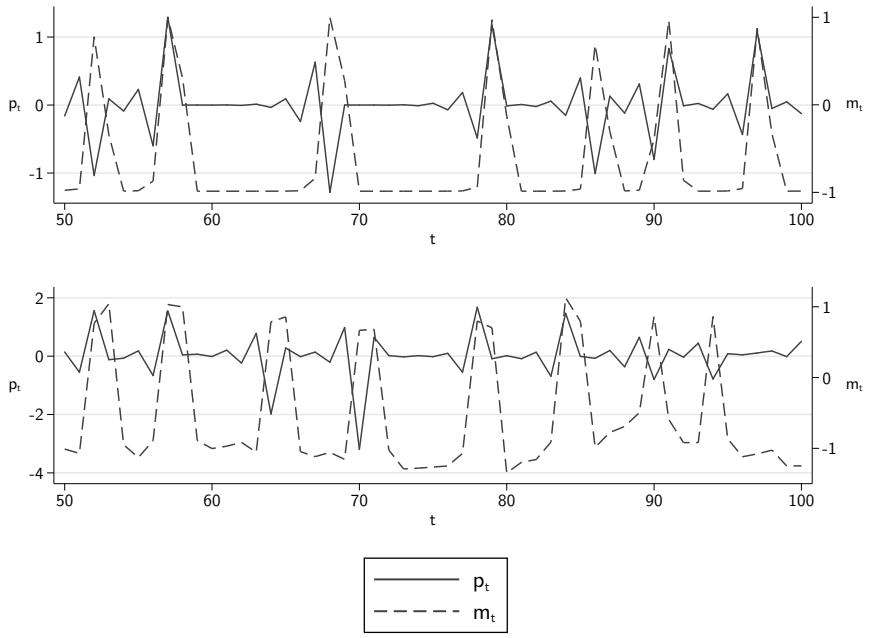

Notes: Simulation of $p_t = \frac{-b(1-m_{t-1})}{2B+b(1+m_{t-1})}p_{t-1}$, $m_t = \tanh\left\{\frac{\beta}{2}\left[\frac{b}{2}\left(\frac{b(1-m_{t-1})}{2B+b(1+m_{t-1})}+1\right)^2 p_{t-1}^2 - C\right]\right\}$ with $b = 1.35$, $B = 0.5$, $C = 1$, and $\beta = 5$. Upper panel: deterministic version. Lower panel: stochastic version with serially correlated error terms $u_t = 0.8u_{t-1} + \epsilon_t$ with $\epsilon_t \sim N(0,1)$.

Table A1 Results of BH97, VAR(1) with AR(0) shocks and VAR(p) with AR(1) shocks, averages

	VAR (1)	VAR(p)
$\hat{\alpha_2\beta_1}$	-0.0064	-0.0047
NRMSE $\hat{a_2b_1}$	1.03	1.02
$(\hat{\alpha_2\beta_1} < 0)$ (% of n)	64.7%	58.7%
$(\hat{\alpha_2\beta_1} < 0)^*$ (% of n)	0.6%	0.95%
$\hat{\lambda}_{dom}$	0.49	0.79
bias $(\hat{\lambda}_{dom})$	-81.2%	-70.1%
Nonnorm (% of n)	72.1%	39%
Lags	1	2.66
Serial corr(% of N)		16%
n	1000	840

Notes: Based on $p_t = \frac{-b(1-m_{t-1})}{2B+b(1+m_{t-1})}p_{t-1}$, $m_t = \tanh\left\{\frac{\beta}{2}\left[\frac{b}{2}\left(\frac{b(1-m_{t-1})}{2B+b(1+m_{t-1})}+1\right)^2 p_{t-1}^2 - C\right]\right\}$ with $b = 1.35$, $B = 0.5$, $C = 1$, and $\beta = 5$. Jacobian matrix: $a_{11} = -2.634$, $a_{12} = a_{21} = a_{22} = 0$. Total number of runs $N = 1,000$; n: number of VARs that did not suffer from serial correlation; sample size $T = 50$; $\hat{a_2b_1}$: estimated cycle condition; bias: % deviation from absolute true value; NRMSE: Root mean squared error normalised by standard deviation; $(\hat{a_2b_1} < 0)$: relative frequency (in n) with which the cycle condition was satisfied; $(\hat{a_2b_1} < 0)^*$: relative frequency (in n) with which the condition was statistically significant at the 10% level; $|\lambda|_{dom}$: 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.

Figure A5: Box plot for $a_2\hat{b}_1$, BH97, different sample sizes, VAR(p) with AR(1) shocks

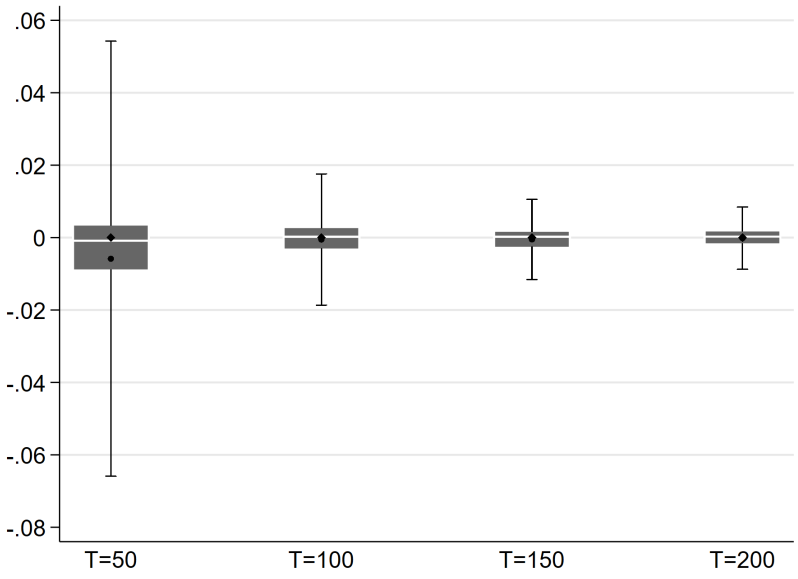

Notes: Boxes represent the first and third quartile, circles mark the mean, and whiskers represent 2 standard deviations around the mean. Diamonds mark the true value.

H Lag selection based on AIC

Table A2 Results for all DGPs, VAR(p) with AR(1) shocks, lag selection based on AIC, averages

	SM39	**BGP20**	**BDRS01**	**NK92**	**WWZ09**	**DW12**	**BH97**
$\hat{a_2 b_1}$	-0.48	-0.31	-0.13	-0.94	-0.19	-0.48	-0.0071
bias ($\hat{a_2 b_1}$)	0.16%	36%	82%	42%	23%	1.4%	
$NRMSE(\hat{a_2 b_1})$	1	1.39	5.58	3.28	1.14	1	1.02
$(\hat{a_2 b_1} < 0)$	100%	97%	90%	100%	97%	100%	55%
$(\hat{a_2 b_1} < 0)^*$	95%	78%	56%	100%	83%	98%	2%
$\lvert\hat{\lambda}\rvert_{dom}$	0.89	0.86	0.84	0.99	0.89	0.82	0.8
bias ($\lvert\hat{\lambda}\rvert_{dom}$)	-0.15%	-16.8%	-19%	-7.73%	-13.7%	-37.2%	-69.7%
Nonnorm (% of N)	5.8%	6.7%	5.3%	15%	6.4%	6.3%	36%
Lags	2.25	2.21	2.24	3.52	2.29	2.25	2.72
Serial corr (% of N)	20%	20%	20%	68%	19%	19%	30%

Notes: Total number of runs $N = 1,000$; sample size $T = 50$; $\hat{a_2 b_1}$: estimated cycle condition; bias: % deviation from absolute true value; NRMSE: Root mean squared error normalised by standard deviation; $(\hat{a_2 b_1} < 0)$: relative frequency (in n) with which the cycle condition was satisfied; $(\hat{a_2 b_1} < 0)^*$: relative frequency (in n) with which the condition was statistically significant at the 10% level; $\lvert\lambda\rvert_{dom}$: dominant modulus; *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.