Learning Macroeconomic Dynamics from Data: Applications of Universal Differential Equations and SINDy

Tato Khundadze*

Working Paper
2026

Abstract

This paper introduces scientific machine learning methods to macroeconomic model discovery, focusing on Sparse Identification of Nonlinear Dynamics (SINDy), Universal Differential Equations (UDEs), and a hybrid UDE–SINDy pipeline in which a neural network embedded within a differential equation learns unknown dynamics that are subsequently distilled into interpretable symbolic equations via sparse regression. The Goodwin growth cycle, a canonical predator–prey formulation of distributional conflict between labour and capital, serves as the testing ground throughout. We survey the relevant methodology, implement each approach in a unified computational framework, and benchmark them on problems calibrated to macroeconomic scales, where data are scarce, noisy, and nonstationary. On stationary synthetic Goodwin data, the hybrid pipeline recovers the bilinear interaction coefficients to four significant figures, with the vu term emerging as the sole active component of the sparse representation. On nonstationary synthetic data, consisting of a four-regime piecewise system with forced oscillations, abrupt parameter shifts, and transition damping, both direct SINDy and UDE–SINDy recover the predator–prey structure with correct structural identification. These experiments establish that the methods can operate under conditions characteristic of macroeconomic time series. When the validated pipeline is applied to 77 years of U.S. quarterly data (1948:Q1–2025:Q3) on employment, capacity utilisation, and the labour share, neither method recovers the vu interaction at any regularisation strength. Direct SINDy returns dense polynomial models; the UDE–SINDy pipeline finds near-zero neural network corrections and no dominant interaction term. The result is consistent across both methods and both empirical specifications, classical (v, u) and structuralist (TCU, u). The Lotka–Volterra functional form that defines the Goodwin model is not present in the data. Several caveats apply, including the restriction to second-order polynomial libraries and the treatment of the system as a two-dimensional autonomous ODE. The methods developed here are not specific to the Goodwin model and can be applied to other macroeconomic dynamical systems for which candidate functional forms can be specified.

Keywords: Goodwin Model, Universal Differential Equations, SINDy, Nonlinear Dynamics, Model Discovery, Macroeconomic Modeling, Falsification
JEL Classification: C63, C61, C45, C52, C88

*The New School for Social Research. Email: `khunt758@newschool.edu`.

There is increasing enthusiasm surrounding the application of Artificial Intelligence across diverse domains, including natural language processing for chatbots, image recognition, and complex engineering tasks. Economics has also witnessed a growing use of AI/ML techniques. Nevertheless, the application of ML for model discovery in macroeconomics remains relatively underexplored. To the best of our knowledge, the only recent contribution in this area is by Li et al. (2025), who employ the SINDy method to identify governing equations in macroeconomic systems. In this paper, we present a comprehensive overview of existing model discovery methods developed within the Scientific ML literature and demonstrate their practical implementation on macroeconomic models using synthetic data.

So far, macroeconomic modeling has been based on heuristics developed by macroeconomic modelers, and then tested through econometric methods. One of the enduring difficulties in the field of macroeconomic modeling lies in the scarcity of sufficiently rich data to support the training of modern machine learning algorithms. In consequence, macroeconomic analysis remains, for the most part, grounded in mechanistic constructs drawn from prior theoretical knowledge. This dependence, while offering a semblance of structure, often limits the predictive power of such models. There exists a fundamental trade-off between adhering to theory-driven models and embracing the flexibility of purely data-driven approaches. Mechanistic models, however refined, are burdened by the assumptions on which they are built, assumptions that are not infrequently shaped as much by ideology as by empirical observation. It is not uncommon, therefore, for data to be employed not in the service of falsification or refinement, but rather to lend post hoc legitimacy to the model itself. This inversion of scientific method, where models dictate the interpretation of data, carries with it the risk of distortion, and in extreme cases, deliberate manipulation. Economists, perhaps more than others, are fond of quoting George Box's dictum that "all models are wrong, but some are useful." Yet in practice, it is often the first clause that is invoked to excuse rather than to challenge, while the second, the demand for utility, is too lightly held. As Box himself remarked, "Models, of course, are never true, but fortunately it is only necessary that they be useful. For this it is usually needful only that they not be grossly wrong" (Box, 1979, p. 2). On the other hand, purely data-driven models could be problematic, in terms of data for fields like macroeconomics, but also in terms of explainability. While ML models could be good in terms of forecasting, frequently, they lack the ability to explain underlying mechanisms of the system.

In this paper we employ two complementary frameworks for equation discovery: (i) sparse identification of nonlinear dynamical systems (SINDy), and (ii) universal differential equations (UDEs), including a hybrid UDE–SINDy pipeline in which a neural network learns unknown dynamics that are subsequently distilled into symbolic equations via sparse regression. A third framework, Kolmogorov–Arnold Network Ordinary Differential Equations (KAN-ODEs), is reviewed for completeness in Appendix A.3; its substantially higher computational cost makes the extensive parameter sweeps conducted in the main text impractical within a single study. The SINDy algorithm was introduced by Brunton et al. (2016) to identify sparse (parsimonious) models directly from data. The key distinction of this framework is its emphasis on parsimony: the best model is not necessarily the one with the most variables, but the one that can adequately capture system dynamics using the fewest terms. To achieve this, the SINDy algorithm employs a library of candidate functions combined with regularized regression, with the goal of balancing model complexity and accuracy, similar to how overfitting is avoided in statistical learning. Over time, several extensions of SINDy have been developed, including those proposed by Fasel et al. (2021), Fasel et al. (2022), and Voina et al. (2024). The method has been widely applied across disciplines such as engineering, physics, and biology (Kaptanoglu, 2021; Corbetta, 2020; Mangan et al., 2016).

Universal Differential Equations provide a second framework for data-driven discovery, which like SINDy is based on discovering models from given measurements. This unifying framework was introduced by Rackauckas et al. (2021). However, there are fundamental differences between the two approaches. Specifically, UDE uses two components for model discovery: (i) a mechanistic part, which is derived from already known physics of the dynamical system, and (ii) a data-driven component discovered via artificial neural networks. As shown by Rackauckas et al. (2021), these two frameworks can be used in conjunction: known and unknown components can be trained via ANN to reduce the complexity of the unknown component, and as a next step, SINDy can be used to generate an interpretable symbolic model. This hybrid UDE–SINDy pipeline is the central methodological tool of the present paper. The framework is relatively new, but there is already active research in terms of further development and extending applications (Schmid et al., 2025; Rooij et al., 2025; de Rooij et al., 2025).

We apply both methods, direct SINDy and the hybrid UDE–SINDy pipeline, to the Goodwin growth cycle model, first validating their capacity to recover known dynamics from synthetic data (including under nonstationarity and measurement noise), and then deploying them on real U.S. macroeconomic time series. The paper's principal finding is negative: a pipeline that recovers Goodwin's predator–prey interaction with high precision from synthetic data finds no such structure in 77 years of quarterly observations. We argue that this constitutes a rigorous, method-independent falsification of the Goodwin model's specific Lotka–Volterra functional form as a description of U.S. distributional dynamics.

1 Motivation and significance of the Goodwin model in macroeconomic dynamics.

Goodwin (1967) developed analytical model of macroeconomic cyclical growth, which is inspired by Marxian analysis(Blecker and Setterfield, 2019). Specifically, profit-squeeze argument has following logical structure: during the period of high profitability and higher rate of investment, there is higher rate of growth for demand of labor. This results in falling profit share, which in Marxian analysis is rate of exploitation, and at some point, reduces profit rate. Diminished profit rate contributes to slower growth rate, and it leads to recession. After the downturn, as in the low growth regime, unemployment rises and with it labor bargaining power goes down and real wages decline. Lower level of real wages, contributes to rising profit share and profit rate, at some point reinstates profitability. This in its turn pushes up accumulation, and increases employment and real wage. This cyclical pattern is summarized on Figure 1, and as it demonstrates, there is clear logical steps how Marx formulated business cycle process, and there is no need exteral shocks, to get this this cyclical patterns.

Goodwin characterizes capitalism as a nonlinear system that is "almost always to be found in non-steady state" (Goodwin, 1987, as cited in Desai and Ormerod, 1998). According to this view, the economy consists of numerous decentralized agents operating under imperfect information, which limits fully coordinated or optimal decision-making. The resulting dynamics generate persistent instability in economic growth and give rise to distinct growth regimes (Desai and Ormerod, 1998).

Work of Goodwin has multiple importance from different perspective. It presents capitalism as non-linear economic system, which is characterized by endogenous shocks. Standard economic perspective is that without exogenous shocks, capitalist economic system is stable. Important to note, that some mainstream macroeconomic theories, such as rational expectations school of business cycle allow existence of cycles, but the cycles are explained by random shock which perturbs inherently stable systems. Different interpretation and extension of works of Goodwin, is important from the standpoint of better understanding of endogenously produced instability within capitalist system. There are numerous works on the empirical estimation of Goodwin model, which are usually done using econometric techniques(Harvie, 2000; Barbosa-Filho and Taylor, 2006; Grasselli and Maheshwari, 2018; Araujo et al., 2019). In this paper, we use machine learning approach, to discover governing equations from macro-informed simulated data, than compare to theoretical patterns the dynamics produced by the system of differential equations.

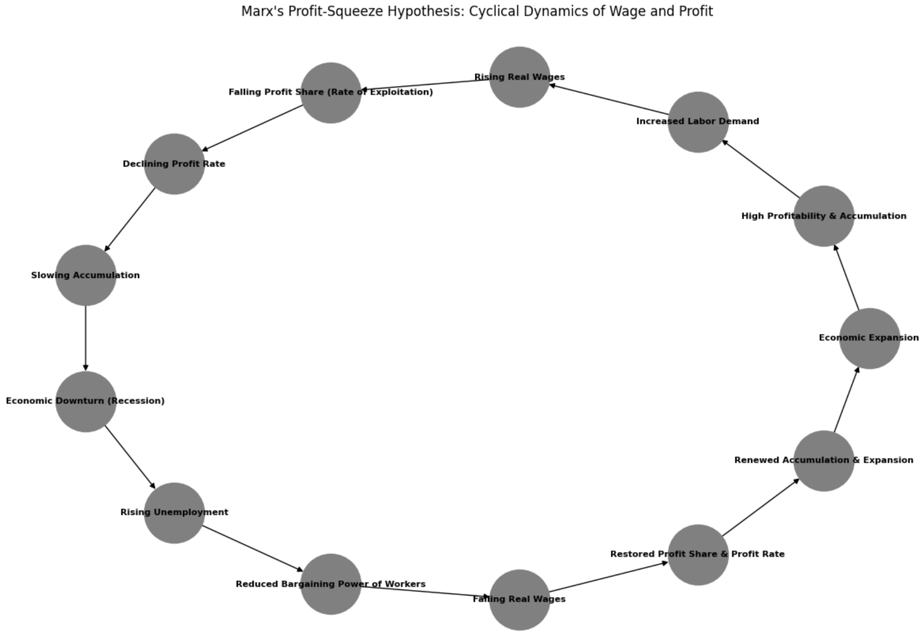

Figure 1: Marx's Profit-Squeeze Hypothesis: Cyclical Dynamics of Wage and Profit

1.1 Basic Set-Up of the Goodwin Model

The Goodwin (1967) model describes a distribution–employment cycle in which labor-market tightness drives wage claims, while the resulting shift in functional income distribution feeds back into accumulation and employment. The model rests on five behavioral and structural assumptions: (i) a real-wage Phillips curve linking real-wage growth to the employment rate, (ii) constant growth of labor productivity at rate $\alpha > 0$, (iii) a fixed capital–output ratio σ, (iv) classical savings behavior in which all profits are invested and depreciation is ignored, and (v) constant labor-force growth at rate $\beta > 0$. Under these assumptions, the dynamics of the two endogenous variables, the employment rate $v = l/n$ and the wage share $u = w/a$, can be expressed as a two-dimensional autonomous nonlinear system. For a step-by-step derivation, we refer the reader to Basu (2024); here we present the model directly in its reparametrized form, following the notation in Basu (2024) and Semmler (1986).

Define the composite parameters

$$\eta_1 = \frac{1}{\sigma} - \alpha - \beta, \qquad \eta_2 = \alpha + \gamma, \qquad \theta_1 = \frac{1}{\sigma}, \qquad \theta_2 = \rho, \tag{1}$$

where $\gamma > 0$ and $\rho > 0$ are the intercept and slope of the real-wage Phillips curve, respectively, and all four composite parameters are assumed strictly positive. The Goodwin system then takes the form

$$\begin{aligned} \dot{v} &= (\eta_1 - \theta_1 \, u) \, v, \\ \dot{u} &= (-\eta_2 + \theta_2 \, v) \, u. \end{aligned} \tag{2}$$

The logic of the cycle runs as follows. When employment is high, workers can bargain for faster real-wage growth, so the wage share rises. A rising wage share compresses profits, slows capital accumulation, and therefore reduces employment growth. Lower employment weakens labor's bargaining position, real-wage growth decelerates, and the wage share begins to fall, eventually restoring profitability and restarting the cycle. In the predator–prey analogy that gives the model its analytical structure, the employment rate v plays the role of the prey and the wage share u acts as the predator (Semmler, 1986).

The parameters η_1 and η_2 govern the autonomous dynamics of each variable: η_1 is the trend growth rate of employment when the entire output accrues to profits ($u = 0$), while η_2 is the rate at which the wage share would decline if the employment rate were zero. The interaction coefficients θ_1 and θ_2 couple the two equations: θ_1 measures how strongly a higher wage share depresses employment growth, and θ_2 captures how much a tighter labor market accelerates wage-share growth. Without these cross-terms the system would reduce to two uncoupled exponential processes; it is precisely the vu interaction that generates the closed-orbit cycles for which the Goodwin model is known.

Symbol	**Description**
v	Employment rate
u	Wage share
σ	Capital–output ratio
α	Growth rate of labor productivity
β	Growth rate of the labor force
γ	Intercept of the real-wage Phillips curve
ρ	Slope of the real-wage Phillips curve
η_1, η_2	Composite trend parameters
θ_1, θ_2	Interaction (cross-term) coefficients

Table 1: Variables and parameters of the Goodwin model.

The system (2) possesses two critical points: the origin $(0, 0)$, which is a saddle, and the interior equilibrium $(\eta_2/\theta_2,\ \eta_1/\theta_1)$, which is a center. At the interior equilibrium the Jacobian has purely imaginary eigenvalues and a classical Lyapunov argument confirms that every trajectory in the positive quadrant is a closed orbit, not merely a spiral (Basu, 2024). Consequently the model produces perpetual, undamped oscillations whose amplitude is pinned down by initial conditions alone. This conservative structure, with no dissipation and no limit cycle, makes the Goodwin system a clean but demanding benchmark: any data-driven method must recover the vu interaction terms that sustain the closed orbits, while correctly returning zero coefficients on all other candidate terms.

2 Methodology: Applying SINDy

2.1 Overview of the SINDy Framework

This section provides a brief overview of the SINDy algorithm proposed by Brunton et al. (2016) and further elaborated in Brunton and Kutz (2022). The major argument provided by the authors of the SINDy algorithm is that, few elements of may be governing the dynamical system, and hence, the idea of sparsity is the key defining concept which underlines this algorithm. For further clarification, we start describing the abstract dynamical system given by the following equation:

$$\frac{d}{dt}\mathbf{x}(t) = \mathbf{f}(\mathbf{x}(t)) \tag{3}$$

In equation (3), $\boldsymbol{x}(t) \in \mathbb{R}^{n\times 1}$ is n-vector at time t and denotes state of a system at time t. $f(\cdot) : \mathbb{R}^{n\times 1} \to \mathbb{R}^{n\times 1}$ is function which describes equations of motions. The goal of the SINDy algorithm is to discover f from the given data. Given function f usually consists of few terms, sparse model identification, balances two goals: model complexity on the one hand, and avoids overfitting on the other hand(Brunton and Kutz, 2022).

The model discovery requires the data of the dependent and indepent variables. Specifically, time-series data can be represented by the following matrix:

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}(t_1) & \mathbf{x}(t_2) & \cdots & \mathbf{x}(t_m) \end{bmatrix}^T \tag{4}$$

Similarly, matrix of derivatives can be compactly represented in the following way:

$$\dot{\mathbf{X}} = \begin{bmatrix} \dot{\mathbf{x}}(t_1) & \dot{\mathbf{x}}(t_2) & \cdots & \dot{\mathbf{x}}(t_m) \end{bmatrix}^T \tag{5}$$

For noisy data, which is most prevalent in economics time-series data, total-variation regularized derivative is best candidate, which provides accurate numerical results. The next component of the algorithm is candidate function library. It is a matrix $\mathbf{\Theta}(\mathbf{X})$ of which columns represent candidate for the f, which the model aims to discover. The library matrix may have trigonometric functions, and polynomials of degree d, as express by $\mathbf{X}^d$ (Brunton and Kutz, 2022).

$$\mathbf{\Theta}(\mathbf{X}) = \begin{bmatrix} \mathbf{1} & \mathbf{X} & \mathbf{X}^2 & \cdots & \mathbf{X}^d & \cdots & \sin(\mathbf{X}) & \cos(\mathbf{X}) & \cdots \end{bmatrix} \tag{6}$$

Finally, the $\mathbf{\Xi}$ matrix is comprised of $\boldsymbol{\xi_k}$. The $\boldsymbol{\xi_k}$ is a vector coefficients and it identifies active terms in k-th rows. The goal is to have parsimonious model, which means the model should be precise and at the same time with few terms, meaning that $\mathbf{\Xi}$ should be sparse. Given all terms are explained, the SINDy model has following representation and the system dynamics is approximated by solving Equation (7):

$$\dot{\mathbf{X}} = \mathbf{\Theta}(\mathbf{X})\mathbf{\Xi} \tag{7}$$

The solution for the Equation (7) is obtained via optimization procedure which is usually presented in the following functional form:

$$\boldsymbol{\xi}_k = \underset{\boldsymbol{\xi}'_k}{\operatorname{argmin}} \left\| \dot{\mathbf{X}}_k - \mathbf{\Theta}(\mathbf{X})\boldsymbol{\xi}'_k \right\|_2 + \lambda \left\| \boldsymbol{\xi}'_k \right\|_1 \tag{8}$$

Equation (8) represents the least absolute shrinkage and selection operator (LASSO), where the dependent variable in Equation (7) corresponds to the time derivative of X. The LASSO formulation is an ℓ_1-regularized regression that promotes sparsity, with the regularization parameter λ controlling the strength of the penalty. A larger value of λ increases the shrinkage effect, driving more coefficients toward zero and setting some exactly to zero (James et al., 2023). As a result, the estimated dynamical system becomes increasingly sparse.

Require: $\Theta \in \mathbb{R}^{m \times L}$ ▷ Library of candidate functions
Require: $\dot{X} \in \mathbb{R}^{m \times n}$ ▷ Time-derivative data
Require: $\lambda > 0$ ▷ Sparsity threshold
Require: $K_{\max} \in \mathbb{N}$ ▷ Max iterations
Ensure: $\Xi \in \mathbb{R}^{L \times n}$ ▷ Sparse coefficient matrix
$\Xi \leftarrow \Theta^\dagger \dot{X}$ ▷ Initial least-squares fit
for $k = 1$ to $K_{\max}$ **do**
  for $j = 1$ to n **do** ▷ Loop over state variables
    $S \leftarrow \{\ell \mid |\Xi_{\ell j}| < \lambda\}$ ▷ Small coefficients
    for all $\ell \in S$ **do**
      $\Xi_{\ell j} \leftarrow 0$
    end for
    $B \leftarrow \{\ell \mid |\Xi_{\ell j}| \geq \lambda\}$ ▷ Retained terms
    if $B \neq \emptyset$ **then**
      $\Xi_{Bj} \leftarrow \arg\min_{\xi \in \mathbb{R}^{|B|}} \|\Theta_{:,B}\, \xi - \dot{X}_{:,j}\|_2^2$
    end if
  end for
end for
return Ξ

Algorithm 1: STLS algorithm, adapted from (Brunton and Kutz, 2022), translated from MATLAB into pseudocode.

Another method for optimization is the sequential thresholded least square (STLS), which is also preferred algorithm by the authors of the algorithm. Unlike LASSO, which relies on convex optimization via ℓ_1 regularization, the Algorithm 1 is a heuristic iterative method. It alternates between least-squares regression and hard thresholding: at each step, coefficients whose absolute values fall below a threshold λ are set to zero, and the model is re-estimated on the remaining terms. This process is repeated to promote sparsity in the solution.

3 Results: Model Recovery and Validation

3.1 Baseline Model

We begin by investigating the ability of the Sparse Identification of Nonlinear Dynamical Systems (SINDy) framework, using the SR3 sparse regression optimizer, to recover the governing equations of the Goodwin growth-cycle model. The Goodwin model represents endogenous distributional cycles arising from the interaction between the employment rate and the wage share. To evaluate the identification performance under controlled conditions, we generate synthetic trajectories from the analytical form of the Goodwin system and estimate a sparse model directly from the simulated data.

The synthetic data are generated using the parameter configuration reported in Table 2. The structural parameters correspond to the canonical multiplicative Goodwin dynamics, while sparsity is enforced through the SR3 optimizer using an ℓ_0 regularizer. Numerical derivatives are estimated using a Smoothed Finite Difference method with a Savitzky–Golay window length of 9.

Table 2: Simulation Setup for SR3-SINDy Identification of the Goodwin Model

Parameter	**Value**	**Description**
η_1	0.1399	Employment growth intercept
θ_1	0.2239	Effect of wage share on employment
η_2	0.2430	Wage share decay rate
θ_2	0.2570	Effect of employment on wage share
Initial condition v_0	0.9416	Initial employment rate
Initial condition u_0	0.6552	Initial wage share
SR3 regularization λ_{SR3}	10^{-4}	Sparsity weight for SR3 optimizer
Smoother window length	9	Window length for Smoothed Finite Difference
Time horizon T	100	Simulation horizon

Figure 2 shows the time-domain recovery of the Goodwin dynamics using SR3-SINDy. The employment rate v and wage share u exhibit stable cyclical behavior over $T = 100$, and the discovered model closely follows the observed trajectory.

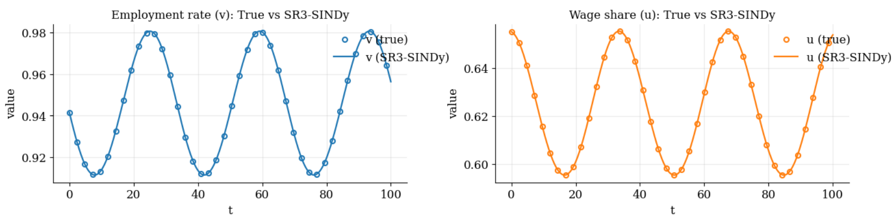

Figure 2: SR3-SINDy reconstruction of $v(t)$ and $u(t)$.

Figure 3 displays the phase portrait of the discovered system. The closed orbit structure is preserved, indicating accurate recovery of the nonlinear interaction between employment and wage share.

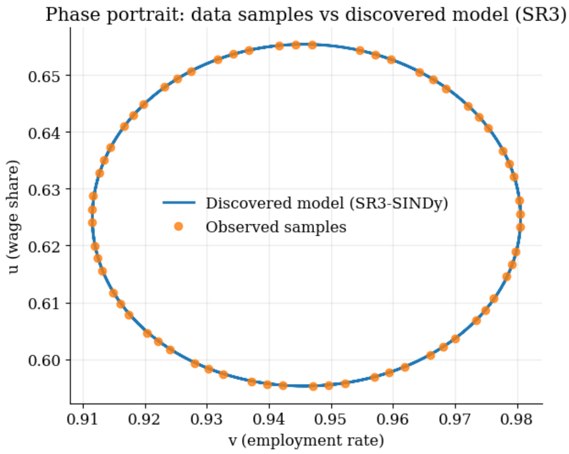

Figure 3: Phase portrait of the SR3-SINDy model.

Macroeconomic data is typically noisy, and this experiment investigates how the addition of noise affects the performance of the SINDy algorithm. Noise is introduced by adding independent Gaussian perturbations to the true solution of the Goodwin model, generating observed data as

$$\mathbf{y}_{\text{noisy}}(t) = \mathbf{y}_{\text{true}}(t) + \boldsymbol{\epsilon}(t),$$

where $\boldsymbol{\epsilon}(t) \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$. This additive noise degrades the signal-to-noise ratio, making the estimation of time derivatives more difficult and increasing the chance of incorrect model identification. By varying the noise level σ, we evaluate the robustness of SINDy to observational errors typically encountered in empirical macroeconomic applications.

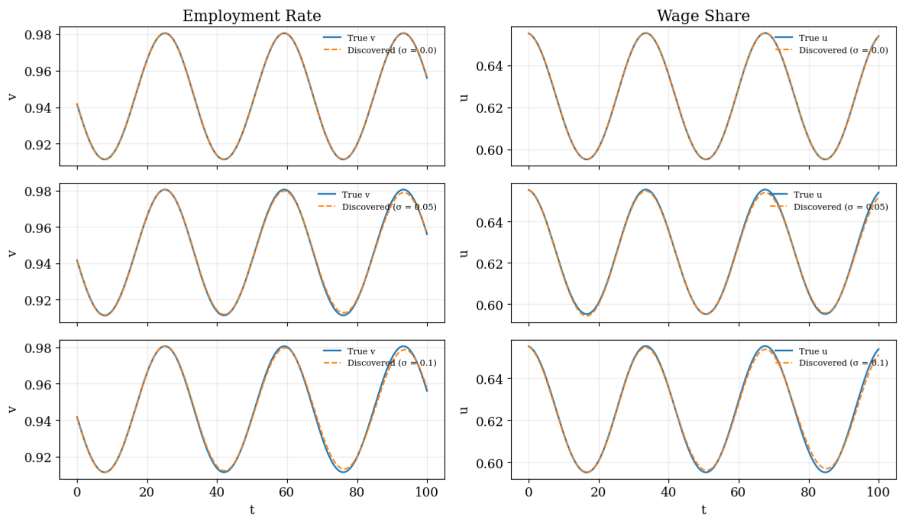

Figure 4: Impact of noise on SINDy model discovery.

This Figure 4 illustrates the impact of increasing noise levels on the ability of the SINDy algorithm to reconstruct the dynamics of the Goodwin model. The left column plots the employment rate $v(t)$, while the right column plots the wage share $u(t)$, both over time. Each row corresponds to a different noise standard deviation σ: 0.0, 0.05, and 0.1, respectively from top to bottom. For each subplot, the solid blue line represents the true trajectory obtained by numerically solving the original Goodwin system, whereas the dashed orange line shows the trajectory simulated from the SINDy-discovered model fitted to the noisy data. When $\sigma = 0.0$, SINDy accurately recovers the underlying dynamics. As noise increases, discrepancies between the true and discovered trajectories become more pronounced, highlighting the degradation in model recovery due to observational errors. The results emphasize the sensitivity of sparse model discovery methods to the quality of the input data, which is particularly relevant for empirical macroeconomic applications where measurement noise is unavoidable.

3.2 SINDy for Discrete Systems: Kaldorian Endogenous Business Cycle

As mentioned above, the key idea of endogenous cycles is that shocks do not originate from the external forces, but they stem from the in-built mechanisms. Kaldor's Business Cycle model belongs to the class of endogenous models. In this section, we clarify underlying concepts. Investment and saving has nonlinear response to saving and to income and output, this is the source of economic cycles. There two regimes: the regime of normal level income and the regime of high/low level of income. At normal level of income investment and savings are very responsive to change in income, and therefore, small change in income could lead large change investment and saving. On the other hand at the high/low level of income the responsivnes of investment/saving declines. This is the key mechanism, which underpins bussiness cycle: at normal levels of income strong sensitivity to income changes amplify economic fluctuaion, on the hand, at the low/high level of income fluctuations dampen. This leads to further claims stability of macroeconomic system: the economy is locally unstable when it is in regime 1 (normal level of income), and it is globally stable, because at low/high level of incomes response level decreases. In other words, after expansion/contraction, the economy reverts back to normalcy. In terms of economic meaning of this process, Kaldor in his 1940 paper provides following reasoning: at low level of economic activity, factories are underutilized. Even if there is sligtly increased economic activity, firms are not eager to build new capacity. In this condition firms already have the opportunity to respond with increasing capacity utilization, instead of making new investment. It means that, the responsivness of investment investment with income is low at low level of income. According to Kaldor's argument similar pattern arises at high level of output: specifically, it is related to increasing cost of borrowing, and also higher cost of construction, considering enrepreneous already have ongoing high capacity projects underway(Kaldor, 1940). While Kaldor's work has been reformulated mathematically by Chang and Smyth, 1971 (Chang and Smyth, 1971). In general form, Kaldors model can be represented in the following form as a system of two differential equations:

$$\begin{aligned} \dot{Y} &= \alpha\big(I(Y,K) - S(Y,K)\big), \\ \dot{K} &= I(Y,K). \end{aligned} \tag{9}$$

In (9), $\dot{Y}$ denotes the rate of change of income, and $\dot{K}$ denotes the rate of change of the capital stock. The functions $I(\cdot)$ and $S(\cdot)$ represent (possibly nonlinear) investment and saving, respectively, and $\alpha > 0$ is the adjustment (reaction) coefficient in the income equation (Semmler, 1986).

Since we are interested in how SINDy works in discrete-time economic systems, this section uses the Kaldor model version presented by Bischi et al. 2001. The model is defined by the following system of equations:

$$Y_t = Y_{t-1} + \alpha\left(I_{t-1} - S_{t-1}\right), \qquad \alpha > 0 \tag{10}$$

$$K_t = (1-\delta)K_{t-1} + I_{t-1}, \qquad \delta \in (0,1) \tag{11}$$

In equation (10), Y_t and K_t denote the income level and capital stock in period t, respectively. The parameter α is the reaction coefficient, also known as the speed of adjustment, which measures firms' reactions to excess demand conditions. If the value of α is small, the reaction to excess demand will be of a smaller magnitude. On the other hand, if α is greater than one, it indicates a more aggressive or rapid reaction. Essentially, the first part of equation (10) indicates that the current period's income, Y_t, is the previous period's income, Y_{t-1}, plus an adjustment.

This adjustment is based on the previous period's excess demand, $(I_{t-1} - S_{t-1})$, scaled by the speed of adjustment, α. Consequently, if investment exceeded savings in the previous period, income will tend to rise in the current period. If savings exceeded investment in the previous period, income will tend to fall. The magnitude of this rise or fall depends on the scaling factor α. The equation (11) represents capital stock dynamic equation, where δ denotes capital stock deprecation rate.

In the formalized Kaldor model, savings are proportional to the current level of income, and is described by following equation:

$$S_t = \sigma Y_t \tag{12}$$

In this equation (12), $\sigma \in (0, 1)$ denotes propensity to save. As noted by Bischi et al. (2001), linear specification of the saving function does not alter the nonlinearity of Kaldor model, and this nonliniearity is achieved via investment function specification which is given below. The investment function, as given in Bischi et al. (2001), follows following form:

$$I_t = \underbrace{\sigma Y^E}_{\text{normal savings}} + \underbrace{\gamma \left(\frac{\sigma Y^E}{\delta} - K_t \right)}_{\text{adjustment to normal capital stock}} + \underbrace{\arctan(Y_t - Y^E)}_{\text{bounded income response}} \tag{13}$$

The equation (13) consists of 3 different components. This first component, which can be called normal savings component is based on Keynesian assumption that expected income is equal to the normal level income μ. The σ parameter represents marginal propensity to save. The second component is adjustment to normal capital stock, where first element $\frac{\sigma\mu}{\delta}$ is normal capital stock level given in steady state where investment equals depreciation ($I_t = \delta K_t$), this equality gives $K^* = \frac{\sigma\mu}{\delta}$. $\gamma > 0$ is the adjustment coefficients, which measures how aggressively firms react to adjust their actual stock of capital K towards it "desired" (normal) level. Finally, the third term $\arctan(Y_t - Y^E)$ is nonlinear component of the income, which captures how deviations of actual income from its normal level influences over investment in short-run period. Investment is increasing function of actual income, which means that higher-than-normal level of income stimulates investment. The arctangent function ensures that the model follows Kaldor's intuition: investment is highly responsive to income near its normal level, but becomes less responsive at the extremes. Following Bischi et al. (2001), the two-dimensional nonlinear map $T : (Y_t, K_t) \to (Y_{t+1}, K_{t+1})$ can be presented in the following form:

$$\begin{aligned} Y_t &= Y_{t-1} + \alpha \left[\sigma Y^E + \gamma \left(\frac{\sigma Y^E}{\delta} - K_{t-1} \right) + \arctan\left(Y_{t-1} - Y^E\right) - \sigma Y_{t-1} \right] \\ K_t &= (1-\delta) K_{t-1} + \sigma Y^E + \gamma \left(\frac{\sigma Y^E}{\delta} - K_t \right) + \arctan\left(Y_{t-1} - Y^E\right) \end{aligned} \tag{14}$$

Below, we test whether the SINDy model is able to recover the governing equations of the system (14). The coefficients used for the simulation follow the parameterisation in (Kohler and Prante, 2025). Specifically, the parameters are given in Table 3.

Table 3: Parameter Values Used in the Simulation Model

Symbol	**Description**	**Value**
α	Adjustment speed of output	1.2
δ	Depreciation rate of capital	0.2
σ	Propensity to save	0.4
Y_E	Normal level of output	10
γ	Sensitivity of investment to capital gap	0.6

Firstly, we simulate the Kaldor model with the given set-up, and afterward we use SINDy, to recover governing equations of the Kaldor discrete nonlinear system. The results are given in the figure.

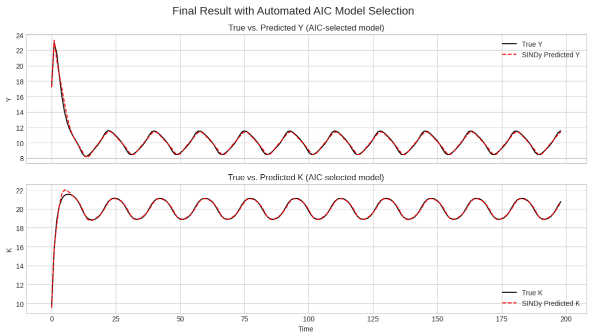

(a) Simulated vs. predicted dynamics of output (Y) and capital (K) using the SINDy model with AIC-selected sparsity.

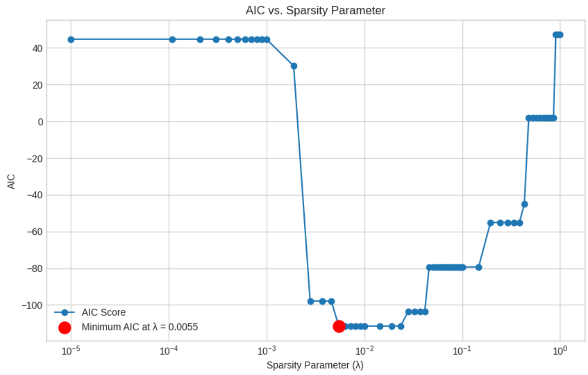

(b) AIC score as a function of the sparsity threshold λ. Minimum AIC at $\lambda = 0.0055$.

Figure 5: Final result of automated model selection and simulation using SINDy.

Figure 5a shows the simulated and predicted dynamics of a nonlinear Kaldor-type business cycle model using Sparse Identification of Nonlinear Dynamical Systems (SINDy). The system evolves in discrete time, where output (Y) and capital stock (K) follow an investment function composed of both polynomial terms and a shifted arctangent nonlinearity. To discover the governing equations, SINDy was applied to synthetic data generated from the model, using a feature library consisting of second-degree polynomials and a single custom nonlinear function $\arctan(Y - c)$.

To balance model accuracy with parsimony, an automated model selection procedure based on the Akaike Information Criterion (AIC) was implemented, as in Mangan et al. (2017). A wide range of sparsity thresholds λ was evaluated, with AIC scores computed using out-of-sample prediction errors on a held-out validation set. The minimum AIC was achieved at $\lambda = 0.0055$, as shown in Figure 5b, resulting in a compact yet expressive model. The final model selected by this procedure accurately reproduces the endogenous cyclical behavior in both Y and K, demonstrating that SINDy, when combined with information-theoretic model selection, can recover interpretable and structurally meaningful equations that capture the stylized dynamics of macroeconomic systems.

However, the SINDy model did not fully recover the underlying analytical structure of the simulated system. The true model equations, derived directly from the data-generating process, are given by:

$$\begin{aligned} Y_{k+1} &= 0.520\,Y_k - 0.720\,K_k + 1.200\,\arctan(Y_k - 10) + 19.200, \\ K_{k+1} &= 0.500\,K_k + 0.625\,\arctan(Y_k - 10) + 10.000. \end{aligned}$$

In contrast, the equations discovered by SINDy using the AIC-selected sparsity parameter are:

$$\begin{aligned} Y_{k+1} &= 15.6641 + 2.231\,Y - 0.604\,K - 0.049\,Y^2 + 7.229\,\arctan(Y - K) \\ K_{k+1} &= 8.1581 + 0.891\,Y + 0.561\,K - 0.025\,Y^2 + 3.765\,\arctan(Y - K). \end{aligned}$$

Although the predicted trajectories of output and capital closely match the ground truth, the recovered model introduces spurious terms—including Y^2 and an incorrect mixed nonlinear function—which deviate from the known analytical structure. This discrepancy highlights a key limitation: good predictive performance alone does not guarantee correct structural recovery. Therefore, domain-informed feature design remains essential when using SINDy for interpretable economic model discovery.

4 Universal Differential Equations (UDE)

4.1 Algorithm Description

The data-driven approach is not new in the field of science, and there are many different approximators that can approximate functions and thus learn models from data. One important distinction is that neural networks are universal approximators and, according to the Universal Approximation Theorem, they can approximate any function f as closely as desired, that is, within a specified criterion of closeness ε, with respect to a defined distance, provided they have sufficient capacity. However, unlike other approximators (e.g., Taylor series, Fourier approximation), neural networks are not subject to the "curse of dimensionality"(Rackauckas et al., 2021). On the other hand, neural networks are "data-hungry" algorithms and therefore require large datasets, which are frequently unavailable in the macroeconomic context. This creates a fundamental tension: the most flexible approximators demand precisely the volume of observations that macroeconomics cannot supply. Universal Differential Equations address this tension by embedding a neural network within a differential equation that already encodes partial mechanistic knowledge, thereby reducing the burden on the data-driven component and balancing the trade-off between purely data-driven and purely mechanistic models (Rackauckas, 2020; Rackauckas et al., 2021).

A **Universal Differential Equation (UDE)** is represented in the following functional form:

$$\frac{du}{dt} = f\left(u, t, U_\theta(u, t)\right) \tag{15}$$

where f is a deterministic function encoding mechanistic knowledge, and $U_\theta(u, t)$ is a universal approximator (e.g., a neural network) parameterized by θ, used to capture unknown dynamics(Rackauckas, 2020).

More compactly, for the ODE version, UDE can be expressed in the following form:

$$\frac{d\mathbf{x}}{dt} = \underbrace{f_M(t, \mathbf{x};\, \theta_M)}_{\text{mechanistic model}} + \underbrace{f_{\text{ANN}}(t, \mathbf{x};\, \theta_{\text{ANN}})}_{\text{data-driven correction}} \tag{16}$$

In this formulation, $\mathbf{x}(t) \in \mathbb{R}^n$ represents the state of the system at time t. The function f_M captures known dynamics derived from theoretical or empirical models, parameterized by θ_M. As mentioned above f_{ANN} is a flexible function approximator—typically a neural network—with parameters θ_{ANN}, used to for fitting for unknown part (Philipps et al., 2025).

In the context of Universal Differential Equations (UDEs), discovering the best solution means minimizing a cost function that measures the distance—according to a chosen metric—between the model-generated trajectory and the observed data. A general formulation can be written as $\rho\left(u_\theta(t_i), y(t_i)\right)$, where ρ denotes the metric used to compare the model prediction $u_\theta(t_i)$ with the observed data $y(t_i)$ at time t_i.

This leads to a cost function of the form

$$C(\theta) = \sum_i \|u_\theta(t_i) - y(t_i)\|_2^2 \quad \text{or} \quad C(\theta) = \sum_i \|u_\theta(t_i) - y(t_i)\|_2$$

depending on the choice of metric. Here, $\|\cdot\|_2$ denotes the Euclidean norm. The goal is to find parameters θ that minimize $C(\theta)$, thereby aligning the model trajectory as closely as possible with the observed data (Devgupta et al., 2024; Silvestri et al., 2023). For the particular optimization (minimizing cost with regards to parameters), various optimization algorithm can be used, including gradient decent or Adaptive Moment Estimation (ADAM) (Devgupta et al., 2024).

4.2 Testing UDE on Macroinformed Parameters

In this subsection, we look how UDE algorithm works on macroinformed parameter space. There are two interesting parts in this case: (i) how UDE could work considering limited data considering macroeconomic data limitation, and (i) how UDE works when the data noisy, which is aslo characteristic of macro variables. The implementation follows the hybrid UDE–SINDy tutorial provided in the SciML documentation (Rackauckas, 2025), adapted to the Goodwin system with macroeconomically calibrated parameters.

The dynamical system follows a Goodwin-type predator-prey structure with macroeconomically calibrated parameters. The governing equations take the form

$$\dot{v} = [\eta_1 - \theta_1 u]\, v \tag{17}$$

$$\dot{u} = [-\eta_2 + \theta_2 v]\, u \tag{18}$$

where v denotes the employment rate (prey) and u denotes the wage share (predator). The parameter η_1 represents the trend in employment growth when the profit share of income is unity, η_2 represents the decline in real wages at zero employment, and θ_1, θ_2 are the interaction coefficients capturing the influence of the wage share on the employment rate and the impact of employment on the wage share, respectively. The calibrated values are $\eta_1 = 0.1399$, $\theta_1 = 0.2239$, $\theta_2 = 0.2570$, $\eta_2 = 0.2430$, with initial conditions $v(0) = 0.94$, $u(0) = 0.65$. The system was integrated over $t \in [0, 10]$ using a seventh-order Verner method (Vern7) at tolerances of 10^{-12} (absolute and relative), with solution snapshots recorded at intervals of $\Delta t = 0.25$. Gaussian noise scaled to 0.1% of the state mean was added to the solution to simulate observational data.

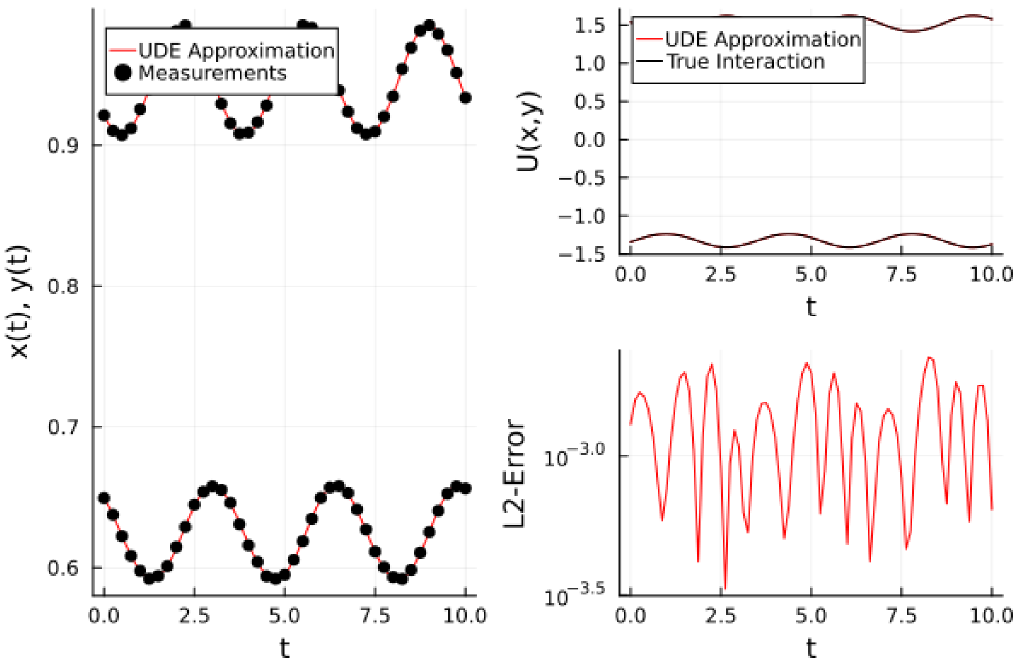

Figure 6: UDE training results. Left: approximated trajectories (red) versus noisy measurements (black dots) for v (top) and u (bottom). Upper right: neural network approximation (red) versus true interaction terms (black). Lower right: L_2 error over time.

4.3 Hybrid UDE–SINDy Pipeline: Synthetic Validation

The purpose of this section is to validate the hybrid UDE–SINDy pipeline on synthetic data with known ground truth before applying it to real macroeconomic observations. The key question is whether the method can recover the functional form and magnitude of the Goodwin interaction terms from noisy measurements alone, without prior specification of vu coupling.

The hybrid universal differential equation retains the linear decay and growth terms as known mechanistic structure while delegating the unknown nonlinear interactions to a neural network:

$$\dot{v} = \eta_1 \, v + \mathcal{N}_1(v, u;\, \theta), \tag{19}$$
$$\dot{u} = -\eta_2 \, u + \mathcal{N}_2(v, u;\, \theta), \tag{20}$$

where v and u denote the employment rate and wage share, respectively, η_1 and η_2 are fixed at their calibrated values, and $\mathcal{N}(\cdot;\, \theta) : \mathbb{R}^2 \to \mathbb{R}^2$ is a feedforward neural network parameterised by θ. If the Goodwin model is correct, the network should learn $\mathcal{N}_1 \approx -\theta_1 \, vu$ and $\mathcal{N}_2 \approx +\theta_2 \, vu$. Crucially, this functional form is not imposed: the network is free to approximate any smooth mapping from the state space to the derivative corrections.

The rationale for fixing η_1 and η_2 rather than treating them as learnable parameters stems from an identifiability constraint. When both the linear coefficients and the neural network weights are optimised simultaneously, the two components compete to explain the same variance in the data, leading to non-unique decompositions. By anchoring the mechanistic terms, we ensure that the network captures only the residual dynamics—precisely the interaction structure we seek to identify.

The network architecture consists of four fully connected layers ($2 \to 5 \to 5 \to 5 \to 2$) with Gaussian radial basis function activations $\sigma(x) = \exp(-x^2)$. This activation was chosen over standard alternatives (ReLU, tanh) because its localised support discourages the network from fitting spurious long-range extrapolations, a desirable property when the training data occupy a narrow region of state space(Rackauckas, 2025).

Training minimises the mean squared error between the numerically integrated UDE trajectory and the noisy observations. ODE sensitivities are computed via the continuous adjoint method with reverse-mode automatic differentiation, which scales favourably with the number of network parameters. Optimisation proceeds in two phases:

1. **Adam** (5 000 iterations, learning rate 0.05): provides rapid initial convergence from a random initialisation to the neighbourhood of the optimum;

2. **BFGS with backtracking line search** (2 000 iterations): refines the solution using second-order curvature information, achieving the precision needed for downstream symbolic regression.

Figure 6 summarises the training outcome. The left panels display the UDE-predicted trajectories for $v(t)$ and $u(t)$ overlaid on the noisy measurements, showing close agreement throughout the integration interval. The upper-right panel compares the neural network outputs $\mathcal{N}_1$ and $\mathcal{N}_2$ against the true interaction terms $-\theta_1 vu$ and $+\theta_2 vu$. The lower-right panel reports the pointwise L_2 reconstruction error, which remains in the range $10^{-3.5}$ to 10^{-3}—small enough to confirm that the network has faithfully captured the missing dynamics, yet large enough to reflect the presence of observation noise rather than overfitting.

The trained UDE is evaluated on a refined time grid ($\Delta t = 0.125$, yielding 81 equally spaced points) to produce smoothed state trajectories and the corresponding neural network outputs for both interaction terms. These serve as inputs to the sparse regression stage, implemented in the PySINDy framework (de Silva et al., 2020) and called from within the Julia environment via PyCall.

An immediate challenge arises from the narrow range of state-space exploration characteristic of macroeconomic variables: $v \in [0.88, 0.98]$ and $u \in [0.60, 0.65]$. Unlike ecological or chemical systems where state variables may span orders of magnitude, macroeconomic rates fluctuate within a few percentage points of their long-run means. This induces severe multicollinearity among polynomial basis functions—the condition number of a degree-3 library matrix exceeds 10^4—causing standard STLSQ regression with polynomial libraries of degree 2 and 3 to return dense, non-unique coefficient sets.

To address this collinearity problem while maintaining algorithmic consistency with the direct SINDy experiments, we employ the SR3 (Sparse Relaxed Regularised Regression) optimiser with ℓ_0 regularisation over a physically motivated candidate library $\{v,\ u,\ v^2,\ vu,\ u^2\}$, as implemented in PySINDy. The ℓ_0 penalty directly targets the number of active terms, promoting maximally sparse solutions without the coefficient shrinkage inherent in ℓ_1 approaches. Column normalisation is disabled (`normalize_columns=False`) to preserve the physical scaling of the basis functions, ensuring that the recovered coefficients are directly interpretable as interaction parameters rather than requiring post-hoc rescaling. The constant term is excluded from the library because the UDE's mechanistic component already accounts for the linear dynamics; any additive constant in the neural network residual would represent an unphysical offset.

The SR3 optimiser identifies the single-term model vu in both equations, with all other candidate terms receiving zero coefficients. The recovered interaction terms are:

$$\hat{\mathcal{N}_1} \approx -0.2241\, vu \qquad (\text{true:}\ \ -\theta_1 vu = -0.2239\, vu), \tag{21}$$

$$\hat{\mathcal{N}_2} \approx +0.2570\, vu \qquad (\text{true:}\ \ +\theta_2 vu = +0.2570\, vu). \tag{22}$$

The recovered values agree closely with the calibrated parameters: θ_2 is recovered to four significant figures (relative error $< 0.004\%$), while θ_1 is recovered to three significant figures (relative error $\approx 0.09\%$).

Figures 7 and 8 compare the SINDy-recovered interaction terms with both the neural network output and the true analytic expressions for $\mathcal{N}_1$ and $\mathcal{N}_2$, respectively. In both cases the three curves are visually indistinguishable, confirming that the sparse regression has successfully distilled the network's learned representation into the correct symbolic form.

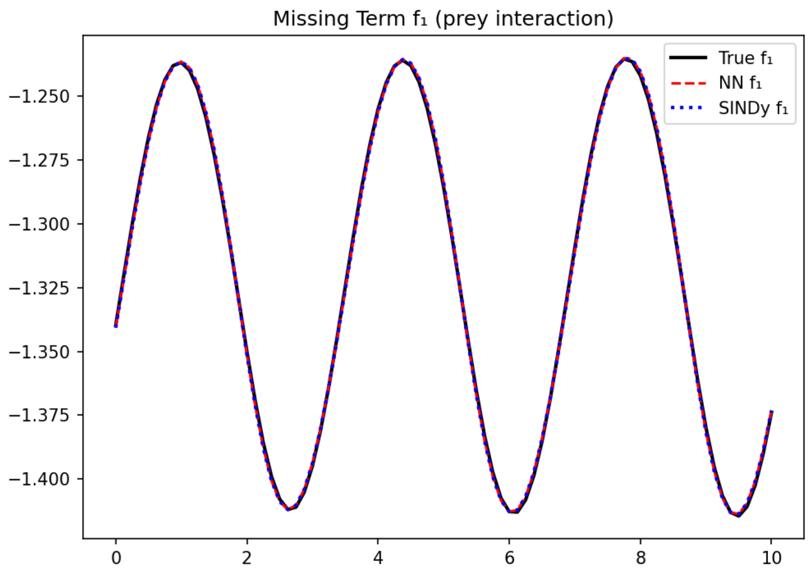

Figure 7: Recovered interaction term $\mathcal{N}_1$ in the employment rate equation. The true term $-\theta_1 vu$ (black solid), neural network approximation (red dashed), and SINDy-identified expression (blue dotted) overlap throughout the integration interval.

The full reconstructed system reads:

$$\dot{v} = 0.1399\, v - 0.2241\, vu, \tag{23}$$

$$\dot{u} = 0.2570\, vu - 0.2430\, u, \tag{24}$$

reproducing the original Goodwin predator–prey structure with all four parameters (η_1, θ_1, θ_2, η_2) recovered from noisy observations. No prior knowledge of the bilinear coupling was supplied

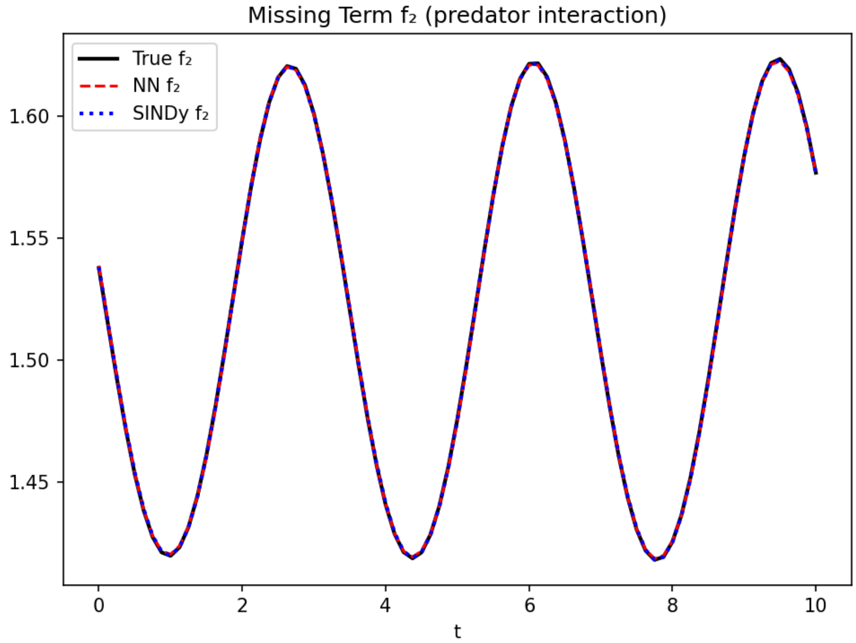

Figure 8: Recovered interaction term $\mathcal{N}_2$ in the wage share equation. The true term $+\theta_2 vu$ (black solid), neural network approximation (red dashed), and SINDy-identified expression (blue dotted) overlap throughout the integration interval.

to the pipeline: the vu interaction term was discovered entirely from the data, passing through the neural network as an intermediate, model-free representation before being crystallised into a parsimonious symbolic equation by sparse regression. Notably, the use of SR3 with ℓ_0 regularisation in the UDE–SINDy pipeline parallels the SR3-based sparse regression employed in the direct SINDy experiments, ensuring that the two methodological pathways share the same model-selection philosophy—differing only in the upstream data-processing stage (neural network smoothing versus direct derivative estimation) and implementation (PySINDy versus Julia's DataDrivenSparse). This validates the pipeline's capacity to detect Goodwin-type limit-cycle dynamics when they are genuinely present in the data, a necessary prerequisite for interpreting negative results on real macroeconomic time series in subsequent sections.

5 Robustness to Nonstationarity

A potential objection to any negative finding on real macroeconomic data is that the discovery methods may simply be unable to cope with the structural breaks and regime changes that characterise actual economies. To pre-empt this objection, we subject both SINDy and the hybrid UDE–SINDy pipeline to a controlled nonstationarity test: a piecewise Goodwin system whose parameters shift abruptly across regimes. If the methods can recover the underlying cyclical structure from such data, then their failure to do so on real observations cannot be attributed to nonstationarity alone.

5.1 Piecewise Goodwin System

The application of SINDy to nonstationary data has been addressed by Voina et al. (2024), who introduce Dynamic SINDy—a combination of SINDy with a variational autoencoder—to uncover the dynamics of non-autonomous systems in which the coefficient matrix becomes a function of time. The authors employ sigmoid functions and regime-switching mechanisms to capture gradual parameter evolution. While this approach is effective for smooth transitions, it is less suited for the type of nonstationarity characteristic of macroeconomic systems, where regime shifts are often abrupt rather than smooth. In such contexts, macroeconomic variables frequently experience sudden shocks followed by periods of instability. Only after passing through this instability window may the variables return—partially or fully—to stable trajectories. For example, in the context of Goodwin models, rather than a single orbit, the system may exhibit multiple orbits with different centres. However, transitions between these orbits are unlikely to be smooth, since they must pass through an instability window caused by abrupt exogenous shocks such as oil crises, wars, or pandemics.

To represent macroeconomic dynamics under changing structural conditions, the model is expressed as a *piecewise Goodwin system* defined over multiple time windows. The state vector

$$x(t) = [v(t),\, u(t)]^\top$$

captures the employment rate $v(t)$ and wage share $u(t)$. The system evolves across consecutive intervals

$$[t_0, t_1),\ [t_1, t_2),\ [t_2, t_3),\ [t_3, t_{\text{end}}),$$

each corresponding to a distinct economic regime. The overall dynamics are defined compactly as

$$\dot{x}(t) = \begin{cases} f_{\text{R1}}(x; \eta_1, \theta_1, \eta_2, \theta_2), & t \in [t_0, t_1), \\ f_{\text{R2}}(t, x; \eta_1(t), \eta_2(t)), & t \in [t_1, t_2), \\ f_{\text{Trans}}(x; x^\dagger, \varepsilon), & t \in [t_2, t_3), \\ f_{\text{R3}}(x; \eta_1', \theta_1', \eta_2', \theta_2'), & t \in [t_3, t_{\text{end}}). \end{cases}$$

The first regime (f_{R1}) represents the baseline autonomous Goodwin dynamics. The second (f_{R2}) introduces short-term non-autonomous fluctuations through small sinusoidal perturbations in $\eta_1(t)$ and $\eta_2(t)$, simulating exogenous shocks. The transition phase (f_{Trans}) includes mild damping toward a new equilibrium $x^\dagger$ to ensure smooth convergence between regimes. The third regime (f_{R3}) establishes a new steady state with shifted parameters reflecting a different macroeconomic configuration. This piecewise formulation allows for multiple regime windows, capturing abrupt shifts in macroeconomic behaviour—such as transitions from stable growth to disruption and subsequent adjustment—within a unified dynamic framework. The full parameter specification is given in Appendix A.1.

Both methods described below are fitted to the *entire* nonstationary trajectory, without any knowledge of the regime boundaries. This is the hardest setting: the algorithms must extract a coherent dynamical description from data that were generated by structurally different equations in different time windows.

5.2 SINDy and UDE+SINDy on Nonstationary Data

Both discovery pipelines—direct SINDy and the hybrid UDE–SINDy—are fitted to the *entire* nonstationary trajectory described in Section 5.1, without any knowledge of the regime boundaries. This is the hardest setting: each algorithm must extract a coherent dynamical description from data generated by structurally different equations in different time windows.

The SR3 optimiser with ℓ_0 regularisation is applied directly to the nonstationary trajectory, using PySINDy's built-in `SmoothedFiniteDifference` (Savitzky–Golay window length 9) for derivative estimation. The candidate library consists of five second-order polynomial terms $\{v,\ u,\ v^2,\ vu,\ u^2\}$, and sparsity is enforced through ℓ_0 penalisation with $\lambda = 10^{-3}$ and column normalisation enabled. The recovered global model is:

$$\dot{v} = 0.037\, v - 0.079\, vu, \tag{25}$$

$$\dot{u} = -0.053\, u + 0.107\, vu. \tag{26}$$

Despite spanning four distinct regimes—including a forced non-autonomous interval and an abrupt parameter shift—SINDy recovers a two-term predator–prey structure with regime-averaged coefficients that lie between the Regime 1 and Regime 3 calibrations (see Appendix for a comparison of three SINDy variants: STLSQ, SR3, and Ensemble-SINDy).

The hybrid pipeline retains the same structural decomposition as in the stationary case (Equations 19–20), with the linear parameters η_1 and η_2 fixed at their Regime 1 values (0.05), while a neural network $\mathcal{N}(\cdot;\theta) : \mathbb{R}^2 \to \mathbb{R}^2$ captures the residual dynamics. The network architecture is widened to $2 \to 15 \to 15 \to 15 \to 2$ with Gaussian RBF activations to accommodate the richer nonstationary dynamics. Training uses a multiple-shooting loss with 50-point windows, and optimisation proceeds in three phases: Adam(0.02) for 5 000 iterations, Adam(0.005) for 5 000 iterations, and BFGS for 2 000 iterations.

After training, the neural network outputs are evaluated at the true states and passed to the SR3 optimiser (ℓ_0, λ selected by sweep, column normalisation disabled). The recovered interaction terms are:

$$\hat{\mathcal{N}_1} \approx -0.106\, vu \qquad (\text{true R1:}\ -0.10\, vu;\ \text{R3:}\ -0.12\, vu), \tag{27}$$

$$\hat{\mathcal{N}_2} \approx +0.096\, vu \qquad (\text{true:}\ +0.10\, vu). \tag{28}$$

The SR3 optimiser correctly identifies the bilinear interaction vu as the sole active term in both equations, with all other candidates (v, u, v^2, u^2) receiving zero coefficients. The recovered values represent an average over regimes rather than a regime-specific identification—an expected consequence of fitting a single model to a piecewise-generated trajectory.

Figure 9 compares the two approaches. Both methods recover a single closed orbit in phase space (centre column), correctly identifying the Goodwin predator–prey structure, while the true system exhibits two distinct orbits separated by an instability region. The time-series panels (left column) show that both discovered systems track the true trajectory in Regime 1 but drift during Regime 2 and Regime 3—an inevitable consequence of approximating a nonstationary system with a single autonomous model. The derivative-residual panels (right column) reveal comparable error magnitudes, with direct SINDy exhibiting higher residuals during the forced regime (R2) and UDE+SINDy producing a smoother error profile.

The central conclusion is that *nonstationarity does not prevent either method from recovering Goodwin-type dynamics when such dynamics are present in the data.* Both pipelines independently identify the bilinear vu interaction structure across regime changes, reinforcing the interpretive strength of subsequent experiments on real macroeconomic data: if neither method finds Goodwin structure in FRED capacity utilisation and labour share series, the explanation cannot be that the methods are unable to handle nonstationarity—it is that the structure is absent from the data.

6 Empirical Application to U.S. Macroeconomic Data

6.1 Data Analysis

The empirical analysis uses three quarterly, seasonally adjusted time series for the United States. The employment rate is constructed as $v = 1 - \text{UNRATE}/100$, where UNRATE is the civilian unemployment rate sourced from the Federal Reserve Economic Data (FRED) database. Capacity utilization is measured by the Total Capacity Utilization Index (series TCU), also obtained from FRED. The wage share u is the nonfarm business sector labor share reported by the Bureau of Labor Statistics (BLS) in the Productivity and Costs release, expressed as a fraction. The sample spans 1948:Q1–2025:Q3 for the employment rate and wage share ($N = 311$ observations), while capacity utilization is available from 1967:Q1 onward ($N = 235$). Following Barbosa-Filho and Taylor (2006), we consider two empirical specifications of the Goodwin sys-

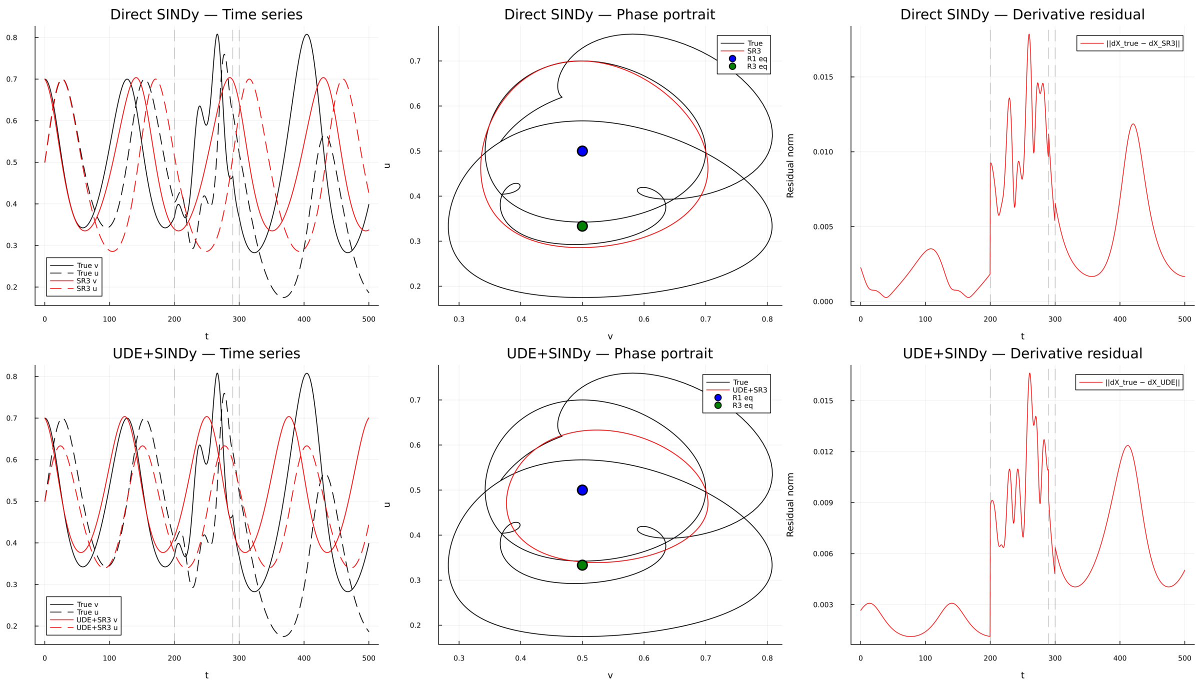

Figure 9: Comparison of direct SINDy (top row) and UDE+SINDy (bottom row) on the nonstationary Goodwin trajectory. Left: time-series overlay of true (black) and discovered (red) dynamics. Centre: phase portrait showing two true orbits (black) and the single recovered orbit (red), with Regime 1 and Regime 3 equilibria marked. Right: pointwise derivative-residual norm $\|\dot{x}_{\text{true}} - \dot{x}_{\text{model}}\|$ over time; vertical dashed lines indicate regime boundaries.

tem: the classical formulation in the (v, u) plane and the structuralist variant in the (TCU, u) plane.

Figures 10 and 11 provide a preliminary visual assessment of the data. The upper panel of Figure 10 plots the employment rate v and capacity utilization (TCU) together; the two series co-move closely, declining during recessions and recovering in tandem. The lower panel displays the wage share u, which fluctuates in the range of 0.62–0.65 from the late 1940s through the early 2000s before undergoing a secular decline to approximately 0.54–0.57 in the most recent decade. Figure 11 decomposes the phase-space trajectories by decade for both the classical (v, u) and the structuralist (TCU, u) specifications. In the earlier decades, the trajectories are concentrated in the upper portion of the wage-share axis, where the post-Keynesian literature has identified counterclockwise cycling consistent with the Goodwin pattern (Barbosa-Filho and Taylor, 2006; Setterfield, 2023). However, the trajectories are noisy and do not trace out the clean closed orbits predicted by the Lotka–Volterra structure. From the 2000s onward, the system migrates to a lower region of the state space, reflecting the secular erosion of the labor share that Setterfield (2023) attributes to the consolidation of neoliberal labor market institutions. The visual evidence alone casts doubt on whether a stable two-dimensional dynamical structure governs these series, a question we now address with formal equation-discovery methods.

6.2 Equation Discovery on Real Data

We apply both pipelines—direct SINDy and the hybrid UDE–SINDy—to the real macroeconomic series described above, using the same candidate library $\{x_1,\ x_2,\ x_1^2,\ x_1x_2,\ x_2^2\}$ and SR3 optimiser with ℓ_0 regularisation. If the Goodwin predator–prey structure governs these data, the interaction term x_1x_2 should emerge as the dominant or sole active term at some regularisation strength.

Derivatives are estimated via smoothed finite differences (Savitzky–Golay, window length 9) and SR3 is applied with column normalisation enabled. Sweeping λ over thirteen values from 10^{-6} to 1.0, the number of active terms per equation jumps directly from five to zero with no intermediate sparsity level; the x_1x_2 interaction never isolates as a dominant term.

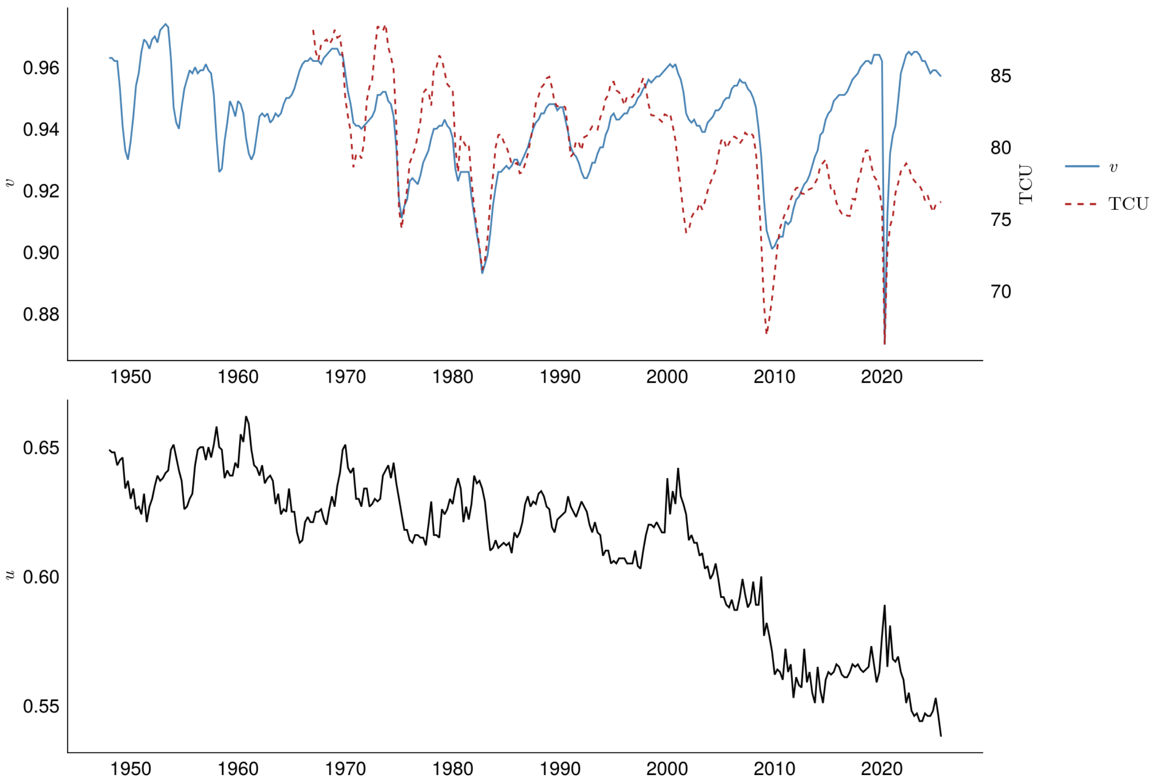

Figure 10: Time series of key variables for the U.S. economy, 1948:Q1–2025:Q3. Upper panel: employment rate v (left axis) and capacity utilization TCU (right axis, available from 1967:Q1). Lower panel: wage share u.

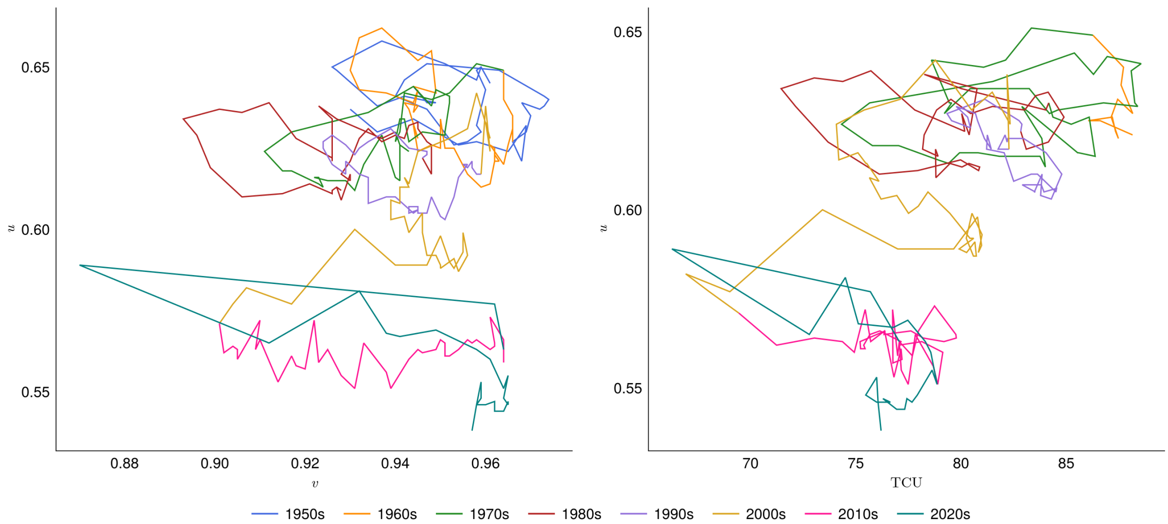

Figure 11: Phase-space trajectories by decade. Left panel: classical specification (v, u). Right panel: structuralist specification (TCU, u). Each decade segment includes the last observation of the preceding decade to ensure visual continuity.

At the selected regularisation strength, the structuralist specification yields a five-term and a three-term equation:

$$\dot{x}_1 = -0.956\,x_1 + 1.339\,x_2 - 0.222\,x_1^2 + 2.168\,x_1x_2 - 2.601\,x_2^2, \tag{29}$$

$$\dot{x}_2 = -0.239\,x_1^2 + 0.661\,x_1x_2 - 0.454\,x_2^2, \tag{30}$$

where $x_1 = \mathrm{TCU}/100$ and $x_2 = u$. Although x_1x_2 appears in both equations, it is accompanied by several other terms—far from the parsimonious Goodwin structure. Figure 12 shows the resulting dynamics. The time-series panel reveals that the SR3 model captures the general level and slow trends of both variables over the first two decades before drifting. The phase portrait traces a large loop around the observed data region, but the trajectory does not close and is quantitatively inconsistent with the actual state-space dynamics. The classical specification (v, u) produces a similarly dense polynomial whose rollout is globally unstable, diverging from the narrow data region almost immediately.

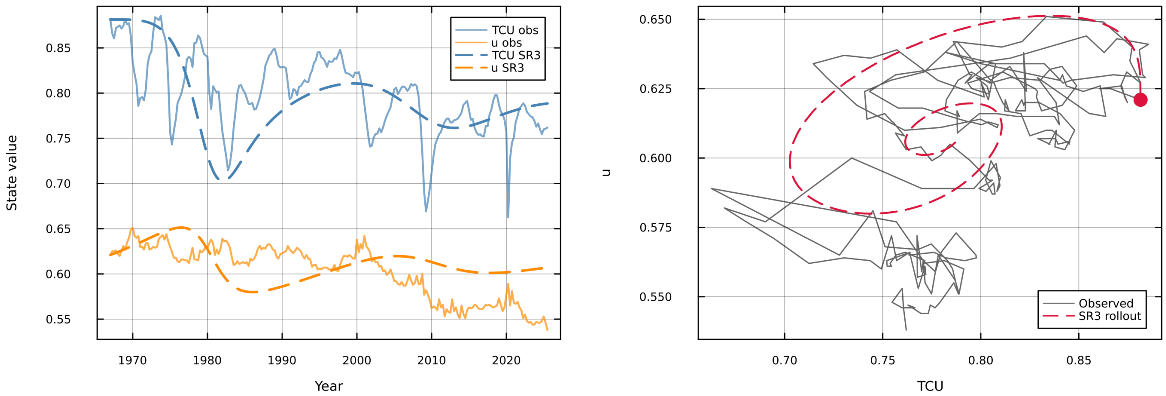

Figure 12: Direct SR3 results for the structuralist specification (TCU, u). Left: time-series overlay of observed data (solid) and SR3 rollout (dashed). Right: phase portrait showing the observed trajectory (black) and the SR3 rollout (red dashed), which traces a closed orbit substantially larger than the data region. The filled circle marks the initial condition.

The hybrid pipeline retains the structural decomposition $\dot{x}_i = \hat{\eta}_i\,x_i + \mathcal{N}_i(x_1, x_2;\,\theta)$, with the linear coefficients estimated from the data via OLS. The estimated values are effectively zero ($|\hat{\eta}| < 10^{-3}$), reflecting the near-stationary character of quarterly macroeconomic data. For the classical specification, SR3 finds no active terms in either equation: the neural network outputs are near-zero, indicating no discoverable nonlinear structure. For the structuralist specification, SR3 identifies linear and quadratic corrections ($-0.163\,x_1 + 0.333\,x_2 - 0.193\,x_2^2$ in the first equation, nothing in the second)—adjustments to individual state variables, not the bilinear x_1x_2 coupling predicted by the Goodwin model.

Neither method finds Goodwin structure at any regularisation strength or in either specification. Direct SINDy produces dense polynomials where x_1x_2 never isolates; the hybrid pipeline finds either nothing or linear corrections unrelated to predator–prey coupling. The failure is symmetric across both methods and both specifications, ruling out a methodological artefact.

7 Conclusion

This paper applies scientific machine learning methods (SINDy, universal differential equations, and their hybrid pipeline) to macroeconomic model discovery, using the Goodwin growth cycle as a testing ground. The primary contribution is methodological: these equation-discovery tools are surveyed, implemented, and benchmarked on problems calibrated to macroeconomic scales, where data are scarce, noisy, and nonstationary.

On stationary synthetic data, the hybrid UDE–SINDy pipeline recovers the Goodwin interaction coefficients to three and four significant figures, with the bilinear vu term as the sole active component of the sparse representation. On nonstationary synthetic data, consisting of a four-regime piecewise system with forced oscillations, abrupt parameter shifts, and transition damping, both direct SINDy and UDE–SINDy recover the predator–prey structure with regime-averaged coefficients that correctly identify vu as the dominant coupling. These experiments establish that the methods can operate under conditions characteristic of macroeconomic data: short time series, quarterly frequency, moderate noise, and structural breaks.

The application to 77 years of U.S. quarterly data yields a negative result on the Goodwin model. Direct SINDy returns dense polynomial models where the vu interaction carries no privileged role at any regularisation strength. The UDE–SINDy pipeline finds near-zero neural network corrections with no bilinear coupling. The result is consistent across both methods and both empirical specifications, classical (v, u) and structuralist (TCU, u), indicating that it does not reflect a methodological artefact. The bilinear predator–prey structure that defines the Goodwin model is not present in the data.

Several caveats apply. The analysis treats the Goodwin system as a continuous-time autonomous ODE, whereas quarterly FRED data are discrete, filtered, and subject to measurement conventions that may obscure continuous-time structure. The candidate library is restricted to second-order polynomials; richer libraries could reveal alternative dynamical structures not expressible within the current framework. The negative result applies to the two-dimensional system in the (v, u) and (TCU, u) planes and does not rule out distributional cycling mediated by additional state variables or operating at different time scales.

The methods developed here are not specific to the Goodwin model and can be applied to other dynamical systems for which candidate functional forms can be specified. Potential applications include testing Kaldorian nonlinear accelerator dynamics, Minsky-type financial instability cycles, or demand-regime switching models against data using the same pipeline.

References

Araujo, R. A., Dávila-Fernández, M. J., and Moreira, H. N. (2019). Some new insights on the empirics of Goodwin's growth-cycle model. *Structural Change and Economic Dynamics*, 51:42–54.

Barbosa-Filho, N. H. and Taylor, L. (2006). Distributive And Demand Cycles In The Us Economy—A Structuralist Goodwin Model. *Metroeconomica*, 57(3):389–411.

Basu, D. (2024). A note on the theory and empirics of the Goodwin model. Working Paper 2024-7, Working Paper.

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.

Blecker, R. A. and Setterfield, M. (2019). *Heterodox Macroeconomics: Models of Demand, Distribution and Growth*. Edward Elgar Publishing.

Box, G. E. P. (1979). Some Problems of Statistics and Everyday Life. *Journal of the American Statistical Association*, 74(365):1–4.

Brunton, S. L. and Kutz, J. N. (2022). Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control. ISBN: 9781009089517.

Brunton, S. L., Proctor, J. L., and Kutz, J. N. (2016). Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937.

Chang, W. W. and Smyth, D. J. (1971). The Existence and Persistence of Cycles in a Non-linear Model: Kaldor's 1940 Model Re-examined. *The Review of Economic Studies*, 38(1):37–44.

Chaudhuri, A. (2021). B-Splines. arXiv:2108.06617 [cs].

Corbetta, M. (2020). Application of sparse identification of nonlinear dynamics for physics-informed learning. In *2020 IEEE Aerospace Conference*, pages 1–8.

de Rooij, M., van Riel, N. A. W., and O'Donovan, S. D. (2025). Conditional universal differential equations capture population dynamics and interindividual variation in c-peptide production. *npj Systems Biology and Applications*, 11(1):84.

de Silva, B. M., Champion, K., Quade, M., Loiseau, J.-C., Kutz, J. N., and Brunton, S. L. (2020). PySINDy: A Python package for the Sparse Identification of Nonlinear Dynamics from Data.

Desai, M. and Ormerod, P. (1998). Richard Goodwin: A Short Appreciation. *The Economic Journal*, 108(450):1431–1435.

Devgupta, R., Dandekar, R. A., Dandekar, R., and Panat, S. (2024). Scientific machine learning in ecological systems: A study on the predator-prey dynamics. arXiv:2411.06858 [cs].

Fasel, U., Kaiser, E., Kutz, J. N., Brunton, B. W., and Brunton, S. L. (2021). SINDy with Control: A Tutorial. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 16–21.

Fasel, U., Kutz, J. N., Brunton, B. W., and Brunton, S. L. (2022). Ensemble-SINDy: Robust sparse model discovery in the low-data, high-noise limit, with active learning and control. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 478(2260):20210904. arXiv:2111.10992 [math].

Grasselli, M. R. and Maheshwari, A. (2018). Testing a Goodwin model with general capital accumulation rate. *Metroeconomica*, 69(3):619–643. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/meca.12204.

Harvie, D. (2000). Testing Goodwin: growth cycles in ten OECD countries. *Cambridge Journal of Economics*, 24(3):349–376.

James, G., Witten, D., Hastie, T., Tibshirani, R., and Taylor, J. (2023). *An Introduction to Statistical Learning: with Applications in Python*. Springer Nature. Google-Books-ID: ygzJEAAAQBAJ.

Kaldor, N. (1940). A Model of the Trade Cycle. *The Economic Journal*, 50(197):78–92.

Kaptanoglu, A. A. (2021). Promoting global stability in data-driven models of quadratic nonlinear dynamics. *Physical Review Fluids*, 6(9).

Koenig, B. C., Kim, S., and Deng, S. (2024). KAN-ODEs: Kolmogorov–Arnold network ordinary differential equations for learning dynamical systems and hidden physics. *Computer Methods in Applied Mechanics and Engineering*, 432:117397.

Kohler, K. and Prante, F. (2025). DIY Macroeconomic Model Simulation.

Li, J., Li, K., and Lin, Z. (2025). Sparse identification of nonlinear economic dynamical model. *Physica A: Statistical Mechanics and its Applications*, 675:130861.

Li, Z. (2024). Kolmogorov-Arnold Networks are Radial Basis Function Networks. arXiv:2405.06721 [cs].

Liu, Z., Wang, Y., Vaidya, S., Ruehle, F., Halverson, J., Soljačić, M., Hou, T. Y., and Tegmark, M. (2025). KAN: Kolmogorov-Arnold Networks. arXiv:2404.19756 [cs].

Mangan, N. M., Brunton, S. L., Proctor, J. L., and Kutz, J. N. (2016). Inferring biological networks by sparse identification of nonlinear dynamics. arXiv:1605.08368 [math].

Mangan, N. M., Kutz, J. N., Brunton, S. L., and Proctor, J. L. (2017). Model selection for dynamical systems via sparse regression and information criteria. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473(2204):20170009.

Philipps, M., Schmid, N., and Hasenauer, J. (2025). Current state and open problems in universal differential equations for systems biology. *npj Systems Biology and Applications*, 11(1):101.

Rackauckas, C. (2020). Universal Differential Equations for Scientific Machine Learning. Publisher: figshare.

Rackauckas, C. (2025). Automatically Discover Missing Physics by Embedding Machine Learning into Differential Equations · Overview of Julia's SciML.

Rackauckas, C., Ma, Y., Martensen, J., Warner, C., Zubov, K., Supekar, R., Skinner, D., Ramadhan, A., and Edelman, A. (2021). Universal Differential Equations for Scientific Machine Learning. arXiv:2001.04385 [cs].

Rooij, M. d., Erdős, B., Riel, N. A. W. v., and O'Donovan, S. D. (2025). Physiology-informed regularisation enables training of universal differential equation systems for biological applications. *PLOS Computational Biology*, 21(1):e1012198.

Schmid, N., Fernandes del Pozo, D., Waegeman, W., and Hasenauer, J. (2025). Assessment of uncertainty quantification in universal differential equations. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 383(2293):20240444.

Schmidt-Hieber, J. (2021). The Kolmogorov–Arnold representation theorem revisited. *Neural Networks*, 137:119–126.

Semmler, W. (1986). On nonlinear theories of economic cycles and the persistence of business cycles. *Mathematical Social Sciences*, 12(1):47–76.

Setterfield, M. (2023). Whatever Happened to the 'Goodwin Pattern'? Profit Squeeze Dynamics in the Modern American Labour Market. *Review of Political Economy*, 35(1):263–286. _eprint: https://doi.org/10.1080/09538259.2021.1921357.

Silvestri, M., Baldo, F., Misino, E., and Lombardi, M. (2023). An analysis of Universal Differential Equations for data-driven discovery of Ordinary Differential Equations. arXiv:2306.10335 [cs].

Vaca-Rubio, C. J., Blanco, L., Pereira, R., and Caus, M. (2024). Kolmogorov-Arnold Networks (KANs) for Time Series Analysis. arXiv:2405.08790 [eess].

Voina, D., Brunton, S., and Kutz, J. N. (2024). Deep Generative Modeling for Identification of Noisy, Non-Stationary Dynamical Systems. arXiv:2410.02079 [cs].

A Appendix

A.1 Appendix: Piecewise-Regime Specification of the Goodwin Model

The state vector is defined as

$$x(t) = \begin{bmatrix} v(t) \\ u(t) \end{bmatrix}, \qquad x(t_0) = \begin{bmatrix} 0.7 \\ 0.5 \end{bmatrix},$$

where $v(t)$ denotes the employment rate and $u(t)$ the wage share. The time domain is partitioned into consecutive intervals

$$[t_0, t_1),\ [t_1, t_2),\ [t_2, t_3),\ [t_3, t_{\text{end}}),$$

with $t_0 = 0$, $t_1 = 200$, $t_2 = 290$, $t_3 = 300$, and $t_{\text{end}} = 500$. Each interval corresponds to a distinct regime of macroeconomic behavior, governed by a particular parameter configuration.

The general Goodwin dynamics are given by

$$\dot{v}(t) = (\eta_1 - \theta_1 u(t))\, v(t), \qquad \dot{u}(t) = (-\eta_2 + \theta_2 v(t))\, u(t),$$

where η_i and θ_i represent growth and distributional response parameters. The model proceeds through four sequential regimes. In the first regime, parameters are constant, $(\eta_1, \theta_1, \eta_2, \theta_2) = (0.05, 0.10, 0.05, 0.10)$, yielding equilibrium $(v^*, u^*) = (0.5, 0.5)$. The second regime introduces short-term shocks by allowing

$$\eta_1(t) = \eta_1 + a \sin(\omega t), \qquad \eta_2(t) = \eta_2 + a \cos(\omega t),$$

with $a = 0.02$ and $\omega = 0.2$, creating mild non-autonomous oscillations that mimic exogenous disturbances. During the transition window $[t_2, t_3)$, a small damping term steers the system toward the next equilibrium $x^\dagger = (v^\dagger, u^\dagger)$ according to

$$\dot{v}(t) = (\eta_1 - \theta_1 u(t))\, v(t) - \varepsilon(v(t) - v^\dagger), \qquad \dot{u}(t) = (-\eta_2 + \theta_2 v(t))\, u(t) - \varepsilon(u(t) - u^\dagger),$$

where $\varepsilon = 0.02$. In the final regime, parameters shift to $(\eta_1, \theta_1, \eta_2, \theta_2) = (0.04, 0.12, 0.05, 0.10)$, yielding a new steady state $(v^*, u^*) = (0.5, 1/3)$.

The full system can thus be expressed compactly as

$$\dot{x}(t) = \begin{cases} f_{\text{R1}}(x; \eta_1, \theta_1, \eta_2, \theta_2), & t \in [t_0, t_1), \\ f_{\text{R2}}(t, x; \eta_1(t), \eta_2(t)), & t \in [t_1, t_2), \\ f_{\text{Trans}}(x; x^\dagger, \varepsilon), & t \in [t_2, t_3), \\ f_{\text{R3}}(x; \eta_1', \theta_1', \eta_2', \theta_2'), & t \in [t_3, t_{\text{end}}). \end{cases}$$

This piecewise-regime structure allows the Goodwin model to capture both smooth and abrupt macroeconomic transitions—such as shocks, instability windows, and post-shock adjustments—within a unified dynamic framework.

A.2 Appendix: Comparison of Discovered Coefficients Across SINDy Variants

Table 4 summarizes the discovered equations for the three implementations of the Sparse Identification of Nonlinear Dynamics (SINDy) method: SINDy with STLSQ, SINDy with SR3, and Ensemble-SINDy. The results are compared to the ground truth derived from the piecewise Goodwin system described in Section 2.

The ground truth Goodwin dynamics across regimes are defined by

$$\dot{v} = (\eta_1 - \theta_1 u)v, \qquad \dot{u} = (-\eta_2 + \theta_2 v)u,$$

Model	Recovered Equations
SINDy (STLSQ)	$\begin{cases} v' = 0.045v - 0.007u - 0.023v^2 - 0.045vu - 0.016u^2, \\ u' = -0.013v - 0.037u + 0.010v^2 + 0.115vu - 0.024u^2. \end{cases}$
Ensemble-SINDy	$\begin{cases} v' = 0.045v - 0.008u - 0.022v^2 - 0.045vu - 0.017u^2, \\ u' = -0.013v - 0.038u + 0.010v^2 + 0.115vu - 0.024u^2. \end{cases}$
SINDy (SR3)	$\begin{cases} v' = 0.037v - 0.079vu, \\ u' = -0.053u + 0.107vu. \end{cases}$

Table 4: Recovered coefficients for three SINDy variants.

with parameters $(\eta_1, \theta_1, \eta_2, \theta_2) = (0.05, 0.10, 0.05, 0.10)$ for Regime 1 and $(0.04, 0.12, 0.05, 0.10)$ for Regime 3.

The coefficients recovered by SINDy (STLSQ) and Ensemble-SINDy closely align with these reference values, both reproducing the expected linear interaction terms and moderate nonlinear corrections in v^2, vu, and u^2. Ensemble averaging slightly stabilizes the coefficients and reduces the variance in small nonlinear terms. The SR3 variant produces a sparser representation, identifying only the dominant cross-term vu and single-variable damping components. While SR3 sacrifices some detail, it effectively captures the core cyclical feedback between employment and wage share that characterizes the Goodwin model. Overall, all three methods reproduce the system's endogenous oscillatory structure, with SR3 providing the most parsimonious yet dynamically consistent form.

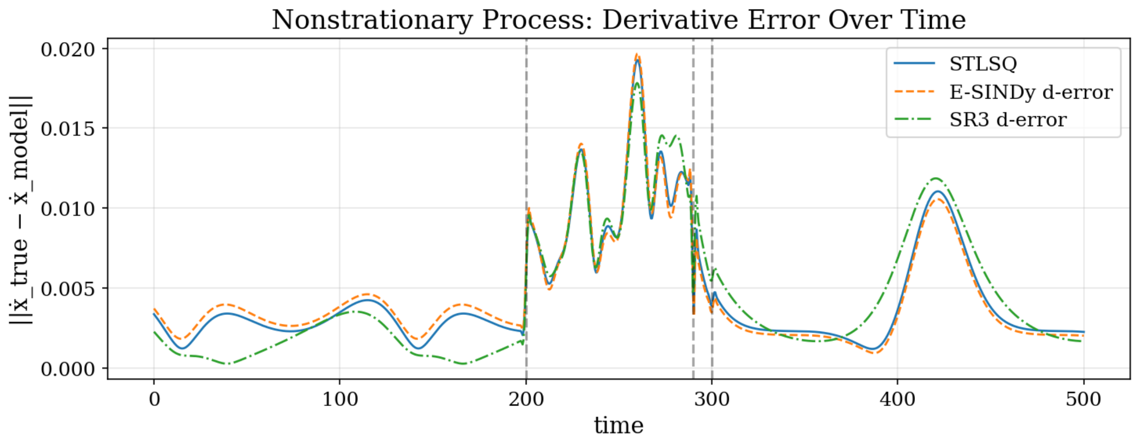

Figure 13: Nonstationary Process: Derivative Error Over Time. Comparison of derivative errors across three methods: STLSQ, E-SINDy, and SR3. Vertical dashed lines denote regime transition points.

Figure 13 compares the derivative errors over time for STLSQ, SR3, and E-SINDy. As shown in the diagram, the SINDy model with the SR3 optimizer achieves the lowest derivative errors among the compared variants.

A.3 Appendix: KAN-ODEs for Discovering Dynamical Systems

This appendix provides a theoretical overview of Kolmogorov–Arnold Network ODEs (KAN-ODEs), the third model discovery framework reviewed in this paper. While the main text focuses on SINDy and UDE+SINDy, which carry the nonstationarity robustness analysis and the real-data falsification exercise, KAN-ODEs represent an emerging alternative that we include for completeness.

There are two different theorems guiding Multilayer Perceptrons (MLPs) and Kolmogorov–Arnold Networks (KANs). While MLPs are based on the universal approximation theorem, as discussed in the section on Universal Differential Equations (UDEs), KANs are founded on the Kolmogorov–Arnold representation theorem (Schmidt-Hieber, 2021). This theorem, developed by Soviet mathematicians Andrey Kolmogorov and Vladimir Arnold, shows that any multivariate function of the form

$$f(\mathbf{x}) : [0,1]^n \to \mathbb{R}$$

can be expressed as a finite composition (finite superposition) of continuous univariate functions:

$$f(\mathbf{x}) = f(x_1, \ldots, x_n) = \sum_{q=1}^{2n+1} \Phi_q \left(\sum_{p=1}^{n} \phi_{q,p}(x_p) \right). \tag{31}$$

In equation (31), the functions $\phi_{q,p} : [0,1] \to \mathbb{R}$ (inner function) are univariate, and each $\Phi_q : \mathbb{R} \to \mathbb{R}$ (outer function) reconstructs the target function f from these univariate components. This implies that if the univariate functions are smooth, then multivariate functions can be learned by combining learnable one-dimensional functions. This property is particularly practical in macroeconomic modeling, for the models where some functional relationships are assumed to be smooth (infinitely differentiable). This makes them well-suited for machine learning approaches, specifically Kolmogorov–Arnold Networks (KANs) (Liu et al., 2025).

The key methodological difference is that the Universal Approximation Theorem is based on approximating functions with neural networks, whereas the Kolmogorov–Arnold Theorem represents multivariate functions as sums of one-dimensional functions. Two important aspects should be highlighted. The first concerns accuracy, efficiency, and interpretability. Since KAN focuses on representing multivariable functions with one-dimensional functions, it can potentially provide more accurate representations while requiring less data. The second concerns architecture. Unlike MLPs, where activation functions are static, in the KAN architecture the activation functions are learnable and placed on the edges of the network. While MLPs have learnable linear weights, KANs use spline-based univariate functions, which differ across edges and are learnable (Vaca-Rubio et al., 2024).

KAN learns one-dimensional activation functions by tuning the positions of control points. Specifically, during the training process, these control points are updated through backpropagation with respect to a given loss function. In this way, the control points are adjusted to minimize the loss, thereby shaping activation functions that provide the best fit for the data. The interpretability trait of the KANs from architecture, one can see what is the discovered structure of activation functions. It should be noted, equation (31) provides KAN with two layers, where the upper index bound on outer summation $2n + 1$ is the width of hidden layer. The generalized form of KAN is represented by matrix-vector form in the following way:

$$\mathbf{y} = \mathrm{KAN}(\mathrm{x}) = (\boldsymbol{\Phi}_L \circ \boldsymbol{\Phi}_{L-1} \circ \ldots \circ \boldsymbol{\Phi}_1)\, \mathrm{x} \tag{32}$$

In the equation (32), L represents numbers of layers, $\boldsymbol{\Phi}$ is the matrix which is composed of ϕ univariate activation functions. It should be noted, in the original work Liu et al. (2025) of KAN, B-spline was incorporated as basis functions, which are formed via orthonormal basis of recursive functions (Chaudhuri, 2021). However, in the model discovery paper by Koenig et al. (2024), the authors instead of B-spline, use Gaussian radial basis function, which was initially proposed by Li (2024). The objective of the KAN is formulated in the similar way, as in SINDy methods. Specifically, the goal is to learn function f which governs the system:

$$\frac{d\mathbf{u}}{dt} = \mathbf{f}(\mathbf{u}, t) \tag{33}$$

However, instead, in the KAN-ODE version, the authors use Kolmogorov–Arnold theorem, to approximate the function f. The KAN-ODE representation then has the following form:

$$\frac{d\mathbf{u}}{dt} = \mathrm{KAN}(\mathbf{u}(t), \boldsymbol{\theta}) \tag{34}$$

In the equation (34), the KAN is approximating governing function f, u is the vector of independent variables, and θ is the vector of weights, which is used for optimizing the loss function $\mathcal{L}(\theta)$. As in the case of UDE, various forms of loss functions can be used for optimization, but the loss function used by the KAN-ODE algorithm is mean squared error (Koenig et al., 2024).

The empirical application of KAN-ODEs to the macro-informed Goodwin system, including nonstationarity robustness tests and comparison with SINDy and UDE+SINDy, is left to a companion paper. The computational cost of training KAN-ODEs is substantially higher than that of either SINDy or UDE+SINDy (a single run on the Goodwin system requires approximately four hours), which makes the extensive parameter sweeps conducted in the main text impractical within a single study.

A.4 SINDy sparse-regression workflow for a Goodwin-type (Lotka–Volterra) system

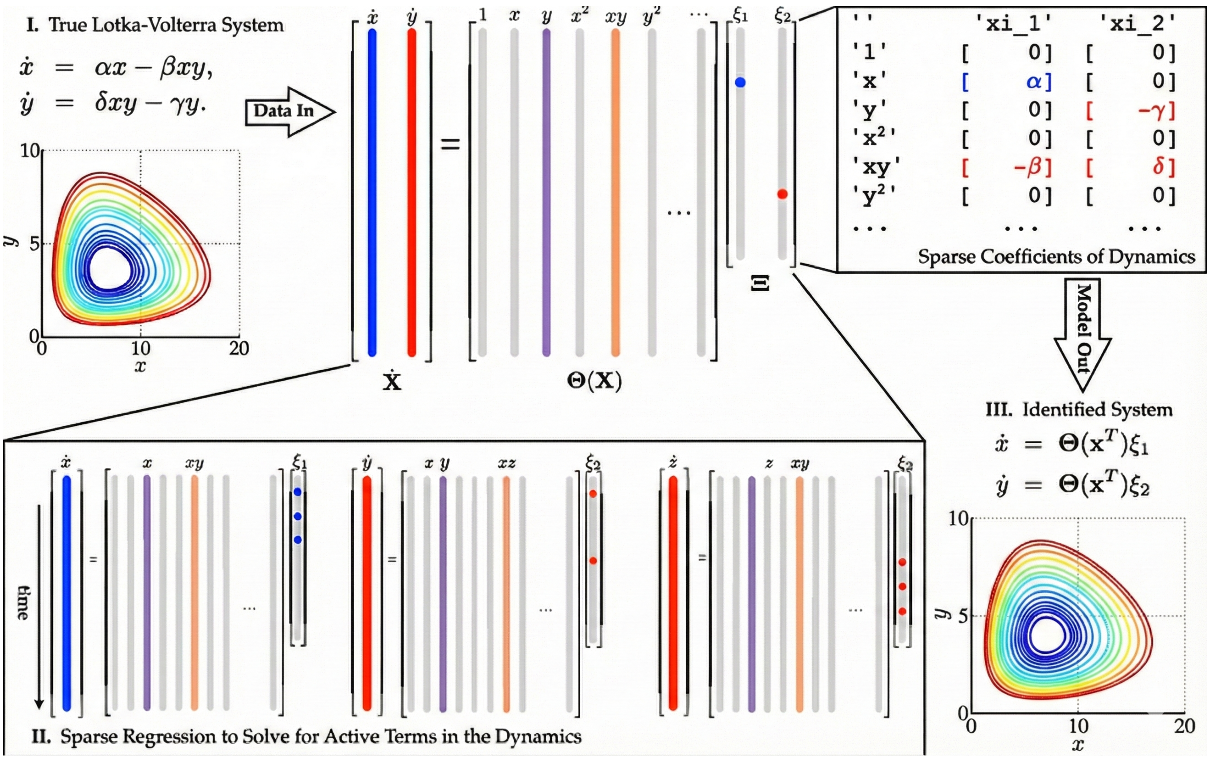

Figure 14: Illustration of the SINDy pipeline for a Goodwin-type (Lotka–Volterra) system: time-series data are used to construct a library of candidate functions $\Theta(X)$ and to solve a sparse regression problem for the active terms (sparse coefficients) that define the governing dynamics. The schematic follows the presentation in Brunton et al. (2016), where the authors provide an analogous figure for the Lorenz system.