SUERF
The European Money and Finance Forum

SUERF Policy Brief
No 1540, August 2026

When tariffs flip the textbook: Deep learning and the non-linear macroeconomics of trade policy

Author(s)

Massimo Ferrari Minesso | European Central Bank (ECB) and Complexity Lab in Economics
Carla Frenzel | European Central Bank (ECB)

Keywords: DSGE models, deep neural networks, solution methods, policy function approximation, non-linear solution

JEL codes: C45, C63, C68, E13, F13

Abstract

Most macroeconomic models used for policy are solved with local methods that assume the economy stays close to its long-run equilibrium. These methods are fast and reliable for understanding business-cycle fluctuations, but the accuracy of their solution deteriorates when shocks are large and dynamics turn non-linear – as with large trade-policy shifts. We develop a deep-learning method that solves DSGE models globally through a sequential training algorithm. Applied to a two-country model with tariffs, it uncovers an effect that local methods miss: while a small tariff appreciates the currency of the imposing country – the textbook result – a large tariff depreciates it, as contractionary effects come to dominate and monetary policy eases. A second-order perturbation, the standard "non-linear" alternative, both understates these effects and misses the direction of the exchange-rate response. These results underline how for questions involving large shocks, the choice of solution method is not a simple technicality but matters for policy answers.

Disclaimer: This Policy Brief is based on Ferrari Minesso, M. and Frenzel, C. (2026), "Sequential solution for DSGE models with deep neural networks", ECB Working Paper Series No 3236. The views expressed are those of the authors and do not necessarily reflect those of the European Central Bank or the Eurosystem.

Why solving general equilibrium models is hard

A macroeconomic model is only useful once it has been "solved": that is once the researcher has computed a rule describing how agents behave in every possible state of the economy and taking into account other agents' expectations and actions. That, in practice, means solving a complex non-linear system of dynamic equations that has not closed-form solution. The "solution" is then approximated through numerical methods, each of which rests on specific assumptions.

The workhorse method in macroeconomics is *perturbation*: decision rules are approximated by the solution of a linearised version of the model, typically a Taylor expansion around a fixed point. This turns a non-linear system of equations into a linearised system, that has a known analytical solution. The method is fast and scales well, which is why it dominates central-bank modelling. But it is inherently *local*. Its accuracy deteriorates as the economy moves away from the steady state, or when shocks are large, or non-linearities matter. Alternative solution methods are based on projections, that are accurate everywhere but become computationally infeasible as the number of state variables grows. A third alternative are perfect-foresight algorithms that can trace non-linear paths but become equally computationally demanding for large models and ignore uncertainty about future shocks.

Figure 1. Structure of a deep neural network

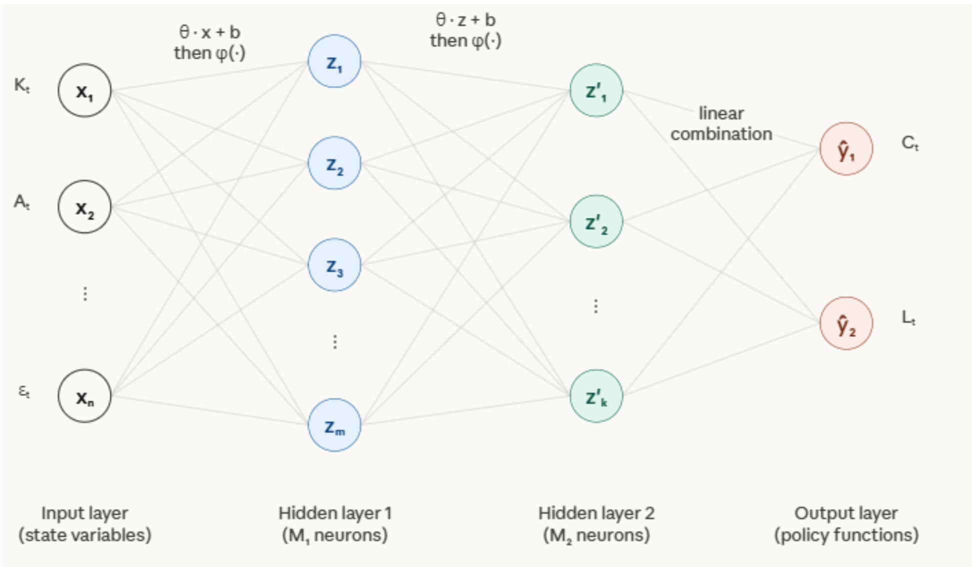

Notes: The figure illustrates a deep neural network used to solve a DSGE model. The input consists of the state variables (K, A) and the shock ε, while the outputs correspond to the model's control variables. The network comprises an input layer, an output layer, and two hidden layers. At each hidden layer, the outputs from the previous layer are combined through weighted connections (and biases), after which an activation function is applied to introduce non-linearity. The first hidden layer contains $M_1 \in [1, m]$ neurons, while the second hidden layer contains $M_2 \in [1, k]$ neurons. Each neuron computes a weighted combination of the outputs from the previous layer and passes the result through the activation function. The blue lines represent the weighted connections between neurons.

Deep neural networks offer a complement. Deep neural networks are a particular class of machine learning models that use a flexible mathematical function, built from many simple interconnected units (neurons), whose internal settings are tuned (trained) until its output matches a desired target, see Figure 1. By stacking these units in layers, deep neural networks can discover the right way to combine its inputs, allowing to approximate even complicated, unknown relationships to a high degree of accuracy. In macroeconomics, they can be used to approximate the unknown solution of a non-linear DSGE model (Maliar et al. 2021, Fernández-Villaverde, 2025, Scheidegger, 2026).

A non-linear solution method

The main practical obstacle to using deep learning for these models has been a circularity between solution and training. The network must be trained on data describing how the economy behaves – but that behaviour is governed by the very rules the network is trying to learn. Existing approaches sidestep this by borrowing a starting guess from a simpler, linear solution or using a grid of feasible points of state variables.

Our procedure addresses this problem from a different perspective. Training begins from the model's long-run equilibrium, which can always be computed analytically, and then introduces dynamics and uncertainty in four progressive phases (Figure 2). The network constructs its own training data as it improves, gradually concentrating effort on the region of the state space the economy actually visits. No auxiliary model and no pre-computed solution are required: a researcher needs only the steady state of the model to begin.

Figure 2. Learning the solution step by step

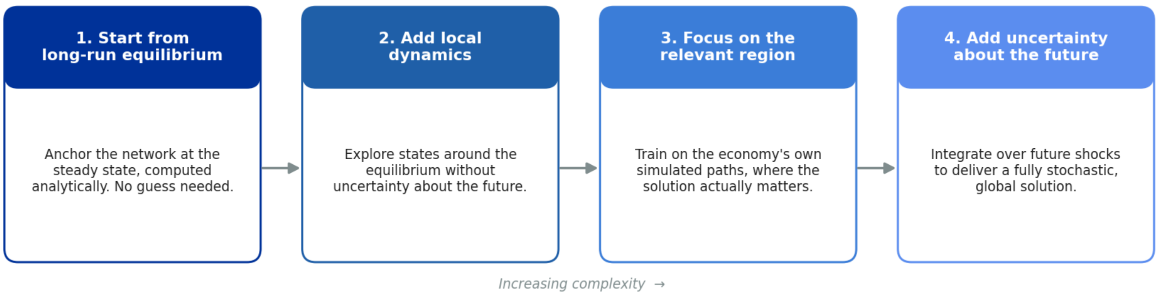

Notes: The method starts from the model's long-run equilibrium and adds dynamics and uncertainty in four stages, building its own training data as it goes – with no starting guess borrowed from a simpler model.

A systematic comparison across network designs (Figure 3) delivers a reassuringly simple recommendation: shallow networks of two or three layers, of moderate width, with an intermediate anchor to the steady state, consistently give the best accuracy at the lowest cost. Bigger is not better – deeper or wider networks bring no systematic gain and can prevent the algorithm from converging at all.

Figure 3. Comparison of penalty across different neural network configurations

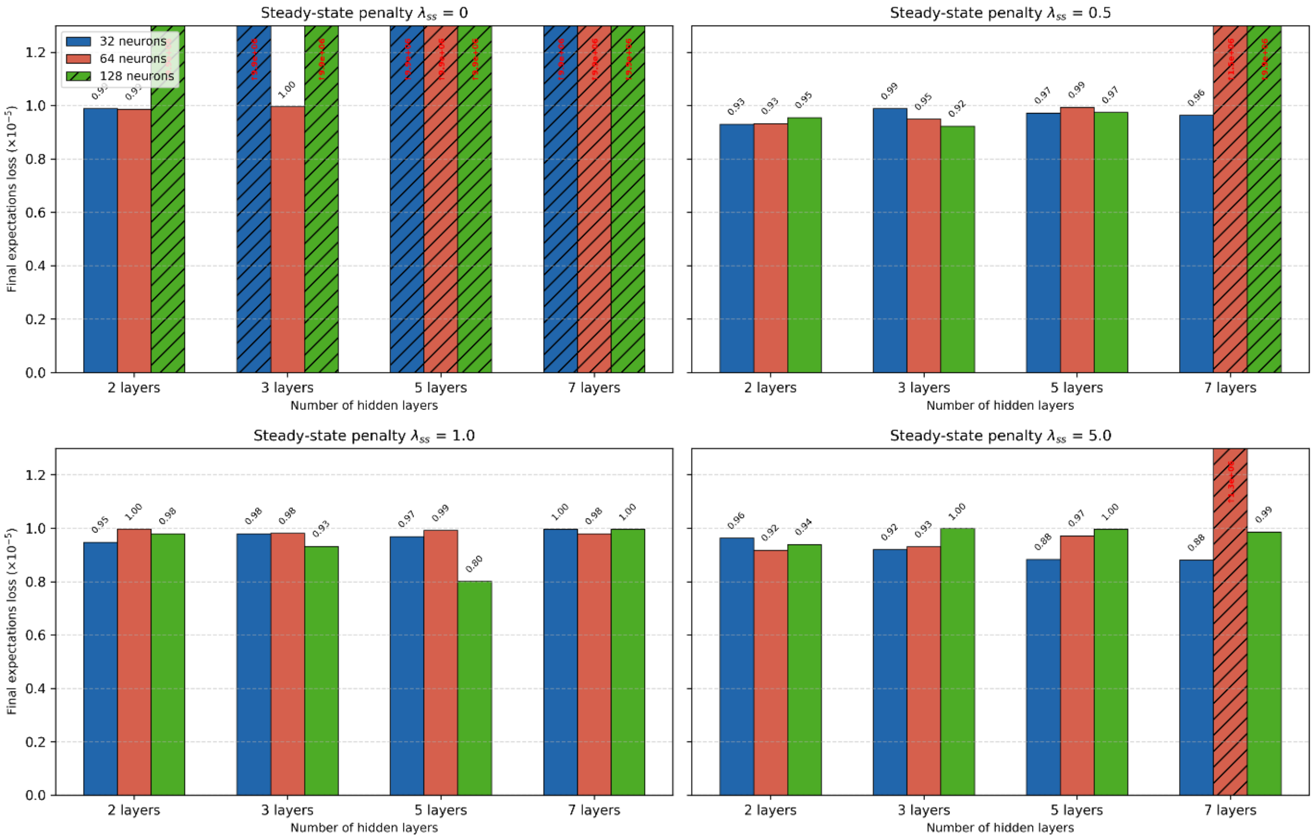

Notes: Each panel reports the final expectations loss (scaled by 10^{-5}) across network architectures for a given steady state penalty weight $\lambda_{ss} \in \{0, 0.5, 1, 5\}$; within each panel, bars are grouped by the number of hidden layers and colours distinguish networks with 32 (blue), 64 (red) or 128 (green) neurons per layer. A lower value indicates a more accurate approximation of the policy function at the end of Phase-4 training state. See Ferrari Minesso and Frenzel (2026) for further details.

Does it work? A benchmark check

Before turning to trade policy, it helps to verify the method on a model whose solution is already well understood: the canonical real business cycle model. Figure 4 compares the responses to a productivity shock obtained from the neural network with those from a standard solution package. The two agree closely, confirming that the network recovers the familiar dynamics.

The one economically meaningful difference is instructive. Because the network solves the model under genuine uncertainty about the future – integrating over the distribution of shocks – forward-looking households build up a precautionary buffer and consume slightly less than under a deterministic benchmark that ignores this risk. This is exactly the kind of effect a fully stochastic, global solution is designed to capture, and it is absent from methods that abstract from future uncertainty.

Figure 4. Impulse responses to a one standard deviation TFP shock in the classic RBC model

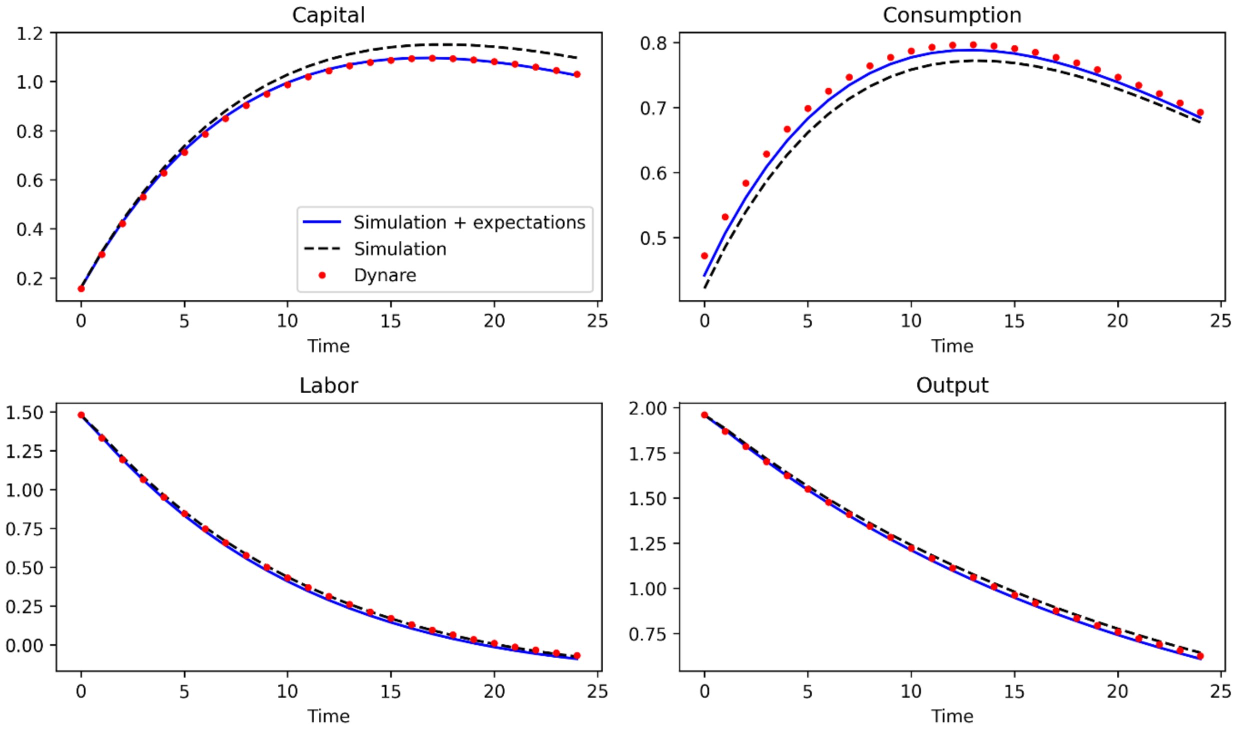

Notes: The red dots show impulse responses computed under Dynare's perfect foresight algorithm, the black dashed line are responses computed with the deep neural network trained in Phase-3 and the blue line are responses under the fully trained neural network, Phase-4. Impulse responses are reported in percent deviation from the deterministic steady state

Tariffs: when the textbook flips

The pay-off comes with more complex models and policy questions. We apply the method to a two-country model with nominal rigidities and an endogenous monetary policy, and study tariffs of the size seen in recent episodes – shocks far too large for local methods to handle reliably.

The central result is a *sign reversal*. Figure 5 compares a small tariff with a large one. If the model was linear, the two responses would look identical up to a scaling factor. They do not. After a small tariff, the currency of the imposing country *appreciates*, in line with the standard open-economy prediction. After a large tariff, it *depreciates* sharply. The logic is somewhat intuitive: under a large shock the contractionary effects of the tariff dominate, domestic demand falls steeply, and monetary policy eases aggressively – pushing the currency down rather than up. The same flip appears in consumption and inflation, whose impact responses change sign between the two scenarios.

Figure 5. Impulse response to a one-percent and ten-percent tariff shocks

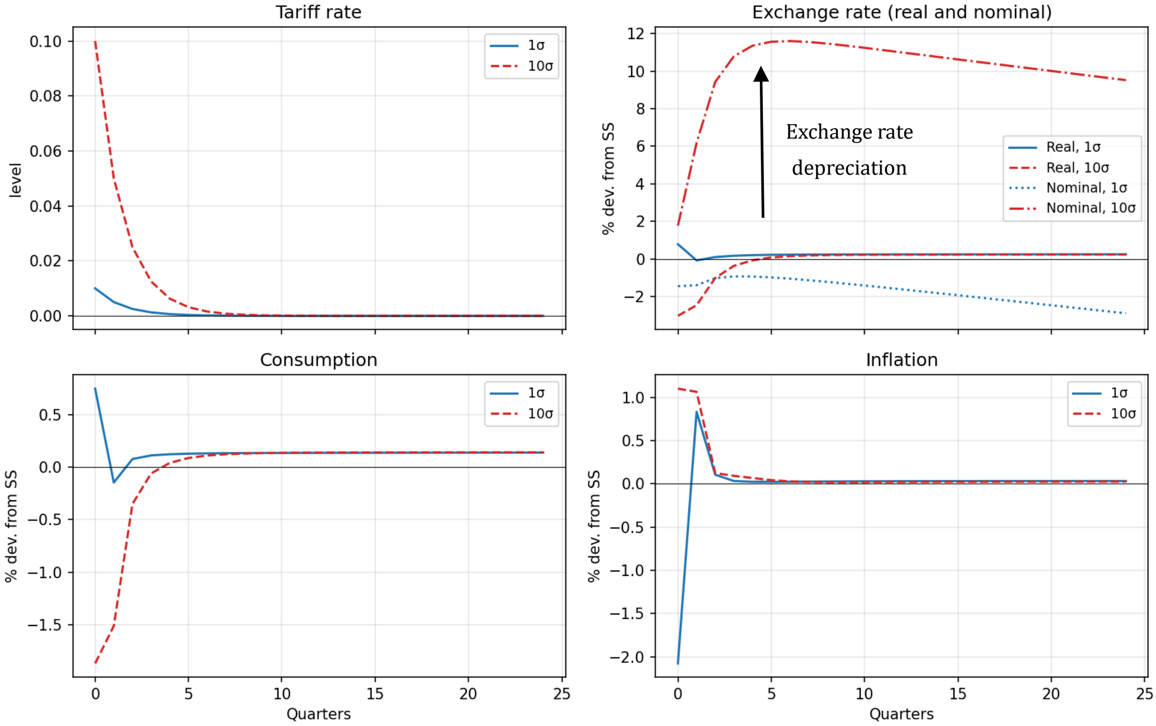

Notes: The chart shows the impulse responses to a one-percent increase in the effective tariff rate imposed by the Home country on imports from the Foreign country. A negative deviation of the nominal exchange rate denotes an appreciation of the Home currency. The model is solved with the deep neural network of Ferrari Minesso and Frenzel (2026).

This pattern matches the recent behaviour of the US dollar after 2 April 2025 and econometric evidence that the US dollar appreciates around unilateral tariff announcements but depreciates when retaliation is anticipated; see Ostry et al. (2026) and Furceri et al. (2022). The policy implications are that the qualitative effect of a tariff on the exchange rate depends on its size and on the impact on the domestic economy and how the domestic monetary policy authority reacts. Conclusions drawn from small calibrated shocks cannot simply be scaled up to the large shocks under discussion today.

Why "second-order" is not enough

A natural objection is that one need not go all the way to deep learning: a second-order perturbation already introduces some curvature and some shock-size dependence, at far lower cost. Figure 4 tests this directly.

For some variables the second-order solution gets the direction of the non-linearity right but *understates its magnitude* substantially. For the exchange rate it fails qualitatively: it simply scales up the small-shock appreciation and never reverses sign. The reason is structural – the mechanisms driving the reversal involve higher-order curvature that a fixed second-order expansion truncates away by construction. A global method commits to no particular order and absorbs that curvature automatically.

The practical implication is that standard "non-linear" solutions can be a useful but unreliable guide to the consequences of large trade shocks: they may compress the size of an effect, and occasionally miss some of the implications of large non-linearities.

Figure 6. Comparison between deep neural network and Dynare second-order solutions second-order

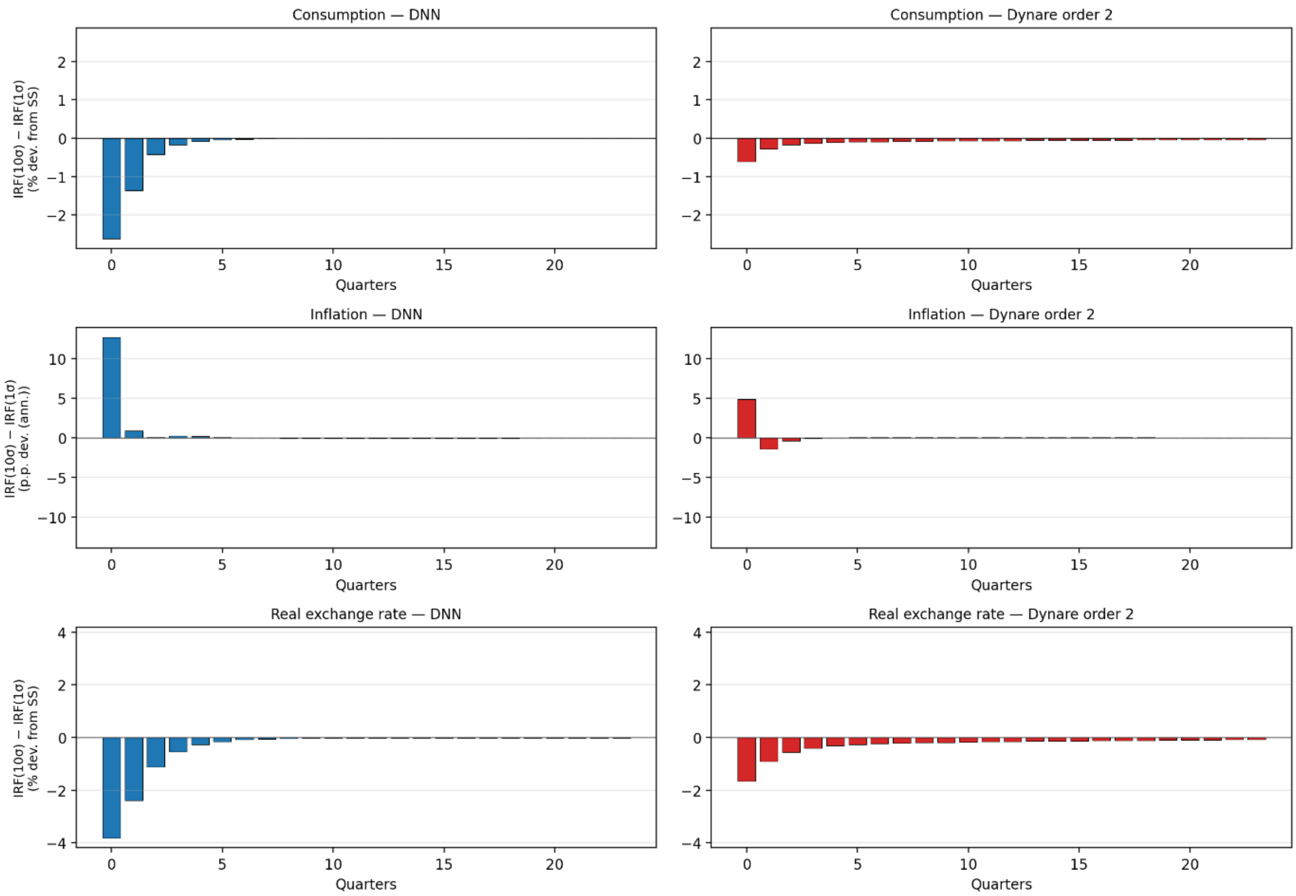

Notes: The chart shows the difference in the impulse responses of selected variables between the deep neural network solution and the second order perturbation solution of Dynare. Bars show the difference in impulse responses between a 1 standard deviation and a 10 standard deviation shock. The model is solved with the deep neural network of Ferrari Minesso and Frenzel (2026).

Takeaways for policy analysis

For policymakers, the central lesson is that the choice of solution method is not a technical detail left to model-builders but can change the policy conclusion itself. For the moderate shocks of routine forecasting, standard linear methods remain fast, reliable and entirely adequate. But many of the questions now reaching central banks – tariffs prominent among them, alongside the effective lower bound, energy-price spikes and financial stress – involve shocks large enough that a local approximation can be not just imprecise but qualitatively misleading. This matters also for policy makers in a country on the receiving end of a tariff, which faces a contractionary loss of export demand and must judge whether the exchange rate is cushioning or amplifying the blow. The sign of the response is decisive: a depreciation of the domestic currency restores some competitiveness but raises import prices, adding an inflationary impulse to an already weakening economy, whereas an appreciation dampens imported inflation but removes the competitiveness cushion just as the contraction bites hardest. A model that delivers the wrong sign does not merely misstate a magnitude – it can point the policy response in the wrong direction. Deep learning addresses this without displacing the existing toolkit, delivering globally accurate, fully stochastic solutions precisely where conventional methods may be least reliable. And because several toolkits are now available to implement these solutions, deep learning becomes now a practical option for the kind of large-shock, non-linear scenario analysis that features ever more prominently in policy work.

References

Adjemian, S., Juillard, M., Karamé, F., Mutschler, W., Pfeifer, J., Ratto, M., Rion, N., and Villemot, S. (2026). Dynare: Reference Manual, Version 7. Dynare Working Papers (number 87), CEPREMAP.

Fernández-Villaverde, J. (2025). Deep learning for solving economic models. NBER Working Paper (number 34250).

Ferrari Minesso, M., and Frenzel, C. (2026). Sequential solution for DSGE models with deep neural networks. ECB Working Paper No. 3236.

Furceri, D., Hannan, S. A., Ostry, J. D., and Rose, A. K. (2022). "The macroeconomy after tariffs". The World Bank Economic Review, 36(2), 361–381.

Maliar, L., Maliar, S., and Winant, P. (2021). "Deep learning for solving dynamic economic models". Journal of Monetary Economics, 122, 76–101.

Ostry, D., Lloyd, S., and Corsetti, G. (2025). Trading blows: The exchange-rate response to tariffs and retaliations. Bank of England Staff Working Paper (number 1,139).

Scheidegger, S. (2026). Deep Learning for Solving and Estimating Dynamic Models in Economics and Finance. arXiv preprint arXiv:2605.14493.

About the author(s)

Massimo Ferrari Minesso is a Lead Economist in the International Policy Analysis Division of the European Central Bank and a research fellow of the Complexity Lab in Economics. He holds a PhD in Economics from the Catholic University of Milan. His research interests are mainly in macro-finance, international macroeconomics and monetary economics.

Carla Fenzel is an analyst in the Market Operations Department at the European Central Bank (ECB), where she works on money markets and banks' liquidity needs. Prior to joining the ECB, she worked at the Deutsche Bundesbank and as a research analyst at the Max Planck Institute. She holds a Master's degree in Computer Science from Goethe University Frankfurt, where she wrote her thesis on quantum computing algorithms. Her research interests include computational methods, quantum algorithms, machine learning, and artificial intelligence, with applications to yield curve modelling, bank-level projections, and money market rates.

Designed by the Information Management and Services Division of the Oesterreichische Nationalbank (OeNB)

SUERF Secretariat
c/o OeNB, Otto-Wagner-Platz 3A-1090 Vienna, Austria
Phone: +43 1 40 420 7206
E-Mail: suerf@oenb.at
Website: https://www.suerf.org/