Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Discussion of "Dynamic Causal Effects in a Nonlinear World: the Good, the Bad, and the Ugly"

Edward P. Herbst, Benjamin K. Johannsen

2025-058

Please cite this paper as:
Herbst, Edward P., and Benjamin K. Johannsen (2025). "Discussion of "Dynamic Causal Effects in a Nonlinear World: the Good, the Bad, and the Ugly"," Finance and Economics Discussion Series 2025-058. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2025.058.

Discussion of "Dynamic Causal Effects in a Nonlinear World: the Good, the Bad, and the Ugly"

Edward P. Herbst and Benjamin K. Johannsen*

March 2025

Abstract

This comment discusses Kolesár and Plagbord-Møller (2025) finding that the standard linear local projection (LP) estimator recovers the average marginal effect (AME) even in nonlinear settings. We apply and discuss a subset their results using a simple nonlinear time series model, emphasizing the role of the weighting function and the impact of nonlinearities on small-sample properties.

1 Introduction

Kolesár and Plagbord-Møller (2025) (hereafter, KP) is an exciting, important advance in the literature on the estimation of dynamic causal effects in the context of local projections (LPs) (see Jordà (2005)). The paper establishes that the "standard" linear LP of an outcome y_{t+h} onto a shock x_t (and possibly a vector of controls) estimates an *average marginal effect* (AME) of the shock on the outcome. This result holds under suitable assumptions even—and perhaps especially—in the case of a nonlinear data generating process for y_t. Deriving the result requires connecting and extending a large literature in microeconometrics.

This comment aims to provide an accessible discussion of some of the results reported in KP that is tailored to macroeconomists. We begin by considering some of the theoretical results in KP under common assumptions in the macroeconomics literature. We devote particular attention to the weighting function, ω, that is used to compute the average in the AME. We then analyze the AME and its LP estimation in the context of the *quadratic autoregressive model* (*QAR(1,1)*) model of Aruoba et al. (2017). This is a stationary, nonlinear

*We thank Cristina Scofield for excellent research assistance. The views expressed here are those of the authors and do not indicate concurrence by the Board of Governors or anyone else associated with the Federal Reserve System. Herbst: Federal Reserve Board, edward.p.herbst@frb.gov, Johannsen: Federal Reserve Board.

time series model designed to mimic the statistical structure of a second-order approximation to the solution of a dynamic stochastic general equilibrium (DSGE) model. In the context of the *QAR(1,1)* model, we relate the population AME to a population nonlinear impulse response function ($NIRF$) defined in Koop et al. (1996). We also discuss small-sample properties of the LP estimator of the AME with a focus on how nonlinearities in the *QAR(1,1)* model affect those properties.

2 What does the standard LP estimate?

In this section, we discuss the LP estimator of the AME. To establish notation, let y_{t+h} be the observed outcome of interest at time $t+h$, and let x_t be the observed shock of interest at time t. Collect all other variables that determine y_{t+h} into a vector $U_{h,t+h}$, which may include past values of y_t, past (and future) values of x_t, and other controls. We require that the vector $U_{h,t+h}$ is independent of x_t. A representation of y_{t+h} based on x_t and $U_{h,t+h}$ is called the *structural function* and is given by

$$y_{t+h} = \psi_h(x_t, U_{h,t+h}). \tag{1}$$

The representation that is used to define the notion of dynamic causal effect used in KP is the *average structural function*, which is given by

$$\Psi_h(x_t) = \mathbb{E}[\psi_h(x_t, U_{h,t+h})] = \mathbb{E}[y_{t+h}|x_t] = g_h(x_t). \tag{2}$$

This function describes the expected outcome y_{t+h} given a specific value of the shock x_t, integrating out all other sources of randomness. Note that because we have assumed that x_t is independent of all other factors affecting y_{t+h}, the average structural function is equal to—and hence can be recovered from—the conditional expectation of y_{t+h} given x_t. This quantity can in principle be estimated from the data.

In macroeconomics, it can be difficult to estimate the average structural function due to small sample sizes. One approach is to impose strong assumptions about the data-generating process for y_t. For example, a researcher could assume that y_t follows an *AR(1)* process. Another option in nonlinear time series analysis is to estimate the AME, defined as

$$\theta_h(\omega) = \int \omega(x_t)\Psi_h'(x_t)dx_t. \tag{3}$$

Here $\Psi_h'(x_t)$ represents the derivative of the average structural function. This derivative captures the effect of an infinitesimal change in x_t on y_{t+h}. The weighting function, $\omega(x_t)$, determines how different values of x_t contribute to the AME, defining the sense in which the AME is an average.

KP study local projections that are indexed by h and given by

$$y_{t+h} = \beta_h x_t + \gamma_h' w_t + e_{h,t+h}. \tag{4}$$

Here, β_h is a parameter, w_t is a vector of controls, γ_h is a vector of parameters, and $e_{h,t+h}$ is an error term. If x_t and w_t have zero covariance, then under standard assumptions β_h is given by

$$\beta_h = \frac{\text{Cov}[g_h(x_t), x_t]}{\text{Var}[x_t]}, \tag{5}$$

where g_h is as defined in equation (2). KP show that

$$\beta_h = \int \omega(x_t) g_h'(x_t) dx_t, \tag{6}$$

where the weighting function, $\omega(x_t)$, is given by

$$\omega(x_t) = \frac{\text{Cov}[\mathbf{1}\{X > x_t\}, X]}{\text{Var}[X]}. \tag{7}$$

Here, x_t is a realization of the random variable X and $\mathbf{1}\{\cdot\}$ is the indicator function.

2.1 The weighting function

As KP explain, the function $\omega(x_t)$ has several desirable properties that make it suitable for computing an average. In particular, $\omega(x_t) \geq 0$ for all x_t and $\int_0^1 \omega(x_t) dx_t = 1$. From the properties of the indicator and covariance functions in equation (7), $\lim_{x_t \to -\infty} \omega(x_t) = 0$ and $\omega(x_t)$ is weakly increasing for $x_t < \mathbb{E}[X]$. Additionally, $\lim_{x_t \to \infty} \omega(x_t) = 0$ and $\omega(x_t)$ is weakly decreasing for $x_t \geq \mathbb{E}[X]$. Taken together, these facts imply that $\omega(x)$ is hump-shaped and that it gives most of the weight to values of x_t near the mean of X. As discussed in the paper, if X follows a Normal distribution, then $\omega(x_t) = \phi(x_t)$ where $\phi(\cdot)$ denotes the (appropriately parameterized) probability density function of the Normal distribution.

It is worth emphasizing that $\omega(x_t)$ depends only on the properties of the random variable X. It does not depend on the outcome y_{t+h}, on the distribution of other random variables that may go into the construction of y_{t+h}, or on nonlinear dependence of y_{t+h} on past realizations of y_t.

We can also deduce some additional properties of ω if we assume that X has a continuous density function with full support on the real line, even if X is not Normally distributed. These are common assumptions about identified shocks in the macroeconomics literature. For simplicity, set $\mathbb{E}[X] = 0$ and $\text{Var}[X] = 1$ so that

$$\omega(x) = \mathbb{E}[\mathbf{1}\{X > x\}X]. \tag{8}$$

We will focus on contrasting ω with f_X, the density function associated with X, and thus, the *AME* with the *expected marginal effect* (*EME*). The *EME* is given by $\theta(f_X)$, which uses the true probability density function $f_X(x_t)$ as the weight function. Here are two additional properties of ω not shown in the paper.

Claim. *If the distribution of X has heavy tails in the sense that as $|x_t| \to \infty$, for some $\alpha > 2$ and $C > 0$*

$$\lim_{x_t \to \infty} \frac{f_X(x_t)}{\frac{C}{x_t^{1+\alpha}}} = 1 \text{ and } \lim_{x_t \to -\infty} \frac{f_X(x_t)}{\frac{C}{x_t^{1+\alpha}}} = 1$$

then $\omega(x_t)$ has heavier tails than $f_X(x_t)$.

To see this, for large x_t, substitute the tail approximation into the definition of ω

$$\omega(x) \approx \int_x^\infty t \frac{C}{t^{1+\alpha}} dt = C \int_x^\infty t^{-\alpha} dt = C \frac{x^{-(\alpha-1)}}{\alpha - 1}.$$

We deduce that

$$\frac{\omega(x)}{f_X(x)} \approx \frac{x^2}{\alpha - 1}.$$

For sufficiently large x this object is greater than 1.

Claim. *Assume that X has finite moments of order $j+2$ and $\lim_{x_t \to \infty} \omega(x_t) x_t^{j+1} = 0$. Then*

$$\int_{-\infty}^{\infty} x^j \omega(x) dx = \frac{1}{j+1} \mathbb{E}[X^{j+2}].$$

The result of this claim follows from integration by parts.

In summary, when X is normally distributed, the LP estimator aligns exactly with the *EME*, making $\omega(x) = f_X(x)$ a convenient theoretical benchmark. For distributions with skewness or heavy tails, the *AME* estimated by the local projection will differ from the *EME*, potentially biasing the interpretation of results. However, this difference also reflects the robustness of the local projection estimator in capturing effects relevant in the tails of the distribution.

To illustrate the weighting function, Figure 1 shows f_X and ω for four distributions. All of the distributions are normalized to have mean zero, so ω peaks at zero in each panel. The

Figure 1: $\omega(x)$ and $f_X(x)$ for four distributions

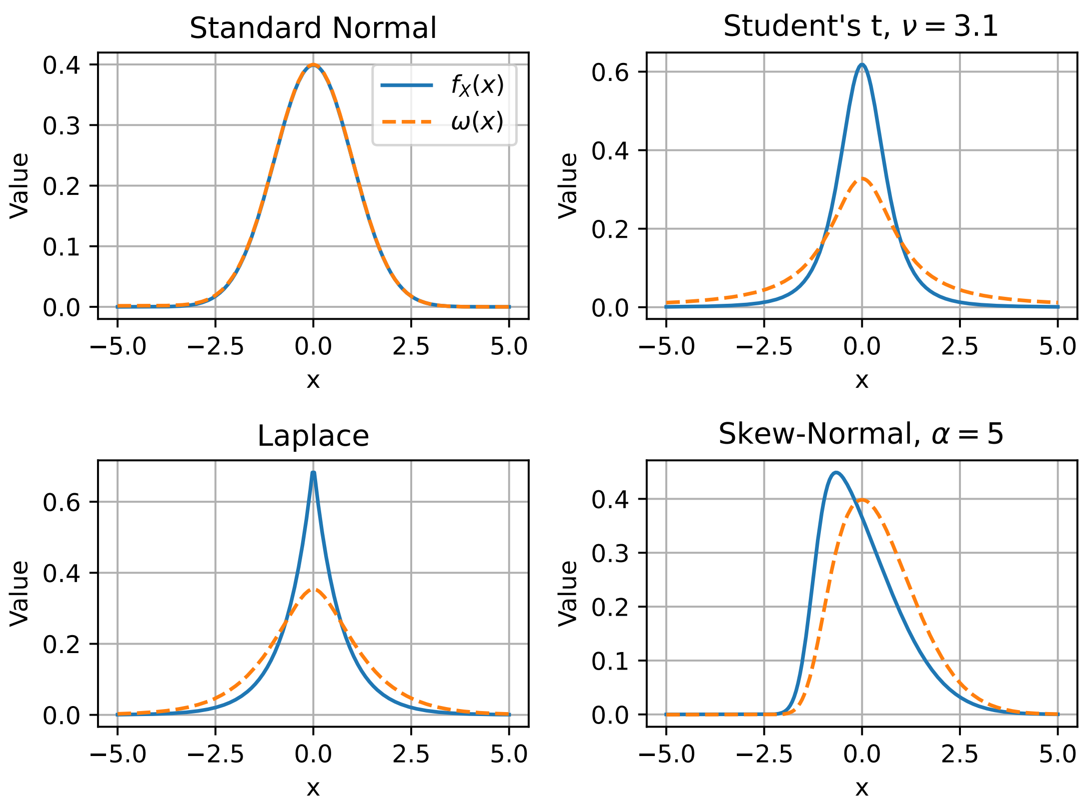

Source: Authors' calculations.

upper-left panel shows that f_X and ω coincide for the standard Normal distribution. The upper-right panel and the lower-left panel show that if f_X has heavy tails, then ω has heavier tails. The lower-right panel shows that if f_X is a skew distribution, then f_X and ω do not peak at the same point.

3 AME in Action: Quadratic Autoregressive Model

In this section we analyze the AME and its estimation using a simple nonlinear, time series model—the first-order *quadratic autoregressive model* or *QAR(1,1)* model—that was developed in Aruoba et al. (2017). This class of time series models was developed to identify nonlinearities in macroeconomic data and to evaluate DSGE models. This model is particularly useful because it nests the familiar *AR(1)* model, which is a common benchmark.

The *QAR(1,1)* is given by

$$y_t = \phi_0 + \phi_1(y_{t-1} - \phi_0) + \phi_2 s_{t-1}^2 + (1 + \gamma s_{t-1})\sigma x_t \quad (9)$$

$$s_t = \phi_1 s_{t-1} + \sqrt{1 - \phi_1^2} x_t. \quad (10)$$

Here, y_t is the observed scalar variable of interest, s_t is an unobserved state variable, and x_t is an observed shock with mean zero and variance unity. We assume that the third moment of x_t is finite. The scalar parameters ϕ_2 and γ control the degree of nonlinearity in the model. Very roughly speaking, γ is associated with conditional heteroskedasticity in y_t, and ϕ_2 is associated with asymmetry and more general state dependence. When $\phi_2 = \gamma = 0$, the model collapses to the *AR(1)* model.

3.1 The NIRF and the AME in the QAR(1,1)

As KP note, in a nonlinear time series setting there are a variety of notions of impulse response. Here, we briefly touch on a nonlinear impulse response function ($NIRF$)—see Koop et al. (1996)—as a definition familiar to most macroeconomists, and later we describe its connection to the objects studied in KP.

In the context of the *QAR(1,1)* model, the $NIRF$ is given by

$$\text{NIRF}(h, x_t, y_{t-1}, s_{t-1}) = \mathbb{E}\left[y_{t+h}|x_t, y_{t-1}, s_{t-1}\right] - \mathbb{E}\left[y_{t+h}|y_{t-1}, s_{t-1}\right].$$

This object can be expressed as

$$\text{NIRF}(h, x_t, y_{t-1}, s_{t-1}) = \phi_1^h (1 + \gamma s_{t-1}) \sigma x_t + \phi_2 \phi_1^{h-1} \frac{1-\phi_1^h}{1-\phi_1} \delta(s_{t-1}, x_t), \tag{11}$$

$$\text{where } \delta(s_{t-1}, x_t) = 2\phi_1 s_{t-1} \sqrt{1-\phi_1^2} x_t + (1-\phi_1^2)(x_t^2 - 1).$$

Notice that the $NIRF$ depends on s_{t-1} but not on y_{t-1}. That is, the $NIRF$ is state dependent and the relevant state is s_{t-1}. The analytical expression for the $NIRF$ makes clear that the model is asymmetric in the sense that

$$NIRF(h, x_t, y_{t-1}, s_{t-1}) + NIRF(h, -x_t, y_{t-1}, s_{t-1}) \neq 0.$$

That is, adding the $NIRF$ from a positive shock to the $NIRF$ from a negative shock of the same size does not equal zero. Allowing for asymmetry is important for analyzing the transmission of macroeconomic shocks—see, for example, Kilian and Vigfusson (2011). Additionally, the analytical expression for the $NIRF$ makes clear that the model displays heteroskedasticity in the sense that if $x_t \neq 0$ and $\kappa \neq 1$ then

$$NIRF(h, \kappa x_t, y_{t-1}, s_{t-1}) - \kappa NIRF(h, x_t, y_{t-1}, s_{t-1}) \neq 0.$$

That is, the $NIRF$ is not homogeneous of degree one in x_t. Interestingly, these nonlinearities do not depend on the level of s_{t-1} even though the $NIRF$ is affected (linearly) by s_{t-1}.

Recall that KP focus on the representation of y_{t+h} given by equation (1), where $y_{t+h} = \psi(x_t, U_{h,t+h})$, and that the variables in the vector $U_{h,t+h}$ are independent of x_t. Taking expectations over $U_{h,t+h}$, we are left with the average structural function

$$\Psi_h(x_t) = \mathbb{E}[\psi_h(x_t, U_t)] = \mathbb{E}[y_{t+h}|x_t].$$

This concept is different from the $NIRF$, in that Ψ_h averages both future and past shocks, while the $NIRF$ defined above explicitly conditions on past information. That is, unlike the $NIRF$, Ψ_h does not feature any state dependence on s_{t-1}. A motivation for focusing on Ψ_h instead of the $NIRF$ is that in most applications, such as the setup for the *QAR(1,1)* model that we consider here, s_{t-1} is unobserved.

Although distinct, Ψ_h and the $NIRF$ are related in that

$$\Psi_h(x_t) = \mathbb{E}[NIRF(h, x_t, y_{t-1}, s_{t-1})|x_t] + \mathbb{E}[y_t].$$

From the parametric expression for the $NIRF$ in the *QAR(1,1)* model, we then have that

$$\Psi_h(x_t) = \phi_0 + \frac{\phi_2}{1-\phi_1} + \phi_1^h \sigma x_t + \phi_2 \phi_1^{h-1}(1-\phi_1^h)(1+\phi_1)(x_t^2 - 1). \tag{12}$$

Notice that the asymmetry and heteroskedasticity of the model are apparent from Ψ_h. However, an indication of the information lost in Ψ_h relative to the $NIRF$ because of the averaging is that γ does not affect Ψ_h even though it contributes importantly to the nonlinearity in the model and appears in the $NIRF$.

KP focus on the estimation of the AME given in equation (3). From the parametric expression for Ψ_h in the *QAR(1,1)* model

$$\Psi_h'(x_t) = \phi_1^h \sigma + 2\phi_2 \phi_1^{h-1}(1-\phi_1^h)(1+\phi_1)x_t. \tag{13}$$

Applying the results related to $\omega(x_t)$ when x_t is continuously distributed gives

$$\theta_h(\omega) = \phi_1^h \sigma + \phi_2 \phi_1^{h-1}(1-\phi_1^h)(1+\phi_1)\mathbb{E}[x_t^3]. \tag{14}$$

Here, we have used our second claim, discussed above. Notice that if x_t is symmetric then $\theta(\omega) = \phi_1^h \sigma$ and the AME in the *QAR(1,1)* model is the same as in the *AR(1)* model. From equation (14), it is also clear that when $\phi_2 = 0$ the AME is the same in these models even if x_t is not symmetric. More generally, in a linear model—that is, a model in which the $NIRF$ is linear in x_t—Ψ_h is linear in x_t. In this case, if x_t follows a standard Normal distribution, then the $NIRF$ for a one standard deviation shock and the AME are equivalent.

3.2 Estimation with small samples

Here we analyze the estimation of the AME in the *QAR(1,1)* model. With an observed set of outcomes and shocks $\{y_t, x_t\}_{t=1}^T$, the estimation of the set of regressions given by equation (4) is straightforward. However, as emphasized in Herbst and Johannsen (2024), even in this idealized setting, finite sample issues can be important, particularly for sample sizes commonly seen in the macroeconomics literature.

Here, we consider how nonlinearities interact with finite sample issues. We simulate the *QAR(1,1)* model $1,000,000$ times with $\phi_0 = 0$, $\phi_1 = 0.95$, and $\sigma = 1$.[1] We think of the model as a quarterly model and use a sample size of 100, which Herbst and Johannsen (2024) argue is typical in the related macroeconomics literature.

We vary ϕ_2 and γ to see how different values—and their associated nonlinearities—affect

[1] We initialize y_{-1000} and s_{-1000} to zero and simulate forward. We begin our sample at y_0 and s_0.

the finite sample properties of the estimator of the AME. We specify the local projection as

$$y_{t+h} = \beta_h x_t + \gamma_h'[1, y_{t-1}]' + e_{h,t+h}.$$

We focus on the case where $h = 6$, which is a relatively short horizon for an impulse response in macroeconomics. We denote the estimator of the AME by $\hat{\beta}_h$.

As a baseline case, we use a standard Normal distribution for x_t. To understand the effects of heavy tails or asymmetries in the distribution of x_t, we also consider the cases where x_t has a (standardized) t distribution with $\nu = 3.1$ degrees of freedom and when x_t has a (standardized) skew-Normal with skew parameter $\alpha = 5$. Note that when x_t follows a Normal distribution or a t distribution the AME in the *QAR(1,1)* model coincides with that in the *AR(1)* model with ϕ_1 as the autoregressive parameter. When x_t has a skew-Normal distribution, the AME differs because of the non-zero third moment of x_t, unless $\phi_2 = 0$.

We compare the small-sample average value of $\hat{\beta}_h$ to β_h (the bias). We know β_h in closed form from the derivations above. The results are shown in Figure 2. When $\phi_2 = 0$ and $\gamma = 0$, the *QAR(1,1)* model reduces to the *AR(1)* model. At that point on the graphs, the bias is the same for each of the three distributions. For different values of ϕ_2 and γ, the bias of the AME estimator depends on the distribution. Interestingly, introducing nonlinearities does not necessarily increase or decrease bias. That is, nonlinearities have unpredictable effects on the small-sample average of $\hat{\beta}_h$

To further explore the small-sample properties of $\hat{\beta}_h$, Figure 3 shows the root-mean-squared-error (RMSE) of $\hat{\beta}_h$, and Figure 4 shows the coverage probability of nominal 95% confidence intervals constructed using $\hat{\beta}_h$ and associated Huber-White standard errors. Notably, as the nonlinearities of the *QAR(1,1)* model increase (ϕ_2 and γ increase in magnitude) the RMSE grows. This increased volatility of $\hat{\beta}_h$ is not fully captured by the standard errors. As a result, the coverage probabilities fall as nonlinearities of the *QAR(1,1)* model increase.

We conclude that although the AME is robust to an array of nonlinearities in population, those nonlinearities may have important implications for the small sample properties of estimators and associated test statistics.

Figure 2: Bias Estimates for the AME in the QAR(1,1) model with $h = 6$ and $T = 100$

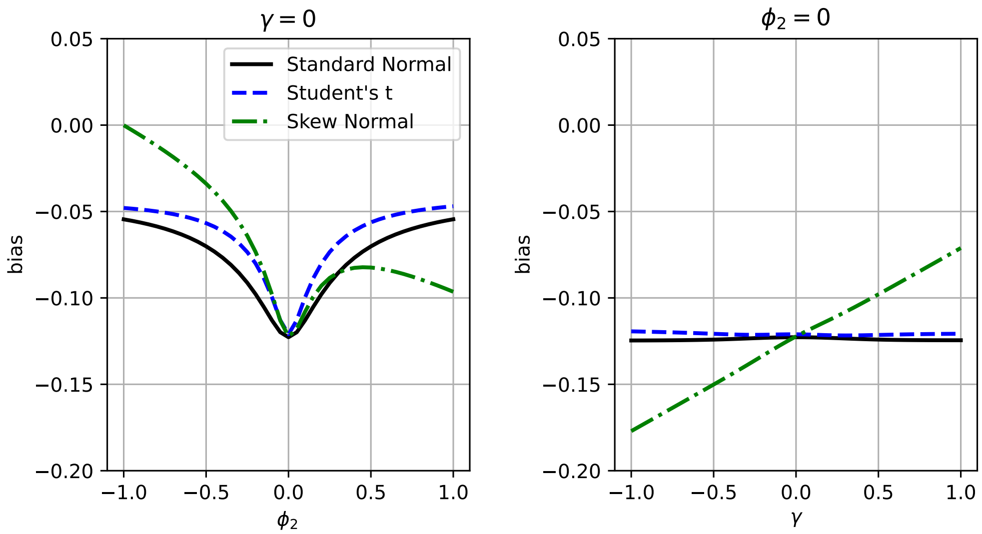

Source: Authors' calculations.

Figure 3: RMSE for the AME in the QAR(1,1) model with $h = 6$ and $T = 100$

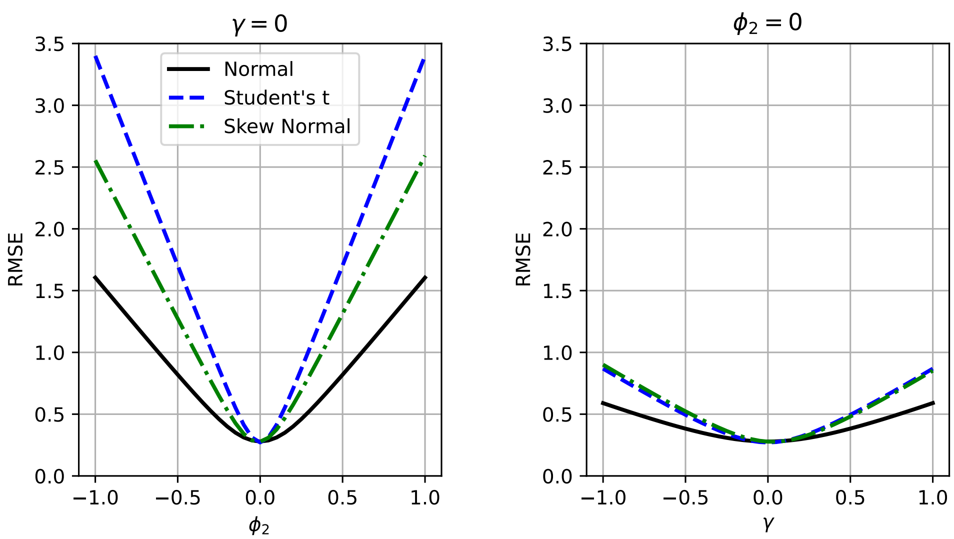

Source: Authors' calculations.

Figure 4: Coverage of 95% CIs for the AME in the QAR(1,1) model with $h = 6$ and $T = 100$

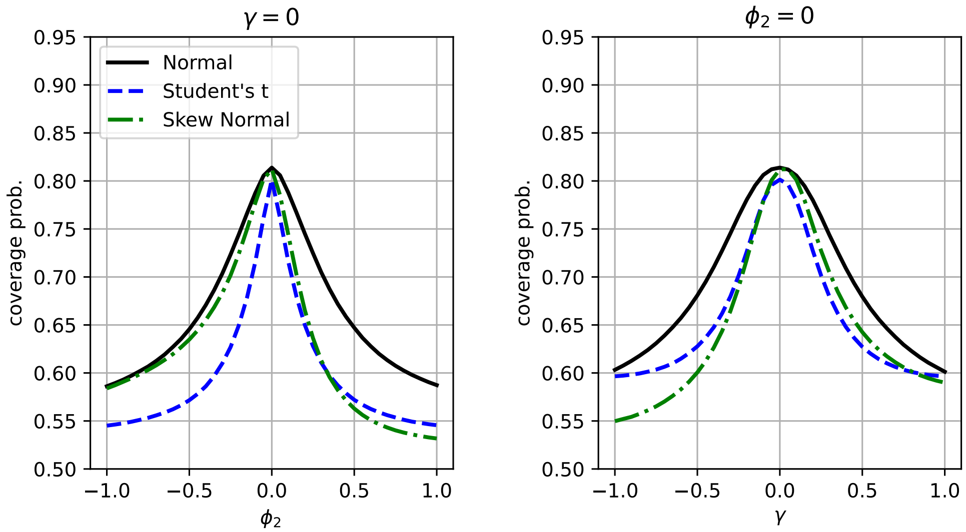

Source: Authors' calculations.

References

Aruoba, S. B., L. Bocola, and F. Schorfheide (2017): "Assessing DSGE model nonlinearities," *Journal of Economic Dynamics and Control*, 83, 34–54.

Herbst, E. P. and B. K. Johannsen (2024): "Bias in local projections," *Journal of Econometrics*, 240.

Jordà, O. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95, 161–182.

Kilian, L. and R. J. Vigfusson (2011): "Are the responses of the U.S. economy asymmetric in energy price increases and decreases?" *Quantitative Economics*, 2, 419–453.

Kolesár, M. and M. Plagbord-Møller (2025): "Dynamic Causal Effects in a Nonlinear World: the Good, the Bad, and the Ugly," *Journal of Business Economics and Statistics*.

Koop, G., M. H. Pesaran, and S. M. Potter (1996): "Impulse response analysis in nonlinear multivariate models," *Journal of econometrics*, 74, 119–147.