UP School of Economics
Discussion Papers

Discussion Paper No. 2026-03 | March 2026

Oil price pass-through in the Philippines: decomposing fuel and non-fuel inflation

by

Jan Carlo B. Punongbayan

University of the Philippines School of Economics

Oil price pass-through in the Philippines: decomposing fuel and non-fuel inflation

Jan Carlo B. Punongbayan*

University of the Philippines School of Economics

March 27, 2026

UPSE Discussion Paper; comments are most welcome

Abstract

This paper estimates the pass-through of world oil price innovations to Philippine fuel prices and headline CPI, and decomposes total CPI pass-through into a fuel-basket component and a residual non-fuel component. Using a structural VAR and 25 years of monthly pump price data, I find that a 10 percentage point increase in year-on-year oil price growth is associated with about a 4.9 percentage point increase in gasoline price growth and a 6.6 percentage point increase in diesel price growth at 12 months, while the corresponding effect on headline CPI inflation is about 0.65 percentage points. The residual non-fuel component accounts for the larger share of the CPI response, though its estimated magnitude is somewhat sensitive to the estimation method. Results are robust to extensions with the exchange rate and rice prices, local projections with HAC inference, alternative data transformations, sub-period splits around the TRAIN Law, and an alternative pump price series.

Keywords: oil price pass-through, consumer prices, structural VAR, Philippines
JEL classification: E31, Q43, C32, F31

*Email: jbpunongbayan@up.edu.ph. Thanks to Marianne Juco-Rodriguez and Sonny Africa for assistance with the time series data. The author used generative AI to assist in this paper and assumes full responsibility for the final output.

1 Introduction

The Philippines imports virtually all of its crude oil requirements. As a small open economy, it is directly exposed to fluctuations in world oil prices, which transmit to domestic consumer prices through multiple channels. Understanding the magnitude, speed, and composition of this pass-through is critical for monetary policy design at the Bangko Sentral ng Pilipinas (BSP), which operates an inflation targeting framework with a 2–4 percent target band.

Despite the obvious policy relevance, the Philippine literature on oil price pass-through is surprisingly thin. Most existing work either examines the oil–inflation relationship at a highly aggregated level or focuses on specific commodities (e.g., rice, transport fares) without providing a structural decomposition of the transmission channels. There is, to my knowledge, no study that cleanly separates the component of oil-to-CPI pass-through attributable to the domestic fuel basket from the residual non-fuel component.

The main empirical objective of this paper is to estimate the pass-through of oil price innovations to Philippine fuel prices and headline CPI, and to decompose the total CPI pass-through into a fuel-basket component and a residual non-fuel component. I apply the structural VAR (SVAR) framework of Yilmazkuday (2021), who estimates oil price pass-through for the United States using weekly data. The approach is a three-variable SVAR with Cholesky (recursive) identification, ordered from world oil prices to domestic fuel prices to headline consumer prices. The decomposition of the total CPI response into a fuel-basket term and a residual is derived formally from the CPI expenditure identity (Appendix A). Extensions incorporating the peso-dollar exchange rate and rice prices, and robustness checks using local projections, alternative transformations, and alternative data sources, are treated as supporting exercises rather than headline contributions.

A distinguishing feature of this paper is the use of actual retail pump prices—monthly prevailing gasoline and diesel prices in Metro Manila from the Department of Energy (DOE)—rather than the fuel CPI sub-index that is standard in the pass-through literature. In a deregulated market like the Philippines, where oil companies adjust pump prices weekly in response to world oil and exchange rate movements, actual pump prices more faithfully capture the crude-to-retail transmission than a CPI sub-index. The pump price series extends from 2000 to 2026, providing 298 monthly observations for estimation after the year-over-year transformation—roughly four times the sample available from the CPI fuel sub-index (available only from 2018). As Kilian (2024) emphasizes, temporal aggregation and index construction choices can have substantive effects on structural inference about oil price transmission.

Using monthly data from January 2001 to October 2025, I find three main results.

First, the fuel pass-through is substantial: gasoline pass-through (PG) reaches 49 percent and diesel pass-through reaches 66 percent at 12 months, consistent with a deregulated fuel market. Second, the residual non-fuel component accounts for the larger share of the total consumer price response in the gasoline specification, while the two components are more balanced in the diesel specification. Third, the total long-run consumer price pass-through is approximately 6.5 percent—a one percentage point increase in year-over-year oil price growth is associated with an increase of about 0.065 percentage points in headline CPI inflation in the long run. Diesel pass-through is consistently higher than gasoline pass-through across all specifications, plausibly reflecting differences in tax structure and diesel's role in freight and public transport. Robustness checks using local projections with HAC-corrected inference, monthly log differences, pre- versus post-TRAIN Law subsamples, rolling windows, and an alternative pump price series from CEIC confirm the broad stability of the fuel pass-through estimates, though the CEIC brand-average data yield systematically lower pass-through and the estimated magnitude of the residual component is somewhat sensitive to the estimation method.

It is important to be explicit about what this paper identifies and what it does not. The Cholesky-identified oil price innovation is whatever variation in oil prices remains after conditioning on lagged values of all variables in the system; it is not a cleanly identified exogenous oil supply shock and may reflect a mix of supply, demand, and speculative forces. The decomposition of CPI pass-through into fuel-basket and residual components is accounting-based, derived from the CPI expenditure identity under the assumption that the fuel CPI weight is known. The contribution is therefore to quantify the magnitude and timing of pass-through under a transparent recursive framework, not to recover fully structural transmission channels. An exploratory proxy-SVAR using Känzig (2021)'s OPEC announcement surprises as an external instrument is reported in Appendix D; the instrument proves weak in the Philippine context ($F \approx 3$), so the exercise is informative about the direction of bias but not definitive.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature. Section 3 describes the data. Section 4 presents the econometric framework, including the formal derivation of the decomposition and the rationale for the baseline year-over-year transformation. Section 5 reports the baseline results and a systematic comparison of gasoline and diesel pass-through. Section 6 presents extensions incorporating the exchange rate and rice prices. Section 7 assesses robustness. Section 8 concludes.

2 Literature review

Yilmazkuday (2021) provides the framework this paper adapts. Using weekly US data in an SVAR with Cholesky identification, he decomposes oil-to-CPI pass-through into a direct channel (via the gasoline CPI basket) and an indirect channel (all other transmission). He finds that gasoline pass-through (PG) reaches about 50 percent in the long run, with the indirect channel slightly exceeding the direct. Corsello and Tagliabracci (2023) arrive at qualitatively similar conclusions for the euro area, though they find that the indirect channel becomes much larger during energy crises, suggesting state dependence. Conflitti and Luciani (2019) sharpen the mechanism using a dynamic factor model on 88 disaggregate price indices for the US and euro area, finding that the "common" (macroeconomic) pass-through component is small but persistent, while the item-specific energy-intensity component is statistically indistinguishable from zero.

The pass-through literature for emerging economies is thinner. Ndou and Gumata (2025) estimate a threshold SVAR for South Africa and find that pass-through is roughly three times lower when inflation is within the target band, consistent with Taylor's (2000) hypothesis about pricing power in low-inflation environments. Ding et al. (2023) introduce the exchange rate as a mediating channel for China. Gagliardone and Gertler (2023) estimate a New Keynesian model in which oil is a complement to labor in production; the low elasticity of substitution amplifies marginal cost effects and is likely relevant for oil-intensive economies like the Philippines.

On methodology, Kilian (2024) emphasizes that temporal aggregation can attenuate estimated pass-through, motivating the use of actual pump price data rather than CPI sub-indices. Känzig (2021) proposes using OPEC announcement surprises as an external instrument in a proxy-SVAR framework (Mertens and Ravn, 2013; Stock and Watson, 2012); I implement this as an exploratory exercise in Appendix D, though the instrument proves weak in the Philippine context. For the Philippines specifically, there is no published study that cleanly decomposes oil pass-through into direct and indirect channels using structural methods.

3 Data

The analysis combines monthly retail fuel prices in Metro Manila with Dubai crude oil prices and the Philippine headline consumer price index.

3.1 Pump prices and oil prices

Monthly prevailing retail gasoline and diesel prices in Metro Manila (PHP per liter) and Dubai crude oil prices (USD per barrel) are obtained from the Department of Energy (DOE). The DOE's "prevailing price" reflects the most common retail price observed across surveyed stations in the National Capital Region during the reference month; it is neither a simple average nor an end-of-month spot price, but rather the modal price at any given time. The series spans January 2000 to February 2026, providing 314 monthly observations. Figure 1 displays all three series. The strong co-movement between Dubai crude and domestic pump prices is visually apparent, with pump prices tracking global oil price movements through the 2008 spike, the 2014–16 decline, the 2020 COVID collapse, and the 2022 energy crisis.

An important caveat is that the pump prices are for Metro Manila only, while the CPI is a national measure. Metro Manila is the price-leading market, and the PSA's CPI sampling is heavily weighted toward the capital region, but some attenuation of the estimated pass-through may result from this geographic mismatch—particularly for the non-fuel component, which partly operates through provincial agriculture and transport costs. Regional pump price data from the DOE, available for a shorter sample, could be used in future work to assess whether the Metro Manila estimates generalize to the rest of the country. As a partial check on the DOE price data, Section 7.4 re-estimates the model using brand-level daily pump prices from CEIC, aggregated to monthly frequency.

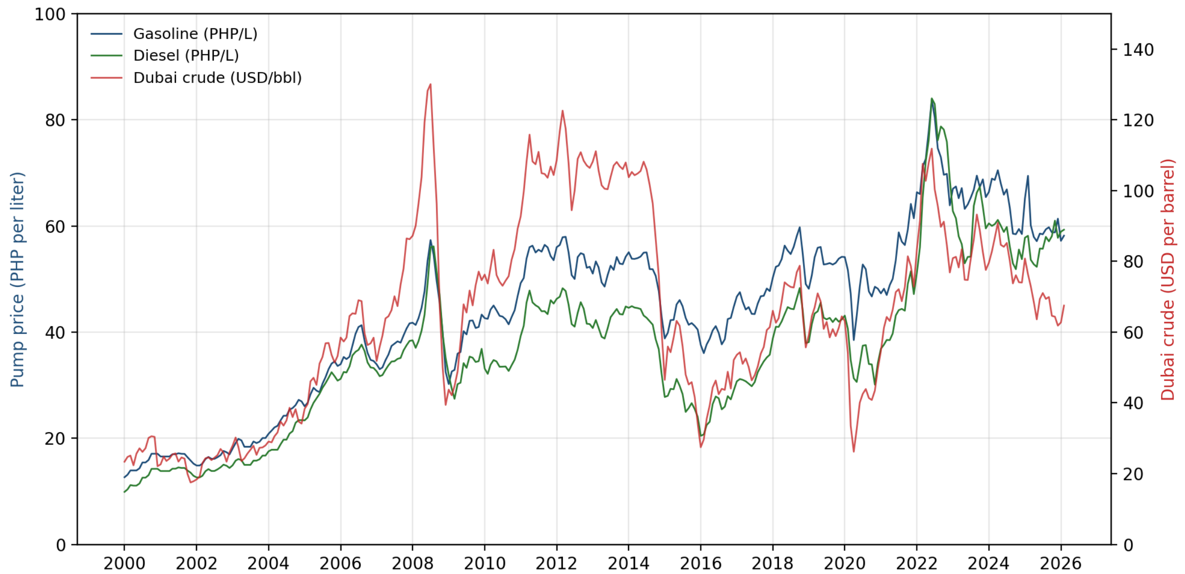

Figure 1: Monthly prevailing prices of gasoline and diesel in Metro Manila (PHP per liter, left axis) and Dubai crude oil (USD per barrel, right axis), 2000–2026. Source: Department of Energy.

The pump prices reflect the fully deregulated retail market established under

RA 8479 (1998). Under this regime, oil companies—principally Petron, Shell, and Chevron Philippines—adjust retail prices weekly based on international product prices, the peso-dollar exchange rate, and applicable taxes. The DOE monitors but does not control these price adjustments.

3.2 Headline CPI

The all-items consumer price index is sourced from the Philippine Statistics Authority (PSA). Since the estimation requires a continuous monthly CPI series from 2000 onward, I splice two PSA series: the 2006-based CPI (available 1994–2018) and the 2018-based CPI (available from 2018 onward), rescaling the older series to match the 2018-based index at the January 2018 overlap. Because all variables enter the SVAR as 12-month log differences, the choice of index base does not affect the estimated dynamics.

3.3 Exchange rate

The monthly average peso-dollar exchange rate is sourced from the Bangko Sentral ng Pilipinas (BSP), spanning January 2000 to February 2026. The exchange rate enters the extended four-variable SVAR as a year-over-year log change.

3.4 Rice CPI sub-index

The rice CPI sub-index is sourced from the PSA. As with headline CPI, I splice two series: the 2012-based rice CPI (PCOICOP 01.1.1.12, available January 2012 to December 2021) and the 2018-based rice CPI (available January 2018 to February 2026), rescaling the 2012-based series at the January 2018 overlap. The spliced series provides 170 monthly observations from January 2012 to February 2026, yielding 156 usable observations in year-over-year log changes (January 2013 to October 2025). Rice is the single largest item in the Philippine CPI basket and is driven by supply-side factors (weather, government policy, imports) that are largely orthogonal to oil. Including it in an extended SVAR helps disentangle food-driven inflation from oil-driven inflation and disciplines the forecast error variance decomposition.

3.5 Variable transformation

All variables are transformed to year-over-year (12-month) log changes for estimation, expressed in percentage points:

$$\Delta x_t = (\ln x_t - \ln x_{t-12}) \times 100 \tag{1}$$

This transformation removes seasonality and renders the series stationary, while preserving the economic interpretation as annual growth rates. The estimation sample covers January 2001 to October 2025 ($N = 298$ observations after the year-over-year transformation and VAR lag requirements).

Table 1: Summary statistics

Variable	N	Mean	Std. dev.	Min	Max
Panel A: Levels (monthly)					
Dubai crude (USD/bbl)	314	64.60	27.55	17.50	130.08
Gasoline pump price (PHP/L)	314	43.84	16.00	12.65	83.85
Diesel pump price (PHP/L)	314	36.74	15.23	9.89	84.04
Headline CPI (2018=100)	310	72.31	18.39	47.95	128.40
PHP/USD exchange rate	310	49.59	4.90	40.43	58.82
Panel B: Year-over-year log changes (percentage points)					
Δ Oil price	298	4.04	34.15	−99.63	87.97
Δ Gasoline pump price	298	5.61	16.84	−37.28	44.87
Δ Diesel pump price	298	6.16	23.60	−52.18	66.69
Δ Headline CPI	298	3.88	1.84	0.35	10.02
Δ Exchange rate	298	1.07	6.69	−17.80	23.17

Notes: Sample period is January 2000–February 2026 (levels) and January 2001–October 2025 (YoY changes). Pump prices are monthly prevailing prices in Metro Manila from the DOE. CPI is a spliced series using the 2006-based and 2018-based PSA indices.

4 Econometric framework

4.1 Why year-over-year inflation is the baseline

The choice to estimate the VAR in year-over-year log changes merits explicit justification, as this transformation shapes the dynamics and interpretation of the results.

First, Philippine CPI and fuel prices are naturally discussed in year-over-year inflation terms. The BSP's inflation target is defined over 12-month CPI growth, and the DOE and PSA report headline inflation as a year-over-year measure. Estimating the VAR in these units aligns the empirical objects with the quantities of direct policy interest.

Second, the year-over-year transformation addresses seasonality transparently, by construction, without requiring a set of seasonal dummies whose specification could influence the results. Monthly log differences, by contrast, are noisier and more sensitive to seasonal specification choices.

Third, this paper is about medium-run pass-through—how oil price movements cumulate in fuel prices and headline CPI over 6–36 months—rather than very short-run weekly pricing microstructure. The year-over-year transformation captures low-frequency dynamics that month-on-month changes may miss, particularly in the residual non-fuel component that builds gradually.

The year-over-year transformation does create overlapping observations (consecutive months share 11 of 12 data points), inducing an MA(11) structure in the VAR residuals. This raises legitimate inference concerns. Section 7.2 addresses this directly by re-estimating the pass-through using local projections with Newey-West (HAC) standard errors, and Section 7 reports a parallel specification using month-on-month log differences with seasonal controls. The qualitative findings are robust across these alternatives (Table 7), confirming that the year-over-year baseline is a disciplined choice rather than a convenience.

4.2 The structural VAR model

I follow Yilmazkuday (2021) in specifying a three-variable structural VAR. The vector of endogenous variables is:

$$\mathbf{y}_t = \begin{bmatrix} \Delta\text{oil}_t \\ \Delta\text{fuel}_t \\ \Delta\text{cpi}_t \end{bmatrix} \tag{2}$$

where Δoil is the year-over-year log change in Dubai crude oil prices, Δfuel is the year-over-year log change in the domestic fuel price variable, and Δcpi is the year-over-year log change in Philippine headline CPI. All variables are expressed in percentage points.

I estimate two versions of the system using pump price data: one with Metro Manila gasoline prices and one with diesel prices. The reduced-form VAR(p) is:

$$\mathbf{y}_t = \mathbf{c} + \mathbf{A}_1\mathbf{y}_{t-1} + \cdots + \mathbf{A}_p\mathbf{y}_{t-p} + \mathbf{u}_t \tag{3}$$

where $\mathbf{c}$ is a constant vector, $\mathbf{A}_1, \ldots, \mathbf{A}_p$ are coefficient matrices, and $\mathbf{u}_t$ is the vector of reduced-form residuals with covariance matrix $\mathbf{\Sigma}$.

4.3 Identification

Structural identification uses the Cholesky decomposition of $\mathbf{\Sigma}$, which imposes a recursive ordering: oil prices are ordered first, domestic fuel prices second, and headline CPI last. The ordering reflects the assumption that the Philippines is a price-taker in the world oil market and that fuel prices adjust to oil before broader consumer prices do.

It is important to be clear about what this identification does and does not achieve. The Cholesky-identified oil price innovation is whatever variation in oil prices remains after conditioning on lagged values of all variables in the system. It is not a cleanly identified oil supply shock: it may capture a mix of supply-driven, demand-driven, and speculative oil price movements. The impulse responses should therefore be interpreted as the dynamic association between recursively identified oil price innovations and domestic prices, not as the causal effect of an exogenous oil supply disturbance. I use the term "oil price innovation" throughout to reflect this distinction.

4.4 Pass-through decomposition

The pass-through measures follow Yilmazkuday (2021). This is an *accounting decomposition*, not a structural identification of distinct economic mechanisms. It separates the CPI response into a component mechanically attributable to the fuel basket and a residual that captures everything else—transport costs, production costs, wage effects, expectation effects, and any other channels—without distinguishing among them.

Let $\mathrm{IRF}(i, j, h)$ denote the orthogonalized impulse response of variable j to a one-standard-deviation Cholesky-identified innovation to variable i at horizon h. The cumulative impulse responses are:

$$\mathrm{CIR}(i, j, h) = \sum_{s=0}^{h} \mathrm{IRF}(i, j, s) \tag{4}$$

The *fuel pass-through* (PG) measures what share of an oil price innovation is transmitted to domestic fuel prices:

$$\mathrm{PG}(h) = \frac{\mathrm{CIR}(\mathrm{oil}, \mathrm{fuel}, h)}{\mathrm{CIR}(\mathrm{oil}, \mathrm{oil}, h)} \tag{5}$$

The *consumer price pass-through* (PC) measures what share reaches headline CPI:

$$\mathrm{PC}(h) = \frac{\mathrm{CIR}(\mathrm{oil}, \mathrm{cpi}, h)}{\mathrm{CIR}(\mathrm{oil}, \mathrm{oil}, h)} \tag{6}$$

The total PC is decomposed into a fuel-basket component (DPC) and a residual non-fuel component (IPC):

$$\mathrm{DPC}(h) = w_{\mathrm{fuel}} \times \mathrm{PG}(h), \qquad \mathrm{IPC}(h) = \mathrm{PC}(h) - \mathrm{DPC}(h) \tag{7}$$

where $w_{\mathrm{fuel}} = 0.0401$ is the expenditure weight of transport fuels (PCOICOP 07.2.2) in the 2018-based CPI basket. Appendix A derives this decomposition formally from the CPI expenditure identity and states the assumptions under which it holds. The fuel-basket component (DPC) captures the mechanical contribution of fuel price changes to headline CPI, weighted by fuel's share in the consumption basket. The residual non-fuel component (IPC) is defined residually after removing the mechanically implied fuel-basket contribution; it should not be interpreted as a separately identified structural channel. It captures whatever remains of the total CPI response that is not attributable to the fuel basket—this could include transport costs, production input costs, wage adjustments, expectation effects, or any other mechanism. The labels "direct" and "indirect" are used as shorthand throughout, following Yilmazkuday (2021), but should be understood strictly as an accounting partition.

Because all variables enter the model as year-over-year log changes expressed in percentage points, PG and PC are elasticity-like objects: PG = 0.50 means a one percentage point increase in year-over-year oil price growth is associated with a 0.50 percentage point increase in fuel price growth by horizon h. Because the CPI fuel weight is only about 4 percent, even large PG translates into small DPC: if PG = 0.50, then DPC $= 0.0401 \times 0.50 = 0.020$, so the fuel-basket channel accounts for only 0.020 percentage points of the CPI inflation response per percentage point of oil price growth.

The year-over-year transformation means the cumulative impulse responses describe accumulated movements in annual inflation rates, not price-level changes. For a one-time oil shock, the year-over-year oil variable rises on impact and reverts as the shock passes the 12-month window, so the pass-through ratios stabilize at a finite value. PG and PC are best interpreted as long-run elasticities of annual growth rates. The overlapping nature of the transformation induces serial correlation in the VAR residuals; Sections 7 and 7.2 re-estimate the model using month-on-month log differences and local projections with HAC-corrected inference to verify that the qualitative findings are not driven by the overlapping-observations structure.

An additional caveat applies to the diesel specification, discussed in detail in Section 5.4: the CPI weight $w_{\text{fuel}} = 0.0401$ reflects household expenditure on transport fuels, which is primarily a gasoline concept, so the diesel DPC/IPC split should be interpreted with particular caution.

4.5 Lag selection, estimation, and inference

The lag order is selected by the Bayesian information criterion (BIC), which selects $p = 2$ in both the gasoline and diesel specifications. The Akaike information criterion (AIC) selects $p = 3$ in the gasoline specification and $p = 2$ in the diesel specification; using $p = 3$ does not materially change the pass-through estimates (gasoline PG at 12 months is 0.49 under both $p = 2$ and $p = 3$). I use $p = 2$ as the baseline, consistent with BIC's preference for parsimony, and note that results are robust to $p \in \{1, 2, 3\}$.

The model is estimated by ordinary least squares equation-by-equation. The effective estimation sample is January 2001 to October 2025 ($T = 296$ usable observations after accounting for 2 lags).

Confidence intervals for the impulse responses are constructed by Monte Carlo simulation with 1,000 replications, drawing from the estimated reduced-form residual covariance matrix. This parametric bootstrap provides valid inference under the assumption of i.i.d. Gaussian innovations, which the year-over-year transformation violates through the induced MA(11) overlap structure. To address this, Section 7.2 supplements the baseline with local projection estimates using Newey-West (HAC) standard errors with bandwidth $\max(h+1, 12)$, providing honest inference under serial correlation. The impulse response horizon is set to 36 months, long enough for the pass-through ratios to stabilize.

4.6 Extension: four-variable SVAR with exchange rate

To assess the exchange rate channel, I extend the baseline to a four-variable SVAR:

$$\mathbf{y}_t = \begin{bmatrix} \Delta\text{oil}_t \\ \Delta\text{fx}_t \\ \Delta\text{fuel}_t \\ \Delta\text{cpi}_t \end{bmatrix} \tag{8}$$

where Δfx is the year-over-year log change in the PHP/USD exchange rate (an increase denotes peso depreciation). The Cholesky ordering places the exchange rate after oil prices and before domestic fuel prices, reflecting the identifying assumption that world oil prices are exogenous to the peso, peso movements contemporaneously affect fuel pricing (since pump price adjustments reflect both crude prices and the exchange rate), and CPI responds to all three within the month.

5 Results

5.1 Impulse responses

Figure 2 shows the orthogonalized impulse response functions to a one-standard-deviation Cholesky-identified oil price innovation in the gasoline specification, with 95 percent Monte Carlo confidence bands. A positive oil price innovation raises oil price growth persistently over the first 6–8 months before reverting toward zero. Gasoline pump prices respond positively and substantially, with the response peaking at approximately 3–4 months. The response is statistically significant and the confidence band excludes zero for the first 10–12 months. Headline CPI shows a small but positive and persistent response, rising gradually and remaining elevated through 36 months.

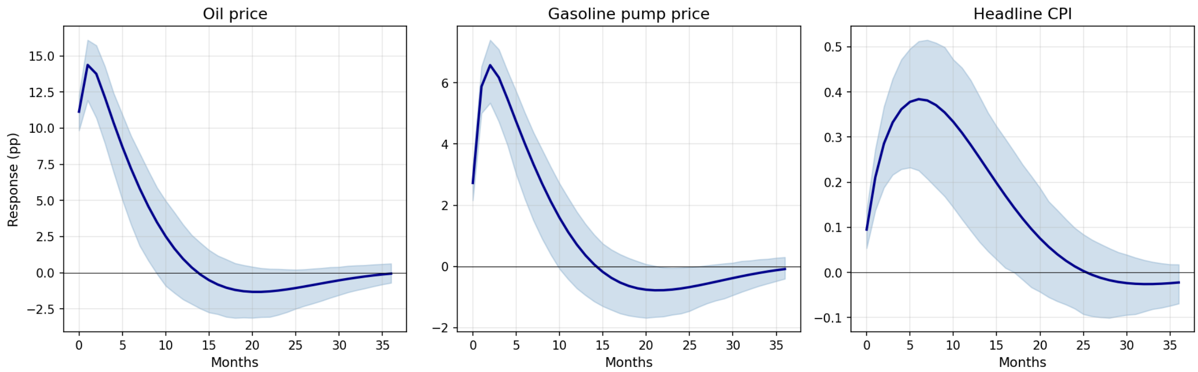

Figure 2: Impulse responses to a one-standard-deviation Cholesky-identified oil price innovation (gasoline specification). Blue line: point estimate. Shaded area: 95 percent Monte Carlo confidence band (1,000 replications).

Figure 3 displays the cumulative impulse responses, which are the basis for the pass-through calculations. Gasoline pump prices accumulate a large positive response to the oil price innovation, consistent with substantial crude-to-retail price transmission. The cumulative response of headline CPI is positive and growing, reaching its long-run level at approximately 24–30 months.

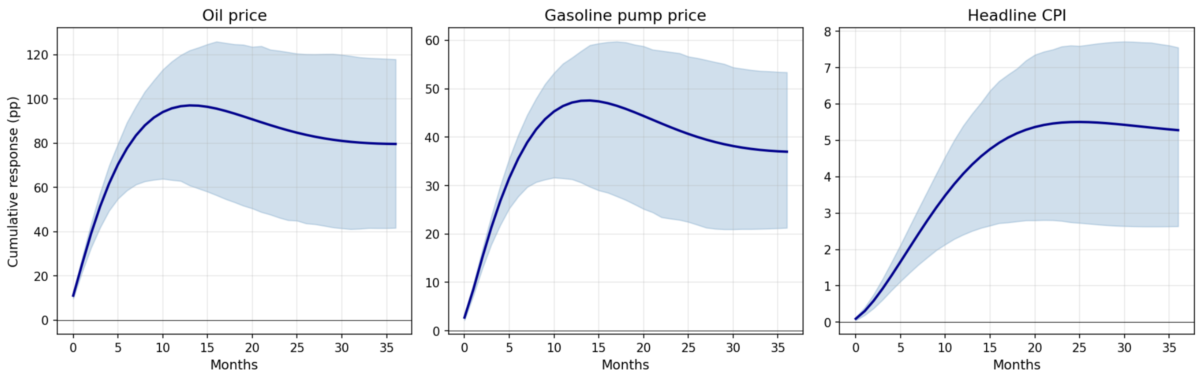

Figure 3: Cumulative impulse responses to a one-standard-deviation Cholesky-identified oil price innovation (gasoline specification). Blue line: point estimate. Shaded area: 95 percent Monte Carlo confidence band.

5.2 Pass-through decomposition

Table 2 reports the pass-through estimates for both specifications. In the gasoline specification, PG rises to 0.49 at 12 months and stabilizes around 0.46 at 36 months—indicating that approximately 46–49 percent of a Cholesky-identified oil price innovation is associated with a corresponding movement in gasoline pump prices. This is in the neighborhood of Yilmazkuday (2021)'s US estimate of approximately 50 percent, though differences in data frequency, variable definitions, tax structures, and market organization limit the comparability of cross-country point estimates. The total consumer price pass-through (PC) rises gradually from 2.6 percent at 6 months to 6.6 percent at 36 months. The residual non-fuel component (IPC) is the larger share at longer horizons, accounting for about 72 percent of total pass-through at 36 months (IPC = 0.048 versus DPC = 0.019). This indicates that the fuel-basket channel alone accounts for only a modest fraction of the total CPI response; the remainder is captured by the residual, which the current framework does not decompose further.

In the diesel specification, PG is higher: 0.66 at 12 months, rising to 0.71 at 18 months before settling at 0.65 at 36 months. The accounting decomposition is more balanced: at 36 months, DPC = 0.026 and IPC = 0.038, with the residual non-fuel component accounting for 60 percent of the total. Section 5.4 discusses the gasoline–diesel difference in detail.

Table 2: Oil price pass-through estimates: Gasoline and diesel pump price specifications

Specification	Horizon	PG	PC	DPC	IPC
Gasoline pump price					
	6 months	0.458	0.026	0.018	0.008
	12 months	0.488	0.042	0.020	0.023
	18 months	0.491	0.056	0.020	0.036
	24 months	0.482	0.064	0.019	0.045
	36 months	0.464	0.066	0.019	0.048
Diesel pump price					
	6 months	0.566	0.025	0.023	0.003
	12 months	0.664	0.040	0.027	0.014
	18 months	0.709	0.054	0.028	0.025
	24 months	0.705	0.063	0.028	0.034
	36 months	0.646	0.064	0.026	0.038

Notes: PG = fuel pass-through (cumulative fuel response / cumulative oil response). PC = consumer price pass-through (cumulative headline CPI response / cumulative oil response). DPC = fuel-basket component ($w_{\text{fuel}} \times$ PG, where $w_{\text{fuel}} = 0.0401$). IPC = residual non-fuel component (PC $-$ DPC). Both specifications use VAR(2) with Cholesky identification and $N = 298$.

Figure 4 visualizes the pass-through decomposition for the gasoline specification. PG is large and positive, rising steeply over the first 12 months and stabilizing near 0.47. The right panel shows that the residual component (IPC) grows steadily and exceeds the fuel-basket component (DPC) at approximately 10 months, ultimately accounting for the larger share of the total CPI response.

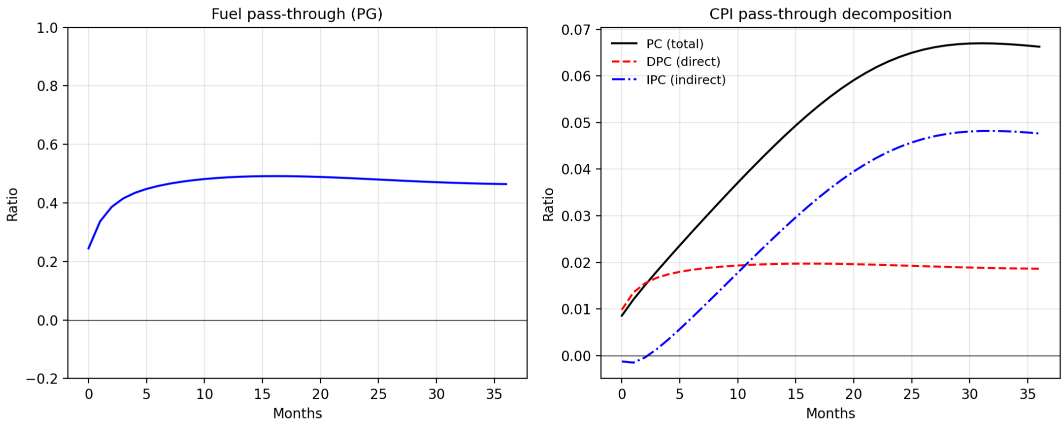

Figure 4: Oil price pass-through decomposition (gasoline pump price specification). Left panel: fuel pass-through (PG). Right panel: consumer price pass-through decomposed into fuel-basket (DPC) and residual non-fuel (IPC) components.

5.3 Forecast error variance decomposition

Appendix Table 10 reports the forecast error variance decomposition (FEVD) for the gasoline specification. FEVD in small recursive systems is sensitive to the variable set and ordering, so these shares should be treated as descriptive rather than as a headline result. Within the recursive identification, oil price innovations account for 43 percent of CPI forecast error variance at 12 months and 46 percent at 36 months. These shares are conditional on the three-variable information set; Section 6.2 shows that expanding the system to include rice prices reduces the oil share only modestly, with rice absorbing variance primarily from the CPI-own component.

Granger causality tests from the gasoline specification confirm that oil prices Granger-cause gasoline pump prices ($F = 62.48$, $p < 0.001$) and headline CPI ($F = 5.28$, $p = 0.005$). The full set of pairwise Granger causality results is reported in Appendix Table 11.

5.4 Why diesel pass-through is higher

The gasoline and diesel specifications are not simply parallel estimations; they capture distinct transmission mechanisms that merit comparison. Diesel PG is consistently 15–20 percentage points higher than gasoline PG across all specifications and horizons. Three factors plausibly explain this gap.

First, the TRAIN Law's per-unit excise tax structure drives a larger wedge between crude oil prices and the gasoline pump price than between crude and diesel. As of 2024, the excise tax on gasoline is PHP 10.00 per liter versus PHP 6.00 per liter for diesel. Because excise taxes are fixed in peso per liter (not ad valorem), they compress the pass-through of proportional crude oil price changes more for gasoline than for diesel. Taxes (excise plus VAT) account for a larger share of the gasoline pump price, leaving a smaller share of the final price responsive to crude oil movements.

Second, diesel is predominantly a commercial fuel used in freight, public transport, agriculture, and power generation. Commercial users face more transparent cost-plus pricing and tighter competitive pressure to pass through input cost changes, compared with the retail gasoline market where oligopolistic firms (Petron, Shell, Chevron Philippines) may absorb some cost variation in margins. The higher diesel pass-through is consistent with a market where cost shocks are transmitted more mechanically.

Third, the DPC/IPC split should be interpreted differently across the two fuels. The CPI weight $w_{\text{fuel}} = 0.0401$ reflects household expenditure on transport fuels, which is primarily a gasoline concept. Diesel's effect on consumer prices operates largely through freight costs, public transport fares, electricity generation, and agricultural inputs—channels that are economically "direct" but not captured by the household fuel weight. The residual non-fuel component in the diesel specification therefore absorbs some transmission that is, in an economic sense, fuel-related. This means that the diesel DPC/IPC split understates the true fuel-related component and overstates the residual. For this reason, the diesel specification is better understood as providing evidence on the total CPI pass-through and fuel pass-through (PG and PC), with the DPC/IPC accounting partition interpreted with particular caution.

6 Extensions

6.1 Exchange rate channel

Table 3 reports the pass-through estimates from the four-variable SVAR that incorporates the peso-dollar exchange rate. The BIC again selects $p = 2$ for both fuel specifications. The fuel pass-through estimates are virtually identical to the baseline: gasoline PG reaches 0.49 at 12 months and diesel PG reaches 0.68 at 12 months, compared with 0.49 and 0.66 in the three-variable specification. This stability indicates that controlling for the exchange rate does not alter the estimated crude-to-retail fuel price transmission.

The exchange rate pass-through (FX-PT, defined as the ratio of the cumulative exchange rate response to the cumulative oil price response) is initially near zero at 6 months but rises to 0.07 at 36 months for both fuel specifications. This pattern is consistent with oil price innovations gradually depreciating the peso—plausibly through the current account channel, as higher oil import costs widen the trade deficit—and this depreciation feeding back into domestic prices.

Table 3: Oil price pass-through: Four-variable SVAR with exchange rate

Specification	Horizon	PG	PC	DPC	IPC	FX-PT
Gasoline						
	6 months	0.461	0.026	0.018	0.008	−0.005
	12 months	0.494	0.042	0.020	0.022	0.018
	18 months	0.502	0.056	0.020	0.036	0.042
	24 months	0.497	0.065	0.020	0.045	0.060
	36 months	0.481	0.066	0.019	0.046	0.069
Diesel						
	6 months	0.574	0.025	0.023	0.002	−0.004
	12 months	0.683	0.039	0.027	0.012	0.021
	18 months	0.742	0.053	0.030	0.023	0.049
	24 months	0.752	0.062	0.030	0.032	0.069
	36 months	0.699	0.063	0.028	0.035	0.071

Notes: FX-PT = exchange rate pass-through (cumulative exchange rate response / cumulative oil response). Ordering: oil $\rightarrow$ PHP/USD $\rightarrow$ fuel $\rightarrow$ CPI. VAR(2), $N = 298$.

The baseline and extended PG estimates are nearly identical, confirming that the exchange rate does not mediate the crude-to-retail fuel price channel. The impulse responses (not shown) indicate that the exchange rate initially appreciates slightly in response to an oil price innovation before depreciating over the medium run, consistent with the exchange rate pass-through turning positive at longer horizons.

6.2 Rice prices and the forecast error variance decomposition

A concern with the baseline three-variable SVAR is that the information set is sparse: oil price innovations may absorb variance that properly belongs to other macroeconomic drivers. To assess this, I extend the system to a four-variable SVAR that includes the rice CPI sub-index, ordered as oil $\rightarrow$ fuel $\rightarrow$ rice $\rightarrow$ CPI. The rice CPI series is available from 2012, yielding an estimation sample of $N = 154$ (January 2013 to October 2025).

Table 4 reports the results. The fuel pass-through is virtually unchanged: gasoline PG = 0.44 at 12 months (versus 0.44 in the three-variable specification on the same sample) and diesel PG = 0.68 (versus 0.68). The stability of PG to the inclusion of rice confirms that the fuel pass-through estimate is robust to expanding the information set.

The rice pass-through—defined as the ratio of the cumulative rice CPI response to the cumulative oil response—is 0.08 at 12 months and 0.17 at 36 months in the gasoline specification. This indicates that oil price innovations are associated with modest but persistent increases in rice price inflation, plausibly through transport and input cost channels.

Table 4: Oil price pass-through: Four-variable SVAR with rice prices

		4-variable (with rice)			3-variable (same sample)		
Specification	Horizon	PG	PC	IPC	PG	PC	IPC
Gasoline							
	6 months	0.446	0.026	0.009	0.449	0.027	0.009
	12 months	0.441	0.039	0.021	0.443	0.039	0.021
	24 months	0.389	0.046	0.030	0.393	0.047	0.032
	36 months	0.371	0.041	0.026	0.380	0.045	0.029
Diesel							
	6 months	0.611	0.022	−0.003	0.608	0.023	−0.001
	12 months	0.678	0.031	0.004	0.678	0.034	0.007
	24 months	0.637	0.042	0.016	0.641	0.046	0.020
	36 months	0.587	0.038	0.014	0.596	0.042	0.018

Notes: Ordering for 4-variable: oil $\rightarrow$ fuel $\rightarrow$ rice $\rightarrow$ CPI. Both specifications use VAR(2) with Cholesky identification. Estimation sample: $N = 154$ (Jan 2013–Oct 2025), determined by rice CPI availability. The 3-variable specification is re-estimated on the same sample for comparability. DPC omitted for space; IPC = PC $-$ DPC where DPC $= 0.0401\times$ PG.

A secondary result concerns the forecast error variance decomposition. In the three-variable system, oil innovations account for 44 percent of CPI forecast error variance at 12 months. Adding rice absorbs 15 percent of CPI variance—drawn primarily from the CPI-own component (which drops from 52 to 38 percent) rather than from the oil share (which remains at 44 percent). Table 5 reports the comparison. The diesel specification shows a larger redistribution, with the oil share declining from 44 to 37 percent and fuel innovations absorbing an additional 20 percent. These results suggest that the three-variable FEVD is not grossly inflating the oil share; rather, it was attributing rice-driven CPI variation to the CPI-own residual. As noted in Section 5, FEVD in small recursive systems is fragile and should be treated as descriptive rather than as evidence that oil is "the" dominant inflation driver.

Table 5: Forecast error variance decomposition for headline CPI: 3-variable versus 4-variable with rice

	3-variable			4-variable (with rice)			
Horizon	Oil	Fuel	CPI	Oil	Fuel	Rice	CPI
Gasoline specification							
$h = 12$	0.442	0.037	0.521	0.435	0.033	0.152	0.380
$h = 24$	0.478	0.037	0.485	0.462	0.036	0.154	0.348
$h = 36$	0.476	0.039	0.485	0.456	0.041	0.160	0.343
Diesel specification							
$h = 12$	0.442	0.037	0.521	0.370	0.204	0.187	0.239
$h = 24$	0.478	0.037	0.485	0.391	0.193	0.188	0.229
$h = 36$	0.476	0.039	0.485	0.383	0.198	0.195	0.224

Notes: Shares of forecast error variance of headline CPI at horizon h attributed to each Cholesky-identified innovation. 3-variable FEVD is re-estimated on the rice-available sample ($N = 154$) for comparability. Ordering for 4-variable: oil $\rightarrow$ fuel $\rightarrow$ rice $\rightarrow$ CPI.

As an alternative identification exercise, I implement a proxy-SVAR using Känzig (2021)'s OPEC announcement surprises as an external instrument for oil supply shocks. The first-stage F-statistics are 2.9 (gasoline) and 2.5 (diesel)—well below the conventional weak-instrument threshold of 10—so no strong quantitative conclusions can be drawn. The proxy-SVAR yields lower fuel pass-through estimates (gasoline PG = 0.37, diesel PG = 0.41 at 12 months), consistent with the Cholesky scheme capturing some demand-driven oil price variation. Full results and a discussion of instrument construction, relevance, and the direction of bias are reported in Appendix D.

7 Robustness

7.1 Monthly log differences with seasonal controls

The baseline specification uses 12-month log differences, which removes seasonality and renders the series stationary but also creates overlapping observations and smooths short-run dynamics. As a robustness check, I re-estimate the three-variable SVAR using month-on-month log changes (multiplied by 100), first residualizing all variables on a full set of monthly seasonal dummies. The BIC selects $p = 1$ for both fuel specifications.

Table 6 reports the results. The fuel pass-through estimates are broadly consistent with the baseline: gasoline PG = 0.47 and diesel PG = 0.56, compared with 0.49 and

0.66 at 12 months in the baseline. A notable difference is that the pass-through stabilizes much faster under monthly differences—essentially by horizon 6—reflecting the faster-decaying dynamics of month-on-month changes. The total CPI pass-through (PC) is lower at 0.02, and the residual component (IPC) is near zero. However, this apparent discrepancy with the baseline is largely an artifact of the VAR's restrictive dynamic structure rather than a genuine absence of non-fuel pass-through: when local projections are used on the same monthly data (Section 7.2), IPC recovers to 0.014–0.028 at 12–24 months as the cumulative responses are allowed to build gradually. These results suggest that the qualitative finding of substantial but incomplete fuel pass-through is not an artifact of the year-over-year transformation. However, the estimated magnitude of the residual non-fuel component is clearly specification-sensitive: the existence of a non-trivial residual is robust across methods, but its size depends on whether the VAR or LP estimator is used and on the data transformation. This should be kept in mind when interpreting the baseline IPC estimates.

Table 6: Robustness: Monthly log differences with seasonal controls

Specification	Horizon	PG	PC	DPC	IPC
Gasoline (VAR(1), $N = 307$)					
	6 months	0.472	0.020	0.019	0.001
	12 months	0.473	0.020	0.019	0.001
	36 months	0.473	0.020	0.019	0.001
Diesel (VAR(1), $N = 307$)					
	6 months	0.557	0.021	0.022	−0.001
	12 months	0.558	0.021	0.022	−0.001
	36 months	0.558	0.021	0.022	−0.001

Notes: All variables are month-on-month log changes $\times 100$, residualized on 11 monthly seasonal dummies. Cholesky identification with ordering: oil $\rightarrow$ fuel $\rightarrow$ CPI.

7.2 Local projections with HAC inference

The year-over-year transformation creates overlapping observations (consecutive months share 11 of 12 data points), inducing an MA(11) structure in the VAR residuals. The parametric Monte Carlo bootstrap used for the baseline confidence bands does not fully account for this dependence. To address this concern, I re-estimate the pass-through using Jordà (2005)'s local projection (LP) method with Newey-West (HAC) standard errors, using a bandwidth of at least 12 to account for the overlapping-observations autocorrelation.

At each horizon h, the cumulative response is estimated by projecting $\sum_{s=0}^{h} j_{t+s}$ on the structural oil shock (from the Cholesky decomposition of the VAR residual covariance), controlling for p lags. This approach does not impose the VAR's dynamic restrictions and the HAC correction provides honest inference under serial correlation.

Table 7 compares the VAR and LP estimates for the gasoline specification using year-over-year data. The LP estimates closely track the VAR: at 12 months, LP yields PG = 0.53 versus VAR PG = 0.49, and LP IPC = 0.018 versus VAR IPC = 0.023. At 36 months, LP gives slightly higher pass-through (PC = 0.072, IPC = 0.050) than the VAR (PC = 0.066, IPC = 0.048). The 95 percent HAC confidence intervals for the cumulative CPI response exclude zero at all horizons through 24 months.

More importantly, the LP approach resolves the apparent discrepancy with the monthly-difference VAR. When local projections are estimated on month-on-month (deseasoned) data, the residual component is no longer zero: gasoline IPC reaches 0.014 at 12 months and 0.028 at 24 months, compared with 0.001 in the monthly VAR (Table 6). The difference arises because LP allows cumulative responses to build gradually across horizons, whereas the monthly VAR's fast-decaying impulse responses flatten the CPI response by horizon 6. The LP monthly estimates are smaller than the year-over-year LP estimates—consistent with the year-over-year transformation capturing some low-frequency dynamics—but they confirm that the non-fuel component is economically non-trivial under both data transformations.

Table 7: Pass-through comparison: VAR versus local projections (gasoline specification)

		VAR (Cholesky)			Local projections (HAC)		
Data	Horizon	PG	PC	IPC	PG	PC	IPC
YoY							
	6 months	0.458	0.026	0.008	0.433	0.024	0.007
	12 months	0.488	0.042	0.023	0.527	0.039	0.018
	24 months	0.482	0.064	0.045	0.605	0.062	0.038
	36 months	0.464	0.066	0.048	0.543	0.072	0.050
Monthly							
	6 months	0.472	0.020	0.001	0.406	0.025	0.008
	12 months	0.473	0.020	0.001	0.436	0.032	0.014
	24 months	0.473	0.020	0.001	0.428	0.045	0.028
	36 months	0.473	0.020	0.001	0.309	0.036	0.023

Notes: LP estimates use Newey-West HAC standard errors with bandwidth $\max(h+1, 12)$ for YoY data and $h+1$ for monthly data. The Cholesky-identified oil shock from a VAR(2) is used as the regressor. Monthly data are residualized on 11 seasonal dummies. PG and PC are computed as cumulative response ratios using the VAR's cumulative oil self-response as the denominator. $N = 298$ (YoY), $N = 307$ (monthly).

7.3 Structural stability: pre- and post-TRAIN subsamples

The TRAIN Law (RA 10963, effective January 2018) imposed per-unit excise taxes on petroleum products, potentially altering the pricing regime. I split the estimation sample at January 2018 and re-estimate the baseline SVAR on each subsample. Table 8 reports the results. The pre-TRAIN subsample (2001–2017, $N = 202$) yields gasoline PG = 0.48 and diesel PG = 0.64 at 12 months. The post-TRAIN subsample (2018–2025, $N = 92$) yields gasoline PG = 0.53 and diesel PG = 0.74. The somewhat higher post-TRAIN estimates for diesel may reflect changes in the tax structure or the particular oil price episodes in the later period (the 2020 collapse and 2022 spike). Both subsamples yield qualitatively similar pass-through dynamics, suggesting that the pooled estimates are not averaging across fundamentally different regimes.

Table 8: Structural stability: Pre- and post-TRAIN subsamples

Period	Fuel	N	PG (12 mo)	PG (36 mo)
Pre-TRAIN (2001–2017)	Gasoline	202	0.480	0.500
	Diesel	202	0.642	0.771
Post-TRAIN (2018–2025)	Gasoline	92	0.525	0.450
	Diesel	92	0.742	0.611
Full sample (2001–2025)	Gasoline	296	0.488	0.464
	Diesel	296	0.664	0.646

Notes: All specifications use VAR(2) with Cholesky identification. The split date is January 2018 (TRAIN Law effective date).

Rolling 10-year estimates of PG at 12 months (not shown) fluctuate within a narrow band for gasoline (roughly 0.35–0.65) with no clear structural break around the TRAIN Law implementation or the COVID-19 pandemic.

7.4 Alternative price data: CEIC brand-level pump prices

As a further robustness check, I replace the DOE prevailing pump prices with an alternative price series constructed from brand-level daily retail prices sourced from CEIC. The CEIC data report high and low pump prices for each major brand (Petron, Shell, Caltex, Seaoil, Unioil, Flying V, Total, and others) for the National Capital Region, at daily frequency. I use the low price for each brand and compute a cross-brand average, then aggregate to monthly frequency by taking the within-month mean across all available trading days. The resulting series thus captures the average lower-bound retail price across brands, as opposed to the DOE's modal prevailing price.

The CEIC gasoline (RON 91) series begins in September 2013 and the regular diesel series in June 2005, yielding estimation samples of $N = 134$ and $N = 233$ respectively after the year-over-year transformation. The monthly correlation between the DOE and CEIC series is 0.97 for both gasoline and diesel in year-over-year log changes.

Table 9 reports the baseline three-variable SVAR estimates using the CEIC data. Diesel fuel pass-through reaches 0.56 at 12 months and 0.52 at 36 months, compared with 0.66 and 0.65 using DOE prevailing prices. Gasoline pass-through is 0.37 at 12 months (versus 0.49 with DOE data). The CEIC estimates are systematically lower by approximately 15–25 percent across both specifications. Two factors likely contribute to this gap: the shorter CEIC gasoline sample misses the pre-2014 period of large oil price movements (including the 2008 spike), and the CEIC "low price" concept measures the lower bound of each brand's retail range rather than the modal market price. The four-variable specification incorporating the exchange rate yields similar PG estimates (diesel PG = 0.59 at 12 months, gasoline PG = 0.41), confirming the stability of the fuel pass-through to the inclusion of the exchange rate regardless of the price source.

Table 9: Pass-through estimates using CEIC brand-average pump prices

		CEIC brand-average			DOE prevailing price		
Specification	Horizon	PG	PC	IPC	PG	PC	IPC
Gasoline							
	6 months	0.336	0.010	-0.003	0.458	0.026	0.008
	12 months	0.365	0.015	0.000	0.488	0.042	0.023
	24 months	0.364	0.020	0.005	0.482	0.064	0.045
	36 months	0.358	0.020	0.006	0.464	0.066	0.048
Diesel							
	6 months	0.440	0.015	-0.003	0.566	0.025	0.003
	12 months	0.556	0.025	0.003	0.664	0.040	0.014
	24 months	0.570	0.040	0.017	0.705	0.063	0.034
	36 months	0.515	0.039	0.019	0.646	0.064	0.038

Notes: Both specifications use VAR(2) with Cholesky identification. CEIC: gasoline $N = 134$ (Oct 2014–Oct 2025), diesel $N = 233$ (Jun 2006–Oct 2025). DOE: $N = 298$ (Jan 2001–Oct 2025). CEIC prices are the monthly mean of daily cross-brand low prices in the NCR. DPC is omitted for brevity; IPC = PC $-$ DPC where DPC = $0.0401\times$ PG.

The CEIC results confirm the qualitative conclusions while suggesting that the DOE-based point estimates may represent an upper bound. Subsample analysis on the CEIC diesel data shows stable PG across pre- and post-TRAIN periods (PG@12 = 0.57 in both), reinforcing the structural stability finding.

7.5 Sensitivity of the DPC/IPC split to the fuel weight

The DPC calculation uses $w_{\text{fuel}} = 0.0401$ from the 2018-based CPI basket throughout the sample, but the effective weight may have differed under the earlier 2006-based and 2012-based CPI baskets for which the disaggregated PCOICOP 07.2.2 weight is not available.[1] To assess sensitivity, I recalculate the decomposition under alternative weights spanning a plausible range (0.030 to 0.060). At $h = 36$ in the gasoline specification, the IPC share of total pass-through ranges from 58 percent ($w = 0.060$) to 79 percent ($w = 0.030$). At $h = 12$, the range is wider: 31 percent ($w = 0.060$) to 65 percent ($w = 0.030$). The qualitative conclusion that the residual component accounts

[1]The 2012-based CPI publishes weights for “Electricity, Gas and Other Fuels” (COICOP 04.5, weight 7.44) and its sub-items, but these cover housing energy, not transport fuels (COICOP 07.2.2). The transport fuel weight under the 2012 basket could not be obtained from available PSA publications.

for the larger share at long horizons is robust across this range; at shorter horizons, the split is more sensitive, and the claim of residual dominance depends on the assumed weight. The 2018-based weight of 0.0401 places the IPC share at 54 percent ($h = 12$) and 72 percent ($h = 36$).

8 Conclusion

This paper estimates the pass-through of Cholesky-identified oil price innovations to Philippine fuel prices and headline CPI, and decomposes the total CPI pass-through into a fuel-basket component and a residual non-fuel component using an accounting identity derived from the CPI expenditure structure. Using 25 years of monthly retail pump price data from the DOE, I find fuel pass-through of 46–49 percent for gasoline and 65–71 percent for diesel, total long-run CPI pass-through of approximately 6.5 percent, and a residual non-fuel component that is economically non-trivial across all specifications. In the gasoline specification, the residual accounts for roughly 72 percent of the total CPI response at 36 months; in the diesel specification the split is more balanced, though the diesel accounting partition is less cleanly interpretable. Diesel pass-through is consistently higher than gasoline pass-through across all specifications, plausibly reflecting differences in excise tax structure and diesel's role in commercial freight and transport. The fuel pass-through estimates are robust to the inclusion of the exchange rate or rice prices, to local projections with HAC inference, to monthly log differences, to pre- versus post-TRAIN Law subsamples, and to CEIC brand-level data. The estimated magnitude of the residual component, however, is somewhat sensitive to the estimation method: the monthly-difference VAR yields a near-zero residual, while local projections on the same data recover a non-trivial one, suggesting that the *existence* of the residual is robust but its precise size should be interpreted with caution.

The total pass-through (PC $\approx$ 6.5 percent at 36 months) implies that a 10 percentage point increase in year-over-year oil price growth is associated with an increase of about 0.65 percentage points in headline CPI inflation—economically meaningful in the context of BSP's 2–4 percent target band. The accounting residual is broadly consistent with Conflitti and Luciani (2019)'s finding that the common (macroeconomic) component of oil pass-through exceeds the item-specific energy-intensity component, though the residual estimated here is not a common factor—it is whatever remains of the CPI response after removing the mechanically implied fuel-basket contribution. This paper does not model inflation expectations, wages, real activity, or the policy rate, so it cannot establish whether the residual reflects persistent cost-push dynamics or whether a specific monetary policy response is warranted.

Several limitations should be noted. First, the Cholesky identification does not isolate a pure oil supply shock; the impulse responses should be interpreted as dynamic associations conditioned on a recursive ordering, not as causal effects of an exogenous supply disturbance. The exploratory proxy-SVAR (Appendix D) yields lower estimates, consistent with the Cholesky scheme capturing some demand-driven variation, but the weak instrument ($F \approx 3$) prevents a definitive comparison. Second, the pump prices are for Metro Manila only, while CPI is national; this geographic mismatch may attenuate the estimated pass-through (Appendix B). Third, the fixed $w_{\text{fuel}} = 0.0401$ weight from the 2018 CPI basket may have differed in earlier years; sensitivity analysis across a plausible range of weights confirms the qualitative dominance of the residual component at long horizons. Fourth, the residual non-fuel component is defined as an accounting residual and should not be read as a separately identified structural channel.

Future work should pursue stronger external instruments for oil supply identification (e.g., geopolitical disruption measures or Dubai futures surprises), threshold SVARs to test regime dependence, disaggregate analysis of CPI sub-indices to shed light on the composition of the residual, and explicit modeling of BSP's monetary policy reaction function.

References

Aastveit, K. A., Bjørnland, H. C., & Cross, J. L. (2023). Inflation expectations and the pass-through of oil prices. *Review of Economics and Statistics*, 105(3), 733–743.

Broer, T., Kramer, J., & Mitman, K. (2025). The distributional effects of oil shocks. *European Economic Review*, 173, 851–889.

Conflitti, C., & Luciani, M. (2019). Oil price pass-through into core inflation. *The Energy Journal*, 40(6), 221–247.

Corsello, F., & Tagliabracci, A. (2023). Assessing the pass-through of energy prices to inflation in the euro area. *Bank of Italy Occasional Papers* No. 745.

Ding, Y., Li, Y., & Song, Z. (2023). The oil price–inflation nexus: The exchange rate pass-through effect. *Energy Economics*, 125, 106861.

Gagliardone, L., & Gertler, M. (2023). Oil prices, monetary policy and inflation surges. NBER Working Paper No. 31263.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1), 161–182.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review*, 111(4), 1092–1125.

Kilian, L. (2024). How to construct monthly VAR proxies based on daily surprises in futures markets. *Journal of Economic Dynamics and Control*, 168, 104966.

López, L., Odendahl, F., Párraga Rodríguez, S., & Silgado-Gómez, E. (2026). The pass-through to inflation of gas price shocks. *Journal of Economic Dynamics and Control*, 182, 105218.

Mertens, K., & Ravn, M. O. (2013). The dynamic effects of personal and corporate income tax changes in the United States. *American Economic Review*, 103(4), 1212–1247.

Ndou, E., & Gumata, N. (2025). Oil price passthrough to consumer price inflation in South Africa: The role of the inflation environment. *Journal of Applied Economics*.

Stock, J. H., & Watson, M. W. (2012). Disentangling the channels of the 2007–2009 recession. NBER Working Paper No. 18094.

Taylor, J. B. (2000). Low inflation, pass-through, and the pricing power of firms. *European Economic Review*, 44(7), 1389–1408.

Yilmazkuday, H. (2021). Oil price pass-through into consumer prices: Evidence from U.S. weekly data. *Journal of International Money and Finance*, 119, 102494.

A Formal derivation of the pass-through decomposition

This appendix derives the accounting decomposition of CPI pass-through into fuel-basket and residual non-fuel components from the CPI expenditure identity.

A.1 The CPI as a weighted average

The headline CPI is a Laspeyres-type index constructed as a weighted average of component price indices. In log levels, the CPI can be approximated as:

$$\ln P_t \approx w_f \ln P_t^f + (1 - w_f) \ln P_t^{nf} \tag{9}$$

where P_t^f is the fuel component price index, P_t^{nf} is the non-fuel composite price index, and w_f is the CPI expenditure weight on fuel. Taking year-over-year log differences:

$$\Delta \ln P_t \approx w_f \Delta \ln P_t^f + (1 - w_f) \Delta \ln P_t^{nf} \tag{10}$$

where Δ denotes the 12-month log difference.

A.2 From the CPI identity to the pass-through decomposition

Consider the impulse response of headline CPI inflation to a Cholesky-identified oil price innovation at horizon h. From equation (10), the response can be written as:

$$\mathrm{IRF}(\text{oil}, \text{cpi}, h) \approx w_f \cdot \mathrm{IRF}(\text{oil}, \text{fuel CPI}, h) + (1 - w_f) \cdot \mathrm{IRF}(\text{oil}, \text{non-fuel CPI}, h) \tag{11}$$

If the fuel component of CPI tracks pump prices one-for-one (i.e., $\Delta \ln P_t^f \approx \Delta \ln \text{fuel}_t$), then the first term on the right-hand side can be proxied by $w_f \cdot \mathrm{IRF}(\text{oil}, \text{fuel}, h)$, where "fuel" denotes the retail pump price in the VAR. Summing over horizons and dividing by the cumulative oil self-response yields the pass-through decomposition:

$$\underbrace{\mathrm{PC}(h)}_{\text{total CPI PT}} \approx \underbrace{w_f \cdot \mathrm{PG}(h)}_{\text{DPC: fuel-basket}} + \underbrace{\mathrm{PC}(h) - w_f \cdot \mathrm{PG}(h)}_{\text{IPC: residual non-fuel}} \tag{12}$$

where $\mathrm{PG}(h) = \mathrm{CIR}(\text{oil}, \text{fuel}, h)/\mathrm{CIR}(\text{oil}, \text{oil}, h)$ is the fuel pass-through and $\mathrm{PC}(h) = \mathrm{CIR}(\text{oil}, \text{cpi}, h)/\mathrm{CIR}(\text{oil}, \text{oil}, h)$ is the total CPI pass-through.

A.3 Assumptions and caveats

This decomposition requires three assumptions:

1. *Fixed weights.* The CPI weight w_f is treated as constant across the sample. In practice, the PSA has re-based the CPI three times during the sample period (2000, 2006, 2012, 2018 bases), and the transport fuel weight may have changed. Section 7 reports sensitivity of the DPC/IPC split to alternative weight values spanning 0.030–0.060.

2. *Fuel CPI tracks pump prices.* The fuel component of CPI is assumed to move one-for-one with the pump price series used in the VAR. In practice, the PSA's fuel CPI sub-index (PCOICOP 07.2.2) is constructed from a broader set of fuel products and outlets than the single Metro Manila prevailing price. Any wedge between the two—due to geographic averaging, product mix, or sampling methodology—introduces measurement error in DPC.

3. *Linearity of the log approximation.* The weighted-average decomposition holds exactly for a geometric index and approximately for a Laspeyres index when price changes are small. For the large oil price movements in the sample (e.g., -100 to $+88$ percentage points in year-over-year terms), the approximation error may be non-trivial for fuel prices but is damped when translated to CPI because w_f is small (0.04).

The residual non-fuel component (IPC) captures everything not attributable to the fuel basket: transport cost pass-through, production input costs, wage adjustments, expectation effects, exchange rate effects, and any measurement error from assumptions 1–2 above. It is an accounting residual, not a separately identified structural object.

B Measurement issues

This appendix documents three measurement concerns and the sensitivity checks conducted to address them.

B.1 Geographic mismatch: Metro Manila pump prices versus national CPI

The pump prices used in the VAR are for Metro Manila, while headline CPI is a national aggregate. Metro Manila is the price-leading market and receives heavy weight in the PSA's CPI sampling, but provincial fuel prices may differ due to transport costs and local market conditions. This mismatch likely attenuates the estimated fuel pass-through (PG) because the VAR's fuel variable is noisier as a proxy for the national fuel price relevant to national CPI. A Manila-area CPI sub-index would provide a cleaner test, but disaggregated regional CPI data at monthly frequency are not available for the full sample.

B.2 Fixed CPI weight across the sample

The decomposition uses $w_{\text{fuel}} = 0.0401$ from the 2018-based CPI basket throughout the sample. The PSA does not publish the PCOICOP 07.2.2 (transport fuel) weight for earlier CPI base years. The broader "Electricity, Gas and Other Fuels" category (COICOP 04.5) had a weight of 7.44 percent in the 2012-based CPI, but this covers housing energy, not transport fuels, and is not comparable.

To assess sensitivity, the DPC/IPC decomposition is recalculated under alternative weights spanning 0.030 to 0.060. At $h = 36$ in the gasoline specification, the IPC share of total pass-through ranges from 58 percent ($w = 0.060$) to 79 percent ($w = 0.030$). The qualitative conclusion that the residual non-fuel component accounts for the larger share at long horizons is robust across this range.

B.3 Pre- and post-TRAIN decomposition

The TRAIN Law (effective January 2018) changed the excise tax structure on petroleum products, potentially altering the fuel weight's relevance. Table 8 reports subsample pass-through estimates split at January 2018. Both subsamples yield qualitatively similar pass-through dynamics, and the fuel pass-through is somewhat higher post-TRAIN for diesel. The sensitivity of the DPC/IPC split to the TRAIN reform is modest: pre-TRAIN gasoline DPC at 12 months is $0.0401 \times 0.480 = 0.019$ versus post-TRAIN DPC $= 0.0401 \times 0.525 = 0.021$. The residual non-fuel component remains the larger share in both periods.

C Additional results

Table 10: Forecast error variance decomposition (gasoline pump price specification)

Variable / Horizon	Oil innovation	Gasoline innovation	CPI innovation
Headline CPI			
$h = 12$	0.427	0.045	0.529
$h = 24$	0.463	0.044	0.493
$h = 36$	0.461	0.046	0.493
Gasoline pump			
$h = 12$	0.834	0.130	0.036
$h = 24$	0.786	0.140	0.074
$h = 36$	0.787	0.140	0.074
Oil price			
$h = 12$	0.880	0.093	0.027
$h = 24$	0.840	0.113	0.047
$h = 36$	0.840	0.113	0.047

Notes: Shares of forecast error variance at horizon h months attributed to each Cholesky-identified innovation. VAR(2) with ordering: oil $\rightarrow$ gasoline $\rightarrow$ CPI, $N = 298$. Shares are conditional on the recursive identification and the three-variable information set.

Table 11: Granger causality tests (gasoline pump price specification)

Null hypothesis	F-statistic	p-value
Oil $\not\rightarrow$ Gasoline pump	62.479	0.000***
Oil $\not\rightarrow$ Headline CPI	5.283	0.005***
Gasoline pump $\not\rightarrow$ Oil	5.749	0.003***
Gasoline pump $\not\rightarrow$ Headline CPI	1.934	0.145
Headline CPI $\not\rightarrow$ Oil	4.836	0.008***
Headline CPI $\not\rightarrow$ Gasoline pump	3.489	0.031**

Notes: F-tests for Granger non-causality based on VAR(2). $\not\rightarrow$ denotes "does not Granger-cause." ***, ** indicate significance at the 1% and 5% levels. The reverse causation from gasoline to oil likely reflects common information content rather than genuine feedback from Philippine pump prices to world oil markets.

D Proxy-SVAR with external instruments

The baseline Cholesky identification does not isolate a pure oil supply shock. To provide an alternative structural lens, I implement a proxy-SVAR following Mertens and Ravn (2013) and Stock and Watson (2012), using Känzig (2021)'s OPEC announcement surprises as an external instrument for oil supply shocks. This is the most credible high-frequency instrument available for oil supply identification in the recent literature.

Instrument construction. The instrument is the first principal component of oil futures price changes (front month through 12-month contracts) in a narrow window around OPEC meeting dates. The raw data contain 123 OPEC events from July 1983 to December 2019. I aggregate to monthly frequency by summing within-month surprises, assigning zero to non-event months. In the estimation sample (January 2001–October 2025), 68 months have non-zero instrument values.

Estimation. The procedure follows Mertens and Ravn (2013): the reduced-form VAR residuals are projected onto the instrument via two-stage least squares to recover the structural impact vector, and impulse responses are propagated through the VAR's moving-average representation. Confidence intervals are computed by wild bootstrap with 500 replications, using the Rademacher distribution for the bootstrap weights.

Instrument relevance. The first-stage F-statistic is 2.9 for the gasoline specification and 2.5 for diesel—well below the conventional weak-instrument threshold of 10 (Stock and Watson, 2012). Two features of the Philippine application contribute to the weak first stage. First, the instrument is constructed from WTI futures contracts, while the Philippine VAR uses Dubai crude prices; the correlation between WTI futures surprises and Dubai spot price innovations is positive but imperfect. Second, OPEC supply decisions may be less informative for Asian benchmark prices than for WTI. Future work could improve instrument relevance by using geopolitical oil supply disruption indicators (e.g., Kilian, 2024) or by constructing high-frequency Dubai futures surprises around OPEC announcements, though the latter requires intraday data that are not readily available.

Results. Table 12 reports the results. The IV-identified contemporaneous impact of oil on gasoline is 0.063 (relative to 0.245 under Cholesky), and the resulting fuel pass-through is lower: gasoline PG = 0.37 at 12 months (versus 0.49 under Cholesky) and diesel PG = 0.41 at 12 months (versus 0.66). Under weak instruments, IV estimates are biased toward the OLS (reduced-form) estimate in just-identified systems, so the lower proxy-SVAR pass-through is unlikely to be an artifact of weak-instrument bias toward zero. The direction of the difference—lower pass-through under IV identification—is consistent with the Cholesky scheme capturing some demand-driven oil price variation that is correlated with global activity and therefore with Philippine inflation through channels other than oil supply.

However, given the weak first stage, no strong quantitative conclusions should be drawn. The exercise is useful primarily in establishing that the baseline Cholesky identification is not the only possible reading of the data, and that the direction of any overstatement is toward higher pass-through. The qualitative finding of substantial but incomplete fuel pass-through is robust across both identification schemes.

Table 12: Oil price pass-through: Proxy-SVAR with Känzig instrument

Specification	Horizon	PG	PC	DPC	IPC
Gasoline ($F = 2.9$)					
	6 months	0.348	0.040	0.014	0.026
	12 months	0.368	0.055	0.015	0.041
	18 months	0.352	0.069	0.014	0.055
	24 months	0.325	0.078	0.013	0.065
	36 months	0.297	0.079	0.012	0.067
Diesel ($F = 2.5$)					
	6 months	0.370	0.038	0.015	0.023
	12 months	0.412	0.051	0.017	0.035
	18 months	0.387	0.062	0.016	0.047
	24 months	0.333	0.068	0.013	0.054
	36 months	0.263	0.066	0.011	0.055

Notes: Proxy-SVAR using Känzig (2021) OPEC announcement surprises as external instrument. F-statistics report the first-stage relevance of the instrument for the reduced-form oil residual. VAR(2), $N = 296$.

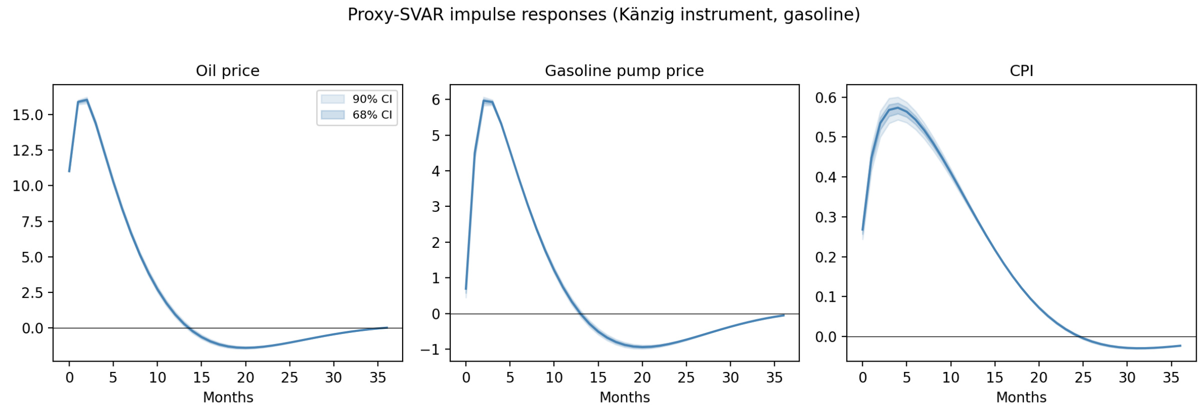

Figure 5: Proxy-SVAR impulse responses (gasoline specification; exploratory, weak instrument). Light shading: 90 percent wild bootstrap confidence band. Dark shading: 68 percent band.

Given the weak first stage, no strong quantitative conclusions should be drawn; the exercise is useful primarily in establishing that the baseline identification is not the only possible reading of the data.