INSTITUTO DE ESTUDIOS AVANZADOS EN DESARROLLO

inesad

SMOG, NOT LIGHTS.
TRACKING BOLIVIA'S 2024–2026 RECESSION FROM SPACE

by:

Werner Hernani-Limarino

Ahmed Eid

Series of Working Papers on Development

No. 6/2026

La Paz, may 2026

INSTITUTO DE ESTUDIOS AVANZADOS EN DESARROLLO

Smog, not lights.
Tracking Bolivia's 2024–2026 recession from space

Werner Hernani-Limarino*

Ahmed Eid

Abstract

Nighttime lights are the standard satellite proxy for economic activity, but they fail to detect recessions. On a 2020–2024 panel of nine Bolivian departments, VIIRS lights match the sign of PIB growth 78 percent of the time during expansions but only 39 percent during contractions —worse than coin-flip. The reason is structural: lit infrastructure is a stock that does not contract when activity falls, while fossil-fuel combustion is a flow that does. Tropospheric NO2, a neardirect measure of combustion retrieved daily by ESA's Sentinel-5P satellite, maintains 72 percent sign concordance during contractions. We exploit this asymmetry to build a monthly recession indicator from TROPOMI NO2 over eleven Bolivian metropolitan areas, July 2018 through April 2026. The indicator peaked in June 2023; INE quarterly PIB first registered negative year-on-year growth five quarters later, in 2024Q4. Through November 2025 the satellite and INE agree to within one log point. A pre-break SARIMA trained through November 2025 projects a 2026 PIB contraction of approximately 2 percent, consistent with the IMF (−3.3 percent), World Bank (−3.2), and IIF (−4) forecasts, though the confidence interval is wide. The centered-12-month March 2026 endpoint becomes fully observable from actual data in October 2026, and we commit to a revisit at that date. The indicator is reproducible from free Copernicus data with a two-month publication lag.

JEL: C32, E01, E32, O11, O54, Q53.
Keywords: TROPOMI, nitrogen dioxide, Sentinel-5P, coincident indicator, recession dating, Bolivia, SARIMA, early warning.

* All opinions are the authors' own and do not reflect the views of any affiliated institution. All remaining errors are ours. Contact: wernerhl@gmail.com.

Resumen

Las luces nocturnas son el indicador satélite estándar de la actividad económica, pero no logran detectar las recesiones. En un panel de nueve departamentos bolivianos correspondiente al periodo 2020-2024, las luces del VIIRS coinciden con el signo del crecimiento del PIB el 78 % de las veces durante las fases de expansión, pero solo el 39 % durante las de contracción —un resultado peor que el de lanzar una moneda al aire— La razón es estructural: la infraestructura iluminada es un stock que no se contrae cuando la actividad cae, mientras que la combustión de combustibles fósiles es un flujo que sí lo hace. El NO_2 troposférico, una medida casi directa de la combustión que se obtiene diariamente mediante el satélite Sentinel-5P de la ESA, mantiene una concordancia de signo del 72 % durante las contracciones. Aprovechamos esta asimetría para construir un indicador mensual de recesión a partir del NO_2 de TROPOMI en once áreas metropolitanas bolivianas, desde julio de 2018 hasta abril de 2026. El indicador alcanzó su punto máximo en junio de 2023; el PIB trimestral del INE registró por primera vez un crecimiento interanual negativo cinco trimestres después, en el cuarto trimestre de 2024. Hasta noviembre de 2025, el satélite y el INE coinciden con un margen de un punto logarítmico. Un modelo SARIMA pre-ruptura entrenado hasta noviembre de 2025 proyecta una contracción del PIB en 2026 de aproximadamente el 2 por ciento, en consonancia con las previsiones del FMI (−3,3 por ciento), el Banco Mundial (−3,2) y el IIF (−4), aunque el intervalo de confianza es amplio. El punto final centrado en 12 meses de marzo de 2026 será totalmente observable a partir de los datos reales en octubre de 2026, y nos comprometemos a revisarlo en esa fecha. El indicador se puede obtener a partir de datos gratuitos de Copernicus, con un retraso de dos meses en su publicación.

Código JEL: C32, E01, E32, O11, O54, Q53.

Palabras clave: TROPOMI, dióxido de nitrógeno, Sentinel-5P, indicador coincidente, datación de recesiones, Bolivia, SARIMA, alerta temprana.

1 Introduction

Bolivia's first non-COVID annual contraction in approximately four decades arrived on 21 April 2026. The National Statistics Office (INE) reported a full-year real PIB contraction of 1.58 percent for 2025 — the first non-pandemic annual contraction since the 1985 stabilization to exceed one percentage point. The 2026 outlook is worse: the IMF projects -3.3 percent, the World Bank -3.2 percent, the Institute of International Finance approximately -4 percent. The Paz government's reformulated 2026 budget projects -1.28 percent. The World Bank has extended its recession forecast into 2027.

The depth and dynamics of the contraction need to be measured at higher frequency than INE delivers, and from a vantage independent of INE itself. INE quarterly PIB arrives with a 60- to 90-day publication lag. The experimental monthly Índice Global de Actividad Económica (IGAE), launched in March 2026 under IMF technical assistance, will not have twenty-four monthly observations until March 2028. Departmental annual PIB arrives twelve to eighteen months after the calendar year closes. The spread across external forecasters — -3.3, -3.2, -4, -1.28 for the same year — is itself a measurement problem. Bolivia is the kind of low-statistical-capacity, fiscal-stress, regime-transition environment that has motivated the broader literature on satellite-based proxies for economic activity.

A growing literature uses satellite observations to construct such proxies. The foundational reference is Henderson et al. (2012), which establishes the cross-country measurement-error framework on DMSP nighttime lights. Donaldson and Storeygard (2016) surveyed the field. The VIIRS-era extensions in Beyer et al. (2022) estimate a quarterly nighttime-lights elasticity of 1.55 for emerging-market economies. Martinez (2022) uses lights to identify autocratic overstatement. Gibson et al. (2021) document that neither DMSP nor VIIRS captures rural low-density activity — a result that motivates the metropolitan focus of the analysis. Tropospheric NO_2 entered more recently: Bauwens et al. (2020) and Liu et al. (2020) document COVID-era drops of 20 to 48 percent across Chinese and European metros within weeks of lockdown imposition. Bricongne et al. (2021) use TROPOMI NO_2 to nowcast industrial production across thirty-three economies. Ezran et al. (2023) find NO_2 outperforms lights for short-term fluctuations. None covers Latin America at metropolitan resolution on a single low-statistical-capacity economy outside crisis windows.

The Bolivian baseline is Bolívar and Cuba (2024), which uses VIIRS nighttime lights but not NO_2. That work narrows the official IGAE release lag using satellite features and finds luminosity is the single most important predictor. It does not use NO_2, does not extend to the eight smaller departmental capitals, and does not benchmark against subnational PIB. The present analysis fills this gap.

Four findings emerge. First, NO_2 tracks Bolivian departmental PIB in both expansion and contraction; VIIRS tracks expansion only. Second, through November 2025 the satellite indicator and INE agree to within one log point. Third, the indicator leads INE on the recession turning point by approximately five quarters at the national-peak level and two to three quarters at the metropolitan first-negative level. Fourth, a pre-break SARIMA projects a 2026 PIB growth rate of approximately -2 percent with a large confidence interval, consistent with the multilateral consensus. The first verifiable test of the forecast arrives in October 2026.

The rest of the paper is organized as follows. Section 2 documents Bolivia's macroeconomic deterioration. Section 3 describes the three satellite and administrative series. Section 4 establishes the NO_2 indicator and the regime asymmetry with VIIRS. Section 5 applies the indicator to the

2024–2026 recession. Section 6 reports robustness exercises. Section 7 concludes.

2 Bolivia 2014–2026

Bolivia's macroeconomic trajectory since 2014 divides into two acts. The first is a decade of slow deterioration: natural-gas production peaked at roughly 60 million cubic metres per day in 2014 and fell to approximately 30 by end-2024, dragging down export revenues, fiscal balances, and international reserves. The government masked the decline through reserve depletion, monetary financing of the treasury, and a currency peg held nominally constant at 6.96 bolivianos per dollar. Between the 2014 commodity peak and end-2024, international reserves fell from $14.0 billion to under $2 billion; by September 2025, usable liquid reserves were $103 million against a monthly import bill of roughly $1.2 billion. Inflation in the official CPI remained in single digits even as imported-goods shortages and fuel queues were binding across the country.

The second act arrived abruptly. The November 2025 inauguration of Rodrigo Paz ended two decades of MAS rule. In December, the new government eliminated fuel subsidies overnight (Decree 5503: gasoline +86 percent, diesel +162 percent), secured multilateral financing of more than $8 billion, and watched sovereign spreads compress from 2002 to 466 basis points by mid-February 2026. The subsidy removal triggered a general strike by the COB from December 22 and highway blockades through January–February 2026, compounding the contractionary impulse. By Q4 2025, oil-and-gas extraction was down 10.51 percent year-on-year, construction 13.57 percent, and commerce 6.05 percent. Agriculture expanded 12.19 percent in the same quarter — the one major sector the NO_2 indicator does not see directly.

The 1988–2024 historical PIB series, constructed by chain-linking the 2017-base departmental series to the 1990-base series following the IMF Stanger (2025) empalme methodology, places the current contraction in perspective. Two contractions appear in 36 years: COVID-19 in 2020 (-12.7 percent) and 2024 (-1.1 percent). The 2025 preliminary of -1.58 percent is the third, and the IMF's 2026 projection of -3.3 percent would complete a three-consecutive-contraction sequence — Bolivia's first since the early-1980s hyperinflation, before the 1985 stabilization. Figure 1 shows the full series.

Figure 1. Bolivia national real PIB, 1988–2026.

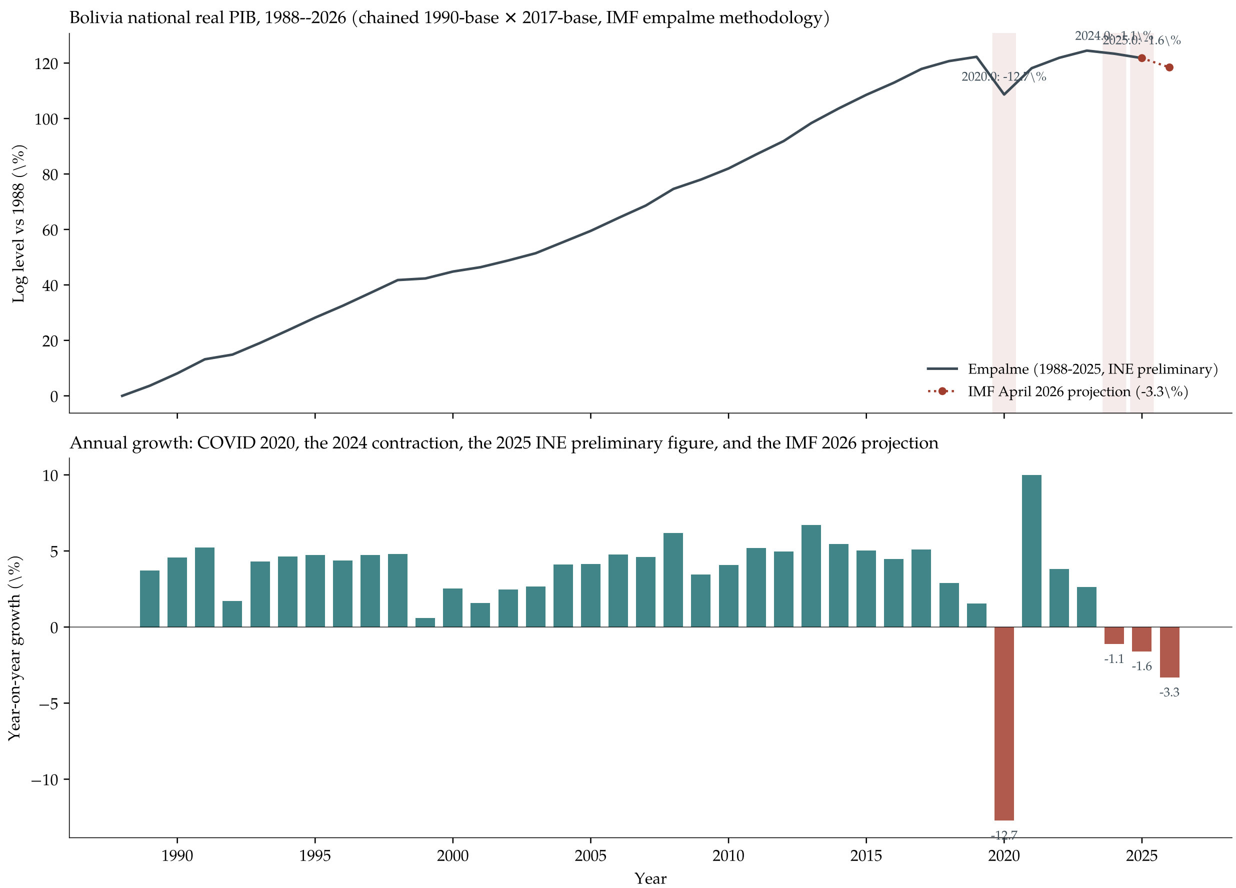

Notes: Chained 1990-base × 2017-base series following the IMF Stanger (2025) empalme methodology. Top: log-level index relative to 1988. Bottom: year-on-year growth. Slate solid: empalme through 2024. Rust dotted: IMF April 2026 projection of −3.3 percent. The 2025 figure (−1.58 percent) is the INE preliminary release of 21 April 2026. Sources: INE Bolivia; IMF April 2026 WEO.

Statistical-agency independence concerns apply to Bolivia in this period. The Fund's May 2025 Article IV consultation called the policy mix "unsustainable" even as INE growth figures remained positive. The priors from Fernald et al. (2021) on manipulation in fixed-regime economies with rapid reserve loss, from Michalski and Stoltz (2013) on strategic misreporting under negative-NFA pressure, and from Martinez (2022) on autocratic overstatement all point in the same direction. The TROPOMI indicator constructed in this paper is independent of INE by construction; whether the gap between the two reflects manipulation or measurement-error differences of other kinds is addressed in Section 6.

3 Data

Three series anchor the analysis: tropospheric NO_2 from two satellite instruments (OMI and TROPOMI), nighttime lights from VIIRS, and official real PIB from INE. Figure 2 plots all three as log deviations from calendar-year 2019. The series co-move tightly through 2019. After 2020, they diverge: NO_2 peaks in June 2023, declines through 2024–2025, and drops sharply in the post-November 2025 wet-season months. VIIRS continues rising through 2024 with no detectable contraction signal. INE registers a small contraction in 2024 (-1.1 percent) and a preliminary -1.58 percent in 2025. The divergence between the two satellite series after 2020 — NO_2 falling while VIIRS rises — is the empirical pattern that organizes Section 4.

Figure 2. Tropospheric NO_2, VIIRS sum-of-lights, and INE national real PIB. Bolivia, 2005–2026.

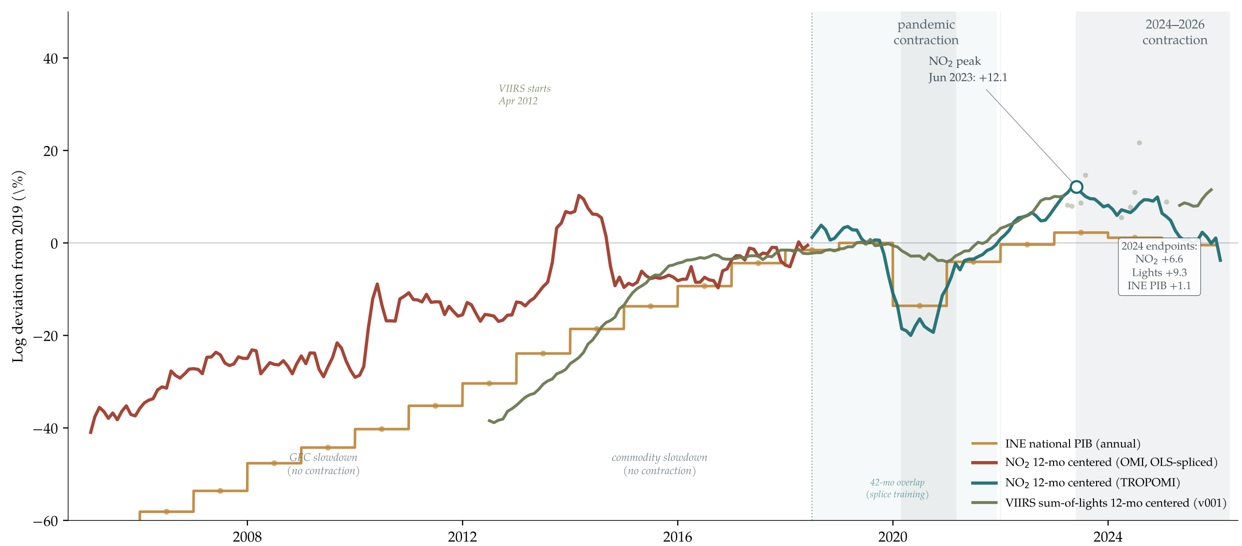

Notes: NO_2: population-weighted aggregate of La Paz–El Alto, Cochabamba, and Santa Cruz, 12-month centered, log deviation from 2019. OMI segment (rust, October 2004 – June 2018) bias-corrected via per-ROI OLS regression on the 2018-07 to 2021-12 overlap (teal band, 41 months); TROPOMI segment (teal, July 2018 – April 2026) direct. VIIRS (sage): VNP46A2 v001 with documented gaps in 2023 and post-2024-Q3 (small markers: available raw months). INE PIB (ochre step): annual, 2025 preliminary. Recession shading: COVID 2020Q2–2021Q1 and 2024–2026 from NO_2 peak June 2023 forward. Sources: Copernicus Sentinel-5P; NASA OMI OMNO2d; NASA VIIRS VNP46A2; INE Bolivia.

3.1 VIIRS sum-of-lights

The comparison nighttime-lights series comes from the Visible Infrared Imaging Radiometer Suite Day-Night Band on Suomi-NPP, which provides 500 m gridded radiance at approximately 1:30 local time. We use the BRDF- and lunar-irradiance-corrected monthly composite VNP46A2 v001, summed over the same eleven metropolitan ROIs as TROPOMI (Table 2). Coverage runs from April 2012 through August 2024, after which two problems interrupt the record. The v001 production pipeline failed to generate composites for seven months in 2023 and three months in 2024 — a processing failure resolved in the v002 reprocessing, which was not yet available at the time of writing.[1] Separately, a May 2024 Suomi-NPP GPS anomaly suspended Level-2 production through November. Together, these gaps prevent the computation of a 12-month centered VIIRS series for any month in 2023 or 2024 — precisely the window where recession dynamics matter most.

[1] NASA initiated the v002 reprocessing in March 2025. A v002-based extraction is in progress.

3.2 Tropospheric NO_2

Tropospheric NO_2 arises predominantly from fossil-fuel combustion: vehicle engines, industrial furnaces, and thermal power plants oxidize atmospheric N_2 at high temperature, producing NO that rapidly converts to NO_2. The atmospheric lifetime is hours to one day, keeping the signal spatially local — a satellite retrieval reflects emissions integrated over the preceding roughly twenty-four hours within the column footprint. This short lifetime is what makes NO_2 useful as an activity proxy: unlike CO_2, which mixes globally, NO_2 stays close to its source.

Two instruments contribute to the record. The Ozone Monitoring Instrument (OMI) on NASA's Aura satellite has been retrieving tropospheric NO_2 at 13×24 km native resolution since October 2004. We use the OMNO2d Level-3 product at 0.25° with cloud-radiance-fraction filter at 0.30, excluding post-2007 row-anomaly cross-track positions. OMI's coarse resolution limits the extraction to the three largest Bolivian metropolitan areas (La Paz–El Alto, Santa Cruz, Cochabamba); the eight smaller ones sit at or below a single OMI pixel.

ESA's TROPOspheric Monitoring Instrument (TROPOMI), aboard Sentinel-5 Precursor (Veefkind et al., 2012), provides an order-of-magnitude improvement: daily global coverage at 5.5×3.5 km since July 2018. We use the offline-reprocessed Level-3 product accessed through Google Earth Engine with quality filter `qa_value` $\geq$ `0.75`. At this resolution, all eleven Bolivian metropolitan areas produce usable retrievals, though Cobija and Montero sit near the pixel boundary and report only one to two valid native pixels per month (their exclusion is robustness-tested in Section 6.5).

Splicing the two instruments requires calibration. We estimate a per-ROI bias regression $\text{TROPOMI}_t = \alpha_r + \beta_r \cdot \text{OMI}_t + \varepsilon_t$ on the 42-month overlap (July 2018 – December 2021) and apply the fitted parameters to the pre-2018 OMI series. Table 1 reports the per-ROI results. The Santa Cruz fit is good ($R^2 = 0.81$), Cochabamba acceptable (0.58), La Paz weak (0.16). La Paz's low correlation reflects the spatial-resolution mismatch: OMI's 0.25° pixel averages the small high-altitude metro with surrounding rural terrain. On the population-weighted national series, the overlap-window correlation between bias-corrected OMI and TROPOMI is 0.85 with mean ratio 1.013. All OMI-period results are reported with and without La Paz in Section 6.

Table 1. OMI–TROPOMI splice calibration parameters by ROI.

ROI	$\hat{\alpha}$	$\hat{\beta}$	R^2	n_{overlap}
Santa Cruz	+5.55	+0.61	0.81	40
Cochabamba	−3.62	+2.27	0.58	23
La Paz–El Alto	+6.07	+0.91	0.16	25
Pop-weighted national	−0.41	+1.01	0.85	41

Notes: Per-ROI OLS regression on the 2018-07 to 2021-12 overlap window: $\text{TROPOMI}_t = \hat{\alpha} + \hat{\beta} \cdot \text{OMI}_t + \varepsilon_t$, with parameters then applied to the pre-2018 OMI series. Units: $\mu\text{mol/m}^2$. Co-sampled monthly observations only. La Paz–El Alto's low overlap correlation reflects the spatial-resolution mismatch between OMI's 0.25° pixel and the small high-altitude metropolitan area; OMI-period results are reported with and without this ROI in robustness exercises (Section 6). Sources: Copernicus Sentinel-5P TROPOMI L3; NASA OMI OMNO2d L3.

Table 2 lists the eleven metropolitan ROIs, their coordinates, buffer sizes, and 2024 census populations. The three large metros account for approximately 55 percent of the population-weighted aggregate. Population weights from the 2024 INE census are held fixed across the sample period; alternatives using 2012 census weights and equal weights are reported in Section 6.5.

Table 2. Eleven metropolitan ROIs for satellite NO_2 extraction.

Metro	Lat	Lon	Buffer (km)	Population (1000s)
La Paz–El Alto	−16.500	−68.150	20	1,800
Santa Cruz	−17.784	−63.180	25	1,900
Cochabamba	−17.394	−66.157	15	700
Sucre	−19.033	−65.263	8	300
Oruro	−17.970	−67.112	7	270
Potosí	−19.583	−65.753	6	190
Tarija	−21.534	−64.729	8	230
Trinidad	−14.832	−64.903	5	130
Cobija	−11.028	−68.768	4	65
Montero	−17.339	−63.253	6	140
Yacuiba	−22.017	−63.685	5	100

Notes: Population from INE 2024 census preliminary tabulations. The three large metros (La Paz–El Alto, Cochabamba, Santa Cruz) are extracted from both OMI 2004–2021 and TROPOMI 2018–2026; the eight smaller metros are TROPOMI-only because their footprint sits at or below the OMI 0.25° native pixel. Buffer is the half-side of the rectangular ROI in kilometers. Sources: INE Bolivia 2024 census; Copernicus Sentinel-5P TROPOMI; NASA OMI OMNO2d.

For each ROI, we aggregate daily retrievals to monthly means, requiring at least fifteen valid days for TROPOMI and ten for OMI (reflecting OMI's post-2007 row-anomaly coverage reduction). Cloud-contaminated months in Bolivia's wet season (November–March) occasionally fall below the threshold; those months are reported as missing rather than imputed. Andean volcanic NO_2 from Sabancaya, Ubinas, and Misti can drift east into the La Paz, Oruro, and Cochabamba ROIs; following Liu et al. (2020), we flag any month-ROI exceeding twice the seasonal baseline and exclude it from the headline aggregation. The TROPOMI-only panel has 1027 valid observations of 1034 possible (July 2018 through April 2026).

Seasonality and the choice of smoothing. The seasonal cycle dominates the raw monthly NO_2 series and determines how the indicator is constructed. Bolivia's NO_2 swings by a factor of approximately three between the dry-winter peak (May–August), when reduced photochemical loss and frequent thermal inversions trap pollution near the surface, and the wet-summer trough (November–March), when convection ventilates the boundary layer and increased UV destroys NO_2. The seasonal amplitude in the three large metros ranges from 20 to 40 percent of the annual mean. A recession signal of a few percent per year is small relative to this cycle and invisible in the raw monthly data.

Two standard approaches remove seasonality. Month-of-year fixed effects (or STL decomposition) subtract an estimated seasonal component and produce a deseasonalized residual, but amplify measurement noise in months where the seasonal estimate is imprecise — particularly during the wet-season cloud-contaminated months when both the retrieval and the seasonal estimate are noisiest. A 12-month rolling mean mechanically eliminates any exactly-annual cycle without estimating a seasonal component, producing a smooth coincident indicator at the cost of six months of endpoint coverage on each end. We use the 12-month centered rolling mean for level visualization and recession dating, and the trailing 12-month mean for growth rates and SARIMA training. Section 6.4 confirms that the recession-dating lead is invariant to the choice among these three methods.

3.3 INE PIB

INE provides three comparison series at progressively higher frequency. Quarterly real PIB at base year 2017, available 1990Q1–2024Q4 ($n = 140$), is the principal benchmark: the most recent observation (2024Q4, released 21 April 2026) is the first quarter to register negative year-on-year growth (-2.6 percent) since the COVID recovery. Annual departmental PIB (2017-base chained from 1990-base at 2017, $n = 333$ dept-year observations) provides the cross-sectional variation used in the elasticity estimation of Section 4.1. Monthly IGAE, available from March 2026 onward, is too short for any of the present exercises but will become the natural comparison series as it accumulates observations.

4 NO_2 versus lights as a recession indicator

If satellite observations are to track a recession, they must move in the same direction as real activity when activity falls. This section tests whether NO_2 and VIIRS meet that requirement on the 2020–2024 Bolivian departmental panel, and finds that they do not perform equally: NO_2 tracks PIB in both regimes while VIIRS tracks it only in expansion.

4.1 Elasticity

We estimate the departmental NO_2–PIB elasticity on the 2020–2024 annual panel ($n = 45$, nine departments $\times$ five years) with department fixed effects and standard errors clustered on department:

$$\log \text{PIB}_{d,t} = \alpha_d + \beta \cdot \log \text{NO}_{2d,t} + \varepsilon_{d,t}. \quad (1)$$

The estimate is $\hat{\beta}_{\text{NO}_2} = 0.531$ (cluster SE 0.098, $p < 0.001$, within $R^2 = 0.44$): a ten-percent deviation in departmental NO_2 from the within-department mean is associated with a five-percent deviation in real PIB. Adding year fixed effects collapses the estimate to 0.090 ($p = 0.25$), as expected — year effects absorb the national common variation the indicator aims to capture.

The same specification on VIIRS sum-of-lights returns $\hat{\beta}_{\text{VIIRS}} = 0.332$ (cluster SE 0.221, $p = 0.133$) — statistically indistinguishable from zero. In a joint specification, NO_2 retains significance at $\hat{\beta} = 0.554$ ($p < 0.001$) while VIIRS turns negative at $\hat{\beta} = -0.167$ ($p = 0.450$). Table 3 reports the full panel.

Table 3. Elasticity of departmental real PIB to tropospheric NO_2 (Panel A) and VIIRS sum-of-lights (Panel B), 2020–2024 panel. Specification: $\Delta \log \text{PIB}_{d,t} = \alpha_d + \gamma_t + \beta \cdot \Delta \log X_{d,t} + \varepsilon_{d,t}$ where X is NO_2 (Panel A) or VIIRS sum-of-lights from the VNP46A2 product, dropping sentinel-zero observations (Panel B). Both panels apply identical specifications to the same 9 departments and 5 first-difference years. Standard errors in parentheses: HC1-robust in columns (1)–(3), cluster-robust by department in columns (4)–(6). Columns (5)–(6) drop one department to test sensitivity. Column (7) reports the joint specification including both NO_2 and VIIRS with department fixed effects and cluster-robust SE; the same regression is shown identically in both panels (the column reads NO_2's coefficient in Panel A and VIIRS's in Panel B). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1) Pooled	(2) Dept FE	(3) Dept+Yr FE	(4) Dept FE	(5) no Cba.	(6) no Tarija	(7) Joint Dept FE
Panel A. NO_2							
$\Delta \log NO_2$	0.512***	0.531***	0.090	0.531***	0.582***	0.594***	0.554***
	(0.097)	(0.102)	(0.078)	(0.098)	(0.121)	(0.085)	(0.121)
R^2	0.395	0.442	0.821	0.442	0.410	0.499	0.447
n	45	45	45	45	40	40	45
Panel B. VIIRS sum-of-lights							
$\Delta \log$ VIIRS	0.314*	0.332	0.061	0.332	0.253	0.451**	-0.167
	(0.181)	(0.202)	(0.210)	(0.221)	(0.226)	(0.188)	(0.222)
R^2	0.027	0.061	0.815	0.061	0.052	0.049	0.447
n	45	45	45	45	40	40	45
Department FE	—	✓	✓	✓	✓	✓	✓
Year FE	—	—	✓	—	—	—	—
Cluster (dept) SE	—	—	—	✓	✓	✓	✓

The panel contains 18 contraction observations, but these arise from only two macroeconomic events: the 2020 pandemic and the 2024 contraction, each affecting all nine departments simultaneously. The clustered standard error understates the true uncertainty about the deep parameter. Section 6.8 reports the elasticity with and without each contraction year.

4.2 Regime asymmetry

The pooled elasticity obscures the result that matters most: the difference in tracking performance between expansions and contractions. Splitting the 45 observations by sign of PIB growth reveals the asymmetry. During the 27 expansion observations, both indicators move with PIB — NO_2 sign concordance is 74 percent, VIIRS 78 percent. During the 18 contraction observations, NO_2 maintains concordance at 72 percent; VIIRS collapses to 39 percent, worse than a coin flip, with VIIRS *rising* on average against falling PIB (indicator-to-PIB ratio -0.16). Figure 3 plots the regime-split scatter.

Figure 3. Regime-asymmetric performance of NO_2 versus VIIRS.

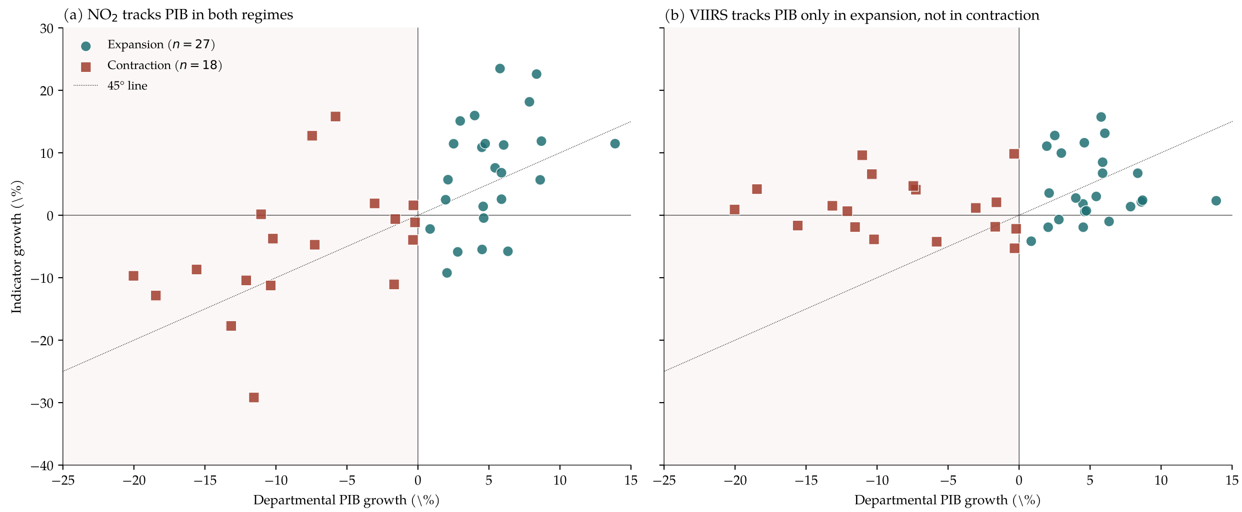

Notes: Each point is a dept-year observation from the 2020–2024 panel ($n = 45$, 9 departments × 5 years). Horizontal: Δ log INE departmental PIB. Vertical: Δ log NO_2 (left) or Δ log VIIRS (right). Diagonal: 45-degree line. NO_2 tracks PIB in both quadrants; VIIRS tracks PIB in expansion (upper-right) but not in contraction (lower-left). Sources: Copernicus Sentinel-5P TROPOMI; NASA VIIRS VNP46A2 v001; INE Bolivia.

The 2020 cross-section provides the cleanest amplitude test. Mean departmental PIB fell 13.6 percent across the nine departments. Mean NO_2 fell 11.5 percent — 84 percent of the PIB drop. Mean VIIRS *rose* 1.8 percent, with five of nine departments registering positive VIIRS growth despite all nine contracting. The asymmetry is structural. VIIRS measures a stock of nighttime infrastructure — lit area, LED-replaced streetlights, peri-urban informal-housing sprawl — that accumulates on multi-year investment cycles and does not contract when activity falls. NO_2 measures a flow of present-day combustion that responds within days to changes in transport, industrial, and residential fuel use. For recession detection, the flow is the binding object.

4.3 Reconciliation with Bolivar-Cuba (2024)

The regime-asymmetry result does not contradict Bolívar and Cuba (2024), who find that luminosity is the single most important satellite predictor in a nowcasting framework for Bolivia's monthly IGAE. The two exercises ask different questions. Nowcasting exploits level and trend variation, where expansion dynamics dominate and lights perform well. Recession detection requires tracking sign reversals in growth, where lights fail. An XGBoost robustness exercise on the same panel confirms the ranking: NO_2 permutation importance is 2.6 times VIIRS, with leave-one-department-out cross-validation producing out-of-sample R^2 of $+0.04$ for the joint model, -0.03 for NO_2-only, and -0.35 for VIIRS-only (Section 6.7).

5 Tracking the 2024–2026 recession

Having established that NO_2 outperforms VIIRS for recession detection on the 2020–2024 panel, we now apply the indicator to the ongoing 2024–2026 episode. The section proceeds chronologically:

what the indicator said before the recession was officially registered, what a pre-break SARIMA projects for 2026, what the first five months of post-break data show, and when the forecast becomes testable.

5.1 Recession dating

The trailing 12-month national NO_2 index peaked in June 2023 at $+12.1$ log points relative to 2019 (visible in Figure 2) and has declined since. INE quarterly PIB first registered negative year-on-year growth in 2024Q4, at -2.6 percent. The NO_2 turning point precedes the INE turning point by approximately five quarters at the national-peak level. At the metropolitan level, trailing 12-month first-negative readings arrived earlier: Santa Cruz in January 2024, La Paz in May 2024, Cochabamba in July 2024 — two to seven months before INE's Q4 2024 reading. The national and metropolitan lead estimates point in the same direction; both are reported.

Through November 2025 — the last date for which the centered 12-month NO_2 window is fully populated by actual data — the satellite implies a PIB level of $+1.2$ log points relative to 2019 under $\hat{\beta} = 0.53$, against INE's $+1.1$ for full-year 2024. The two agree to within one log point.

The agreement does not hold uniformly across departments. December-on-December trailing 2025 growth rates are -15.5 percent in Santa Cruz, -8.2 percent in Cochabamba, and $+4.1$ percent in La Paz. La Paz's positive reading reflects sustained government-employment combustion and motorization growth in the metropolitan area; it does not indicate that La Paz escaped the recession. Figure 4 plots the three large metropolitan area–department pairs. The Cochabamba gap — INE departmental PIB above the satellite-implied level by 8.7 log points at 2024 — is the largest positive discrepancy. Santa Cruz and La Paz show reverse gaps of 10 and 4.6 log points respectively, with NO_2 above INE. Tarija is excluded: the urban ROI does not capture the rural gas-extraction collapse (-21.9 percent) that dominates Tarija departmental PIB.

Figure 4. NO_2 versus INE departmental real PIB, three large metropolitan areas.

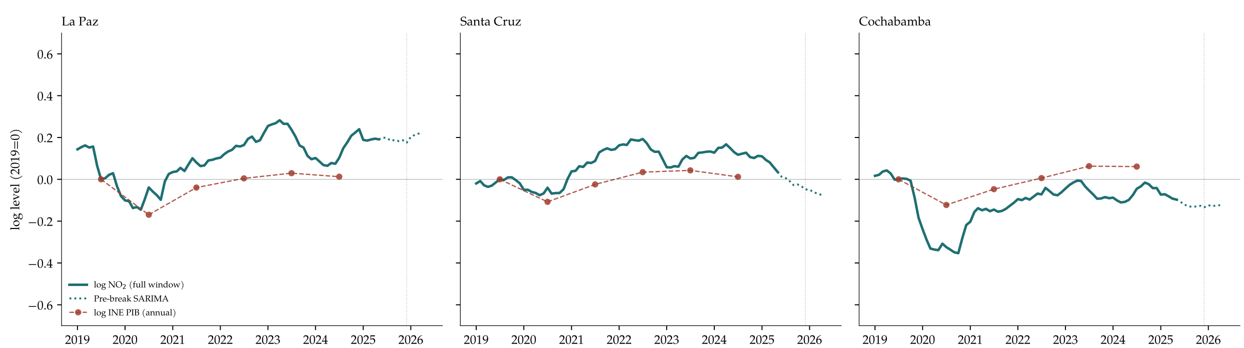

Notes: NO_2 (12-month centered, full-window only) versus INE departmental real PIB, both in log levels relative to $2019 = 0$. Solid teal: actual NO_2, terminating November 2025. Dotted teal: pre-break SARIMA point extrapolation through March 2026. Rust dashed with markers: INE annual departmental PIB. 95-percent forecast confidence intervals at March 2026, log points: La Paz $[-4.8, +49.4]$; Santa Cruz $[-24.2, +10.8]$; Cochabamba $[-43.4, +19.2]$. Tarija is excluded: the urban ROI does not capture the rural gas-extraction collapse (-21.9 percent) that dominates departmental PIB. Sources: Copernicus Sentinel-5P TROPOMI; INE Bolivia.

5.2 SARIMA pre-break forecast

To project the NO_2 trajectory beyond the last fully observable centered window, we fit the airline model SARIMA(0,1,1)(0,1,1)$_{12}$ to log monthly national NO_2, trained on data through November 2025 (89 months). The training window ends before the December 2025 fuel-subsidy elimination so that the forecast represents a no-policy-change counterfactual.

The centered-12-month forecast at March 2026 is -4.8 log points relative to 2019, with a 95-percent confidence interval of $[-28.6, +22.2]$. Under the departmental elasticity of 0.53, the implied PIB level is -2.5 log points $[-15.2, +11.8]$. Comparing to INE's preliminary 2025 level of -0.5 log points yields an implied 2026 growth rate of approximately -2.0 percent $[-14.7, +12.3]$. The point estimate falls within the multilateral consensus of -1.28 to -4 percent; the wide interval reflects the high seasonal amplitude of monthly NO_2 compounded over a five-month forecast window from a short training sample. The SARIMA is an early-warning signal, not a competitor to multilateral point forecasts. Figure 5 plots the forecast alongside post-break actuals.

Figure 5. Pre-break SARIMA forecast and post-break actuals.

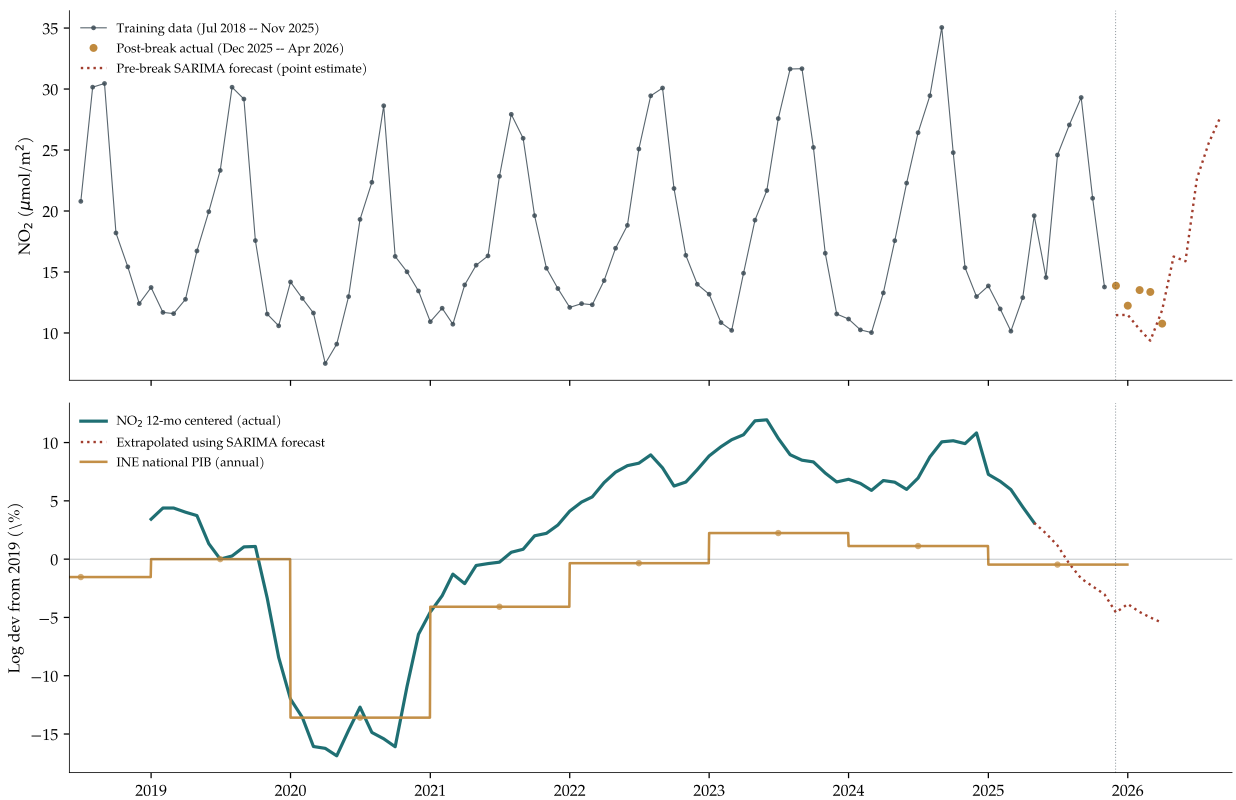

Notes: Top: raw monthly population-weighted national NO_2 (μmol/m^2). Slate solid: training data (July 2018 – November 2025). Ochre markers: post-break actuals (December 2025 – April 2026). Rust dotted: pre-break SARIMA point forecast. Bottom: 12-month centered rolling mean (log dev vs 2019). Teal solid: actual (full-window only, terminating November 2025). Rust dotted: SARIMA extension to March 2026. Ochre step: INE annual PIB. 95-percent CI at the March 2026 endpoint: $[-28.6, +22.2]$ log points (NO_2); implied 2026 PIB growth -2.0 percent $[-14.7, +12.3]$ under $\hat{\beta} = 0.53$. Sources: Copernicus Sentinel-5P TROPOMI; INE Bolivia.

5.3 Post-break observations

The first five months after the fuel-subsidy elimination (December 2025 – April 2026) provide an out-of-sample check, though not a clean one — the SARIMA was trained in a policy regime that no longer exists. Four of five months sit above the pre-break point forecast, and one (March 2026) reaches the upper 95-percent bound. On average, post-break actuals exceed the forecast by 15.7 percent. Two non-exclusive mechanisms contribute: subsidized diesel previously diverted to contraband markets may be returning to formal Bolivian consumption now that the price gap with neighboring countries has narrowed, and the December–April wet-to-dry meteorological transition raises column NO_2 through boundary-layer dynamics.

The post-break response splits along the highland–lowland axis. Highland metropolitan areas (Sucre $+28$, Oruro $+25$, Potosí $+10$, Cochabamba $+9$, La Paz -1 percent) show a group mean of $+14$ log points. Lowland and frontier areas (Tarija $+12$, Santa Cruz $+11$, Cobija $+5$, Trinidad $+3$, Yacuiba $+1$, Montero -7) show a group mean of $+4$. ERA5 boundary-layer-height controls reduce the highland mean from $+14$ to $+10$ log points, a one-third reduction (Section 6.6); the remainder is consistent with differential fuel-fleet composition and narrowed cross-border arbitrage.

Formally identifying whether the subsidy elimination constitutes a structural break in the NO_2 generating process is not yet possible. Three alternative shock dummies — step, pulse, and geometric decay — all fail to reach significance ($p = 0.83$ to 0.90), as reported in Section 6.3. Five months of post-break data cannot be separated from a seasonal cycle whose amplitude is 40 percent.

5.4 Verification timeline

The centered-12-month NO_2 average at March 2026 becomes fully populated by actual data when August 2026 NO_2 is released — which, given TROPOMI's two-month publication lag, will be around October 2026. INE Q1 2026 PIB will publish around July 2026 in INE's standard cadence. Both releases fall within six months of this writing, and we commit to a revisit at that horizon. The pre-break SARIMA's forecast of approximately -2 percent 2026 growth with a large confidence interval is reported here; we will update it when the data arrive.

6 Robustness

The headline results rest on choices about sensor calibration, time-series specification, seasonality removal, population weighting, and the treatment of meteorological confounders. This section varies each in turn. The recession-dating lead survives every alternative examined. The SARIMA point forecast is more sensitive to specification, as expected for a model-based extrapolation, but the qualitative direction is stable.

6.1 OMI splice sensitivity

Extending the SARIMA training sample from 89 months (TROPOMI-only, July 2018 – November 2025) to 254 months (OMI–TROPOMI spliced, October 2004 – November 2025) preserves the centered-12-month March 2026 point estimate within four log points (-4.8 on TROPOMI-only versus -0.9 on the spliced series) while narrowing the 95-percent confidence interval from 44.3 to 31.5 log points — a 29 percent reduction. The tighter interval reflects the additional seasonal cycles available for parameter estimation, not a change in the underlying signal. The headline forecast is reported on the TROPOMI-only sample because splice quality varies materially across ROIs (Table 1); the spliced result is presented as a bound on what a longer training window delivers.

6.2 Alternative SARIMA orders

The airline model SARIMA(0,1,1)(0,1,1)$_{12}$ used in the headline forecast is not the AIC-minimizing specification. An exhaustive grid search over $(p, q) \in \{0, 1, 2\}^2$ and $(P, Q) \in \{0, 1\}^2$ on the spliced sample with COVID and sensor dummies selects (0,1,2)(0,1,1)$_{12}$ at AIC $= -140.4$, an 18-point improvement over the airline's -122.0. The gain is driven by a statistically significant second-order moving-average term ($\hat{\theta}_2 = -0.31$, $p < 0.001$). Table 4 reports the top seven specifications. The recession-dating result — a turning-point detection on smoothed actuals — does not depend on the SARIMA order.

Table 4. SARIMA grid search on OMI–TROPOMI spliced sample.

Specification	AIC	BIC	$\hat{\sigma}^2$	k
$(0,1,2)(0,1,1)_{12}$	−140.44	−119.78	0.029	3
$(1,1,1)(0,1,1)_{12}$	−138.96	−118.28	0.029	3
$(1,1,2)(0,1,1)_{12}$	−138.82	−114.72	0.028	4
$(2,1,2)(0,1,1)_{12}$	−137.52	−109.98	0.028	5
$(2,1,1)(0,1,1)_{12}$	−136.56	−112.43	0.029	4
$(2,1,2)(1,1,1)_{12}$	−135.20	−104.22	0.028	6
$(0,1,1)(0,1,1)_{12}$ *(airline)*	−122.00	−104.73	0.032	2

Notes: Information criteria for SARIMA$(p, 1, q)(P, 1, Q)_{12}$ specifications fit on the OMI–TROPOMI spliced sample (October 2004 – November 2025, $n = 254$ valid months) with COVID dummy and sensor dummy as exogenous regressors. The AIC-minimizing specification is (0,1,2)(0,1,1)$_{12}$ with statistically significant second-order MA term $\hat{\theta}_2 = -0.31$ ($p < 0.001$). The airline (0,1,1)(0,1,1)$_{12}$ used in the headline forecast (Section **??**) is reported in the last row for reference. Sources: Copernicus Sentinel-5P TROPOMI; NASA OMI OMNO2d.

6.3 Gasolinazo dummies

Bolivia eliminated fuel subsidies on December 17, 2025 (Decree 5503). Whether this policy break constitutes a structural shift in the NO_2 generating process is testable by adding a shock dummy to the SARIMA training window. Table 5 reports three specifications: a step dummy for a permanent level shift ($D = 1$ for $t \geq$ December 2025), a pulse dummy for a one-month outlier ($D = 1$ at December 2025 only), and a geometric-decay dummy ($D_t = 0.7^k$ for k months after December 2025). None of the three reaches conventional significance: p-values range from 0.83 to 0.90, and none improves AIC over the baseline. With only five post-break months in a series whose seasonal cycle has 40-percent amplitude, the shock parameter cannot be separated from the seasonal pattern.

Table 5. Gasolinazo dummy specifications, SARIMA(0,1,2)(0,1,1)$_{12}$ on spliced sample.

	S0 cov+sen	S1 +step	S2 +pulse	S3 +decay
COVID	-0.230^{***}	-0.229^{***}	-0.231^{***}	-0.231^{***}
	(0.051)	(0.051)	(0.051)	(0.051)
Sensor	-0.004	-0.006	-0.000	$+0.001$
	(0.072)	(0.073)	(0.073)	(0.074)
Gasolinazo shock	—	$+0.014$	$+0.117$	$+0.053$
	—	(0.109)	(0.558)	(0.337)
	—	[0.90]	[0.83]	[0.88]
$\hat{\theta}_1$ (ma.L1)	-0.674^{***}	-0.674^{***}	-0.699^{***}	-0.704^{***}
$\hat{\theta}_2$ (ma.L2)	-0.307^{***}	-0.308^{***}	-0.322^{***}	-0.318^{***}
$\hat{\Theta}_{12}$ (ma.S.L12)	-0.719^{***}	-0.720^{***}	-0.721^{***}	-0.719^{***}
$\hat{\sigma}^2$	0.0286	0.0286	0.0276	0.0277
Log-likelihood	73.22	73.23	73.48	73.27
AIC	-140.44	-138.45	-138.96	-138.55
BIC	-119.78	-114.36	-114.86	-114.46
n	259	259	259	259

Notes: SARIMA(0,1,2)(0,1,1)$_{12}$ refit with three alternative gasolinazo shock specifications. Step dummy: $D_t = 1$ for $t \geq$ Dec 2025. Pulse dummy: $D_t = 1$ at $t =$ Dec 2025, zero elsewhere. Geometric-decay dummy: $D_t = 0.7^k$ for $t \geq$ Dec 2025, where k is months since Dec 2025. Standard errors in parentheses; p-values for the gasolinazo shock coefficient in brackets. Significance: *** $p < 0.01$. None of the three shock specifications reaches conventional statistical significance; none improves AIC over the COVID-and-sensor-only baseline (column 1). Sources: Copernicus Sentinel-5P TROPOMI; NASA OMI OMNO2d.

6.4 Alternative seasonality methods

The 12-month centered rolling mean used in the headline is one of three standard approaches to removing the seasonal cycle. Under STL decomposition (seasonal smoother window 13), the post-2020 NO_2 peak shifts from June 2023 to May 2023; under month-of-year fixed-effect adjustment it remains at June 2023. The last centered reading is $+3.5$ log points under the headline and month-FE methods. The STL trend component reads $+9.9$ at April 2026 because STL extrapolates to the series edge without the six-month endpoint loss that rolling-mean aggregation imposes. The five-quarter recession-dating lead is invariant across all three.

6.5 Alternative weighting and ROI sets

The population-weighted national aggregate uses 2024 INE census weights held fixed across the sample. Under 2012-census weights, the last centered reading shifts from $+3.5$ to $+3.7$ log points; excluding Cobija and Montero (the two ROIs with the smallest valid-pixel counts) shifts it to $+3.6$. The June 2023 peak date is invariant across all three census-based alternatives. Equal weighting assigns 73 percent of total weight to the eight smaller metropolitan areas and shifts the peak to October 2024 — a consequence of heterogeneous trajectories across small ROIs, not a challenge to the headline.

6.6 Meteorological controls

The post-break highland–lowland asymmetry documented in Section 5.3 could reflect meteorological seasonality rather than differential activity responses. To test this, we regress log NO_2 on ROI fixed effects, month-of-year fixed effects, and a quadratic in ERA5 10-meter wind speed and boundary-layer height ($R^2 = 0.75$; boundary-layer height significant at $p < 0.001$, wind speed at $p = 0.055$). The highland-mean post-break increase falls from $+14.2$ to $+9.6$ log points after meteorological control — a 33 percent reduction. The lowland-mean response falls from $+4.0$ to $+1.8$. Meteorology explains part of the highland increase but not all of it; the remaining $+9.6$ log points in highland metropolitan areas after controlling for wind and boundary-layer conditions are consistent with differential fuel-fleet composition, reduced cross-border arbitrage, and residual thermal-inversion dynamics not captured by ERA5's boundary-layer parameterization.

6.7 XGBoost cross-validation

The regime-asymmetry finding from the linear specification carries over to a nonlinear model. An XGBoost regressor trained on the growth-on-growth panel with two features ($\Delta \log NO_2$, $\Delta \log$ VIIRS) ranks NO_2 permutation importance 2.6 times above VIIRS. Out-of-sample R^2 values in leave-one-department-out cross-validation are small in absolute terms — $+0.04$ for the joint model, -0.03 for NO_2-only, -0.35 for VIIRS-only — reflecting the limited panel size rather than a failure of the satellite signal. The ranking, not the absolute fit, is the substantive finding.

6.8 Identification limits

The 2020–2024 departmental panel contains 18 contraction observations, but these arise from only two macroeconomic events: the 2020 pandemic (nine departments simultaneously) and the 2024 contraction (nine departments again). Re-estimating Equation (1) without the 2020 observations returns $\hat{\beta} = 0.23$ (cluster SE 0.23, $p = 0.32$); without 2024 it returns $\hat{\beta} = 0.51$ (SE 0.10). The pooled estimate of 0.53 draws most of its identification from the COVID year. This is a limitation of the Bolivian data, not a defect of the specification: a country with two contractions in five years cannot bound a deep parameter tightly from contraction observations alone. Future contractions — including the ongoing 2025–2026 episode, once INE departmental data become available — will tighten or shift the estimate.

7 Conclusion

Bolivia's national NO_2 indicator peaked in June 2023 and has declined since. INE quarterly PIB first registered a negative year-on-year reading five quarters later, in 2024Q4. The lead is robust to alternative seasonality treatments, weighting schemes, and aggregation choices; it is identified from completed actuals on smoothed series and does not depend on the elasticity estimate, the SARIMA specification, or the post-break observations.

The lead exists because NO_2 measures a combustion flow that contracts when activity falls, while VIIRS measures an infrastructure stock that does not. On the 2020–2024 departmental panel, VIIRS sign concordance during contractions is 39 percent — worse than coin-flip. NO_2 maintains 72 percent. This asymmetry explains why light-based nowcasts (Bolívar and Cuba, 2024) track Bolivian trends well but miss the recession that the present indicator detects.

A pre-break SARIMA projects a 2026 PIB contraction of approximately 2 percent, consistent with the IMF, World Bank, and IIF forecasts, though the confidence interval is wide — an honest reflection of what monthly NO_2 from a single small economy can identify on a six-month forecast horizon. Five months of post-break actuals sit above the projected path, consistent with the December 2025 fuel-subsidy elimination introducing a combustion shock the model never trained on. The centered-12-month March 2026 endpoint becomes fully observable from actual data in October 2026; INE Q1 2026 PIB arrives around July 2026. Both releases fall within six months of this writing, and we commit to a revisit at that horizon.

The broader implication is methodological. The regime asymmetry between combustion proxies and luminosity proxies is unlikely to be Bolivia-specific: LED rollout, peri-urban sprawl, and the stock character of nighttime infrastructure are global phenomena. If NO_2 outperforms lights for recession detection in Bolivia, it plausibly does so elsewhere — particularly in low-statistical-capacity economies where the measurement gap is largest and the stakes of early warning are highest. The indicator is reproducible from free Copernicus data with a two-month publication lag at `wernerhl.github.io/bolivia-ews`.

References

M. Bauwens, S. Compernolle, T. Stavrakou, J.-F. Müller, J. van Gent, H. Eskes, P. F. Levelt, R. van der A, J. P. Veefkind, J. Vlietinck, H. Yu, and C. Zehner. Impact of coronavirus outbreak on NO_2 pollution assessed using TROPOMI and OMI observations. *Geophysical Research Letters*, 47(11): e2020GL087978, 2020.

Robert C. M. Beyer, Yingyao Hu, and Jiaxiong Yao. Measuring quarterly economic growth from outer space. IMF Working Paper WP/22/109, International Monetary Fund, 2022.

Omar Bolívar and Paúl Cuba. Nowcasting Bolivian GDP using machine learning methods. Working paper, Banco Central de Bolivia, 2024.

Jean-Charles Bricongne, Baptiste Meunier, and Thomas Pical. Can satellite data on air pollution predict industrial production? Working Paper 847, Banque de France, 2021.

Dave Donaldson and Adam Storeygard. The view from above: Applications of satellite data in economics. *Journal of Economic Perspectives*, 30(4):171–198, 2016.

Mathilde Ezran, Stephen D. Morris, Martín Rama, and Daniel Riera-Crichton. Tropospheric NO_2 as a measure of economic activity: A new global indicator. Technical Report 10445, World Bank Policy Research Working Paper, 2023.

John G. Fernald, Eric Hsu, and Mark M. Spiegel. Is China fudging its GDP figures? Evidence from trading partner data. *Journal of International Money and Finance*, 114:102406, 2021.

John Gibson, Susan Olivia, Geua Boe-Gibson, and Chao Li. Which night lights data should we use in economics, and where? *Journal of Development Economics*, 149:102602, 2021.

J. Vernon Henderson, Adam Storeygard, and David N. Weil. Measuring economic growth from outer space. *American Economic Review*, 102(2):994–1028, 2012.

Fei Liu, Ashley Page, Sarah A. Strode, Yasuko Yoshida, Sungyeon Choi, Bo Zheng, Lok N. Lamsal, Can Li, Nickolay A. Krotkov, Henk Eskes, Ronald van der A, Pepijn Veefkind, Pieternel F. Levelt, Oliver P. Hauser, and Joanna Joiner. Abrupt decline in tropospheric nitrogen dioxide over China after the outbreak of COVID-19. *Science Advances*, 6(28):eabc2992, 2020.

Luis R. Martinez. How much should we trust the dictator's GDP growth estimates? *Journal of Political Economy*, 130(10):2731–2769, 2022.

Tomasz Michalski and Gilles Stoltz. Do countries falsify economic data strategically? Some evidence that they might. *Review of Economics and Statistics*, 95(2):591–616, 2013.

Michael Stanger. Bolivia: Informe de la misión de cuentas nacionales trimestrales (cnt). Informe de asistencia técnica, Fondo Monetario Internacional, Departamento de Estadística, September 2025. Apéndice A: Clasificador de las Actividades Económicas de la Base 1990 y Referencia 2017.

J. P. Veefkind, I. Aben, K. McMullan, H. Förster, J. de Vries, G. Otter, J. Claas, H. J. Eskes, J. F. de Haan, Q. Kleipool, M. van Weele, O. Hasekamp, R. Hoogeveen, J. Landgraf, R. Snel, P. Tol, P. Ingmann, R. Voors, B. Kruizinga, R. Vink, H. Visser, and P. F. Levelt. TROPOMI on the ESA Sentinel-5 Precursor: A GMES mission for global observations of the atmospheric composition for climate, air quality and ozone layer applications. *Remote Sensing of Environment*, 120:70–83, 2012.