

Article

External Inflation Exposure and Fiscal Policy Under the Dollar Peg: Evidence from GCC Economies

Muna Husain 

Department of Economics, College of Business Administration, Kuwait University, P.O. Box 5969, Safat 13060, Kuwait; m.safar@ku.edu.kw

Abstract

The six Gulf Cooperation Council (GCC) economies share a dollar peg and heavy hydrocarbon dependence, yet their inflation paths diverge sharply over time. Because the peg imports US monetary policy and forecloses nominal adjustment, GCC inflation plausibly reflects imported price pressure, oil revenue cycles working through fiscal capacity, and domestic real activity. We quantify the associations between these channels and headline inflation in country fixed-effects regressions estimated on an annual panel of the six economies over 1991–2023. Trading partner inflation is associated with domestic inflation with a contemporaneous, within-year coefficient of about 0.46—an association, not an identified causal effect, which strengthens to about 0.65 under country-specific, backward-looking trade weights—while an oil-orthogonal proxy for discretionary fiscal stance is negatively associated with inflation. Both results are stable in leave-one-country-out checks and are robust to lagged inflation; the external association holds under wild cluster bootstrap inference appropriate to six clusters, under Driscoll–Kraay inference robust to cross-sectional dependence, and in first-difference and distributed lag specifications, while two-stage bootstraps propagating the fiscal proxy’s construction uncertainty place the fiscal association at the margin of conventional significance. After the 2014 oil price collapse, the external association weakens and domestic activity and financial volatility carry the largest standardized weights. The cyclically adjusted fiscal residual—our methodological contribution—tracks published IMF non-oil balances more closely than the raw balance does and requires only an overall balance and an oil price.

Keywords: inflation pass-through; dollar peg; fiscal–monetary interaction; GCC economies; cyclically adjusted fiscal balance; oil-exporting economies

JEL Classification: E31; E63; F31; F41; Q43; C38



Academic Editor: Robert Czudaj

Received: 4 July 2026

Revised: 26 July 2026

Accepted: 29 July 2026

Published: 4 August 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

The Gulf Cooperation Council (GCC) economies present an unusual macroeconomic configuration. Six small open economies share a dollar peg, depend heavily on hydrocarbon exports, and operate similar fiscal institutions built around oil revenue stabilization. By construction, they have limited independent monetary policy: their interest rates track the US Federal Reserve, and their nominal exchange rates against the dollar barely move. One might therefore expect their inflation paths to look alike. They do not. In 2008, at the peak of the commodity boom, headline inflation reached 15 percent in Qatar and 12 percent in the United Arab Emirates, while inflation in Bahrain stayed below 4 percent throughout the 2006–2008 boom years. Between 2014 and 2019, the picture flipped: Qatar and the UAE

settled into near-zero inflation while Saudi Arabia recorded mild deflation in 2017 and 2019. Same peg, same oil dependence, very different inflation outcomes.

This heterogeneity is informative. Because the peg ties the policy rate to the US Federal Funds rate and forecloses independent nominal adjustment—interest rate and bank-lending channels still operate, but their stance is imported rather than chosen (Cevik & Teksoz, 2013)—GCC inflation is plausibly generated by some combination of imported price pressure transmitted through the dollar, oil revenue cycles working through fiscal capacity, and domestic real activity in the non-oil economy. The combination is rare globally. Norway and Chile share some elements but not all; the GCC is the largest cohort of dollar-pegged hydrocarbon exporters in the world, and their data span more than three decades of varied global conditions. Understanding which of these channels dominate, and how their relative importance has shifted over time, is the goal of this paper. It is also of direct policy relevance: if the association with imported inflation is strong, GCC authorities must watch external price developments closely even though the peg ties their hands on interest rates; if the discretionary fiscal stance is systematically associated with inflation, fiscal authorities hold the one macroeconomic instrument the peg leaves in domestic hands.

The empirical literature on GCC inflation has converged on an externally driven account, mostly under linear and time-invariant transmission. Espinoza and Prasad (2012) document the role of trading partner inflation and import prices in a vector autoregression; Hasan and Alogeel (2008) emphasize imported inflation in Saudi Arabia and Kuwait; Kandil and Morsy (2011) examine domestic and external determinants across the GCC; and Murshed and Nakibullah (2015) find trading partner inflation to be the principal determinant of GCC consumer prices, with exchange rate effects close to zero. More recent work widens the lens: Fareed et al. (2023) trace GCC inflation to imported inflation and the effective exchange rate in a quarterly GVAR through 2022; Alsamara et al. (2018) and Alsamara et al. (2020) document asymmetric import cost pass-through in GCC data; Cevik and Miryugin (2025) show that the inflationary impact of fiscal shocks in a 140-country panel varies with fiscal space and the cycle—with effects concentrated in developing economies—though their analysis does not separately treat the GCC; and Sadraoui and Mili (2025) model the oil–fiscal nexus in GCC data through 2024 with fiscal, rather than inflation, outcomes. Two gaps remain. First, to our knowledge, no study combines this full set of external drivers—trading-partner inflation, oil prices, US monetary surprises, global real activity, and financial volatility—with a fiscal stance measure in a single annual panel of the six dollar-pegged GCC economies. Second, the measurement of fiscal stance in these economies deserves a treatment suited to annual GCC data: the headline fiscal balance mixes discretionary policy with the mechanical movement of oil revenues, and the non-oil balance measures standard in the resource economy literature (Barnett & Ossowski, 2003) require consistent oil and non-oil revenue decompositions that are not available for all six countries back to 1990. This paper addresses both.

We assemble an annual panel of the six GCC economies constructed for 1990–2024—the effective estimation sample, after listwise deletion described in Section 3.2, spans 1991–2023 with $N = 165$ —and estimate country fixed-effects regressions of headline CPI inflation on external drivers (trading partner inflation, the nominal effective exchange rate, Brent prices, US monetary policy surprises following Bauer and Swanson (2023), the Kilian (2009) real activity index, and the VIX), domestic monetary and real activity controls (M2 growth, non-oil GDP growth), and two fiscal variables: government consumption as a share of GDP, a measure of fiscal size, and a cyclically adjusted fiscal balance that we construct as the country-specific residual of the overall balance regressed on contemporaneous and lagged Brent prices. The cyclically adjusted measure strips the budget’s mechanical response to oil prices, leaving a residual we read as a proxy for discretionary fiscal stance. Because

aggregate external regressors induce cross-sectional dependence and the panel has only six countries, we report wild cluster bootstrap inference alongside conventional clustered standard errors and confirm the headline results in leave-one-country-out checks.

This paper makes three contributions. First, to our knowledge, we provide the first annual panel analysis of the six dollar-pegged GCC economies to combine trading partner inflation, US monetary surprises, global real activity, financial volatility, and an oil-adjusted fiscal stance measure in a single country fixed-effects framework over 1990–2024, showing that trading partner inflation is associated with domestic inflation with a contemporaneous coefficient of about 0.46 while a tighter discretionary fiscal stance is associated with lower inflation. We emphasize at the outset that the partner inflation estimate captures same-year co-movement: in a lagged specification, the association vanishes at one year, which is consistent with rapid within-year transmission through import prices but cannot, by itself, rule out common shocks or reverse causation—a distinction we treat as central to interpreting the coefficient throughout. Second, we document a structural shift in the GCC inflation process around the 2014 oil price collapse, after which the external association weakens and domestic real activity and financial volatility carry the largest standardized weights. Third, we construct an oil-orthogonal fiscal-balance residual for commodity exporters by purging oil price cyclicity from the budget country by country—a transparent, replicable proxy for discretionary stance that complements the non-oil balance framework where consistent revenue decompositions are unavailable. The remainder of the paper is organized as follows. Section 2 reviews the literature; Section 3 describes the data; Section 4 sets out the methodology; Section 5 presents the results; Section 6 discusses interpretation, an exploratory analysis of state dependence, and limitations; and Section 7 concludes.

2. Literature Review

2.1. Inflation in the GCC

The empirical literature on inflation in the Gulf Cooperation Council has converged on an external-driven account. [Hasan and Alogoskoufis \(2008\)](#) decompose Saudi Arabian and Kuwaiti inflation using a structural vector autoregression and attribute the bulk of variation to imported inflation operating through the dollar peg. [Espinoza and Prasad \(2012\)](#) extend the framework to a panel of GCC and other Middle East economies and report that trading partner inflation and oil prices dominate the conditional variance of CPI. [Kandil and Morsy \(2011\)](#) reach broadly similar conclusions for individual GCC economies, with limited estimated roles for fiscal policy or domestic credit. [Murshed and Nakibullah \(2015\)](#) extend the imported inflation account with an ARDL treatment of the six countries, finding trading partner prices to be the principal long-run determinant of the consumer price level and exchange rate effects close to zero. [Basher and Elsamadisy \(2012\)](#), by contrast, assign the leading role to the money supply in both the short and long runs—a reminder that the external-driven consensus is not unanimous, and one to which our pre-2014 M2 result speaks directly. [Mohanty and Klau \(2004\)](#), in a wider emerging-market context, document that central banks in countries with substantial external exposure react strongly to exchange rate movements in their interest-rate-setting behavior—a complementary observation rather than a direct GCC inflation result, but consistent with the picture of external dominance under pegged regimes. Three recent contributions round out the landscape. [Fareed et al. \(2023\)](#) estimate a quarterly Global VAR through 2022 and attribute GCC inflation principally to imported inflation and the effective exchange rate, with direct commodity price pass-through limited by subsidies and administered prices. [Alsabhan \(2026\)](#) applies a panel PMG-ARDL model to GCC data over 1998–2023 and likewise weighs internal against external determinants, though without a fiscal stance measure or state-conditional analysis.

The most recent contribution to the GCC PVAR literature, [Sadraoui and Mili \(2025\)](#), focuses on the fiscal side: they estimate a panel vector autoregression on annual GCC data over 2000–2024 to analyze how oil price volatility transmits to government spending and budget balances. Their fiscal focus complements, rather than overlaps with, the inflation focus of the earlier consensus literature; together, these papers provide the GCC empirical landscape against which our results should be read.

Two features of this literature are worth noting. First, most of the existing work imposes linear, time-invariant transmission: coefficients on imported inflation, oil prices, and fiscal aggregates are estimated as constants. The exceptions are [Alsamara et al. \(2018\)](#) and [Alsamara et al. \(2020\)](#), who document asymmetric import cost pass-through in GCC data with nonlinear ARDL methods—asymmetry in the sign of the shock, however, is a different object from the state dependence in the macroeconomic environment that we examine, and neither paper conditions on fiscal stance. Second, the literature treats fiscal policy almost exclusively through government consumption or budget balance as a share of GDP, without separating the cyclical component driven by oil prices from the discretionary component reflecting policy choice. Both features leave open questions that the broader pass-through and resource economy literatures have already moved past.

2.2. Exchange Rate and Import Price Pass-Through Under Pegged Regimes

The pass-through literature has, since the early 2000s, increasingly emphasized that the relationship between external prices and domestic inflation is not constant. [Taylor \(2000\)](#) was the first to argue formally that pass-through is endogenous to the inflation environment: in low and stable inflation regimes, firms expect price changes to be transitory and absorb them into margins, while in high or volatile inflation regimes, firms pass external shocks through quickly to defend real revenues. The Taylor hypothesis predicts state-dependent pass-through directly, and a substantial body of empirical work has confirmed it.

[Choudhri and Hakura \(2006\)](#) test the Taylor hypothesis on a panel of 71 countries and document a strong positive relationship between average inflation and pass-through magnitude. Their cross-country result has since been refined into within-country state-dependent estimates. [Caselli and Roitman \(2019\)](#) study commodity-exporting emerging economies specifically and find that pass-through coefficients more than double during high-inflation episodes. [Forbes et al. \(2018\)](#) emphasize that the source of an exchange rate shock matters: pass-through is materially larger when shocks originate in monetary policy than when they originate in demand or supply, even when the size of the exchange rate movement is similar. [Bauer and Swanson \(2023\)](#) provide the cleanest identification of the monetary policy channel by constructing high-frequency monetary surprises from FOMC announcement windows that are plausibly exogenous to GCC conditions and therefore well suited as instruments for studying transmission under dollar pegs.

The pegged regime case is theoretically sharp. Under a credible peg, the nominal exchange rate against the anchor currency is mechanically fixed, so pass-through to domestic prices runs through the relative price adjustment between traded and non-traded goods, the volume and composition of imports, and the indirect effect of monetary conditions transmitted via the anchor's policy rate. [Shambaugh \(2004\)](#) shows that countries with fixed exchange rates effectively import the monetary policy of the anchor country, leaving domestic interest rates highly correlated with the anchor's rate. Pegs do not eliminate inflation transmission; they rechannel it through the anchor country interest rate path and the relative price margins available within the traded–non-traded structure. [Brandão-Marques et al. \(2020\)](#) document, in a panel of 40 emerging and developing economies, that monetary policy transmission to inflation operates substantially through the exchange

rate channel—a finding directly relevant to dollar-pegged contexts where the channel is mechanically fixed but the underlying transmission still operates through prices.

The dominant currency pricing literature sharpens the prediction for our setting. [Gopinath \(2015\)](#) and [Gopinath et al. \(2020\)](#) document that the overwhelming majority of world trade is invoiced in a small number of dominant currencies—the dollar above all—so that import prices respond to the dollar exchange rate and to dollar-denominated foreign prices rather than to bilateral rates. For an economy pegged to the dollar, the implication is direct: with the dollar rate fixed, the operative external price channel is movement in trading partner prices themselves, not exchange rate fluctuation. Our headline finding—a strong association with trade-weighted partner inflation alongside a small and insignificant NEER coefficient—is exactly the configuration dominant currency pricing predicts for a dollar-pegged importer.

The GCC's presence in this literature is narrower than its theoretical interest warrants. The GCC-specific pass-through papers—[Murshed and Nakibullah \(2015\)](#) on the exchange rate channel, [Alsamara et al. \(2018\)](#) and [Alsamara et al. \(2020\)](#) on import cost asymmetry, and [Fareed et al. \(2023\)](#) on external factors at quarterly frequency—each isolate one dimension of transmission, and none combines the full external-driver set with a fiscal stance measure or asks whether transmission varies across macroeconomic states. Quantifying these channels jointly for the GCC, and providing a fiscal stance proxy suited to the setting, is the gap this paper fills.

2.3. Fiscal Policy and Inflation in Resource-Rich Economies

The fiscal–inflation nexus in resource-rich economies poses a measurement problem that pass-through analysis does not. Government consumption as a share of GDP captures the size of the public sector but not its *stance*: a country running a 25 percent consumption share with a stable surplus is different from one running 25 percent with a widening deficit, yet both look identical in this metric. The overall fiscal balance is similarly contaminated in oil-exporting economies because it moves mechanically with hydrocarbon revenue. A negative fiscal–inflation coefficient on the raw balance can reflect either discretionary tightening or a passive oil revenue windfall.

For oil exporters specifically, the standard answer is the non-oil balance. [Barnett and Ossowski \(2003\)](#) argue that because oil revenue originates abroad and moves with a price outside the authorities' control, the non-oil (primary) balance is the more reliable indicator of discretionary stance, and it is the measure IMF surveillance routinely reports for the GCC. The applied literature follows this convention: [Villafuerte and López-Murphy \(2010\)](#) track oil producers' fiscal stance through the 2003–2009 price cycle using non-oil indicators, [Erbil \(2011\)](#) tests fiscal procyclicality across 28 oil producers with non-oil revenue and the non-oil primary balance among her fiscal measures, and [Sturm et al. \(2009\)](#) apply the framework to the GCC. The constraint on this approach is data: it requires a consistent decomposition of revenues into oil and non-oil components, which is not available on a comparable basis for all six GCC countries back to 1990.

Beyond the oil exporter setting, the general literature has converged on cyclically adjusted measures. [Cevik and Miryugin \(2025\)](#) study fiscal policy shocks in a panel of 140 countries using the cyclically adjusted primary balance as the stance variable and document that the inflationary impact of fiscal shocks is state-dependent—varying with fiscal space, recession versus expansion, monetary policy framework, and exchange rate regime—with effects concentrated in developing economies. [Tariq et al. \(2024\)](#) document a nonlinear threshold relationship between government size and inflation in MENA economies, with an estimated turning point at 12.46 percent of GDP. Neither the non-oil balance nor a cyclically adjusted measure has, to our knowledge, been applied in a GCC panel inflation

regression. We construct our measure by regressing each country's overall fiscal balance on contemporaneous and lagged Brent prices and using the residual: an oil-orthogonal proxy for discretionary stance that requires only the overall balance and the oil price, which we read as a data-light approximation to the non-oil balance where consistent revenue decompositions are unavailable. It is a simpler construction than the IMF's structural balance methodology (Bornhorst et al., 2011) adapted to the data we have.

Theory also offers a competing framework. The Fiscal Theory of the Price Level (FTPL), developed by Leeper (1991) and Sims (1994), and recently consolidated by Cochrane (2023), argues that the price level adjusts to satisfy the government's intertemporal budget constraint; in a regime where the fiscal authority does not respond to debt accumulation, deficits are inflationary regardless of monetary policy. This prediction—that fiscal expansion raises inflation—is the opposite sign to the demand channel literature surveyed above, in which a tighter discretionary fiscal stance lowers inflation. Our discretionary fiscal stance results speak to which reading better fits the GCC sample, though a full FTPL test, which requires government debt and surplus process data at sufficient frequency to identify the fiscal–monetary policy mix, is beyond the scope of our annual panel.

2.4. Regime Identification in Macroeconomic Panels

The empirical literature on macroeconomic regime identification splits into two traditions. The model-based tradition, beginning with Hamilton (1989) and extended by Kim and Nelson (1999) and Sims and Zha (2006), uses Markov-switching or time-varying parameter vector autoregressions to recover regime probabilities and transition rates from the parametric model itself. These methods deliver structural interpretability and probabilistic regime assignment but require substantial within-country time-series length—typically 80–100 observations per series at minimum—to estimate transition matrices reliably. With $T = 35$ annual observations per GCC country, model-based regime identification is statistically infeasible at the country level and impractical at the panel level.

The model-free tradition uses unsupervised clustering on macroeconomic indicators directly. Its econometric foundation is the grouped fixed-effects framework of Bonhomme and Manresa (2015), which establishes the large-sample properties of estimators that assign panel units to data-driven groups; recent applied work—Nishi (2025) on k-means panel data clustering in the presence of small groups, alongside a growing applied machine learning in macro literature—extends that framework in the direction most relevant to our setting. Nishi's results are particularly directly applicable: he establishes consistency and asymptotic normality of least-squares estimators for grouped fixed-effects panel models in the presence of asymptotically negligible small groups (our R3 regime, with $N = 11$, is exactly such a case) and proposes modified information criteria that perform well in this small-group setting. The advantages of the model-free approach are flexibility and modest sample requirements; the disadvantage is that clustering is correlational by construction and can be unstable if not diagnosed properly. Hennig (2007) develops the bootstrap–Jaccard stability assessment we use to detect persistent partitions. Hennig and Liao (2013) provide a broader framework for matching clustering methodology to the substantive aim of an application, arguing that there is no universally optimal clustering method—the right method depends on what the regimes are meant to represent. We adopt both perspectives. The k-means classification we estimate is best read not as a competitor to Markov-switching estimation but as a descriptive companion to the panel regression, used to test whether regression coefficients vary across data-driven states.

2.5. Hypotheses

The literature reviewed above yields three testable hypotheses and one exploratory conjecture, which structure the empirical work that follows. H1: Under the dollar peg, trading partner inflation is positively associated with GCC headline inflation, conditional on oil cycle, US monetary, global activity, and domestic controls—the configuration dominant currency pricing predicts for a dollar-pegged importer. H2: A tighter discretionary fiscal stance, measured by the cyclically adjusted balance, is negatively associated with inflation, holding structural fiscal size fixed. H3: These associations are not constant over the sample; in particular, the external price association weakens after the 2014 oil price collapse as the inflation process becomes more domestically driven. Finally, and exploratorily, C1: transmission may be state-dependent, amplifying in boom states—a conjecture the Taylor (2000) hypothesis motivates but that a six-country annual panel can examine only descriptively. Section 5.1 evaluates H1 and H2, Section 5.2 evaluates H3, and Section 5.3 examines C1; Section 7 returns to each.

3. Data

3.1. Panel Construction and Coverage

The analysis is built on an annual panel of six Gulf Cooperation Council economies—Bahrain, Kuwait, Oman, Qatar, Saudi Arabia, and the United Arab Emirates—covering 1990 through 2024. The country-year frame is balanced (210 country-year observations), although several external series are unavailable for parts of this window, beginning after 1990 or ending before 2024. The variables fall into six groups: the dependent variable (CPI inflation), domestic macroeconomic controls, fiscal variables, exchange rate measures, external drivers (commodity prices, partner inflation, US monetary policy, global activity, and financial volatility indices), and a regime classification input set used in the unsupervised analysis. Table 1 catalogs each of these variables, together with their units, sources, and roles in the empirical work. Table 2 then reports country-by-year coverage for the series with non-trivial missingness; the remaining variables are observed for all 210 country-years.

Table 1. Variable inventory.

Variable	Units	Source	Time Coverage	Role
country	string	—	1990–2024	Identifier
code	ISO-3	World Bank	1990–2024	Identifier
year	integer	—	1990–2024	Identifier
inflation	%	World Bank WDIs	1990–2024 (OMN, UAE partial)	Dependent variable
gdp_growth	%	World Bank WDIs	1990–2024	Control
gov_cons_gdp	% of GDP	World Bank WDIs	1990–2024	Fiscal structure (level)
capital_form_gdp	% of GDP	World Bank WDIs	1990–2024	Control
non_oil_gdp_growth_pp	%	GCC central banks	1990–2024	Domestic activity
m2_growth_pp	%	GCC central banks/IFS	1990–2024	Domestic monetary
fiscal_balance	% of GDP	IMF Article IV (GCC)	1990–2024	Fiscal stance (raw)
fiscal_balance_caf	% of GDP	Authors (residual on Brent_t, Brent_{t−1})	1991–2024	Fiscal stance (cyclically adjusted)
neer	index	GCC source	1990–2024	Exchange rate (level)
reer	2007 = 100	Bruegel/Darvas (29 April 2026)	1992–2024	Robustness/regime validation
partner_inflation_pp	pp	Authors' calculation from IMF IMTS (May 2026 vintage) and World Bank WDIs (20 July 2026 extract)	1990–2024	External driver

Table 1. Cont.

Variable	Units	Source	Time Coverage	Role
brent	USD/bbl	EIA/FRED	1990–2024	External driver (level)
brent_yoy_pp	%	Authors	1990–2024	External driver (YoY change)
fed_funds	%	FRED FEDFUNDS	1990–2024	External driver (level)
fed_funds_change_bps	bps	Authors	1990–2024	External driver (change)
bs_mps_bps	bps	Bauer and Swanson (2023)	1990–2023	US monetary surprise
bs_mps_n	count	Bauer and Swanson (2023)	1990–2023	Number of FOMC meetings
kilian_rea	index	Kilian (2009), updated	1990–2024	Global real activity
vix	index	CBOE	1990–2024	Financial volatility

Note: Units, source, and time coverage for each variable in the panel. “Partial” indicates a series with less than full 1990–2024 country coverage.

Table 2. Panel coverage by variable and country.

Variable	BHR	KWT	OMN	QAT	SAU	UAE
bs_mps_bps	34/35	34/35	34/35	34/35	34/35	34/35
fiscal_balance	35/35	35/35	35/35	35/35	34/35	34/35
fiscal_balance_caf	34/35	34/35	34/35	34/35	34/35	34/35
inflation	35/35	35/35	24/35	35/35	35/35	17/35
m2_growth_pp	35/35	33/35	35/35	35/35	35/35	35/35
reer	33/35	33/35	33/35	33/35	33/35	33/35

Note: Cells report observed / total country-years (total = 35 per country; 210 overall). BHR = Bahrain, KWT = Kuwait, OMN = Oman, QAT = Qatar, SAU = Saudi Arabia, UAE = United Arab Emirates. Counts reflect genuine source coverage. The Kuwait 1990–1993 observations are retained in the regression panel and set aside only from the macroeconomic state clustering (Section 5.3).

Two coverage gaps deserve note up front. The Bauer–Swanson monetary policy surprise series ends in December 2023 in the version we use, leaving 2024 unobserved for that variable. The Bruegel real effective exchange rate series begins in 1992 for all six GCC countries, leaving 1990 and 1991 unobserved for REER. Neither gap is binding for the headline specification: REER is not an HM5 regressor, and the HM5 estimation sample spans 1991–2023 ($N = 165$), with each country entering once its own covariates are observed (Kuwait, for example, enters in 1993 because M2 and NEER are unobserved in the early war years). The lagged-driver specification HM6 is described in Section 4.1.

3.2. Dependent Variable and Core Macroeconomic Controls

The dependent variable is annual CPI inflation, measured as the headline percentage change in the consumer price index, taken from the World Bank’s World Development Indicators (WDIs). All six GCC economies report CPI annually for the full sample period, including Kuwait, whose 1990 and 1991 inflation observations are available in the WDI source and are retained. Genuine source gaps arise instead for Oman and the United Arab Emirates, whose CPI series begin later (see Table 2). The mean of inflation across the 181 observed country-years is 2.30 percent, with a standard deviation of 2.90 percentage points; within the HM5 estimation sample, the corresponding figures are 2.26 and 2.90. Inflation ranges from a minimum of −4.86 percent (Qatar in 2009) to a maximum of 15.05 percent (Qatar in 2008)—the two extremes fall in consecutive Qatari years, at the peak of the 2008 commodity boom and the trough of its 2009 reversal.

Core macroeconomic controls also come from the WDIs: real GDP growth, gross capital formation as a share of GDP, and general government final consumption expenditure as a share of GDP. Government consumption is our measure of fiscal structure—the long-run

size of the public sector relative to the economy—and enters the regressions as a level. Real GDP growth captures the aggregate cycle, while non-oil real GDP growth (described in Section 3.3 below) captures the component of the cycle most relevant to non-traded inflation. Two domestic monetary aggregates enter the panel: M2 growth, taken from national source compilations through the GCC central banks, and non-oil real GDP growth, drawn from the same source. Table 3 then reports descriptive statistics on the dependent variable and these macroeconomic controls.

Table 3. Descriptive statistics.

Variable	N	Mean	SD	Min	p25	Median	p75	Max
CPI inflation (%)	181	2.296	2.902	−4.863	0.612	2.007	3.255	15.05
Real GDP growth (%)	210	4.516	8.543	−41.008	1.538	3.914	6.284	82.809
Non-oil real GDP growth (%)	208	4.525	3.464	−5.8	3.0	4.0	5.5	25.0
Gov consumption (% of GDP)	191	19.885	7.913	6.733	14.869	18.847	24.200	76.222
Gross capital formation (% GDP)	191	25.175	8.049	6.952	19.076	25.199	29.673	48.869
Fiscal balance (% of GDP, raw)	208	3.231	17.070	−144.537	−4.957	1.727	10.078	52.267
Fiscal balance (CAF, % of GDP)	204	0.000	13.710	−151.565	−4.169	−0.363	5.083	40.902
M2 growth (%)	208	9.068	7.962	−6.1	3.875	7.4	12.2	44.0
Partner inflation (%)	210	3.979	1.659	1.40	2.81	3.33	5.29	9.23
Brent price (USD/bbl)	210	52.911	31.833	12.72	20.67	52.32	79.61	111.63
Brent YoY change (%)	204	7.424	27.961	−47.135	−12.164	1.353	31.447	68.875
Fed funds rate (%)	210	2.855	2.336	0.08	0.36	2.16	5.30	8.10
BS-MPS, annual sum (bps)	204	−2.417	24.067	−81.733	−7.975	3.917	11.04	36.759
Kilian REA index	210	3.050	52.282	−89.488	−32.115	−5.492	42.494	144.446
VIX index	210	19.467	5.740	11.09	15.388	17.536	24.203	32.695
NEER index	208	106.077	8.677	85.0	99.725	105.35	113.8	123.9
REER index (2007 = 100)	198	107.147	14.004	64.609	98.739	107.611	119.556	131.176

Note: Summary statistics computed on the full 1990–2024 panel (210 country-years); per-variable N reflects source coverage (Table 2). The Kuwait 1990–1993 war-period observations are included and account for the extreme values of real GDP growth, government consumption, and the two fiscal balance measures; these observations are retained in the regression panel and set aside only from the macroeconomic state clustering (Section 3.2). The cyclically adjusted fiscal balance has mean zero by construction. Estimation samples are smaller than the counts shown because of listwise deletion on the headline covariates ($N = 165$ for HM5; Section 3.1).

The estimation sample is determined by listwise deletion on the headline covariates rather than by any deliberate exclusion of observations. Every country loses its first sample year because the nominal effective exchange rate percentage change is undefined at the start of each series, and every country loses 2024 because the Bauer–Swanson monetary policy surprise series ends in 2023. Countries with later-starting source series (notably Oman and the United Arab Emirates) contribute correspondingly fewer years. This yields the headline sample of 165 country-year observations. One deliberate exclusion applies, and it affects only the unsupervised macroeconomic state clustering in Section 5.3, not the regressions: the Kuwait 1990–1993 observations, distorted by the Iraqi invasion and reconstruction (August 1990 to early 1994) and the post-COVID-19 year 2020, are set aside from the clustering because those macroeconomic states would otherwise distort the state assignment. Both are retained throughout the regression analysis and in the construction of the cyclically adjusted fiscal balance (Appendix B).

3.3. External Drivers

The empirical novelty of the panel is the joint inclusion of five external drivers, several of which have not previously been applied together in GCC inflation regressions. We describe each in turn.

Trading partner inflation is constructed as a trade-weighted average of CPI inflation in each country's import partners. Trade weights are derived, for each country separately, from its average bilateral goods imports (CIF) over 2010–2019 reported in the IMF Inter-

national Trade in Goods by Partner Country database (May 2026 vintage); partner CPI inflation is taken from the World Bank World Development Indicators (July 2026 extract). Four construction rules apply. First, the other five GCC members are excluded from each country's partner set so that the measure captures external rather than intra-bloc price pressure. Second, partner-year observations with CPI inflation above 50 percent are excluded before weighting, with the remaining weights renormalized, preventing hyperinflation episodes among minor partners from dominating the series. Third, partners absent from the WDI inflation series in a given year are dropped with renormalization; the minimum resulting basket coverage is 0.855 (Bahrain, early 1990s); coverage below 0.90 is confined to 26 of the 210 country-years, all in 1990–1995, and coverage exceeds 0.90 in every country-year from 1996 onward. Fourth, the series is expressed in percentage points. Because import structures are similar across the six economies, the resulting country-specific series co-move very tightly—pairwise correlations across countries range from 0.98 to 1.00—but they are not identical: each reflects its own import basket. Two limitations of the fixed-weight construction motivate a robustness exercise, which is reported in Section 5.1: weights estimated on 2010–2019 embed forward-looking information into earlier observations, and they cannot capture the pronounced shift in GCC import baskets toward Asia over the sample. Section 5.1 therefore re-estimates the headline specification with an alternative series in which each country's weights are computed from its trailing five-year import shares from the same sources so that no forward-looking information enters any observation.

Brent crude oil prices are taken from the EIA spot price series via FRED, averaged to annual frequency. The variable enters the regressions as a year-on-year percentage change; the price's one-year lag is also used in the construction of the cyclically adjusted fiscal balance described below.

The *US Federal Funds rate* is the annual mean of the daily effective federal funds rate from FRED. We also construct the year-on-year change in basis points and a categorical FOMC stance label (Easing, Easing-ZLB, Hold, Hold-ZLB, Tightening) that is used descriptively in the regime validation analysis.

Bauer–Swanson monetary policy surprises (henceforth BS-MPS) come from the dataset of Bauer and Swanson (2023). The dataset reports the unanticipated component of the federal funds rate at each FOMC meeting, identified as the change in the first eurodollar futures contract within a tight window around each FOMC announcement. We aggregate to annual frequency by summing per-meeting surprises within each calendar year and converting them to basis points. The number of FOMC meetings underlying each annual sum is also reported in the panel; the count is eight in most years but ranges from nine to thirteen in years containing intermeeting actions (for example, 2001, 2008, and 2020) and is higher still in the early 1990s—up to twenty-two in 1991—when the dataset records frequent intermeeting policy moves. The Bauer–Swanson dataset in the version we use ends in December 2023; consequently, BS-MPS is unobserved for 2024 in our panel, and any specification that includes it is estimated on the 1990–2023 sub-sample.

Kilian's index of global real economic activity (Kilian, 2009) measures the demand-driven component of commodity price movements as a deviation from the trend in dry bulk shipping rates. We use the updated monthly series from the author's website and aggregate to annual frequency by averaging monthly values within each calendar year. The variable enters the regressions in level form (it is already demeaned by construction).

VIX is the CBOE volatility index measured at daily close; we average daily values within each calendar year. It enters the regressions in level form. The mean across our sample is 19.5 points, with the high points in 2008 (32.7) and 2020 (29.3) reflecting the global financial crisis and the COVID-19 panic, respectively.

We also include the *real effective exchange rate* (REER) from the Bruegel database, constructed and maintained by Darvas (2021), updated to April 2026, on the 170-partner basket. We do *not* use REER as a primary regressor in the headline specification because REER is mechanically related to own inflation through its construction ($\text{REER} = \text{NEER} \times \text{home CPI}/\text{foreign CPI}$); including it on the right-hand side of an inflation regression creates simultaneity bias. We use REER instead in two roles: as an external validation indicator for the regime classification (see Section 4) and as a robustness check (Appendix A). The nominal effective exchange rate (NEER) does not have this construction problem and enters the regressions as a percentage change.

3.4. Fiscal Variables and the Cyclically Adjusted Measure

The fiscal balance variable in our panel is the overall fiscal balance as a share of GDP, taken from IMF Article IV reports compiled by GCC source institutions. In raw form, this measure mixes discretionary policy choice with the mechanical effect of oil price movements: a country whose oil revenues rise by 5 percent of GDP because of a Brent price rally has a larger fiscal balance regardless of any policy decision. To isolate the discretionary component, we construct a cyclically adjusted fiscal balance (henceforth CAF) as the country-specific residual from a regression of the overall fiscal balance (FB) on the contemporaneous Brent price and its one-year lag.

Country-by-country, the regression R^2 ranges from 0.12 (Kuwait) to 0.46 (Saudi Arabia), the contemporaneous Brent coefficient is positive in all six countries (statistically significant at the 1% level in Oman, Saudi Arabia, and the UAE; marginally significant in Bahrain, $p = 0.057$), and the residual series is uncorrelated with the Brent price by construction. We report the country-specific regression statistics in Appendix B. The CAF residual is our preferred measure of fiscal *stance* and enters the regressions as a level. Where we report results using the raw fiscal balance, we describe the variable explicitly.

3.5. Unit Conventions

To ensure comparability across specifications, all rate and share variables enter the regressions in percentage points; the two index variables—the Kilian real activity index and the VIX—enter in level form in their native index units. Most other series are expressed in percentage points, including CPI inflation, GDP growth, the fiscal balance, and government consumption, with the latter two measured as shares of GDP. Two series—M2 growth and non-oil real GDP growth—are reported in decimal form in their source database (for example, 0.030 rather than 3.0 percent); we rescale these to percentage points so that all variables share a common unit in the reported specifications, and we apply the same convention to the year-on-year change in Brent; trading partner inflation is constructed in percentage points directly (Section 3.3). The original series are retained alongside the rescaled versions to preserve reproducibility. The full panel and replication code were made available at submission in accordance with the journal’s data availability policy (see Supplementary Materials).

4. Methodology

4.1. Panel Regression Specification

Our headline empirical model is a country fixed-effects panel regression of CPI inflation on a set of external drivers, fiscal variables, domestic monetary aggregates, and real-activity controls. For country i in year t , the specification is

$$\pi_{it} = \alpha_i + \beta' X_{it} + \gamma' F_{it} + \delta' M_{it} + \varepsilon_{it}$$

where π_{it} is annual CPI inflation, α_i is a country fixed effect, X_{it} is the vector of external drivers (trading partner inflation, NEER percentage change, Brent year-on-year change, Bauer–Swanson monetary surprise, Kilian REA, VIX), F_{it} is the vector of fiscal variables (cyclically adjusted balance and government consumption as a share of GDP), and M_{it} contains the domestic monetary and real activity controls (M2 growth and non-oil GDP growth). All variables enter in percentage points where applicable, ensuring the coefficients are directly interpretable as the change in inflation (in percentage points) per one-percentage-point change in the regressor.

Two methodological choices in this specification deserve explanation. First, we use country fixed effects only, not country-plus-year fixed effects. The reason is that several of our external drivers—Brent year-on-year change, Bauer–Swanson monetary surprise, Kilian REA, VIX, partner inflation—vary across years but are constant across countries within each year. Adding year fixed effects would absorb these aggregate regressors and eliminate the very channels we want to estimate. The country fixed effects suffice to absorb time-invariant country heterogeneity (long-run inflation differentials, country-specific data conventions, structural fiscal institutional features), which is the relevant nuisance variation in our setting.

Second, inference in this panel must confront two distinct problems, and we treat them with distinct tools reported as co-primary. The first problem is the small number of clusters: with six countries, cluster-robust asymptotics are unreliable, and we therefore confirm the significance of our two headline coefficients—the partner inflation association and the cyclically adjusted fiscal balance—with a wild cluster bootstrap using the six-point weights of Webb (2014) and verify their stability in a leave-one-country-out exercise (Section 5.1). The second problem is contemporaneous cross-sectional dependence: because several regressors are aggregate within each year, common shocks induce correlation across countries that country-level clustering does not, by itself, accommodate. We therefore report Driscoll–Kraay standard errors (Driscoll & Kraay, 1998), which are robust to cross-sectional and serial correlation, alongside the entity-clustered standard errors for every headline model in Table 4, and we treat the two as jointly co-primary: the wild cluster bootstrap addresses the few-cluster problem, and the Driscoll–Kraay correction addresses common shock dependence, each covering the other’s weakness (Driscoll–Kraay asymptotics rely on the time dimension, which at 35 years is adequate but not long). Entity-clustered standard errors remain the tabulated baseline, consistent with Brandão-Marques et al. (2020) and similar panel work on emerging-market inflation transmission; two-way clustering is degenerate here since, with aggregate regressors, the time dimension has insufficient within-cluster variance, producing undefined or unreliable standard errors.

We progressively augment the specification across six headline models, denoted HM1 through HM6, to assess coefficient stability as variables enter. HM1 includes only the baseline external drivers (partner inflation, NEER, Brent). HM2 adds the Bauer–Swanson monetary surprise. HM3 adds Kilian REA and VIX. HM4 adds the cyclically adjusted fiscal balance and government consumption. HM5—the headline specification—adds the domestic monetary and real activity controls. The pattern of coefficient stability across HM1–HM5 is the central diagnostic of whether our headline finding is driven by a particular variable selection.

HM6 is the one specification of substantive methodological character rather than progressive augmentation, and it warrants explanation. Our external drivers are aggregate within each year and could in principle be affected, contemporaneously, by GCC inflation through global commodity price feedbacks or correlated sentiment in financial markets. To address this simultaneity concern, HM6 substitutes one-year-lagged versions of all six external drivers for their contemporaneous counterparts while leaving the fiscal variables

and the domestic controls in contemporaneous form. The trade-off is twofold. We lose the first year of the panel for each country (six observations), and we sacrifice contemporaneous transmission—co-movement that occurs within the same calendar year as the external shock. What HM6 can identify clearly is delayed transmission: the part of inflation transmission that takes a year to manifest. If our headline finding is driven entirely by within-year contemporaneous correlation, HM6 will return weak or insignificant external driver coefficients. If transmission has a substantial delayed component, HM6 will return coefficients of comparable magnitude to HM5. This is a more demanding test than commonly applied in GCC inflation work, and we report it explicitly so readers can assess the simultaneity concern themselves.

Table 4. Panel regression results: HM1–HM6 specifications.

Variable	HM1	HM2	HM3	HM4	HM5	HM6
Partner inflation (pp)	+0.351 *** (0.093) [0.265]	+0.369 *** (0.096) [0.254]	+0.285 *** (0.085) [0.240]	+0.457 *** (0.094) [0.224]	+0.464 *** (0.085) [0.197]	
NEER% change	−0.110 (0.082)	−0.088 (0.077)	+0.007 (0.049)	−0.045 * (0.027)	−0.049 (0.047)	
Brent YoY change (pp)	+0.011 *** (0.004)	+0.017 *** (0.006)	−0.002 (0.005)	−0.010 ** (0.004)	−0.009 (0.007)	
BS-MPS (bps)		−0.016 * (0.009)	−0.005 (0.008)	−0.011 (0.008)	−0.008 (0.007)	
Kilian REA			+0.022 *** (0.004)	+0.016 *** (0.006)	+0.006 *** (0.002)	
VIX			+0.020 (0.042)	+0.001 (0.045)	+0.053 ** (0.025)	
Fiscal balance, CAF				−0.025 ** (0.012) [0.022]	−0.050 *** (0.018) [0.024]	−0.040 *** (0.010) [0.026]
Gov consumption (% GDP)				−0.203 *** (0.057)	−0.151 * (0.077)	−0.089 (0.060)
Non-oil GDP growth (pp)					+0.187 (0.157)	+0.288 ** (0.118)
M2 growth (pp)					+0.086 * (0.044)	+0.026 (0.035)
Partner inflation, lag (pp)						+0.044 (0.059) [0.132]
NEER % change, lag						+0.019 (0.033)
Brent YoY change, lag (pp)						−0.023 *** (0.005)
BS-MPS, lag (bps)						+0.008 * (0.005)
Kilian REA, lag						+0.024 *** (0.005)
VIX, lag						−0.091 * (0.053)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Within R ²	0.100	0.112	0.240	0.319	0.399	0.470
N	175	169	169	165	165	165

Notes: Dependent variable is annual CPI inflation. All specifications include country fixed effects. Entity-clustered standard errors in parentheses; Driscoll–Kraay standard errors (Bartlett kernel, bandwidth two) in square brackets for the two headline coefficients, reported as co-primary inference alongside the wild cluster bootstrap in Section 5.1. Stars follow the entity-clustered standard errors. HM6 replaces the six contemporaneous external drivers with their one-year lags, while the fiscal and domestic controls remain contemporaneous (Section 4.1). HM5 is the headline specification. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$. Parentheses denote entity-clustered SEs; square brackets denote Driscoll–Kraay SEs (headline coefficients); stars follow entity-clustered SEs.

4.2. Construction of the Cyclically Adjusted Fiscal Balance

The fiscal balance obtained from the Article IV reports conflates discretionary fiscal stance with the mechanical effect of oil price movements on government revenues. To isolate the discretionary component, we construct a cyclically adjusted fiscal balance, country by country.

For each country i , we estimate

$$fiscal_balance_{it} = a_i + b_i \cdot Brent_t + c_i \cdot Brent_{t-1} + v_{it}$$

by ordinary least squares on country-level time series, then take the residual v_{it} as our cyclically adjusted measure. The residual is uncorrelated with Brent prices by construction; this is what we mean when we describe it as “oil-orthogonal.” The country-level regression diagnostics are reported in Appendix B: the R^2 ranges from 0.12 (Kuwait) to 0.46 (Saudi Arabia), and the contemporaneous Brent coefficient is positive in all six countries, statistically significant at the 1% level in three and marginally significant in a fourth.

This construction is simpler than the IMF’s structural balance methodology, which would adjust for the output gap, decompose oil revenues into permanent and transitory components, and apply hydrocarbon-specific elasticity assumptions. We sacrifice that richness for transparency and replicability: every reader can reproduce our CAF series from the panel using two regressions per country. The cost is that the residual will absorb any non-oil cyclical movement in fiscal stance—a counter-cyclical fiscal stimulus during the COVID-19 recession, for instance, would show up partly in the residual rather than entirely in the cyclical component. We discuss the implications in Section 6 and report results using the raw fiscal balance as a robustness check.

One further caveat concerns inference. Because the CAF is itself an estimated residual, it enters the inflation regressions as a generated regressor (Pagan, 1984): the second-stage standard errors do not account for sampling error in the six first-stage regressions, and the reported precision of the CAF coefficient is correspondingly optimistic. Two features of our results limit the practical force of this concern. The raw fiscal balance—which involves no first stage—delivers the same sign and comparable significance in an otherwise identical specification (Appendix A.2), and the CAF coefficient itself is confirmed by wild cluster bootstrap and leave-one-country-out inference (Section 5.1). We address this issue directly in Section 5.1, where two two-stage bootstrap designs propagate first-stage sampling uncertainty into the inference and are reported alongside the second-stage-only benchmark: a parametric design, drawing the six first-stage coefficient vectors from their estimated sampling distributions under a Newey–West covariance appropriate for annual time-series dependence, and a nonparametric block-dependent wild design, in which year-level weights, common across countries and constant within four-year blocks, perturb the first-stage residuals so that the re-estimated coefficient vectors genuinely vary across replications while common-year dependence across countries is preserved.

4.3. Regime Classification via Unsupervised Clustering

We classify country-year observations into macroeconomic regimes using k-means clustering on a three-variable feature set: CPI inflation, government consumption as a share of GDP, and real GDP growth. These three variables capture the dimensions along which GCC macroeconomic states differ—inflation outcome, fiscal structural position, and aggregate cycle—without including any of the external drivers we want to test against the regimes in Section 4.4.

One property of this feature set deserves emphasis up front: CPI inflation—the dependent variable of the regression analysis—is itself one of the three clustering features. The states are therefore partly generated from the outcome, and any regression estimated within or interacted with them conditions, in part, on the dependent variable. This is a deliberate choice—a macroeconomic state that ignored the inflation outcome would be a strange object in a paper about inflation regimes—but it carries an inferential cost: coefficients estimated within an outcome-selected group are descriptive of that group rather than identified estimates of state-dependent transmission, which is one reason, alongside the small

boom-state sample, that we present the entire state-conditional analysis as exploratory. As a robustness check, we re-run the classification on the two non-outcome features alone (government consumption share and real GDP growth); the resulting partition agrees with the three-feature solution at an adjusted Rand index of 0.81 and recovers a high-growth state ($N = 15$, mean real GDP growth of 16.1 percent) whose mean inflation, 4.3 percent, exceeds that of the other two states (1.7 and 2.5 percent; ANOVA $F = 5.61$, $p = 0.004$) and which contains six of the eleven boom-state observations. The boom pattern is thus partly—though not fully—recoverable without using inflation as a feature.

The clustering uses k-means with k-means++ initialization, run with $n_{init} = 100$ to reduce sensitivity to starting values, on standardized features (mean zero, unit variance). The Kuwait observations for 1990–1993 and all observations for 2020 are excluded for the reasons described in Section 3.2: war-period and pandemic-period dynamics reflect macroeconomic shocks that the regime classification is not designed to capture. The clean sample for clustering contains 163 country-year observations.

To select the number of clusters K , we report four diagnostics: the within-cluster sum of squares (inertia), the silhouette score, the Calinski–Harabasz index, and Tibshirani’s gap statistic. Calinski–Harabasz peaks at $K = 3$, the silhouette score is maximized at $K = 4$, with $K = 3$ a close second (0.331 against 0.320), and inertia shows the standard elbow at $K = 3$. The gap statistic results, however, are at variance: the gap statistic decision rule prefers $K = 1$, indicating no statistically detectable cluster structure relative to a uniform null distribution.

We treat this disagreement transparently. The gap statistic is conservative against the partition structure when clusters are not sharply separated in feature space: it tests whether the data depart from a uniform distribution null, and overlapping clusters within a single mass can fail the test even when persistent partition structure exists. Bootstrap–Jaccard stability—the diagnostic of Hennig (2007), which assesses whether cluster assignments are reproduced under data resampling—is more sensitive to persistent partitions and is the criterion better suited to our descriptive aim of identifying regimes for use in subsequent regression conditioning. The bootstrap–Jaccard scores at $K = 3$ are 0.89, 0.89, and 0.79 for the three regimes (300 bootstrap replications), all above the 0.75 threshold for stability proposed by Hennig and Liao (2013). We therefore proceed with $K = 3$, while disclosing the gap statistic preference for $K = 1$ as a counter-indicator that future work with longer time series or different clustering aims may want to revisit.

We also assess robustness to the clustering method by comparing the $K = 3$ k-means solution against Gaussian mixture models. With spherical covariance—the assumption that most closely matches the equal variance implicit assumption of k-means—the GMM solution agrees with k-means at an adjusted Rand index of 0.79, indicating high method-of-method agreement. With full covariance, GMM produces a partition with one large cluster (148 of 163 observations) and two small high-variance clusters; the adjusted Rand index against k-means is 0.17, reflecting the genuinely different clustering model that full-covariance GMM implements. Both observations are honest. The GMM-full result is not a contradiction of our k-means classification but a reminder that allowing an arbitrary covariance structure can identify a different—and in our setting, less interpretable—partition. We report both and rely on the methodologically comparable spherical GMM result for our robustness conclusion.

The three clusters that emerge will be characterized and labeled in Section 5.

4.4. Testing for State-Dependent Transmission (Exploratory)

The exploratory question of this subsection is whether the regression coefficients in HM5 differ across the macroeconomic states in Section 4.3. To examine it, we estimate a pooled regression with regime indicators and regime $\times$ variable interactions:

$$\pi_{it} = \theta_0 + \theta_1 R_{2it} + \theta_2 R_{3it} + \sum_k \phi_k Z_{it,k} + \sum_k \psi_{k,2} (R_{2it} \times Z_{it,k}) + \sum_k \psi_{k,3} (R_{3it} \times Z_{it,k}) + \text{Country dummies} + \xi_{it}$$

where R_2 and R_3 are dummies for the second and third regimes (with the first regime as the reference) and k indexes a parsimonious four-variable specification (partner inflation, Kilian REA, government consumption, M2 growth). The parsimonious specification is required by the small size of the third regime (eleven country-year observations); a full ten-variable HM5 specification within each regime would not converge with usable degrees of freedom. Inference follows the conventions in Section 4.2: Driscoll–Kraay standard errors (Bartlett kernel, bandwidth two) as the baseline, with wild cluster bootstrap p -values (Cameron et al., 2008; Cameron & Miller, 2015) computed with the six-point weights of Webb (2014)—the appropriate choice with as few as six clusters—treated as the more conservative and more reliable guide wherever the two disagree.

We report two tests. First, the joint Chow-type F-test of the null hypothesis that all eight interaction terms equal zero—that is, that transmission is regime-invariant. Second, we report per-driver F-tests: for each variable Z_k , we test the joint null $\psi_{k,2} = \psi_{k,3} = 0$, which identifies which specific drivers vary across regimes. As Section 5.3 makes clear, we interpret rejection of these nulls cautiously: with only six clusters and a small boom-state group, conventional rejection is suggestive rather than decisive, and we confirm it against the wild cluster bootstrap before drawing any conclusion.

We report regime-specific coefficient estimates as a presentation device only (Appendix B.2); the pooled regression with interactions delivers the same point estimates, plus the joint tests, within a single internally consistent framework.

4.5. Sub-Period Analysis

The 2014 oil price collapse—Brent fell from \$109 in 2013 to \$52 in 2015—was the largest peacetime structural shock to GCC economies in our sample. We test whether the inflation generation process changed around this break by estimating the HM5 specification separately on a pre-2014 sub-sample (1990–2013) and a post-2014 sub-sample (2014–2023, with 2024 excluded for BS-MPS coverage as discussed in Section 3.3). Differences in coefficient magnitude, sign, and statistical significance between the two sub-samples constitute evidence of a structural break independent of the regime classification. We do not impose the break date statistically (e.g., via a Chow break test), preferring the substantive choice of 2014 as the visible structural moment. This choice also matches the convention adopted in other recent GCC papers.

The sub-period analysis and the macroeconomic state analysis answer related questions: the former asks whether the structural relationship changed at a known historical moment (2014), the latter whether coefficients differ across data-driven states. As we show in Section 5, the two are not independent in our sample—the high-inflation states all fall in the pre-2014 period—so the sub-period split and the state analysis are largely two views of the same underlying variation rather than separate dimensions of evidence. Section 5 reports both and is explicit about this overlap.

5. Results

5.1. Average Inflation Transmission: Panel Regression Results

Table 4 reports the panel regression results for the six headline specifications, HM1 through HM6. We take the headline specification HM5—which includes the full set of

external drivers, fiscal variables, and domestic controls—as the reference column for interpretation; HM1 through HM4 trace the stability of the coefficients as variables enter, and HM6 serves as the lagged-external-driver robustness check. All specifications employ country fixed effects with entity-clustered standard errors. The estimation sample for HM5 comprises $N = 165$ country-years.

Six drivers are statistically significant at conventional levels in HM5 with economically sensible signs. Trading partner inflation enters with a coefficient of $+0.464$ ($t = +5.46$, $p < 0.001$), the most precisely estimated coefficient in the model and, in standardized terms, the largest among the external drivers (Table A8): a one-percentage-point rise in trade-weighted partner inflation is associated with GCC headline inflation higher by 0.46 percentage points within the same year. Kilian's index of global real activity enters with a coefficient of $+0.006$ ($t = +3.40$, $p = 0.001$), confirming a small but robust contribution from the global demand component of commodity price movements. The VIX enters with a coefficient of $+0.053$ ($t = +2.16$, $p = 0.03$), indicating that years of elevated global financial volatility are associated with higher GCC inflation. The cyclically adjusted fiscal balance enters with a coefficient of -0.050 ($t = -2.81$, $p = 0.005$): a one-percentage-point increase in discretionary fiscal surplus is associated with inflation 0.05 percentage points lower. Government consumption as a share of GDP enters with a coefficient of -0.151 ($t = -1.96$, $p = 0.05$): structurally larger public sectors are associated with somewhat lower inflation; we defer interpretation of this coefficient to Section 6.2, where a denominator effect complicates any structural reading. M2 growth enters at the borderline ($+0.086$, $t = +1.97$, $p = 0.05$), suggesting a marginal role for domestic monetary expansion.

Three drivers do not survive HM5 at the 5% level. The Bauer–Swanson monetary surprise enters with a small negative coefficient (-0.008 , $t = -1.11$)—the wrong sign for direct interest rate transmission to GCC inflation, which, on its own, is unsurprising given that the peg imports the US interest rate path rather than setting an independent one. Brent year-on-year change enters with a coefficient close to zero (-0.009 , $t = -1.24$); this reflects the inclusion of partner inflation, Kilian REA, and the cyclically adjusted fiscal balance, all of which absorb different oil cycle channels. NEER percentage change enters with a coefficient of -0.049 ($t = -1.04$), small and imprecisely estimated, consistent with the fact that NEER moves only modestly under a dollar peg.

The coefficient stability across HM1 through HM5 is informative. The contemporaneous partner inflation coefficient is positive in all five specifications, HM1 through HM5, and statistically significant at the 1 percent level in each, ranging from $+0.285$ (HM3, where the Kilian index first enters and absorbs part of the shared oil cycle variation) to $+0.464$ (HM5). Under the co-primary Driscoll–Kraay inference (square brackets in Table 4), significance consolidates as the specification fills out: the coefficient is imprecisely estimated under Driscoll–Kraay in the sparse specifications HM1 through HM3 (t between $+1.2$ and $+1.5$) and significant at the 5 percent level in HM4 and HM5 ($t = +2.04$ and $+2.35$), a pattern consistent with the common shock component of the sparse specifications' residuals being absorbed by the global controls. The lagged specification HM6 replaces the variable with its one-year lag and is discussed below. The Kilian REA coefficient, once it enters in HM3, is similarly robust across HM3–HM5. The cyclically adjusted fiscal balance, once introduced in HM4, retains its sign and significance through HM5. By contrast, Brent year-on-year and NEER coefficients shift sign or fall to insignificance as additional controls enter, suggesting that their apparent role in simpler specifications is largely subsumed by partner inflation, Kilian REA, and the cyclically adjusted fiscal balance.

The HM6 specification, which substitutes one-year lags for the six contemporaneous external drivers, delivers a nuanced picture. The lagged Kilian REA enters with a coefficient of $+0.024$ ($t = +5.27$, $p < 0.001$)—substantially larger and more precisely estimated than

the contemporaneous coefficient, suggesting that the global real activity channel transmits to GCC inflation with a one-year delay. The lagged Brent change enters negatively and precisely (-0.023 , $t = -4.47$, $p < 0.001$): conditional on lagged global activity and lagged partner prices, a year of rising oil prices is followed by lower GCC inflation the next year. Two features of the setting make this reversal intelligible rather than anomalous. It matches the sign pattern of the CAF construction itself, where the lagged Brent coefficient is negative in five of six countries (Appendix B.1)—fiscal positions partially retrench in the year after an oil price movement—and it is consistent with the administered price and subsidy adjustments that follow, rather than coincide with, the oil cycle. We flag it as a finding that the annual frequency cannot decompose further. The cyclically adjusted fiscal balance retains its negative sign and significance under lagged external drivers ($t = -4.03$, $p < 0.001$), and non-oil real GDP growth is also significant ($t = +2.44$, $p = 0.01$). By contrast, the lagged partner inflation coefficient is not statistically significant ($+0.044$, $t = +0.74$, $p = 0.46$), indicating that the contemporaneous $+0.464$ estimate captures within-year transmission that does not persist into the following year; the remaining lagged drivers (NEER, BS-MPS, VIX) are individually insignificant. HM6 therefore identifies a clean distinction between channels: the global real activity channel transmits with delay and is robust to anti-simultaneity controls; the partner inflation channel transmits within-year, and we cannot rule out same-year reverse causation from GCC inflation to global price aggregates. We read this distinction as descriptive—the imported price and global activity channels operate on different timescales—and it sharpens the interpretation of external transmission we develop in Sections 5.3 and 6.

Because the panel has only six clusters, we subject the two coefficients that carry the paper's argument—the partner inflation association and the cyclically adjusted fiscal balance—to inference that does not rely on six-cluster asymptotics. Under a wild cluster bootstrap with the six-point weights of Webb (2014), the appropriate choice with as few as six clusters (9999 replications), the partner inflation coefficient has a bootstrap p -value of 0.009 and the cyclically adjusted balance coefficient has a p -value of 0.070—the former comfortably significant, the latter just beyond the conventional 5 percent threshold under this design, while remaining significant at the 5 percent level under entity-clustered, Driscoll–Kraay (Table 4), and lagged-driver (HM6) inference. Both are also stable in a leave-one-country-out exercise: the partner inflation estimate ranges from $+0.42$ to $+0.56$ across the six re-estimations, and the cyclically adjusted balance estimate from -0.039 to -0.096 , with no single country driving either result. These two findings are therefore robust to the small cross-section, with the stated caveat on the fiscal coefficient's bootstrap margin; we rely on them as the paper's principal results. Appendix C completes the picture: the progressive specifications re-estimated on the common sample (Table A9) confirm that coefficient movement across Table 4 reflects the introduction of controls rather than sample composition, and stationarity, multicollinearity, correlation, serial correlation, and observation-level influence diagnostics (Tables A10 and A11 and accompanying text) raise no concern that alters the reading of the headline estimates. On integration properties specifically, the headline partner inflation series rejects the unit-root null decisively (Maddala–Wu $\chi^2(12) = 27.47$, $p = 0.007$; Table A10); the one series that does not reject is the Kilian real activity index ($p = 0.142$), a business cycle frequency measure for which the annual sample test has limited power. Rather than rest on that argument, we verify directly that the central association is not an artifact of regressor persistence: a first-difference specification of the full headline model—in which every variable is unambiguously stationary—delivers a partner inflation coefficient of $+0.695$ (Driscoll–Kraay $t = +4.74$; wild cluster bootstrap $p = 0.014$), and an ARDL(1,1) formulation delivers a long-run partner association of $+0.485$ (delta-method $t = +2.95$), nearly identical to the levels estimate (Table A12). Because the

cyclically adjusted balance is a generated regressor, we additionally report a two-stage bootstrap that propagates first-stage sampling uncertainty into the inference (Section 4.2; Table A5). Two valid designs bracket the answer: drawing the first-stage coefficient vectors from their estimated sampling distributions, with a Newey–West covariance appropriate to annual time-series dependence, yields $p = 0.076$, and a nonparametric block-dependent wild design—year-level weights common across countries, constant within four-year blocks, under which the first-stage coefficients genuinely vary—yields $p = 0.037$, against $p = 0.070$ when only the second stage is re-sampled. Thus, across the corrected designs, the fiscal coefficient sits at or just beyond the conventional 5 percent threshold—a position we report as it stands: the association is marginal under the most conservative valid inference, survives with essentially unchanged magnitude in every design, and is corroborated by the raw balance specification, which involves no first stage (Appendix A.2, $t = -2.26$), and by the lagged-driver specification (Table 4, HM6).

Next, we address the construction of the partner inflation series itself. The headline series uses fixed 2010–2019 import weights, which embed forward-looking information into earlier observations and cannot reflect the pronounced shift in GCC import baskets toward Asia over the sample—a mismeasurement that, under classical assumptions, attenuates the estimated association toward zero. We therefore rebuild the series with country-specific weights computed from trailing five-year import shares (Section 3.3) and re-estimate the full specification (Table A6). The association strengthens rather than weakens: the full-sample coefficient rises to $+0.655$ (SE 0.173), and the sub-period contrast sharpens— $+2.074$ (SE 0.444) before 2014 and -0.132 (SE 0.349, not significant) after—while the fiscal coefficient is essentially unchanged (-0.042). Both the headline static ex-GCC series and this rolling-weight alternative are built from the same sources under the same construction rules, so the comparison varies only the weighting scheme. This pattern is exactly what attenuation under fixed weights predicts, and it indicates that our headline estimate is, if anything, conservative.

Turning to dynamics, the static specification omits lagged inflation on Nickell bias grounds; at our panel length, however, that bias is of order $1/T$ and small, so we report the sensitivity directly. Adding one lag of inflation (Table A7; $N = 163$) yields a persistence coefficient of $+0.372$ (SE 0.078) and leaves both headline associations intact: the partner coefficient is $+0.423$ (SE 0.123) and the fiscal coefficient is -0.037 (SE 0.011), with implied long-run associations of $+0.674$ and -0.058 , respectively. Thus, the contemporaneous estimates are not artifacts of omitted persistence.

5.2. Structural Break Around the 2014 Oil Price Collapse

Table 5 reports the HM5 specification estimated separately on the pre-2014 sub-sample (1990–2013) and the post-2014 sub-sample (2014–2023). The two sub-samples deliver economically distinct pictures.

In the pre-2014 sub-sample, the inflation-generating process is dominated by external and monetary channels. The partner inflation association is $+0.909$ ($t = +4.25$, $p < 0.001$), nearly twice the full-sample HM5 estimate. M2 growth enters with a positive coefficient ($+0.159$, $t = +3.31$, $p = 0.001$), suggesting that domestic monetary expansion contributed materially to inflation during the pre-collapse era. The VIX coefficient is also strong ($+0.142$, $t = +2.75$, $p = 0.006$). The cyclically adjusted fiscal balance, government consumption, and non-oil GDP growth do not enter significantly in the pre-2014 sub-sample.

Table 5. Sub-period regressions (pre- and post-2014).

Variable	Pre-2014 (1990–2013)	Post-2014 (2014–2023)
Partner inflation (pp)	+0.909 *** (0.214)	−0.430 (0.364)
NEER % change	−0.028 (0.059)	+0.065 (0.270)
Brent YoY change (pp)	−0.002 (0.013)	+0.024 (0.017)
BS-MPS (bps)	+0.003 (0.007)	+0.045 (0.040)
Kilian REA	+0.010 *** (0.003)	−0.022 (0.017)
VIX	+0.142 *** (0.052)	+0.211 *** (0.057)
Fiscal balance, CAF	−0.015 (0.026)	+0.067 (0.045)
Gov consumption (% GDP)	−0.123 (0.114)	+0.023 (0.103)
Non-oil GDP growth (pp)	+0.086 (0.149)	+0.377 ** (0.164)
M2 growth (pp)	+0.159 *** (0.048)	−0.074 (0.044)
Country FE	Yes	Yes
Within R^2	0.495	0.527
N	106	59

Note: HM5 specification estimated separately on the pre-2014 (1990–2013) and post-2014 (2014–2023) sub-samples. Country fixed effects included. Entity-clustered standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$.

In the post-2014 sub-sample, the picture changes. The partner inflation association is no longer statistically significant (-0.430 , $t = -1.18$); its coefficient even shifts sign at the point estimate, although the standard error is wide. M2 growth becomes insignificant (-0.074 , $t = -1.67$). What matters post-2014 is non-oil real GDP growth ($+0.377$, $t = +2.30$, $p = 0.02$) and the VIX ($+0.211$, $t = +3.72$, $p < 0.001$). The pre-2014 within R^2 is 0.495; the post-2014 within R^2 is 0.527. Both periods deliver well-fitting models, but the drivers are different.

The qualitative reading is straightforward: before 2014, GCC inflation was largely an externally driven phenomenon—partner inflation, monetary expansion, and global financial conditions accounted for most of the variation. After 2014, with oil prices roughly halving and the GCC fiscal authorities tightening their discretionary stance, the inflation process became more domestically driven, with non-oil real activity emerging as the leading cyclical correlate. The VIX is the only variable significant in both sub-samples, which suggests that global financial volatility transmits to GCC inflation through a channel—perhaps capital flow spillovers or risk premium effects on imported goods—that is independent of the oil and monetary regimes.

We read the sub-period contrast as evidence that the inflation-generating process changed around the 2014 oil price collapse, the largest commodity shock in our sample. A formal test supports this: interacting partner inflation with a post-2014 indicator in the full-panel HM5 specification yields a negative, significant interaction under clustered standard errors (-1.06 , $t = -3.33$, $p = 0.001$), indicating that the association was materially weaker after 2014 than before, and the interaction survives the wild cluster bootstrap ($p = 0.010$), so the structural break stands as a result in its own right rather than only as corroboration of the broader account. The post-2014 window also contains discrete policy shocks—value-added taxes introduced in Saudi Arabia and the UAE in 2018 (with the Saudi rate tripled in 2020), in Bahrain in 2019, and in Oman in 2021, alongside the 2020 pandemic year. Controlling for these directly leaves the account intact: indicator variables for the VAT-introduction years and the pandemic year enter the full-panel HM5 specification with the expected signs—the VAT indicator at $+1.76$ percentage points ($t = +2.43$), a price-level step, and the pandemic indicator at -1.50 ($t = -1.79$, marginal)—while the partner inflation and cyclically adjusted balance coefficients are essentially unchanged ($+0.480$ and -0.053 , both $p < 0.01$). Within the post-2014 sub-sample, the qualitative pattern in Table 5 likewise survives the VAT indicators: partner inflation remains insignificant while non-oil growth and the VIX remain significant. We also note that the pre-2014 and post-2014 sub-samples are not independent

of the macroeconomic state analysis in Section 5.3: the high-inflation episodes that the clustering isolates all fall in the pre-2014 period, so the sub-period and state views are two windows onto the same underlying variation rather than two separate findings.

5.3. Exploratory Analysis: Macroeconomic States and the Possibility of State-Dependent Transmission

Having established the sub-period contrast in Section 5.2, we ask, exploratorily, whether transmission coefficients differ across data-driven macroeconomic states—a descriptive cross-check on the same underlying variation rather than an independent line of evidence, since the high-inflation states the clustering recovers all fall in the pre-2014 window. We classify the 163 clean country-year observations using k-means on three features—inflation, government consumption share, and real GDP growth (Section 4.3)—and stress at the outset that the third group is small and the state-dependence it appears to generate is not robust; the analysis is reported as motivation for future work, not as an established result. Table 6 characterizes the three states. Regime 1 (High Fiscal Stability; 61 country-years) pairs low inflation (1.66 percent on average) with a large public sector (25.2 percent of GDP) and modest growth; Regime 2 (Moderate Stability; 91 country-years) pairs low inflation (1.84 percent) with a smaller public sector (15.8 percent) and stronger growth; and Regime 3 (Boom Inflation; 11 country-years) combines elevated inflation (10.18 percent) with very rapid growth (13.7 percent).

Table 6. Exploratory macroeconomic state characterization.

Regime	N	Mean Inflation (%)	Mean Gov Cons (% GDP)	Mean GDP Growth (%)	Bootstrap–Jaccard
R1 (High Fiscal Stability)	61	1.66	25.21	2.54	0.889
R2 (Moderate Stability)	91	1.84	15.81	5.01	0.886
R3 (Boom Inflation)	11	10.18	15.0	13.74	0.793

Note: Means computed on the clean sample of 163 country-years used for clustering. Bootstrap–Jaccard scores are based on 300 re-samples; values above ~0.75 indicate stable clusters. R1, R2, and R3 denote the three states ordered by mean inflation. This analysis is exploratory; see Section 5.3.

The classification is stable and externally validated, with disclosed counter-indicators. The bootstrap–Jaccard scores are 0.889, 0.886, and 0.793 for the three regimes (300 replications), all above the 0.75 threshold of Hennig and Liao (2013), and a spherical covariance Gaussian mixture—the model comparable to k-means—recovers the same partition (adjusted Rand index of 0.79), though a full-covariance mixture does not (0.17), and the gap statistic and the Bayesian information criterion prefer one and two groups, respectively: the boom group is a stable but not sharply separated feature of the data. Analysis of variance tests on thirteen indicators not used in the clustering show that eight differ across regimes at the 1 percent level, led by the oil cycle, exchange rate, and US monetary measures that constitute genuinely independent validation. Notably, partner inflation and both fiscal balance measures do not differ across regimes: the states are not defined by the size of the imported price shock or by discretionary stance, so any cross-state variation in coefficients would reflect differences in transmission rather than in shocks.

Table 7 reports the state-conditional estimates: the parsimonious four-variable specification, estimated as a pooled panel regression with country fixed effects, regime indicators, and regime $\times$ variable interactions, with implied per-regime coefficients, delta-method standard errors, and per-driver interaction F-tests.

Table 7. Exploratory state-conditional regressions.

Variable	R1	R2	R3	F (State-Dep)
Partner inflation (pp)	+0.173 (0.155)	+0.280 (0.180)	+1.384 * (0.752)	2.42 *
Kilian REA	+0.001 (0.007)	+0.002 (0.005)	+0.011 (0.013)	0.22
Gov consumption (% GDP)	+0.047 (0.059)	−0.175 * (0.093)	−0.199 (0.190)	3.29 **
M2 growth (pp)	+0.144 *** (0.055)	−0.005 (0.045)	+0.157 *** (0.043)	4.19 **
Joint F-test (eight interactions)				6.46 ***
N (regime)	61	91	11	163
Country FE	Yes	Yes	Yes	Yes

Note: Pooled specification with regime $\times$ variable interactions, country fixed effects, and Driscoll–Kraay standard errors (Bartlett kernel, bandwidth 2) in parentheses. F-tests are joint state-dependence tests of the R2 and R3 interactions per driver; the joint test covers all eight interactions, $F(8, 143)$. Inference on the boom-state (R3) interaction is not robust to the wild cluster bootstrap ($p = 0.22$) or to leave-one-country-out re-estimation; see Section 5.3. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

The pattern, read descriptively, locates the apparent state-dependence in the price and monetary channels: the partner inflation coefficient is modest in the two stable states (+0.173 in R1, +0.280 in R2) and much larger in the boom group (+1.384), the M2–growth interaction behaves similarly ($F = 4.19$), and the joint test of regime invariance rejects under conventional clustered inference ($F(8, 143) = 6.46$)—though the per-driver partner inflation test is itself only marginal ($F = 2.42$, $p = 0.09$)—while government consumption and the Kilian index show little cross-state variation. This apparent amplification is not robust, and the reason is concentration: six of the boom group’s eleven observations are from Qatar and most of the remainder fall in 2008, so removing Qatar alone destabilizes the estimate in a leave-one-country-out exercise, a wild cluster bootstrap appropriate to six clusters does not reject the null of no amplification ($p = 0.22$), and, because inflation is itself a clustering feature, the within-group coefficients condition partly on the outcome (Section 4.3). Thus, the data are consistent with elevated boom-period transmission but cannot establish it; we carry this pattern into Section 6.3 as a hypothesis for future research, with within-regime estimates reported for completeness in Appendix B.2.

5.4. Discretionary Fiscal Stance Versus Structural Fiscal Size

The empirical results support a distinction between the *stance* and *size* dimensions of fiscal policy in GCC inflation transmission. The cyclically adjusted fiscal balance enters HM5 with a coefficient of -0.050 and a t-statistic of -2.81 . The raw fiscal balance, in a comparable specification reported in Appendix A, also enters significantly (coefficient -0.046 , $t = -2.26$), but the interpretation is different. The raw balance coefficient conflates discretionary policy with the mechanical movement of revenues against the oil price; the CAF measure isolates the discretionary component. The two measures deliver similar precision in our sample, but the CAF measure is the right object for testing demand channel predictions, which speak to discretionary policy rather than to the passive accounting of oil revenues.

Government consumption as a share of GDP enters HM5 with a coefficient of -0.151 and a t-statistic of -1.96 —the structural fiscal effect that prior GCC evidence has documented. Our headline regression confirms this finding while showing it survives the inclusion of the discretionary stance measure. The two fiscal variables operate through different channels: government consumption captures the long-run influence of public sector size on aggregate prices, while the cyclically adjusted balance captures discretionary policy choice on top of that structure. Both signs are negative. We discuss the interpretation of these two fiscal coefficients in Section 6.

6. Discussion

6.1. *The External Price Association and Its Decline After 2014*

Our central empirical result is that trading partner inflation is associated with GCC inflation, with an average contemporaneous coefficient of about 0.46 (Table 4, HM5), an estimate that is positive and significant across the progressive specifications HM1–HM5, survives Driscoll–Kraay and Webb-weight wild cluster bootstrap inference (bootstrap $p = 0.009$), and barely moves when any single country is removed (leave-one-country-out range 0.42 to 0.56). For a set of economies whose nominal exchange rate against the dollar is fixed, this is the expected channel: with the monetary instrument committed to the peg, imported price pressure transmitted through dollar-invoiced trade is expected to be a first-order driver of domestic inflation. The magnitude is economically meaningful without being implausible—a one-percentage-point rise in trade-weighted partner inflation is associated with GCC inflation roughly half a point higher, conditional on oil prices, US monetary conditions, global activity, and domestic fiscal and monetary stance.

This average, however, conceals an important change over time. Estimating HM5 separately before and after the 2014 oil price collapse, the partner inflation association falls from about +0.91 in the pre-2014 sample to near zero (indeed, negative and insignificant) afterward, and a formal interaction test rejects coefficient equality at conventional levels (the partner inflation $\times$ post-2014 interaction is significant at $p = 0.001$ under clustered standard errors and at $p = 0.010$ under the wild cluster bootstrap). After 2014, the variables most associated with GCC inflation are domestic non-oil activity and global financial volatility (the VIX) rather than imported price pressure, and the VAT indicator robustness reported in Section 5.2 confirms that this shift is not an artifact of the tax and pandemic events concentrated in the later window. We read this as a structural shift in the inflation-generating process around the largest peacetime commodity shock in our sample: as oil revenues collapsed and the GCC moved through fiscal consolidation and (in several countries) the introduction of value-added taxation, the external price channel that had dominated the boom years gave way to domestic and financial channels. We present this as corroborating evidence of a changing transmission structure rather than as an independently decisive finding, given the bootstrap result on the interaction term.

6.2. *Discretionary Fiscal Stance and the Measurement of Fiscal Policy in Oil Economies*

Our second principal result is that discretionary fiscal stance is negatively associated with inflation. The cyclically adjusted fiscal balance enters HM5 with a coefficient of -0.050 ($t = -2.81$ clustered; Webb-weight bootstrap $p = 0.070$, with the two-stage designs in Table A5 that propagate the proxy's construction uncertainty bracketing this at 0.037 and 0.076—at or just beyond the margin of conventional significance under the most conservative valid inference we apply—and significant at the 5 percent level under clustered, Driscoll–Kraay, and lagged-driver inference; external validation of the measure itself is reported in Appendix B.1): a tighter discretionary stance, holding external drivers and structural fiscal size fixed, is associated with lower inflation. Like the partner inflation estimate, this coefficient is stable in the leave-one-country-out exercise. The sign is consistent with a conventional aggregate demand reading: discretionary fiscal expansion adds to demand and, in economies where the monetary stance is fixed by the peg, that demand pressure shows up in prices rather than being offset by an interest rate response.

The result also speaks to a measurement problem that runs throughout the resource economy literature. The headline fiscal balance in an oil exporter mixes two very different things: discretionary policy choice and the mechanical movement of revenues with the oil price. A coefficient on the raw balance therefore conflates the policy effect with the passive accounting of oil receipts and is correspondingly hard to interpret. Our cyclically adjusted

measure—the country-specific residual of the fiscal balance regressed on contemporaneous and lagged Brent prices—removes the oil-driven component by construction and leaves a residual we read as a proxy for the discretionary stance. In our data, the construction does substantial work for Saudi Arabia, Qatar, the United Arab Emirates, and Oman, where current and lagged oil prices explain between 28 and 46 percent of fiscal balance variation, and less for Kuwait, where they explain only about 12 percent (Appendix B.1). The raw balance enters a comparable specification with a similar sign and significance (-0.046 , $t = -2.26$; Appendix A.2), but the cyclically adjusted coefficient is the policy-relevant one because it maps onto fiscal choice rather than onto the budget's mechanical response to commodity prices. We offer this construction as a transparent, replicable tool: any reader can reproduce the series from two regressions per country, without the data demands of the IMF's full structural balance methodology (Bornhorst et al., 2011) or the oil and non-oil revenue decompositions the non-oil balance framework requires (Barnett & Ossowski, 2003).

Government consumption as a share of GDP enters HM5 with a separately negative coefficient (-0.155 , $t = -1.97$). The two fiscal variables thus operate through distinct channels—structural size and discretionary stance—and both point in the same direction. We do not press a structural interpretation of the government consumption coefficient for two reasons: it is less robust than the cyclically adjusted balance result under bootstrap inference, and the ratio carries a mechanical denominator effect—nominal GDP in these economies moves with the oil price, so the consumption share can rise in an oil bust with no change in spending, generating variation correlated with the oil cycle rather than with fiscal structure. A robustness specification that replaces the consumption share with the growth rate of real government consumption supports this caution: the real growth coefficient is small and insignificant (-0.020 , SE 0.022) on a reduced sample of 98 observations, in which Qatar drops out entirely for lack of real consumption coverage. We accordingly treat the consumption share coefficient as a control rather than a finding.

6.3. Exploratory Analysis: Does Transmission Vary by Macroeconomic State?

The exploratory state analysis in Section 5.3 suggests, without establishing, that transmission may amplify in boom conditions: the boom-state partner inflation coefficient is several times its stable-state counterpart, a pattern consistent with the Taylor (2000) mechanism—firms defend real revenues by passing shocks through more quickly in high-inflation environments—and with the emerging-market evidence of Caselli and Roitman (2019). For the reasons documented in Section 5.3—a boom group of eleven concentrated observations, non-robustness to the wild cluster bootstrap and to the removal of Qatar, and a classification that conditions partly on the outcome—we treat this as a hypothesis for future work with longer panels or a broader cross-section of dollar-pegged commodity exporters, not as a result of the present paper. What the hypothesis would imply, if confirmed, is nonetheless worth stating: the quiescence of imported price transmission since 2014 would be state-contingent rather than permanent, and monitoring frameworks should not extrapolate it into the next boom.

6.4. Limitations

Four limitations deserve explicit acknowledgement.

First, the panel is annual and the cross-section is small (six countries), which limits the precision of any state- or regime-conditional inference and is the reason we treat the boom-period analysis in Section 6.3 as exploratory. A longer or higher-frequency panel, or a broader set of dollar-pegged commodity exporters, would allow sharper tests of whether transmission is genuinely state-dependent.

Second, our identification of external drivers does not fully rule out within-year simultaneity: global commodity prices and financial market sentiment could respond to GCC conditions within the calendar year. The lagged-driver specification (HM6) mitigates but does not eliminate this concern, and a higher-frequency design with external instruments would be a natural extension. The economic case that this simultaneity is quantitatively minor is nonetheless strong: the six GCC economies are small relative to the aggregates on the right-hand side—their combined output is on the order of two percent of the world economy—and while global oil prices plausibly respond to Gulf supply conditions, it is far less plausible that the consumer price indices of the GCC's trading partners respond within the year to GCC demand; the oil price channel, where feedback is most credible, is separately controlled by Brent and the Kilian index.

Third, the cyclically adjusted fiscal balance, by construction, absorbs any non-oil cyclical movement in fiscal stance—a counter-cyclical stimulus during the COVID-19 downturn, for example, would appear partly in the residual; we exclude 2020 from the clustering for this reason, but the caveat applies to the fiscal coefficient throughout.

Fourth, the headline specifications are static, although annual CPI inflation is persistent. The lagged dependent variable sensitivity in Table A7 shows that both headline associations survive the inclusion of one lag of inflation with modest attenuation, and at our panel length, the Nickell bias such a specification introduces is of order $1/T$ and small; the static form remains the baseline because it keeps the progressive specifications comparable, and the persistence evidence is a further reason to read the estimated coefficients as conditional associations rather than structural parameters.

7. Conclusions

This paper has asked how inflation is transmitted in the Gulf Cooperation Council—through which external, fiscal, and domestic channels—and how their relative importance has shifted over time. We brought to bear an annual panel of six dollar-pegged hydrocarbon exporters spanning 1990–2024—an effective estimation sample of 1991–2023—and estimated country fixed-effects panel regressions of inflation on external drivers, fiscal variables, and domestic monetary and real activity aggregates, with inference confirmed by wild cluster bootstrap and leave-one-country-out checks appropriate to the panel's small cross-section.

We reach three principal findings. First, trading partner inflation is associated with GCC inflation with an average contemporaneous coefficient of about 0.46—the expected first-order channel in economies whose monetary stance is committed to a dollar peg—and this estimate is robust to wild cluster bootstrap inference, to the removal of any single country, to the inclusion of lagged inflation, and to re-weighting the partner basket with trailing country-specific trade shares, under which it strengthens to about 0.65. Second, discretionary fiscal stance is negatively associated with inflation: a tighter cyclically adjusted balance is associated with lower inflation, conditional on external drivers and structural fiscal size, and this result is stable across countries and inference methods, at or just beyond the margin of conventional significance under the most conservative bootstrap designs, including a two-stage bootstrap that accounts for the proxy's construction. Third, the inflation-generating process shifted around the 2014 oil price collapse: estimated separately by sub-period, the external price association of the pre-2014 years gives way to domestic real activity and financial volatility afterward—a contrast the macroeconomic state classification in Section 5.3 corroborates descriptively rather than independently. In exploratory analysis, we also find suggestive evidence that transmission amplifies during commodity revenue booms, but because the relevant episodes are few and concentrated in a small

number of country-years, this pattern is not robust to small-cluster inference, and we report it as a hypothesis for future research rather than as an established result.

Two implications follow. The first is for policy, and it is period-dependent. Over the full sample, the external price association is the leading external correlate of GCC inflation even under the peg; after 2014, domestic real activity and financial volatility carry the largest standardized weights. The prudent reading is not that imported inflation has ceased to matter: the full-sample association is strong and precisely estimated—larger still under trailing trade weights—and the pre-2014 evidence shows how forcefully external prices and domestic inflation have moved together in the past. Imported price pressure therefore warrants monitoring even in quiet periods; whether its force re-intensifies in boom conditions, as our exploratory evidence weakly suggests, is a question the present data cannot settle. Meanwhile, fiscal authorities retain the one discretionary instrument the peg leaves in domestic hands, an instrument our results show is systematically associated with inflation outcomes. The second is methodological: any panel analysis of fiscal policy in commodity-exporting economies should distinguish discretionary stance from structural size, and the country-by-country cyclical adjustment we propose is a tractable complement to the non-oil balance framework wherever consistent oil and non-oil revenue decompositions are unavailable.

These findings raise several questions for future research. First, the exploratory evidence of amplified transmission during commodity revenue booms deserves a sharper test than our six-country annual panel can provide; a longer panel, or a broader cross-section of dollar-pegged commodity exporters, such as Brunei or Trinidad and Tobago, would establish whether boom-period state dependence is a genuine and general feature or an artifact of a few episodes. Second, can higher-frequency identification or external instruments sharpen the channel-by-channel simultaneity question that our lagged specification (HM6) only partly resolves? Third, with another decade of data, would the GCC panel support a Markov-switching estimation that delivers transition probabilities the present descriptive analysis cannot? Each of these is a natural extension of the transmission framework set out here.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/economics14080308/s1>. Replication package (raw data extracts and the staged code pipeline reproducing all numerical results reported in the paper).

Funding: This research received no external funding.

Data Availability Statement: The dataset was compiled by the author from publicly available sources, including the World Bank World Development Indicators; the U.S. Energy Information Administration and the Federal Reserve Economic Data (FRED) database for the Brent crude price and the Federal Funds rate; Bauer and Swanson (2023) for monetary policy surprises; Kilian (2009) for the real economic activity index; the IMF International Trade in Goods by Partner Country database for bilateral import weights; the Bruegel real effective exchange rate database (Darvas, 2021); and the Chicago Board Options Exchange (CBOE) for the VIX. The cyclically adjusted fiscal balance was constructed by the author as described in Section 3 and Appendix B.

Acknowledgments: Claude-4.8 (Anthropic) and ChatGPT-5.5 (OpenAI) were used to assist with manuscript editing, including checking language clarity and consistency. All content was reviewed, verified, and revised by the author, who takes full responsibility for the integrity of the work.

Conflicts of Interest: The author declares no conflicts of interest.

Appendix A. Robustness Checks for HM5

This appendix reports two robustness checks for the headline HM5 specification: (i) Driscoll–Kraay standard errors, robust to cross-sectional and serial correlation, in place

of entity-clustered standard errors, and (ii) HM5 estimated with the raw fiscal balance in place of the cyclically adjusted (CAF) measure. Both specifications retain country fixed effects.

Appendix A.1. Driscoll–Kraay Standard Errors

Table A1 compares the headline HM5 specification (entity-clustered standard errors, the baseline in Table 4 of the main text) against the same specification re-estimated with Driscoll–Kraay (DK) standard errors. DK errors are robust to arbitrary cross-sectional correlation and serial correlation up to a specified lag (we use a Bartlett kernel with bandwidth 2). They are the natural robustness check for panels with aggregate regressors like ours, where cross-sectional correlation is induced mechanically by variables that are constant across countries within each year.

Table A1. HM5 with entity-clustered vs. Driscoll–Kraay standard errors.

Variable	HM5 (Entity-Clustered SE)	HM5 (Driscoll–Kraay SE)
Partner inflation (pp)	+0.464 *** (0.085)	+0.464 ** (0.197)
NEER % change	−0.049 (0.047)	−0.049 (0.088)
Brent YoY change (pp)	−0.009 (0.007)	−0.009 (0.010)
BS-MPS (bps)	−0.008 (0.007)	−0.008 (0.020)
Kilian REA	+0.006 *** (0.002)	+0.006 (0.005)
VIX	+0.053 ** (0.025)	+0.053 (0.046)
Fiscal balance, CAF	−0.050 *** (0.018)	−0.050 ** (0.024)
Gov consumption (% GDP)	−0.151 * (0.077)	−0.151 ** (0.074)
Non-oil GDP growth (pp)	+0.187 (0.157)	+0.187 (0.142)
M2 growth (pp)	+0.086 * (0.044)	+0.086 ** (0.038)
N	165	165
Within R ²	0.399	0.399
Country FE	Yes	Yes

Note: Both columns estimate the identical HM5 specification on the same sample; only the covariance estimator differs. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

Both specifications are estimated on the same sample ($N = 165$); the point estimates are therefore identical across the two columns, and only the standard errors differ. The pattern of statistical significance is broadly preserved: trading partner inflation, the cyclically adjusted fiscal balance, and government consumption remain significant under both error structures, while the partner inflation t-statistic falls from +5.46 (clustered) to +2.35 (Driscoll–Kraay)—a decline that reflects the wider standard errors Driscoll–Kraay delivers when cross-sectional correlation is genuinely present, as it is in our panel of small economies sharing external drivers. The VIX and the Kilian real activity index, significant under clustered errors, are no longer individually significant under Driscoll–Kraay. The substantive interpretation of HM5 is unchanged: a robust positive partner inflation association and a robustly negative discretionary fiscal stance association survive both error structures.

Appendix A.2. Raw Fiscal Balance in Place of Cyclically Adjusted

Table A2 compares HM5 with the cyclically adjusted fiscal balance (CAF, the headline measure) against an otherwise identical specification using the raw fiscal balance. The two measures deliver similar precision: the CAF coefficient is -0.050 ($t = -2.81$), while the raw coefficient is -0.046 ($t = -2.26$). Both are statistically significant at conventional levels.

Table A2. HM5 with cyclically adjusted (CAF) vs. raw fiscal balance.

Variable	HM5 with CAF	HM5 with Raw
Partner inflation (pp)	+0.464 *** (0.085)	+0.501 *** (0.081)
NEER % change	−0.049 (0.047)	−0.046 (0.040)
Brent YoY change (pp)	−0.009 (0.007)	−0.005 (0.006)
BS-MPS (bps)	−0.008 (0.007)	−0.008 (0.007)
Kilian REA	+0.006 *** (0.002)	+0.005 ** (0.002)
VIX	+0.053 ** (0.025)	+0.050 ** (0.024)
Fiscal balance, CAF	−0.050 *** (0.018)	—
Fiscal balance, raw	—	−0.046 ** (0.021)
Gov consumption (% GDP)	−0.151 * (0.077)	−0.195 ** (0.086)
Non-oil GDP growth (pp)	+0.187 (0.157)	+0.181 (0.163)
M2 growth (pp)	+0.086 * (0.044)	+0.084 * (0.043)
N	165	165
Within R ²	0.399	0.393
Country FE	Yes	Yes

Note: Entity-clustered standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

The two coefficients are not directly comparable in interpretation. The raw fiscal balance moves mechanically with hydrocarbon revenues; its coefficient mixes the discretionary policy effect with the passive accounting of oil receipts. The CAF residual purges this mechanical correlation, leaving only the discretionary stance component. The negative sign on both measures is consistent with the demand channel reading of the fiscal–inflation nexus discussed in Section 6.3, but the policy-relevant interpretation belongs to the CAF coefficient because that coefficient maps onto fiscal policy choice rather than onto the budget’s mechanical response to commodity prices.

One additional finding worth noting from this comparison: the partner inflation coefficient is slightly larger in the raw fiscal specification (+0.501 versus +0.464). This is because the raw fiscal balance does not absorb as much of the oil-cycle-driven inflation pressure as the CAF measure does, leaving more of the variation to be attributed to imported price pressure. The CAF construction therefore implicitly disciplines the partner inflation coefficient as well—another reason to prefer it as the headline measure.

Appendix B. Construction of the CAF Measure and Within-Regime Robustness

This appendix reports two pieces of supporting empirical material. B.1 documents the country-by-country regressions used to construct the cyclically adjusted fiscal balance (CAF) introduced in Sections 3.4 and 4.2. Appendix B.2 reports within-regime separate regressions, providing a methodologically comparable robustness check to the pooled-with-interactions specification reported in Table 7 of the main text.

Appendix B.1. Country-by-Country CAF Construction Regressions

Table A3 reports the regression $fiscal_balance_{it} = a_i + b_i \cdot Brent_t + c_i \cdot Brent_{t-1} + v_{it}$ estimated separately for each of the six GCC countries. The CAF measure used throughout the paper is the residual $\hat{v}_{it}$ from this regression, estimated on each country’s full available sample (34 country-years each). The Kuwait 1990–1993 observations, distorted by the invasion and reconstruction, are retained in the CAF construction but excluded from the macroeconomic state clustering in Section 5.3, as described in Section 3.2.

Table A3. Country-by-country CAF construction regressions.

Country	Intercept	Brent _t	Brent _{t-1}	R ²	Adj. R ²	N	RMSE
Bahrain	−3.001	+0.1165 * (1.98)	−0.1656 *** (2.82)	0.221	0.170	34	5.46
Kuwait	+0.941	+0.4354 (1.27)	−0.1105 (0.32)	0.115	0.058	34	31.83
Kuwait, with 1991–1992 indicator	+18.708	+0.2375 (1.12)	−0.1266 (0.60)	0.677	0.645	34	19.54
Oman	−2.081	+0.2775 *** (3.50)	−0.2299 *** (2.90)	0.284	0.238	34	7.35
Qatar	−6.291	+0.0530 (0.71)	+0.1242 (1.66)	0.407	0.368	34	6.94
Saudi Arabia	−7.134	+0.3998 *** (4.87)	−0.2791 *** (3.41)	0.460	0.425	34	7.61
United Arab Emirates	+1.182	+0.2438 *** (4.29)	−0.1928 *** (3.40)	0.378	0.338	34	5.26

Note: Absolute t-statistics in parentheses. The Kuwait variant adds an indicator for the 1991–1992 invasion-and-reconstruction years; the indicator enters at −106.9 ($t = -7.23$). The CAF series used throughout the paper is the residual from the two-regressor specification, estimated on each country's full available sample. *** $p < 0.01$ and * $p < 0.10$.

The contemporaneous Brent coefficient is positive in all six countries. It is statistically significant at the 5% level in Oman (+0.278, $t = +3.50$), Saudi Arabia (+0.400, $t = +4.87$), and the United Arab Emirates (+0.244, $t = +4.29$); marginally significant in Bahrain (+0.117, $t = +1.98$, $p = 0.057$); and not significant in Kuwait (+0.435, $t = +1.27$) or Qatar (+0.053, $t = +0.71$). The lagged Brent coefficient is negative in five of six countries—a partial reversion of the contemporaneous response in the year following an oil price movement, consistent with the typical pattern of one-year-ahead fiscal adjustment in commodity-exporting economies.

The within-country R^2 ranges from 0.12 (Kuwait) to 0.46 (Saudi Arabia). The lower R^2 values for Kuwait suggest that a substantial share of its fiscal balance variation is not explained by current and lagged Brent prices, in part because the early-1990s war years introduce large fiscal swings unrelated to the oil price. This does not undermine the validity of the CAF construction: the residual is, by construction of the OLS procedure, orthogonal to both $Brent_t$ and $Brent_{t-1}$ in every country. What the residual measures is the within-country variation in fiscal balance that is not captured by the two-Brent specification—variation that, for the purposes of our inflation regressions, we treat as the discretionary stance component.

As a diagnostic on the war-period leverage, Table A3 also reports a variant of the Kuwait regression that adds an indicator for the 1991–1992 invasion-and-reconstruction years. The indicator enters at −106.9 ($t = -7.23$), the R^2 rises from 0.115 to 0.677, and the RMSE falls from 31.8 to 19.5—confirming that the war years, not the oil price, account for the bulk of Kuwait's unexplained fiscal variation. The choice between the two constructions is immaterial for the regression analysis: the two residual series correlate at 0.875 on non-war observations, and the extreme war-year residuals never enter the estimation sample because missing M2 and NEER observations exclude Kuwait before 1993 under listwise deletion. We retain the two-regressor construction as the headline CAF for uniformity across countries and report the indicator-augmented variant for transparency.

The CAF coefficient in the headline HM5 regression (−0.050, $t = -2.81$) therefore identifies the inflation response to the country-specific, year-specific residual movement in fiscal balance that cannot be explained by the contemporaneous or lagged oil price. This is the empirical content of “discretionary fiscal stance” in our setting.

As external validation, we compare the CAF with the IMF's published non-oil fiscal balance for the GCC, assembled from four vintages of the Regional Economic Outlook for the Middle East and Central Asia (131 country-years over 2000–2024; general government non-oil balance in percent of non-oil GDP—the non-oil primary variant is not tabulated in that source). The relevant dimension for validation is variation rather than level: the regressions use only within-country variation, and the published series differs from the CAF

in both denominator and trend. In first differences, the pooled correlation between the CAF and the published non-oil balance is +0.44, against +0.28 for the raw fiscal balance—the CAF tracks year-to-year movements in the IMF’s own measure appreciably better than the unadjusted balance does. The country-level pattern is informative: the correlation is strongest for Kuwait (+0.64) and weakest for Oman (+0.14). We read the exercise as supporting the CAF’s role as a complement to, not a substitute for, full non-oil accounting.

Appendix B.2. Within-Regime Separate Regressions

Table A4 reports the parsimonious four-variable regression (partner inflation, Kilian REA, government consumption, M2 growth) estimated separately within each regime, with country dummies included where degrees of freedom allow. Standard errors are White-robust (HC1). The right-hand columns of Table A4 reproduce the implied per-regime coefficients from the pooled-with-interactions specification in Table 7 in the main text, for direct comparison.

Table A4. Within-regime separate regressions compared to pooled-with-interactions specification.

Variable	R1 (Within)	R2 (Within)	R3 (Within)	R1 (Pooled)	R2 (Pooled)	R3 (Pooled)
Partner inflation (pp)	+0.229 (0.172)	+0.216 (0.148)	+3.287 *** (1.065)	+0.173	+0.280	+1.384
Kilian REA	−0.001 (0.006)	+0.000 (0.003)	+0.002 (0.034)	+0.001	+0.002	+0.011
Gov. consumption (% GDP)	−0.049 (0.098)	−0.256 *** (0.068)	−0.165 (0.190)	+0.047	−0.175	−0.199
M2 growth (pp)	+0.134 ** (0.053)	−0.006 (0.036)	+0.176 (0.155)	+0.144	−0.005	+0.157
Country dummies	Yes	Yes	Yes	Yes (FE)	Yes (FE)	Yes (FE)
R^2	0.296	0.277	0.897	—	—	0.674 (within)
N	61	91	11	61	91	11

Note: Within-regime columns report White-robust (HC1) standard errors in parentheses. Pooled columns reproduce the implied per-regime coefficients from Table 7 for comparison. *** $p < 0.01$, ** $p < 0.05$.

The within-regime estimates broadly track the pooled-with-interactions estimates, with two patterns worth flagging. First, for partner inflation transmission, the qualitative finding is the same in both specifications: the association is modest and statistically marginal in R1 (+0.229 within-regime, +0.173 pooled) and R2 (+0.216, +0.280) and large and significant in R3 (+3.287 within-regime, +1.384 pooled). The within-regime point estimate for R3 is larger than the pooled-with-interactions estimate because the pooled specification shrinks the R3 coefficient toward the common baseline structure, while the within-regime specification lets R3 stand alone. Both specifications therefore show the same descriptive pattern—a much larger point estimate in the boom group—but, as Section 5.3 establishes, this amplification is not robust: it rests on eleven observations concentrated in Qatar and does not survive the wild cluster bootstrap ($p = 0.22$) or leave-one-country-out re-estimation. We report the within-regime results as descriptive corroboration of that exploratory pattern, not as independent evidence of state dependence.

Second, for government consumption, the within-regime estimates are significantly negative in R2 (−0.256, $t = -3.75$) but not in R1 (−0.049, $t = -0.50$) or R3 (−0.165, $t = -0.87$). The pooled-with-interactions specification gives a comparable picture, with the R1 coefficient close to zero (+0.047) and the R2 and R3 coefficients negative (−0.175 and −0.199). The headline finding from the main text—that the government consumption share enters with a negative coefficient—is sustained by both specifications, although in the within-regime view, it is concentrated in R2, the largest regime.

The R^2 values from the within-regime regressions are 0.30 (R1), 0.28 (R2), and 0.90 (R3). The R3 R^2 is unusually high because of the small sample size ($N = 11$) relative to the parameter count (four features + five country dummies + intercept): the R3 specification

has a single residual degree of freedom, and the tight fit should be read as a property of the small sample rather than as evidence of a particularly well-specified model. The pooled-with-interactions specification reported in Table 7, with its within R^2 of 0.67 and joint state-dependence test $F(8, 143) = 6.46$, provides the more reliable inferential framework. We report the within-regime results here as a robustness check, not as a competing specification.

Appendix C. Additional Robustness and Diagnostics

Table A5 reports the two-stage bootstrap for the cyclically adjusted fiscal balance; Table A6 re-estimates the headline specification with the trailing-weight partner inflation series; Table A7 adds lagged inflation; Table A8 reports standardized coefficients; Table A9 repeats the progressive specifications on the common estimation sample; Tables A10 and A11 report stationarity, multicollinearity, and correlation diagnostics; and Table A12 reports the stationary transformation robustness of the headline model. Serial correlation and influence results are reported in the text below Table A11. Inference throughout matches the main text: country fixed effects with entity-clustered standard errors, with significance assessed against the t distribution with five degrees of freedom.

Table A5. Bootstrap inference for the cyclically adjusted fiscal balance: second-stage-only and two-stage designs.

Bootstrap Design	Webb WCB p -Value
Second stage only (as in Section 5.1)	0.070
Two-stage, parametric first-stage draws (Newey–West covariance)	0.076
Two-stage, block-dependent wild first stage (common year draws)	0.037

Note: Restricted (null-imposed) wild cluster bootstrap; Webb (2014) six-point weights at the country level; $B = 9999$ replications. Observed coefficient -0.050 (entity-clustered $t = -2.53$, $N = 165$). The parametric design draws the six first-stage coefficient vectors from their estimated sampling distributions under a Newey–West (Bartlett, three-lag) covariance appropriate for annual time-series dependence and recomputes the cyclically adjusted balance in every replication. The block-dependent wild design, in the spirit of the dependent wild bootstrap of Shao (2010), draws Webb weights at the year level—common across countries and constant within four-year calendar blocks—perturbs the first-stage residuals, and re-estimates each first stage so that the first-stage coefficient vectors genuinely vary while common year dependence across countries is preserved; an independent country-level draw then perturbs the restricted second stage. Two degenerate designs are excluded: applying a common draw to both stages mechanically reimports the observed covariance, and applying one weight to a country’s entire first-stage residual vector leaves the re-estimated first-stage coefficients unchanged—because the residual is orthogonal to its regressors, the perturbed residual is an exact rescaling of the original—and therefore re-samples no first-stage parameter uncertainty.

Table A6. HM5 with the trailing-weight partner inflation series: full sample and sub-periods.

	Full Sample	Pre-2014	Post-2014
Trading partner inflation (trailing weights)	+0.654 (0.173) **	+2.074 (0.444) ***	−0.132 (0.349)
NEER (% change)	−0.071 (0.058)	−0.001 (0.065)	−0.003 (0.314)
Brent YoY change	−0.015 (0.008)	−0.020 (0.011)	+0.025 (0.021)
BS-MPS (annual sum)	−0.0065 (0.0086)	+0.0267 (0.0053) ***	+0.0281 (0.0427)
Kilian real activity	+0.009 (0.002) ***	+0.014 (0.003) **	−0.026 (0.021)
VIX	+0.040 (0.029)	+0.142 (0.041) **	+0.183 (0.066) **
Cyclically adjusted fiscal balance	−0.0416 (0.0185) *	+0.0170 (0.0269)	+0.0668 (0.0540)
Government consumption (% GDP)	−0.111 (0.091)	−0.047 (0.107)	+0.041 (0.108)
Non-oil GDP growth	+0.199 (0.170)	+0.061 (0.163)	+0.401 (0.197) *
M2 growth	+0.080 (0.050)	+0.140 (0.065) *	−0.073 (0.055)
N	165	106	59

Note: The partner inflation series is rebuilt with country-specific weights computed from trailing five-year import shares (IMF International Trade in Goods by Partner Country, May 2026 vintage), replacing the fixed 2010–2019 weights of the headline series; all other variables are as in Table 4. Country fixed effects; entity-clustered standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

Table A7. HM5 with lagged inflation.

	Coefficient (SE)
Inflation ($t - 1$)	+0.372 (0.078) ***
Trading partner inflation	+0.423 (0.123) **
NEER (% change)	−0.080 (0.049)
Brent YoY change	+0.009 (0.012)
BS-MPS (annual sum)	−0.0149 (0.0116)
Kilian real activity	+0.000 (0.003)
VIX	−0.001 (0.042)
Cyclically adjusted fiscal balance	−0.0366 (0.0111) **
Government consumption (% GDP)	−0.095 (0.050)
Non-oil GDP growth	+0.103 (0.130)
M2 growth	+0.078 (0.029) *
Long-run partner inflation association	+0.674
Long-run fiscal balance association	−0.058
N	163

Note: HM5 augmented with one lag of the dependent variable. Long-run associations are computed as the contemporaneous coefficient divided by one minus the persistence coefficient. Country fixed effects; entity-clustered standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

Table A8. Standardized (beta) coefficients: full sample and sub-periods.

	Full Sample	Pre-2014	Post-2014
Trading partner inflation	+0.246 ***	+0.421 **	−0.359
NEER (% change)	−0.056	−0.031	+0.086
Brent YoY change	−0.089	−0.011	+0.483
BS-MPS (annual sum)	−0.064	+0.021	+0.476
Kilian real activity	+0.125 **	+0.161 **	−0.503
VIX	+0.110	+0.267 *	+0.636 **
Cyclically adjusted fiscal balance	−0.152 *	−0.038	+0.225
Government consumption (% GDP)	−0.312	−0.239	+0.068
Non-oil GDP growth	+0.213	+0.086	+0.568
M2 growth	+0.232	+0.406 **	−0.180

Note: All variables standardized to unit variance within the relevant estimation sample before estimation; country fixed effects; significance from entity-clustered standard errors. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

Table A9. Progressive specifications HM1–HM5 estimated on the common HM5 sample ($N = 165$).

	HM1	HM2	HM3	HM4	HM5
Trading partner inflation	+0.390 **	+0.410 **	+0.316 **	+0.457 ***	+0.464 ***
NEER (% change)	−0.106	−0.090	+0.010	−0.045	−0.049
Brent YoY change	+0.010 *	+0.016 *	−0.003	−0.010 *	−0.009
BS-MPS (annual sum)	—	−0.0159	−0.0054	−0.0110	−0.0081
Kilian real activity	—	—	+0.022 ***	+0.016 *	+0.006 **
VIX	—	—	+0.015	+0.001	+0.053
Cyclically adjusted fiscal balance	—	—	—	−0.0250	−0.0499 *
Government consumption (% GDP)	—	—	—	−0.203 **	−0.151
Non-oil GDP growth	—	—	—	—	+0.187
M2 growth	—	—	—	—	+0.086

Note: Each column re-estimates the corresponding progressive specification in Table 4 on the fixed HM5 estimation sample, so coefficient movement across columns reflects the introduction of controls rather than changes in sample composition. Country fixed effects; significance from entity-clustered standard errors. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

Table A10. Stationarity tests and variance inflation factors.

	Test	<i>p</i> -Value	VIF
CPI inflation	MW $\chi^2(12) = 75.38$	0.000	—
Trading partner inflation	MW $\chi^2(12) = 27.47$	0.007	1.30
NEER (% change)	MW $\chi^2(12) = 85.32$	0.000	1.59
Brent YoY change	ADF (aggregate series)	0.000	2.11
BS-MPS (annual sum)	ADF (aggregate series)	0.000	1.69
Kilian real activity	ADF (aggregate series)	0.142	2.62
VIX	ADF (aggregate series)	0.002	1.50
Cyclically adjusted fiscal balance	MW $\chi^2(12) = 84.84$	0.000	1.25
Government consumption (% GDP)	MW $\chi^2(12) = 21.83$	0.040	1.55
Non-oil GDP growth	MW $\chi^2(12) = 68.70$	0.000	2.07
M2 growth	MW $\chi^2(12) = 28.24$	0.005	2.47

Note: Country-varying series are tested with the Maddala–Wu Fisher combination of country-level ADF statistics (constant, lag order by AIC up to two); series that are aggregate within each year are tested with a single ADF on the annual series. Variance inflation factors are computed on the within-transformed HM5 regressors, matching the fixed-effects estimator. All series are tested in the form in which they enter the regressions.

Table A11. Pairwise correlations among the external drivers (HM5 estimation sample).

	Partner Infl.	Brent YoY	BS-MPS	Kilian	VIX	NEER %
Partner infl.	1.000					
Brent YoY	0.224	1.000				
BS-MPS	0.161	0.459	1.000			
Kilian	0.241	0.512	0.089	1.000		
VIX	−0.063	−0.033	−0.396	0.128	1.000	
NEER %	−0.225	−0.491	−0.119	−0.492	0.071	1.000

Note: Lower triangle shown; computed on the $N = 165$ estimation sample.

Two further diagnostics complete the battery. First, a Wooldridge-type test for first-order serial correlation in the fixed-effects residuals—regressing the first-differenced HM5 residuals on their first lag and testing the coefficient against the value of minus one-half implied by serially uncorrelated errors—rejects the null (coefficient -0.129 , standard error 0.136 ; t against -0.5 of 2.73 , $p = 0.041$). Serial correlation is therefore present; it is accommodated by the entity-clustered and Driscoll–Kraay covariance estimators used throughout, and the lagged dependent variable specification in Table A7 addresses its source directly, with both headline associations intact.

Second, observation-level influence. Re-estimating HM5 dropping each of the 165 observations, in turn, moves the partner inflation coefficient within $[+0.388, +0.498]$ and the fiscal balance coefficient within $[-0.067, -0.042]$; no single observation is decisive for either result. Excluding the five most influential observations jointly leaves the partner coefficient positive at $+0.272$ and strengthens the fiscal coefficient to -0.072 . The single most influential observation for the fiscal coefficient is Kuwait 1993—the earliest Kuwaiti observation in the estimation sample—and its removal strengthens, rather than weakens, the result, complementing the war year discussion of Appendix B.1.

Finally, because the number of FOMC meetings varies across years, we re-estimated HM5 with the per-meeting mean surprise in place of the annual sum; the headline coefficients are unchanged ($+0.469$ and -0.050 , both significant at conventional levels), and the monetary surprise term remains insignificant under either aggregation.

Table A12. Stationary transformation robustness: first differences and ARDL(1,1).

	First Differences	ARDL(1,1)
Trading partner inflation (contemporaneous)	+0.695 (0.201) **	+0.807 (0.287) *
Trading partner inflation (one-year lag)	–	–0.542 (0.249) *
Inflation ($t - 1$)	–	+0.454 (0.096) **
Cyclically adjusted fiscal balance	+0.015 (0.031)	–0.037 (0.011) *
Long-run partner inflation association	–	+0.485 (0.164) *
Partner inflation: Driscoll–Kraay t	+4.74	–
Partner inflation: Webb WCB p	0.014	0.013
N	159	163

Note: The first-difference column estimates the full HM5 specification with every variable in one-year changes, country intercepts retained; the ARDL(1,1) column adds lagged inflation and lagged partner inflation to the levels HM5, with the long-run association $(b_0 + b_1)/(1 - \rho)$ and its delta-method standard error computed from the entity-clustered covariance. Entity-clustered standard errors in parentheses; significance against the t distribution with five degrees of freedom. Wild cluster bootstrap p -values use Webb (2014) six-point weights, $B = 9999$; in the ARDL column the lagged dependent variable is treated as predetermined within replications, so primary inference for that column is the clustered t and the delta method long-run test. The central external association survives both stationary transformations with magnitude comparable to the levels estimate; the fiscal coefficient, which strongly rejects a unit root in levels (Table A10), is a levels relationship and is correspondingly absent in differences. ** $p < 0.05$, * $p < 0.10$.

References

- Alsabhan, T. H. (2026). Internal and external determinants of inflation in GCC countries: Evidence from a panel PMG-ARDL model. *Economics*, 14(4), 107. [\[CrossRef\]](#)
- Alsamara, M., Mrabet, Z., & Dombrecht, M. (2018). Asymmetric import cost pass-through in GCC countries: Evidence from nonlinear panel analysis. *Economic Modelling*, 75, 432–440. [\[CrossRef\]](#)
- Alsamara, M., Mrabet, Z., & Hatemi-J, A. (2020). Pass-through of import cost into consumer prices and inflation in GCC countries: Evidence from a nonlinear autoregressive distributed lags model. *International Review of Economics & Finance*, 70, 89–101. [\[CrossRef\]](#)
- Barnett, S., & Ossowski, R. (2003). Operational aspects of fiscal policy in oil-producing countries. In J. M. Davis, R. Ossowski, & A. Fedelino (Eds.), *Fiscal policy formulation and implementation in oil-producing countries* (pp. 45–81). International Monetary Fund.
- Basher, S. A., & Elsamadisy, E. M. (2012). Country heterogeneity and long-run determinants of inflation in the Gulf Arab states. *OPEC Energy Review*, 36(2), 170–203. [\[CrossRef\]](#)
- Bauer, M. D., & Swanson, E. T. (2023). A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37, 87–155. [\[CrossRef\]](#)
- Bonhomme, S., & Manresa, E. (2015). Grouped patterns of heterogeneity in panel data. *Econometrica*, 83(3), 1147–1184. [\[CrossRef\]](#)
- Bornhorst, F., Dobrescu, G., Fedelino, A., Gottschalk, J., & Nakata, T. (2011). *When and how to adjust beyond the business cycle? A guide to structural fiscal balances* (IMF Technical Notes and Manuals No. 2011/02). International Monetary Fund.
- Brandão-Marques, L., Gelos, G., Harjes, T., Sahay, R., & Xue, Y. (2020). *Monetary policy transmission in emerging markets and developing economies* (IMF Working Paper No. 20/35). International Monetary Fund.
- Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3), 414–427. [\[CrossRef\]](#)
- Cameron, A. C., & Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2), 317–372. [\[CrossRef\]](#)
- Caselli, F. G., & Roitman, A. (2019). Nonlinear exchange-rate pass-through in emerging markets. *International Finance*, 22(3), 279–306. [\[CrossRef\]](#)
- Cevik, S., & Miryugin, F. (2025). It's never different: Fiscal policy shocks and inflation. *Comparative Economic Studies*, 67, 186–220.
- Cevik, S., & Teksoz, K. (2013). Lost in transmission? The effectiveness of monetary policy transmission channels in the GCC countries. *Middle East Development Journal*, 5(3), 1350018. [\[CrossRef\]](#)
- Choudhri, E. U., & Hakura, D. S. (2006). Exchange rate pass-through to domestic prices: Does the inflationary environment matter? *Journal of International Money and Finance*, 25(4), 614–639. [\[CrossRef\]](#)
- Cochrane, J. H. (2023). *The fiscal theory of the price level*. Princeton University Press.
- Darvas, Z. (2021). *Timely measurement of real effective exchange rates* (Bruegel Working Paper 15/2021). Bruegel.
- Driscoll, J. C., & Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, 80(4), 549–560. [\[CrossRef\]](#)

- Erbil, N. (2011). *Is fiscal policy procyclical in developing oil-producing countries?* (IMF Working Paper No. 11/171). International Monetary Fund.
- Espinoza, R. A., & Prasad, A. (2012). *Monetary policy transmission in the GCC countries* (IMF Working Paper No. 12/132). International Monetary Fund.
- Fareed, F., Rezghi, A., & Sandoz, C. (2023). *Inflation dynamics in the Gulf Cooperation Council (GCC): What is the role of external factors?* (IMF Working Paper No. 23/263). International Monetary Fund.
- Forbes, K., Hjortsoe, I., & Nenova, T. (2018). The shocks matter: Improving our estimates of exchange rate pass-through. *Journal of International Economics*, 114, 255–275. [\[CrossRef\]](#)
- Gopinath, G. (2015). *The international price system* (NBER Working Paper No. 21646). National Bureau of Economic Research.
- Gopinath, G., Boz, E., Casas, C., Díez, F. J., Gourinchas, P.-O., & Plagborg-Møller, M. (2020). Dominant currency paradigm. *American Economic Review*, 110(3), 677–719. [\[CrossRef\]](#)
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357–384. [\[CrossRef\]](#)
- Hasan, M., & Alogoskoufis, H. (2008). *Understanding the inflationary process in the GCC region: The case of Saudi Arabia and Kuwait* (IMF Working Paper No. 08/193). International Monetary Fund.
- Hennig, C. (2007). Cluster-wise assessment of cluster stability. *Computational Statistics & Data Analysis*, 52(1), 258–271. [\[CrossRef\]](#)
- Hennig, C., & Liao, T. F. (2013). How to find an appropriate clustering for mixed-type variables with application to socio-economic stratification. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 62(3), 309–369. [\[CrossRef\]](#)
- Kandil, M., & Morsy, H. (2011). Determinants of inflation in GCC. *Middle East Development Journal*, 3(2), 141–158. [\[CrossRef\]](#)
- Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3), 1053–1069. [\[CrossRef\]](#)
- Kim, C.-J., & Nelson, C. R. (1999). *State-space models with regime switching: Classical and Gibbs-sampling approaches with applications*. MIT Press.
- Leeper, E. M. (1991). Equilibria under ‘active’ and ‘passive’ monetary and fiscal policies. *Journal of Monetary Economics*, 27(1), 129–147. [\[CrossRef\]](#)
- Mohanty, M. S., & Klau, M. (2004). *Monetary policy rules in emerging market economies: Issues and evidence* (BIS Working Paper No. 149). Bank for International Settlements.
- Murshed, H., & Nakibullah, A. (2015). Price level and inflation in the GCC countries. *International Review of Economics & Finance*, 39, 239–252. [\[CrossRef\]](#)
- Nishi, M. (2025). K-means panel data clustering in the presence of small groups. *arXiv*. [\[CrossRef\]](#)
- Pagan, A. (1984). Econometric issues in the analysis of regressions with generated regressors. *International Economic Review*, 25(1), 221–247. [\[CrossRef\]](#)
- Sadraoui, T., & Mili, H. (2025). Dynamic interactions between oil price volatility and fiscal policy: Empirical evidence from GCC countries. *Economics*, 13(11), 334. [\[CrossRef\]](#)
- Shambaugh, J. C. (2004). The effect of fixed exchange rates on monetary policy. *Quarterly Journal of Economics*, 119(1), 301–352. [\[CrossRef\]](#)
- Shao, X. (2010). The dependent wild bootstrap. *Journal of the American Statistical Association*, 105(489), 218–235. [\[CrossRef\]](#)
- Sims, C. A. (1994). A simple model for study of the determination of the price level and the interaction of monetary and fiscal policy. *Economic Theory*, 4(3), 381–399. [\[CrossRef\]](#)
- Sims, C. A., & Zha, T. (2006). Were there regime switches in U.S. monetary policy? *American Economic Review*, 96(1), 54–81. [\[CrossRef\]](#)
- Sturm, M., Gurtner, F., & Gonzalez Alegre, J. (2009). *Fiscal policy challenges in oil-exporting countries: A review of key issues* (ECB Occasional Paper No. 104). European Central Bank.
- Tariq, A., Amin, A., & Ahmad, M. (2024). Deciphering the non-linear nexus between government size and inflation in MENA countries: An application of dynamic-panel threshold model. *Journal of Economic Structures*, 13(1), 14. [\[CrossRef\]](#)
- Taylor, J. B. (2000). Low inflation, pass-through, and the pricing power of firms. *European Economic Review*, 44(7), 1389–1408. [\[CrossRef\]](#)
- Villafuerte, M., & López-Murphy, P. (2010). *Fiscal policy in oil producing countries during the recent oil price cycle* (IMF Working Paper No. 10/28). International Monetary Fund.
- Webb, M. D. (2014). *Reworking wild bootstrap based inference for clustered errors* (Queen’s Economics Department Working Paper No. 1315). Queen’s University.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.