Inflation and Exchange Rate Shocks in the Presence of Parallel Markets*

Toni Oki†
Olalekan Bello‡

May 2026

Abstract

We study exchange-rate pass-through to consumer prices in dual-rate economies, where an official exchange rate coexists with a parallel market rate at a substantial premium, using monthly Nigerian data from 2014 to 2023. Prices respond strongly to the parallel rate but weakly to the administered official rate. In a structural VAR, a one-standard-deviation parallel-rate depreciation raises cumulative CPI by about 0.5% over the following year, implying a twelve-month structural pass-through of about 0.29. By contrast, a shock that raises the official rate does not raise prices; the central estimate is weakly negative. The asymmetry appears across major CPI sub-indices, is more pronounced when the parallel premium is wide, and is corroborated by an instrumental-variables local projection using Nigeria-specific oil-supply disruptions and by a joint OLS local projection. We interpret these findings through a framework with foreign-exchange rationing and repressed goods markets, in which the parallel rate is the marginal price of imports and the official rate operates indirectly through fiscal financing and subsequent money creation; the weakly negative response is consistent with this fiscal channel. These results suggest that central banks that focus primarily on the official rate may underestimate the sensitivity of inflation to exchange-rate movements in dual-rate regimes.

JEL Codes: D31, E31, F31, F61, O11, O19, O24.
Keywords: Exchange Rates, Parallel Markets, Pass-Through, Large Devaluations, Emerging Markets.

*We would like to thank Kenneth Rogoff, Laura Alfaro, Lawrence Summers and Jeremy Stein, for their invaluable support. We would also like to thank Isaiah Andrews, Alberto Cavallo, Gabriel Chodorow-Reich, Joseph Gagnon, Gita Gopinath, Ricardo Hausmann, Cullen Hendrix, Adnan Mazarei, Stephen Morris, Pablo Andrés Neumeyer, Marcus Noland, Adam Posen, Carmen Reinhart, Andrei Shleifer, Hillary Stein, James Stock, Elie Tamer, David Wilcox and Diego Perez, as well as participants of the Peterson Institute for International Economics Research Staff Lunch and two anonymous referees, for helpful suggestions and comments. We are also grateful to Winiboya Aboyure, Mark Ayiah, Eyad Elsafoury, and Lingbo Wang for outstanding research assistance, and to Molly and Domenic Ferrante Economics Research Fund and Pathways to Research and Doctoral Careers for financial assistance, as well as Nigeria's National Bureau of Statistics and AbokiFX for providing data used in this study. Declarations of interest: none.

†toni.oki@gmail.com
‡New York University, ob708@stern.nyu.edu

1 Introduction

Exchange rate pass-through to retail prices shapes the real effects of nominal exchange rate movements, the welfare consequences of devaluation, and the transmission of monetary policy in open economies. A literature beginning with Mussa (1986) studies the channels linking exchange rates to domestic prices, including dollar invoicing, distribution margins, and pricing-to-market (Burstein et al., 2005; Atkeson and Burstein, 2008; Burstein and Gopinath, 2014; Gopinath et al., 2020).

This literature has largely abstracted from parallel foreign exchange markets.[1] Parallel markets are secondary markets in which foreign exchange trades outside official channels, typically at a premium over the central bank rate. They arise under managed regimes when official supply is rationed. They are common: in 2021, 23% of countries operated some form of multiple currency practice or exchange restriction conducive to parallel-market activity (IMF, 2022). When an official rate coexists with a parallel rate at a meaningful premium, two questions follow. Which rate is the relevant price of foreign exchange for the transactions that drive retail prices? And how should pass-through be measured when parallel-rate movements may already have transmitted to inflation before the official rate adjusts?

We answer these questions for Nigeria, where the gap between the parallel and official rates has been persistent. We assemble a daily dataset of parallel-market quotes covering 2014 to 2023 and combine it with official rates, central-bank external reserves, and money and consumer prices from the Central Bank of Nigeria and the National Bureau of Statistics. Over the pre-unification estimation sample (November 2014 to May 2023), the parallel premium averages 30% and peaks above 77%.

We organize the analysis around the general-equilibrium framework of inflation under foreign exchange rationing and repressed goods markets developed by Cerra (2016). When dollars are rationed at an overvalued official rate, the marginal price of foreign exchange is set in the parallel market. A binding cash-in-advance constraint pins down contemporaneous prices through money chasing the available quantity of imports. Because importers obtain incremental dollars at the parallel rate, marginal import pricing reflects the parallel market even when inframarginal dollars are allocated at the official rate (i.e., the parallel rate equals the scarcity price of those imports). The official rate enters only through future money creation: a higher official rate raises the naira value of dollar revenues sold by the central bank, narrows the fiscal deficit, and reduces subsequent

[1] Some prior work considers parallel markets in the context of real exchange rate misalignment and deviations from purchasing power parity (e.g., Edwards, 1988; Sekkat and Varoudakis, 2000; Cavallo and Rigobon, 2016); our focus is the dynamics of pass-through to consumer prices.

monetary financing. Reserves operate the same way: when the central bank accumulates reserves, fewer dollars are absorbed from the private sector and money creation rises. The framework predicts that distinct underlying shocks—a parallel-market shock, an official-rate shock, an adverse oil-revenue shock, and a non-FX inflation shock—should leave different qualitative responses across the parallel rate, the official rate, money, prices, and reserves.

The two rates are jointly determined by common shocks, so we identify the inflation response to each using two complementary strategies. A structural VAR uses the qualitative responses implied by the model to recover the inflation response to each shock. An instrumental-variable local projection uses Niger Delta militant attacks on Nigerian oil infrastructure as an external supply shifter for the parallel rate. Both designs deliver the same asymmetry. A shock that depreciates the parallel rate raises cumulative CPI by about 0.5% over the following year, implying a structural pass-through from the parallel rate to consumer prices of about 0.29 at twelve months. A shock that raises the official rate produces zero or weakly-negative inflation, with the central estimate consistent with the fiscal-financing channel in the framework. An adverse oil-revenue shock raises both rates, expands money, drains reserves by roughly 3.6% at twelve months, and raises prices through both channels. The asymmetry survives at the food sub-index, holds across divisional CPI sub-indices, and is sharper in months with a wide parallel premium.

The paper contributes to two literatures. To the literature on exchange-rate pass-through in emerging markets, it shows that estimates relying on a single official rate can mismeasure pass-through in dual-rate regimes, where the relevant transaction price of foreign exchange is set in the parallel market. To the literature on parallel currency markets, it provides a quantitative reading of the fiscal-channel mechanism in a dual-rate economy, including its dependence on the size of the parallel premium.

The findings speak to a recurring policy question in emerging markets: whether central banks should delay an official devaluation to protect against inflation. When a meaningful parallel premium exists, much of the inflation has already passed through by the time the official rate moves, blunting the protective effect of waiting. The structural results sharpen this conclusion. An official-rate movement that is driven by CBN policy does not raise prices on net at intermediate horizons; the central estimate is weakly contractionary, consistent with the fiscal-financing channel in the framework. An official-rate movement that arrives through oil-revenue pressure raises prices through both the direct and money channels. The two events look similar in the data—in both cases the CBN devalues—but their consequences for inflation differ sharply.

Although the empirical analysis focuses on Nigeria, the mechanism is potentially rel-

evant to a much broader class of economies characterized by foreign-exchange rationing, segmented currency markets, and administered exchange rates. Parallel-market premia emerged repeatedly in economies such as Argentina, Egypt, Lebanon, Iran, and Venezuela during periods of reserve scarcity and capital controls. The Nigerian case therefore provides a tractable setting in which to study a broader question: which exchange rate matters for inflation when official and market exchange rates diverge persistently?

Related Literature. An empirical literature studies exchange rate pass-through to consumer prices in economies without a parallel foreign exchange market. Goldberg and Knetter (1997) survey the early evidence, and Campa and Goldberg (2005, 2010) provide cross-country estimates of pass-through into import and consumer prices. Burstein et al. (2005) show that large nominal devaluations in emerging markets produce limited consumer-price inflation, attributing the gap to local distribution margins and substitution toward lower-quality goods. A complementary theoretical literature studies pricing-to-market, currency choice, and dollar invoicing as drivers of pass-through (Engel, 2006; Atkeson and Burstein, 2008; Gopinath et al., 2010). Boz et al. (2019) and Gopinath et al. (2020) document the empirical centrality of the dollar in international price-setting. Itskhoki and Mukhin (2021) revisit the joint behavior of nominal and real exchange rates and rationalize key pass-through puzzles in a unified framework. Auer et al. (2021) report low retail pass-through in Switzerland after the 2015 franc appreciation. Cravino and Levchenko (2017) find that pass-through is heterogeneous across income deciles, with poorer households facing larger price increases.

A second literature, on currency misalignment and purchasing power parity, treats the parallel rate as the relevant price of foreign exchange. Dornbusch et al. (1983) provides an early case study of the Brazilian black market; Greenwood and Kimbrough (1987) and Edwards (1988) model the determinants of the parallel rate under foreign exchange controls. Ghura and Grennes (1993), Sekkat and Varoudakis (2000), and Luintel (2000) link parallel premia to real exchange rate misalignment and to deviations from PPP. Cavallo and Rigobon (2016) use Argentine online prices to document persistent failures of the law of one price when the official rate is overvalued. Our contribution differs in focusing on dynamic CPI pass-through and on distinguishing the inflationary effects of parallel- and official-rate movements.

A third literature addresses the macroeconomic and fiscal consequences of parallel currency markets. Agénor (1992), Kiguel and O'Connell (1995), and Reinhart and Rogoff (2004) survey the historical prevalence of dual-rate regimes. Pinto (1991) and Morris

(1995) build models in which the parallel rate is the marginal price of foreign exchange and official devaluations have no direct contemporaneous effect on inflation. Schmitt-Grohé and Uribe (2024) solve for the optimal mix of exchange controls and inflation finance in a general-equilibrium economy with chronic fiscal deficits. Cerra (2016), on whose framework we build, develops a general-equilibrium model of inflation under foreign exchange rationing in Venezuela. Most directly related, Lariau et al. (2016) estimate pass-through in Angola and Nigeria using official and interbank rates and find no significant pass-through to Nigerian headline inflation; we revisit this finding by adding the parallel rate.

We also build on a methodological literature on identifying structural shocks in macroeconomic time series. We trace impulse responses using local projections following Jordà (2005). The sign-restricted SVAR adapts Uhlig (2005) and Rubio-Ramírez et al. (2010): structural shocks are recovered up to a set of rotations consistent with qualitative restrictions, and the restrictions in our case are read directly off the model in Section 3. The instrumental-variable local projection uses an external shifter to identify a single structural shock, as in Stock and Watson (2018) and Mertens and Ravn (2013). Our use of Nigeria-specific oil-supply disruptions is closest in spirit to Känzig (2021), who isolates global oil supply news from high-frequency surprises around OPEC announcements; we instead exploit physical disruption events documented in Rexer and Hvinden (2026).

Roadmap. The rest of the paper proceeds as follows. Section 2 provides background on parallel foreign exchange markets in Nigeria. Section 3 sets out the conceptual framework. Section 4 describes the data and the empirical strategy. Section 5 presents the main results. Section 6 concludes.

2 Parallel Markets for Foreign Exchange

Parallel markets arise when a managed exchange rate coexists with restrictions on foreign exchange access. Agénor (1992) traces their emergence in developing countries to trade-flow restrictions in low-income economies and to capital controls in middle-income ones. The parallel market fills the residual demand left unmet by the official window.

Such restrictions are common across the developing world. According to the IMF's *Annual Report on Exchange Arrangements and Exchange Restrictions* (IMF, 2022), 69% of countries operated some form of managed exchange arrangement in 2021. Of these, 11% maintain multiple currency practices and a further 12% apply foreign exchange restrictions that can give rise to parallel-market activity during periods of distress. All countries

in the latter two groups are emerging-market and developing economies. Given the unofficial nature of many parallel markets, these counts likely understate prevalence.

Where parallel markets exist, the premium over the official rate is often substantial. Table 1 reports official and parallel rates as of March 2023 for the 25 countries with documented multiple currency practices, compiled by the World Bank (Malpass, 2023). Premiums vary widely. Crisis and conflict economies such as Lebanon, Yemen, Syria, and Iran show premiums above 100%. A second group of managed-regime economies—Argentina, Ethiopia, Zimbabwe, Burundi, Nigeria, and Algeria—operates with persistent premiums of 40% to 90%. Nigeria, the focus of this paper, sits in this group at 62% on the reference date. A third group maintains parallel markets with premiums under 15%.

2.1 Nigeria's Foreign Exchange Market

Nigeria has operated a parallel foreign exchange market continuously since the 1986 Structural Adjustment Programme replaced a fixed-rate regime with an auction-based one. Through the 1990s and 2000s, the Central Bank of Nigeria (CBN) cycled through a series of managed-rate arrangements. A parallel market filled residual demand whenever official supply was rationed.

Our sample covers an active phase of this regime. The 2014–2016 oil price collapse prompted the CBN to impose administrative rationing on a managed official rate. In June 2015 it issued a list of imports (the "41-item list," later expanded to 43 items) that could not be funded through any official window. These restrictions effectively forced importers in excluded categories to source incremental foreign exchange through the parallel market. The key economic intuition is simple. Even when some firms receive foreign exchange at the official rate, the relevant price for retail pricing is often the cost of obtaining the next dollar rather than the average historical acquisition cost of existing inventories. Firms unable to secure sufficient official allocations must source incremental foreign exchange in the parallel market, and retail prices adjust accordingly. In this sense, the parallel rate functions as the marginal exchange rate for import pricing under rationing, while the official rate primarily determines the distribution of rents across firms with privileged access to the official window.

In June 2016 the CBN devalued the official rate from N197 to N305 per dollar and announced a transition to a flexible interbank market, but re-pegged shortly afterward. From the April 2017 launch of the Investors and Exporters (I&E) window onward, the CBN operated several official windows with distinct access criteria. In June 2023, the CBN collapsed its multiple windows into a single willing-buyer-willing-seller market,

Table 1: Countries with Multiple Currency Practices, March 2023

Country	Official	Parallel	Premium (%)
Iran (IFS rate)[a]	42,000	544,000	1,195.2
Lebanon	15,000	107,500	616.7
Yemen (Sana'a vs. Aden)	250	1,230	392.0
Syria	3,015	7,550	150.4
Iran (NIMA rate)[a]	285,000	544,000	90.9
Argentina	209	391	87.1
Ethiopia	54.4	100.2	84.1
Zimbabwe	930	1,600	72.1
Burundi[b]	2,061	3,359	63.0
Nigeria	461	745	61.7
Algeria	136	209	53.5
Malawi	1,028	1,495	45.4
Myanmar	2,100	2,857	36.0
Congo, Democratic Republic	2,036	2,323	14.1
Angola[c]	504	560	11.1
Bangladesh	106	113.3	6.9
Lao PDR[d]	16,221	17,327	6.8
Ghana	11.01	11.75	6.7
Libya	4.79	5.09	6.2
Mozambique	64.5	67.4	4.4
Ukraine	36.6	37.7	3.2
Sri Lanka	327	337	2.8
Sudan	590	605	2.6
Venezuela	24.5	24.7	0.9
South Sudan	851	850	-0.2

Notes: Official and parallel exchange rates in local currency per US dollar on March 31, 2023, unless otherwise noted. The premium is the percentage gap of the parallel rate over the official rate.
[a] Iran reports two official rates: the baseline IMF IFS rate and the NIMA rate.
[b] Burundi as of December 31, 2022.
[c] Angola as of January 27, 2023.
[d] Lao PDR as of February 28, 2023.
Sources: World Bank compilation (Malpass, 2023).

and the naira moved from N463 per dollar to N750 over the following weeks. Our estimation sample ends in May 2023, immediately before this break, to isolate the dual-rate regime. This institutional setting makes Nigeria a natural environment in which to distinguish between the administered official exchange rate and the marginal market-clearing exchange rate relevant for retail pricing.

3 Conceptual Framework

This section sets out the framework we use to interpret the empirical results. The framework follows Cerra (2016), drawing directly on her nontraded-good extension; we write out the two-good case in a form that lines up with the empirical specifications, and we make explicit the role of central-bank foreign exchange reserves. The mechanism is Cerra's: under foreign exchange rationing, current prices are pinned down by money and goods availability rather than by the official exchange rate, while the official rate operates indirectly through fiscal financing.

3.1 Environment

Consider a small open economy that exports oil and imports one consumption good, indexed by T. Oil export receipts X_t are exogenous and accrue to the government. The government sells these receipts to the central bank at the official exchange rate s_t (domestic currency per dollar) and makes transfers T_t to households in domestic currency. The central bank allocates foreign exchange through rationing. A fraction $(1-\phi)X_t$ is sold directly to importers at the official rate, while the remaining ϕX_t is allocated to arbitrageurs, who resell those dollars on the parallel market at rate b_t.

Central bank balance sheet. The central bank's assets are international reserves R_t (in domestic-currency terms) and net domestic credit to the government N_t; its liability is base money M_t:

$$R_t + N_t = M_t. \tag{1}$$

We allow for the possibility that the central bank does not fully distribute its dollar receipts each period. Let $\rho_t \in [-1, 1]$ denote the share of X_t that the central bank retains as reserves: $\rho_t > 0$ means accumulation, $\rho_t < 0$ means active depletion (selling reserves to defend the official rate beyond what current oil receipts allow), $\rho_t = 0$ is the Cerra (2016) baseline of full pass-through to the private sector. The reserves stock evolves as

$$R_{t+1} = R_t + s_t \rho_t X_t. \tag{2}$$

The volume of dollars actually sold to importers and arbitrageurs is $(1-\rho_t)X_t$ rather than X_t.

Money creation. Combining the consolidated government budget with the central-bank balance sheet (1) yields the money-supply equation

$$M_{t+1} = M_t + T_t - s_t(1 - \rho_t)X_t. \tag{3}$$

When $\rho_t = 0$, equation (3) reduces to the standard Cerra (2016) expression $M_{t+1} = M_t + T_t - s_t X_t$: money growth equals the fiscal deficit. When $\rho_t > 0$, the central bank absorbs fewer naira from the private sector, and money growth is higher than the fiscal deficit by $s_t \rho_t X_t$. When $\rho_t < 0$, the central bank sells reserves to absorb more naira, and money growth is lower. Equation (3) is equivalent to $\Delta M_{t+1} = T_t - s_t X_t + \Delta R_{t+1}$.

Equations (1)–(3) carry implications we use later for shock identification. A higher official rate raises the naira value of every dollar transaction with the central bank, narrowing the fiscal deficit and reducing money creation, even if reserves are unchanged. An adverse oil shock that reduces X_t raises money creation in the standard Cerra channel *and*, if the central bank defends the official rate by drawing down reserves ($\rho_t < 0$), raises it further. The two cases generate opposite-signed reserves responses, which we exploit in Section 4.

3.2 Production, preferences, and pricing

There are two consumption goods. The first, T, is the imported good purchased at world price P_t^* and sold domestically at price P_t^T. For the short horizons relevant to our pass-through exercise, we treat P_t^* as exogenous and approximately constant. The second, N, is locally produced using only inelastic labor:

$$Y_t^N = L_t, \tag{4}$$

with $L_t = 1$ for all t. A representative competitive producer chooses labor to maximize $P_t^N Y_t^N - w_t L_t$, implying the zero-profit condition

$$w_t = P_t^N. \tag{5}$$

The representative household has preferences

$$U(C_t^T, C_t^N) = \omega \log C_t^T + (1 - \omega) \log C_t^N, \qquad \omega \in (0, 1), \tag{6}$$

and faces the budget constraint

$$M_{t+1} + P_t^T C_t^T + P_t^N C_t^N = M_t + T_t + \Pi_t + w_t, \tag{7}$$

where Π_t denotes profits from import and arbitrage firms. Purchases require cash in advance,

$$M_t \geq P_t^T C_t^T + P_t^N C_t^N. \tag{8}$$

With positive expected inflation, the cash-in-advance constraint binds.

3.3 Imported and local goods under rationing

Let Q_t denote the quantity of the imported good available for sale in period t. Under a binding cash-in-advance constraint and Cobb-Douglas preferences, the household allocates constant expenditure shares across the two goods:

$$P_t^T C_t^T = \omega M_t, \qquad P_t^N C_t^N = (1-\omega) M_t. \tag{9}$$

Imposing goods-market clearing, $C_t^T = Q_t$ and $C_t^N = Y_t^N$, yields

$$P_t^T = \omega \frac{M_t}{Q_t}, \tag{10}$$

and

$$P_t^N = (1-\omega) \frac{M_t}{Y_t^N}. \tag{11}$$

Equations (10)–(11) are the key departure from a standard cost-push view of inflation. Current prices depend on money and on the quantities of goods available for sale; they do not depend directly on the official exchange rate. An official devaluation changes firms' rents but not the contemporaneous pricing equations.

Import availability is tied to the dollars actually sold by the central bank, which equals $(1-\rho_t)X_t$ when the central bank retains $\rho_t X_t$ as reserves. Hence

$$Q_t = \frac{(1-\rho_t)X_t}{P_t^*}. \tag{12}$$

Reserves accumulation tightens the import supply for given oil receipts; reserves drawdown relaxes it. This second channel of reserves on inflation, operating through quantities rather than through money creation, sharpens an oil-revenue shock's effect on prices: when X_t falls and the central bank simultaneously runs reserves down to defend the rate,

Q_t falls less than in proportion, but the money-creation channel (3) is amplified.

Import firms obtain dollars from two sources: the central bank, at rate s_t, and arbitrageurs in the parallel market, at rate b_t. With Q_t units imported at unit world price P_t^*, total dollar needs are $P_t^* Q_t$; the central bank supplies the inframarginal share $(1-\phi)$ at s_t and arbitrageurs supply the marginal share ϕ at b_t. The import firm's profit is

$$\Pi_t^{\text{imp}} = P_t^T Q_t - [(1-\phi)s_t + \phi b_t]\, P_t^* Q_t. \tag{13}$$

The bracketed term is the firm's average acquisition cost of foreign exchange. The parallel rate, however, is determined by the value of one additional dollar acquired at the margin. An extra dollar purchased on the parallel market allows the firm to import $1/P_t^*$ additional units of the good, which can be sold domestically at revenue P_t^T/P_t^*. As long as $b_t < P_t^T/P_t^*$, importers bid more for parallel-market dollars; competition among importers with arbitrageurs supplying the parallel market implies

$$b_t = \frac{P_t^T}{P_t^*}. \tag{14}$$

Combining (10) and (14) gives

$$b_t = \omega \frac{M_t}{P_t^* Q_t}. \tag{15}$$

The parallel rate is the scarcity price of imported consumption goods. It rises when money rises relative to import availability, when imported supply tightens (because X_t falls, ρ_t rises, or both), or both.

3.4 The role of the official exchange rate

The official exchange rate does not appear in (10), (11), or (15). Its role is instead fiscal. Figure 1 illustrates the imported-good market under rationing. Aggregate supply is vertical at the rationed quantity Q_t that import firms can finance with the dollars they obtain at s_t, and aggregate demand is the cash-in-advance schedule $P_t^T = \omega M_t/Q_t$ traced out as Q_t varies. Their intersection pins down the retail price P_t^T. The unit cost of supplying the good is $s_t P_t^*$, and the shaded rectangle between P_t^T and $s_t P_t^*$ is the rent earned by import firms on dollars purchased at the official rate.

A higher s_t raises the domestic-currency counterpart of the dollars sold by the central bank, reducing the need for monetary financing in (3). It also compresses the rents earned by importers that receive dollars at the official rate. Importers' and arbitrageurs'

Figure 1: The Market for the Imported Good Under Rationing

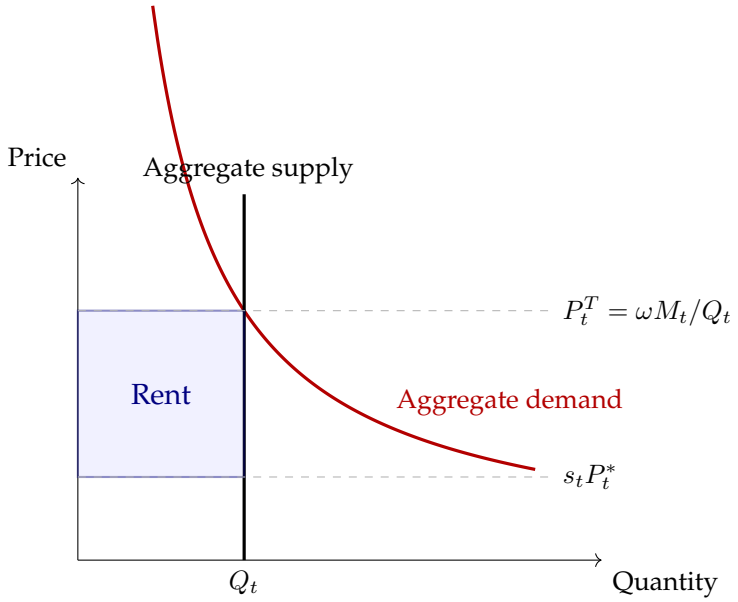

Notes: The market for the imported good in any period t. Aggregate supply is vertical at the quantity $Q_t = (1-\rho_t)X_t/P_t^*$ that import firms can finance with the dollars the central bank actually releases. Aggregate demand is the cash-in-advance schedule $P_t^T = \omega M_t/Q_t$ that household spending traces out as Q_t varies. The intersection pins down the retail price P_t^T. The unit cost of importing one unit at the official rate is $s_t P_t^*$; the shaded rectangle is the rent earned by importers on dollars they buy through the official window. A rise in M_t shifts demand outward and raises P_t^T proportionately; a tightening of Q_t (through lower X_t or higher ρ_t) rotates the supply schedule inward and also raises P_t^T. A devaluation of the official rate raises unit cost without affecting demand or supply: it leaves P_t^T unchanged on impact and shrinks the rent, leading to a fiscal correction that operates through future money creation, equation (3).

combined profits are

$$\Pi_t = P_t^T Q_t - s_t P_t^* Q_t. \tag{16}$$

Using (10), this becomes

$$\Pi_t = \omega M_t - s_t P_t^* Q_t. \tag{17}$$

Holding M_t and Q_t fixed, an official devaluation lowers rents one-for-one with the dollar amount sold through the official window.

The local good leaves Cerra's money-creation channel intact. Combining the household budget constraint (7), the binding cash-in-advance constraint (8), market clearing, and the wage identity (5), and substituting $Q_t = (1-\rho_t)X_t/P_t^*$, reduces to (3). The presence of a local good changes relative prices but does not alter the fiscal mechanism: official devaluations affect inflation through future money creation rather than through current pricing.

3.5 Aggregate prices and structural shocks

The Cobb-Douglas cost-of-living index is

$$P_t = \kappa (P_t^T)^{\omega} (P_t^N)^{1-\omega}, \tag{18}$$

for a constant $\kappa > 0$. Substituting (10)–(11) yields

$$P_t = \tilde{\kappa} \frac{M_t}{Q_t^{\omega} (Y_t^N)^{1-\omega}}, \tag{19}$$

where $\tilde{\kappa}$ is constant. If P_t^* and Y_t^N are constant over the horizon of interest, then (15) and (19) imply the decomposition

$$\Delta \log P_t = \omega \Delta \log b_t + (1-\omega) \Delta \log M_t. \tag{20}$$

The decomposition is useful as an interpretive identity, not as a basis for identification or a strict empirical restriction.

Implications for shock identification. The model implies a small set of structural shocks with distinct qualitative signatures across the variables $(b_t, s_t, M_t, P_t, R_t)$. Three shocks bear directly on the empirical exercise:

1. A *parallel-market scarcity shock* ε^{par}: an exogenous tightening of the dollars available in the parallel market (capital flight, demand pressure from the import-banned categories, sentiment shifts). The parallel rate b_t rises through (15); s_t, M_t, and R_t are unaffected at impact since these are administratively or institutionally determined; P_t rises through pass-through to the tradable price.
2. A *CBN policy devaluation* ε^{cbn}: an exogenous increase in s_t. From (3), money creation falls; from (15), the parallel rate is broadly unaffected at impact since M_t and Q_t are predetermined. The aggregate price P_t has no direct response from s_t in the same period and a negative cumulative response through the slower money channel.
3. An *oil-revenue shock* ε^{oil}: a fall in X_t. Reserves R_t fall as the central bank defends the official rate ($\rho_t < 0$). The dollars released to importers, $(1-\rho_t)X_t$, fall by less than X_t in proportion but still fall, tightening Q_t. From (3), fewer dollars sold to the private sector means money creation rises. Both rates depreciate: b_t through scarcity (15) and s_t as the central bank eventually relaxes its peg. The price level rises through both channels.

A fourth, non-FX, inflation shock ε^P enters (19) via Y_t^N or external prices and primarily affects P_t on impact, with the FX rates and reserves responding passively. Section 4 formalizes these signatures as sign restrictions in a structural VAR.

Two implications for the empirical work. The framework yields two implications that we take to the data. First, the parallel rate should contain more information about contemporaneous inflation than the official rate, because it is the market-clearing scarcity price of imported goods while the official rate affects current prices only indirectly through future money creation. Second, official-rate movements should be associated with subsequent changes in money creation through the fiscal channel: a higher official rate raises the domestic-currency value of dollar revenues, narrows the fiscal deficit, and reduces money creation. Both line up with the Nigerian data, as documented in Section 5.

4 Data and Empirical Strategy

4.1 Data

Central Bank of Nigeria. Three CBN-published series anchor the analysis. The official rate s_t is the CBN reference rate reported in the IMF *International Financial Statistics*, at which the government's oil-dollar revenues are converted to naira. This is the rate that the framework in Section 3 predicts should affect prices through the fiscal channel, because it determines the naira value of government FX inflows; it is not the rate at which most retail import transactions actually clear. The relevant transaction prices show up in the interbank rate (NAFEX/I&E) and, where rationing binds, in the parallel rate. We use the CBN reference rate as s_t for this reason. External reserves R_t are the CBN's liquid foreign-currency holdings in US dollars, drawn from the CBN's daily published series and aggregated to end-of-month values. Money supply M_t is M2 from the CBN's *Money and Credit Statistics*. We aggregate daily rate quotes to monthly averages throughout.

Parallel rate. The parallel rate b_t is from AbokiFX, a Nigerian website that publishes a daily parallel-market rate based on surveys of foreign exchange dealers. These data are widely used by market participants and in the applied literature on Nigeria. We supplement the series after its September 2021 suspension with an alternative daily dealer-survey series from a financial institution.

Consumer prices. Headline and food CPI from the World Bank, originally compiled by Nigeria's National Bureau of Statistics (NBS), together with the full panel of NBS divisional sub-indices.

Macroeconomic controls. Brent crude oil prices, the FAO food price index, and the U.S. Federal Funds rate from FRED.

Niger Delta militancy. Monthly counts of militant attacks against Nigerian oil infrastructure—pipelines, terminals, and offshore platforms—compiled by Rexer and Hvinden (2026) from ACLED event records. The series runs through December 2019.

The estimation sample spans November 2014 to May 2023, ending immediately before the unification of Nigeria's exchange rate windows in June 2023. Figure 2 plots the rate series. The CBN reference rate adjusts in discrete devaluations with long flat plateaus, the interbank rate moves continuously, and the parallel rate sits persistently above both. The premium widens during major depreciation episodes, exceeding 70% in 2016–2017 and again in 2022–2023, and narrows in the intervening period of relative stability.

Figure 2: Nigerian Official and Parallel Exchange Rates, November 2014 to May 2023

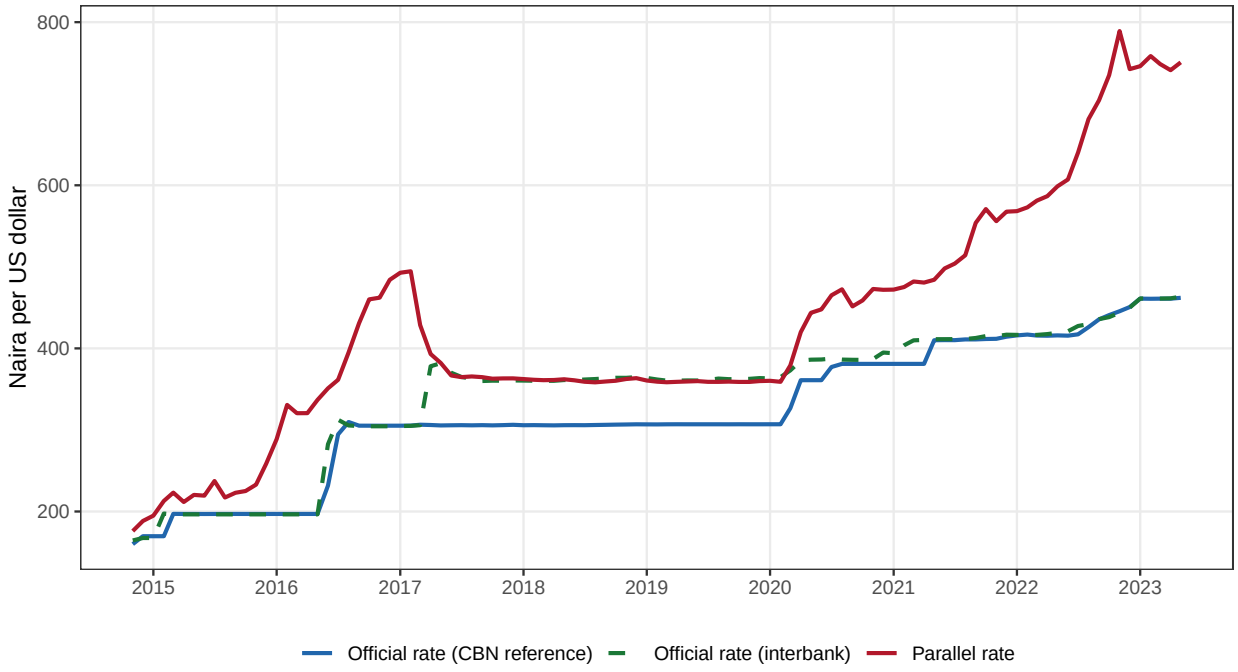

Notes: Monthly averages of daily naira-per-dollar quotes. The CBN reference rate is reported in the IMF *International Financial Statistics*. The interbank rate is the NAFEX/I&E rate. The parallel rate is from AbokiFX through September 2021 and from an alternative dealer-survey dataset thereafter.

4.2 From Model to Identification

The framework in Section 3 predicts that distinct structural shocks produce distinct qualitative responses across the parallel rate, the official rate, money, prices, and reserves. A

parallel-market shock raises the parallel rate and prices but leaves the administered official rate and reserves unmoved. A CBN devaluation shock raises the official rate and lowers money creation through the fiscal channel. An adverse oil-revenue shock raises both rates, expands money creation, drains reserves, and raises prices. A non-FX inflation shock raises prices alone. We use these patterns to identify the shocks in two ways.

A structural VAR with model-implied sign restrictions. We estimate a reduced-form VAR with two lags on $y_t = \left(\Delta \log b_t,\ \Delta \log s_t,\ \Delta \log M_t,\ \Delta \log P_t,\ \Delta \log R_t\right)'$ and impose the model-implied sign patterns on the structural impact matrix.[2] Table 2 lists the restrictions. Impact restrictions ($h = 0$) apply to the parallel rate, official rate, and CPI; the slower fiscal-channel restrictions on money and reserves apply to the cumulative response at $h = 6$ months. The unidentified fifth column of the structural matrix absorbs residual variation that does not match any of the four labeled shock patterns. Of 30,000 candidate rotations, 10,596 satisfy the restrictions (a 35.3% acceptance rate); the cumulative impulse responses below are summarized over these accepted draws.

Table 2: Identifying Sign Restrictions

	$\Delta \log b$ (parallel)	$\Delta \log s$ (official)	$\Delta \log M$ (M2)	$\Delta \log P$ (CPI)	$\Delta \log R$ (reserves)
ε^{par}: parallel-market shock	> 0	~ 0	—	> 0	—
ε^{cbn}: CBN devaluation	—	> 0	< 0 at h=6	—	—
ε^{oil}: adverse oil revenue	> 0	≥ 0	> 0 at h=6	> 0	< 0 at h=6
ε^{P}: non-FX inflation	—	—	—	> 0	—

Notes: Sign restrictions used to identify the four labeled structural shocks. "~ 0" for ε^{par} on $\Delta \log s$ is implemented as $|\Delta \log b| > |\Delta \log s|$ at impact. "—" means no restriction. Restrictions on $\Delta \log M$ and $\Delta \log R$ for ε^{cbn} and ε^{oil} are imposed on the cumulative response at horizon six months; all other restrictions are imposed at impact ($h = 0$). A fifth, unlabeled column of the structural impact matrix absorbs residual unidentified shocks.

An external-instrument local projection. As a complementary identification of the parallel-market channel, we estimate

$$\Delta \log P_{t+h} - \Delta \log P_{t-1} = \alpha_h + \beta_h \Delta \log b_t + \gamma_h \Delta \log s_t + \boldsymbol{\delta}_{\boldsymbol{h}}' X_t + \nu_{t+h}, \tag{21}$$

[2]The framework leaves the structural matrix identified up to an orthogonal rotation. Following Rubio-Ramírez et al. (2010), we draw random orthogonal rotations and retain those that satisfy the sign restrictions, reporting posterior medians and credible bands across accepted draws. Implementation details are in Appendix D.

instrumenting $\Delta \log b_t$ with the monthly count of Niger Delta militant attacks. The identifying assumption is that the attacks reduce Nigerian oil production, lower dollar inflows to the CBN, and tighten parallel-market rationing, while leaving global oil prices and Nigerian retail prices unaffected except through the foreign-exchange channel. We discuss residual threats below; the placebo test in Section 5 confirms that Nigeria-specific attacks do not move world oil prices. The CBN rate, Brent, FAO food prices, the U.S. Federal Funds rate, and M2 enter as exogenous controls. Newey-West standard errors are computed at lag four. The militancy data restrict the sample to $N = 61$ months ending in December 2019.

The two strategies target complementary objects. The SVAR identifies each structural shock up to a posterior over admissible rotations and traces the cumulative CPI response for each. The LP-IV identifies one of those shocks—a Nigeria-specific dollar-inflow shock acting through the parallel rate—using a single external shifter, and delivers a point estimate of the structural pass-through coefficient. Both recover the inflation response to underlying shocks that move the two rates, in a setting where the rates are jointly determined.

What the CBN devaluation shock identifies. CBN devaluations in Nigeria are not random; they are responses to reserves pressure, inflation expectations, IMF program timing, or political circumstances. What the sign restrictions identify is the residual component of official-rate variation orthogonal to the three other shocks: the parallel-market shock, which would push the parallel rate up by more than the official rate; the oil-revenue shock, which would deplete reserves; and the non-FX inflation shock, which would move CPI but not the rates. What remains for ε^{cbn} is the component of official-rate variation that is not explained by parallel-market pressure, oil revenue, or other inflation drivers. We are not claiming that the CBN devalues at random; we are claiming that the data identify a part of official-rate variation whose responses are consistent with the model's fiscal channel.

5 Results

5.1 Aggregate Pass-Through to Headline CPI

A central distinction in the paper is between movements in the official exchange rate and movements in the scarcity value of foreign exchange. In a dual-rate regime, these need not coincide. A policy-driven adjustment in the official rate may leave the marginal availabil-

ity of foreign exchange largely unchanged if rationing persists, while a parallel-market depreciation directly reflects tighter dollar scarcity facing importers and households. The empirical analysis therefore distinguishes between shocks that primarily move the parallel market and shocks that primarily move the administered official rate, even though the two exchange rates remain jointly endogenous in equilibrium.

The headline finding is an asymmetry between two structural shocks. A parallel-market shock raises cumulative CPI. A CBN devaluation shock does not; the central estimate is weakly contractionary, consistent with the model's fiscal channel. Figure 3 plots the two cumulative responses.

Figure 3: Cumulative CPI Response to a Parallel-Market vs. CBN Devaluation Shock

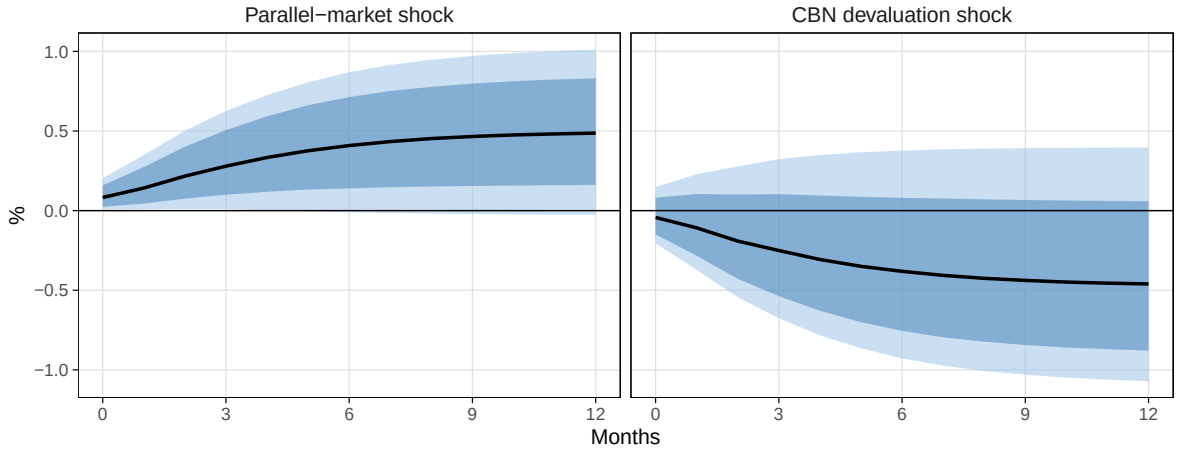

Notes: Posterior median (black line) and 68%/90% credible bands (shaded darker/lighter) of the cumulative log-CPI response to one-standard-deviation realizations of the parallel-market shock and the CBN devaluation shock. Estimated from the sign-restricted SVAR described in Section 4 on five Nigerian monthly variables, November 2014 to May 2023.

Table 3 reports the cumulative responses at impact and at three-month intervals through one year. The parallel-market response grows steadily. It is near zero on impact ($+0.001$), turns positive by $h = 3$ ($+0.003$), reaches $+0.004$ at $h = 6$, and stabilizes near $+0.005$ for horizons of nine to twelve months. At twelve months the median cumulative response is $+0.005$ with a 68% credible band of $[+0.002, +0.008]$ that lies entirely above zero. A one-standard-deviation parallel-market shock raises consumer prices by about half a percent over the year following the shock.

The official-rate response runs the other way. It is small at impact, drifts negative through the first quarter, and reaches its trough at intermediate horizons. The median cumulative response is -0.003 at $h = 3$, -0.004 at $h = 6$, and -0.005 at $h = 12$, with a 68% credible band of $[-0.009, +0.001]$ at twelve months. A majority of the posterior mass

Table 3: Cumulative CPI Response by Structural Shock, Sign-Restricted SVAR

Horizon	ε^{par} (parallel)	ε^{cbn} (CBN deval.)	ε^{oil} (oil revenue)	ε^{P} (non-FX inflation)
$h = 0$	0.001	-0.000	0.001	0.001
	[0.000, 0.002]	[-0.001, 0.001]	[0.000, 0.002]	[0.000, 0.002]
$h = 3$	0.003	-0.003	0.002	0.003
	[0.001, 0.005]	[-0.005, 0.001]	[0.001, 0.005]	[0.001, 0.005]
$h = 6$	0.004	-0.004	0.003	0.004
	[0.001, 0.007]	[-0.008, 0.001]	[0.001, 0.007]	[0.001, 0.007]
$h = 9$	0.005	-0.004	0.004	0.004
	[0.002, 0.008]	[-0.008, 0.001]	[0.001, 0.007]	[0.001, 0.008]
$h = 12$	0.005	-0.005	0.004	0.004
	[0.002, 0.008]	[-0.009, 0.001]	[0.001, 0.008]	[0.001, 0.008]

Notes: Posterior median and 68% credible bands of the cumulative log-CPI response to a one-standard-deviation realization of each structural shock, from the sign-restricted SVAR described in Section 4. Bracketed numbers are the 16th and 84th percentiles across 10596 accepted rotation draws (30000 proposals, 35.3% acceptance rate). The SVAR is estimated on monthly Nigerian data, November 2014 to May 2023.

lies below zero at every horizon past impact, and the 68% upper bound is at most $+0.001$. A CBN devaluation shock does not raise prices; the central estimate is weakly contractionary. The weakly negative central estimate is consistent with the model's fiscal channel: a higher official rate raises the naira value of dollar receipts sold by the central bank, narrows the fiscal deficit, and reduces subsequent money creation. Section 5.4 shows that the weakly negative response sharpens in months with a wide parallel premium, where the rent in the official window is largest.

The two remaining shocks raise prices monotonically. The oil-revenue shock and the residual non-FX inflation shock each deliver cumulative twelve-month CPI responses near $+0.004$ with bands that exclude zero. We return to the oil-revenue shock in Section 5.3.

Implied structural pass-through. Structural pass-through from the parallel rate is the ratio of the CPI to the parallel-rate response under the parallel-market shock. At twelve months the median ratio is $0.005/0.017 \approx 0.29$. The OLS-LP estimate of the same coefficient on the full sample is $\widehat{\beta}_{12} = 0.176$ (Appendix B). The structural identification recovers a pass-through estimate roughly 60% larger than the conditional correlation. We document a third, point-identified estimate in Section 5.5. Both structural estimates exceed the OLS conditional correlation, consistent with attenuation in OLS from measurement error in the dealer-survey parallel-rate quotes.

The estimate of 0.29 falls between the industrial-economy benchmarks compiled by Campa and Goldberg (2005); Burstein and Gopinath (2014) and Auer et al. (2021)—one- to two-year pass-through to tradable consumer prices in the range of 0.10 to 0.36, with an OECD average closer to 0.20—and the larger pass-through documented in high-inflation emerging-market devaluations (Burstein et al., 2005; Cravino and Levchenko, 2017). The natural reference set for Nigeria is the latter: dollar-invoicing-heavy economies with episodic large devaluations, where Gopinath et al. (2020) predicts above-OECD pass-through. The OLS conditional correlation of 0.18 understates the structural elasticity.

5.2 Heterogeneity Across Sub-Indices

The asymmetry holds category by category. We re-estimate the SVAR sub-index by sub-index, substituting each NBS sub-index for headline CPI in the five-variable system and leaving the sign restrictions unchanged.

Food CPI. Food is the natural place to start. It accounts for a large share of the Nigerian consumption basket, has the cleanest import-share data, and is the category in which the framework's tradable-versus-non-tradable distinction has the sharpest content. The cumulative twelve-month food-CPI response is $+0.006$ to the parallel-market shock, with 68% credible band $[+0.002, +0.010]$ lying entirely above zero, and -0.003 to the CBN devaluation, with band $[-0.009, +0.003]$ that places most posterior mass below zero. The headline asymmetry reappears at the food level, with magnitudes in the same range and a slightly noisier CBN response.

Divisional sub-indices. Figure 4 plots cumulative CPI responses for three NBS sub-indices: food, core (ex-farm), and transport. Appendix Figure A2 reports all fifteen.

The parallel-market shock raises consumer prices across all three categories, building from impact through $h = 6$ and stabilizing at cumulative twelve-month responses of $+0.6\%$ for food, $+0.4\%$ for core ex-farm, and $+0.4\%$ for transport, with 68% bands above zero in each case. The CBN devaluation shock lowers prices across the three categories, with central estimates between -0.3% and -0.5%, less precisely estimated than the parallel-rate response. The structural asymmetry that the headline SVAR identifies thus appears category by category. Appendix Figure A2 reports the full set of fifteen sub-indices: the parallel-market shock raises prices in every category, the CBN devaluation shock lowers them in every category, and the credible bands exclude zero for tradable-content sub-indices more often than for services and administered prices.

Figure 4: Cumulative CPI Response Across Sub-Indices, Sign-Restricted SVAR

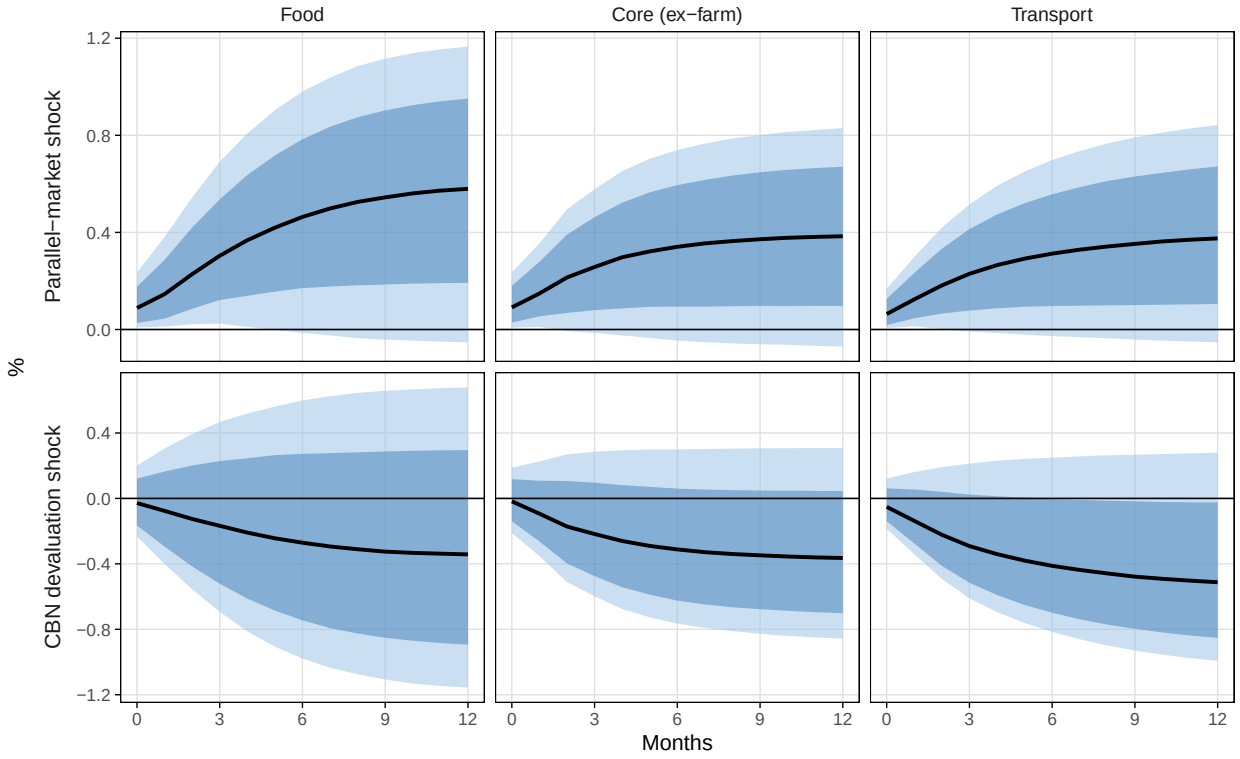

Notes: Cumulative log-CPI response (in percent) to a one-standard-deviation realization of the parallel-market shock (top row) and the CBN devaluation shock (bottom row), for three NBS sub-indices. Black lines are posterior medians; darker shading is the 68% credible band and lighter shading is the 90% band. Each column is a separate SVAR with the indicated sub-index substituted for headline CPI; identifying sign restrictions and other variables match the headline specification.

The OLS-LP joint specification yields the conditional-correlation analogue, reported in Appendix B.

5.3 The Fiscal Channel and Reserves

At first glance, a weakly negative inflation response to an official devaluation may appear counterintuitive. In the framework developed above, however, the sign of the inflation response depends on the channel through which the devaluation occurs. The model splits the transmission from the official rate to prices into two channels: a contemporaneous direct pass-through muted by foreign-exchange rationing, and a slower money-creation channel governed by the consolidated government budget. The CBN-devaluation-shock CPI response and the corresponding M2 response together trace the model's predicted footprint of the second channel.

Table 4 reports the twelve-month cumulative response of each variable to each struc-

Table 4: Cumulative Twelve-Month Response by Structural Shock, All Variables

Variable	ε^{par}	ε^{cbn}	ε^{oil}	ε^{P}
Parallel rate	0.017	-0.007	0.043	-0.019
	[-0.006, 0.038]	[-0.032, 0.019]	[0.027, 0.057]	[-0.042, 0.006]
Official rate	0.011	0.005	0.037	-0.004
	[-0.001, 0.021]	[-0.007, 0.018]	[0.030, 0.042]	[-0.019, 0.014]
M2	-0.001	-0.007	0.010	-0.003
	[-0.009, 0.009]	[-0.013, -0.002]	[0.004, 0.016]	[-0.011, 0.005]
CPI	0.005	-0.005	0.004	0.004
	[0.002, 0.008]	[-0.009, 0.001]	[0.001, 0.008]	[0.001, 0.008]
USD reserves	0.006	0.010	-0.036	-0.011
	[-0.024, 0.031]	[-0.019, 0.036]	[-0.054, -0.016]	[-0.038, 0.020]

Notes: Cumulative response at horizon $h = 12$ months of each variable to a one-standard-deviation realization of each structural shock. Bracketed numbers are 68% credible bands. Sample: November 2014 to May 2023. See Section 4 for the SVAR specification and identifying sign restrictions.

tural shock. Three patterns stand out. The parallel-market shock raises the parallel rate by 1.7% at twelve months but leaves the official rate, money, and reserves unmoved within sampling uncertainty. The CBN devaluation shock raises the official rate by 0.5%, lowers money by 0.7% over six months, and yields the weakly negative CPI response in Table 3. The oil-revenue shock moves both rates up together (4.3% on parallel, 3.7% on official), expands money by 1.0%, drains reserves by 3.6%, and raises prices through both channels.

Figure 5: Cumulative Reserves Response to Each Structural Shock

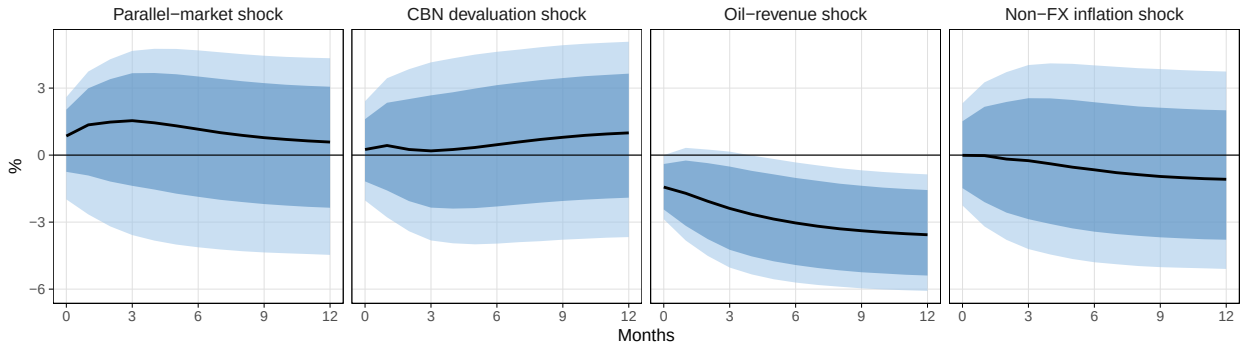

Notes: Cumulative log USD-reserves response to each of the four structural shocks. Only the oil-revenue shock produces a robust reserves drain.

The reserves response in Figure 5 separates the oil-revenue shock from the others. Only the oil-revenue shock produces a robust reserves drain: -0.04 at twelve months

with a 68% band of $[-0.05, -0.02]$ that lies well below zero. The reserves responses to the parallel-market shock, the CBN devaluation shock, and the non-FX inflation shock are centered near zero throughout. This asymmetry is what allows the identification scheme to separate the oil shock from a CBN devaluation shock, even though both raise the official rate.

Robustness. Appendix D reports three robustness exercises. The cumulative twelve-month CPI responses to ε^{par} and ε^{cbn} are stable across VAR lag lengths from one to three, across cumulative-restriction horizons from three to twelve months, and across alternative restriction sets, including the most aggressive variant that drops the M2 and reserves restrictions entirely. The headline asymmetry is not driven by specific identification choices. The estimation sample is short by macroeconometric standards (103 monthly observations for the SVAR and 61 for the LP-IV), so the credible bands and confidence sets reported throughout likely understate the full uncertainty; the asymmetry survives the pre-2020 subsample reported in the appendix.

5.4 State-Dependence in the Fiscal Channel

The state-dependent results provide one of the clearest links between the empirical evidence and the mechanism in the model. If the fiscal channel drives the negative CBN-devaluation response, the effect should be larger when the rent in the official window is larger: the fiscal correction triggered by a devaluation scales with the gap between the official rate and the scarcity price of foreign exchange. The log premium $\log(b_t/s_t)$ is a direct proxy for that gap. We augment the joint OLS specification with interactions between each rate and an indicator that the lagged log premium exceeded its sample median (a gap of 22%).[3]

Table 5 reports the results. The conditional correlation between the official rate and CPI is essentially zero in low-premium months and negative in high-premium months: in the low-premium regime, point estimates oscillate between -0.05 and -0.10 at intermediate horizons and are not statistically significant; in the high-premium regime, the point estimate drops to $\widehat{\gamma}_6^{\text{hi}} = -0.20$ at six months. The high-premium coefficients are unusually large in magnitude relative to a pass-through interpretation, reflecting that the high-premium subsample captures sharp episodes such as the June 2016 devaluation; we read them as directionally consistent with the model rather than as cleanly identified

[3]We do not estimate the state-dependent specification with the militancy IV. The LP-IV sample is short ($N = 61$), and splitting further by premium regime weakens the first stage below acceptable thresholds.

Table 5: State-Dependent Pass-Through to Headline CPI by Premium Regime

	Parallel rate			Official rate			
h	Low prem.	High prem.	Diff.	Low prem.	High prem.	Diff.	N
0	0.005	-0.019	-0.024	0.029	-0.011	-0.041*	102
	(0.010)	(0.012)	(0.016)	(0.023)	(0.010)	(0.022)	
3	0.091*	-0.036	-0.127**	0.034	-0.104***	-0.138	99
	(0.047)	(0.052)	(0.058)	(0.111)	(0.032)	(0.100)	
6	0.191*	0.061	-0.130	-0.060	-0.204***	-0.144	96
	(0.108)	(0.055)	(0.106)	(0.201)	(0.046)	(0.197)	
9	0.183	0.130**	-0.052	-0.048	-0.203***	-0.155	93
	(0.127)	(0.059)	(0.141)	(0.252)	(0.076)	(0.263)	
12	0.204	0.128*	-0.076	-0.292	-0.136*	0.156	90
	(0.130)	(0.072)	(0.149)	(0.286)	(0.071)	(0.302)	

Notes: Cumulative pass-through coefficients from a joint local projection of the cumulative log change in headline CPI on contemporaneous changes in the parallel and official rates, both interacted with an indicator that the lagged log parallel premium $\log(b_{t-1}/s_{t-1})$ was above its sample median (median premium $\approx 22\%$). "Low prem." reports the coefficient at $\mathbb{1}\{\text{high}\} = 0$; "High prem." is the sum of the level coefficient and the interaction; "Diff." is the interaction coefficient. Sample: monthly Nigerian data, November 2014 to May 2023. Newey-West standard errors with lag 4. *, **, *** denote significance at the 10%, 5%, and 1% levels.

elasticities. The data are more consistent with a contractionary fiscal channel in high-premium months than in low-premium months, which is the model's prediction: a wider premium implies a larger rent in the official window, and closing that gap through devaluation produces a larger reduction in the financing need and in subsequent money creation.

The parallel-rate coefficient is similar across regimes, as the framework predicts. Rationing binds throughout the pre-unification sample, so the parallel rate carries comparable pricing information when the premium is wide or narrow.

5.5 Bias-Corrected Pass-Through with an External Instrument

The SVAR identifies the parallel-market shock up to a set of admissible rotations. We point-identify the same channel using Niger Delta militant attacks on Nigerian oil infrastructure as an external instrument for the parallel rate, following the local-projection IV in equation (21).

The first-stage relationship is positive and economically meaningful. An additional militant attack per month is associated with a 0.9% depreciation of the parallel rate that month, significant at the 1% level, controlling for the CBN rate, Brent, FAO food prices,

and the U.S. Federal Funds rate.[4] The first-stage robust F-statistic on the excluded instrument is 7.7. This is below the Stock-Yogo 10%-bias threshold of 16.4 recently advocated by Lee et al. (2022), so standard 2SLS asymptotics are not fully reliable; we report a weak-instrument-robust Anderson-Rubin (1949) confidence set alongside the conventional Newey-West interval. A placebo regression of the contemporaneous Brent log change on the same instrument yields a small negative coefficient with $R^2 = 0.002$, indicating that the instrument captures a Nigeria-specific dollar-inflow shock rather than a global oil-price movement.

Figure 6: LP-IV vs. OLS-LP, Cumulative Pass-Through to Headline CPI

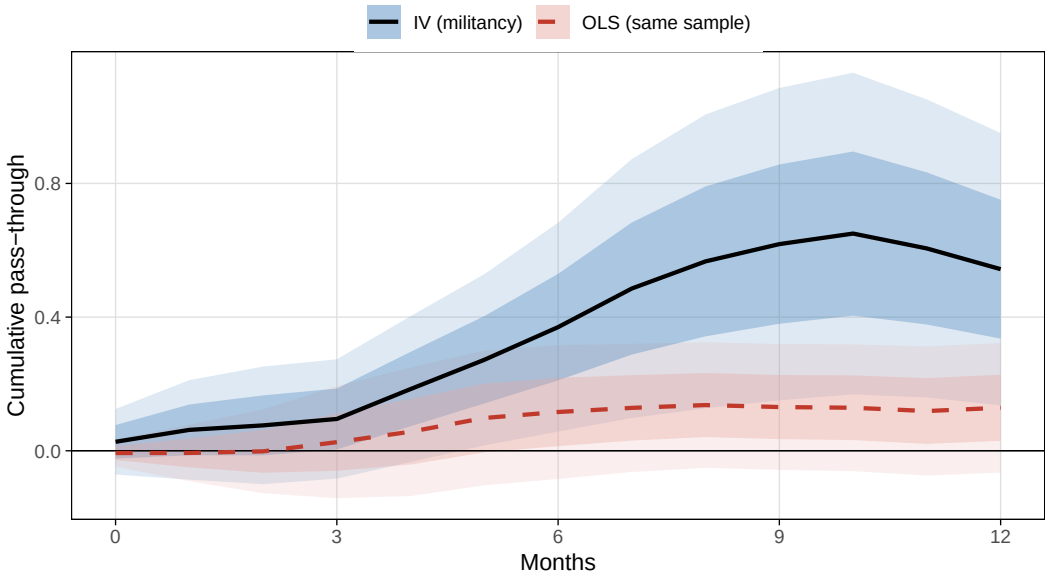

Notes: Cumulative pass-through to headline CPI from a one-percent depreciation of the parallel rate. The IV line instruments the parallel rate with the monthly count of Niger Delta militant attacks; the OLS line is the conditional correlation on the same subsample. Darker shading is the 68% confidence band; lighter shading is the 95% confidence band. Sample: November 2014 to December 2019.

Figure 6 and Table 6 report the cumulative IV pass-through estimates. The response builds steadily. It is near zero at impact ($\widehat{\beta}_0^{IV} = 0.03$), turns positive by six months ($\widehat{\beta}_6^{IV} = 0.37$), peaks near $\widehat{\beta}_9^{IV} = 0.62$, and settles at $\widehat{\beta}_{12}^{IV} = 0.54$. The Anderson-Rubin 90% confidence set, robust to weak identification, is $[0.19, 0.95]$ at six months, $[0.36, 1.51]$ at nine months, and $[0.30, 1.27]$ at twelve months: the set excludes zero from $h = 6$ onward. A one-percent depreciation of the parallel rate driven by a Nigeria-specific dollar-inflow shock is associated with cumulative CPI inflation in the range of 0.3% to 1.3% over

[4]We do not condition on M2 in the LP-IV. The model in Section 3 attributes part of the parallel-rate channel to money creation, so including M2 as a regressor would partial out a piece of the mechanism we are trying to estimate.

Table 6: Cumulative Pass-Through to Headline CPI: LP-IV vs. OLS-LP

	LP-IV (Niger Delta militant attacks)				OLS-LP, same sample		
Horizon	$\widehat{\beta}_h^{IV}$	SE	AR 90% CI	First-stage F	$\widehat{\beta}_h^{OLS}$	SE	N
$h = 0$	0.027	0.050	[-0.07, 0.13]	7.67	-0.007	61	
	(0.050)				(0.021)		
$h = 3$	0.095	0.091	[-0.05, 0.32]	7.67	0.026	61	
	(0.091)				(0.086)		
$h = 6$	0.370**	0.159	[0.19, 0.95]	7.67	0.116	61	
	(0.159)				(0.102)		
$h = 9$	0.618***	0.238	[0.36, 1.51]	7.67	0.131	61	
	(0.238)				(0.096)		
$h = 12$	0.543***	0.208	[0.30, 1.27]	7.67	0.129	61	
	(0.208)				(0.099)		

Notes: Local-projection coefficient on the cumulative log change in the parallel rate. The LP-IV columns instrument $\Delta \log b_t$ with the monthly count of militant attacks on Nigerian oil infrastructure compiled by Rexer and Hvinden (2026). Exogenous controls are the CBN reference rate, Brent crude prices, the FAO global food price index, and the U.S. Federal Funds rate; M2 is excluded from the controls because part of the parallel-rate channel operates through money creation. The Anderson-Rubin 90% confidence set is computed by grid inversion and is robust to weak instruments. The OLS columns report the conditional correlation on the same sub-sample. Newey-West standard errors with lag four in parentheses. Sample: November 2014 to December 2019 (limited by attack-series coverage). *, **, *** denote significance at the 10%, 5%, and 1% levels.

the following year. The corresponding OLS estimate on the same $N = 61$ subsample is $\widehat{\beta}_{12}^{OLS} = 0.13$. The Anderson-Rubin confidence set lies entirely above the OLS point estimate at every horizon past $h = 4$, indicating that the conditional correlation understates the instrument-identified pass-through.

The LP-IV and SVAR estimates of structural pass-through differ in magnitude (0.50 versus 0.29) but agree in sign and direction relative to OLS. The two designs identify related but distinct objects. The SVAR averages over admissible rotations and traces the response to a parallel-market shock with the other shocks held at zero. The LP-IV isolates the response to a Nigeria-specific dollar-inflow shock. They also cover different samples: the LP-IV ends in December 2019, while the SVAR runs through May 2023, and the post-2020 period broadens the parallel rate's distribution. Both exceed the OLS conditional correlation, consistent with measurement-error attenuation in the dealer-survey quotes.

Threats to exclusion. The identifying assumption is that militant attacks affect Nigerian retail prices only through the parallel rate. Four alternative channels are worth weighing. Pipeline attacks could disrupt domestic refined-fuel distribution and raise transport-CPI directly; Nigerian refineries were largely non-functional over the sample and Brent enters

as a control. Niger Delta unrest could disrupt agricultural supply chains running through Port Harcourt and Onne, but Nigerian food production is concentrated in the geographically separate north. Reverse causality from inflation to attack timing is possible, but the monthly attack pattern in our sample is dominated by discrete events—pipeline bombings and terminal blockades—whose timing is strategic rather than a smooth function of monthly inflation. A capital-flight risk-premium effect operates through the parallel rate and is part of the identified channel rather than a violation. None of these channels is fully shut down by the controls. We read the LP-IV estimates as informative about an upper-end magnitude rather than a clean LATE, and rely on the SVAR for the headline asymmetry.

5.6 Post-Unification Check

The framework predicts that once the official rate becomes the marginal price of foreign exchange, it should carry the pass-through that the parallel rate carried under the dual-rate regime. We examine this directly. We re-estimate the local projection in equation (A1) with only the official rate on the right-hand side, separately on the pre-unification sample (November 2014 – May 2023, $N = 103$) and the post-unification sample (June 2023 – December 2024, $N = 17$ at impact). We drop the parallel rate from the specification because the two rates are nearly collinear once the premium has collapsed. Appendix Table A6 reports the results. The pre-unification official-rate coefficient is small and drifts negative through intermediate horizons ($\widehat{\gamma}_3 = -0.012$, $\widehat{\gamma}_6 = -0.049$). The post-unification official-rate coefficient is positive throughout ($\widehat{\gamma}_3 = 0.023$, $\widehat{\gamma}_6 = 0.010$), and the post-unification three-month coefficient is close in magnitude to the pre-unification parallel-rate coefficient at the same horizon in the same specification ($\widehat{\beta}_3 = 0.021$). We have only nineteen months of post-unification parallel-rate and CPI data and cannot draw definitive conclusions, but this is suggestive evidence that the role of the marginal price of foreign exchange has transferred from the parallel rate to the official rate post-unification.

6 Implications and Conclusion

This paper documents an asymmetry in exchange-rate pass-through in a dual-rate economy: consumer prices respond to the parallel rate, the marginal price of foreign exchange, but not to the administered official rate. Using monthly data for Nigeria from 2014 to 2023, we identify the asymmetry through a sign-restricted SVAR whose restrictions are read off a Cerra-style model with foreign-exchange reserves, and we corroborate the

parallel-rate channel with a local projection that uses Niger Delta militant attacks as an external instrument and with a joint OLS local projection.

The headline finding is the asymmetry. A parallel-market shock raises cumulative CPI by about 0.5% over the following year, with the credible band excluding zero. A CBN devaluation shock does not raise prices; the central estimate is weakly contractionary. The weakly negative central estimate is consistent with the model's fiscal channel: a higher official rate raises the naira value of dollar receipts sold by the central bank, narrows the fiscal deficit, and reduces subsequent money creation. The weakly negative response sharpens in months with a wide parallel premium, where the rent in the official window is largest. An oil-revenue shock raises both rates, expands M2, drains reserves by 3.6%, and raises prices through both channels. A non-FX inflation shock raises prices alone.

The asymmetry between the two rates is the paper's main fact. The structural identification narrows the interpretation: it tells us that the parallel rate carries the pass-through and that the official rate, taken on its own, does not. The weakly negative central estimate to a CBN devaluation shock is consistent with the model's fiscal channel, but we read it as suggestive rather than as a sharply identified causal effect. The militancy IV-LP is a complement to the SVAR rather than a substitute, and is subject to exclusion-restriction concerns that the controls only partially address. The framework also abstracts from two channels that could shape the empirical responses: official-rate announcements may shift inflation expectations, and the parallel rate may clear speculative and hedging demand for dollars alongside residual demand for imports.

Two implications follow for central banks in dual-rate regimes. First, a monetary authority that tracks only the official rate understates the inflation pressure embedded in a depreciation. By the time an official adjustment is announced, much of the pass-through has already occurred through the parallel market. Second, an official-rate movement that is driven by CBN policy is not inflationary on impact; the central estimate is weakly contractionary at intermediate horizons. An official-rate movement that arrives through oil-revenue pressure is inflationary, because the same underlying shock that drives the central bank to devalue also expands money creation through reduced dollar sales. Treating all devaluations symmetrically misses both the underlying mechanism and the implications for inflation.

Two extensions are natural. The first is a sharper test of the rationing mechanism using direct data on central-bank allocations across the official window. The second is comparative analysis across other dual-rate economies—Argentina, Egypt, Iran, Lebanon, Venezuela—where daily parallel-market data would test whether the asymmetry is a general feature of foreign-exchange rationing or specific to Nigeria.

References

Agénor, P.-R. (1992). Parallel Currency Markets in Developing Countries: Theory, Evidence and Policy Implications. Essays in International Finance. 188.

Atkeson, A. and Burstein, A. (2008). Pricing-to-Market, Trade Costs, and International Relative Prices. *American Economic Review*, 98(5):1998–2031.

Auer, R., Burstein, A., and Lein, S. M. (2021). Exchange Rates and Prices: Evidence from the 2015 Swiss Franc Appreciation. *American Economic Review*, 111(2):652–686.

Boz, E., Gopinath, G., and Plagborg-Møller, M. (2019). Dollar Invoicing and the Heterogeneity of Exchange Rate Pass-Through. In *AEA Papers and Proceedings*, volume 109, pages 527–532. American Economic Association.

Burstein, A., Eichenbaum, M., and Rebelo, S. (2005). Large Devaluations and the Real Exchange Rate. *Journal of Political Economy*, 113(4):742–784.

Burstein, A. and Gopinath, G. (2014). International Prices and Exchange Rates. In *Handbook of International Economics*, volume 4, pages 391–451. Elsevier.

Campa, J. M. and Goldberg, L. S. (2005). Exchange rate pass-through into import prices. *Review of Economics and Statistics*, 87(4):679–690.

Campa, J. M. and Goldberg, L. S. (2010). The Sensitivity of the CPI to Exchange Rates: A Distribution Margin Story for Seven OECD Countries. *Review of Economics and Statistics*, 92(2):392–407.

Cavallo, A. and Rigobon, R. (2016). The Billion Prices Project: Using Online Prices for Measurement and Research. *Journal of Economic Perspectives*, 30(2):151–178.

Cerra, V. (2016). Inflation and the black market exchange rate in a repressed market: A model of venezuela. IMF Working Paper 16/159, International Monetary Fund.

Cravino, J. and Levchenko, A. A. (2017). The Distributional Consequences of Large Devaluations. *American Economic Review*, 107(11):3477–3509.

Dornbusch, R., Dantas, D. V., Pechman, C., de Rezende Rocha, R., and Simões, D. (1983). The Black Market for Dollars in Brazil. *The Quarterly Journal of Economics*, 98(1):25–40.

Edwards, S. (1988). Real and Monetary Determinants of Real Exchange Rate Behavior: Theory and Evidence from Developing Countries. *Journal of Development Economics*, 29(3):311–341.

Engel, C. (2006). Equivalence Results for Optimal Pass-Through, Optimal Indexing to Exchange Rates, and Optimal Choice of Currency for Export Pricing. *Journal of the European Economic Association*, 4(6):1249–1260.

Ghura, D. and Grennes, T. J. (1993). The Real Exchange Rate and Macroeconomic Performance in Sub-Saharan Africa. *Journal of Development Economics*, 42(1):155–174.

Goldberg, P. K. and Knetter, M. M. (1997). Goods Prices and Exchange Rates: What Have We Learned? *Journal of Economic Literature*, 35(3):1243–1272.

Gopinath, G., Boz, E., Casas, C., Díez, F. J., Gourinchas, P.-O., and Plagborg-Møller, M. (2020). Dominant Currency Paradigm. *American Economic Review*, 110(3):677–719.

Gopinath, G., Itskhoki, O., and Rigobon, R. (2010). Currency Choice and Exchange Rate Pass-Through. *American Economic Review*, 100(1):304–336.

Greenwood, J. and Kimbrough, K. P. (1987). Foreign Exchange Controls in a Black Market Economy. *Journal of Development Economics*, 26(1):129–143.

IMF (2022). Annual Report on Exchange Arrangements and Exchange Restrictions 2021. *Report*.

Itskhoki, O. and Mukhin, D. (2021). Mussa Puzzle Redux. *NBER Working Paper*.

Jordà, Ò. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182.

Känzig, D. R. (2021). The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements. *American Economic Review*, 111(4):1092–1125.

Kiguel, M. and O'Connell, S. A. (1995). Parallel exchange rates in developing countries. *The World Bank Research Observer*, 10(1):21–52.

Lariau, A., El-Said, M., and Takebe, M. (2016). An Assessment of the Exchange Rate Pass-Through in Angola and Nigeria. *IMF Working Paper*.

Lee, D. S., McCrary, J., Moreira, M. J., and Porter, J. (2022). Valid t-Ratio Inference for IV. *American Economic Review*, 112(10):3260–3290.

Luintel, K. B. (2000). Real Exchange Rate Behaviour: Evidence from Black Markets. *Journal of Applied Econometrics*, 15(2):161–185.

Malpass, D. (2023). The Parallel Exchange Rate Problem: The World Bank's Approach to Helping People in Developing Countries. World Bank Voices Blog.

Mertens, K. and Ravn, M. O. (2013). The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States. *American Economic Review*, 103(4):1212–1247.

Morris, S. (1995). Inflation Dynamics and the Parallel Market for Foreign Exchange. *Journal of Development Economics*, 46(2):295–316.

Mussa, M. (1986). Nominal Exchange Rate Regimes and the Behavior of Real Exchange Rates: Evidence and Implications. In *Carnegie-Rochester Conference Series on Public Policy*, volume 25, pages 117–214. Elsevier.

Pinto, B. (1991). Black Markets for Foreign Exchange, Real Exchange Rates and Inflation. *Journal of International Economics*, 30(1-2):121–135.

Reinhart, C. M. and Rogoff, K. S. (2004). The Modern History of Exchange Rate Arrangements: A Reinterpretation. *The Quarterly Journal of Economics*, 119(1):1–48.

Rexer, J. M. and Hvinden, E. S. (2026). Violence and Black Markets: Evidence from the Niger Delta Conflict. *Journal of Development Economics*, 182:103765.

Rubio-Ramírez, J. F., Waggoner, D. F., and Zha, T. (2010). Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference. *Review of Economic Studies*, 77(2):665–696.

Schmitt-Grohé, S. and Uribe, M. (2024). Exchange controls as a fiscal instrument. Columbia University Working Paper.

Sekkat, K. and Varoudakis, A. (2000). Exchange Rate Management and Manufactured Exports in Sub-Saharan Africa. *Journal of Development Economics*, 61(1):237–253.

Stock, J. H. and Watson, M. W. (2018). Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments. *Economic Journal*, 128(610):917–948.

Uhlig, H. (2005). What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure. *Journal of Monetary Economics*, 52(2):381–419.

Appendices

A Time-Series Diagnostics

This appendix reports time-series diagnostics that support OLS-LP and reduced-form VAR estimation on the pre-unification sample.

Unit roots and stationarity. Table A1 reports augmented Dickey-Fuller, Phillips-Perron, and KPSS tests on the levels and first differences of log CPI, the log parallel rate, and the log CBN rate. The three levels are non-stationary; the three first differences are stationary at the 1% or 5% level.

Table A1: Unit-Root and Stationarity Tests

Series	Form	ADF (p)	PP (p)	KPSS (p)	N
log CPI	level	0.990	0.990	0.010	103
log CPI	first difference	0.010	0.010	0.031	102
log Parallel rate	level	0.546	0.670	0.010	103
log Parallel rate	first difference	0.014	0.010	0.100	102
log Official rate	level	0.562	0.465	0.010	103
log Official rate	first difference	0.010	0.010	0.100	102

Notes: Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) tests of the null of a unit root; Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test of the null of level stationarity. Sample: November 2014 to May 2023. The KPSS p-values for the first-difference series are reported as the smallest p for which the test statistic falls in the right tail; a value of 0.10 in the table means the test fails to reject stationarity at conventional levels.

Granger causality. Table A2 reports pairwise Granger-causality tests at six monthly lags. The parallel rate Granger-causes consumer prices ($F = 2.45$, significant at the 5% level), while the official rate does not—consistent with the structural asymmetry recovered in the main text.

Multicollinearity. Variance inflation factors on the joint specification regressors at $h = 12$ range from 1.11 for the parallel rate to 1.67 for Brent crude, well below the conventional threshold of 5.

Table A2: Granger Causality Tests

Direction	Lags	F	p-value	N
Panel A: Main spec (s = CBN reference rate)				
$\Delta \log b \to \Delta \log s$	6	3.03	0.010	103
$\Delta \log s \to \Delta \log b$	6	3.37	0.005	103
$\Delta \log b \to \Delta \log P$	6	2.45	0.031	103
$\Delta \log P \to \Delta \log b$	6	1.63	0.149	103
$\Delta \log s \to \Delta \log P$	6	1.11	0.365	103
$\Delta \log P \to \Delta \log s$	6	4.96	< 0.001	103
Panel B: Alternative spec (s^{NAFEX} = interbank rate)				
$\Delta \log b \to \Delta \log s^{\text{NAFEX}}$	6	3.60	0.003	103
$\Delta \log s^{\text{NAFEX}} \to \Delta \log b$	6	1.15	0.343	103
$\Delta \log s^{\text{NAFEX}} \to \Delta \log P$	6	1.96	0.080	103
$\Delta \log P \to \Delta \log s^{\text{NAFEX}}$	6	7.33	< 0.001	103

Notes: Each row reports a Granger causality test of the null that lags of the variable on the left do not predict the variable on the right, conditional on the right-hand variable's own lags. Variables are first differences of log CPI, the log parallel rate, and the indicated official rate. The conditional rows in Panel A also include lags of the third variable on the right-hand side. The lag length of six matches the intermediate horizon at which the main asymmetry in Section B becomes statistically significant. Sample: November 2014 to May 2023.

B OLS-LP Joint Specification

This appendix reports the OLS-LP joint specification that the structural results in Section 5 reinterpret. The OLS-LP estimates a horizon-by-horizon regression of cumulative log CPI on contemporaneous changes in the parallel and official rates, plus macro controls:

$$\Delta \log P_{t+h} - \Delta \log P_{t-1} = \alpha_h + \beta_h \Delta \log b_t + \gamma_h \Delta \log s_t + \boldsymbol{\delta}_{\boldsymbol{h}}' X_t + \varepsilon_{t+h}. \tag{A1}$$

Newey-West standard errors are computed at lag four. The coefficients β_h and γ_h are conditional correlations, not structural objects: they measure the partial association between each rate and CPI holding the other rate at its sample value. The SVAR and LP-IV results in the main text indicate that these conditional correlations understate the structural pass-through.

Table A3 reports the cumulative pass-through coefficients at impact and at three-month intervals through one year. The parallel-rate estimate at twelve months is $\widehat{\beta}_{12} = 0.176$, significant at the 5% level. The official-rate estimate $\widehat{\gamma}_{12} = -0.011$ is indistinguish-

Table A3: Cumulative Pass-Through to Headline CPI, by Horizon

Horizon	Parallel rate (β_h)	Official rate (γ_h)	N
$h = 0$	-0.005	0.011*	102
	(0.013)	(0.007)	
$h = 3$	0.032	-0.018	99
	(0.056)	(0.028)	
$h = 6$	0.147**	-0.087**	96
	(0.073)	(0.035)	
$h = 9$	0.195***	-0.068	93
	(0.071)	(0.049)	
$h = 12$	0.176**	-0.011	90
	(0.076)	(0.048)	

Notes: Cumulative pass-through coefficients from the local-projection specification in Equation (A1), estimated on Nigerian data from November 2014 to May 2023. Each row reports the joint regression of the log change in the consumer price index between $t-1$ and $t+h$ on contemporaneous log changes in the parallel and official naira-dollar rates. Newey-West standard errors with a lag of 4 in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels.

able from zero, with a negative dip of $\widehat{\gamma}_6 = -0.087$ at six months. The qualitative pattern aligns with the structural results in the main text.

Food CPI. Table A4 reports the OLS-LP coefficients with the food sub-index as the outcome. The parallel-rate response builds through $\widehat{\beta}_6 = 0.09$ and $\widehat{\beta}_7 = 0.11$, peaks at $\widehat{\beta}_{10} = 0.15$, and settles at $\widehat{\beta}_{12} = 0.114$. A one-percent depreciation of the parallel rate is associated with cumulative food inflation of about 0.11% over the following year, about two-thirds of the headline estimate. The official-rate coefficient drifts negative through the first half-year, reaches $\widehat{\gamma}_6 = -0.085$ (significant at the 5% level), and fades to zero by twelve months.

Divisional sub-indices. Table A5 reports the joint OLS-LP estimates for a representative subset of NBS sub-indices. Table A12 reports the full sixteen-sub-index set. The parallel-rate coefficient is positive in twelve of the fifteen non-headline sub-indices. The two largest are administered or rebased-component categories: housing and utilities ($\widehat{\beta}_{12} = 0.43$) and education ($0.29$). The three sub-indices with negative point estimates are predominantly services or administered prices. The official-rate intermediate-horizon dip is more pronounced in the cross-section than at the headline level; at $h = 6$, ten of fifteen non-headline sub-indices carry official-rate coefficients significant at the 5% level or

Table A4: Cumulative Pass-Through to Food CPI, by Horizon

Horizon	Parallel rate (β_h)	Official rate (γ_h)	N
$h = 0$	-0.023*	0.006	102
	(0.014)	(0.008)	
$h = 3$	-0.028	-0.024	99
	(0.056)	(0.035)	
$h = 6$	0.090	-0.085**	96
	(0.074)	(0.043)	
$h = 9$	0.140	-0.032	93
	(0.086)	(0.061)	
$h = 12$	0.114	0.038	90
	(0.083)	(0.067)	

Notes: Cumulative pass-through coefficients from the joint local-projection specification in Equation (A1), with the food sub-index of the consumer price index as the outcome variable. Sample: Nigerian data from November 2014 to May 2023. Newey-West standard errors with a lag of 4 in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels.

better.

Table A5: Cumulative Twelve-Month Pass-Through Across Selected CPI Sub-Indices

Sub-index	Parallel rate ($\widehat{\beta}_{12}$)	Official rate ($\widehat{\gamma}_{12}$)	N
Imported food	0.039	-0.123***	81
	(0.083)	(0.041)	
Core (ex-farm)	0.167*	-0.011	81
	(0.093)	(0.051)	
Clothing	0.170*	0.072	81
	(0.096)	(0.058)	
Transport	0.139	-0.016	81
	(0.085)	(0.049)	
Communication	-0.131*	-0.153***	81
	(0.076)	(0.044)	
Restaurants	-0.042	-0.034	81
	(0.031)	(0.037)	

Notes: Cumulative twelve-month pass-through to each CPI sub-index from the joint local-projection specification in Equation (A1). Sample: monthly Nigerian data from November 2014 to August 2022 ($N = 81$ at $h = 12$). Newey-West standard errors with a lag of 4 in parentheses. The full table covering all sub-indices and intermediate horizons is in Appendix Table A12. *, **, *** denote significance at the 10%, 5%, and 1% levels.

C Post-Unification Pass-Through

This appendix reports the single-rate variant of the OLS local projection in equation (A1), estimated separately on the pre-unification sample (November 2014 – May 2023) and the post-unification sample (June 2023 – December 2024). The specification keeps only one of the two rates on the right-hand side; otherwise it matches the joint OLS-LP. The exercise supports the post-unification check in Section 5.6.

Table A6: Single-Rate OLS-LP, Pre- and Post-Unification

	$h = 0$	$h = 3$	$h = 6$	N
Panel A: Pre-unification (November 2014 – May 2023)				
Official rate	0.011	-0.012	-0.049	103
	(0.012)	(0.046)	(0.074)	
Parallel rate	-0.003	0.021	0.094	103
	(0.010)	(0.038)	(0.062)	
Panel B: Post-unification (June 2023 – December 2024)				
Official rate	0.010*	0.023	0.010	17
	(0.006)	(0.017)	(0.010)	
Parallel rate	0.005	0.005	0.017	17
	(0.010)	(0.029)	(0.015)	

Notes: Cumulative log-CPI response from a single-rate version of equation (A1), with only the indicated rate on the right-hand side together with the macro controls (Brent, FAO food prices, U.S. Federal Funds rate). Each row reports a separate regression. Standard errors in parentheses are OLS standard errors; we report OLS rather than HAC standard errors because the post-unification sample is too short for reliable HAC adjustment. *, **, *** denote significance at the 10%, 5%, and 1% levels. N at $h = 0$; longer horizons drop observations at the end of the post-unification sample.

D SVAR Robustness

This appendix reports three robustness exercises for the sign-restricted SVAR: varying the reduced-form VAR lag length, varying the horizon at which the cumulative M2 and reserves restrictions are imposed, and varying the restriction set itself. Each table reports the cumulative twelve-month log-CPI response to each structural shock under the alternative specification.

Lag length. Table A7 reports the SVAR with VAR(1) and VAR(3) reduced forms alongside the VAR(2) main specification. The qualitative asymmetry between ε^{par} and ε^{cbn} survives in every case. VAR(3) yields the sharpest result, with the 68% credible band for ε^{cbn} lying entirely below zero.

Table A7: SVAR Robustness: VAR Lag Length

Specification	ε^{par}	ε^{cbn}	ε^{oil}	ε^{P}	Accept rate	N_{acc}
VAR(1)	0.004	-0.003	0.003	0.005	39.8%	11941
	[0.001, 0.008]	[-0.008, 0.002]	[0.000, 0.007]	[0.001, 0.008]		
VAR(2) (main)	0.005	-0.005	0.004	0.004	35.3%	10596
	[0.002, 0.008]	[-0.009, 0.001]	[0.001, 0.008]	[0.001, 0.008]		
VAR(3)	0.005	-0.006	0.005	0.003	35.9%	10773
	[0.000, 0.008]	[-0.010, -0.002]	[0.001, 0.008]	[-0.001, 0.007]		

Notes: Posterior median (top row) and 68% credible band (bottom row) of the cumulative twelve-month log-CPI response to each structural shock, under three lag specifications for the reduced-form VAR. All other identification choices match the main spec. Acceptance rate is the share of 30,000 rotation draws that satisfy the sign restrictions.

Restriction horizon. The main specification imposes the M2 fiscal-channel restriction (for ε^{cbn}) and the M2 and reserves restrictions (for ε^{oil}) on cumulative responses at $h = 6$ months. The choice of horizon is reasonable but not pinned down by the model. Table A8 reports the SVAR with the cumulative restrictions imposed at $h \in \{3, 6, 9, 12\}$. The cumulative twelve-month CPI responses are essentially unchanged across all four horizons; the acceptance rate rises from 31% at $h = 3$ to 38% at $h = 12$ because longer-horizon cumulative responses are less restrictive in finite samples.

Restriction set. The most aggressive test drops the M2 and reserves restrictions and identifies the four shocks only from the impact responses of the parallel rate, the official rate, and CPI. To distinguish ε^{cbn} from ε^{oil} without the M2 discriminator, we add the dominance restriction $|\Delta \log s| > |\Delta \log b|$ at impact for ε^{cbn}, the policy analogue of the parallel-dominance restriction on ε^{par}. Table A9 also reports an intermediate specification that imposes the M2 and reserves restrictions at impact only (rather than cumulatively at $h = 6$). The headline twelve-month CPI responses are within sampling uncertainty of the main specification under both alternatives. When the M2 and reserves restrictions are dropped, the M2 and reserves IRFs are centered near zero with wide bands across all four shocks—as expected, since the disciplining restrictions have been removed—but the CPI asymmetry survives.

Table A8: SVAR Robustness: Cumulative Restriction Horizon

Specification	ε^{par}	ε^{cbn}	ε^{oil}	ε^{P}	Accept rate	N_{acc}
$h_{\text{restr}} = 3$	0.005	-0.005	0.004	0.004	31.6%	9484
	[0.002, 0.008]	[-0.009, 0.000]	[0.001, 0.008]	[0.001, 0.008]		
$h_{\text{restr}} = 6$ (main)	0.005	-0.005	0.004	0.004	35.3%	10596
	[0.002, 0.008]	[-0.009, 0.001]	[0.001, 0.008]	[0.001, 0.008]		
$h_{\text{restr}} = 9$	0.005	-0.005	0.004	0.004	37.1%	11145
	[0.002, 0.008]	[-0.009, 0.001]	[0.001, 0.008]	[0.001, 0.008]		
$h_{\text{restr}} = 12$	0.005	-0.005	0.004	0.004	38.1%	11436
	[0.002, 0.008]	[-0.009, 0.001]	[0.001, 0.008]	[0.001, 0.008]		

Notes: Posterior median (top row) and 68% credible band (bottom row) of the cumulative twelve-month log-CPI response to each structural shock, under four choices of the horizon h_{restr} at which the cumulative M2 and reserves restrictions are imposed. All other identification choices match the main spec. The reduced-form VAR has two lags.

Table A9: SVAR Robustness: Restriction Set

Specification	ε^{par}	ε^{cbn}	ε^{oil}	ε^{P}	Accept rate	N_{ac}
Main spec ($h_{\text{restr}} = 6$)	0.005	-0.005	0.004	0.004	35.3%	1059
	[0.002, 0.008]	[-0.009, 0.001]	[0.001, 0.008]	[0.001, 0.008]		
Impact-only M2 & reserves	0.005	-0.005	0.004	0.004	31.0%	928
	[0.002, 0.008]	[-0.009, -0.000]	[0.001, 0.008]	[0.001, 0.008]		
Drop M2 & reserves entirely	0.006	-0.005	0.003	0.004	66.4%	1991
	[0.002, 0.009]	[-0.009, 0.000]	[0.000, 0.007]	[0.000, 0.008]		

Notes: Posterior median (top row) and 68% credible band (bottom row) of the cumulative twelve-month log-CPI respons each structural shock under three identification schemes. The main spec imposes cumulative M2 and reserves restriction $h = 6$. The impact-only spec imposes the same sign patterns at $h = 0$ rather than cumulative. The drop-M2-and-reser spec drops the M2 and reserves restrictions entirely and adds an impact dominance restriction $|\Delta \log s| > |\Delta \log b|$ for ε^{cb} distinguish it from ε^{oil}.

Pre-2020 sub-sample. Re-estimating the SVAR on the November 2014 – December 2019 window (61 monthly observations) yields a cumulative twelve-month CPI response of $+0.005$ to ε^{par} and -0.003 to ε^{cbn}, close to the full-sample estimates. The asymmetry is not a post-2020 artifact.

Are ε^{cbn} and ε^{oil} the same shock? A reasonable concern is that major CBN devaluations in Nigeria were triggered by oil-revenue pressure, so the SVAR may not actually separate the two shocks. We address this with a historical decomposition: at each accepted rotation we extract the implied structural-shock time series $\varepsilon_t = (B_0' B_0)^{-1} B_0' u_t$ and summarize across rotations. The cross-rotation median correlation between ε_t^{cbn} and ε_t^{oil} is 0.03—essentially zero. Three of the ten largest-magnitude months are shared between the two

shocks (Table A10); the remaining seven are distinct. The June 2016 devaluation and the December 2015 reserves crisis register as both kinds of shock, as one would expect for episodes in which oil pressure and policy response coincided. But the SVAR pulls them apart over most of the sample.

Table A10: Largest-Magnitude Months for ε^{cbn} and ε^{oil}

	ε^{cbn} (CBN devaluation)		ε^{oil} (oil revenue)	
Rank	Date	Shock	Date	Shock
1	2015-03	3.08	2015-12	3.07
2	2016-07	2.90	2016-06	3.00
3	2016-06	2.69	2020-04	2.53
4	2015-12	-1.85	2020-05	-2.08
5	2015-05	1.46	2015-02	1.80
6	2021-05	1.45	2020-03	1.67
7	2015-04	-1.44	2018-11	-1.59
8	2017-02	1.41	2015-04	-1.54
9	2015-07	1.26	2022-12	-1.44
10	2017-12	-1.20	2016-11	-1.43

Notes: Months ranked by $|\varepsilon_t^{cbn}|$ and $|\varepsilon_t^{oil}|$ at the posterior-median rotation. Shocks are extracted from the sign-restricted SVAR by computing $\varepsilon_t = (B_0'B_0)^{-1}B_0'u_t$ for each accepted rotation and taking the cross-rotation median in each month. Of the ten largest-magnitude months in each column, 3 are shared between the two lists; the remainder are distinct. The pairwise correlation between ε_t^{cbn} and ε_t^{oil} at the median rotation is 0.03.

Identification uncertainty. Posterior credible bands reflect only rotation uncertainty given the reduced-form point estimates. A fully Bayesian implementation would also sample over the reduced-form parameters; this would widen the bands modestly but not change the qualitative pattern.

Full IRF grid. Figure A1 plots cumulative impulse responses of all five variables to each of the four structural shocks.

E Full Sub-Index Tables

Figure A2 reports cumulative CPI responses across all fifteen sub-indices to the parallel-market and CBN devaluation shocks, complementing the seven-sub-index version in Sec-

Figure A1: Cumulative IRFs of All Variables to Each Structural Shock

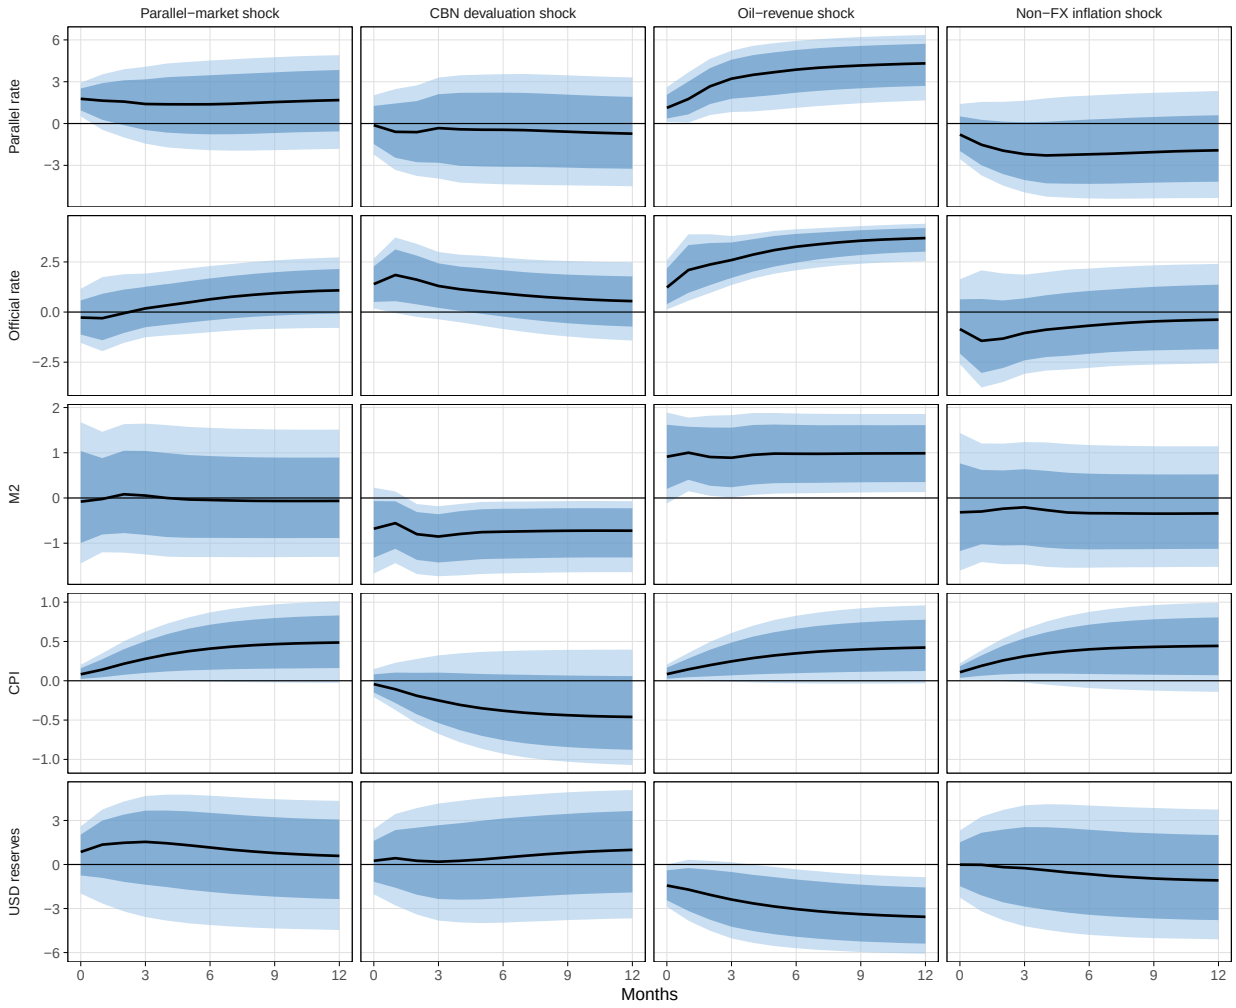

Notes: Cumulative impulse responses of each of the five variables to one-standard-deviation realizations of each of the four labeled shocks. Solid lines are posterior medians; shaded bands are 68% / 90% credible intervals.

tion 5.2. Table A11 reports the same SVAR estimates in numerical form. Table A12 reports the corresponding OLS-LP joint specification.

Table A11: Sign-Restricted SVAR by CPI Sub-Index, Full Set

	$h = 6$		$h = 12$		
Sub-index	$\widehat{\mathsf{IRF}}^{par}$	$\widehat{\mathsf{IRF}}^{cbn}$	$\widehat{\mathsf{IRF}}^{par}$	$\widehat{\mathsf{IRF}}^{cbn}$	N_{accepted}
Food	**0.005**	-0.003	**0.006**	-0.003	3430
	[0.002, 0.008]	[-0.007, 0.003]	[0.002, 0.010]	[-0.009, 0.003]	
Imported food	**0.003**	**-0.003**	**0.003**	**-0.003**	4232
	[0.001, 0.005]	[-0.005, -0.000]	[0.001, 0.005]	[-0.006, -0.001]	
Core (ex-farm)	**0.003**	-0.003	**0.004**	-0.004	4193
	[0.001, 0.006]	[-0.006, 0.001]	[0.001, 0.007]	[-0.007, 0.000]	
Core (ex-farm, ex-energy)	**0.003**	-0.002	**0.003**	-0.003	3630
	[0.001, 0.005]	[-0.005, 0.001]	[0.001, 0.005]	[-0.006, 0.001]	
Alcohol & tobacco	**0.003**	**-0.004**	**0.004**	**-0.006**	4103
	[0.001, 0.005]	[-0.007, -0.000]	[0.001, 0.007]	[-0.009, -0.001]	
Clothing	**0.003**	-0.002	**0.004**	-0.003	3523
	[0.001, 0.005]	[-0.006, 0.002]	[0.002, 0.007]	[-0.008, 0.002]	
Housing & utilities	**0.006**	-0.005	**0.007**	-0.006	5159
	[0.002, 0.011]	[-0.010, 0.001]	[0.002, 0.011]	[-0.011, 0.001]	
Furnishings	**0.003**	-0.003	**0.003**	-0.004	3121
	[0.001, 0.005]	[-0.006, 0.000]	[0.001, 0.005]	[-0.006, 0.000]	
Health	**0.002**	-0.003	**0.003**	-0.004	4440
	[0.001, 0.004]	[-0.006, 0.000]	[0.001, 0.005]	[-0.007, 0.000]	
Transport	**0.003**	**-0.004**	**0.004**	**-0.005**	3214
	[0.001, 0.006]	[-0.007, -0.000]	[0.001, 0.007]	[-0.009, -0.000]	
Communication	**0.001**	**-0.002**	**0.002**	**-0.004**	4842
	[0.000, 0.003]	[-0.003, -0.001]	[0.001, 0.004]	[-0.006, -0.001]	
Recreation	**0.003**	-0.003	**0.003**	-0.003	3281
	[0.001, 0.005]	[-0.006, 0.001]	[0.001, 0.006]	[-0.006, 0.001]	
Education	**0.003**	-0.001	**0.004**	-0.002	5011
	[0.001, 0.006]	[-0.005, 0.002]	[0.001, 0.007]	[-0.006, 0.002]	
Restaurants	0.001	-0.002	0.001	-0.002	4744
	[-0.000, 0.003]	[-0.004, 0.001]	[-0.000, 0.003]	[-0.004, 0.001]	
Miscellaneous	**0.002**	-0.003	**0.002**	-0.003	3429
	[0.001, 0.004]	[-0.005, 0.001]	[0.001, 0.004]	[-0.005, 0.001]	

Notes: Cumulative log-CPI response to one-standard-deviation realizations of the parallel-market shock (ε^{par}) and the CBN devaluation shock (ε^{cbn}) at horizons six and twelve months, for each NBS CPI sub-index. Bracketed numbers are 68% credible bands across accepted rotation draws. Median estimates whose 68% band excludes zero appear in bold. Each row is a separate SVAR with the indicated sub-index substituted for headline CPI; identifying sign restrictions and other variables match the headline specification (Section 4).

Figure A2: Full Set of Cumulative CPI Responses Across Sub-Indices, Sign-Restricted SVAR

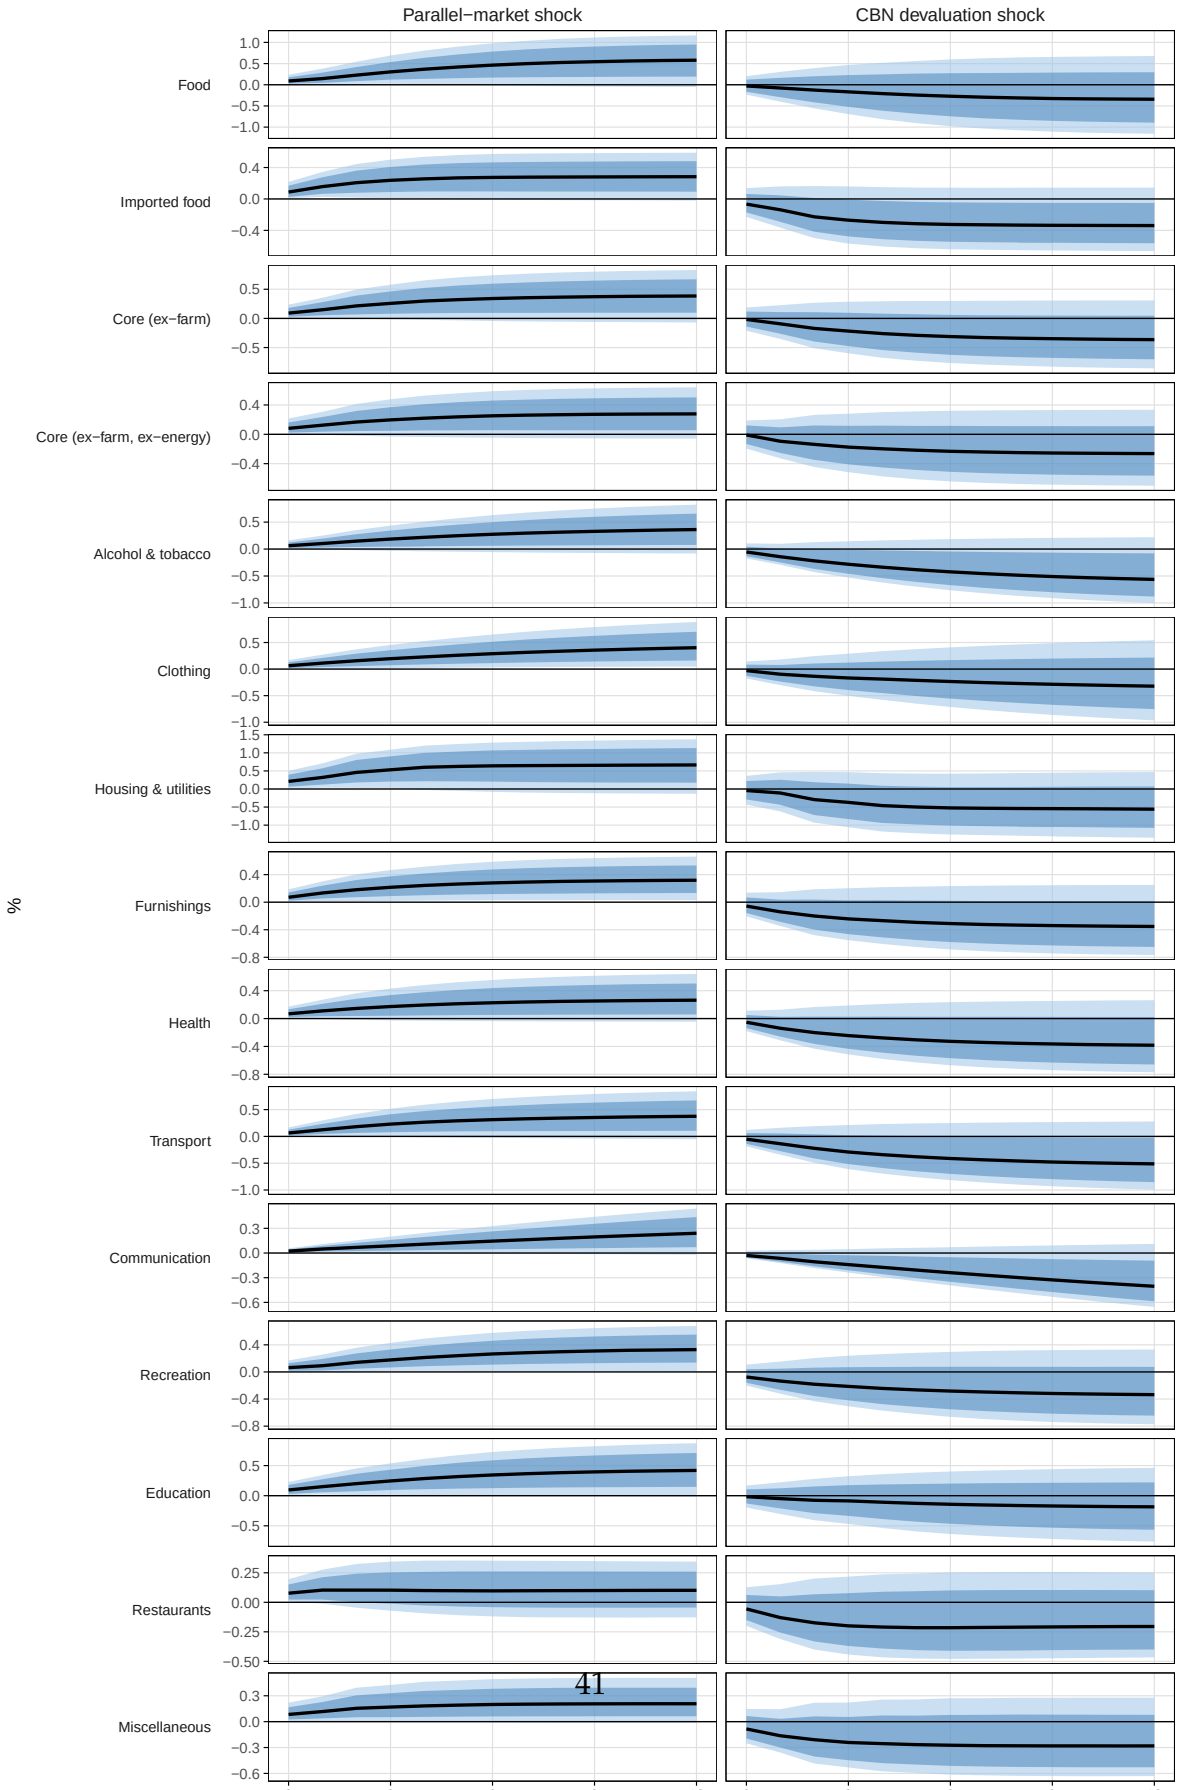

Table A12: Cumulative Pass-Through to COICOP Sub-Indices, Joint LP

	$h=0$		$h=6$		$h=12$		
Sub-index	$\widehat{\beta}$	$\widehat{\gamma}$	$\widehat{\beta}$	$\widehat{\gamma}$	$\widehat{\beta}$	$\widehat{\gamma}$	N
All items	0.001	0.012*	0.114	-0.089**	0.146*	0.023	81
	(0.015)	(0.006)	(0.089)	(0.039)	(0.085)	(0.051)	
Food	-0.018	0.006	0.049	-0.088**	0.080	0.074	81
	(0.016)	(0.007)	(0.079)	(0.037)	(0.084)	(0.066)	
Imported food	0.005	0.006	0.088	-0.137***	0.039	-0.123***	81
	(0.011)	(0.008)	(0.092)	(0.044)	(0.083)	(0.041)	
Core (ex-farm)	0.014	0.019***	0.136	-0.074*	0.167*	-0.011	81
	(0.015)	(0.007)	(0.098)	(0.043)	(0.093)	(0.051)	
Core (ex-farm, ex-energy)	0.004	0.013*	0.060	-0.067**	0.071	-0.008	81
	(0.009)	(0.007)	(0.066)	(0.028)	(0.059)	(0.033)	
Alcohol & tobacco	0.015**	0.017***	0.112*	-0.071**	0.097	-0.057	81
	(0.007)	(0.006)	(0.057)	(0.029)	(0.065)	(0.039)	
Clothing	0.001	0.019***	0.096	-0.027	0.170*	0.072	81
	(0.013)	(0.007)	(0.088)	(0.040)	(0.096)	(0.058)	
Housing & utilities	0.048	0.027	0.366*	-0.133	0.429**	-0.081	81
	(0.038)	(0.018)	(0.204)	(0.101)	(0.213)	(0.121)	
Furnishings	-0.005	0.004	0.042	-0.087***	0.049	-0.008	81
	(0.010)	(0.009)	(0.074)	(0.031)	(0.072)	(0.034)	
Health	0.003	0.005	0.017	-0.094**	-0.018	-0.050	81
	(0.009)	(0.009)	(0.052)	(0.037)	(0.038)	(0.034)	
Transport	0.006	0.023***	0.142*	-0.074*	0.139	-0.016	81
	(0.009)	(0.008)	(0.086)	(0.039)	(0.085)	(0.049)	
Communication	0.004	-0.006	-0.014	-0.102***	-0.131*	-0.153***	81
	(0.008)	(0.005)	(0.045)	(0.022)	(0.076)	(0.044)	
Recreation	-0.003	-0.008	0.045	-0.093***	0.018	-0.013	81
	(0.013)	(0.007)	(0.038)	(0.026)	(0.034)	(0.030)	
Education	0.029*	0.017*	0.183**	0.050	0.285**	0.064	81
	(0.017)	(0.009)	(0.087)	(0.060)	(0.122)	(0.080)	
Restaurants	0.003	-0.000	-0.006	-0.091**	-0.042	-0.034	81
	(0.008)	(0.006)	(0.040)	(0.036)	(0.031)	(0.037)	
Miscellaneous	-0.002	0.001	0.057	-0.095***	0.035	-0.031	81
	(0.011)	(0.007)	(0.058)	(0.032)	(0.046)	(0.033)	

Notes: Cumulative pass-through to the indicated CPI sub-index at horizons $h = 0, 6, 12$ months, from the joint local-projection specification in Equation (A1). $\widehat{\beta}$ is the coefficient on the contemporaneous log change in the parallel rate; $\widehat{\gamma}$ is the coefficient on the contemporaneous log change in the official rate. Sample: monthly Nigerian data from November 2014 to August 2022 (COICOP series end August 2022 due to NBS rebasing). Newey-West standard errors with lag 4 in parentheses. N reports the sample size at $h = 12$. *, **, *** denote significance at the 10%, 5%, and 1% levels.