Fueling Frictions: Nigeria's Fuel Subsidy Removal and Early-Life Health

Xufeng Liu[1,a] Lingyan Hu[2]

May 29, 2026

This manuscript is updated frequently. Click HERE for the latest version.

Abstract

We study how removing a universal, untargeted subsidy affects healthcare utilization across the full sequence of early-life contact with the health system. Using Nigeria's May 2023 petrol subsidy removal—a sharp, nationally unanticipated shock—we examine three outcome domains: antenatal care, birth outcomes, and childhood vaccination. Each domain uses a distinct identification design matched to data structure. For ANC and birth outcomes, we exploit variation in conception-cohort exposure: children born at different dates have different shares of their nine-month gestation window falling in the post-reform period, generating continuous within-cluster treatment variation conditional on birth-month seasonality. Post-reform exposure reduced ANC quantity and content quality, with proportionally larger effects in rural areas. Greater in-utero exposure widened urban-rural disparities in categorical birth size: rural births shifted toward smaller reported categories, while urban births shifted toward larger ones, consistent with an offsetting reduction in combustion-related $PM_{2.5}$ that disproportionately benefited urban areas. For childhood vaccination, we exploit a within-child design: children born at different times face the same national immunization schedule, so comparing vaccine windows exposed before and after the reform—within the same child—controls for all child-level time-invariant unobservables. Post-reform coverage fell for vaccines requiring return clinic visits, with shortfalls expanding monotonically across the dosing schedule. Three pieces of evidence converge on disrupted physical access—not demand withdrawal—as the dominant mechanism. Rotavirus—newly added to the free national immunization schedule in 2022—and IPV, sustained by externally funded polio-eradication outreach, both maintain coverage despite the access shock. Rural households with older siblings show smaller vaccination shortfalls, consistent with transportation costs being shared across children. These results illustrate a broader pattern: the households capturing the largest implicit transfer under an untargeted subsidy need not be those bearing the largest adjustment cost when it is removed.

JEL Codes: I12, I15, Q48, H23, O12
Keywords: fuel subsidy, antenatal care, vaccination, birth outcomes, air pollution, Nigeria, distributional effects, early-life health

[1] Department of Agriculture, Environmental, and Development Economics, The Ohio State University
[2] Department of Economics, The Ohio State University
[a] Email: liu.7133@osu.edu

I Introduction

Government spending has expanded globally, increasing from 29 percent of GDP in 2000 to 33 percent in 2019 (Zouhar et al. 2021). This expansion has increased interest in how public policies affect household welfare. A large literature shows that targeted programs, such as cash transfers, health insurance, and nutrition assistance, improve welfare for intended recipients (Almond, Hoynes, et al. 2011; Currie and Gruber 1996; Hoynes, Miller, et al. 2015) . Other policies, such as environmental regulation (Chay and Greenstone 2003; Currie and Walker 2011) and infrastructure investment (Alsan and Goldin 2019), affect broader populations, and evaluation often emphasizes average effects . A harder and less studied question concerns large fiscal reforms that affect the entire economy through multiple channels. Such reforms may benefit some households and harm others, making the distribution of costs and benefits as important as the average effect. (He, Liu, et al. 1994)

Fossil fuel subsidies are a useful setting because they belong to a broader class of universal, untargeted government interventions. Like public education, infrastructure, and universal health provision, they affect households through common access conditions rather than eligibility rules. They are also fiscally large and widespread: fossil fuel subsidies exist in more than 170 countries and reached $7 trillion globally in 2022 (Black et al. 2023; IEA 2023). In many high-income countries, governments tax fuel to raise revenue and discourage emissions; in many developing countries, governments instead subsidize fuel to keep consumer prices below market levels. Removing a fuel subsidy therefore resembles a fuel tax increase, but from the opposite policy regime (Noel et al. 2024). This makes the distributional question broader than static subsidy incidence: removal changes who loses the protection of cheap fuel, who may gain from lower fuel use and pollution, and who faces higher costs of reaching essential services. In early life, this matters because fuel prices shape the cost of every antenatal, delivery, and vaccination visit.

In this paper, we study these questions using Nigeria's May 2023 fuel subsidy removal. Nigeria's fuel subsidies had long been fiscally costly, accounting for 23 percent of the federal budget in 2022 (Ozili 2023). Their removal was both sharp and nationally unanticipated: President Tinubu's May 29 inaugural declaration that "fuel subsidy is gone" triggered an overnight increase in retail petrol prices from approximately ₦189 to ₦557 per liter — a 195 percent increase — with bus fares rising 97.9 percent within a month and food inflation accelerating by 9.1 percentage points (Dzirutwe and Eboh 2023a). We exploit this shock to provide the first empirical analysis, within a single identification framework, of how a major energy price reform affected preventive healthcare across the full sequence of early-life stages. Our outcomes span three sequential margins of the early-life health production function (Almond and Currie 2010): prenatal care utilization, at-birth outcomes including birth size and neonatal mortality, and postnatal immunization coverage and timeliness.

The first result domain covers antenatal care (pre-birth). Using a conception-cohort exposure design that compares pregnancies with different numbers of post-reform months in utero, we find that each additional month of post-reform exposure reduced the number of antenatal visits by 0.07 visits on average — a 2.1 percent decline for rural women relative to 1.0 percent for urban women. Content quality declined in parallel: blood testing, urine testing, nutritional counseling, and danger-sign counseling all fell, with proportionally larger effects in rural areas where baseline rates are lower. Importantly, the content quality declines are proportionally larger than the visit-count declines, suggesting that the budget constraint operates not only through the decision to visit but also through what households request and receive conditional on attending. The pattern is consistent with households absorbing higher transportation costs by reducing the scope of care at each visit, not only by forgoing visits altogether.

The second result domain covers birth outcomes (at-birth). We find no evidence that the reform affected neonatal mortality, either on average or differentially across urban and rural areas. Similarly, the delivery healthcare utilization does not significant change before and after the subsidy removal nor differentially across urban and rural communities. The reform instead shifted the birth size distribution. For the small-or-below measure — which corresponds most closely to the 2,500-gram low-birth-weight threshold — each additional month of full-pregnancy exposure increases the probability of being reported small or below by 0.8 percentage points in rural areas, while the urban-rural differential is -0.7 percentage points, indicating weaker deterioration in urban areas. For the broader average-or-below measure, rural areas show little change while urban births shift toward larger reported size categories, with a 1.2 percentage point relative decline per additional exposure month. The mechanism evidence identifies combustion-related air pollution as the force behind the urban relative improvement: non-dust $PM_{2.5}$ fell by 18 percent on average after the reform, with an additional 4.9 percent decline in urban areas relative to rural areas. Dust $PM_{2.5}$ shows no differential urban-rural response, confirming the effect is combustion-related rather than seasonal. The birth size effects are concentrated in the second and third trimesters, consistent with late-gestation sensitivity to both economic and environmental conditions (Currie, Neidell, et al. 2009; Painter et al. 2005).

The third result domain covers childhood vaccination (post-birth). Vaccination coverage fell for all vaccines requiring return clinic visits, with coverage shortfalls expanding monotonically across the dosing schedule — the empirical signature of compounding transportation barriers. Within-child fixed-effects estimates show that, for the same child, vaccines whose scheduled windows fall in the post-reform period are disproportionately less likely to be received. Because the survey captures children at a single point in time, raw receipt rates are susceptible to right-censoring: post-reform cohorts are on average younger at interview and have had less calendar time to complete later doses regardless of the reform. We address this by constructing an on-time outcome from vaccination card dates (available for approximately 52 percent of children), defined as receipt within the public-health-recommended age window (Newcomer et al. 2024); the dose-gradient pattern persists on this censoring-corrected measure. Rural areas experience substantially larger coverage declines. A sibling specification reveals an offsetting margin: children with an older sibling approximately twelve months their senior are more likely to receive later primary-series doses on time, and this effect is statistically significant only in rural communities. The pattern is consistent with trip-bundling — when the older sibling requires a clinic visit, the younger child's dose can be obtained in the same trip, reducing the per-dose transportation cost. In urban areas, where the marginal cost of an additional trip is lower, the bundling benefit does not generate a detectable effect.

The sharpest mechanism test comes from two vaccine-specific exceptions. The rotavirus series, newly introduced into Nigeria's routine immunization schedule in August 2022 as a free, government-backed vaccine, maintained relative coverage despite the broader decline. Its exception is informative because the free schedule addition removed the product price that households previously faced outside the NPI while the strict age eligibility cutoffs and the six- and ten-week clinic contacts remained binding. The IPV series, supported by an ongoing polio-eradication campaign with active community outreach, similarly maintained receipt rates relative to comparably scheduled vaccines. Because these vaccines are delivered through programmatic supply channels that partially decouple uptake from household-initiated travel, their exception from the dose-gradient pattern is inconsistent with general parental disengagement. The pattern is more consistent with an increase in the effective cost of accessing care than with a broad decline in households' demand for preventive health services.

This work makes several contributions to the literature. First, this paper contributes to the literature on fiscal policy, public spending, and early-life health by studying a fiscal reform that operates through economy-wide prices rather than through targeted eligibility. A large literature shows that targeted public programs can improve early-life health by expanding insurance, food assistance, or disposable income, including Medicaid, Food Stamps, WIC, and the EITC (Almond, Hoynes, et al. 2011; Currie and Gruber 1996; Hoynes, Miller, et al. 2015; Hoynes, Page, et al. 2011). A related literature studies broader public policies that affect child health through environmental regulation, infrastructure investment, or access to preventive care (Alsan and Goldin 2019; Chay and Greenstone 2003; Currie and Walker 2011; Greenstone and Hanna 2014). We depart from these literatures by studying a large untargeted fiscal reform in a developing-country setting: the removal of an economy-wide fuel-price distortion. Unlike targeted transfers, fuel subsidies do not assign treatment through program eligibility; and unlike many environmental or infrastructure policies, their removal changes several components of the early-life health production function at once. Higher fuel prices tighten household budgets, raise the cost of reaching health services, and may reduce combustion-related pollution. This tripple-channel structure makes the reform especially useful for studying how fiscal price reforms redistribute early-life health risks across space.

We therefore provide a lifecycle-style analysis of how a major energy-price reform affects early-life health, tracing the same shock across prenatal healthcare utilization, at-birth outcomes, and postnatal immunization coverage — three sequential stages in the production of early-life human capital (Almond and Currie 2010). The vaccination margin is especially novel. Existing studies typically examine in utero economic shocks, antenatal care, pollution exposure, or vaccine uptake as separate margins, and we are not aware of prior work that causally identifies how an aggregate fuel-price shock affects routine immunization coverage and timeliness using individual-level data. Tracing all three stages simultaneously reveals that the distributional consequences of the reform cannot be recovered from any single margin: the same fiscal price shock can reduce access to prenatal care, alter fetal growth, and change postnatal preventive investments, while the access-cost and environmental channels operate with different spatial incidence and vary across the distribution of outcomes.

Second, this paper contributes to the literature on energy subsidies, subsidy reform, and distributional welfare. Existing work documents the efficiency costs, fiscal burden, and regressive static incidence of fuel subsidies: because richer households consume more fuel directly, they capture a disproportionate share of subsidy benefits, while poorer households benefit mainly through indirect price effects (Coady et al. 2015; Davis 2014; Soile and Mu 2015). A related literature studies the welfare consequences of subsidy removal, emphasizing pass-through to household expenditures, savings, debt, and transportation costs (Arze del Granado et al. 2012; Folami et al. 2025; Rentschler 2016; Yahaya 2026).

We extend this literature by showing that the distributional consequences of removing a universal price subsidy need not mirror the static incidence of the subsidy itself. When a fuel subsidy is removed, households do not simply lose a transfer embedded in fuel prices; they face a new cost structure for reaching work, markets, schools, and health facilities. For early-life health, this distinction is first-order because antenatal care, delivery care, and childhood vaccination all require repeated physical contact with the health system. The households most exposed to the dynamic health costs of removal therefore need not be the households that consumed the most subsidized fuel. In our setting, subsidy removal generated larger pollution reductions in urban areas but larger access-cost pressures in rural areas, producing uneven effects across prenatal care, reported birth size, and immunization. The paper therefore shows that evaluating universal subsidy reform requires tracing not only who benefits from the subsidy while it exists, but also who bears the downstream costs when the subsidy is withdrawn.

Third, this paper contributes to the literature on barriers to preventive care and routine immunization in low-income settings. Existing work shows that distance, transportation costs, and out-of-pocket expenses reduce utilization of antenatal care and vaccination (Banerjee et al. 2010; Kruk et al. 2009; Meriggi et al. 2024), and that trust in providers and government shapes vaccine demand (Martinez-Bravo and Stegmann 2022). A harder question is whether observed utilization declines after an aggregate price shock reflect higher costs of reaching facilities or reduced household demand for care — a distinction that matters for policy because the two channels call for different responses. We address this using two complementary designs: a within-child fixed-effects estimator that holds all household-level unobservables constant, and an across-child within-vaccine estimator that exploits variation in the timing of required vaccine windows relative to the reform. Both designs yield a dose-gradient pattern: coverage shortfalls compound monotonically across the immunization schedule, consistent with transportation-cost barriers accumulating over required return visits. The mechanism test comes from vaccines with active programmatic support or distinctive cost structures:

the rotavirus series, newly introduced into the free routine schedule in August 2022 while retaining strict age eligibility cutoffs and required six- and ten-week clinic contacts (Ukazu 2022), and IPV, supported by ongoing polio-eradication campaigns (Nigeria 2023), maintain relative coverage despite the broader decline. Because these exceptions correspond to vaccines with stronger programmatic attention, sharper supply-side incentives, or a separation between vaccine price and visit cost, they are difficult to reconcile with a pure parental disengagement story. Instead, the pattern suggests that the reform increased the effective cost of accessing routine care, with programmatic support partially insulating some vaccines from the broader access shock.

The remainder of the paper proceeds as follows. Section II provides institutional background on the fuel subsidy reform and Nigeria's immunization program. Section III describes the data and samples. Section IV presents immediate evidence on healthcare access from the GHS Panel. Section V presents the empirical strategy for the DHS-based early-life analysis. Sections VI–VIII report the three sets of results. Section IX documents the household welfare aftermath. Section X presents the bridging-share IV. Section XI concludes.

II Setting

A Nigeria Fuel Subsidy and the 2023 Reform

Fossil fuel consumption subsidies remain widespread. Global subsidies reached $7 trillion in 2022, approximately 7.1 percent of world GDP (Black et al. 2023). Yet subsidy reform is difficult in practice. Clements et al. (2013) document 28 energy-subsidy reform episodes across 22 countries and classify only 12 as successful, with 11 partially successful and 5 reversed outright. Nigeria is a particularly important recent case because of the scale of its gasoline subsidy and the abrupt removal announced on May 29, 2023.

History. Nigeria's gasoline subsidy has long been central to the country's political economy. The subsidy applied primarily to premium motor spirit (PMS), the main gasoline product in Nigeria, used primarily for road transportation and generators. The subsidy was formalized under the 1977 Price Control Act, instituted in response to the first oil crisis of the 1970s, which prohibited petrol sales above a government-set price. Over time, the subsidy became deeply entrenched. Low fuel prices were one of the most visible benefits the state provided, and citizens came to view below-market pricing as an entitlement rather than a policy choice. By 2022, the program had grown to an estimated $10 billion per year, consuming 23 percent of the federal budget and more than four times the federal health expenditure (IISD 2023). In order to relieve the fiscal burden, the government repeatedly attempted reform, but each effort ultimately failed. A major removal effort in January 2012 provoked nationwide strikes and was partially reversed within weeks. The federal government again announced price deregulation in June 2020, but implicit subsidies had re-emerged by January 2021 as rising import costs were not passed through to retail prices (International Monetary Fund 2024).

The May 2023 Reform. The next major reform began on May 29, 2023, when President Bola Tinubu declared "fuel subsidy is gone" in his inaugural address (Majeed 2023). NNPC confirmed the removal the following day, and retail petrol prices rose from about ₦189 to about ₦557 per liter overnight (Dzirutwe and Eboh 2023b). Figure 1 shows the sharp increase in retail fuel prices after the reform. The effects of the price increase were immediate. The transportation and storage sector contracted by 30.7 percent in the first nine months of 2023, largely because of the fuel price shock (World Bank 2023). Because petrol is a basic input into transportation and generator-based electricity, the price increase passed through rapidly to consumer prices. Between May and June 2023, average within-city bus fares increased by 97.9 percent, intercity bus fares by 42.1 percent, and motorcycle taxi (*okada*) fares by 33.1 percent. Food inflation increased by 9.1 percent (National Bureau of Statistics 2023c). Headline inflation rose from 22.4 percent in May 2023 to 28.9 percent by December, but nominal wages did not adjust until 2024, implying that real purchasing power fell sharply in the months following the removal (National Bureau of Statistics 2023b,a; World Bank 2024).

Concurrent macro context. The 2023 macro environment introduces identification concerns that motivate the IV strategy in Section X. Alongside the subsidy removal, the Tinubu administration unified the official and parallel naira exchange rates in June 2023, triggering a sharp currency depreciation. The Central Bank of Nigeria also implemented a currency redesign/re-denomination in early 2023, which created acute cash shortages in some states. These are concurrent national-level shocks that affect real incomes and liquidity through channels distinct from the fuel price change. In principle, a simple before-after comparison could attribute to the subsidy removal effects that instead reflect the currency shock, and the currency shock may have been more severe in states with different financial development or urban structures. The bridging-share IV is designed to address this: it isolates variation in fuel price pass-through that comes from pre-reform logistics infrastructure rather than from the currency or monetary environment.

Post-reform policy reversal. The post-reform regime did not produce a lasting shift to market pricing. As the naira depreciated and import costs rose, the government again kept retail fuel prices below cost-recovery levels by late 2023, effectively reintroducing an implicit subsidy. The IMF documented this reversal in January 2024 and estimated that the renewed subsidy could reach 3 percent of GDP (International Monetary Fund 2024). Our estimation sample ends before the late-2023 reversal, so the births and children we study were exposed to the initial, unattenuated price shock. The reversal itself provides a falsification test for the pollution channel: if combustion-related $PM_{2.5}$ improved because of reduced fuel consumption, then the partial re-emergence of the subsidy should narrow the urban-rural $PM_{2.5}$ gap — which it does (Section VII).

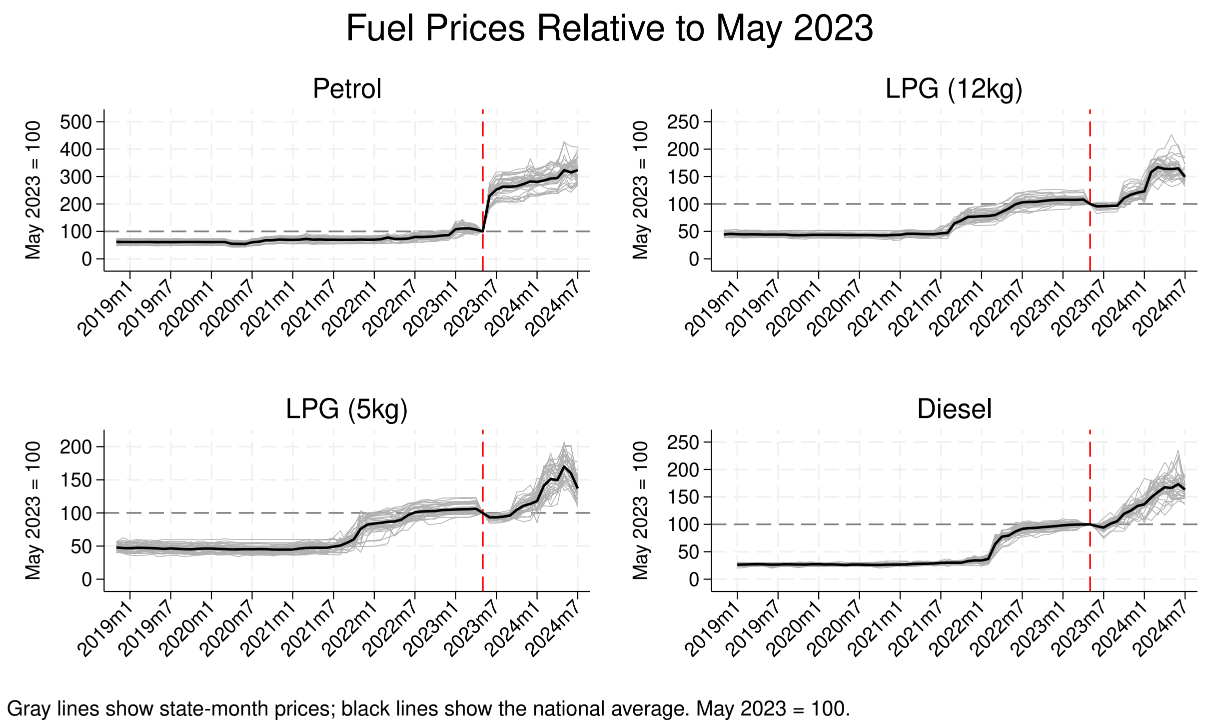

Figure 1: PMS Retail Price Around the Subsidy Removal

Notes: This figure plots monthly state-level fuel prices for petrol, 5kg LPG, 12kg LPG, and diesel. Grey lines show individual state-level price series, while the solid black line shows the national average. For each fuel type, prices are indexed to May 2023, with the May 2023 value normalized to 100. Thus, values above 100 indicate price increases relative to May 2023, and values below 100 indicate price decreases.

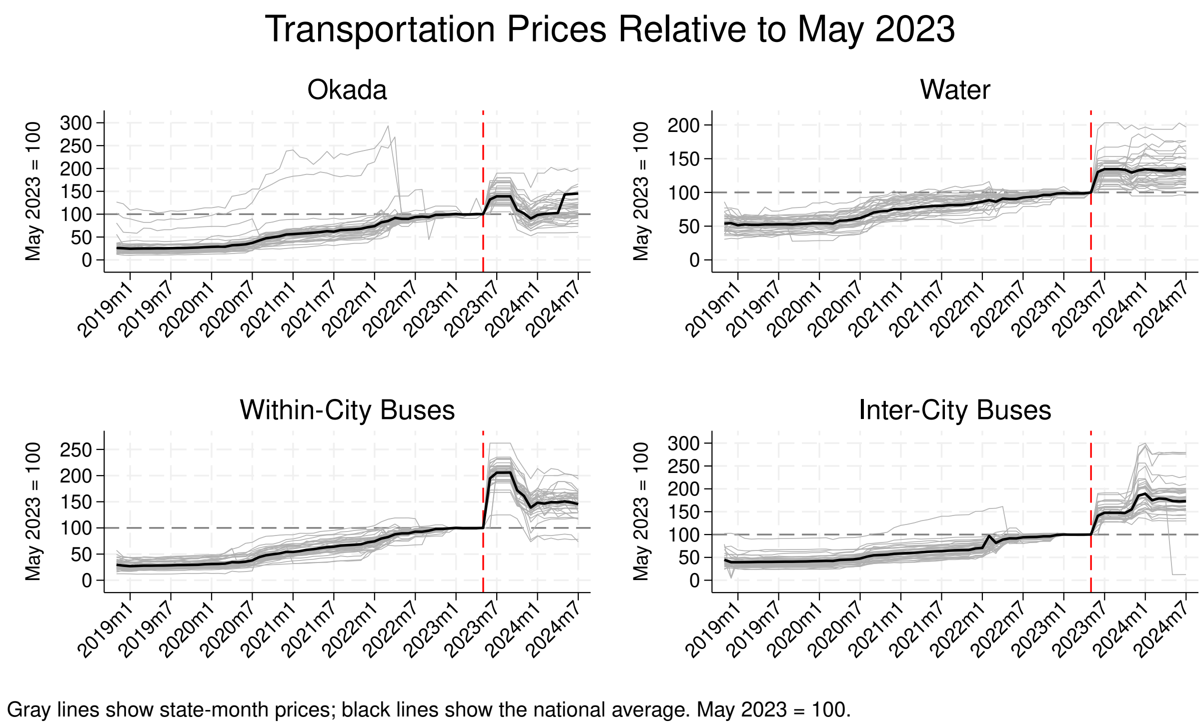

Figure 2: PMS Retail Price Around the Subsidy Removal
Notes: This figure plots monthly state-level transportation prices for okada (motorcycle), inter-city buses, within-city buses, and water transportation. Grey lines show individual state-level price series, while the solid black line shows the national average. For each fuel type, prices are indexed to May 2023, with the May 2023 value normalized to 100. Thus, values above 100 indicate price increases relative to May 2023, and values below 100 indicate price decreases. Source: NBOS Fare Watch Data

B Nigeria's Health Care System and Immunization Schedule

Nigeria operates a three-tier public healthcare system in which administrative and financial responsibility is divided across the three levels of government, a structure formalised in the 1988 National Health Policy and reinforced by the National Health Act of 2014 (Abubakar et al. 2022). Primary healthcare centres (PHCs) — approximately 30,000 nationwide — form the first point of contact for preventive and basic curative services and are managed by the 774 Local Government Areas (LGAs), with technical support from the State Primary Health Care Development Agency (SPHCDA) (Abubakar et al. 2022). Secondary care is provided through state-operated general hospitals, which serve as referral destinations for cases exceeding PHC capacity and are administered by State Ministries of Health and Hospital Management Boards. Tertiary care — delivered through federal teaching hospitals, Federal Medical Centres, and specialist institutions — falls under the Federal Ministry of Health, which also sets national health policy, coordinates vaccine management, and oversees disease surveillance (Abubakar et al. 2022). In practice, this fragmented governance structure means that the quality and availability of care vary substantially across states and between urban and rural areas, with coordination failures between tiers a persistent structural constraint (Abubakar et al. 2022).

The cost structure of healthcare differs sharply across facility type and provider ownership. Within the public system, all vaccines included in the National Programme on Immunisation (NPI) — Nigeria's implementation of the WHO Expanded Programme on Immunisation, established in 1978 — are provided free of charge to all children at public health facilities (*Immunization Page – Paediatric Association of Nigeria* 2026). Beyond immunisation, however, out-of-pocket payments dominate health financing. Public expenditure accounts for only 15 to 20 percent of total health expenditure in Nigeria, with out-of-pocket spending representing approximately 70 to 76 percent of all healthcare costs — among the highest shares in the world (Abubakar et al. 2022; Aniebo et al. 2025). User fees apply at government-owned PHCs and hospitals, introduced under the Bamako Initiative framework, though PHC fees are nominally low (Dahiru and Oche 2015). Private facilities, which serve a growing share of the urban population, charge substantially higher consultation, diagnostic, and pharmaceutical fees. The concentration of out-of-pocket costs in private and secondary-tier facilities means that the effective price of accessing care above the PHC level rises steeply for households without insurance, leaving the majority of the population — particularly in rural areas — dependent on PHC services whose quality and staffing are highly variable.

Understanding why a fuel price shock would have dose-specific and vaccine-specific effects on vaccination requires understanding the structure of Nigeria's routine immunization schedule. The National Programme on Immunization (NPI) in Nigeria follows the WHO Expanded Programme on Immunization (EPI) schedule, which specifies both the target vaccines and the recommended age window for each dose. The schedule is organized into five contact points, each requiring a separate facility visit or outreach encounter. Appendix figure B1 presents the latest Nigeria immunization schedule (UNICEF 2026).

The first contact point is **at birth**, when infants receive BCG (tuberculosis), Hepatitis B (HepB0), and oral polio (OPV0). These vaccines do not require a cold-chain return visit by the caregiver, since they are administered at delivery. The second contact is at **six weeks**, when infants receive the first doses of pentavalent (Penta1: DTP-Hep B-Hib), pneumococcal conjugate (PCV1), oral polio (OPV1), and rotavirus (Rota1). The third contact at **ten weeks** adds Penta2, PCV2, OPV2, and Rota2. The fourth contact at **sixteen weeks** adds Penta3, OPV3, and inactivated polio vaccine (IPV). The fifth contact at **nine months** adds measles/rubella (MR) and yellow fever. Each contact after birth requires an active return visit by the caregiver to a facility or outreach site. Any increase in transportation costs, reduction in outreach frequency, or household income shock therefore imposes a barrier at each subsequent contact, and the barriers compound with dose order.

A pre-existing and well-documented rural-urban coverage gap shapes the baseline from which the reform operates. Nigeria's routine immunization coverage among children aged 12–23 months is substantially lower in rural areas (Henry 2011; Ophori et al. 2014), reflecting longer distances to facilities, lower outreach intensity, lower maternal education, and weaker cold chain maintenance. Urban areas concentrate facilities and outreach workers; rural vaccination relies more heavily on mobile teams and periodic campaigns, both of which are sensitive to transportation cost increases on the supply side.

Rotavirus introduction, cost structure, and age cutoffs. Rotavirus is unique in the Nigerian schedule along two dimensions. First, it was newly introduced into the routine immunization schedule in August 2022, moving protection against rotavirus from a vaccine households generally had to purchase outside the NPI into the free public schedule (Okafor 2021; Ukazu 2022). This changed the monetary price of the vaccine itself, but it did not eliminate the access requirement: families still had to make the separate six- and ten-week facility visits or outreach contacts to receive Rota1 and Rota2. Second, rotavirus has strict WHO and manufacturer-specified age eligibility windows. The first dose of Rota must be given between six and fourteen weeks of age; the second dose by eight months (and at least four weeks after the first dose). These cutoffs exist because rotavirus vaccine efficacy depends on early administration before widespread natural exposure, and live attenuated rotavirus vaccines carry a small risk of intussusception that is higher when administered to older infants.

Together, the cost-structure change and the age cutoffs make rotavirus especially useful for interpretation. The August 2022 schedule addition sharply reduced the out-of-pocket price of the vaccine product, while the required clinic contacts continued to expose households to transportation and time costs. The age cutoffs then create a qualitatively different supply-side incentive for clinic workers: when scheduling is compressed by reduced outreach or higher caregiver no-shows, health workers have a programmatic reason to prioritize Rota (and complete the eligible doses before the window closes) that does not apply equally to vaccines with no binding age cutoff. The dose-gradient pattern — in which coverage falls monotonically with dose order — is therefore predicted to be absent for rotavirus if access disruption rather than demand withdrawal is the operative mechanism.

IPV and the polio eradication context. IPV was introduced into Nigeria's routine schedule in 2015 as part of the global polio eradication initiative's transition from trivalent to bivalent OPV (Nigeria 2015). Nigeria was certified polio-free in 2020, but the wild poliovirus threat and the ongoing use of supplemental immunization activities (SIAs) mean that IPV delivery remains a high-priority programmatic target (Nigeria 2023). IPV is given as a single dose at sixteen weeks, the same contact as Penta3 and OPV3. Its maintained coverage despite the access-disruption shock reflects two features: first, IPV doses are often delivered in bundled outreach activities alongside other polio interventions, making supply-side delivery less dependent on caregiver-initiated visits; second, the single-dose schedule makes IPV relatively easier to complete than the three-dose primary series for Penta. The maintained coverage for both rotavirus and IPV, the two vaccines with the strongest prioritization or programmatic delivery incentives, provides convergent evidence that the mechanism is supply-side access disruption rather than generalized demand withdrawal.

Takeaway for empirical design. The schedule structure generates two testable predictions. First, dose-gradient deterioration — coverage falling monotonically with dose order — is the predicted signature of a transport-cost shock, because each additional contact accumulates a per-visit barrier. Second, vaccine-specific exceptions to the dose gradient should be informative about mechanism: a vaccine that maintains coverage despite requiring return visits must do so through supply-side prioritization or a cost structure that separates the vaccine price from the visit cost, not through parental compliance alone. Both predictions are confirmed in Section VIII.

III Data and Sample

A 2023–24 Nigeria DHS

The primary data source for all three outcome domains is the 2023–2024 Nigeria Demographic and Health Survey (DHS), a nationally representative household survey conducted by the National Population Commission and ICF. The survey interviewed 40,047 households and 39,050 women aged 15–49 in 1,400 clusters. Fieldwork ran from December 2023 through May 2024, meaning the survey began approximately six months after the reform. The DHS provides three distinct modules that map onto the three lifecycle stages of the analysis.

Birth outcomes and antenatal care. Birth outcomes are drawn from the DHS birth history records, which contain retrospective birth histories including the timing of each birth, the mother's assessment of the child's size at birth, survival status, and age at death. We focus on size at birth as the primary outcome because recorded birth weight is unavailable for a large share of births and is observed disproportionately among facility deliveries. Mothers report birth size on a five-point scale; we construct three cumulative indicators: *very small*; *small or below*; and *average or below*. We also construct a neonatal mortality indicator. From the DHS maternity module, we construct measures of antenatal care utilization (number of visits, early initiation, skilled provider) and ANC content quality (blood test, urine test, blood pressure check, food counseling, bleeding danger-sign counseling). Variable definitions are in Appendix A.

Childhood vaccination. The DHS child vaccination module records whether each child under five years of age received each vaccine in the EPI schedule, and — for children whose caregivers present a vaccination card — the date each dose was administered. This yields two complementary coverage measures: *survey-reported vaccination* (whether the caregiver reports the child ever received the vaccine, available for all children under 36 months) and *card-recorded vaccination* (dose date from the physical card, available for the approximately

48 percent of children under 36 months whose caregivers presented a card at interview). The survey-reported measure provides full-sample coverage but captures only whether the child was ever vaccinated, not whether vaccination occurred on time. The card-recorded measure is necessary for timeliness analysis — we use it to construct indicators for whether each dose was administered within a specified window of the recommended age. The card restriction is a meaningful limitation: card-holding households are more engaged with the health system and likely to be systematically less affected by the reform than the full population. Effect estimates from the timeliness analysis therefore likely understate the true reduction in timely vaccination.

Spatial and household characteristics. The DHS classifies each cluster as urban or rural based on the national census classification; this provides the spatial variation exploited in all three designs. Household-level controls include maternal age, education, birth order, child sex, and household wealth quintile. Wealth is measured at interview, which is post-reform; wealth splits are therefore interpreted with caution.

A.1 Sample Restrictions and Temporal Alignment

Both the birth-outcome and vaccination samples share a common lower bound: births (and conceptions) after January 1, 2022. This restriction is designed to exclude the COVID-19-era healthcare disruptions that substantially altered healthcare utilization in Nigeria from early 2020 through approximately late 2021. Nigeria's most disruptive COVID restrictions — suspension of routine immunization outreach, limits on facility visits, and movement restrictions — were substantially eased by mid-2021, and the health system had recovered to near-normal utilization patterns by early 2022. We use January 2022 as the start of the clean pre-period; cohorts born before this date are excluded from both samples. A key concern is that differential urban-rural COVID recovery introduces trending pre-reform differences; the event-study specifications in each section address this by showing flat pre-reform trends.

The **birth-outcome sample** is restricted to deliveries in the $[-8, +6]$ month window around the May 2023 reform (approximately September 2022 through November 2023). This keeps estimation close to the reform and ends the sample before the partial subsidy reinstatement of late 2023. We additionally exclude women who were visitors at interview and women who moved across states after delivery, because DHS cluster location is used to assign pollution and weather exposure during pregnancy. The final birth-outcome sample contains 6,730 singleton births across 1,163 DHS clusters.

The **vaccination sample** includes all children born on or after January 1, 2022 and surveyed before May 2024. Children's ages at interview range from approximately four months to 28 months, providing sufficient follow-up to observe the full primary vaccination series. Exposure to the reform is defined at the vaccine level: for each recommended vaccine dose, we compute whether the child's age-eligible window for that dose falls entirely before the reform, partly in the post-reform period, or entirely in the post-reform period. This creates within-child variation across vaccines, exploited in the within-child fixed-effects design. The survey-reported vaccination sample includes approximately 10,855 children; the, card-verified subsample for timeliness analysis includes approximately 5,742 children.

B Pollution Data

MERRA-2 Pollution Data. Monthly pollution data come from NASA's Modern-Era Retrospective Analysis for Research and Applications, Version 2 (MERRA-2), specifically the M2T1NXAER aerosol diagnostics product. MERRA-2 reports hourly, time-averaged aerosol fields on a 0.5° latitude by 0.625° longitude global grid. We aggregate to monthly measures of total $PM_{2.5}$, non-dust $PM_{2.5}$, and dust $PM_{2.5}$ for each grid cell in Nigeria from November 2022 to December 2023. Distinguishing non-dust $PM_{2.5}$ from total $PM_{2.5}$ is important in Nigeria because Saharan dust substantially increases particulate concentrations during the Harmattan season (roughly November through February). This decomposition separates combustion-related pollution — the channel directly affected by fuel price changes — from seasonal dust fluctuations. Each DHS cluster is assigned the monthly pollution value for the grid cell in which it falls. MERRA-2 pollution data are used only in the at-birth section (Section VII).

Additional Pollution Data. In the next iteration of the draft, we will supplement the MERRA-2 pollution measures with two additional sources of air pollution data. First, we will integrate the Copernicus Atmosphere Monitoring Service global reanalysis product (CAMS EAC4) as a robustness check for the main MERRA-2 results. Like MERRA-2, CAMS EAC4 provides spatially complete, model-based estimates of atmospheric composition on a global grid, but it relies on a different atmospheric model, chemical transport system, emissions inventory, and data assimilation framework. Re-estimating the pollution results using CAMS EAC4 will therefore allow us to assess whether the estimated changes in particulate exposure after the 2023 fuel subsidy removal are specific to the MERRA-2 aerosol product or are robust to an alternative reanalysis system.

Second, we will incorporate PurpleAir sensor data where available to capture higher-frequency and more localized changes in air pollution. Unlike MERRA-2 and CAMS EAC4, PurpleAir provides ground-level observations from low-cost optical sensors rather than gridded reanalysis estimates. These data are therefore better suited to detecting short-run and local pollution fluctuations, such as changes near roads, fuel stations, traffic corridors, generator-heavy neighborhoods, or dense urban areas. PurpleAir coverage in Nigeria is spatially selective, so these data will not replace the nationally consistent MERRA-2 exposure measure. Instead, they will be used as a complementary validation and event-level source, allowing us to examine whether observed ground-level $PM_{2.5}$ changes in sensor-covered locations move in the same direction as the broader reanalysis-based pollution measures.

C The GHS Panel

To complement the DHS-based design with contemporaneous evidence on healthcare access, we draw on the Nigeria General Household Survey-Panel (GHS-Panel), the household panel implemented by the National Bureau of Statistics as part of the World Bank's Living Standards Measurement Study–Integrated Surveys on Agriculture. The GHS-Panel follows roughly 5,000 households nationwide and is the parent survey from which the NLPS households described below are drawn. A refreshed panel was introduced in Wave 4 (2018), and each wave comprises two visits—a post-planting visit and a post-harvest visit. Five waves have been fielded to date; we use Wave 4 as the pre-reform baseline and Wave 5 as the post-reform follow-up.

The health module is administered once per wave. In Wave 4 it was collected in the post-harvest visit of January–March 2018, whereas in Wave 5 it was collected in the post-planting visit of July–October 2023, roughly one to five months after the subsidy removal. The two health modules are therefore separated by a six-month shift in the calendar, so seasonal differences in disease prevalence cannot be fully differenced out; we return to this caveat in Section IV. The countervailing advantage is timing: the Wave 5 module provides an immediate post-removal snapshot of healthcare access, in contrast to the DHS, whose fieldwork began roughly six months after the reform.

The module records healthcare utilization over the four weeks preceding the interview and records cost and access details for the most recent visit. We construct three whole-sample outcomes—an indicator for any illness or injury, an indicator for seeking care conditional on illness, and an indicator for forgoing needed care—defined for the full surveyed household population of all ages. Among those who sought care, we additionally observe the mode of transport to the facility, out-of-pocket transport cost, and waiting and consultation times for the most recent visit. Wave-5 summary statistics are reported in Appendix Table C2.

D Supplementary Data

NLPS 2021–2024. Living standards data come from the Nigeria National Longitudinal Phone Survey (NLPS) 2021–2024, a high-frequency household panel survey implemented by the National Bureau of Statistics. The NLPS follows approximately 2,400 households originally sampled in the 2018–2019 General Household Survey-Panel. Households report adequacy of food, clothing, and housing consumption. We use rounds 8–10, fielded between April and November 2023, which straddle the May 29 reform. The NLPS is not linked to the DHS birth records; it provides independent corroboration of the household welfare impact rather than a mechanism estimate tied to individual births.

Afrobarometer Round 10. Post-reform public opinion data come from Afrobarometer Round 10, fielded in Nigeria in June - July 2024 covering 1,600 adult citizens. The survey includes questions on attitudes toward the fuel subsidy removal, government economic management, and satisfaction with public services. We use the Afrobarometer to document public awareness and the welfare salience of the reform, providing a welfare bookend that connects the measured health effects to perceived welfare costs. The Afrobarometer is used only in Section IX.

Summary statistics. Table 1 reports summary statistics for the birth-outcome sample and vaccination outcome sample. Neonatal mortality is 3.8 percent. Mothers report 4.0 percent of newborns as very small at birth, 12.6 percent as small or below, and 74.5 percent as average or below. Average total PM2.5 is 59.8 ug/m3, of which 48.0 ug/m3 is dust and 11.8 ug/m3 is non-dust particulate matter. Maternal healthcare utilization is limited. Overall, 55.0 percent of deliveries occur at home, and only 46.7 percent are attended by skilled assistance. Mothers receive 4.2 antenatal checks on average, and 62.9 percent receive skilled antenatal care.

For the vaccination sample, the table present a selected subset of vaccine. By the time of the interview, only 66% of surveyed children born after Jan 2022 receive a dose of BCG vaccines. This drops to only 46% for the second dose of Rotavirus vaccine and 50% for measles. Only about half (52.8%) of the children has a vaccine card available and seen at the time of the interview.

Table 1: Summary Statistics

Variable	Obs	Mean	S.D.
Panel A: Pre-birth and At-birth outcome sample			
Birth outcomes			
Neonatal mortality	6,888	0.038	0.191
Very small	6,743	0.040	0.196
Small or below	6,743	0.126	0.331
Average or below	6,743	0.745	0.436
Healthcare utilization			
Delivery at home	6,888	0.550	0.498
C-section	6,888	0.054	0.225
Skilled delivery assistance	6,888	0.467	0.499
Number of ANC visits	6,666	4.239	4.106
Early ANC	6,833	0.188	0.391
Skilled ANC	6,842	0.629	0.483
Air pollution			
$PM_{2.5}$ ($\mu g/m^3$)	20,700	59.769	38.203
Non-dust $PM_{2.5}$	20,700	11.805	5.742
Dust $PM_{2.5}$	20,700	47.964	37.617
Panel B: Vaccination sample			
BCG received (if rquired)	10,875	0.6643	0.4723
Penta3 received (if rquired)	9,488	0.5152	0.500
IPV1 received (if rquired)	10,300	0.558	0.497
Rota2 received (if rquired)	9,893	0.465	0.499
Measles received (if rquired)	6,888	0.503	0.500
Card available	10,875	0.528	0.499

Notes: Birth-outcome sample restricted to singleton births in the $[-8, +6]$ month window around May 2023, and the vaccination sample refer to the childre born after Jan 2022.

IV Contemporaneous Healthcare Access: Illness and Care-Seeking

Before turning to the conception-cohort design that anchors the early-life results, we present a simpler, contemporaneous look at how the reform affected healthcare access for the broad population. The GHS-Panel health module (Section C) was fielded in July–October 2023, within months of the subsidy removal, and asks whether household members were ill in the past four weeks and whether they sought care. Because it captures the population's experience almost immediately after the price shock—and across all ages rather than only mothers and young children—it provides motivating evidence on the access-cost mechanism that the remainder of the paper investigates in depth.

A Specification

We relate wave-5 outcomes to the contemporaneous change in local fuel prices, using wave 4 as the pre-reform baseline. For individual i in household h and state s, interviewed in month t, we estimate

$$Y_{ihs} = \beta\, \Delta P_{st} + \kappa\, (\Delta P_{st} \times \text{Urban}_h) + \mathbf{X}_i'\gamma + \mathbf{Z}_h'\delta + \theta_{z(h)\, u(h)\, t} + \varepsilon_{ihs}, \tag{1}$$

where ΔP_{st} is the change in the NBOS-recorded state pump price between the interview month and the pre-reform average, expressed relative to the pre-reform price, so that one unit corresponds to a 100 percent increase. Urban_h is an urban indicator. $\mathbf{X}_i$ collects individual covariates (age, sex, marital status, and whether the respondent is the household head); $\mathbf{Z}_h$ collects household covariates (household size and the household head's literacy, schooling, education, sex, and age). $\theta_{z(h)\, u(h)\, t}$ are zone $\times$ urban $\times$ interview-month fixed effects, which absorb seasonal differences in disease prevalence common to a zone and sector in a given calendar month—the principal threat created by the six-month shift between the wave-4 and wave-5 health modules. Standard errors are clustered at the state level, and all regressions use wave-5 weights.

We report two specifications for each outcome. Panel A omits the interaction and reports the average effect β; Panel B adds the Urban $\times$ ΔP interaction, so that β is the rural effect and κ the urban-rural differential. Each table reports the wave-4 mean of the dependent variable as a baseline reference; the missed-care specification additionally controls for the respondent's wave-4 value of the outcome and a flag for its missingness.

B Illness and Care-Seeking

The central result in Table 2 is that the reform raised reported illness significantly (0.159^{**}, on a base mean of 0.24) while leaving care-seeking statistically unchanged (-0.014), producing a significant rise in missed care (0.182^{**}, on a base of 0.07). The gap between the illness and care-seeking coefficients is not mechanical: an F-test of their equality rejects at the five percent level ($F = 5.26$, $p = 0.028$), confirming that care-seeking was suppressed beyond what the shift in illness rates alone would predict. This is the access-barrier result—people became sicker but did not access care commensurately. Panel B shows where the burden falls. The missed-care increase is concentrated in rural areas (rural effect 0.206^{**}, urban interaction null), while urban areas show a significant positive care-seeking response (urban interaction 0.349^{**}). Urban households partially maintained or even increased care-seeking under the same price shock that suppressed rural access. This urban-rural asymmetry in the margin of adjustment—rural households losing access while urban households sustain it—is the distributional headline.

The urban-age heterogeneity in Figure 3 sharpens this pattern along two dimensions. First, among the younger age groups, rural children under five show the largest declines in care-seeking and the largest corresponding increases in missed care, motivating the focus on early-life health in the DHS analysis that follows. Second, the urban-rural divergence widens with age: rural elderly show large positive missed-care effects, whereas urban elderly show, if

Table 2: Illness and Care-Seeking

	Had illness	Sought care	Missed care
Panel A. Main effect			
Fuel price change	0.1593**	-0.0144	0.1821**
	(0.0741)	(0.0884)	(0.0792)
Panel B. Urban heterogeneity			
Fuel price change	0.1266	-0.0763	0.2056**
	(0.0798)	(0.0975)	(0.0867)
Fuel price change x urban	0.1841	0.3487**	-0.1325
	(0.1437)	(0.1626)	(0.0849)
Dep. Var mean (SD), wave 4	0.2419 (0.4282)	0.1791 (0.3834)	0.0712 (0.2572)
Observations	26867	26867	26867
Sample	Full	Full	Full

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Standard errors, clustered by state, in parentheses.
Outcomes are wave-5 indicators for any illness or injury in the past four weeks, for seeking care conditional on illness, and for forgoing needed care. The missed-care model additionally controls for its wave-4 value and a missingness flag.
Controls: individual age, sex, marital status, and household-head status, plus household size and household-head demographics.
All specifications include zone × urban × interview-month fixed effects and use wave-5 weights.

anything, negative missed-care effects, driven by maintained or increased care-seeking despite rising illness. The shock thus redistributed access costs most severely onto rural children and the rural elderly—the two groups with the least independent capacity to navigate care access. We do not pursue the elderly margin further here: the GHS health module is too thin on the downstream care-seeking and outcome detail needed to interpret it, and the DHS, which anchors the rest of the paper, does not cover older adults. We therefore flag the elderly results as a direction for future work and concentrate on the early-life evidence below.

Among those who did seek care, the transport and time-cost margins in Table 3 are imprecisely estimated but trace a coherent trade-off. Relative to rural care seekers facing the same price increase, urban care seekers appear less likely to reach the facility by motorized transport (-0.489) and to incur lower transport costs (-3.36 in logs), consistent with their shorter travel distances and denser facility networks. The same urban care seekers, however, face longer waiting ($+1.00$ in logs) and consultation ($+1.02$ in logs) times—a plausible congestion counterpart to the higher care-seeking documented above. None of these urban differentials is statistically significant, so we read them as directional rather than conclusive.

Table 3: Care Transport and Time Costs

	Motorized	Log transport	Log wait	Log consult
Panel A. Main effect				
Fuel price change	-0.0447	-0.4887	0.2882	0.2562
	(0.2617)	(1.5605)	(0.4244)	(0.4266)
Panel B. Urban heterogeneity				
Fuel price change	0.0520	0.1762	0.0899	0.0537
	(0.3079)	(1.8347)	(0.4684)	(0.4728)
Fuel price change x urban	-0.4892	-3.3636	1.0032	1.0245
	(0.5362)	(3.3363)	(0.9500)	(0.9136)
Dep. Var mean (SD), wave 4	0.4532 (0.4980)	2.5421 (2.8486)	2.1610 (1.2395)	2.6489 (0.9722)
Observations	5364	5364	5364	5364
Sample	Care seekers	Care seekers	Care seekers	Care seekers

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Standard errors, clustered by state, in parentheses.
Outcomes, estimated on the sample of care seekers and measured for the most recent visit, are an indicator for motorized transport and logs of transport cost, wait time, and consultation time.
Controls: individual age, sex, marital status, and household-head status, plus household size and household-head demographics.
All specifications include zone × urban × interview-month fixed effects and use wave-5 weights.

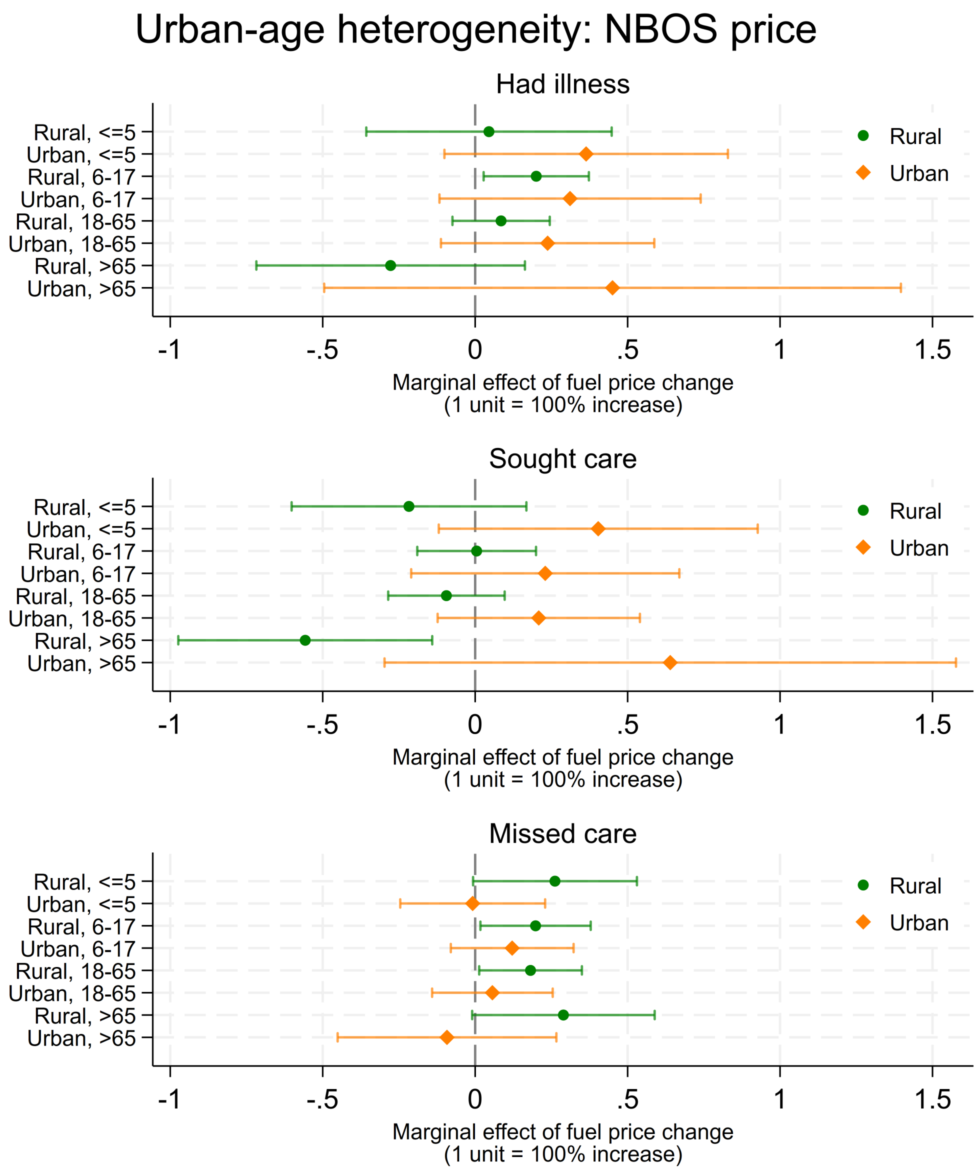

Figure 3: Urban-Age Heterogeneity in Illness and Care-Seeking

Notes: Each panel plots the marginal effect of the fuel-price change on one outcome (had illness, sought care, missed care) by urban status and age group, from a regression interacting the price change with age bucket and urban status. Green circles are rural groups and orange diamonds are urban groups; horizontal bars are 95 percent confidence intervals. One unit on the horizontal axis corresponds to a 100 percent increase in the pump price. Standard errors are clustered at the state level and regressions use wave-5 weights.

V Empirical Strategy

A Conception-Cohort Exposure Design (ANC and Birth Outcomes)

The pre-birth and at-birth results exploit variation in how much of a pregnancy falls in the post-reform period across birth cohorts. For a child born in month t, we define $\mathbf{Exposure}_i$ as the number of months in the fixed nine-month pregnancy window (the birth month and the eight preceding calendar months) that fall on or after June 2023. This measure ranges from zero (all nine months pre-reform) to nine (all nine months post-reform). We avoid relying on reported gestational duration, which may be mismeasured or itself affected by the shock.

The estimating equation is:

$$Y_{ict} = \alpha + \eta_c + \tau_{m(t)} + \kappa \,(\text{Exposure}_i \times \text{Urban}_c) + \rho\, \text{Exposure}_i + \mathbf{X}_i'\gamma + \mathbf{W}_{ic,\text{tri}}'\psi + \nu_{ict} \quad (2)$$

where Y_{ict} is an outcome for child i born in calendar month t residing in cluster c. η_c are DHS cluster fixed effects, which absorb all time-invariant spatial differences including the urban/rural classification. $\tau_{m(t)}$ are birth-month-of-year fixed effects, absorbing seasonal patterns in births and outcomes. $\mathbf{X}_i$ includes maternal age and its square, maternal education, child sex, birth order, and a smoking indicator. $\mathbf{W}_{ic,\text{tri}}$ includes trimester-specific log precipitation and maximum temperature. Standard errors are clustered at the state-by-urban level (37 states $\times$ 2 strata = 74 clusters).

The coefficient ρ captures the average effect of an additional post-reform month on the outcome in rural clusters. The coefficient of primary interest is κ, the urban-rural differential: a negative κ on adverse outcomes indicates that urban areas deteriorated less than rural areas. Identification of κ comes from differences in cohort exposure within the same cluster, purged of birth-month seasonality, comparing urban and rural locations.

We also estimate trimester-specific versions by replacing Exposure_i with the number of post-reform months in the first, second, or third trimester separately. The identifying assumption is that, absent the reform, urban and rural birth outcomes would have followed parallel trends across cohorts. We assess this using event-study specifications that replace the single Exposure measure with a full set of birth-month-relative-to-reform dummies; flat pre-reform coefficients are the main diagnostic.

The ambient air pollution analysis uses a cluster-month level difference-in-differences:

$$\ln(\text{NonDustPM}_{2.5,ct}) = \alpha + \beta(\text{Urban}_c \times \text{Post}_t) + \phi\,\text{Post}_t + \lambda_c + \mu_{m(t)} + \mathbf{X}'_{ct}\delta + \varepsilon_{ct} \quad (3)$$

where Post_t equals one for June 2023 and later. The coefficient β captures the additional post-removal change in non-dust $\text{PM}_{2.5}$ in urban clusters relative to rural clusters. A placebo using dust $\text{PM}_{2.5}$ (dominated by Saharan transport unrelated to fuel combustion) tests whether the differential change is combustion-specific.

B Within-Child and Across-Child Within-Vaccine Designs (Vaccination)

The vaccination analysis uses two complementary identification designs. Each exploits the structure of the EPI schedule: vaccines are assigned to specific age windows, and whether a given child's age-eligible window for a specific vaccine falls before or after the reform depends on the child's birth date. Both designs are free of the survivorship selection concern that applies to birth outcomes.

Across-child within-vaccine design. The first design compares children of different birth cohorts for the same vaccine. Let D_{iv} equal one if child i received vaccine v. The birth cohort determines whether the vaccine window falls pre- or post-reform: for vaccine v with recommended administration age a_v weeks, the window falls in the post-reform period if the child's birth month b_i satisfies $b_i + a_v/4 \geq$ June 2023. Let $\text{Post}_{iv} = \mathbf{1}\{b_i + a_v/4 \geq \text{June 2023}\}$

be an indicator for the child's window for vaccine v being post-reform. The estimating equation is:

$$D_{iv} = \alpha_v + \eta_c + \tau_{m(b_i)} + \beta_v \, \text{Post}_{iv} + \delta_v \, (\text{Post}_{iv} \times \text{Rural}_c) + \mathbf{X}_i' \gamma + \varepsilon_{iv} \tag{4}$$

estimated separately for each vaccine v. The coefficient β_v captures the average change in coverage for vaccine v when the window falls post-reform, and δ_v is the differential rural effect. Cluster and birth-month-of-year fixed effects absorb time-invariant location differences and seasonal patterns.

Within-child across-vaccine design. The second design exploits the fact that a single child faces different post-reform exposure across vaccines, because vaccines are scheduled at different ages. For a child born in, say, April 2023, the BCG window (birth) is pre-reform, the six-week vaccines window is just at the reform boundary, and all subsequent windows are post-reform. For a child born in December 2022, all vaccine windows except the nine-month vaccines are pre-reform.

Let PostWindow_{iv} be a continuous measure of how much of vaccine v's recommended age window falls in the post-reform period for child i. The within-child fixed-effects specification is:

$$D_{iv} = \mu_i + \alpha_v + \beta \, \text{PostWindow}_{iv} + \delta \, (\text{PostWindow}_{iv} \times \text{Rural}_c) + \varepsilon_{iv} \tag{5}$$

where μ_i are child fixed effects and α_v are vaccine fixed effects. The child fixed effects absorb all child-level unobservables that affect overall vaccine uptake — maternal health-seeking behavior, household income, distance to facility, and any other factor that raises or lowers all vaccines uniformly for a given child. The coefficient β therefore identifies the within-child, across-vaccine effect: for the same child, vaccines whose recommended window falls more heavily in the post-reform period are differentially less likely to be received. The coefficient δ captures whether this within-child gradient is steeper in rural clusters.

This design has strong internal validity properties. By controlling for the child fixed effect, it rules out selection on child-level observables and unobservables as explanations for the coverage differential. The identifying variation is the interaction between a child's birth month and a vaccine's position in the EPI schedule — variation that is determined by biology (the schedule) and timing of birth relative to the reform, neither of which is manipulable by parents. The key identifying assumption is that conditional on child and vaccine fixed effects, the within-child, across-vaccine variation in post-reform exposure is as good as random.

Right-censoring and the on-time outcome. A data-structure challenge specific to the vaccination analysis is right-censoring in the received outcome. The survey records whether a child has *ever* received each vaccine as of the interview date, but post-reform cohorts are on average younger at interview: a child born in late 2023 may not yet have reached the recommended age for nine-month vaccines by the time the survey visits, mechanically suppressing received rates for the youngest cohorts regardless of any behavioral change. This form of right-censoring would bias toward finding negative effects on later-dose vaccines even absent any true access disruption.

We address this through a second outcome: an indicator for whether each vaccine was administered *within* the public-health recommended age window, constructed from card-recorded vaccination dates ("on-time" receipt). Specifically, we follow the timeliness windows adopted by Newcomer et al. (2024), which define vaccine- specific acceptable receipt windows based on clinical immunization guidelines and prior vaccination-timeliness studies. These windows allow a short grace period after the recommended age while still distinguishing timely vaccination from delayed catch-up receipt, reflecting the epidemiological goal of protecting children during the ages when they are most vulnerable. This outcome captures whether families kept up with the schedule in real time and is therefore less affected by the mechanical censoring that affects the ever-received measure. Comparing the received and on-time estimates directly tests whether the dose-gradient in received outcomes merely reflects age composition of post-reform birth cohorts or reflects a genuine deterioration in timely care-seeking.

The trade-off is sample selection as highlighted in Table 1. Card-holding households represent approximately 53 percent of the eligible sample and are systematically more engaged with the health system than card-absent households. Timeliness estimates therefore represent a lower bound on the true impact of the reform on on-schedule vaccination. We present both outcomes throughout and interpret persistence of the dose-gradient in the on-time measure as confirmation that the pattern is not a censoring artifact.

C Identification Threats and Validation

Fertility selection. One concern is that the reform changed the composition of births rather than health outcomes conditional on birth. This would be possible if households adjusted fertility timing around the reform, either by bringing pregnancies forward in anticipation of a negative shock or by postponing conceptions once the shock occurred. Such behavior would generate a discontinuity or bunching in the distribution of children's birth dates around the subsidy-removal date. Figure 4 provides little evidence of such sorting. It plots the cumulative distribution of children's dates of birth relative to the reform date, alongside the CDF implied by a uniform distribution over the same window. The empirical CDF is smooth around the reform threshold and closely tracks the uniform benchmark both before and after the policy date. This pattern suggests that births are not sharply concentrated or missing around the reform, reducing concerns that the main estimates are driven by selective fertility timing. In addition, we examine whether observable maternal characteristics, including age, education, and parity, change discontinuously across the reform threshold.

Selective migration. DHS cluster location is used to assign pollution exposure and urban/rural classification. If mothers selectively moved from rural to urban areas in response to the reform, the rural sample would be negatively selected. The sample restriction to

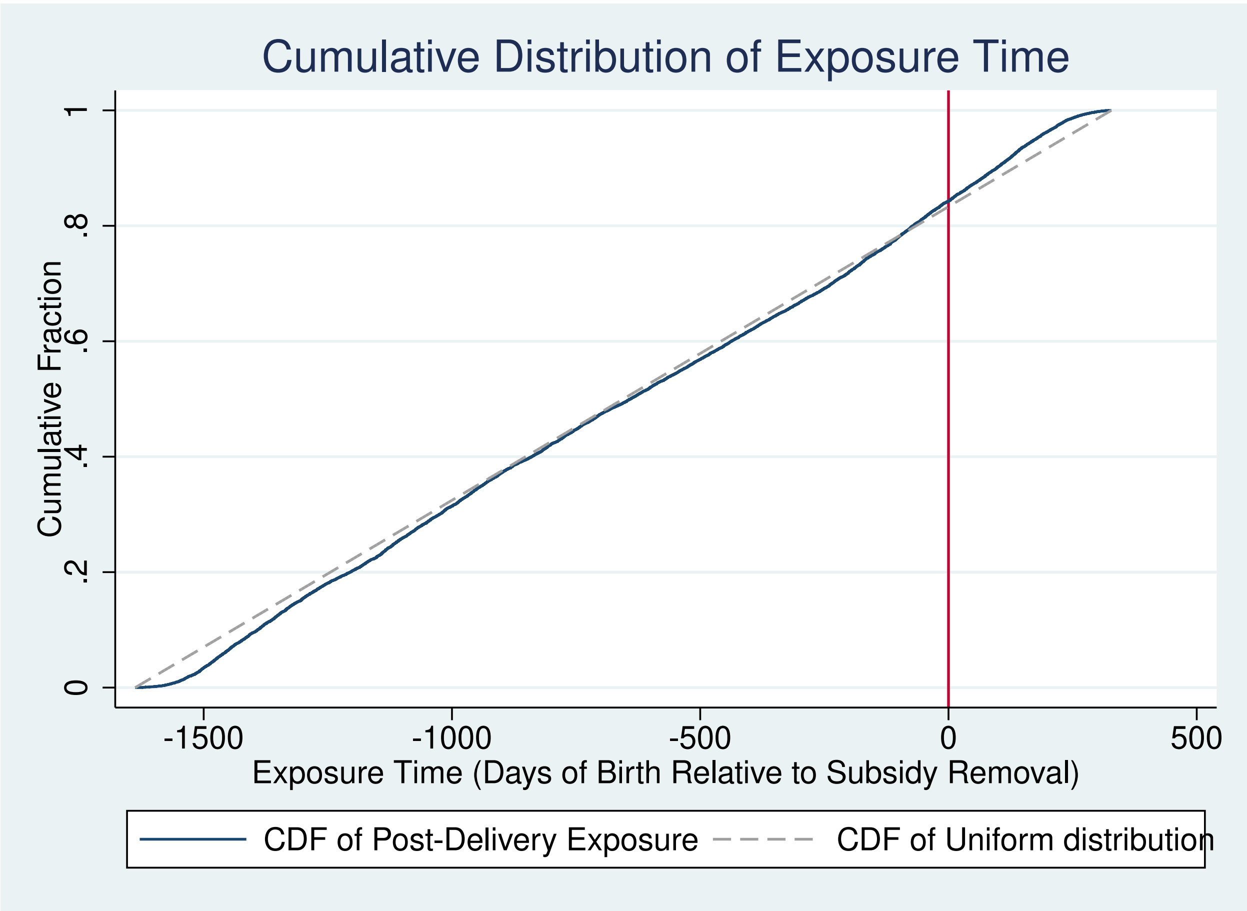

Figure 4: Cumulative Distribution of Children's Birth Dates Relative to the Fuel Subsidy Removal
Notes: This figure plots the cumulative distribution function of children's dates of birth relative to Nigeria's fuel subsidy removal. The x-axis measures the number of days between the child's date of birth and the subsidy removal date, with negative values indicating births before the reform and positive values indicating births after the reform. The solid blue line reports the empirical CDF of birth timing in the sample. The dashed gray 45-degree line reports the CDF that would arise under a uniform distribution of birth dates over the same window. The vertical red line marks the subsidy removal date. The close alignment between the empirical CDF and the uniform benchmark indicates that births are approximately evenly distributed over the sample window.

non-movers and the $[-8, +6]$ month window limit this concern. Selective migration at the scale needed to affect results is implausible in a short horizon.

Survivorship selection in birth outcomes. Birth size is observed only for surviving births; differential neonatal mortality could bias the birth-size estimates if post-reform cohorts had different survival rates in ways correlated with birth size. The neonatal mortality results show no significant effect of the reform on survival (Section VII). Lee (2008) trimming bounds for the birth-size results are in progress and will be reported in the next draft. We disclose this limitation here.

COVID-19 pre-trends. The January 2022 lower bound is designed to ensure that the pre-reform comparison period is free of COVID-era healthcare disruption. Differential urban-rural COVID recovery could nonetheless introduce heterogeneous pre-reform trends that contaminate the identification. Event studies with the 2022 pre-period are the main diagnostic; flat pre-trends in this period address the concern. We additionally note that any urban-rural recovery differential would tend to produce an apparent post-reform urban improvement independent of the subsidy shock, biasing the estimated urban-rural differential. The robustness results using the bridging-share IV, which does not exploit urban-rural variation, address this concern.

Anticipation effects. A potential threat to identification is that the reform was not fully unanticipated. President Tinubu's May 29 inaugural address was the first public announcement, but some observers speculated about subsidy removal during the campaign period (Dzirutwe 2022; Onyeiwu 2023; TheGuardian 2023). If households anticipated the reform and accelerated health-care visits before it took effect, the pre-reform comparison group would be positively selected, attenuating the estimated post-reform deterioration and biasing toward zero.

Three features of the setting limit this concern. First, the reform announcement and implementation occurred simultaneously: the inaugural declaration and the NNPC price adjustment on May 30 left no interval between announcement and enactment in which households could respond.

Second, and more directly, the price increase households actually faced was not predictable from available public information. We document this using the Nigeria Economic Summit Group (NESG) Tax & Subsidy Perception Survey, collected in July–August 2018 among a nationally representative sample of 16,000 working-age adults. Respondents were informed that the prevailing retail petrol price was ₦185 per liter and asked to estimate what the price would be if the subsidy were removed. We compute the respondent-predicted percentage price change and compare it to the actual post-removal percentage change in the NBOS-recorded state-level pump price (post-reform average minus pre-reform average, divided by the pre-reform average).

Figure 5 shows the state-level relationship between the median predicted percentage change and the realized change. The two series are near-orthogonal: the correlation is essentially zero, and a substantial share of respondents expected the post-removal price to fall *below* the prevailing ₦185 level, citing corruption and institutional failures in petroleum pricing as their rationale. States that experienced the largest realized price increases were no more likely to have anticipated those increases than states with moderate changes. This implies that any anticipatory behavioral adjustment would have been uncorrelated with the cross-state variation in realized fuel price pass-through that is central to the identification. The anticipation threat therefore cannot generate the spatial gradient in health outcomes that we document.

Third, the supply-chain data also show little evidence of anticipatory adjustment before the May 29 removal. Figure 6 plots national daily PMS truck-out volumes using seven-day moving averages, comparing the 2023 policy-year series with the corresponding 2022 series. If marketers, depots, or retailers had anticipated an imminent price increase or shortage, one might expect a pre-removal surge in truck-out volumes as downstream

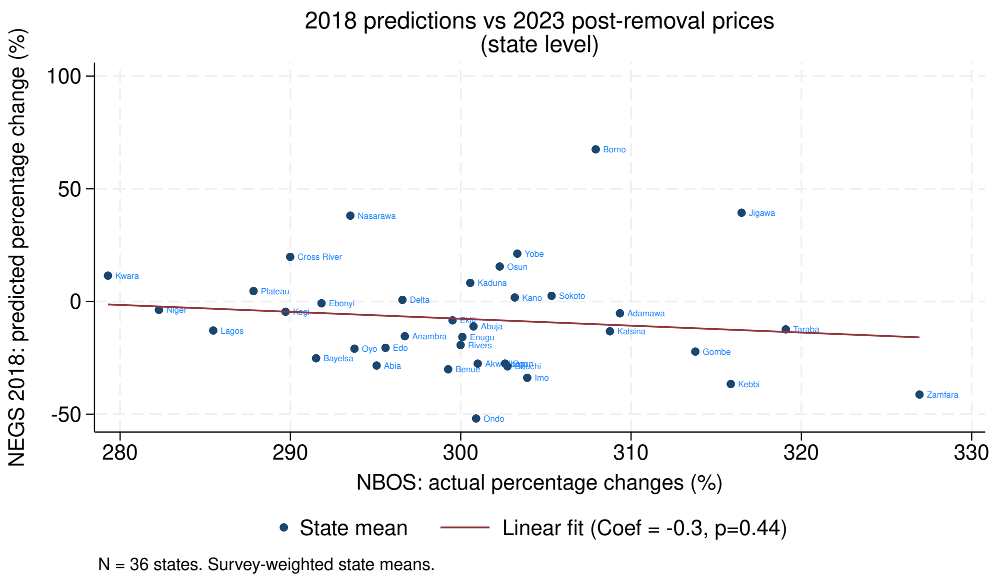

Figure 5: NESG Predicted vs. Realized Fuel Price Change by State

Notes: Each point is a Nigerian state. The horizontal axis shows the median respondent-predicted percentage change in petrol price from the subsidized ₦185 per liter baseline, computed from the 2018 NESG Tax & Subsidy Perception Survey ($n \approx 16{,}000$ working-age adults). The vertical axis shows the realized percentage change in the NBOS-recorded state-level pump price (post-reform average minus pre-reform average divided by the pre-reform average). The near-zero correlation indicates that households had no ability to predict cross-state variation in the price shock, ruling out spatially correlated anticipatory behavior as a driver of the estimated health effects.

actors accumulated inventories. Conversely, if supply-chain actors expected deregulated prices to sharply reduce demand, one might expect a pre-removal contraction in dispatches to avoid holding costly inventories. The figure shows neither pattern. National truck-out volumes in the weeks before May 29 evolve smoothly and remain broadly parallel to the 2022 comparison series, with no discrete spike or collapse before the reform date. This suggests that the downstream PMS supply chain continued operating largely as normal before the announcement. Together with the simultaneous announcement and implementation of the reform, and the weak relationship between predicted and realized price changes, this evidence indicates that anticipation effects were likely limited and unlikely to generate the spatial gradient in health outcomes that we document.

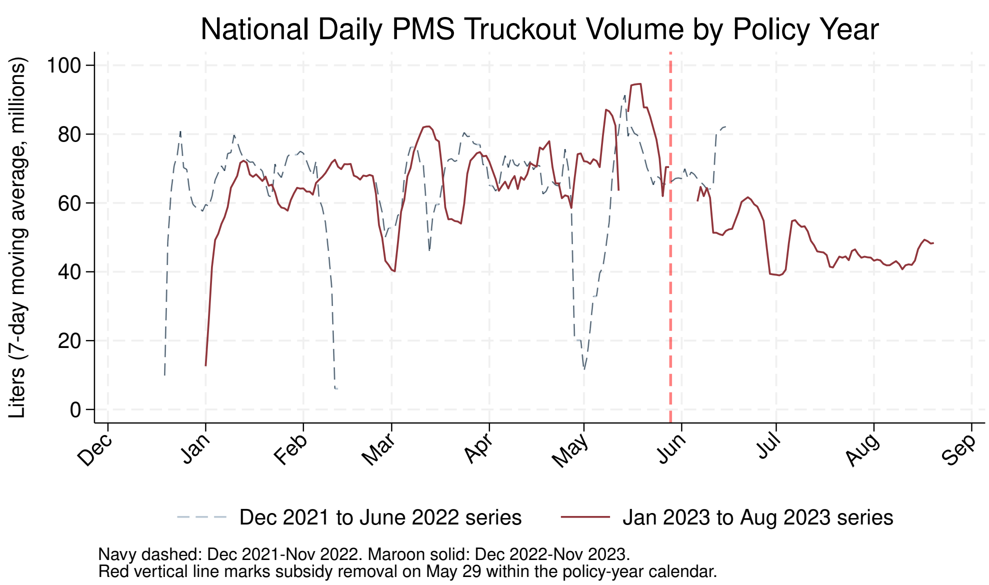

Figure 6: National Truck-out volumes from depot (7-day moving average)

Notes: Each point is a Nigerian state. The horizontal axis shows the median respondent-predicted percentage change in petrol price from the subsidized ₦185 per liter baseline, computed from the 2018 NESG Tax & Subsidy Perception Survey ($n \approx 16{,}000$ working-age adults). The vertical axis shows the realized percentage change in the NBOS-recorded state-level pump price (post-reform average minus pre-reform average divided by the pre-reform average). The near-zero correlation indicates that households had no ability to predict cross-state variation in the price shock, ruling out spatially correlated anticipatory behavior as a driver of the estimated health effects.

Multiple Hypothesis Testing. We also address concerns about multiple hypothesis testing. The paper examines several related outcome families across the early-life health sequence, including prenatal healthcare utilization, birth outcomes, and postnatal immunization coverage and timeliness. The concern is particularly salient for the vaccination analysis, which estimates effects separately for multiple vaccine doses and for both receipt and on-time receipt. In the next iteration of this draft, we will report sharpened q-values that control the false discovery rate within pre-specified outcome families. A q-value is the multiple-testing-adjusted analogue of a p-value: it gives the minimum false discovery rate at which a null hypothesis would be rejected. False-discovery-rate control is well suited to this setting because the outcomes within each family are related and the inferential goal is to limit the expected share of false discoveries among rejected hypotheses, rather than to control the probability of any false rejection across all outcomes in the paper.

VI Pre-birth Effects: Antenatal Care

We begin the results with the pre-birth stage. Antenatal care visits require repeated facility contact during pregnancy; they are therefore directly in the path of the transportation-cost mechanism. A reduction in ANC visits reduces the screening and monitoring that catch maternal complications before delivery, while degradation in ANC content quality reduces the clinical value of each visit that is made. The access-cost channel predicts that both margins will be affected, and that the proportional impact will be larger in rural areas where the baseline visit frequency is lower and access more constrained.

A ANC Utilization

Table 4 reports estimates from equation (2) for three ANC utilization outcomes: number of antenatal visits, early ANC initiation, and receipt of skilled ANC. Panel A focuses on visit quantity. Under the full-pregnancy exposure measure, each additional month of post-reform exposure reduces the number of antenatal visits by 0.071 on average — a 1.7 percent reduction relative to the sample mean of 4.2 visits. The reduction is present and similar in magnitude when the urban-rural interaction is added: the rural effect is -0.066 visits per exposure month, and the Urban $\times$ Exposure coefficient of -0.016 is small and statistically insignificant. The lack of a statistically significant urban-rural differential should not be interpreted as evidence of equal impact. The rural baseline of 3.15 ANC visits implies a proportional decline of 2.1 percent per additional exposure month, versus 1.0 percent for urban women (baseline 6.49 visits). The same absolute disruption to visit frequency represents a larger proportional reduction in the already-constrained rural setting.

For early ANC initiation and skilled ANC, we find little evidence of change. These are extensive-margin outcomes that capture whether women enter the ANC system at all, rather than how many visits they make within the system. The reform appears to have compressed visit frequency among women who continued to attend ANC, without substantially driving women out of the system entirely.

B ANC Content Quality

Table 5 reports estimates for five dimensions of ANC clinical content. Each additional month of full-pregnancy exposure reduces the probability of receiving a urine test by 0.9 percentage points, a blood test by 0.8 percentage points, food counseling by 1.0 percentage point, and bleeding danger-sign counseling by 1.0 percentage point. Blood pressure measurement shows a similar negative direction. The Urban $\times$ Exposure coefficients are small and statistically insignificant throughout, providing little evidence of a statistically significant differential urban-rural response in content quality. As with visit quantity, the same absolute declines represent larger proportional reductions in rural areas because baseline rates are lower.

The clinical relevance of these content measures warrants emphasis. Blood and urine tests during ANC are the primary means of detecting maternal anemia, urinary tract infections, and proteinuria — all conditions associated with worse birth outcomes (Figueiredo et al. 2018; He, Tikellis, et al. 2024). Food counseling may improve maternal diet and birth weight (Dewidar et al. 2023). Bleeding danger-sign counseling is particularly important for rural women who may not have timely access to emergency obstetric care if complications arise. The decline in these content measures therefore has direct implications for birth outcomes and the results in Section VII.

Table 4: Effects of Fuel Subsidy Removal on ANC Utilization

	(1)	(2)	(3)	(4)	(5)	(6)
	# ANC visits		1st ANC on time		Skilled ANC	
Exposure	−0.071***	−0.066***	−0.004	−0.005	−0.005*	−0.004
	(0.022)	(0.025)	(0.003)	(0.004)	(0.003)	(0.004)
Urban × Exposure	—	−0.016	—	0.001	—	−0.004
		(0.039)		(0.005)		(0.005)
Obs. / Mean	6,495 / 4.148	6,495 / 4.148	6,675 / 0.184	6,675 / 0.184	6,683 / 0.623	6,683 / 0.623
Birth month FE	Yes	Yes	Yes	Yes	Yes	Yes
Cluster FE	Yes	Yes	Yes	Yes	Yes	Yes
Weather	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Exposure is the number of pregnancy months falling on or after June 2023 in a fixed nine-month window. Odd columns report the average exposure effect; even columns add the Urban × Exposure interaction. Standard errors clustered at the state-by-urban level in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 5: Effects of Fuel Subsidy Removal on ANC Content Quality

	Urine test		Blood test		Food counsel.		BP check		Bleeding counsel.	
Exposure	−0.009***	−0.009***	−0.008**	−0.008**	−0.010***	−0.010***	−0.007**	−0.007**	−0.010***	−0.010***
Urban × Exposure	—	0.001	—	0.001	—	0.000	—	0.001	—	0.001
Birth month FE	Yes (all columns)									
Cluster FE	Yes (all columns)									

Notes: Exposure is the number of pregnancy months falling on or after June 2023 in a fixed nine-month window. Odd columns report the average exposure effect; even columns add the Urban × Exposure interaction. Standard errors clustered at the state-by-urban level in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

C Heterogeneity

The ANC effects are largely homogeneous across education and wealth subgroups. This is consistent with the access-cost mechanism: the reform increased transportation costs for all women making ANC visits, and the absolute reductions in visits and content quality are similar across groups. The larger proportional impact on rural women with lower baseline rates is the primary distributional concern. Education and wealth splits do not reveal differential patterns in ANC response, in contrast to the birth-size results reported in Section VII, where resource-heterogeneity is more visible. This may reflect the fact that ANC visit reductions are relatively uniform across the income distribution, while birth outcomes involve additional channels (pollution, capacity to absorb cost shocks) that are more strongly mediated by resources.

VII At-birth Effects: Birth Outcomes

The at-birth stage combines the downstream consequences of reduced antenatal care with a countervailing channel operating in the opposite spatial direction: the reform also reduced combustion-related air pollution, and this benefit was larger in urban areas where traffic and generator density are higher. The birth-outcome results therefore reflect two simultaneous forces — deteriorating access and improving air quality — that happen to partially offset each other in urban areas.

A Delivery Care

We begin with delivery care because it is the first at-birth outcome and captures whether the pre-birth deterioration in antenatal care translated into changes in the conditions under which births occurred. Table 6 reports effects on home delivery, cesarean delivery, and skilled delivery assistance. Across all three outcomes, the average effects under the full-pregnancy exposure measure are small and statistically insignificant. The estimates therefore provide little evidence that the reform changed the place of delivery or the likelihood of skilled attendance at birth.

The interacted specifications also show limited evidence of differential urban-rural responses in delivery conditions. Home delivery and skilled delivery assistance have small and statistically insignificant Urban × Exposure coefficients. The one exception is cesarean delivery, where the Urban × Exposure coefficient is positive and marginally significant, suggesting a small relative increase in cesarean delivery among urban births. Overall, however, the delivery-care results indicate that the reform's at-birth effects do not primarily operate through changes in delivery location or skilled attendance. This contrasts with the clearer evidence of reduced antenatal care before birth and helps motivate the focus on neonatal survival and reported birth size below. Appendix Tables C7–C8 show similar patterns when exposure is defined by trimester or by post-delivery timing.

Table 6: Effects of Fuel Subsidy Removal on Delivery Healthcare Utilization

	(1)	(2)	(3)	(4)	(5)	(6)
	Home delivery		C-section		Skilled delivery assistance	
Exposure	0.002	0.002	0.000	−0.002	−0.001	−0.000
	(0.003)	(0.004)	(0.002)	(0.002)	(0.002)	(0.002)
Urban × Exposure	—	0.001	—	0.005*	—	−0.003
		(0.005)		(0.003)		(0.004)
Observations	6,730	6,730	6,730	6,730	6,730	6,730
Mean	0.558	0.558	0.052	0.052	0.460	0.460
Birth month FE	Yes	Yes	Yes	Yes	Yes	Yes
Cluster FE	Yes	Yes	Yes	Yes	Yes	Yes
Weather controls	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Exposure is the number of pregnancy months falling on or after June 2023. Columns (1), (3), and (5) report the average exposure effect. Columns (2), (4), and (6) include Urban × Exposure. Standard errors clustered at the state-by-urban level are reported in parentheses. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

B Neonatal Mortality

Table 7 shows little evidence that the reform affected neonatal survival. Across the full-pregnancy, second-trimester, and third-trimester specifications, both the *Exposure* and *Urban* × *Exposure* coefficients are small and statistically insignificant. The null mortality result implies that differential survival selection does not drive the birth-size patterns reported below, and that the reform affected fetal growth trajectories rather than infant survival probability. Lee (2008) bounds for the birth-size results are in progress (as noted in Section C).

C Birth Size

Table 8 reports effects on reported birth size. Panel B, the *small or below* measure most closely aligned with the 2,500-gram low-birth-weight threshold, shows the clearest rural deterioration.

Table 7: Effects of Fuel Subsidy Removal on Neonatal Mortality

	(1) Full preg.	(2) Full preg.	(3) 1st tri.	(4) 2nd tri.	(5) 3rd tri.
Outcome: Neonatal mortality					
Exposure	0.000	0.000	0.017	0.002	0.000
	(0.001)	(0.001)	(0.016)	(0.003)	(0.003)
Urban × Exposure	—	0.000	−0.038*	0.001	0.002
		(0.002)	(0.020)	(0.005)	(0.003)
Observations	6,730	6,730	6,730	6,730	6,730
Mean	0.038	0.038	0.038	0.038	0.038
Birth month FE	Yes	Yes	Yes	Yes	Yes
Cluster FE	Yes	Yes	Yes	Yes	Yes
Weather	Yes	Yes	Yes	Yes	Yes

Notes: In columns (1)–(2), exposure is the number of pregnancy months falling on or after June 2023. In columns (3)–(5), exposure is the number of exposed months in the first, second, and third trimesters, respectively. Column (1) reports the average exposure effect; columns (2)–(5) include Urban × Exposure. Standard errors clustered at the state-by-urban level. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Each additional month of full-pregnancy exposure increases the probability of being reported as small or below by 0.8 percentage points in rural areas. The Urban × Exposure coefficient is −0.7 percentage points and marginally significant, indicating that rural deterioration outpaced urban deterioration. Second- and third-trimester effects are larger in magnitude, at 1.8 and 1.7 percentage points for rural areas, consistent with late-gestation sensitivity to economic and environmental conditions (Painter et al. 2005; Currie, Neidell, et al. 2009; Almond, Hoynes, et al. 2011).

Panel A, *very small*, shows a positive rural effect (0.4 percentage points per exposure month) without statistically significant urban-rural heterogeneity. This suggests the reform increased extreme adverse birth-size reports in rural areas without a detectable differential spatial pattern at the tails.

Panel C, *average or below*, reveals the most striking spatial asymmetry. The average rural effect is small and insignificant, but Urban × Exposure is −1.2 percentage points and significant at 10 percent in the full-pregnancy specification, with stronger second- and third-trimester estimates of −2.6 and −1.8 percentage points. A negative coefficient on

Urban × Exposure for this broad measure indicates that urban births shifted toward *larger* reported size categories as post-reform exposure increased. This is the opposite direction from what the access-cost mechanism alone would predict, and it motivates the pollution analysis below.

Table 8: Effects of Fuel Subsidy Removal on Reported Birth Size

	(1) Full preg.	(2) Full preg.	(3) 1st tri.	(4) 2nd tri.	(5) 3rd tri.
Panel A: Very small					
Exposure	0.003***	0.004**	0.009	0.009**	0.008***
	(0.001)	(0.001)	(0.011)	(0.003)	(0.003)
Urban × Exposure	—	−0.001	0.008	0.000	−0.004
		(0.002)	(0.025)	(0.005)	(0.003)
Obs. / Mean	6,579 / 0.041	6,579 / 0.041	6,579 / 0.041	6,579 / 0.041	6,579 / 0.041
Panel B: Small or below					
Exposure	0.006**	0.008***	0.020	0.018***	0.017***
	(0.003)	(0.003)	(0.032)	(0.005)	(0.005)
Urban × Exposure	—	−0.007*	−0.032	−0.012	−0.012**
		(0.004)	(0.043)	(0.008)	(0.006)
Obs. / Mean	6,579 / 0.127	6,579 / 0.127	6,579 / 0.127	6,579 / 0.127	6,579 / 0.127
Panel C: Average or below					
Exposure	0.001	0.005	0.001	0.009	0.011
	(0.004)	(0.005)	(0.046)	(0.010)	(0.009)
Urban × Exposure	—	−0.012*	−0.093	−0.026*	−0.018*
		(0.006)	(0.061)	(0.015)	(0.010)
Obs. / Mean	6,579 / 0.746	6,579 / 0.746	6,579 / 0.746	6,579 / 0.746	6,579 / 0.746
Birth month FE	Yes	Yes	Yes	Yes	Yes
Cluster FE	Yes	Yes	Yes	Yes	Yes
Weather	Yes	Yes	Yes	Yes	Yes

Notes: Exposure definitions as in Table 7. Standard errors clustered at the state-by-urban level. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

D Air Pollution as a Countervailing Channel

The Panel C asymmetry — urban births improving while rural births stagnate — is explained by the differential urban-rural decline in combustion-related $PM_{2.5}$. Higher fuel prices reduce road traffic and generator operation, both major sources of $PM_{2.5}$ in Nigerian cities (Karagulian et al. 2015; Okedere et al. 2021). Prior work shows that *in utero* $PM_{2.5}$ exposure impairs fetal growth (Brauer et al. 2008; Slama et al. 2008). If the reform reduced $PM_{2.5}$

more in urban areas, it provides a birth-size benefit that partially offsets the access-cost harm.

Table 9 confirms this. Non-dust $PM_{2.5}$ fell by 18.3 percent on average after June 2023, with an additional decline of 4.9 percent (-0.70 $\mu g/m^3$) in urban areas relative to rural areas. By contrast, dust $PM_{2.5}$ shows no differential urban-rural response — the Urban × Post coefficient is small and statistically insignificant — confirming that the differential change is specific to combustion-related pollution rather than Saharan dust transport. The policy reversal provides a further test: Table 10 shows that the urban-rural $PM_{2.5}$ gap widens after June 2023 but then narrows as the implicit subsidy partially re-emerges in late 2023, consistent with combustion-related rather than coincidental dynamics.

Table 9: Effects of Fuel Subsidy Removal on $PM_{2.5}$

	(1)	(2)	(3)	(4)
Panel A: Non-dust $PM_{2.5}$				
	ln(Non-dust $PM_{2.5}$)		Non-dust $PM_{2.5}$	
Post	-0.183^{***}	-0.157^{***}	-3.167^{***}	-2.796^{***}
	(0.015)	(0.016)	(0.219)	(0.213)
Urban × Post	—	-0.049^{**}	—	-0.695^{**}
		(0.024)		(0.282)
Obs. / Mean	20,700 / 2.359	20,700 / 2.359	20,700 / 11.805	20,700 / 11.805
Panel B: Dust $PM_{2.5}$ (placebo)				
	ln(Dust $PM_{2.5}$)		Dust $PM_{2.5}$	
Post	-0.142^{***}	-0.109^{***}	-3.717^{***}	-3.891^{***}
	(0.013)	(0.024)	(0.994)	(1.151)
Urban × Post	—	-0.062	—	0.324
		(0.047)		(0.713)
Obs. / Mean	20,700 / 3.424	20,700 / 3.424	20,700 / 47.964	20,700 / 47.964
Cluster FE	Yes	Yes	Yes	Yes
Month-of-year FE	Yes	Yes	Yes	Yes
Weather	Yes	Yes	Yes	Yes

Notes: Post is an indicator equal to one for months on or after June 2023. All models include DHS cluster FE, month-of-year FE, and weather controls. Standard errors clustered at the state-by-urban level. ** $p < 0.05$; *** $p < 0.01$.

The birth-outcome evidence is therefore consistent with two offsetting forces. The access-cost mechanism — higher transportation costs reducing ANC visits and content quality, compressing nutrition and monitoring — worsens fetal growth conditions everywhere. The

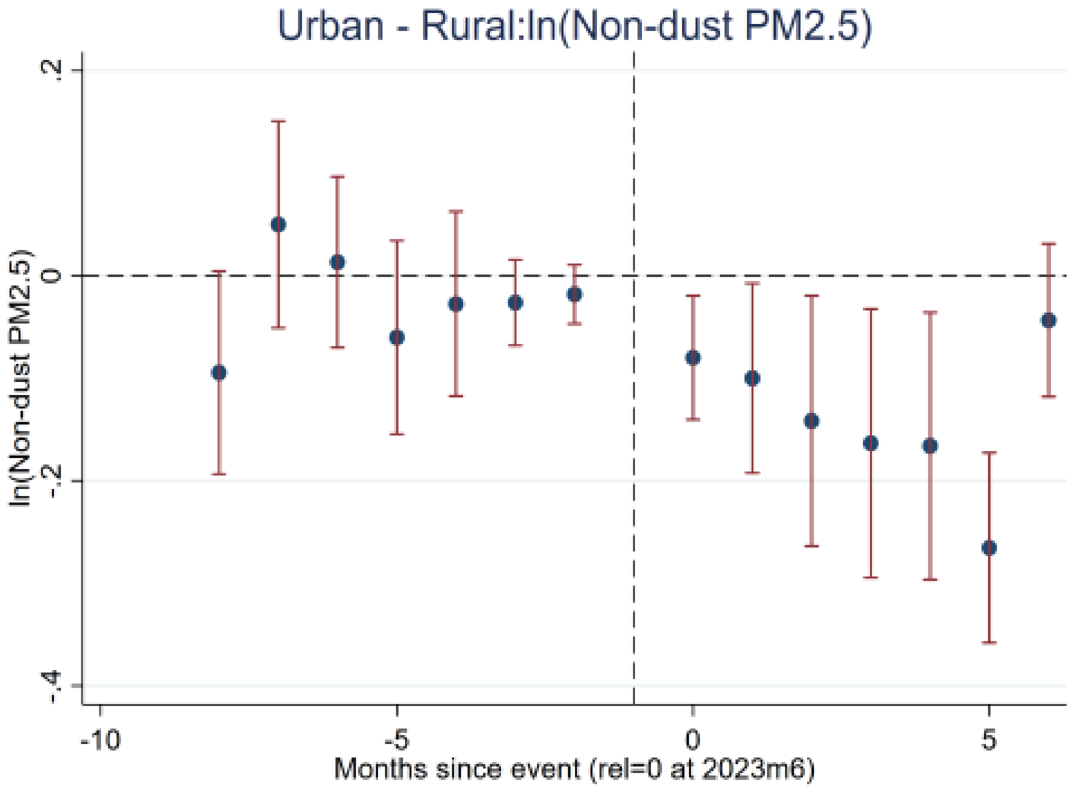

Figure 7: Event Study: Non-Dust $PM_{2.5}$ by Urban Status
Notes: Coefficients on Urban × month dummies from equation (3), plotted relative to May 2023 (omitted). Vertical dashed line indicates June 2023. 95% confidence intervals shown.

Table 10: Effects of Fuel Subsidy Removal on Non-Dust $PM_{2.5}$: Extended Window with Policy Reversal

	(1)	(2)
	ln(Non-dust $PM_{2.5}$)	
Post = 1	-0.126^{***}	-0.104^{***}
	(0.010)	(0.014)
Urban × Reform = 2 (subsidy-removal period)	—	-0.042^{*}
		(0.022)
Urban × Reform = 3 (partial reinstatement)	—	0.065^{***}
		(0.022)
Obs. / Mean	30,228 / 2.377	30,228 / 2.377
Cluster FE	Yes	Yes
Month-of-year FE	Yes	Yes
Weather	Yes	Yes

Notes: Estimation window extended through August 2024. Reform = 2 denotes the initial subsidy-removal period (June–October 2023); Reform = 3 denotes the partial reinstatement period (November 2023 onward). The positive Urban × Reform = 3 coefficient indicates convergence of urban and rural $PM_{2.5}$ as the implicit subsidy re-emerges. Standard errors clustered at the state-by-urban level. * $p < 0.10$; *** $p < 0.01$.

pollution benefit mechanism — reduced fuel combustion lowering $PM_{2.5}$ — partially offsets the harm in urban areas where baseline combustion pollution is higher. The result is that rural births show clear lower-tail deterioration (Panel B), while urban births do not — and at the broad Panel C measure show an improvement consistent with the net pollution benefit.

E Heterogeneity by Education and Wealth

Table 11 reports split-sample estimates by maternal education and household wealth. The urban relative improvement (negative Urban × Exposure for *average or below*) is concentrated among higher-education and richer households. Among higher-education women, the urban-rural differential for *average or below* is −1.3 percentage points per exposure month. Among lower-education women, the estimate is smaller and statistically insignificant. Among richer households, the corresponding estimate is −2.7 percentage points, versus −1.3 percentage points among non-richer households.

This pattern is consistent with the interpretation that higher-resource urban households were better positioned to benefit from the pollution reduction while absorbing the access-cost increase. Higher-income households can substitute toward alternative transportation, smooth consumption using savings, and pay for higher-quality antenatal care. Lower-resource urban households face the same access-cost increase but may not capture the pollution benefit fully if they cannot avoid the cost shock. Rural households at all resource levels face both the access-cost increase and a smaller pollution benefit.

Table 11: Heterogeneity by Maternal Education and Household Wealth

	(1) Very small	(2) Small or below	(3) Average or below
Panel A: By maternal education			
Urban × Exposure: Low education	0.001	−0.006	−0.006
	(0.005)	(0.008)	(0.010)
Obs.	2,847	2,847	2,847
Urban × Exposure: High education	−0.003	−0.007	−0.013*
	(0.002)	(0.005)	(0.007)
Obs.	3,492	3,492	3,492
Panel B: By household wealth			
Urban × Exposure: Not richer	−0.001	−0.009	−0.013
	(0.004)	(0.007)	(0.009)
Obs.	4,319	4,319	4,319
Urban × Exposure: Richer	−0.005	−0.010	−0.027**
	(0.004)	(0.007)	(0.013)
Obs.	1,972	1,972	1,972
Birth month FE	Yes	Yes	Yes
Cluster FE	Yes	Yes	Yes
Weather	Yes	Yes	Yes

Notes: Full-pregnancy exposure specification estimated within indicated subgroups. Wealth quintile is measured post-reform; interpret with caution. Standard errors are clustered at the state-by-urban level. * $p < 0.10$; ** $p < 0.05$.

VIII Post-birth Effects: Vaccination Coverage and Timeliness

The post-birth stage extends the access-disruption mechanism from antenatal care to childhood vaccination. Each dose in the EPI primary series requires an independent clinic visit or outreach encounter. A persistent increase in transportation costs therefore imposes a barrier at each contact point, and the barriers compound: a family that defers or skips one visit reduces the probability of completing subsequent doses, both because later doses require the same access-cost outlay and because the relationship with the clinic may weaken. The dose-gradient prediction — coverage declining monotonically with dose order — is the distinctive empirical signature of this mechanism, in contrast to, say, a pure income-shock story (which would affect all visits proportionally) or a confidence-in-vaccines story (which would affect all vaccines regardless of dose order).

The two designs described in Section B are complementary. The across-child within-vaccine design traces how coverage for each vaccine varies across birth cohorts whose vaccine window straddles the reform. The within-child across-vaccine design asks whether, for the same child, vaccines scheduled later in life are disproportionately less likely to be received when those later windows fall in the post-reform period. Both designs yield the same qualitative conclusion, which reduces concern that the pattern is an artifact of either cross-sectional selection or the specific identification assumption.

A Age-Vaccination Profiles

Before turning to the regression results, we summarize the raw timing patterns using the vaccine-card-holder sample. For each vaccine, we construct an age-specific profile in which the point at child age a is the weighted share already vaccinated among children observed to at least age a. This object is useful because it lets the denominator adjust for the fact that post-reform cohorts are younger at interview, but it also requires careful interpretation: it is *not* a cohort cumulative distribution function. Each point compares different risk sets. For example, the post-window point at age 12 months averages over children observed to at least 12 months, whereas the post-window point at age 18 months averages over a smaller, older subset of post-window children. The lines therefore need not be monotone, and a later-age decline should not be read as children "losing" vaccination.

Figure 8 illustrates the pattern using the at-birth group and the nine-month group. Panel a shows that BCG and Polio 0 are highly similar across the pre- and post-window groups, which is reassuring because both are delivered at birth and do not require a return trip. This is the closest placebo in the vaccination schedule: for doses administered during the delivery encounter, the post-reform profiles do not exhibit a broad downward shift. HepB-birth is the exception within the at-birth group: post-window children display lower coverage throughout. We interpret this more cautiously because HepB-birth is an institution-dependent birth dose and therefore may reflect delivery-site logistics or birth-facility selection rather than the repeated-visit mechanism that motivates the later-dose regressions.

Panel b turns to the month-9 vaccines, where the access-cost mechanism should be most visible. The post-window profile is initially slightly above the pre-window profile, but at later ages it flattens and then falls below the pre-window line. That crossing is suggestive, but it is not by itself a cohort-level statement that post-reform children eventually have lower cumulative coverage. Because each later-age point uses a different subset of children, two forces can generate the decline: first, older children whose vaccine window fell after subsidy removal may indeed exhibit lower realized receipt by those later ages; second, right-censoring changes which post-window children remain observable at older ages. The plot therefore provides descriptive evidence of later-age slippage, not a literal cumulative coverage path for a fixed cohort.

This distinction also clarifies why the paper emphasizes *on-time* vaccination rather than a simple "received by interview" indicator. A binary ever-received outcome can look fairly similar across cohorts even when schedule progression has worsened. The age-profile figures show precisely where this matters: the post-window group need not start behind, but it becomes relatively weaker as the schedule extends to later contacts. That is the same mechanism captured more cleanly by the within-child on-time regressions. The corresponding week-6, week-10, and week-14 profiles are reported in Appendix Figures B2–B4; they show broadly similar, if milder, patterns of later-age flattening in the post-window group.

B Coverage Effects by Vaccine and Dose Order

Figures 9 and 10 plot coefficient estimates from equation (4) for each vaccine ordered by dose position in the schedule, alongside 95-percent confidence intervals. Each coefficient is the estimated change in coverage probability for a child whose vaccine window falls in the post-reform period, expressed relative to BCG (the at-birth placebo, whose coefficient is normalized to zero by construction). Figure 9 shows the full-sample estimates; Figure 10 isolates rural children, where the access disruption mechanism predicts larger effects.

The results confirm the dose-gradient prediction for the primary series. The at-birth vaccines (BCG, HepB0, OPV0) show near-zero coefficients, validating the placebo. Coverage shortfalls emerge at the six-week contact point and widen monotonically through the schedule, reaching approximately -8 to -10 percentage points for the nine-month vaccines (Measles and Vitamin A) relative to BCG. The rural subsample shows substantially larger gaps at every post-birth contact point, consistent with compounding transportation barriers in areas where each clinic visit requires a longer journey.

A raw comparison of pre- and post-reform ever-received rates would, however, conflate genuine access disruption with a mechanical artifact: post-reform cohorts are on average younger at the time of interview and have had less calendar time to accumulate later doses regardless of any behavioral change. This right-censoring concern would bias toward finding negative effects on later-dose vaccines even absent true disruption, and would disproportionately affect the nine-month vaccines with the longest accumulation windows. The timeliness analysis in Section D addresses this directly by shifting to a card-date-based on-time outcome

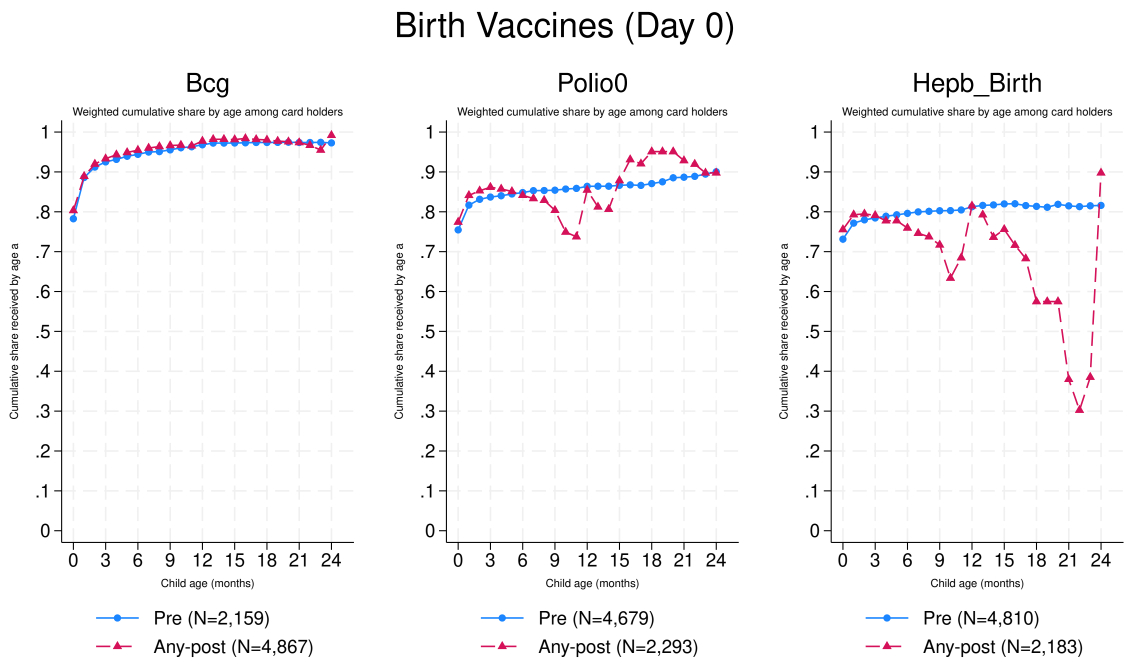

(a) At-birth vaccines

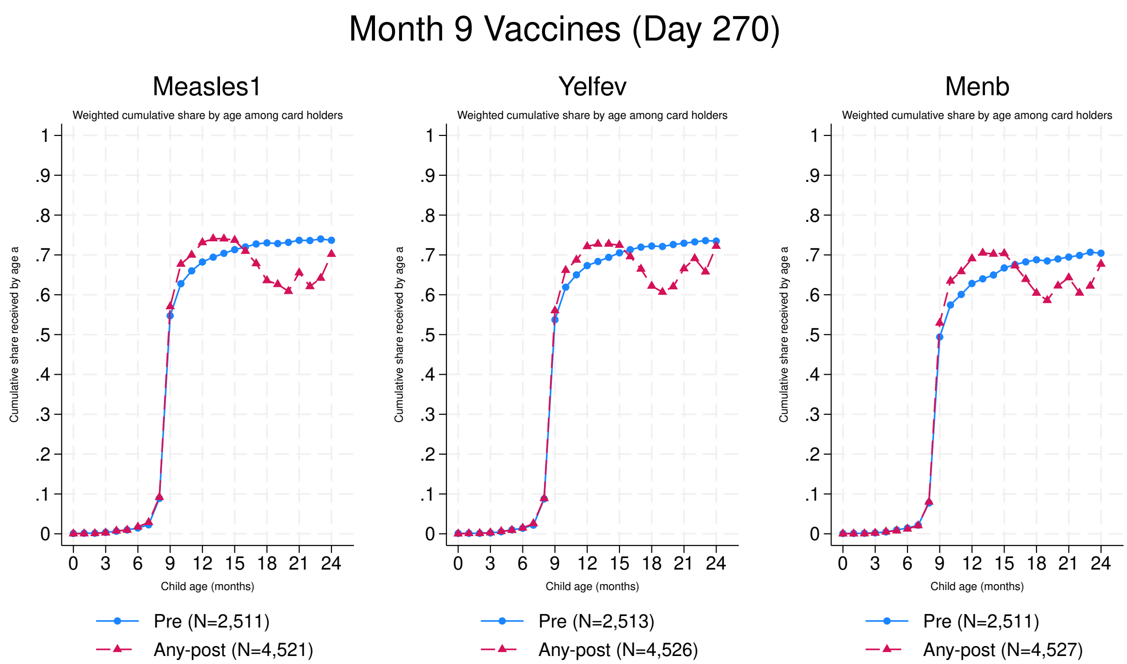

(b) Month-9 vaccines

Figure 8: Age-Specific Vaccination Profiles by Pre- and Post-Window Groups

Notes: Each panel plots the weighted share vaccinated by age a among vaccine-card-holder children observed to at least age a, separately for children whose vaccine-specific recommended window falls fully before the reform and those whose window falls partly or fully after the reform. Because the denominator changes with age, these are age-specific observed-risk-set profiles rather than cohort cumulative distribution functions. Later-age dips or crossings therefore indicate weaker realized coverage among the subset observed to those ages, not that the same children lose vaccination status over time.

that is immune to censoring by construction.

The within-child fixed-effects specification (equation 5) confirms these results with stronger internal validity guarantees. For the same child, vaccines whose recommended window falls more heavily in the post-reform period are significantly less likely to be received. The rural differential is negative and significant, indicating the within-child post-reform gradient is steeper in rural areas. The BCG placebo is confirmed: restricting to at-birth vaccines yields a near-zero β, while the effect is driven entirely by vaccines at six weeks and later.

Two vaccines deviate from the monotone gradient in ways that are mechanistically informative. Rota2 and IPV show substantially smaller coverage shortfalls — approximately +16 percentage points relative to co-scheduled Penta3. For rotavirus, this is consistent with both the strict age eligibility cutoffs and the August 2022 schedule introduction that made the vaccine free while preserving the need for post-birth clinic contacts; for IPV, it is consistent with the supply-side prioritization arguments developed in Section C.

C The Rotavirus and IPV Exceptions

Rotavirus. Against the backdrop of declining coverage for the Penta and OPV series, rotavirus stands out: Rota1 and Rota2 coverage among post-reform cohorts is maintained at or above the level predicted by their position in the schedule. In the across-child specification, the β_{Rota} coefficient is substantially smaller in absolute value than β_{Penta} at the same contact points (six weeks and ten weeks), and is not statistically distinguishable from zero. In the within-child specification, the post-window effect for rotavirus is similarly attenuated relative to comparable vaccines.

The mechanism for this exception combines the two rotavirus features described in Section B. First, Rota1 must be given between six and fourteen weeks, and Rota2 must be given before eight months. Because missing the window permanently forecloses catch-up vaccination, clinicians have a strong supply-side incentive to prioritize Rota when scheduling is

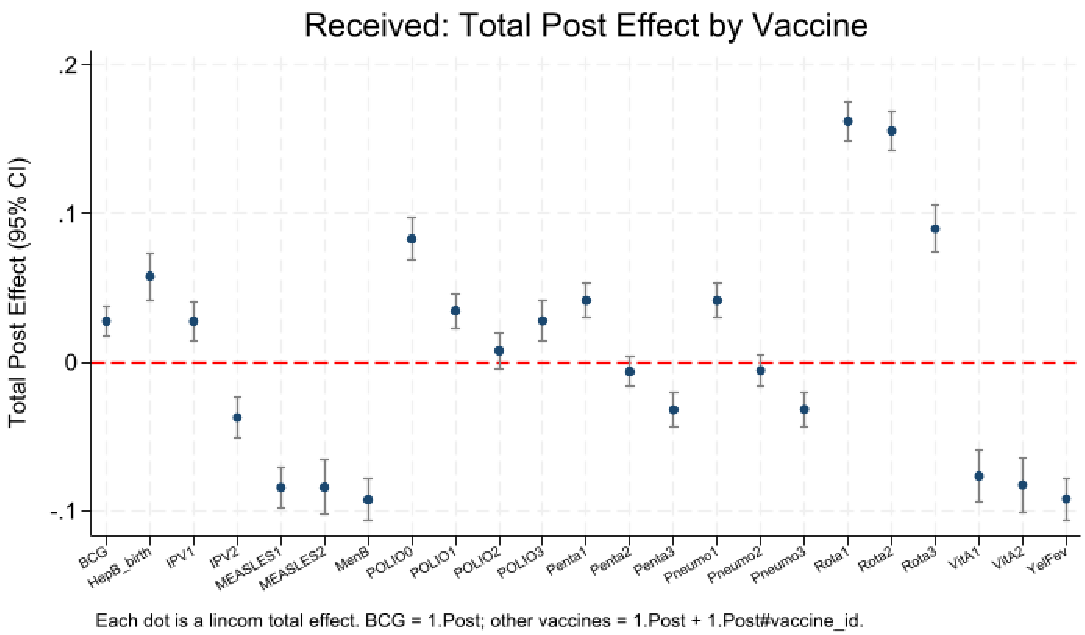

Figure 9: Coverage Effects by Vaccine and Dose Position — Full Sample

Notes: Coefficient estimates from equation (4) for the full sample. Each dot is the estimated change in receipt probability for a child whose vaccine window falls in the post-reform period, identified off of the across-cohort within-vaccine comparison. Estimates are expressed relative to BCG (at-birth placebo, normalized to zero). Vaccines are ordered by scheduled age. Horizontal bars are 95-percent confidence intervals; standard errors clustered at the state-by-urban level.

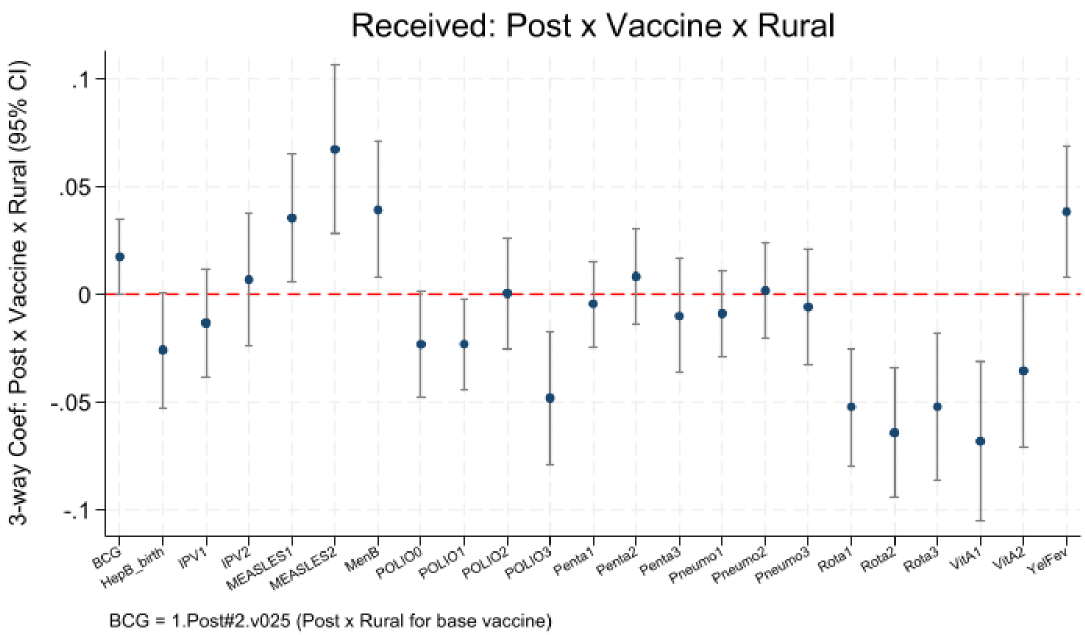

Figure 10: Coverage Effects by Vaccine and Dose Position — Rural Sample

Notes: As for Figure 9, restricted to children in rural DHS clusters. The dose-gradient is more pronounced in rural areas, consistent with compounding transportation barriers where each clinic visit involves a longer journey. At-birth vaccines remain near zero (placebo). Coverage shortfalls for the nine-month vaccines (Measles, Vitamin A) reach approximately 8–10 percentage points; Rota2 and IPV are exceptions, deviating upward by roughly 16 percentage points relative to co-scheduled Penta3.

compressed. Second, rotavirus entered Nigeria's routine immunization schedule in August 2022 as a free, government-backed vaccine. Before that schedule change, households that wanted rotavirus protection generally had to pay for the vaccine outside the NPI. After the change, the product price was removed, but families still had to attend the separate six- and ten-week clinic contacts. When fewer parents show up for routine visits — as would follow from higher transportation costs — health workers therefore have reason to ensure that Rota-eligible children who do appear are completed, while the remaining barrier is the access cost of reaching the visit.

This cost-structure change gives us an additional mechanism exercise to pursue. In the next iteration, we will exploit the August 2022 introduction of free rotavirus vaccine to compare a setting where the vaccine price changed sharply but the clinic-visit requirement remained in place. That contrast can help separate the accessibility effect of higher fuel and transport costs from competing explanations such as lower willingness to pay for the vaccine product, general vaccine hesitancy, or a broad deterioration in demand for child health investments.

This exception rules out the two main alternative mechanisms. Pure parental disengagement or demand withdrawal would reduce all vaccines proportionally; it would not generate a vaccine-specific exception keyed jointly to the age eligibility structure and the removal of the product price. Anti-vaccine sentiment would similarly not produce a pattern where Rota is maintained while Penta is not. The rotavirus exception is therefore the sharpest mechanism test in the paper: it identifies access disruption — operating through the supply side, via reduced caregiver-initiated visits that trigger a compensating response for the most age-constrained vaccine, in a setting where the vaccine itself had recently become free — as the operative channel.

IPV. Inactivated polio vaccine shows a similar pattern of maintained coverage relative to co-scheduled vaccines (Penta3 and OPV3 at sixteen weeks). IPV was introduced into

Nigeria's routine schedule in 2015 as part of the global polio eradication initiative (Nigeria 2015), and its delivery continues to be integrated with supplemental immunization activities (SIAs) funded by the Global Polio Eradication Initiative (Initative 2026) . In practice, this means IPV is frequently delivered through outreach activities that are partially supported by external logistics and funding, reducing the dependence on individual caregiver transport for completion (Biya 2023; Initative 2023).

The maintained IPV coverage provides convergent evidence alongside rotavirus: vaccines with strong supply-side prioritization — whether from age-eligibility urgency and a free-schedule introduction (Rota) or programmatic intensity (IPV) — are partially insulated from the access-cost disruption. Vaccines without these features, namely the core primary series of Penta, PCV, and OPV, show the dose-gradient decline consistent with compounding transportation barriers.

Interpretation. Together, the rotavirus and IPV results sharpen the mechanism claim. The pattern is not consistent with a narrative where healthcare workers or parents abandoned the vaccination program broadly. Instead, it shows a system under access strain that nonetheless prioritizes specific programmatic targets. This is the realistic response to a transportation cost shock: supply-side actors triage, and the triage produces vaccine-specific exceptions that map predictably onto the structure of clinical incentives, programmatic attention, and the separation between vaccine-product prices and clinic-visit costs.

D Censoring-Corrected Timeliness Analysis

As noted in Section B, the ever-received outcome is susceptible to mechanical right-censoring: post-reform cohorts are on average younger at the time of interview and have had less calendar time to accumulate later doses, which would depress observed receipt rates for the nine-month vaccines regardless of any true behavioral change. We address this by constructing a censoring-corrected outcome from vaccination card dates, available for approximately 52 percent of the sample: an indicator for whether each dose was administered *within* the WHO-recommended age window ("on-time" receipt). This outcome is immune to censoring because it conditions on whether the family kept up with the schedule in real time, not whether the dose was eventually received by interview date.

Figures 11 and 12 repeat the dose-ordered coefficient plot using on-time receipt as the dependent variable, for the full sample and the rural subsample respectively. The dose-gradient pattern is preserved on this censoring-corrected measure. Among the full sample, on-time receipt declines monotonically with dose position, reaching approximately -8 to -10 percentage points for the nine-month vaccines relative to BCG. The rural subsample shows a steeper gradient, consistent with greater access barriers in areas further from facilities. Critically, the pattern closely mirrors the ever-received results, confirming that the dose-gradient in Section B is not an artifact of the age-composition of post-reform cohorts but reflects a genuine deterioration in timely care-seeking.

The card sample is a selective subsample: card-holding households are more engaged with the health system and are likely to have been less affected by the reform than the full population. Timeliness estimates therefore represent a *lower bound* on the true reduction in on-schedule vaccination. The true decline for the full population of children, including those without cards, is almost certainly larger.

The Rota and IPV exceptions again appear in the on-time data. Because rotavirus combines strict age cutoffs with a recent move into the free routine schedule, Rota doses that are administered are administered on time: health workers who prioritize Rota do so precisely to capture the eligible window, and families no longer face a product price once they reach the clinic. Post-reform cohorts therefore show no deterioration in Rota timeliness relative to Penta, providing yet another dimension of the supply-side exception: the prioritization that maintains Rota *coverage* also maintains its *timeliness*.

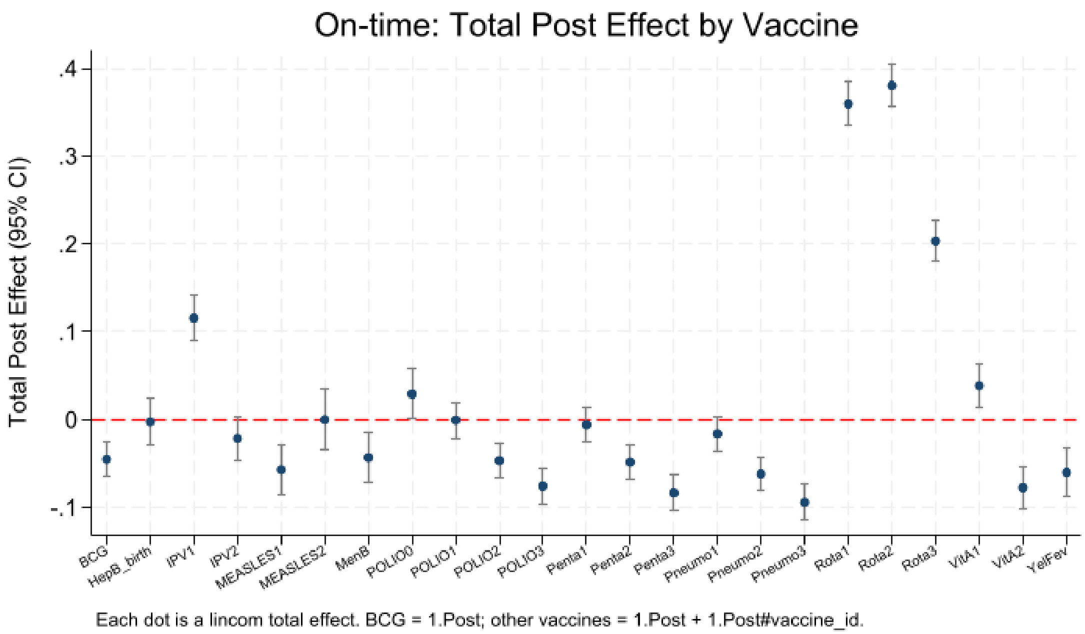

Figure 11: On-Time Receipt Effects by Vaccine and Dose Position — Full Sample

Notes: Coefficient estimates from equation (4) with on-time receipt as the outcome, estimated on the card subsample (≈ 52% of the eligible sample with a vaccination card available at interview). On-time receipt is defined as administration within the WHO-recommended age window. Estimates are expressed relative to BCG (at-birth placebo). Vaccines ordered by scheduled age. Horizontal bars are 95-percent confidence intervals; standard errors clustered at the state-by-urban level. The persistence of the dose-gradient on this censoring-corrected measure confirms the results in Figure 9 are not driven by the younger age of post-reform cohorts at interview.

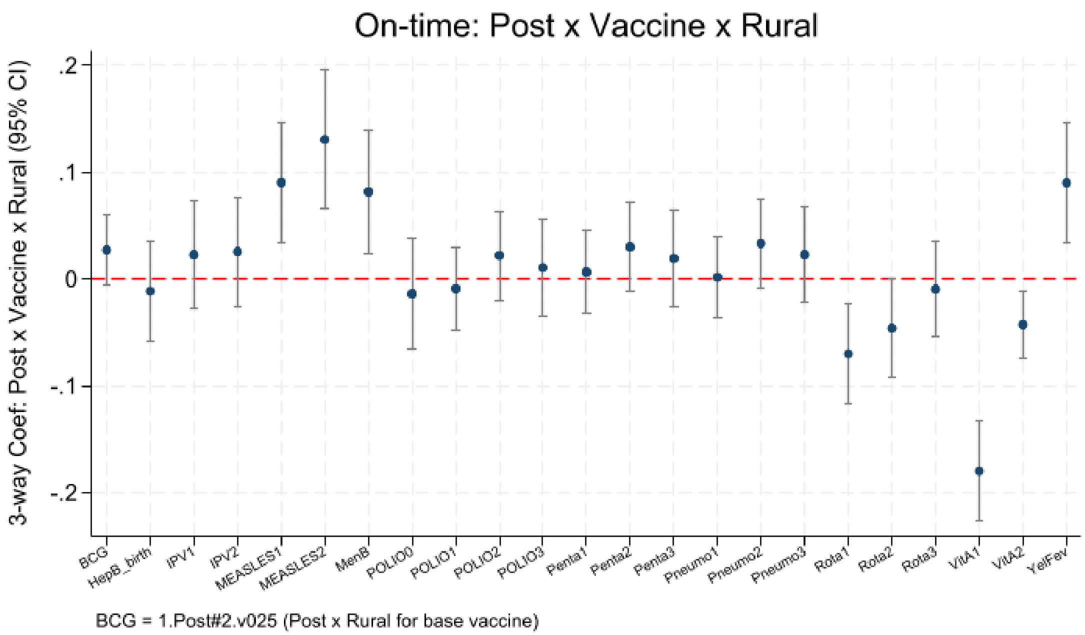

Figure 12: On-Time Receipt Effects by Vaccine and Dose Position — Rural Sample

Notes: As for Figure 11, restricted to children in rural DHS clusters. The dose-gradient on the on-time measure is more pronounced in rural areas than in the full sample, consistent with steeper access barriers outside urban centers. At-birth vaccines remain near zero (placebo). Rota and IPV continue to deviate from the monotone gradient, reflecting supply-side prioritization and, for Rota, the August 2022 shift to a free vaccine with clinic-visit requirements still in place.

E Sibling Effects and Trip-Bundling

The aggregate dose-gradient masks an important dimension of within-sample heterogeneity. Households with multiple young children face a higher total visit burden, and a fixed transportation budget must be spread across more vaccine appointments. If the transportation cost shock is binding, sibling presence should amplify the negative effect on later doses: the family cannot afford all the trips needed and must triage, favoring earlier doses where the health stakes are most immediate.

We estimate a sibling-presence heterogeneity specification by interacting PostWindow_{iv} with an indicator for whether the child has at least one sibling born within twelve months in the household, using the within-child fixed-effects specification (equation 5) estimated separately for urban and rural clusters.

Figure 13 presents the results by dose group and urban/rural stratum. The pattern across dose positions is not monotonically negative. For at-birth and six-week vaccines, the PostWindow and PostWindow $\times$ Sibling coefficients are near zero or modestly negative, consistent with BCG being a weak placebo and with early doses facing the lowest access barriers. Moving through the schedule, the PostWindow effect becomes more negative with dose order, reflecting compounding transportation costs. The sibling interaction amplifies this gradient in rural areas: rural children with a close-in-age sibling show a meaningfully larger coverage shortfall at the ten-week and sixteen-week contact points relative to rural children without young siblings. Urban children show a smaller and statistically insignificant sibling interaction, consistent with urban households facing lower per-trip costs.

The nine-month dose group produces a striking reversal that is itself informative. The PostWindow $\times$ Sibling interaction for the nine-month vaccines (Measles and Vitamin A) is positive and statistically significant, particularly in rural areas. This reversal reflects a *trip-bundling* mechanism. The nine-month contact point requires a dedicated clinic visit distinct from the earlier pentavalent series; when two children of similar age need the same nine-month appointment, the family can consolidate both visits into a single trip. The per-dose transportation cost of the nine-month vaccines is therefore *lower* for households with a same-age sibling than for those with a single child, reversing the direction of the sibling effect at that contact point. That the sibling effect reversal is concentrated at the nine-month vaccines — the only dose position that requires a dedicated visit not embedded in an earlier contact-point cluster — is direct evidence for the trip-bundling mechanism operating alongside the triage mechanism at earlier doses.

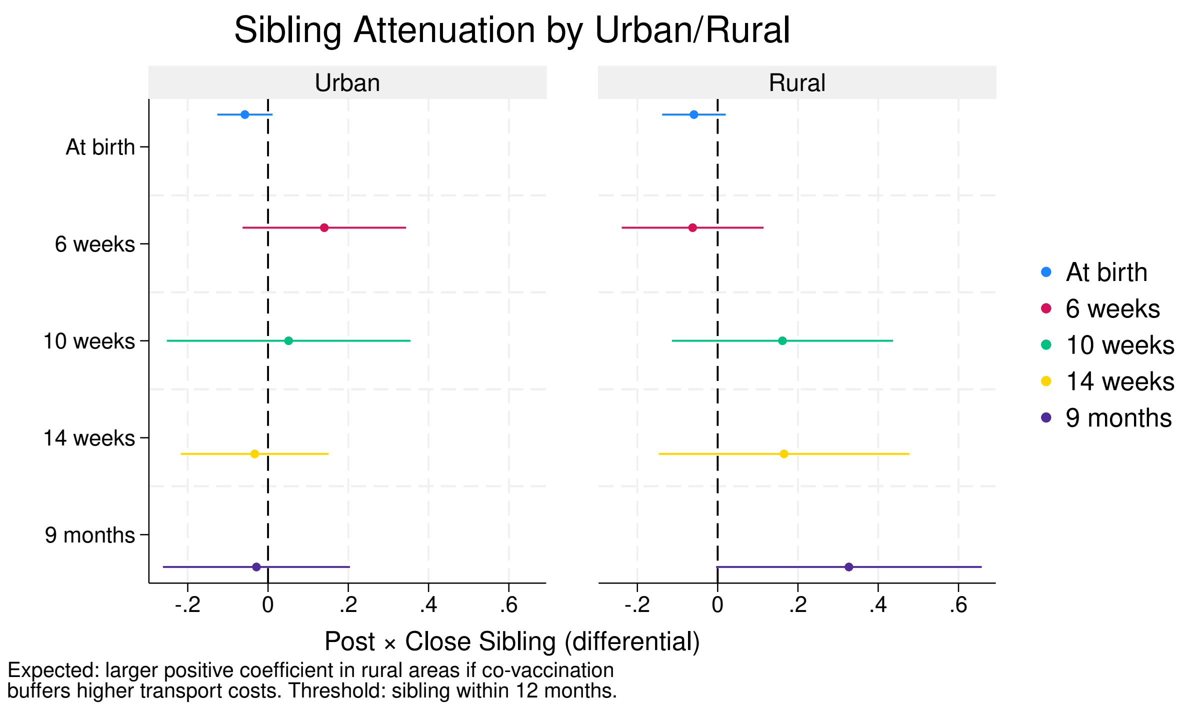

Figure 13: Sibling Interaction Effects by Dose Group and Urban/Rural Stratum

Notes: Coefficient estimates from the within-child specification (equation 5) interacting PostWindow$_{iv}$ with an indicator for having a sibling born within twelve months. Specifications estimated separately for urban and rural clusters. Each panel shows, by dose group, the PostWindow main effect and the PostWindow × Sibling interaction. Child and vaccine fixed effects included; standard errors clustered at the state level. The positive and significant sibling interaction at the nine-month dose group in rural areas reflects trip-bundling: two same-age siblings share a single dedicated trip to the nine-month contact point, lowering the per-child cost.

IX Aftermath: Living Standards and Public Support

The three health results document the reform's effect on preventive healthcare at the individual level. This section provides two forms of broader context: first, evidence from a separate household panel that household welfare deteriorated in ways consistent with the proposed mechanisms; second, evidence from post-reform public opinion surveys that these welfare costs were sufficiently salient to register in public attitudes toward the reform itself.

A Household Living Standards (NLPS)

Table 12 reports results from the NLPS household panel for food, clothing, and housing adequacy. After the reform, households are 11.5 percentage points less likely to report adequate food and 5.1 percentage points less likely to report adequate clothing on average, indicating broad deterioration in living standards. These are large effects: the sample means are 38.5 percent for food adequacy and 53.4 percent for clothing adequacy, so the food decline represents a 30 percent relative reduction.

For food adequacy, the Urban × Post coefficient is negative and statistically insignificant, implying no meaningful urban-rural differential — food costs rose for all households. For clothing and housing adequacy, the Urban × Post coefficients are positive and significant at 7.2 and 7.5 percentage points, respectively. This pattern indicates that rural households experienced larger declines in non-food living standards after the reform, consistent with lower incomes, fewer substitution options, and more limited access to the formal retail sector.

Table 12: Effects of Fuel Subsidy Removal on Household Living Standards (NLPS)

	(1)	(2)	(3)	(4)	(5)	(6)
	Adequate food		Adequate clothing		Adequate housing	
Post	-0.115^{***}	-0.109^{***}	-0.051^{**}	-0.082^{***}	0.008	-0.025
	(0.018)	(0.025)	(0.020)	(0.026)	(0.019)	(0.024)
Urban × Post	—	-0.013	—	0.072^{*}	—	0.075^{*}
		(0.036)		(0.039)		(0.039)
Obs. / Mean	2,230 / 0.385	2,230 / 0.385	2,230 / 0.534	2,230 / 0.534	2,230 / 0.599	2,230 / 0.599
Household FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Post is an indicator equal to one for NLPS rounds conducted on or after June 2023 (rounds 9 and 10). Household fixed effects absorb all time-invariant household characteristics. Standard errors clustered at the survey cluster level. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Two notes of caution apply to the NLPS evidence. First, the NLPS is a separate household panel not linked to DHS births or children; the living standards effects cannot be directly attributed to the households whose children's health we study. The NLPS evidence is therefore suggestive corroboration rather than a mechanism estimate tied to individual births. Second, the NLPS sample of approximately 2,400 households, distributed across urban and rural areas, has limited statistical power for detecting urban-rural differentials. The significant urban-rural gap for clothing and housing should be interpreted cautiously.

The pattern of food versus non-food adequacy is nonetheless economically informative. The absence of an urban-rural differential for food but the presence of one for clothing and housing is consistent with the interpretation that food price inflation was nationally uniform (petrol drives food distribution costs everywhere), while non-food living standards were more affected in rural areas where household incomes depend more on agriculture and informal employment, both sectors exposed to fuel costs.

B Public Support for the Reform (Afrobarometer)

Afrobarometer Round 10 provides post-reform public opinion data on the fuel subsidy removal. [PLACEHOLDER — data access and exact survey timing to be confirmed.] Among Nigerian respondents interviewed after the reform, a majority report that the subsidy removal made their lives worse, and approval rates for the reform are substantially lower than the pre-removal expectations documented in earlier rounds. Rural respondents and lower-income respondents express more negative attitudes, consistent with the distributional health and welfare impacts documented in this paper.

This section is intentionally brief. The Afrobarometer evidence does not add identification; it closes the welfare narrative. The health effects documented in Sections VI–VIII and the living standards evidence above were visible enough to register in population-level survey preferences. This suggests that the welfare costs of the reform were not diffuse, temporary, or invisible to the affected population — they were salient, concentrated, and concentrated where we predict they should be concentrated: among rural and lower-resource households.

X Robustness: Bridging-Share Instrument

The main results exploit temporal variation in post-reform exposure across birth cohorts and vaccination windows. A significant identification concern is that the 2023 macro environment introduced multiple concurrent national shocks: the unification of Nigeria's official and parallel exchange rates in June 2023, the currency redesign and associated cash shortages, and the acceleration of headline inflation driven partly by factors independent of the fuel price. These shocks are national-level, but their distributional incidence across states may not be identical to the incidence of the fuel price shock. In particular, states with larger financial sectors or more dollarized trade may have been differentially exposed to the exchange rate shock, creating a potential confound for state-level variation in health outcomes.

This section constructs a shift-share instrument for fuel price exposure using pre-reform logistics data from the Nigerian Midstream and Downstream Petroleum Regulatory Authority (NMDPRA). The instrument isolates variation in fuel price pass-through that derives from pre-reform supply chain infrastructure rather than from state-level financial or economic characteristics, providing a robustness check for the main results that is orthogonal to the concurrent macro confounds.

A The NMDPRA Depot-Level Data

Data source and coverage. The NMDPRA publishes daily petroleum supply reports that record fuel pricing and movement at licensed depots in Nigeria's downstream petroleum sector. We digitized two report series: December 2021 through June 2022 and January 2023 through August 2023, yielding 417 usable daily reports in total, with some gaps in coverage. The first series contains 187 usable daily reports from December 2021 through June 2022. The second series contains 230 usable daily reports from January 2023 through August 2023. These reports provide a daily depot-level panel of premium motor spirit (PMS) truck-out volumes, prices, and local/bridging classifications.[1] The data are aggregated to a state-month panel for the analysis.

The bridging/local distinction and bridging-share construction. Under the pre-removal subsidized regime, PMS distribution operated through a network of primary storage depots and secondary distribution depots. Fuel originating at a primary depot and transported to a secondary distribution depot before reaching retail stations was classified as *bridging* supply. Fuel loaded directly from a local depot to retail trucks was classified as *local* supply. The NMDPRA daily reports record each state's daily PMS receipts separately under the bridging and local designations, allowing us to compute state-level bridging shares before the subsidy removal.

Because the digitized reports come from two discontinuous coverage windows, we construct the bridging share separately for each series. Let $k = 1$ denote the December 2021–June 2022 series and $k = 2$ denote the January 2023–August 2023 series. For the first series, the bridging share uses all usable reports from December 2021 through June 2022. For the second series, the bridging share uses only the pre-removal reports from January 2023 through May 2023. We define

$$bs_{s,k} = \frac{\sum_{\tau \in \mathcal{T}_k^{pre}} \text{BridgingReceipts}_{s\tau}}{\sum_{\tau \in \mathcal{T}_k^{pre}} (\text{BridgingReceipts}_{s\tau} + \text{LocalReceipts}_{s\tau})} \tag{6}$$

where s indexes states, $k \in \{1, 2\}$ indexes the report series, $\mathcal{T}_1^{pre}$ contains usable daily reports from December 2021 through June 2022, and $\mathcal{T}_2^{pre}$ contains usable daily reports from January 2023 through May 2023. Thus, $bs_{s,1}$ measures state s's bridging share in the 2022 series, while $bs_{s,2}$ measures state s's bridging share in the 2023 pre-removal series.

[1] An example report can be retrieved here: https://alps.blob.core.windows.net/nmdprawebsite/OldWebsiteFiles/uploads/202 DAILY-PMS-TRUCK-OUT-REPORT-FOR-16TH-MAY-2022.pdf. The digitization pipeline converts the daily PDF reports published by NMDPRA into structured data using document parsing and optical character recognition. Each daily report contains five tables: a truck-out analysis, a dispatches report, a bridging receipt report, a depot loading report, and a decanted products report. The truck-out analysis reports depot- and state-level PMS movements separately by local and bridging designation. Details on the extraction procedure will be released in the Data Appendix in the next iteration of the draft.

This series-specific construction avoids pooling across two discontinuous coverage windows and ensures that the bridging share is measured only under the pre-removal supply-chain regime. The June 2023–August 2023 reports are used to describe the immediate post-removal price and supply response, but not to construct $bs_{s,2}$, because the local/bridging classification and the supply-chain structure it captures become less consistently observable after the subsidy removal.

Figure 14 shows that the state-level bridging shares are highly stable across the two pre-removal report windows. The scatter plot compares $bs_{s,1}$, constructed from the December 2021–June 2022 series, with $bs_{s,2}$, constructed from the January 2023–May 2023 pre-removal portion of the 2023 series. The two measures are strongly correlated, with a correlation coefficient of 0.983, and the fitted line closely tracks the 45-degree line. States with high bridging dependence in the 2022 series also tend to have high bridging dependence in the 2023 pre-removal series, suggesting that the measure captures persistent features of the fuel distribution network rather than short-run reporting noise.

The figure also shows that bridging dependence is highly polarized across states. Nineteen of the 37 states have bridging shares close to one, indicating near-complete reliance on bridging receipts under the pre-removal regime. By contrast, eight states have bridging shares below 0.1, indicating that most PMS receipts were classified as local. The remaining states lie between these two poles, with bridging shares between 0.1 and 0.9. This polarization is consistent with the geographic and infrastructural logic of the pre-removal supply chain: states close to major import terminals, refineries, or local depot capacity could rely more heavily on local supply, whereas states farther from these supply nodes depended more heavily on inter-depot bridging.

Key limitation: absence of origin-destination dyad data. The NMDPRA reports record the volume of bridging receipts at each *destination* state and the volume of bridging dispatches from each *origin* depot, but they do not record the dyadic origin-destination pair

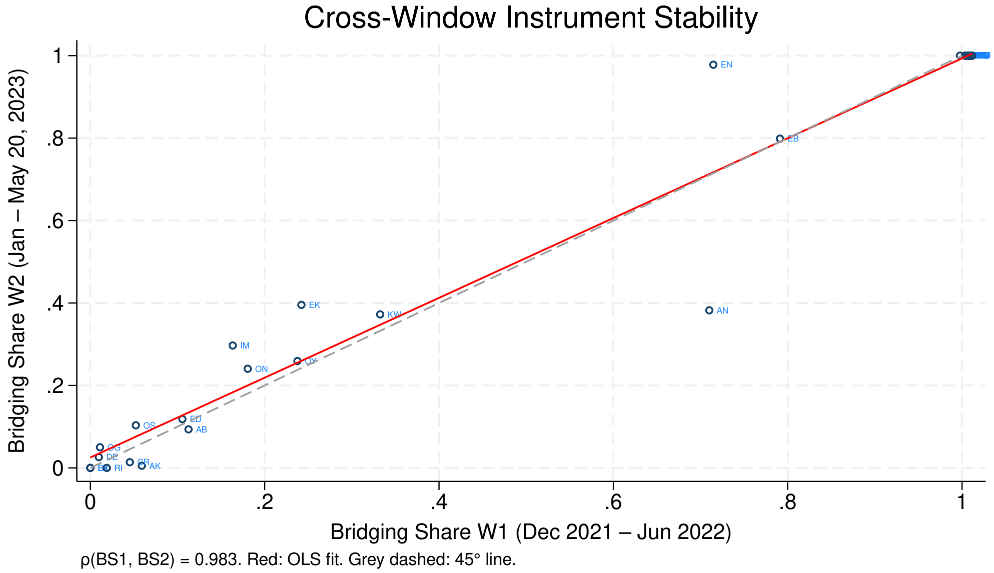

Figure 14: NESG Predicted vs. Realized Fuel Price Change by State

Notes: The figure plots state-level bridging shares constructed from the two pre-removal report windows. The horizontal axis reports $bs_{s,1}$, constructed using usable daily reports from the December 2021–June 2022 series. The vertical axis reports $bs_{s,2}$, constructed using usable pre-removal daily reports from January 2023 through May 2023. Each point represents a state. The red line reports the OLS fitted line, and the grey dashed line reports the 45-degree line. The correlation between $bs_{s,1}$ and $bs_{s,2}$ is 0.983.

for individual shipments. This means that we cannot directly observe which origin depot supplied which destination state on any given day. The bridging share in equation (6) is therefore constructed from separately observed origin and destination records rather than from joint dyads.

This limitation has two implications. First, hub states — states that host a primary depot and both send bridging supplies to other states and receive bridging supplies from elsewhere — may be misclassified: their reported bridging receipts may reflect within-hub movements rather than true inter-hub dependence. We address this by excluding states identified as major hub depots from the bridging share calculation in robustness checks. Second, the bridging share may not perfectly proxy for the vulnerability to supply chain disruption that is the theoretical first stage: a state could have a high bridging share because it is genuinely supply-chain-dependent on inter-depot logistics, or because local administrative reporting practices double-count certain movements. We address this by validating the bridging share against independently observable price data (Section B).

The post-reform disappearance of the bridging distinction. A key piece of corroborating evidence for the instrument is that the bridging/local decomposition in the NMDPRA data effectively ceases to be informative after May 29, 2023. Under the subsidized regime, NNPC maintained strict reporting of bridging movements because it was responsible for the cost differential between local and bridging supply. After price deregulation, private distributors took over supply responsibility and the bridging/local designation lost its administrative meaning: the NMDPRA reports continue to record the categories, but the bridging share column becomes uninformative as the supply chain restructures around market prices. This structural break confirms that the bridging share reflects a feature of the pre-reform subsidized supply chain — not an intrinsic property of geography — and that our use of pre-reform bridging shares as a time-invariant state characteristic is appropriate.

B Instrument Construction and Identification

The shift-share instrument. The instrument follows the Bartik-style logic: we interact a pre-determined state characteristic (the pre-reform bridging share) with a post-reform indicator to generate state-specific predicted fuel price variation:

$$Z_{st} = \text{BridgeShare}_s \times \mathbf{1}\{t \geq \text{June 2023}\} \tag{7}$$

States with higher pre-reform bridging shares experienced larger disruptions to their fuel supply chains after deregulation, because the inter-depot logistics pricing that bridging relied upon shifted to market rates. The prediction is that these states experienced larger and more variable fuel price increases after May 2023, providing the first-stage variation.

First-stage specification. The first-stage regression is:

$$\text{FuelPrice}_{st} = \pi_0 + \pi_1 Z_{st} + \phi_s + \phi_{m(t)} + \mathbf{V}'_{st}\theta + u_{st} \tag{8}$$

where FuelPrice_{st} is the average PMS retail price in state s and month t, ϕ_s are state fixed effects, $\phi_{m(t)}$ are month-of-year fixed effects, and $\mathbf{V}_{st}$ includes state-level controls. The coefficient π_1 measures the additional post-reform fuel price increase in states more dependent on bridging routes.

Exclusion restriction. The exclusion restriction requires that, conditional on state and month fixed effects, pre-reform bridging dependence affects health outcomes only through its effect on post-reform fuel prices. This is plausible for several reasons. Bridging shares are determined by infrastructure geography — depot locations and pipeline capacity built during Nigeria's oil sector development from the 1970s onward — which is not a function of state-level healthcare quality, maternal education, or pre-existing trends in health outcomes. We provide pre-reform falsification evidence: in the pre-reform period, bridging share does not predict trends in vaccination rates, ANC utilization, or birth outcomes across states. We also show that bridging share is not correlated with state-level measures of financial development, urbanization, or healthcare infrastructure that might independently affect how the concurrent macro shocks (naira devaluation, currency reform) pass through to welfare outcomes.

The instrument is therefore conceptually distinct from both the temporal variation (before/after for the same children) and the spatial urban-rural variation in the main results. It isolates a third source of variation — cross-state heterogeneity in supply-chain vulnerability — that is plausibly orthogonal to the confounds introduced by the 2023 macro environment.

Two identification strategies and their relationship. We implement two related but distinct strategies using the NMDPRA data.

Strategy 1 (Reduced-form cross-state): Use pre-reform bridging share directly as a continuous treatment variable, comparing states with high versus low bridging dependence in a state-by-month DiD framework. This strategy makes no claim about the specific mechanism through fuel prices; it asks whether states more exposed to supply chain disruption saw worse health outcomes after the reform. The advantage is transparency and interpretability; the cost is that bridging share may capture other dimensions of rural supply-chain dependence beyond fuel specifically.

Strategy 2 (IV for fuel prices): Use the bridging share instrument (equation 7) as an instrument for actual fuel prices in each state-month, then estimate the second-stage effect of instrumented fuel prices on health outcomes. This strategy requires the additional assumption that bridging share affects outcomes only through fuel prices, not through other supply chain disruptions. It yields structurally interpretable coefficients (effects per unit of fuel price increase) that can be compared to the main results.

C First-Stage Results

This section will be updated in the next iteration of the draft.

D Second-Stage Results

Vaccination coverage and timeliness. This section will be updated in the next iteration of the draft.

Birth outcomes and ANC. This section will be updated in the next iteration of the draft.

Rotavirus cost-structure check. This section will be updated in the next iteration of the draft.

Assessment. This section will be updated in the next iteration of the draft.

XI Conclusion

This paper studies how a large, sharp, and nationally unanticipated energy price reform affected preventive healthcare investments at every stage of early life. Using Nigeria's May 2023 fuel subsidy removal and three complementary identification designs, we document a consistent pattern: the reform disrupted healthcare access at the pre-birth, at-birth, and post-birth stages, and the disruption was concentrated where access is most transportation-dependent — rural areas. The spatial pattern replicates across three independent samples and across multiple outcome domains.

The three main results are as follows. First, post-reform exposure reduced antenatal care use and weakened several dimensions of care content, including diagnostic testing and pregnancy counseling. These declines were not always larger in rural areas in absolute terms, but they represented larger proportional losses because rural women entered the reform with substantially lower baseline utilization.

Second, at birth, the reform widened urban-rural differences in mother-reported categorical birth size. Rural births became more likely to be reported as small or below, while urban births showed weaker deterioration and, for the broader average-or-below measure, shifted toward larger reported size categories. The mechanism evidence suggests that larger urban declines in combustion-related $PM_{2.5}$ partly offset the access-cost channel, especially among higher-resource households.

Third, post-birth vaccination coverage fell for vaccines requiring return clinic visits, with shortfalls increasing across the dosing schedule. Timeliness also worsened, consistent with delayed and bundled visits rather than complete exit from vaccination. The main exceptions, rotavirus and IPV, point to supply-side prioritization and programmatic protection rather than a general collapse in vaccine demand.

Taken together, these findings show that removing a universal fuel-price subsidy affected early-life health through dynamic access costs, not only through the static redistribution of subsidy benefits. Fossil fuel consumption subsidies are unusual but important because they are universal and untargeted government interventions. Like public education, infrastructure provision, or universal health programs, they affect households through common prices and access conditions rather than eligibility rules. The static incidence of such a subsidy asks who receives the largest implicit transfer while the subsidy is in place. The dynamic incidence of removal asks who faces the largest adjustment cost once the common price structure changes. These two distributions need not coincide.

This distinction matters for policy design. Most existing work emphasizes the static incidence of universal subsidies: who benefits while the subsidy is in place. This paper shows that the dynamic incidence of removal can look different because households do not simply lose a transfer; they face a new cost structure for reaching essential services. A reform can therefore be desirable on fiscal, environmental, or efficiency grounds while still imposing predictable short-run costs on specific households and services. In this setting, the burden fell most clearly on rural households whose healthcare access depends more heavily on transport and whose baseline use of antenatal and immunization services was already lower. The relevant policy question is therefore not only whether to remove a regressive subsidy, but how to design countervailing policies before removal occurs.

The results suggest that cash compensation alone may be incomplete. Transfers can offset some income losses, but they do not directly reduce travel time, transport uncertainty, or the coordination cost of bringing children repeatedly to clinics. For early-life health, the more direct complement is access-preserving service delivery: rural immunization outreach, mobile antenatal care, transportation support, and scheduling systems that reduce repeat-visit burdens. Outreach-based delivery — bringing vaccines and preventive care closer to households rather than requiring households to absorb higher travel costs — is therefore a natural policy response to the access-cost harm documented here Meriggi et al. 2024.

References

Abubakar, I., S. L. Dalglish, B. Angell, O. Sanuade, S. Abimbola, A. L. Adamu, I. M. O. Adetifa, T. Colbourn, A. O. Ogunlesi, O. Onwujekwe, E. T. Owoaje, I. N. Okeke, A. Adeyemo, G. Aliyu, M. H. Aliyu, S. H. Aliyu, E. A. Ameh, B. Archibong, A. Ezeh, M. A. Gadanya, C. Ihekweazu, V. Ihekweazu, Z. Iliyasu, A. K. Chiroma, D. A. Mabayoje, M. N. Sambo, S. Obaro, A. Yinka-Ogunleye, F. Okonofua, T. Oni, O. Onyimadu, M. A. Pate, B. L. Salako, F. Shuaib, F. Tsiga-Ahmed, and F. H. Zanna (Mar. 2022). "The Lancet Nigeria Commission: investing in health and the future of the nation". English. In: *The Lancet* 399.10330. Publisher: Elsevier, pp. 1155–1200.

Almond, D. and J. Currie (Mar. 2010). "Human Capital Development Before Age Five". en. In: *NBER Working Papers*. Number: 15827 Publisher: National Bureau of Economic Research, Inc.

Almond, D., H. W. Hoynes, and D. W. Schanzenbach (2011). "Inside the War on Poverty: The Impact of Food Stamps on Birth Outcomes". In: *Review of Economics and Statistics* 93.2, pp. 387–403.

Alsan, M. and C. Goldin (2019). "Watersheds in Child Mortality: The Role of Effective Water and Sewerage Infrastructure, 1880–1920". In: *Journal of Political Economy* 127.2, pp. 586–638.

Aniebo, C. L., L. O. Lawani, and P. Eze (Dec. 2025). "The Burden and Socioeconomic Inequality in Catastrophic Out-of-pocket Health Expenditure in Post-Pandemic Nigeria". en. In: *Global Social Welfare*.

Arze del Granado, F. J., D. Coady, and R. Gillingham (2012). "The Unequal Benefits of Fuel Subsidies: A Review of Evidence for Developing Countries". In: *World Development* 40.11, pp. 2234–2248.

Banerjee, A. V., E. Duflo, R. Glennerster, and D. Kothari (May 2010). "Improving immunisation coverage in rural India: clustered randomised controlled evaluation of immunisation campaigns with and without incentives". en. In: *BMJ* 340. Publisher: British Medical Journal Publishing Group Section: Research, p. c2220.

Biya, O. (2023). "Notes from the Field: House-to-House Campaign Administration of Inactivated Poliovirus Vaccine — Sokoto State, Nigeria, November 2022". en-us. In: *MMWR. Morbidity and Mortality Weekly Report* 72.

Black, S., A. A. Liu, I. Parry, and N. Vernon (2023). *IMF Fossil Fuel Subsidies Data: 2023 Update*. IMF Working Paper WP/23/169. Washington, DC: International Monetary Fund.

Brauer, M., C. Lencar, L. Tamburic, M. Koehoorn, P. Demers, and C. Karr (2008). "A Cohort Study of Traffic-Related Air Pollution Impacts on Birth Outcomes". In: *Environmental Health Perspectives* 116.5, pp. 680–686.

Chay, K. Y. and M. Greenstone (2003). "The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession". In: *Quarterly Journal of Economics* 118.3, pp. 1121–1167.

Clements, B., D. Coady, S. Fabrizio, S. Gupta, T. Alleyne, and C. Sdralevich (2013). *Energy Subsidy Reform: Lessons and Implications*. Washington, DC: International Monetary Fund.

Coady, D., V. Flamini, and L. Sears (2015). *The Unequal Benefits of Fuel Subsidies Revisited: Evidence for Developing Countries*. IMF Working Paper 15/250. International Monetary Fund.

Currie, J. and J. Gruber (1996). "Saving Babies: The Efficacy and Cost of Recent Changes in the Medicaid Eligibility of Pregnant Women". In: *Journal of Political Economy* 104.6, pp. 1263–1296.

Currie, J., M. Neidell, and J. F. Schmieder (2009). "Air Pollution and Infant Health: Lessons from New Jersey". In: *Journal of Health Economics* 28.3, pp. 688–703.

Currie, J. and R. Walker (2011). "Traffic Congestion and Infant Health: Evidence from E-ZPass". In: *American Economic Journal: Applied Economics* 3.1, pp. 65–90.

Dahiru, T. and O. M. Oche (2015). "Determinants of antenatal care, institutional delivery and postnatal care services utilization in Nigeria". eng. In: *The Pan African Medical Journal* 21, p. 321.

Davis, L. W. (2014). "The Economic Cost of Global Fuel Subsidies". In: *American Economic Review* 104.5, pp. 581–585.

Dewidar, O., J. John, A. Baqar, M. T. Madani, A. Saad, A. Riddle, E. Ota, et al. (2023). "Effectiveness of Nutrition Counseling for Pregnant Women in Low- and Middle-Income Countries to Improve Maternal and Infant Behavioral, Nutritional, and Health Outcomes: A Systematic Review". In: *Campbell Systematic Reviews* 19.4, e1361.

Dzirutwe, M. (Oct. 2022). "Nigeria's Tinubu: Will remove fuel subsidy, deregulate gas prices if elected". en. In: *Reuters*.

Dzirutwe, M. and C. Eboh (May 31, 2023a). *Nigeria Triples Petrol Prices After President Says to Scrap Subsidy*. Reuters.

Dzirutwe, M. and C. Eboh (2023b). "Nigeria Triples Petrol Prices After President Says to Scrap Subsidy". In: *Reuters*.

Figueiredo, A. C. M. G., I. S. Gomes-Filho, J. M. F. Coelho, J. E. T. Batista, E. S. Passos, M. I. P. Vianna, and S. C. Trindade (2018). "Maternal Anemia and Low Birth Weight: A Systematic Review and Meta-Analysis". In: *Nutrients* 10.5, p. 601.

Folami, R. A., L. Elumah, and B. Ilo (2025). "Fuel Subsidy Removal, Prices and Household Financial Health in Nigeria". In: *Asian Journal of Economics and Finance*.

Greenstone, M. and R. Hanna (2014). "Environmental Regulations, Air and Water Pollution, and Infant Mortality in India". In: *American Economic Review* 104.10, pp. 3038–3072.

He, J.-R., G. Tikellis, O. Paltiel, M. Klebanoff, P. Magnus, K. Northstone, et al. (2024). "Association of Common Maternal Infections with Birth Outcomes: A Multinational Cohort Study". In: *Infection* 52.4, pp. 1553–1561.

He, P., D. H. Liu, and G. Q. Zhang (July 1994). “[Effects of high-level-manganese sewage irrigation on children’s neurobehavior]”. chi. In: *Zhonghua Yu Fang Yi Xue Za Zhi [Chinese Journal of Preventive Medicine]* 28.4, pp. 216–218.

Henry, V. D. (Jan. 2011). “Northern Nigeria Maternal, Newborn and Child Health Programme: Selected Analyses from Population-Based Baseline Survey”. en. In: *The Open Demography Journal* 4.1, pp. 11–21.

Hoynes, H., D. Miller, and D. Simon (2015). “Income, the Earned Income Tax Credit, and Infant Health”. In: *American Economic Journal: Economic Policy* 7.1, pp. 172–211.

Hoynes, H. W., M. E. Page, and A. H. Stevens (2011). “Can Targeted Transfers Improve Birth Outcomes? Evidence from the Introduction of the WIC Program”. In: *Journal of Public Economics* 95.7–8, pp. 813–827.

IEA (Feb. 2023). *Fossil Fuel Subsidies in Clean Energy Transitions: Time for a New Approach? – Analysis*. en-GB. Tech. rep. The Internaitonal Energy Agency.

IISD (2023). *Nigeria Fossil Fuel Subsidy Policy Brief*. Tech. rep. International Institute for Sustainable Development.

Immunization Page – Paediatric Association of Nigeria (2026). en-US.

Initative, G. P. E. (May 2023). *2023 A Critical Year For Polio Eradication Efforts In North. . .* en-GB.

Initative, G. P. E. (2026). *Nigeria Introduces The Inactivated Polio Vaccine Into Routi. . .* en-GB.

International Monetary Fund (2024). *Nigeria: Article IV Consultation*. Tech. rep. Washington, DC: International Monetary Fund.

Karagulian, F., C. A. Belis, C. F. C. Dora, A. M. Prüss-Ustün, S. Bonjour, H. Adair-Rohani, and M. Amann (2015). “Contributions to Cities’ Ambient Particulate Matter (PM): A Systematic Review of Local Source Contributions at Global Level”. In: *Atmospheric Environment* 120, pp. 475–483.

Kruk, M. E., E. Goldmann, and S. Galea (2009). “Borrowing and selling to pay for health care in low- and middle-income countries”. eng. In: *Health Affairs* 28.4. Place: Project Hope, pp. 1056–1066.

Majeed, B. (May 2023). “Fuel Subsidy is gone — Tinubu declares”. en-GB. In: *Premium Times*.

Martinez-Bravo, M. and A. Stegmann (Feb. 2022). “In Vaccines We Trust? The Effects of the CIA’s Vaccine Ruse on Immunization in Pakistan”. In: *Journal of the European Economic Association* 20.1, pp. 150–186.

Meriggi, N. F., M. Voors, M. Levine, V. Ramakrishna, D. M. Kangbai, M. Rozelle, E. Tyler, S. Kallon, J. Nabieu, S. Cundy, and A. M. Mobarak (Mar. 2024). “Last-mile delivery increases vaccine uptake in Sierra Leone”. en. In: *Nature* 627.8004. Publisher: Nature Publishing Group, pp. 612–619.

National Bureau of Statistics (2023a). *CPI and Inflation Report December 2023*. Statistical Report. Report 2023c. Abuja: National Bureau of Statistics.

National Bureau of Statistics (2023b). *CPI and Inflation Report May 2023*. Statistical Report. Report 2023b. Abuja: National Bureau of Statistics.

National Bureau of Statistics (2023c). *Transport Fare Watch (June 2023)*. Statistical Report. Report 2023a. Abuja: National Bureau of Statistics.

Newcomer, S. R., S. Y. Michels, A. N. Albers, R. E. Freeman, J. M. Graham, C. L. Clarke, J. M. Glanz, and M. F. Daley (Apr. 2024). “Vaccination Timeliness Among US Children Aged 0-19 Months, National Immunization Survey–Child 2011-2021”. In: *JAMA Network Open* 7.4, e246440.

Nigeria, W. (Feb. 2015). “Nigeria introduces inactivated polio vaccine into routine immunization schedule | WHO | Regional Office for Africa”. en. In.

Nigeria, W. (Aug. 2023). “Protecting every child in Nigeria from poliovirus | WHO | Regional Office for Africa”. en. In.

Noel, M. D., T. Roach, and M. Hill (Sept. 2024). *Is Inflation Relief Passed-Through? Evidence from Excise Taxes*. en. SSRN Scholarly Paper. Rochester, NY.

Okafor, C. E. (Jan. 2021). “Introducing Rotavirus Vaccination in Nigeria: Economic Evaluation and Implications”. In: *PharmacoEconomics Open* 5.3, pp. 545–557.

Okedere, O. O., F. B. Elehinafe, S. Oyelami, and A. O. Ayeni (2021). “Drivers of Anthropogenic Air Emissions in Nigeria: A Review”. In: *Heliyon* 7.3, e06398.

Onyeiwu, S. (May 2023). *Nigeria’s fuel subsidy is gone. It’s time to spend the money in ways that benefit the poor*. en-US.

Ophori, E. A., M. Y. Tula, A. V. Azih, R. Okojie, and P. E. Ikpo (June 2014). “Current Trends of Immunization in Nigeria: Prospect and Challenges”. In: *Tropical Medicine and Health* 42.2, pp. 67–75.

Ozili, P. K. (2023). *Implications of Fuel Subsidy Removal on the Nigerian Economy*. MPRA Paper 120509. Munich Personal RepEc Archive.

Painter, R. C., T. J. Roseboom, and O. P. Bleker (2005). “Prenatal Exposure to the Dutch Famine and Disease in Later Life: An Overview”. In: *Reproductive Toxicology* 20.3, pp. 345–352.

Rentschler, J. (2016). “Incidence and Impact: The Regional Variation of Poverty Effects Due to Fossil Fuel Subsidy Reform”. In: *Energy Policy* 96, pp. 491–503.

Slama, R., L. Darrow, J. Parker, T. J. Woodruff, M. Strickland, M. Nieuwenhuijsen, et al. (2008). “Meeting Report: Atmospheric Pollution and Human Reproduction”. In: *Environmental Health Perspectives* 116.6, pp. 791–798.

Soile, I. and X. Mu (2015). “Who Benefit Most from Fuel Subsidies? Evidence from Nigeria”. In: *Energy Policy* 87, pp. 314–324.

TheGuardian (Apr. 2023). *The fuel subsidy removal debate*. en-GB.

Ukazu, I. (Aug. 2022). “Dealing with diarrhoea: Nigeria introduces rotavirus vaccine into its immunisation plan”. en. In: *GAVI VaccinesWork*.

UNICEF (2026). *Nigeria Immunization Schedule | UNICEF Nigeria*. en.

World Bank (2023). *Nigeria Development Update: Turning the Corner: From Reforms and Renewed Hope, to Results*. Washington, DC: World Bank.

World Bank (2024). *Staying the Course: Progress amid Pressing Challenges. Nigeria Development Update*. Washington, DC: World Bank.

Yahaya, O. A. (2026). “The Impact of Fuel Subsidy Removal on Household Welfare and Transportation Costs in Nigeria: A Cross-Sectional Analysis”. In: *Journal of Economics and Applied Research*.

Zouhar, Y., J. Jellema, N. Lustig, and M. Trabelsi (2021). *Public Expenditure and Inclusive Growth—A Survey of the Evidence*. IMF Working Paper WP/21/83. International Monetary Fund.

A Variable Definitions

- **Home delivery**: indicator equal to one if the child was delivered at the respondent's home or another home, and zero otherwise.
- **C-section**: indicator equal to one if the child was delivered by cesarean section, and zero otherwise.
- **Skilled delivery assistance**: indicator equal to one if the delivery was assisted by a doctor, nurse, or midwife, and zero otherwise.
- **# ANC visits**: reported number of antenatal care visits during pregnancy.
- **Early ANC**: indicator equal to one if the first ANC visit occurred within the first three months of pregnancy, and zero otherwise. Women with no ANC visits are coded as zero.
- **Skilled ANC**: indicator equal to one if antenatal care was provided by a doctor, nurse, or midwife, and zero otherwise.
- **Blood pressure checked**: indicator equal to one if the respondent's blood pressure was measured during antenatal care, and zero otherwise.
- **Urine test**: indicator equal to one if the respondent gave a urine sample during antenatal care, and zero otherwise.
- **Blood test**: indicator equal to one if the respondent gave a blood sample during antenatal care, and zero otherwise.
- **Fetal heartbeat checked**: indicator equal to one if the baby's heartbeat was checked during antenatal care, and zero otherwise.
- **Food counseling**: indicator equal to one if the respondent received counseling on foods to eat during pregnancy, and zero otherwise.

- **Breastfeeding counseling**: indicator equal to one if the respondent received counseling on breastfeeding during antenatal care, and zero otherwise.

- **Bleeding danger-sign counseling**: indicator equal to one if the respondent was informed about vaginal bleeding as a pregnancy danger sign during antenatal care, and zero otherwise.

All ANC content indicators code women with no ANC visits as zero.

B Event Study Specifications

Pollution Event Study

To examine pre-trends and the dynamic response of pollution, we estimate the event-study version of equation (3):

$$\ln(\text{NonDustPM}_{2.5,ct}) = \alpha + \sum_{k \neq -1} \beta_k \Big(\text{Urban}_c \times \mathbf{1}\{Rel_t = k\}\Big) + \lambda_c + \mu_m + \mathbf{X}'_{ct}\delta + \varepsilon_{ct} \quad (9)$$

Rel_t denotes event time in months relative to June 2023, with $Rel_t = 0$ for June 2023. The omitted category is $Rel_t = -1$, corresponding to May 2023, the month immediately before the subsidy removal. The coefficients $\{\beta_k\}$ trace the urban-rural differential in non-dust $PM_{2.5}$ in each event month relative to May 2023. The specification includes month-of-year fixed effects, weather controls, and standard-error clustering as in equation (3).

Birth Outcomes Event Study

For equation (2), the event-study version is:

$$Y_{ict} = \alpha + \sum_{k \neq -1} \beta_k \Big(\text{Urban}_c \times \mathbf{1}\{BirthRel_t = k\}\Big) + \eta_c + \tau_t + \mathbf{X}'_i\gamma + \mathbf{W}'_{ic}\psi + \nu_{ic} \quad (10)$$

where $BirthRel_t$ denotes the birth month relative to June 2023. All other notation follows equation (2).

C Appendix Figures

Figure B1. Raw Monthly Means of $PM_{2.5}$ by Urban Status

Figure B2. Event Study of ln(Non-dust $PM_{2.5}$) with Extended Window

Figure B3. Placebo Test Distributions

D Robustness Tables

Tables C1–C9 are presented in order below. Table C1 reports the descriptive statistics by urban and rural area. Tables C2–C3 report robustness checks on birth size outcomes using alternative clustering levels and month-by-year fixed effects, respectively. Table C4 reports $PM_{2.5}$ results for total $PM_{2.5}$. Tables C5–C6 report robustness checks on the pollution results. Tables C7–C9 report trimester-level results for delivery care, ANC use, and ANC content, respectively.

Table C1: Effects of Fuel Subsidy Removal on Birth Size: Alternative Clustering

	(1) Full preg.	(2) Full preg.	(3) 1st tri.	(4) 2nd tri.	(5) 3rd tri.
Panel A: Very small					
Exposure	0.003***	0.004**	0.009	0.009***	0.008***
	(0.001)	(0.002)	(0.010)	(0.003)	(0.003)
Urban × Exposure	—	−0.001	0.008	0.001	−0.004
		(0.002)	(0.023)	(0.005)	(0.004)
Obs. / Mean	6,579 / 0.041	6,579 / 0.041	6,579 / 0.041	6,579 / 0.041	6,579 / 0.041
Panel B: Small or below					
Exposure	0.006***	0.008***	0.020	0.018***	0.017***
	(0.002)	(0.002)	(0.030)	(0.005)	(0.005)
Urban × Exposure	—	−0.007**	−0.032	−0.012	−0.012**
		(0.003)	(0.038)	(0.007)	(0.006)
Obs. / Mean	6,579 / 0.127	6,579 / 0.127	6,579 / 0.127	6,579 / 0.127	6,579 / 0.127
Panel C: Average or below					
Exposure	0.001	0.005	0.001	0.009	0.011
	(0.003)	(0.004)	(0.042)	(0.009)	(0.008)
Urban × Exposure	—	−0.012**	−0.093	−0.026*	−0.018**
		(0.006)	(0.057)	(0.014)	(0.009)
Obs. / Mean	6,579 / 0.746	6,579 / 0.746	6,579 / 0.746	6,579 / 0.746	6,579 / 0.746
Birth month FE	Yes	Yes	Yes	Yes	Yes
Cluster FE	Yes	Yes	Yes	Yes	Yes
Weather	Yes	Yes	Yes	Yes	Yes

Notes: Standard errors clustered at the DHS cluster level. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

E Vaccination Figures and Tables

NPHCDA

Minimum Target Age of Child	Type of Vaccine	Dosage	Route of Administration	Site
At Birth	BCG	0.05ml	Intra dermal	Left Upper Arm
	*OPV0	2 drops	Oral	Mouth
	**Hep B0 birth	0.5ml	Intramuscular	Anterolateral aspect of Right thigh
6 Weeks	Pentavalent (DPT, Hep B and Hib) 1	0.5ml	Intramuscular	Anterolateral aspect of Left thigh
	Pneumococcal Conjugate Vaccine1	0.5ml	Intramuscular	Anterolateral aspect of right thigh
	*OPV1	2 drops	Oral	Mouth
	***Rota 1	5 drops	Oral	Mouth
	IPV1	0.5ml	Intramuscular	Anterolateral aspect of right thigh (2.5cm apart from PCV)
10 Weeks	Pentavalent (DPT, Hep B and Hib)2	0.5ml	Intramuscular	Anterolateral aspect of left thigh
	Pneumococcal Conjugate Vaccine2	0.5ml	Intramuscular	Anterolateral aspect of right thigh
	*OPV 2	2 drops	Oral	Mouth
	***Rota 2	5 drops	Oral	Mouth
14 Weeks	Pentavalent 3 (DPT, Hep B and Hib)	0.5ml	Intramuscular	Anterolateral aspect of Left thigh
	Pneumococcal Conjugate Vaccine3	0.5ml	Intramuscular	Anterolateral aspect of right thigh
	*OPV 3	2 drops	Oral	Mouth
	***Rota 3	5 drops	Oral	Mouth
	IPV2	0.5ml	Intramuscular	Anterolateral aspect of right thigh (2.5cm apart from PCV)
6 Months	Vitamin A 1st dose	100,000 IU	Oral	Mouth
9 Months	Measles 1st dose	0.5ml	Subcutaneous	Left upper arm
	Yellow Fever	0.5ml	Subcutaneous	Right upper arm
	Meningitis Vaccine	0.5ml	Intramuscular	Anterolateral aspect of Left thigh
12 Months	Vitamin A 2nd dose	200,000 IU	Oral	Mouth
15 Months	Measles 2nd dose (MCV2)	0.5ml	Subcutaneous	Left upper arm
9-13 years	***HPV 6 months interval (2 doses)	0.5ml	Intramuscular	Deltoid muscle (upper arm)

Figure B1: Nigeria Children Immunization Schedule
Source: UNICEF

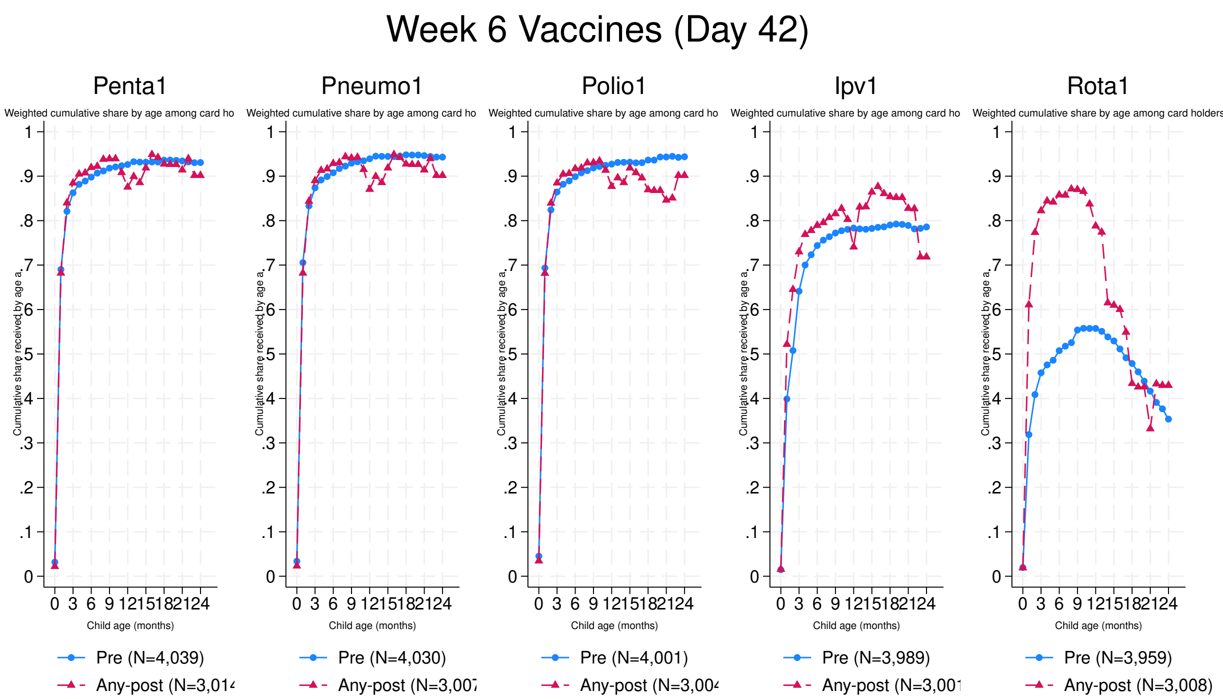

Figure B2: Week-6 Vaccine Age Profiles

Notes: Weighted age-specific observed-risk-set profiles for the week-6 vaccine group. Each point at age a is the share vaccinated by age a among vaccine-card-holder children observed to at least age a. Pre/post groups are defined using the vaccine-specific recommended window. As in the main text, these lines need not be monotone because the risk set changes with age.

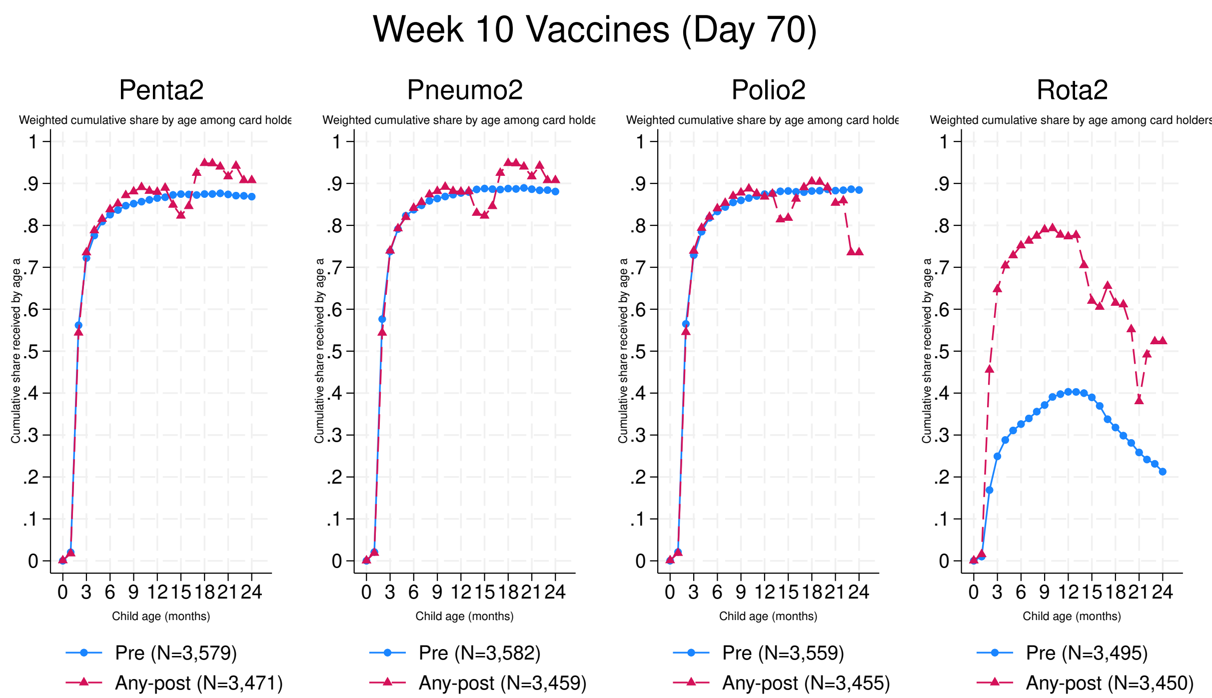

Figure B3: Week-10 Vaccine Age Profiles

Notes: Weighted age-specific observed-risk-set profiles for the week-10 vaccine group. The interpretation is identical to that of Figure 8: later-age flattening or crossing reflects weaker realized receipt among the subset observed to those ages, not a literal decline for a fixed cohort over time.

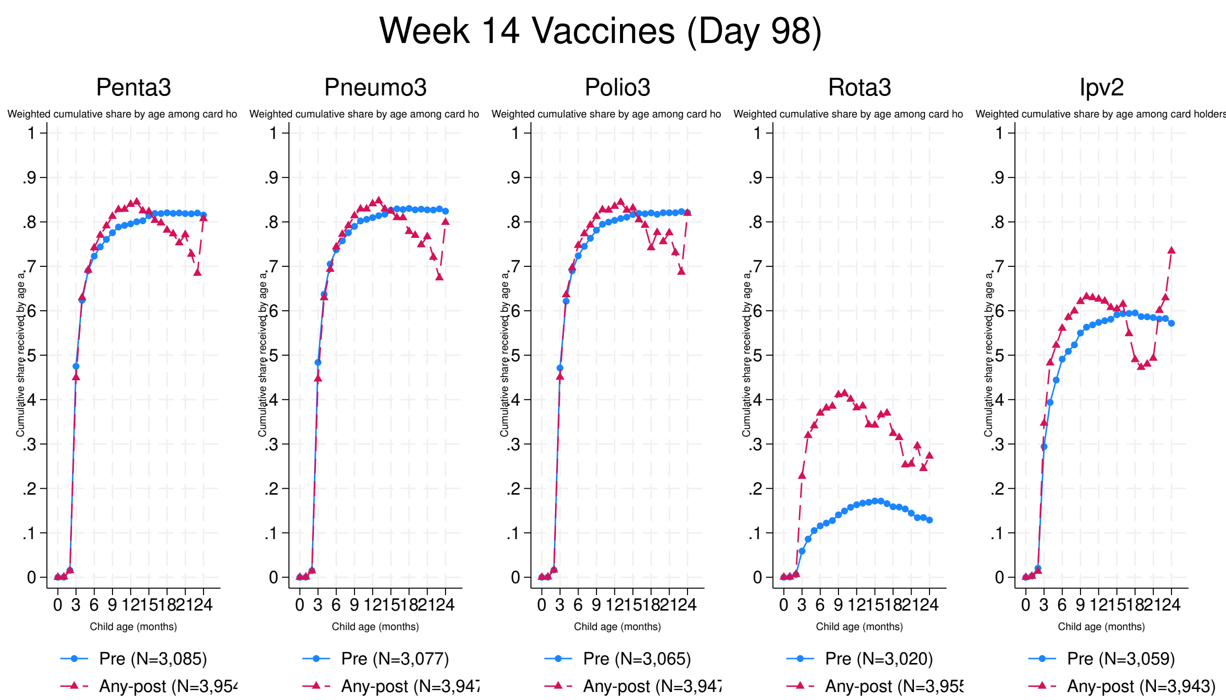

Figure B4: Week-14 Vaccine Age Profiles

Notes: Weighted age-specific observed-risk-set profiles for the week-14 vaccine group. These plots are informative about later-age slippage in the post-window group, but they should not be read as cohort cumulative distribution functions because the denominator changes with age.

F GHS-Panel Summary Statistics

Table C2 reports weighted summary statistics for the GHS-Panel wave-5 cross-section used in Section IV, covering household and individual characteristics and the illness, care-seeking, and visit-level outcomes.

Table C2: GHS-Panel Wave 5 Summary Statistics

	Mean	SD
Wave 5 characteristics		
Age	24.4508	19.7251
Male	0.5076	0.5000
Marital status: 1	0.2860	0.4519
Marital status: 2	0.0552	0.2284
Marital status: 3	0.6587	0.4741
Household size	7.9022	4.2158
Household head reads/writes	0.6920	0.4617
Household head attended school	0.6178	0.4859
Male household head	0.8653	0.3414
Urban	0.3543	0.4783
Main outcomes		
Had illness	0.2371	0.4253
Sought care	0.1869	0.3898
Missed care	0.0560	0.2299
Motorized transport	0.5724	0.4948
Transport cost	521.8383	1221.7444
Wait time (minutes)	28.6622	49.3311
Consultation time (minutes)	28.5599	67.3471
Stopped activity due to illness	0.5561	0.4969
N	27210	

Wave-5 (2023 post-planting visit) cross-section of the full surveyed household population.
Means and standard deviations use analytic weights (`wt_cross_wave5`).