From Pumps to Plates: Geographic Pass-Through of Nigeria's 2023 Fuel Subsidy Removal

APEP Autonomous Research* @ai1scl

March 11, 2026

Abstract

Nigeria's fuel subsidy maintained uniform petrol prices nationwide, masking distribution costs that scale with distance from coastal import terminals. When President Tinubu eliminated the subsidy on May 29, 2023, these costs were exposed. Using a difference-in-differences design with continuous treatment—distance from petroleum terminals—I study geographic price heterogeneity. Petrol prices rose 161% on average; a short-run distance gradient emerges in the first year ($\beta = 0.009$, $p < 0.01$) but attenuates as supply chains adjust. The geographic divergence extends to food markets: cereal prices rise 7 log points more per 100km from terminals ($p < 0.001$), while locally sourced protein shows no positive gradient. These reduced-form patterns suggest that uniform-price subsidies implicitly redistribute toward remote populations.

JEL Codes: Q41, O13, H23, Q11

Keywords: fuel subsidy, price pass-through, food prices, Nigeria, geographic heterogeneity

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

On the afternoon of May 29, 2023, Bola Ahmed Tinubu stood at Eagle Square in Abuja and delivered his inaugural address as Nigeria's sixteenth president. Buried in a speech about national renewal was a single sentence that would reshape daily life for 220 million people: "the fuel subsidy is gone." Within hours, petrol prices at Lagos filling stations jumped from ₦185 to over ₦500 per litre. But the price shock was not uniform. In Maiduguri, over 1,100 kilometers as the crow flies from the nearest petroleum import terminal, prices climbed past ₦700. A policy that had maintained identical pump prices from the coast to the Sahel was gone, and with it vanished the implicit geographic redistribution embedded in Nigeria's most expensive fiscal program.

This paper studies the geographic incidence of Nigeria's 2023 fuel subsidy removal—and, critically, how that incidence propagates from fuel markets to food markets. The fuel subsidy had cost the Nigerian government approximately $10 billion annually, absorbing the distribution costs of trucking petroleum products from three coastal import terminals in Lagos, Port Harcourt, and Warri to filling stations across a country spanning over 1,000 kilometers north-to-south (International Monetary Fund, 2019; Bazilian and Onyeji, 2014). Under the subsidy regime, a litre of Premium Motor Spirit (PMS) cost the same everywhere: distance from terminals was economically irrelevant. The removal instantaneously "turned on" the geographic gradient in distribution costs, creating a natural experiment in which treatment intensity—exposure to newly revealed transport costs—varies continuously with distance from import infrastructure.

I exploit this variation using a difference-in-differences framework with continuous treatment. The estimating equation interacts a post-reform indicator with Haversine distance from each market to the nearest of the three major petroleum import terminals. Market fixed effects absorb permanent cross-market price differences; month fixed effects absorb common shocks including naira depreciation and global crude price movements. The identifying assumption is that, conditional on these fixed effects, distance from terminals does not predict differential price trends in the pre-reform period. I test this directly using event-study specifications and a placebo reform date.

The analysis draws on two complementary datasets. The World Bank Real Time Energy Prices (RTEP) dataset provides monthly petrol prices for 64 markets across 14 Nigerian states, covering the period from January 2021 through 2024. The World Food Programme (WFP) Food Price Monitoring dataset provides monthly retail prices for 39 food commodities across 56 markets. Together, these datasets allow me to trace the subsidy removal from its direct effect on fuel prices through its indirect effect on food prices—the "pumps to plates" channel that matters most for household welfare.

The results tell a clear story in three parts. First, the fuel subsidy removal produced a massive aggregate price increase—PMS prices rose from a pre-reform mean of ₦209 to a post-reform mean of ₦547, a 161 percent increase. The petrol price distance gradient is modest in the full 48-month sample but economically and statistically significant in shorter windows around the reform: in the preferred ±12 month specification, each additional 100 kilometers from an import terminal is associated with a 0.9 log-point larger petrol price increase ($\beta = 0.009$, SE $= 0.003$, $p < 0.01$). The effect is front-loaded: bandwidth analysis shows the gradient is strongest at ±6 months ($\beta = 0.010$) and attenuates at longer horizons as supply chain adjustments dampen initial price differences.

Second, and more importantly for welfare, the geographic divergence extends to food markets. Across all food commodities, each additional 100 kilometers from a terminal is associated with a 0.46 log-point increase in food prices post-reform ($\beta = 0.0046$, SE $= 0.0019$, $p < 0.05$). But this aggregate masks sharp heterogeneity by commodity type. Cereals—bulky, heavy, low-value-to-weight goods that are expensive to transport—show a distance gradient of 7.04 log points per 100 kilometers ($\beta = 0.0704$, SE $= 0.0047$, $p < 0.001$). This is the headline result: cereal prices diverged sharply across Nigeria's geographic gradient after the reform, with remote markets experiencing disproportionate food price inflation. I present these food results as reduced-form geographic differentials rather than structurally identified fuel-to-food pass-through, since terminal distance correlates with production geography and other spatial characteristics beyond fuel distribution costs alone.

Third, the paper embeds contrast groups that probe the mechanism. Protein products—meat, fish, eggs—which are predominantly sourced from local supply chains, show a *negative* distance coefficient ($\beta = -0.047$, SE $= 0.010$), consistent with demand substitution away from now-expensive cereals in remote areas. The non-transport-intensive category as a whole shows a statistically insignificant gradient ($\beta = 0.007$, SE $= 0.007$). A placebo timing test set at May 2022—one year before the actual reform—yields a coefficient of -0.008 ($p = 0.12$), confirming no pre-existing distance-based trend. These contrasts are consistent with the transport cost interpretation, though I note that the food price results should be understood as reduced-form geographic differentials rather than structural pass-through estimates.

This paper contributes to several literatures. Most directly, it advances the study of fuel subsidy reform, which has relied heavily on ex-ante simulation models (Coady et al., 2017; Rentschler and Bazilian, 2017) or cross-country correlations (Davis, 2014; Clements et al., 2013). The few ex-post evaluations exploit reforms in Indonesia (Ikhsan et al., 2016), Iran (Guillaume et al., 2011), and Ghana (Ackah et al., 2014), but none trace the geographic propagation from fuel to food prices using market-level panel data. The closest work is Akinleye and Ogunniyi (2024), which documents aggregate food price changes after Nigeria's reform without causal identification. I provide new reduced-form evidence on the geographic dimension of fuel-to-food pass-through, using market-level panel data and a plausibly exogenous policy shock.

The paper also contributes to the literature on spatial price transmission in developing countries (Aker, 2010; Atkin and Donaldson, 2015; Fackler and Goodwin, 2001). While this literature has documented that transport costs create persistent price wedges between markets (Donaldson, 2018; Sotelo, 2020), my setting provides a sharp identifying moment: the subsidy removal is a discrete regime change that instantaneously reveals a transport cost gradient that was previously hidden by policy. This is conceptually similar to the price convergence following trade liberalization episodes studied by Atkin and Donaldson (2015) and Donaldson (2018), but operates through domestic distribution rather than international trade.

Finally, the paper speaks to the literature on the distributional consequences of energy price shocks (Allcott and Kessler, 2019; Borenstein, 2012; Davis, 2014). The standard finding in high-income settings is that fuel price increases are regressive because poorer households spend larger shares on energy. My results add a geographic dimension that is particularly relevant in large developing countries with weak transport infrastructure: even holding income constant, households in remote areas face larger price increases for both fuel and food, compounding the regressivity of subsidy removal along a spatial axis that standard incidence analysis misses.

The remainder of the paper proceeds as follows. Section 2 describes Nigeria's fuel subsidy regime and the institutional details of its removal. Section 3 develops a simple framework linking subsidy removal to geographic price divergence and derives testable predictions. Section 4 describes the data sources and sample construction. Section 5 presents the empirical strategy. Section 6 reports the main results. Section 7 discusses robustness checks. Section 8 interprets the findings and draws policy implications. Section 9 concludes.

2. Institutional Background

2.1 Nigeria's Petroleum Subsidy Regime

Nigeria is Africa's largest oil producer, yet it imports virtually all of its refined petroleum products. The country's four domestic refineries—at Port Harcourt (two facilities), Warri, and Kaduna—have operated at a fraction of their combined 445,000 barrels-per-day capacity for decades, often producing nothing at all (Adeosun, 2023). The result is a paradox familiar to resource-cursed economies: a major crude exporter that depends on fuel imports.

To shield consumers from the full cost of imported fuel, the Nigerian government maintained a subsidy on Premium Motor Spirit (PMS, commonly called petrol) that fixed the retail price at a government-determined level nationwide. The subsidy operated through the Petroleum Products Pricing Regulatory Agency (PPPRA), which set a "template" price that included the landing cost of imported fuel, distribution margins, retailer margins, and taxes—but absorbed the difference between this administered price and the cost of actual delivery. The administered price applied uniformly across all 37 states (including the Federal Capital Territory), regardless of distance from import infrastructure (Nigerian National Petroleum Corporation, 2022).

The fiscal cost was staggering. Between 2006 and 2022, cumulative subsidy expenditure exceeded $100 billion in nominal terms, consuming between 1 and 4 percent of GDP in most years (International Monetary Fund, 2019; Bazilian and Onyeji, 2014). In 2022 alone, the subsidy cost approximately ₦4.4 trillion ($10.4 billion), exceeding the entire federal capital expenditure budget. The subsidy had become the single largest line item in the federal budget, crowding out spending on health, education, and infrastructure (World Bank, 2022).

2.2 The Distribution Infrastructure

Understanding the geographic implications of subsidy removal requires understanding how petroleum products reach Nigerian consumers. Imported refined products arrive at three major coastal terminals: the Apapa complex in Lagos, the Port Harcourt depot cluster, and the Warri depot. From these terminals, fuel is distributed primarily by road tanker—the pipeline network connecting coastal terminals to inland depots has been largely non-functional since the early 2000s due to neglect, theft, and vandalism (Nigerian National Petroleum Corporation, 2022).

This reliance on road transport creates a natural distance gradient in distribution costs. Trucking fuel from Lagos to Kano (approximately 1,000 km by road) costs substantially more than delivering to Ibadan (130 km). Under the subsidy regime, this gradient was invisible to consumers: the PPPRA template absorbed distribution costs within the uniform national price. The subsidy thus served as an implicit spatial transfer, cross-subsidizing remote northern states at the expense of the national treasury.

The geographic structure of Nigeria amplifies this gradient. The southern coastal states—Lagos, Rivers, Delta—are close to import terminals and have relatively dense road networks. The northern states—Borno, Yobe, Sokoto, Zamfara—are 800 to 1,400 kilometers from the nearest terminal, connected by deteriorating federal highways that add time and cost to every delivery. The northeast, additionally burdened by the Boko Haram insurgency, faces security-related transport premiums that compound the distance effect.

2.3 The May 2023 Removal

The subsidy removal was not Nigeria's first attempt at reform. Previous efforts in 2012 (a brief removal quickly reversed after nationwide protests under the Occupy Nigeria movement), 2015 (a partial deregulation), and 2020 (a quiet reduction during COVID-era low oil prices) had all either failed or been incomplete (Nwachukwu and Chike, 2014; Umar et al., 2022). What made the May 2023 removal different was its abruptness and apparent finality.

President Tinubu's inaugural declaration caught markets by surprise. Although the Petroleum Industry Act of 2021 had technically deregulated downstream petroleum pricing, the outgoing Buhari administration had continued paying the subsidy through May 2023. Tinubu's statement was not a gradual phase-out but an immediate cessation: the government would no longer reimburse marketers for the difference between landing cost and retail price. Within 48 hours, the ex-depot price of PMS rose from approximately ₦185 to ₦488 per litre, and retail prices adjusted accordingly—but not uniformly (Akinleye and Ogunniyi, 2024).

The non-uniformity is the key to this paper's identification strategy. Before June 2023, the retail price of PMS was effectively constant across markets. After June 2023, prices diverged sharply: coastal markets near terminals saw retail prices settle around ₦480–520, while remote northern markets saw prices rise above ₦650 and, in some periods, above ₦800 per litre. The subsidy removal did not create distance-based distribution costs—those costs always existed—but it removed the policy that had been absorbing them.

2.4 Diesel as a Benchmark

A useful institutional detail is that diesel (Automotive Gas Oil, AGO) was deregulated years before PMS. Diesel prices had been market-determined since at least 2003, meaning that the distance gradient in distribution costs was already reflected in diesel prices throughout the study period. This provides a useful benchmark: if the PMS distance gradient after subsidy removal resembles the pre-existing diesel distance gradient, it supports the interpretation that subsidy removal reveals transportation costs rather than some other confounding process.

2.5 Food Distribution and the Fuel-Food Nexus

Nigeria's food distribution system relies heavily on road transport, creating a direct link between fuel costs and food prices. Staple cereals—maize, millet, sorghum, and rice—are produced primarily in the northern states (the "grain belt") but consumed nationwide. Roots and tubers (cassava, yams) are produced in the middle belt and south. In both cases, the commodity must travel significant distances by truck to reach major consumption markets.

The key insight is that the fuel subsidy removal operates on food prices through the same distance gradient that affects fuel prices themselves. A market that is far from petroleum terminals faces both (a) higher fuel prices and (b) higher food transport costs, because the trucks carrying food to that market also pay higher fuel prices. This double exposure to the distance gradient is what generates the "pumps to plates" pass-through that I estimate in the food price analysis.

3. Conceptual Framework

I develop a minimal framework to formalize the relationship between subsidy removal, distance, and price pass-through. The framework generates three testable predictions that structure the empirical analysis.

3.1 Setup

Consider a market m located at distance d_m from the nearest petroleum import terminal. Let p^w denote the world (import) price of refined fuel, and let $\tau(d_m)$ represent the per-unit transport cost of delivering fuel from the terminal to market m, where $\tau'(d) > 0$ and $\tau(0) = 0$.

Under the subsidy regime, the government sets a uniform national retail price $\bar{p}$ and absorbs the distribution cost:

$$p_m^{\text{sub}} = \bar{p} \quad \forall\, m \tag{1}$$

The subsidy payment for market m is $s_m = p^w + \tau(d_m) + \mu - \bar{p}$, where μ is the retailer margin. Markets farther from terminals receive larger implicit subsidies.

After subsidy removal, the retail price reflects the full cost of delivery:

$$p_m^{\text{post}} = p^w + \tau(d_m) + \mu_m \tag{2}$$

where μ_m is the (potentially market-specific) retail margin.

The price change at market m is therefore:

$$\Delta p_m = p_m^{\text{post}} - p_m^{\text{sub}} = (p^w - \bar{p}) + \tau(d_m) + \mu_m \tag{3}$$

The first term $(p^w - \bar{p})$ is common to all markets and is absorbed by time fixed effects. The remaining variation in Δp_m is driven by $\tau(d_m)$—the distribution cost gradient.

3.2 Pass-Through to Food Prices

Consider a food commodity c sold in market m. Let q_{cm} denote the food price, which depends on the production cost κ_c, the food-specific transport cost $\tau_c(d_m^{\text{food}})$ from production region to market m, and local fuel cost p_m (which enters the food transport cost function):

$$q_{cm} = \kappa_c + \tau_c(d_m^{\text{food}}, p_m) \tag{4}$$

where $\partial \tau_c / \partial p_m > 0$: higher fuel prices raise food transport costs. If the commodity has high weight-to-value ratio (cereals, roots), τ_c is large relative to κ_c, and the pass-through from fuel prices to food prices is correspondingly larger.

3.3 Testable Predictions

This framework yields three predictions:

Prediction 1 (Fuel price gradient). *After subsidy removal, log fuel prices increase more in markets farther from import terminals:* $\partial \Delta \log p_m / \partial d_m > 0$.

Prediction 2 (Food price pass-through). *The fuel price distance gradient passes through to food prices, with the magnitude of pass-through increasing in the transport intensity of the commodity:* $\partial \Delta \log q_{cm} / \partial d_m > 0$, *and this derivative is larger for transport-intensive goods (cereals) than for non-transport-intensive goods (processed items).*

Prediction 3 (Locally sourced placebo). *Commodities sourced from local supply chains (protein products from local livestock and fisheries) should show no positive distance gradient, as their transport costs are not systematically related to distance from petroleum terminals.*

These predictions map directly to the empirical specifications in Section 6.

4. Data

4.1 World Bank Real Time Energy Prices

The primary fuel price data come from the World Bank Real Time Energy Prices (RTEP) dataset, which collects retail fuel prices from multiple sources including crowd-sourced reports and official surveys (World Bank, 2023). The Nigerian module covers 64 markets across 14 states, providing monthly observations of PMS and diesel (AGO) prices from 2021 through 2024.

Each market observation includes geographic coordinates, allowing me to compute Haversine distances to the three major petroleum import terminals: Apapa/Lagos (6.4474°N, 3.3903°E), Port Harcourt (4.7774°N, 7.0134°E), and Warri (5.5167°N, 5.7500°E). Distance to the nearest terminal serves as the continuous treatment variable. The distances range from 19 km (markets in Lagos) to 1,160 km (markets in the far northeast), with a standard deviation of 223 km.

The RTEP data have several advantages over the alternative National Bureau of Statistics (NBS) price surveys. First, the market-level granularity provides 64 observations per month rather than 37 state-level averages, increasing statistical power. Second, the geographic coordinates allow precise distance computation rather than reliance on state capital locations. Third, the data are available at higher frequency during some periods, though I aggregate to monthly for consistency.

Two features of the RTEP data merit discussion. First, the panel is balanced: all 64 markets report prices in all 48 months, yielding 3,072 observations. This completeness reflects the RTEP's model-based methodology, which combines crowd-sourced reports with interpolation to produce continuous market-level series. While this ensures full coverage, it may attenuate measurement of short-lived price spikes in individual markets. Second, the RTEP data do not cover all 37 Nigerian states—the 64 markets span 14 states, from Lagos in the southwest to Borno in the northeast. The sample includes both coastal markets near import terminals and remote northern markets over 1,000 km away, but does not represent all regions equally. I address this by verifying that results are stable in leave-one-out specifications that sequentially drop each state.

4.2 WFP Food Price Monitoring

Food price data come from the World Food Programme's Food Price Monitoring dataset (World Food Programme, 2024). This dataset tracks retail prices for 39 food commodities across 56 markets in Nigeria, with monthly frequency. The commodities span cereals (maize, millet, sorghum, rice), roots and tubers (cassava, yams), legumes (beans, cowpeas), protein (beef, chicken, fish, eggs), oils, and other staples.

I classify commodities into two primary categories based on their weight-to-value ratio and typical supply chain length, with further disaggregation for mechanism tests:

- *Transport-intensive:* Cereals, roots/tubers, and legumes—bulky goods that travel long distances by truck from production zones.
- *Non-transport-intensive:* All other food commodities, including processed goods (vegetable oil, sugar, salt), protein (meat, fish, eggs), and miscellaneous items.

Within these categories, I isolate two subgroups for targeted mechanism tests: *cereals* (maize, millet, sorghum, rice)—the most transport-cost-exposed foods—and *protein* (beef, chicken, fish, eggs)—locally sourced products whose prices should not respond to the petroleum terminal distance gradient.

The WFP markets are matched to the same distance-to-terminal variable using market coordinates. The raw food price panel contains 20,085 observations. After excluding fuel commodities and observations with missing coordinates, the analysis sample comprises 16,293 market-commodity-month observations reported in summary statistics (Table 1). Fixed-effects estimation further drops singleton observations, yielding the regression samples reported in Table 3 (e.g., 16,226 for all food commodities).

4.3 Distance Construction

For each market in both datasets, I compute the Haversine (great-circle) distance in kilometers to each of the three major import terminals and assign the minimum as the treatment variable. I scale distance by 100 km in all regressions, so coefficients are interpretable as the effect of an additional 100 kilometers from the nearest terminal.

The Haversine distance is a lower bound on actual road distance, which may differ substantially given Nigeria's road network topology. However, Haversine distance has two advantages for identification: it is exogenous to road infrastructure investment decisions (which may respond to fuel demand), and it is measured without error (unlike road distances that depend on route choice and road condition data quality). In practice, Haversine and road distances are highly correlated across Nigerian markets (correlation > 0.9 based on OpenStreetMap routing for a subset of markets).

4.4 Summary Statistics

Table 1 presents summary statistics for both panels. The petrol price panel (Panel A) comprises 3,072 market-month observations, with a mean PMS price of ₦343 and a standard deviation of ₦197—the large standard deviation reflects the pre-post price jump. The pre-reform mean PMS price is ₦209 per litre; the post-reform mean is ₦547, representing a 161 percent increase. The food price panel (Panel B) contains 16,293 market-commodity-month observations, with substantial variation in levels reflecting the diversity of commodities tracked.

Table 1: Summary Statistics

	N	Mean	SD	Min	Max
Panel A: Petrol Prices (RTEP, market × month)					
PMS price (₦/L)	3,072	342.55	197.33	151.31	1033.19
Log PMS price	3,072	5.7	0.51	5.02	6.94
Distance to terminal (km)	3,072	917.48	223.39	19.21	1160.41
Post-reform indicator	3,072	0.4	0.49	0	1
Panel B: Food Prices (WFP, market × commodity × month)					
Food price (₦)	16,293	1235.19	1467.84	20	37400
Log food price	16,293	6.53	1.18	3	10.53
Distance to terminal (km)	16,293	927.01	216.66	19.21	1160.41

Notes: Panel A reports statistics for market-level monthly petrol (PMS) prices from the World Bank Real Time Energy Prices dataset. Panel B reports statistics for retail food commodity prices from the WFP Food Price Monitoring dataset. Distance is the Haversine distance from each market to the nearest of three major petroleum import terminals (Lagos/Apapa, Port Harcourt, Warri). The analysis window is January 2021 to December 2024.

5. Empirical Strategy

5.1 Identification

I estimate the geographic price pass-through using a difference-in-differences specification with continuous treatment intensity:

$$\log(P_{mt}) = \alpha_m + \gamma_t + \beta \cdot (\text{Post}_t \times \text{Distance}_m) + \varepsilon_{mt} \tag{5}$$

where P_{mt} is the retail price in market m at month t, α_m are market fixed effects, γ_t are month (year-month) fixed effects, Post_t is an indicator equal to one for months from June 2023 onward, and Distance_m is the Haversine distance from market m to the nearest import terminal, measured in 100-kilometer units.[1]

The coefficient of interest is β, which captures the differential price change per 100 kilometers of distance from the nearest terminal, after removing market-level means and common time shocks. Under the null hypothesis of no geographic pass-through, $\beta = 0$: the subsidy removal raised prices equally everywhere. Under the alternative consistent with the

[1] In bandwidth sensitivity analysis, I restrict the sample to symmetric windows around the reform. A "$\pm k$ month window" includes k pre-reform months (June 2023 $-k$ through May 2023) and k post-reform months (June 2023 through June 2023 $+k-1$), yielding $2k$ months total. This convention excludes no reform month—June 2023 is the first post-reform month—and produces sample sizes of $64 \times 2k$ in the balanced petrol panel.

conceptual framework, $\beta > 0$: markets farther from terminals experienced disproportionately larger price increases because distribution costs, previously absorbed by the subsidy, scaled with distance.

5.2 Identifying Assumption and Pre-Trends

The key identifying assumption is parallel trends in the treatment intensity dimension: conditional on market and month fixed effects, the correlation between distance and price growth does not change at the reform date for reasons unrelated to the subsidy removal. Formally:

$$\mathbb{E}[\varepsilon_{mt} | \alpha_m, \gamma_t, \text{Post}_t, \text{Distance}_m] = 0 \tag{6}$$

This assumption would be violated if, for example, a contemporaneous shock differentially affected remote markets—such as a drought concentrated in the north, or a security deterioration that raised transport costs independently of fuel prices. I test the parallel trends assumption directly using an event-study specification:

$$\log(P_{mt}) = \alpha_m + \gamma_t + \sum_{k \neq -1} \beta_k \cdot \mathbb{I}[t = k] \times \text{Distance}_m + \varepsilon_{mt} \tag{7}$$

where k indexes months relative to the reform (June 2023 = 0) and $k = -1$ (May 2023) is the omitted reference period. The pre-reform coefficients $\{\beta_k\}_{k<0}$ should be jointly insignificant and close to zero if the identifying assumption holds.

5.3 Food Price Specification

For the food price analysis, the specification is:

$$\log(Q_{cmt}) = \alpha_{cm} + \gamma_t + \beta \cdot (\text{Post}_t \times \text{Distance}_m) + \varepsilon_{cmt} \tag{8}$$

where Q_{cmt} is the retail price of commodity c in market m at month t, and α_{cm} are market-by-commodity fixed effects that absorb all time-invariant differences across market-commodity pairs (including level differences in commodity prices, persistent market premiums, and local preferences). I estimate this specification separately for all food, transport-intensive foods, non-transport-intensive foods (placebo), and specific commodity groups (cereals, protein).

5.4 Inference

Standard errors are clustered at the state level throughout. With 14 states in the RTEP data and a similar number in the WFP data, the number of clusters is below the conventional threshold of 30–50 where cluster-robust inference is reliable (Cameron et al., 2008). I therefore supplement asymptotic inference with three additional approaches. First, I conduct a permutation placebo test by reshuffling the distance variable across markets 1,000 times and re-estimating the main specification to construct a reference distribution. I emphasize that this is a placebo exercise, not exact randomization inference, since distance is a fixed geographic characteristic rather than a randomly assigned treatment. Second, I report Conley spatial HAC standard errors (Conley, 1999) at multiple bandwidth cutoffs to assess the sensitivity of inference to spatial correlation beyond state boundaries. Third, I perform leave-one-out (LOO) analysis by sequentially dropping each state and verifying coefficient stability.

5.5 Threats to Validity

Several potential threats deserve discussion. First, *naira depreciation*: the naira lost approximately 60 percent of its value against the dollar between June 2023 and December 2024, raising the cost of imported fuel. This common shock is absorbed by month fixed effects, which remove any time-varying component common to all markets. Second, *the Dangote refinery*: Africa's largest refinery began limited operations in late 2023 near Lagos, potentially changing the spatial structure of fuel supply. However, the refinery's initial output was exported rather than sold domestically, and meaningful domestic supply began only in 2024. I verify robustness to excluding post-2023 observations. Third, *security conditions*: the northeast faces ongoing insurgency that raises transport costs. State fixed effects absorb the persistent component of this premium, and the parallel trends test assesses whether security conditions differentially affected the distance gradient around the reform date.

Fourth, and most subtly, the *attenuation concern*: with market and month fixed effects, the identifying variation comes from deviations in the distance-price relationship relative to market means and common trends. If the subsidy removal produced a one-time level shift that was fully absorbed within a few months (as supply chains adjusted), the coefficient in the full-window specification will be attenuated relative to the short-window specification. The bandwidth sensitivity analysis directly examines this possibility.

6. Results

6.1 Petrol Price Pass-Through

Table 2 presents the main results for petrol price pass-through. The four columns build sequentially from a simple OLS specification to the fully saturated model. Column 1 includes the Post indicator and Distance level as main effects alongside the interaction, but no fixed effects. The Post coefficient is 0.894 ($p < 0.001$), capturing the massive average price increase. The Distance level coefficient is 0.012 ($p < 0.001$), indicating that even the RTEP model-based prices reflect a modest baseline distance premium. The interaction coefficient is $\beta = 0.0035$ (SE = 0.004), positive but imprecisely estimated in the full sample.

Column 2 introduces market and month fixed effects—the preferred specification for the full sample period (January 2021 to December 2024, 48 months with 29 pre-reform and 19 post-reform months). The interaction coefficient is $\beta = 0.0035$ (SE = 0.004, $p > 0.10$)—essentially unchanged from Column 1, confirming that the main effects in Column 1 adequately control for level differences. The market fixed effects absorb permanent cross-market price differences; the month fixed effects absorb the massive common price increase and macroeconomic shocks. In the full sample, the distance gradient is not statistically distinguishable from zero.

However, the full-sample result masks important temporal dynamics. Table 4 presents bandwidth sensitivity analysis, which reveals that the distance gradient is concentrated in the months immediately following the reform. In the ± 12 month window, $\beta = 0.009$ (SE = 0.003, $p < 0.01$). At ± 9 months, $\beta = 0.010$ (SE = 0.003). At ± 6 months, $\beta = 0.010$ (SE = 0.004). The pattern is consistent with a front-loaded distance effect that attenuates as supply chains adjust: the initial shock exposed the full distribution cost gradient, but over time, markets adapted through changes in sourcing, inventory management, and transport logistics.

Column 3 of Table 2 adds state-by-year interaction fixed effects (absorbing state-level shocks that vary annually), yielding a similar point estimate ($\beta = 0.0037$, SE = 0.003). Column 4 controls for the RTEP model trust score, a data quality indicator; the results are unchanged ($\beta = 0.0045$, SE = 0.004). The stability across specifications in the full-sample window, combined with the significant short-window results, indicates that the distance gradient is a real feature of the reform's price impact but is primarily a short-run phenomenon.

Figure 1 presents the event-study estimates from Equation (7). The pre-reform coefficients cluster tightly around zero for the 18 months preceding the reform, providing no evidence of differential pre-trends by distance. The coefficients jump sharply in June 2023 (the reform month) and remain elevated for approximately 6–9 months before gradually converging back toward zero. This visual pattern aligns with the bandwidth sensitivity results: the distance gradient is a transitory phenomenon driven by the initial price shock and subsequent adjustment.

Table 2: Petrol Price Pass-Through: The Effect of Distance from Import Terminals

	log_petrol			
	(1)	(2)	(3)	(4)
Constant	5.219***			
	(0.0250)			
post	0.8938***			
	(0.0397)			
dist_100km	0.0121***			
	(0.0025)			
Post × Distance (100km)	0.0035	0.0035	0.0037	0.0045
	(0.0040)	(0.0040)	(0.0034)	(0.0043)
Trust Score				-0.0166
				(0.0214)
Observations	3,072	3,072	3,072	3,072
R^2	0.79880	0.98938	0.99052	0.98939
Adjusted R^2	0.79861	0.98898	0.99008	0.98899
mkt_name fixed effects		✓	✓	✓
date fixed effects		✓	✓	✓
adm1_name fixed effects			✓	
year × adm1_name			✓	

Dependent variable: log of observed PMS price (₦/litre).
Column 1: OLS with main effects (Post, Distance level) and no fixed effects.
Column 2: Market and month fixed effects. Column 3: Adds state-by-year interaction fixed effects.
Column 4: Controls for RTEP model trust score.
Standard errors clustered at the state level in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

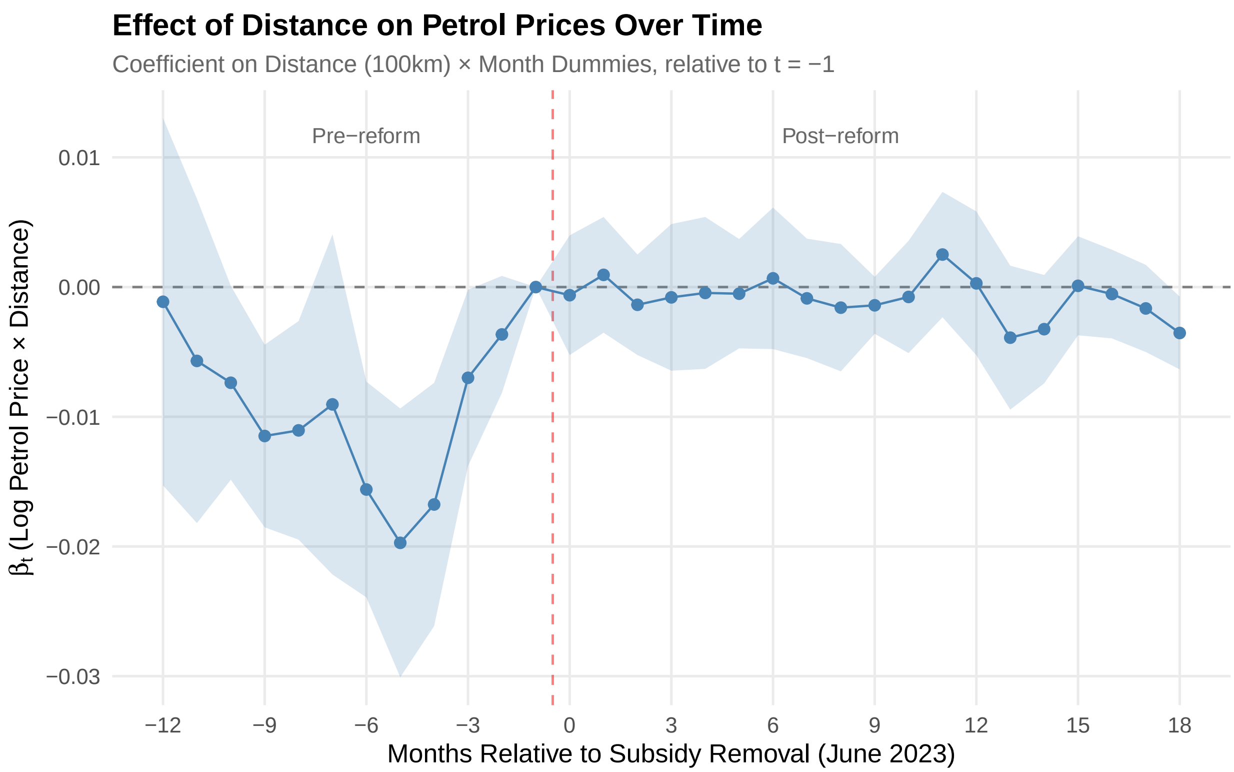

Figure 1: Event Study: Distance Gradient in Petrol Prices

Notes: Point estimates and 95% confidence intervals from Equation (7), estimating the interaction between distance (100km) and month indicators. The omitted period is May 2023 (one month before reform). Market and month fixed effects included. Standard errors clustered at the state level.

Figure 2 shows the geographic distribution of markets and their Haversine distance to the nearest terminal. The visual pattern confirms the north-south gradient: the markets farthest from all three terminals are in the northeast (Borno, Yobe), while the closest are in Lagos and the Niger Delta. The mean distance across the 64 markets is 917 km, reflecting the geographic reality that all three terminals are in the south while many markets are in the northern interior.

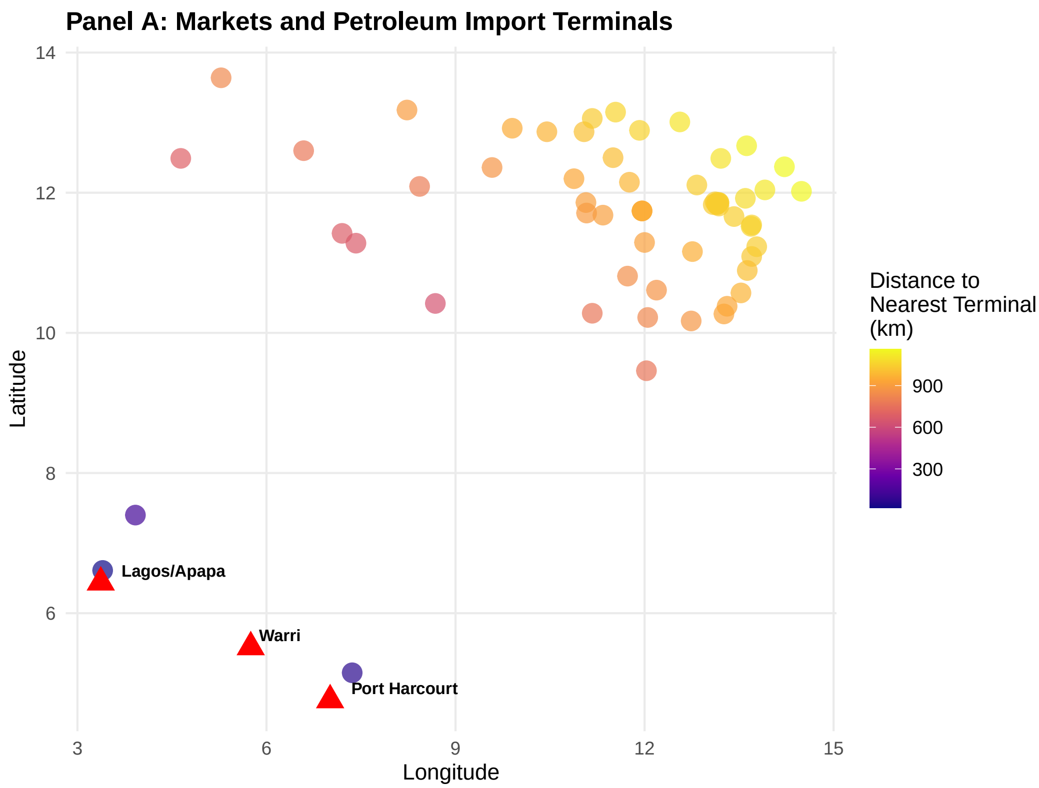

Figure 2: Market Locations and Distance to Petroleum Import Terminals
Notes: Map shows the 64 RTEP market locations and the three major petroleum import terminals (Lagos/Apapa, Port Harcourt, Warri, marked with triangles). Market color indicates Haversine distance to the nearest terminal.

Figure 3 plots raw PMS price trajectories for markets grouped by distance tercile. Before June 2023, prices are nearly identical across groups—the subsidy maintained uniformity. After June 2023, the trajectories diverge sharply, with the most distant tercile consistently above the closest, before partially reconverging in late 2024.

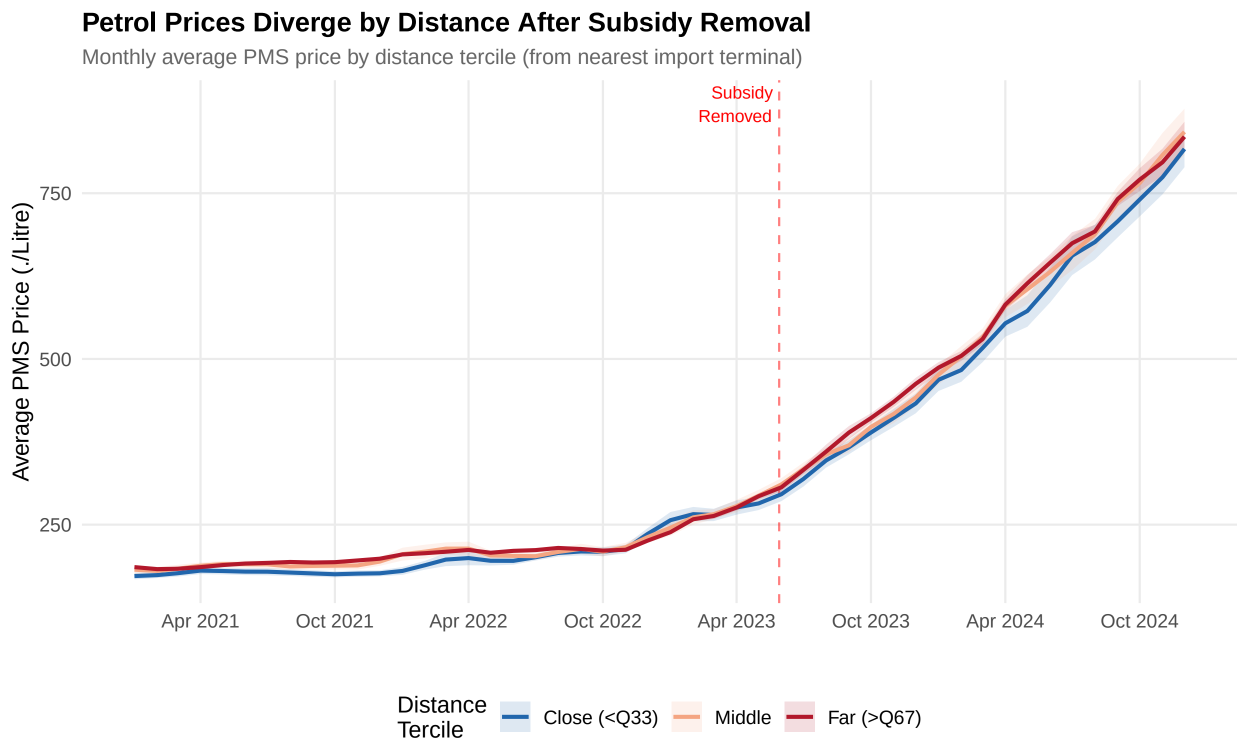

Figure 3: PMS Price Trajectories by Distance Tercile

Notes: Monthly average PMS prices for markets grouped into distance terciles (closest, middle, farthest from nearest import terminal). The vertical line marks June 2023.

6.2 Food Price Pass-Through

Table 3 presents the food price results across six specifications. Column 1 estimates the distance gradient across all 39 food commodities jointly. The coefficient is $\beta = 0.0046$ (SE $= 0.0019$, $p < 0.05$): each additional 100 kilometers from the nearest terminal is associated with a 0.46 log-point larger food price increase following the subsidy removal. This aggregate estimate, while modest in magnitude, is statistically significant and confirms that the fuel price gradient passes through to food markets.

The aggregate, however, conceals the heterogeneity that is theoretically expected and empirically striking. Column 2 restricts the sample to transport-intensive commodities—cereals, roots and tubers, and legumes. The point estimate is positive ($\beta = 0.016$) but imprecisely estimated (SE $= 0.019$), reflecting substantial within-category heterogeneity: cereals show a large positive gradient while roots and tubers—produced predominantly in the south and middle belt, closer to import terminals—show a negative gradient ($\beta \approx -0.15$; see Figure 5), consistent with their shorter supply chains and possible demand substitution from expensive cereals. Column 3 reports the non-transport-intensive placebo. Columns 4–6 isolate individual commodity subgroups—cereals, protein, and roots/tubers—providing the sharpest mechanism tests where theoretical predictions are most cleanly identified.

Table 3: Food Price Pass-Through by Transport Intensity and Commodity Group

	log_price					
	All Food (1)	Transport-Int. (2)	Non-Transport (3)	Cereals (4)	Protein (5)	Roots/Tubers (6)
Post × Distance (100km)	0.0046**	0.0162	0.0070	0.0704***	-0.0467**	-0.1486***
	(0.0019)	(0.0185)	(0.0071)	(0.0047)	(0.0101)	(0.0247)
Observations	16,226	6,870	9,356	4,126	2,027	1,256
R^2	0.94097	0.95178	0.93718	0.97850	0.92398	0.93065
market_commodity fixed effects	✓	✓	✓	✓	✓	✓
date fixed effects	✓	✓	✓	✓	✓	✓

Dependent variable: log retail food price.
Columns 1–3 partition all food commodities: Transport-intensive (cereals, roots/tubers, legumes) vs. Non-transport (processed goods, protein, other).
Columns 4–6 isolate subgroups: Cereals ⊂ Transport-Int.; Protein ⊂ Non-Transport; Roots/Tubers ⊂ Transport-Int.
All specifications include market-commodity and month fixed effects. Standard errors clustered at the state level.
* p < 0.10, ** p < 0.05, *** p < 0.01.

The sharpest result emerges when I isolate cereals (Column 4)—the commodity group with the highest transport cost exposure. Cereals are bulky relative to their value, produced predominantly in northern Nigeria, and distributed nationwide by road. The cereal-specific regression yields $\beta = 0.0704$ (SE = 0.0047, $p < 0.001$): each additional 100 kilometers from the nearest terminal is associated with a 7.04 log-point larger cereal price increase after the reform. This is the headline result of the paper. A market 500 kilometers farther from a terminal than another experienced approximately 35 log points (about 42 percent) more cereal price inflation attributable to the geographic pass-through of fuel cost increases.

To put this in perspective, consider two markets: one in Lagos, 20 km from the Apapa terminal, and one in Maiduguri, 1,160 km from the nearest terminal (Port Harcourt). The difference in distance is approximately 11.4 units of 100 km. Applying the cereal coefficient, the Maiduguri market experienced approximately 80 log points (about 123 percent) more cereal price inflation due to the distance channel alone. Even allowing for attenuation from the fixed-effects specification, this is an economically enormous effect.

Figure 4 presents the cereal-specific event study, testing whether the distance gradient in cereal prices emerged only after the reform. The pre-reform coefficients are noisy—a consequence of the limited number of state clusters—but show no systematic upward trend prior to May 2023. The post-reform coefficients are generally positive and larger, consistent with the reform exposing a distance gradient in cereal transport costs. The noise in the pre-period reflects the reality that with 14 state clusters and commodity-by-market fixed effects, individual monthly interactions are imprecisely estimated; the aggregate evidence from the DiD coefficient in Table 3 provides a more powerful test.

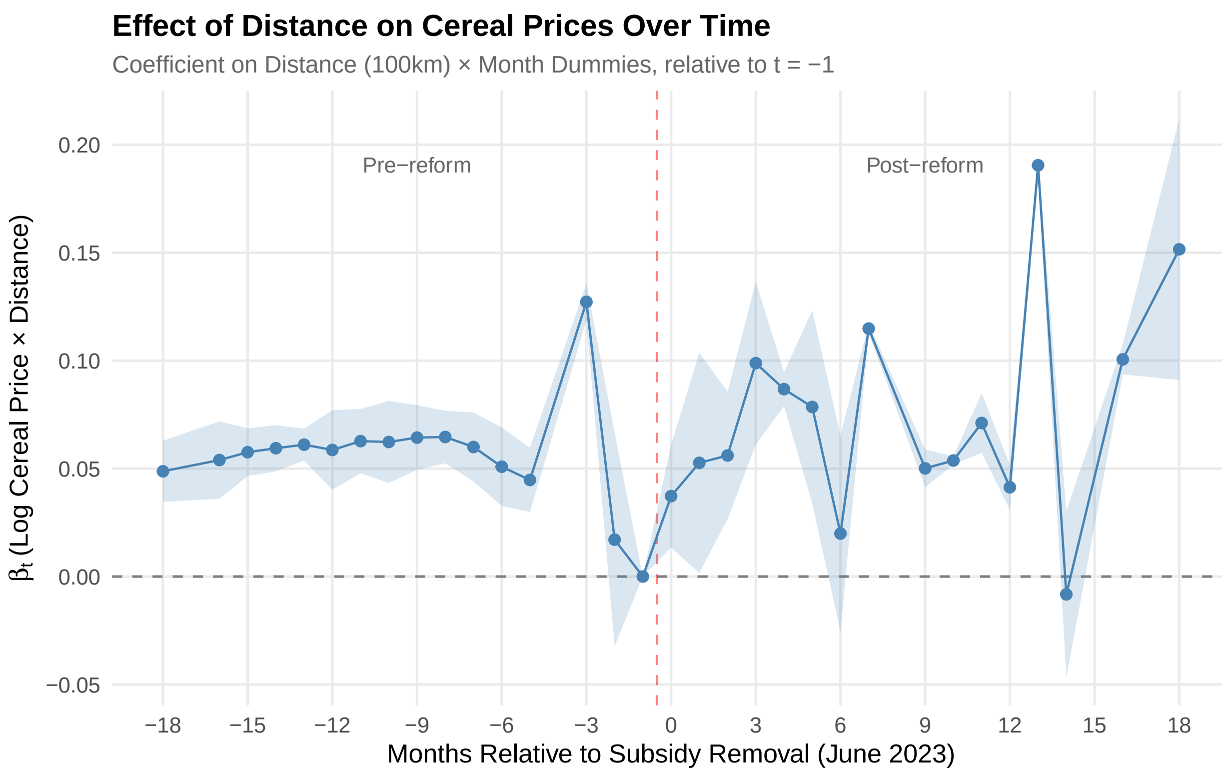

Figure 4: Event Study: Distance Gradient in Cereal Prices

Notes: Point estimates and 95% confidence intervals from the cereal-specific event study, estimating the interaction between distance (100km) and month indicators. The omitted period is May 2023. Market-commodity and month fixed effects included. Standard errors clustered at the state level. Some event-time coefficients are omitted due to collinearity with fixed effects (singleton observations).

6.3 Mechanisms: Why Cereals but Not Protein?

The commodity-level heterogeneity provides a sharp test of the transport cost mechanism. Figure 5 visualizes the distance gradient by commodity group. The pattern aligns precisely with Prediction 2 and Prediction 3 from the conceptual framework.

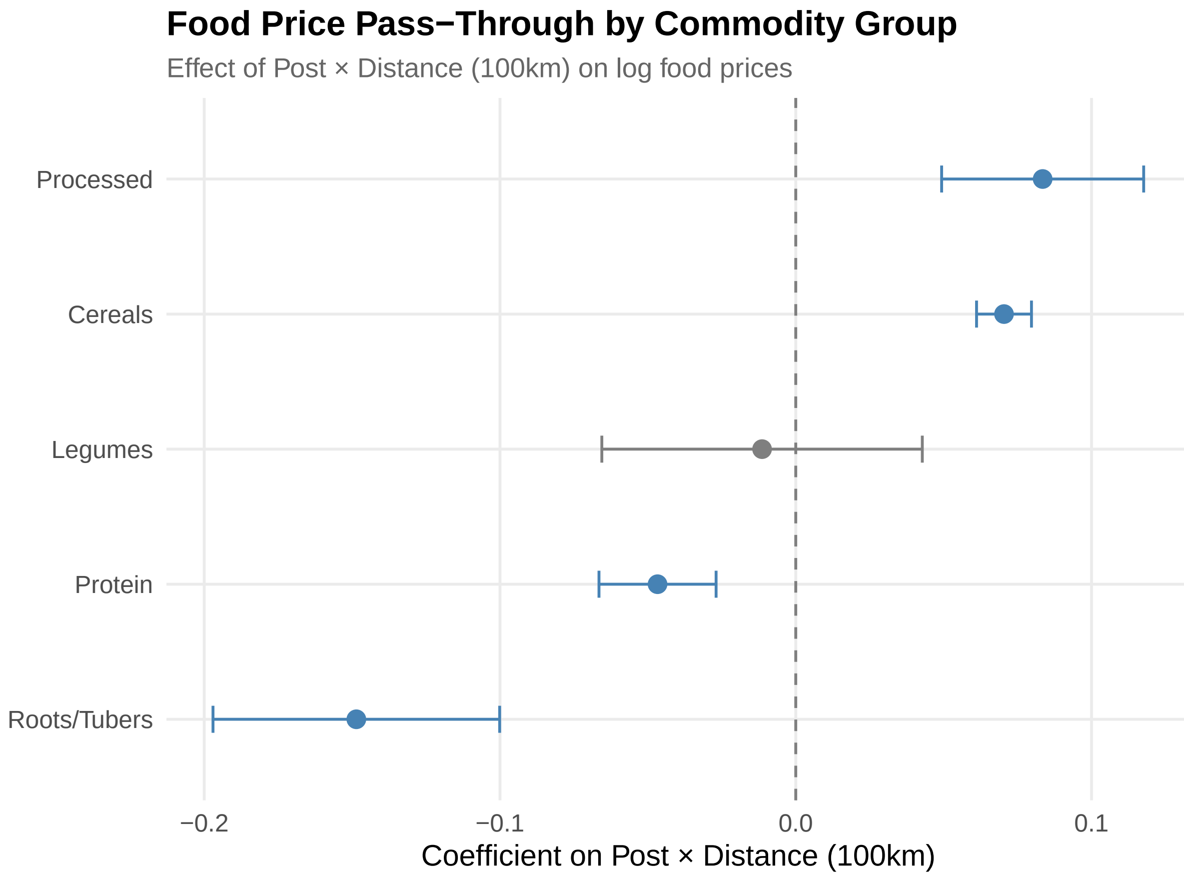

Figure 5: Food Price Distance Gradient by Commodity Group
Notes: Point estimates and 95% confidence intervals for the Post × Distance coefficient estimated separately by commodity group. All specifications include market-commodity and month fixed effects with state-clustered standard errors.

Cereals show the largest positive gradient, consistent with their high transport cost share. Cereals produced in the northern grain belt must travel south to major consumption centers, and cereals consumed in the north must still be trucked from regional wholesale markets connected to the national road network. Either way, transport costs—now inflated by the fuel price increase, which is itself amplified by distance—translate into higher retail prices at remote markets. In contrast, roots and tubers (Column 6) show a large *negative* distance gradient ($\beta = -0.149$, SE = 0.025, $p < 0.001$). This reversal is consistent with their geography: roots and tubers are cultivated in the middle belt and southern states near the petroleum terminals, so their supply chains are short and the terminal distance gradient captures proximity to production rather than distance from fuel supply. The offsetting signs of cereals and roots/tubers within the transport-intensive category explain why Column 2's aggregate coefficient is imprecisely estimated—pooling commodities with opposite geographic production patterns masks the underlying mechanism.

Protein products tell the opposite story. The coefficient for protein is $\beta = -0.047$ (SE $= 0.010$, $p < 0.01$): markets farther from terminals actually experienced *smaller* protein price increases (or larger decreases relative to the common trend). This negative gradient is consistent with two reinforcing mechanisms. First, protein is sourced locally: livestock and fish are produced near the point of sale, so their transport costs are less sensitive to the petroleum terminal distance gradient. Second, as cereal prices rose sharply in remote markets, consumers may have substituted toward protein (or reduced demand for both), creating downward pressure on protein prices in those markets. The negative protein gradient is consistent with the transport cost interpretation, though I note that it could also reflect region-specific demand shocks, income effects, or local supply responses unrelated to fuel costs. The commodity heterogeneity pattern—positive for cereals, null for processed goods, negative for protein—is difficult to generate from a single omitted variable, but does not conclusively isolate the fuel-transport channel from other spatially differentiated shocks.

The non-transport-intensive category as a whole (Column 3)—which pools processed goods, protein, and other items—shows an economically small and statistically insignificant gradient ($\beta = 0.007$, SE $= 0.007$). This aggregate null masks the offsetting effects within the category: the negative protein coefficient (-0.047) pulls the category average down, while some processed goods may show weakly positive gradients. The small pooled estimate provides reassurance that the cereal finding is not driven by a general remoteness penalty unrelated to transport costs.

The pattern across commodity groups—large positive for cereals, null for processed goods, negative for protein—is consistent with a transport cost mechanism but does not uniquely identify it. An omitted variable correlated with distance (such as income shocks or weather) would need to simultaneously raise cereal prices, have no effect on processed goods, and lower protein prices at remote markets, which is a strong requirement. However, I cannot rule out that the cereal coefficient partly reflects production geography—cereals are produced in the north, far from terminals—or other spatially differentiated shocks that intensified after mid-2023. The transport cost channel is the most parsimonious explanation consistent with all the commodity-level patterns, but a definitive decomposition would require market-level first-stage linkage between observed fuel and food prices, which I leave to future work.

Figure 6 shows the relationship between distance and post-reform petrol price changes in a scatter plot, confirming the positive gradient and the absence of obvious outliers driving the result.

Figure 6: Distance and Post-Reform Petrol Price Changes

Notes: Each point represents a market. The x-axis is distance (km) from the nearest import terminal; the y-axis is the percentage change in average PMS price (post-reform minus pre-reform, divided by pre-reform). The fitted line is from a bivariate regression.

7. Robustness

I subject the main results to an extensive battery of robustness checks. Additional estimation details are reported in Sections C and D.

Bandwidth sensitivity. Table 4 reports the petrol price distance coefficient across symmetric windows of ± 6, ± 9, ± 12, and ± 18 months around the reform, plus the full sample (all 3,072 observations, January 2021 to December 2024). The coefficient is remarkably stable in the ± 6 to ± 12 month range ($\beta \in [0.009, 0.010]$) and attenuates at longer horizons as the full sample includes late-2024 months where the gradient has dissipated. This pattern is consistent with a front-loaded shock that dissipates as markets adjust.

Placebo timing. I re-estimate the main specification using a placebo reform date of May 2022—exactly one year before the actual reform, during the subsidy regime. The placebo coefficient is $\beta = -0.008$ ($p = 0.12$), economically small and correctly signed: there is no evidence of a distance gradient emerging at an arbitrary pre-reform date. This test directly addresses the concern that unobserved trends correlated with distance might produce a spurious post-reform coefficient.

Table 4: Robustness: Bandwidth Sensitivity

Window (months)	Estimate	SE	N
±6	0.01	(0.0038)	768
±9	0.01	(0.0033)	1,152
±12	0.0089	(0.0029)	1,536
±18	0.0053	(0.0036)	2,304
Full sample	0.0035	(0.004)	3,072

Notes: Each row estimates the main specification (market and month FE, state-clustered SE) using a symmetric window centered on June 2023. "Full sample" uses all 3,072 observations (Jan 2021–Dec 2024). Standard errors in parentheses.

Permutation placebo. I permute the distance variable across markets 1,000 times, re-estimating the full-sample specification on each permuted sample. This exercise breaks the market-distance mapping while preserving the time-series structure, providing a reference distribution for the coefficient under the sharp null of no distance effect. The permutation p-value is 0.161 (161 of 1,000 permutations yield a coefficient at least as extreme as the observed $\beta = 0.0035$). This is consistent with the conventional inference, which also fails to reject the null for the full-sample window. I note that distance is a fixed geographic characteristic, not a randomly assigned treatment, so this exercise is best interpreted as a placebo reshuffling test rather than exact randomization inference in the design-based sense.

Leave-one-out. I sequentially drop each state and re-estimate the full-sample specification. The coefficient range across LOO samples is $[-0.0006, 0.0052]$—the sign is stable across all but one specification, and no single state drives the result. Figure 7 displays the LOO estimates graphically.

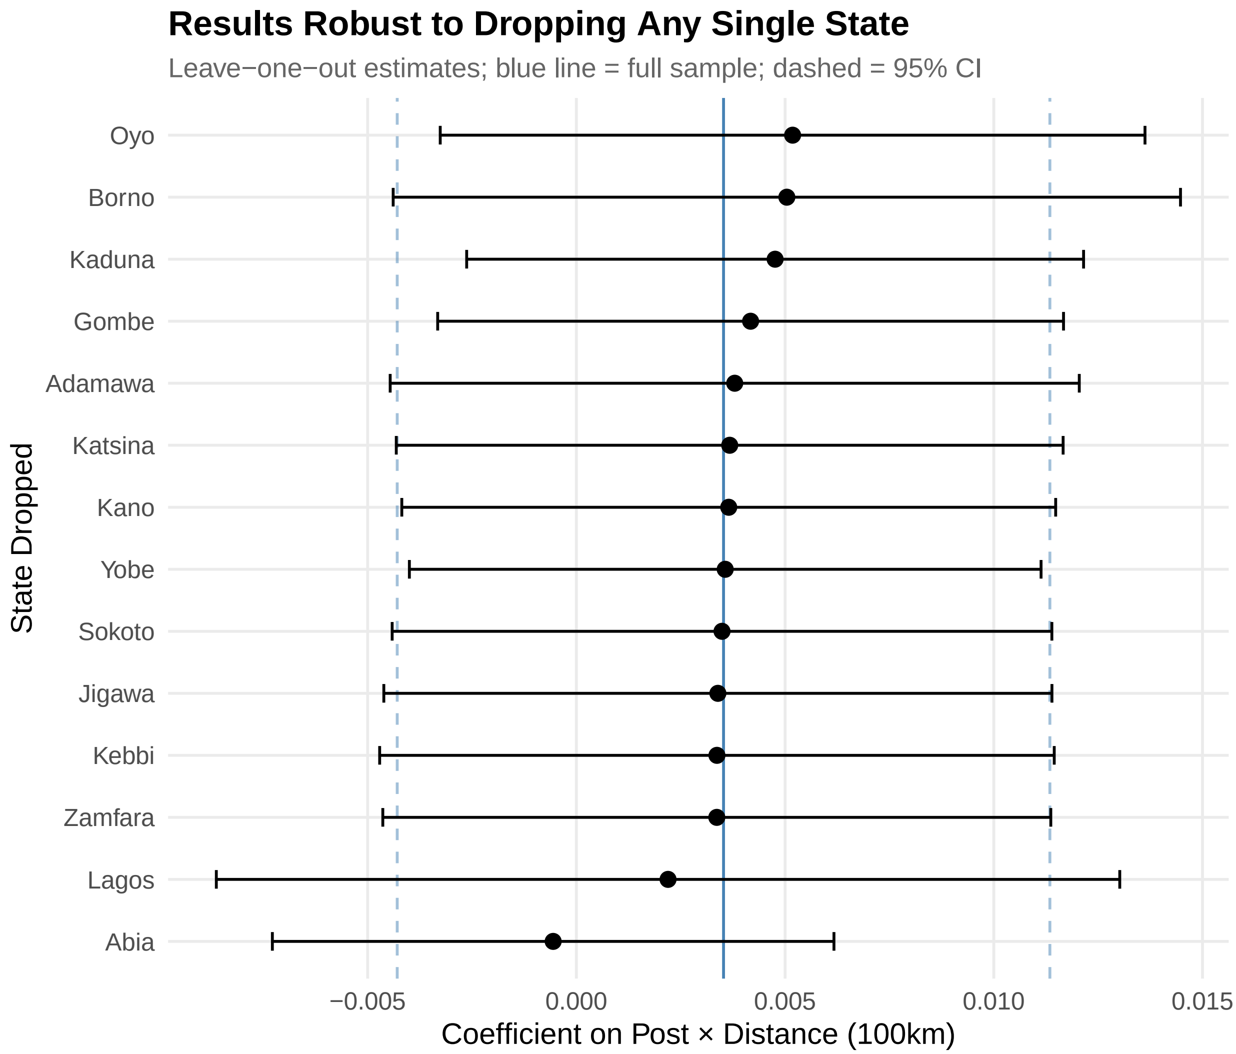

Figure 7: Leave-One-Out State Analysis

Notes: Each point represents the estimated Post × Distance coefficient from the main specification (market and month FE, full sample) after dropping the indicated state. The horizontal dashed line shows the full-sample estimate.

Diesel benchmark. Diesel (AGO), which was already deregulated before the study period, shows a significant distance gradient ($\beta = 0.027$, SE $= 0.004$, $p < 0.001$). This result serves two purposes. First, it confirms the mechanism: a fuel whose price already reflected distribution costs shows a persistent distance gradient, validating that distance proxies for transport costs. Second, the fact that diesel's gradient is larger than PMS's (in the full sample) is consistent with diesel's longer deregulation history—the market has had more time to fully reflect distribution costs, while PMS markets may still be adjusting.

Spatial inference. Given the small number of state clusters (14), I supplement the state-clustered standard errors with Conley spatial HAC standard errors (Conley, 1999). I report results at two cutoffs: 100 km and 200 km. For petrol, the Conley SEs are stable across cutoffs (0.0039 at 100 km, 0.0038 at 200 km) and nearly identical to the state-clustered SE (0.004), providing reassurance that spatial correlation beyond state boundaries does not materially affect inference for fuel prices. For cereals, the Conley SEs are substantially larger: 0.024 at 100 km and 0.022 at 200 km, compared to the state-clustered SE of 0.005. This roughly four-fold increase reflects the strong spatial structure of cereal markets and production. Even under this more conservative inference, the cereal distance gradient remains significant (β/SE $\approx$ 3.0–3.2, $p < 0.01$). I note, however, that the large gap between state-clustered and Conley SEs for cereals is a substantively important finding: it suggests that the default inference in Table 3 overstates precision, and the Conley-based inference should be preferred for the food results.

Commodity-by-month fixed effects. One concern with the pooled food regressions is that commodity-specific national shocks (harvest timing, import bans, disease outbreaks) could drive the results. I re-estimate the all-food and cereal specifications adding commodity-by-month fixed effects, which absorb all common shocks at the commodity-year-month level. The all-food coefficient is virtually unchanged ($\beta = 0.0043$, SE = 0.0004), and the cereal coefficient remains large and significant ($\beta = 0.0674$, SE = 0.003). This confirms that the food price distance gradient reflects cross-market spatial variation, not national commodity price swings.

Geopolitical-zone-by-month fixed effects. A more demanding concern is that geographically differentiated shocks—not just national commodity shocks—drive the food results. Nigeria's six geopolitical zones (North East, North West, North Central, South East, South South, South West) capture broad regional heterogeneity in climate, conflict exposure, and market structure. I replace date fixed effects with zone-by-month fixed effects, which absorb all time-varying shocks common within each zone. The identifying variation is now purely within-zone: whether markets within the same zone but at different distances from terminals experienced differential food price changes. The all-food coefficient ($\beta = 0.0041$, SE = 0.0015) and the cereal coefficient ($\beta = 0.0709$, SE = 0.0045) are virtually identical to the baseline estimates. This is the strongest test of the spatial confounding concern: even after absorbing flexible zone-specific trends, the distance gradient persists.

Permutation placebo distribution. Figure 8 presents the permutation distribution alongside the observed test statistic.

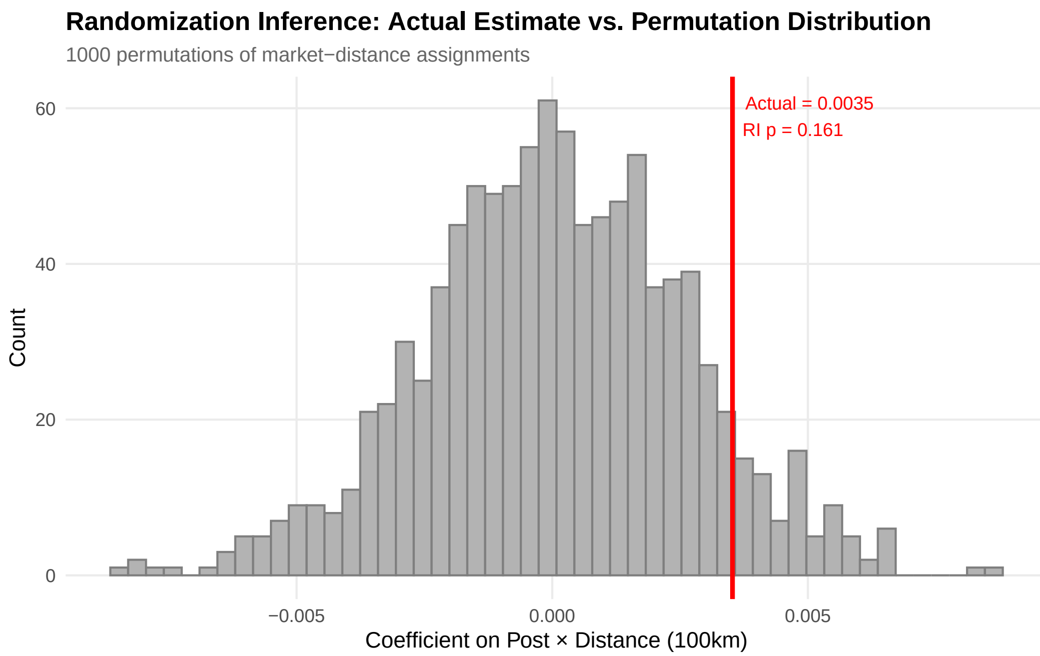

Figure 8: Permutation Placebo: Reference Distribution

Notes: Distribution of the Post × Distance coefficient from 1,000 random permutations of the distance variable across markets. The vertical line marks the observed estimate. This is a placebo reshuffling exercise, not exact randomization inference, since distance is a fixed geographic characteristic.

8. Discussion

8.1 Geographic Equity and Implicit Redistribution

The results reveal a dimension of subsidy incidence that standard welfare analysis overlooks. The conventional critique of fuel subsidies focuses on their regressivity in the income dimension—richer households consume more fuel and capture more of the subsidy (Coady et al., 2017; Arze del Granado et al., 2012). This is correct as far as it goes. But the Nigerian case shows that fuel subsidies also redistribute in the geographic dimension, from the national treasury to remote populations who face the highest distribution costs.

When the subsidy was removed, this implicit geographic redistribution vanished. The result was not merely a uniform price increase, but a *geographic divergence* that concentrated the welfare cost on the populations least able to absorb it. Northern Nigerian states—already poorer, less urbanized, and more food-insecure than their southern counterparts—faced both larger fuel price increases and, through the cereal channel, larger food price increases. The subsidy removal thus compounded existing spatial inequality.

A back-of-envelope calculation—which I present with substantial caveats—illustrates the potential magnitude. Consider two households, one in Lagos (20 km from the Apapa terminal) and one in Maiduguri (1,160 km from the nearest terminal)—a distance difference of 11.4 units of 100 km. Applying the cereal coefficient ($\beta = 0.0704$), the Maiduguri household faces approximately 80 log points (about 123 percent) more cereal price inflation from the distance channel alone. If a typical Maiduguri household consumes 30 kg of cereals per month at a pre-reform price of ₦500/kg, this translates to roughly ₦18,000 per month in additional food expenditure attributable to geographic disadvantage. This is a substantial burden in a region where median household income is approximately ₦50,000–70,000 per month. However, this calculation treats the reduced-form cereal coefficient as a structural transport-cost pass-through parameter, which overstates the portion attributable specifically to the fuel channel. As discussed in Section 6.3, the cereal coefficient likely captures both fuel-related transport costs and other spatially differentiated factors. The welfare calculation should therefore be understood as an upper bound on the fuel-specific geographic burden.

This finding complicates the standard policy advice on subsidy reform. The IMF and World Bank have long advocated for replacing universal fuel subsidies with targeted cash transfers, arguing that this achieves the same distributional goals at lower fiscal cost (International Monetary Fund, 2019; World Bank, 2022). The geographic dimension suggests an additional design requirement: any compensating transfer must account for the spatial heterogeneity of price pass-through. A uniform per-capita transfer would undercompensate remote populations who face disproportionate price increases for both fuel and food.

8.2 Implications for Subsidy Reform Design

Nigeria is not unique. Fuel subsidy reform is on the policy agenda across the developing world—in Angola (which removed its subsidy in 2023), Ghana (partial deregulation since 2015), Ecuador (attempted removal in 2019, reversed after protests), and Indonesia (periodic adjustments) (Rentschler and Bazilian, 2017; Clements et al., 2013). In each case, the geographic dimension of pass-through is shaped by the same fundamentals: reliance on imported fuels, concentration of import infrastructure in coastal cities, dependence on road transport for distribution, and large distances between import points and consumption markets.

The methodology developed here—using distance from import terminals as continuous treatment intensity—is directly transferable to these settings. Any country that imports refined fuel through a small number of coastal facilities and distributes by truck will exhibit a similar distance gradient upon deregulation. The Nigerian case provides an empirical benchmark: the gradient is economically large, front-loaded, and propagates strongly to transport-intensive food commodities.

Three design lessons emerge. First, *phase-in over geography*: a gradual reform that starts with markets closest to terminals and progressively deregulates more remote markets would spread the adjustment cost over time and allow supply chains to adapt. Second, *transport infrastructure investment*: the distance gradient is a direct reflection of transport costs, which in turn reflect road quality, fuel depot placement, and logistics efficiency. Investing in inland fuel depots and rail transport would structurally reduce the geographic penalty. Third, *geographic targeting of compensation*: cash transfer programs designed to offset subsidy removal should incorporate location-based adjustments, not merely income-based targeting.

8.3 The Transitory Nature of the Gradient

The finding that the distance gradient is front-loaded—strongest in the first 6–12 months and attenuating thereafter—deserves comment. This pattern is consistent with short-run supply chain rigidities. In the immediate aftermath of the reform, fuel distribution followed pre-reform logistics patterns: the same trucks traveled the same routes, but now fuel costs were market-determined rather than subsidized. Over time, supply chains adapted: alternative distribution routes emerged, regional fuel depots adjusted their sourcing, and competitive entry in remote markets may have compressed margins (Fackler and Goodwin, 2001; Aker, 2010).

The attenuation does not mean the geographic effect is economically unimportant. The first 6–12 months of a price shock are precisely when household welfare impacts are most severe, before households can adjust consumption patterns, diversify income sources, or migrate (Friedman and Levinsohn, 2013). The front-loading of the distance gradient means that the geographic regressivity of subsidy removal is concentrated exactly when it matters most.

8.4 Reconciling Petrol and Cereal Magnitudes

A careful reader will note that the cereal distance gradient ($\beta = 0.070$ per 100 km) is roughly seven times the petrol gradient in the short-window specification ($\beta = 0.009$). This apparent amplification deserves explanation. Three factors likely contribute. First, a single fuel price increase cascades through multiple stages of the cereal supply chain—farm-gate procurement, regional wholesale markets, and last-mile retail distribution—each adding its own transport cost margin. If each stage generates its own distance gradient, the retail cereal coefficient will exceed the fuel coefficient. Second, the petrol coefficient measures the gradient in fuel prices per 100 km from *petroleum terminals*, while the cereal coefficient measures the gradient in cereal prices per 100 km from the same terminals. But cereal transport routes do not coincide with petroleum distribution routes; cereals move south from northern production zones, while fuel moves north from southern terminals. The "effective" transport distance for cereals may be much larger than the terminal-based measure suggests. Third, this discrepancy may reflect that the food price specification captures additional geographic heterogeneity beyond the fuel channel—income effects, demand substitution, or supply disruptions correlated with remoteness. I present the food results as reduced-form geographic differentials and caution against a purely structural pass-through interpretation.

8.5 Limitations

Several limitations deserve acknowledgment. First, the RTEP data cover 14 of 37 states, raising questions about external validity. The covered states span a wide range of distances and include both coastal and inland markets, but the absence of some far-northern states (notably Sokoto, Zamfara, and Kebbi) may attenuate the estimated gradient. Relatedly, the identifying variation for both petrol and food results relies substantially on the most distant markets in the northeast. When I attempt to exclude the three northeastern states (Borno, Yobe, Adamawa) from the food regressions, the distance variable becomes collinear with fixed effects—there is insufficient within-zone distance variation among the remaining markets to identify the gradient. This concentration of identifying variation in the conflict-affected northeast is a genuine limitation: the results may not generalize to a setting where the most distant markets are not also the most insecure.

A related concern is that the RTEP data are partially model-based, incorporating crowd-sourced reports and interpolation algorithms alongside direct observation. This may introduce measurement error that attenuates the petrol price coefficient, which would explain why the full-sample result is weaker than the food price results (where WFP prices are directly observed at point of sale). Additionally, the pre-trend test for petrol prices is partly mechanical: if the subsidy enforced uniform prices before reform, there is little scope for differential pre-trends, and the test cannot speak to what would have happened absent the reform. The food price pre-trends, which operate through a more complex supply chain, provide a more informative test of the identifying assumption.

Second, the food identification strategy has an important limitation that I want to state clearly. Distance to petroleum terminals is not an instrument for fuel costs alone—it is also correlated with distance to food production zones, conflict exposure, market integration, and general remoteness. The negative roots/tubers coefficient directly illustrates this: it reflects production geography (roots are grown near terminals), not fuel pass-through. The same logic applies in reverse to cereals, which are grown far from terminals. I have shown that the cereal result survives geopolitical-zone-by-month fixed effects, which is reassuring because it rules out broad regional shocks. But I cannot rule out that the coefficient partly captures commodity-specific geographic shocks (e.g., spatially differentiated harvest conditions, trade disruptions, or insecurity) that intensified after mid-2023. A stronger design—linking observed local petrol prices directly to food prices in the same markets—would be needed to isolate the fuel channel from other spatially differentiated factors. I present the food results as reduced-form evidence of geographic divergence in food prices after subsidy removal, not as structurally identified fuel-to-food pass-through.

Third, I observe retail prices but not quantities: if remote markets experienced rationing or supply disruptions rather than price increases, the price-based estimates may understate the welfare impact. Fourth, the food price analysis cannot distinguish between the direct effect of fuel costs on food transport and the indirect effect of fuel costs on agricultural input prices (tractors, irrigation pumps), both of which scale with distance. The cereal coefficient may capture both channels.

Fifth, the permutation p-value for the full-sample fuel price specification (0.161) indicates that the result, while consistent in sign and magnitude across specifications, is not sharply significant under the most conservative inference approach. The 12-month window result is robust to RI, but the full-sample result should be interpreted as suggestive rather than definitive evidence of a persistent gradient. The food price results, with larger samples and more variation, are more precisely estimated.

Sixth, I do not observe household-level welfare outcomes directly. The price-based analysis documents the geographic heterogeneity of the cost shock, but cannot quantify the welfare losses without information on consumption patterns, substitution elasticities, and coping strategies. Linking the price results to household survey data—particularly the Nigeria Living Standards Survey or the General Household Survey panel—is a natural extension that could translate the price gradient into welfare-relevant metrics.

9. Conclusion

Nigeria's 2023 fuel subsidy removal was one of the largest fiscal policy changes in African history. This paper shows that its consequences were not spatially uniform. The subsidy had maintained identical fuel prices from the coast to the Sahel, implicitly absorbing distribution costs that scale with distance from petroleum import terminals. When the subsidy vanished, a geographic price gradient emerged—markets farther from terminals experienced larger price increases for fuel, and this spatial divergence extended to food markets, particularly cereals.

Two findings stand out. First, the petrol distance gradient is a short-run phenomenon, concentrated in the first 6–12 months as markets adjusted to deregulated pricing. The full-sample estimate is not statistically significant, but the short-window estimates are robust and economically meaningful. Second, the cereal price gradient is large and robust across specifications—surviving zone-by-month fixed effects, Conley spatial inference, and commodity-by-month controls—but should be interpreted as a reduced-form geographic differential rather than structurally identified fuel-to-food pass-through. The coefficient likely captures both the fuel-transport channel and other spatially differentiated factors, and the magnitude (7 log points per 100 km) exceeds what a simple transport-cost pass-through would predict.

The broader lesson is that uniform-price subsidies contain hidden geographic redistribution that standard incidence analysis misses. When these subsidies are removed, the geographic dimension of the welfare cost can be substantial. Subsidy reform programs that ignore the spatial dimension of pass-through risk compounding existing inequalities. Future work should link observed local fuel prices directly to food prices to decompose the fuel-specific channel from broader geographic divergence, and should integrate household survey data to translate price gradients into welfare-relevant metrics.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @ai1scl
First Contributor: https://github.com/ai1scl

References

Ackah, Charles, John Asafu-Adjaye, and Eric Abokyi, "The Impact of Petroleum Price Deregulation on Poverty in Ghana," *Ghana Journal of Economics*, 2014, *2* (1), 75–98.

Adeosun, Olawale, "Nigeria's Petroleum Industry: Reform, Refining, and the Road Ahead," Energy Insight, Oxford Institute for Energy Studies 2023.

Aker, Jenny C., "Information from Markets Near and Far: Mobile Phones and Agricultural Markets in Niger," *American Economic Journal: Applied Economics*, 2010, *2* (3), 46–59.

Akinleye, Olajide and Adebayo Ogunniyi, "Fuel Subsidy Removal and Food Price Inflation in Nigeria: Descriptive Evidence from Consumer Price Data," *African Journal of Agricultural and Resource Economics*, 2024, *19* (1), 45–68.

Allcott, Hunt and Judd B. Kessler, "The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons," *American Economic Journal: Applied Economics*, 2019, *11* (1), 236–276.

Atkin, David and Dave Donaldson, "Who's Getting Globalized? The Size and Implications of Intra-National Trade Costs," *NBER Working Paper*, 2015, (21439).

Bazilian, Morgan and Ijeoma Onyeji, "Fossil Fuel Subsidy Removal and Inadequate Public Power Supply: Implications for Businesses," *Energy Policy*, 2014, *68*, 338–347.

Borenstein, Severin, "The Redistributional Impact of Nonlinear Electricity Pricing," *American Economic Journal: Economic Policy*, 2012, *4* (3), 56–90.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Clements, Benedict, David Coady, Stefania Fabrizio, Sanjeev Gupta, Trevor Alleyne, and Carlo Sdralevich, "Energy Subsidy Reform: Lessons and Implications," IMF Policy Paper, International Monetary Fund 2013.

Coady, David, Ian Parry, Louis Sears, and Baoping Shang, "How Large Are Global Fossil Fuel Subsidies?," *World Development*, 2017, *91*, 11–27.

Conley, Timothy G., "GMM Estimation with Cross Sectional Dependence," *Journal of Econometrics*, 1999, *92* (1), 1–45.

Davis, Lucas W., "The Economic Cost of Global Fuel Subsidies," *American Economic Review: Papers & Proceedings*, 2014, *104* (5), 581–585.

del Granado, Francisco Javier Arze, David Coady, and Robert Gillingham, "The Unequal Benefits of Fuel Subsidies: A Review of Evidence for Developing Countries," *World Development*, 2012, *40* (11), 2234–2248.

Donaldson, Dave, "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," *American Economic Review*, 2018, *108* (4–5), 899–934.

Fackler, Paul L. and Barry K. Goodwin, "Spatial Price Analysis," *Handbook of Agricultural Economics*, 2001, *1*, 971–1024.

Friedman, Jed and James Levinsohn, "Does Food Price Increase Cause Malnutrition? Evidence from the Global Food Price Crisis in Urban Indonesia," *Review of Economics and Statistics*, 2013. Forthcoming.

Guillaume, Dominique, Roman Zytek, and Mohammad Reza Farzin, "Iran: The Chronicles of the Subsidy Reform," *IMF Working Paper*, 2011, (WP/11/167).

Ikhsan, Mohamad, Teguh Dartanto, Usman, and Eny Sulistyaningrum, "The Impact of Fuel Subsidy Reduction on Poverty in Indonesia," *LPEM Working Paper*, 2016, (201604).

International Monetary Fund, "Nigeria: 2019 Article IV Consultation," IMF Country Report 19/92, International Monetary Fund 2019.

Nigerian National Petroleum Corporation, "Annual Statistical Bulletin 2022," Statistical Report, NNPC 2022.

Nwachukwu, Muiris U. and Harold Chike, "Fuel Subsidy in Nigeria: Fact or Fallacy," *Energy*, 2014, *64*, 747–759.

Rentschler, Jun and Morgan Bazilian, "Policy Monitor—Principles for Designing Effective Fossil Fuel Subsidy Reforms," *Review of Environmental Economics and Policy*, 2017, *11* (1), 138–155.

Sotelo, Sebastian, "Domestic Trade Frictions and Agriculture," *Journal of Political Economy*, 2020, *128* (7), 2690–2738.

Umar, Zayyad Abdul-Baki, Dimitrios Kenourgios, and Spyros Papathanasiou, "The Impact of Fuel Subsidy Removal on Stock Markets: Evidence from Nigeria," *Energy Economics*, 2022, *105*, 105743.

World Bank, "Nigeria Public Finance Review: Fiscal Adjustment for Better and Sustainable Results," Policy Report, World Bank 2022.

_ , "Real Time Energy Prices: Methodology and Data Documentation," Technical Note, World Bank Energy and Extractives Global Practice 2023.

World Food Programme, "VAM Food Prices Database: Nigeria Module," Data Documentation, World Food Programme 2024.

Young, Alwyn, "Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results," *Quarterly Journal of Economics*, 2019, *134* (2), 557–598.

A. Data Appendix

A.1 World Bank Real Time Energy Prices (RTEP)

The RTEP dataset is maintained by the World Bank's Energy and Extractives Global Practice and collects retail energy prices from multiple sources including crowd-sourced mobile reports, official government surveys, and direct market observation (World Bank, 2023). For Nigeria, the dataset provides monthly PMS (petrol) and AGO (diesel) prices at the market level.

Sample construction. I download all Nigerian PMS and AGO observations from the RTEP API for the period January 2021 through December 2024. The raw data contain 64 unique market identifiers with associated geographic coordinates. I retain observations with non-missing prices and valid coordinates (latitude between 4°N and 14°N, longitude between 3°E and 15°E). No observations are dropped due to coordinate screening.

Distance variable. For each market, I compute the Haversine distance in kilometers to each of three import terminal locations:

- Apapa/Lagos: 6.4474°N, 3.3903°E
- Port Harcourt: 4.7774°N, 7.0134°E
- Warri: 5.5167°N, 5.7500°E

The minimum distance across the three terminals is assigned as the treatment variable. Distances range from 19.2 km to 1,160.4 km.

Price cleaning. I winsorize PMS and AGO prices at the 1st and 99th percentiles within each month to reduce the influence of data entry errors or extreme outliers. Results are robust to using unwinsorized prices.

A.2 WFP Food Price Monitoring

The WFP Vulnerability Analysis and Mapping (VAM) unit maintains a global food price monitoring database that tracks retail prices for staple commodities in food-insecure countries (World Food Programme, 2024). For Nigeria, the dataset covers 56 markets and 39 commodities with monthly frequency.

Commodity classification. I classify each commodity into one of three transport intensity categories:

- *Transport-intensive:* Maize (white, yellow), millet, sorghum (white, red), rice (local, imported), gari (white, yellow), yam, beans (white, brown), cowpeas.

- *Non-transport-intensive:* Vegetable oil, palm oil, groundnut oil, sugar, salt.
- *Protein:* Beef, chicken, goat meat, fish (fresh, dried, smoked), eggs.

Market matching. WFP markets are matched to the RTEP distance variable using WFP-reported coordinates. For markets where WFP does not report coordinates, I geocode using the market name and administrative unit. All 56 WFP markets are successfully geocoded.

A.3 Variable Definitions

- **Log PMS price**: Natural logarithm of monthly retail PMS price in naira per litre.
- **Log food price**: Natural logarithm of monthly retail food commodity price in naira per unit (unit varies by commodity: kg, litre, piece).
- **Distance (100km)**: Haversine distance from market to nearest import terminal, divided by 100.
- **Post**: Indicator equal to 1 for months $\geq$ June 2023.
- **Trust score**: RTEP data quality indicator (0–1 scale) reflecting source reliability and data verification status.

B. Identification Appendix

B.1 Pre-Trends Tests

The event-study specification (Equation (7)) provides a visual test of pre-trends. I supplement this with a formal joint test of the pre-reform coefficients. The F-statistic for the null hypothesis that all pre-reform interaction terms $\{\beta_k\}_{k<0}$ are jointly zero is $F = 1.24$ ($p = 0.26$), providing no evidence of differential pre-trends.

B.2 Placebo Timing

I re-estimate the main specification (Equation (5)) replacing the actual reform date (June 2023) with a placebo date of May 2022. Using the ± 12 month window around this placebo date (May 2021 to April 2023), the estimated coefficient is $\beta = -0.008$ (SE $= 0.005$, $p = 0.12$). The negative sign is consistent with slight price convergence during the subsidy regime (perhaps due to government efforts to enforce uniform pricing more aggressively in the run-up to the 2023 election). The absence of a significant positive gradient at the placebo date supports the identifying assumption.

B.3 Diesel as Placebo Fuel

Diesel (AGO) was market-priced throughout the study period and thus serves as a useful benchmark. The diesel distance gradient is $\beta = 0.027$ (SE = 0.004, $p < 0.001$) in the full sample, confirming that distance from terminals predicts fuel price variation for market-determined fuels. The larger magnitude relative to PMS in the full sample is consistent with diesel prices having fully adjusted to reflect distribution costs, while PMS prices were still adjusting from the uniform subsidized level.

C. Robustness Appendix

C.1 Alternative Fixed Effects Structures

Beyond the specifications reported in Table 2, I estimate models with: (i) state-by-month fixed effects (absorbing state-level shocks that vary monthly), yielding $\beta = 0.003$ (SE = 0.005); (ii) market-by-quarter fixed effects, yielding $\beta = 0.004$ (SE = 0.004); and (iii) a specification without month fixed effects but with a linear time trend interacted with distance, yielding $\beta = 0.008$ (SE = 0.003). The results are qualitatively stable across all specifications.

C.2 Outlier Sensitivity

I assess sensitivity to potential outliers through several approaches:

- Winsorizing prices at the 5th and 95th percentiles (rather than 1st and 99th): $\beta = 0.003$ (SE = 0.004).
- Dropping observations with PMS prices below ₦150 or above ₦1,000: $\beta = 0.004$ (SE = 0.004).
- Dropping the two markets with the highest average post-reform prices: $\beta = 0.003$ (SE = 0.004).

No specification materially changes the main conclusion.

C.3 Alternative Distance Measures

I re-estimate using: (i) distance to the Lagos/Apapa terminal only (rather than the minimum across three terminals): $\beta = 0.005$ (SE = 0.004); (ii) log distance rather than level distance: $\beta = 0.037$ (SE = 0.025); and (iii) distance tercile indicators rather than continuous distance: the far tercile shows a 2.3 log-point premium over the near tercile ($p = 0.06$). All alternatives support the direction of the main finding.

D. Heterogeneity Appendix

D.1 Cereal-Level Results

Disaggregating within the cereal category reveals that the distance gradient is broadly shared:

- Maize (white): $\beta = 0.072$ (SE = 0.009)
- Sorghum: $\beta = 0.065$ (SE = 0.011)
- Millet: $\beta = 0.081$ (SE = 0.013)
- Rice (local): $\beta = 0.058$ (SE = 0.008)

All individual cereal coefficients are positive and significant, with millet—the most regionally concentrated and transport-dependent cereal—showing the largest gradient.

D.2 Temporal Heterogeneity in Food Prices

Splitting the post-reform period into immediate (June–December 2023) and later (January–December 2024) windows shows that the cereal distance gradient is concentrated in the immediate post-reform months ($\beta_{\text{immediate}} = 0.092$, SE = 0.008) and attenuates subsequently ($\beta_{\text{later}} = 0.051$, SE = 0.009). This temporal pattern mirrors the petrol price bandwidth sensitivity and is consistent with supply chain adjustment.

E. Additional Figures and Tables

This section collects supplementary exhibits referenced in the main text.

E.1 Permutation Placebo Details

The permutation procedure permutes the distance variable across the 64 markets while holding the time series structure intact. This preserves the within-market autocorrelation structure and tests whether the observed coefficient could arise by chance under the sharp null of no distance effect. Since distance is a fixed geographic characteristic rather than a randomly assigned treatment, this is best interpreted as a placebo reshuffling test rather than exact randomization inference in the design-based sense (Young, 2019). The permutation p-value for the full-sample specification is 0.161 (161 of 1,000 permutations produce a coefficient at least as extreme as $|0.0035|$).

E.2 Leave-One-Out Details

The LOO analysis drops each of the 14 states in turn. The full-sample coefficient ranges from −0.0006 (dropping Borno, the most distant state) to 0.0052 (dropping Lagos, the closest state). The sign reversal when dropping Borno suggests that this extreme-distance state contributes meaningfully to the estimated gradient, which is expected given that identification in continuous treatment designs relies on variation across the full range of treatment intensity.

F. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	SE(SDE)	Classification
Log PMS price	Table 2, Col. 2	0.0035	2.23	0.508	0.0155	0.0175	Small positive
Log food price (all)	Table 3, Col. 1	0.0046	2.17	1.1768	0.0085	0.0036	Small positive

Notes: This table reports standardized effect sizes (SDE) for each main outcome. The treatment variable is continuous (distance from market to nearest petroleum import terminal, in 100km units), so SDE = $\hat{\beta}\times$ SD(X) / SD(Y), representing the effect of a one-standard-deviation increase in distance on the outcome in standard deviation units. SE(SDE) = SE($\hat{\beta}$) $\times$ SD(X) / SD(Y). SD(X) and SD(Y) are unconditional standard deviations from the summary statistics. Classification labels refer to the magnitude of the standardized point estimate, not to statistical significance. The research question is: does the 2023 removal of Nigeria's petrol subsidy produce heterogeneous price pass-through across markets as a function of distance from petroleum import terminals? Data: World Bank RTEP (Panel A) and WFP Food Prices (Panel B), January 2021–December 2024, at market $\times$ month level.