



Federal Reserve
Bank of Dallas

The Incidence of Fuel-Price Shocks and Tax Holidays: Evidence from the 2026 Oil Shock

Jacob T. Bradt and Reid B. Taylor

Working Paper 2624

August 2026

Research Department

<https://doi.org/10.24149/wp2624>

Working papers from the Federal Reserve Bank of Dallas are preliminary drafts circulated for professional comment. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

The Incidence of Fuel-Price Shocks and Tax Holidays: Evidence from the 2026 Oil Shock*

Jacob T. Bradt[†] and Reid B. Taylor[‡]

July 23, 2026

Abstract

We measure the distributional incidence of US motor-fuel tax holidays using transaction records from ~13,200 gasoline stations linked to neighborhood income. The 2026 Iran War raised gasoline expenditure shares 2.9 times more in the lowest- than highest-income census tracts. Pre-shock exposure accounts for 89% of the gap while the residual heterogeneity widens rather than offsets it. State-level tax holiday lowered retail prices but offset the same fraction (28%) of the per-gallon burden across quintiles. A counterfactual federal holiday preserves this incidence. Per-gallon relief is burden-proportional as it attenuates the shock's level without correcting its regressive income gradient.

JEL: H22, H23, Q41, Q48

Keywords: gasoline tax holidays; pass-through; tax incidence; energy price shocks; distributional analysis

*We thank Erich Muehlegger and Lutz Kilian for helpful comments and discussion. We also thank Nick Hansen for excellent research assistance and PDI Technologies for access to the daily store-level transaction microdata used in this study. Bradt acknowledges support from the McCombs School of Business Faculty Research Excellence Grant. The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Dallas, the Federal Reserve System, the Federal Reserve Board of Governors, The University of Texas at Austin, The State of Texas, PDI Technologies, the US Census Bureau, or the US Energy Information Administration. The authors declare no competing interests. All errors are our own.

[†]Jacob T. Bradt, McCombs School of Business, University of Texas at Austin, jacob.bradt@mcombs.utexas.edu.

[‡]Reid B. Taylor, Federal Reserve Bank of Dallas, reidbtaylor@gmail.com.

1 Introduction

Governments frequently respond to gasoline-price shocks by suspending motor-fuel taxes. Within sixty days of the 2026 Iran oil shock, two US states enacted statutory motor-fuel-tax suspensions ([Office of the Governor of Georgia, 2026a](#); [Office of the Governor of Indiana, 2026a](#)) and a federal motor-fuel-tax holiday is under active consideration in the 119th Congress. These policies are attractive because they are visible to consumers, administratively simple, and can be rapidly implemented. However, their distributional logic is less straightforward. Gasoline-price shocks are typically regressive: although high-income households consume more fuel in total, low-income households spend a larger share of their income on fuel and therefore bear a larger burden. Gasoline-tax holidays, by contrast, return relief in proportion to gallons purchased.

This distinction matters because the same price change can produce very different welfare consequences across the income distribution. To first order, the change in a household's gasoline expenditure share is the product of its pre-shock share and the change in gasoline spending. Incidence can therefore depart from what ex ante budget shares alone predict only if retail pass-through or the demand response varies with income; and because relief is tied to consumption volume rather than relative burden, even substantial pass-through may leave the shock's regressivity largely intact. The central policy question is therefore not only whether or by how much tax holidays reduce retail prices, but whether a per-gallon instrument can rebalance a burden that is measured as a share of household income.

In this paper, we provide novel evidence on the distributional consequences for US households of the 2026 Iran War oil shock and the realized incidence of the fuel-tax holidays adopted in response. We overcome prior data and identification challenges by using transaction records from $\sim 13,200$ US gasoline stations nationwide, directly observing prices and quantities for each station. We match each station to the median household income of its census tract, attributing expenditures and behavioral responses to neighborhoods of known income. We use the unexpected surge in oil price in the wake of the Iran conflict to estimate the dynamic response in fuel prices, quantity sold, and total expenditures compared to each station's seasonally-adjusted counterfactual level. We then use fuel-tax holidays in Georgia and Indiana as quasi-experimental settings to identify the pass-through of statutory relief to retail prices and its variation across the income distribution. Finally, we use our estimate

of Georgia’s pass-through to anchor a counterfactual federal holiday whose aggregate cost and per-quintile incidence depend on in-sample parameters rather than external elasticity estimates.

We find that the shock increased fuel expenditures and that the burden was sharply regressive. At the start of hostilities, fuel expenditures rose an average of 22% above the counterfactual for gasoline and 27% for diesel, before declining to 12% and 3% by week 17 as gradual quantity adjustment partially offset the more immediate and large price spike. Using within-store variation, we estimate a short-run price elasticity of demand for gasoline of -0.23 , increasing to -0.36 in the medium-term by the end of our sample, similar to state-of-the-art estimates in [Kilian and Zhou \(2024\)](#).

The increase in gasoline expenditure shares was 2.9 times larger in the lowest-income tracts than in the highest-income tracts. Using the transaction-level data, we find that lower-income areas adjusted primarily by reducing gallons per transaction — initially purchasing fuel in smaller, more frequent increments — while sustained reductions in trip counts were concentrated in higher-income areas, consistent with a limited ability to eliminate gasoline-consuming trips at the bottom of the income distribution. The larger consumption reduction at the bottom of the distribution, however, does not offset this regressivity. We estimate that pre-shock budget shares account for 89% of the gap as implied fuel-spending growth is close to uniform (22.7% to 25.6%), with the residual heterogeneity widening it.

Next, we find that the policy response was not redistributive. While state-level gasoline-tax holidays substantially reduced retail prices with pass-through estimates of 64%–76%, they offset nearly the same fraction of the per-gallon burden across income quintiles, leaving the regressivity of the initial price shock intact. We show that a counterfactual federal holiday calibrated to the same level of pass-through covers only 16% of the shock’s aggregate burden while again leaving the cross-quintile gradient intact. The policy is therefore burden-proportional rather than redistributive: it reduces the level of the shock while preserving its regressive income gradient. Since per-unit relief subtracts the same amount from every gallon, this pattern is general in that it lowers the level of a burden governed by income shares without correcting the gradient.

Finally, our results connect several literatures: on the pass-through of crude and wholesale-cost shocks to retail prices ([Borenstein et al., 1997](#); [Knittel et al., 2017](#); [Tappata, 2009](#)), on short-run gasoline demand elasticities ([Coglianese et al., 2017](#); [Kilian and Zhou, 2024](#); [Levin](#)

et al., 2017; Li et al., 2014), and on the regressive incidence of fuel taxes (Bento et al., 2009; Chouinard and Perloff, 2004). Whether per-gallon relief can rebalance a shock’s regressivity relies on three objects these literatures measure separately: how much relief reaches the pump (pass-through), how many gallons it applies to (the demand response), and whose burden it offsets (income). Capturing all three together presents a measurement trilemma, as conventional data sources rarely observe more than two. Expenditure surveys and credit card spending data record expenditure amounts and household characteristics, but not the prices and quantities that separate pass-through from demand response; high-frequency station pricing data resolve stations finely but lack proprietary quantity data, which can move first-order even when prices do not (Pizzi, 2026); and transaction data record prices and quantities but not buyer income.

We improve upon prior studies as we are able to observe, in a national panel, station-level transactions linked to geographically granular, neighborhood income. These data allow for improved price elasticity estimates identified from station-level variation that are not attenuated by aggregation and let us test whether the shock’s incidence departs from the mechanical benchmark that ex ante budget shares imply and whether the tax holidays’ pass-through and burden offset are even across the income distribution. Prior tax holiday evaluations that rely on geographically and temporally aggregated data (Doyle and Samphantharak, 2008; Kilian and Zhou, 2024; Kopczuk et al., 2016; Marion and Muehlegger, 2011) recover pass-through using average prices but lack the pertinent income data, while the closest transaction-level study (Bonnet et al., 2025) observes household income and gasoline expenditures, but imputes the quantities on which the demand response is calculated.

2 Data and methodology

2.1 Gasoline Data

Our primary data source is daily store-level transaction summaries from PDI Technologies, a point-of-sale processor for the US convenience-store retail fuel channel which accounts for approximately 80% of US retail gasoline sales (National Association of Convenience Stores, 2025). PDI records fuel sales at 22,219 stores nationally, and our balanced-panel estimation sample comprises $\sim 13,200$ of these sites, which transact approximately 8.3 billion gallons of

gasoline annually, roughly 6.5% of US on-highway gasoline volume.¹ For each store-day-grade observation, PDI reports total revenue, total volume sold (gallons), and transaction count for six motor-vehicle fuel grades: regular, premium, and mid-grade gasoline; low-blend and flex-fuel ethanol; and diesel. Our headline gasoline series is a volume-weighted aggregate of the five gasoline grades while diesel is kept separate throughout. The panel covers January 2024 through the week of 28 June 2026 at daily frequency, which we aggregate to the store-week-grade level for the empirical analysis. We validate the resulting retail-price series against the US Energy Information Administration’s weekly national retail gasoline price. The two series are highly correlated (Pearson’s $r = 0.97$; Fig. E3, Appendix E).

2.2 Estimation sample

The estimation panel is built from the raw daily records via seven restrictions detailed in Appendix A. The binding one, a store-grade panel-coverage threshold, retains store-grade combinations appearing in at least 80% of the calendar weeks in the panel window; sensitivity to alternative thresholds (60%, 70%, 90%) leaves headline price and quantity coefficients within ± 0.5 percentage points of the respective baselines.

We geocode stores from PDI’s monthly store-status snapshots and spatially join them to state, core-based statistical area, county, and 2020 census tract using TIGER/Line boundary files from the US Census Bureau. We then merge tract-level economic and demographic estimates from the 2020–2024 five-year American Community Survey. Tract median household income is the primary stratification variable for the distributional analysis.²

We assign stores to one of five national quintiles of tract median household income, with breakpoints computed once from the final estimation panel’s cross-section. Quintiles are therefore time-invariant by construction, so the distributional results in Section 3.2 cannot be driven by post-event movement of stores across the income distribution. Stores in tracts that lack an ACS income observation enter pooled analyses but are excluded from quintile-stratified analyses.

¹We assess the panel’s coverage of the national market and the income representativeness of its store locations in Appendix A.3.

²We include a multivariate decomposition, presented in Appendix A, that additionally uses tract-level controls for race composition, housing tenure, and commute time.

2.3 Seasonal adjustment

Throughout the analysis, we use seasonally-adjusted outcomes to account for the seasonal variation inherent to retail gasoline sales. For each fuel grade g , we fit

$$y_{ist}^{(g)} = \alpha_i + \gamma_w + \beta_1 t + \beta_2 t^2 + \sum_{h \in \mathcal{H}} \delta_h \cdot \mathbf{1}\{t \in h\} + \varepsilon_{ist} \quad (1)$$

on all available pre-event data (January 2024 through 2026-02-27), where $y_{ist}^{(g)}$ is log price, log quantity, or log expenditure for store i in state s in week t ; α_i are store fixed effects; γ_w are week-of-year fixed effects ($w \in \{1, \dots, 52\}$); $\beta_1 t + \beta_2 t^2$ is a quadratic time trend in weeks since the start of the panel; and $\mathcal{H}$ is the set of date-anchored moving-holiday windows (US Thanksgiving and Easter). For log quantity, the specification additionally includes a calendar-year-by-week fixed effect over the year-end window (weeks 52–53 and 1–4) to absorb annual inventory-reconciliation cycles that the common week-of-year fixed effect cannot otherwise distinguish. We apply the fitted model to the full panel via prediction. The model is fit on data post-2023, when retail fuel demand was still stabilizing from post-pandemic commuting shifts and the unwinding of the 2022 price spike. Expenditure residuals are constructed as the sum of price and quantity residuals, which is exact for the additive linear specification. Appendix B reports the full specification and the residual-construction derivation; Appendix E (Fig. E7) shows robustness to alternative trend, holiday, and calendar-year specifications.

3 The shock’s distributional burden

3.1 Demand response

We estimate the dynamic response of retail prices and quantities to the 28 February 2026 escalation using an event-study leads-and-lags specification. For each fuel grade, we regress

$$y_{ist}^{\text{sa},(g)} = \sum_{k \in [-52, 17] \setminus \mathcal{R}} \beta_k \cdot \mathbf{1}\{t - t^* = k\} + u_{ist} \quad (2)$$

where $y_{ist}^{\text{sa},(g)}$ is the seasonally-adjusted log residual from (1), $t^* = 28\text{February}2026$ is the event week, k indexes weeks relative to t^* , and $\mathcal{R} = \{-4, -3, -2, -1\}$ is the reference

period. Coefficients $\hat{\beta}_k$ for $k < -4$ are pre-event leads and expected to cluster near zero if the seasonal adjustment has captured pre-event dynamics while coefficients for $k \geq 0$ trace the post-event response. Standard errors are clustered at the store level.³

Figure 1 displays the resulting weekly dynamic response of seasonally-adjusted log retail prices (panel a), log retail quantities (panel b), and log retail expenditures (panel c) for aggregate gasoline and diesel, by weeks since the event; Table 1 (Panel A) reports the post-event window-average and most-recent values. Importantly, coefficients in the pre-event periods cluster near zero across all three panels. Following the escalation, retail prices for both fuel classes rose rapidly, consistent with the literature on the propagation of crude oil price shocks to retail gasoline (Barkley et al., 2022; Baumeister and Kilian, 2016; Borenstein et al., 1997; Tappata, 2009). For aggregate gasoline, prices rose 8% above the seasonally-adjusted baseline in the week of the escalation and peaked at 43% before partially retreating to 21% by the week of 28 June 2026 as the initial shock eased with peace negotiations. Averaged across the 18 post-event weeks, gasoline prices were 30% above the counterfactual. Diesel prices moved more rapidly and to higher peaks (47%), averaging 36% above the counterfactual across the post-event window before a parallel late-window retreat to 23% by week 17. The gasoline–diesel disparity is consistent with the concentration of diesel demand among industrial users with limited short-run substitution capacity.

Quantity responses were substantially smaller than price responses for both fuel classes. Averaged across the post-event window, gasoline volumes declined by approximately 6% and diesel volumes by 7% relative to the seasonally-adjusted baseline (panel b). The post-event averages mask a deepening gradient over the first three post-event months, with gasoline quantities reaching 13% below the baseline and diesel quantities 18%. As prices retreated late in the window, quantities recovered only partially to 7% and 17% below baseline, respectively, by the week of 28 June 2026.

There is a long literature estimated the price elasticity of demand for gasoline (Coglianese et al., 2017; Davis and Kilian, 2011; Levin et al., 2017; Li et al., 2014), based on less granular data which can attenuate the results (Levin et al., 2017). Combining our station-level estimated responses for prices and quantities, the implied short-run price elasticities are both approximately -0.23 for gasoline and diesel. This result is inline with the most recent, state-of-the-art, estimate from Kilian and Zhou (2024) who use monthly state-level data with an

³The identifying assumptions and placebo bounds are discussed in Appendix B.

instrumental variables approach. Demand was essentially inelastic as prices jumped. In the medium-run, the implied elasticity became less inelastic, reaching -0.36 for gasoline and -0.49 for diesel by the final month. These late-horizon values may be confounded, as the seasonally-adjusted counterfactual is extrapolated further from the estimation window as the horizon lengthens (Appendix B).

Store-week expenditures rose an average of 22% above the baseline for gasoline and 27% for diesel, before declining to 12% and 3% by week 17 as gradual quantity adjustment partially offset the more immediate price response (panel c). These aggregate measures mask heterogeneity in the margin of response across the income distribution (§3.2), though the regressive incidence we document below arises primarily from pre-shock exposure rather than from differential consumption responses.

Two robustness exercises support the findings of the event-study design. First, as noted above, the store-level price series closely tracks the US Energy Information Administration’s national retail gasoline price. The largest discrepancies appear in the pre-shock baseline rather than around the event (Appendix E, Fig. E3), suggesting the event-window price response is not an artifact of the two samples diverging around the event. Second, we apply the specification to two placebo dates (February 28, 2025 and August 31, 2024), where no shock occurred. The exercise yields maximum log-price deviations of at most 5%, roughly $6.7\times$ smaller than the 43% peak around the 2026 escalation (Appendix E, Fig. E1).⁴

3.2 Distributional incidence and mechanism

Next, we show that the price shock did not distribute its burden uniformly across households. Figure 2 displays the unadjusted weekly expenditures in the PDI data scaled by the number of households in the station’s tract separately by census-tract income quintile, from one year before the event through the week of 28 June 2026 (panel a). The price shock amplified the gradient sharply, with the expenditure shares rising much higher for Q1 than for Q5.

We formally measure the distributional incidence as the pre-to-post change in the per-household gasoline expenditure share within the fixed census-tract income quintiles. For

⁴The headline price and quantity responses are also similar across urbanicity and chain-size strata and common across individual fuel grades, confirming the aggregate result is not driven by a particular retail segment or grade composition.

each income quintile q , we estimate:

$$\omega_{jt} = \alpha_q + \beta_q \text{Post}_t + \varepsilon_{jt}, \quad j \in q, \quad (3)$$

where $\omega_{jt} = (\sum_{i \in j} \text{rev}_{it} / H_j) / (m_j / 52)$ is tract j 's per-household weekly regular-gasoline expenditure as a share of per-household weekly income (H_j households, tract median income m_j), and $\text{Post}_t = \mathbf{1}\{t \geq t^*\}$. We keep ω_{jt} in levels rather than as a seasonally-adjusted residual so that it retains its interpretation as a fraction of income. The pre-event reference window spans the full year before t^* . Standard errors are clustered at the tract level. β_q is the pre-to-post change reported in Figure 2 b and Table 1 (Panel B).

The numerator, ω_{jt} , aggregates revenue at sampled stations rather than observed household spending on fuel. As we are interested in the change over time, within station and tract, our results are robust to sampling differences by tract. ω_{jt} also attributes spending that flows through a tract's stations to its resident households. Appendix B examines two ways this can fail: commuter inflows and shock-induced shifts of customers across stations. Commuter inflow into low-income tracts would inflate the gradient, but restricting to tracts whose inflow does not exceed resident employment leaves the Q1–Q5 ratio at 3.0 against 2.9, so it is not an artifact of this attribution.

Expenditure shares rose by 0.87 percentage points in Q1 but only 0.30 in Q5, a rate $2.9\times$ higher at the bottom of the distribution, with intermediate quintiles tracing a monotonic income gradient. The size of this gradient is close to what pre-shock exposure alone implies. To first order, a quintile's change in expenditure share is the product of its pre-shock share and the proportional change in its per-household fuel spending, $\Delta\omega_q = \omega_q^{\text{pre}} g_q$. Inverting this identity at the estimated $\Delta\omega_q$ implies spending growth of 25.6% in Q1 and 22.7% in Q5, a range of 22.7% to 25.6% across the distribution around a pooled 24.6%. Applying that pooled rate to each quintile's pre-shock share alone reproduces 89% of the observed Q1–Q5 gap (0.52 of 0.58 percentage points), but we can reject exact uniformity in g_q ($\chi^2(4) = 10.22$, $p = 0.04$). The implied spending growth is lowest in the highest-income quintile, so the response heterogeneity that remains widens the predicted gradient from pre-shock shares, from $2.6\times$ to $2.9\times$, rather than offsetting it. Our results align with Bonnet et al. (2025) who find, using evidence from French household data points, that fuel-price sensitivity varies little with income.

For robustness, we report a multivariate decomposition controlling jointly for race composition, homeownership, commute time, and urbanicity in Appendix E (Fig. E4). Income remains the dominant standalone predictor of the post-event share change after these joint controls. We again report a placebo test at two non-event dates and estimate no systematic per-quintile drift in the pre- or post-placebo windows (Appendix E, Fig. E2), supporting the required parallel-trends assumption. Extending the price and quantity event-study specification of (2) to each income quintile separately (Appendix E, Fig. E2) shows within-quintile price and quantity dynamics of similar shape and magnitude across Q1–Q5, consistent with the near-uniform implied spending growth documented above.

To understand the mechanisms behind the change in the expenditure share, we decompose the within-quintile quantity adjustment into per-transaction fill size (intensive margin) and transaction count (extensive margin), measured at two post-shock windows: weeks 1–4 and weeks 13–18 (panel c). Total quantity declined by similar amounts across the income distribution, approximately 14% for both Q1 and Q5, but the margin of adjustment differed both by income and over the course of the shock. In weeks 1–4, households in the lowest-income tracts cut gallons per transaction by -13.6% but actually increased transaction counts by +7.6%. In contrast, the highest-income tracts cut fills by -9.0% with essentially unchanged trip counts (+0.4%). The trip-frequency increase declines monotonically across the income quintiles. By weeks 13–18, transaction counts fell in every quintile while maintaining the cross-income gradient (-9.6% in Q5 versus -6.1% in Q1), while per-fill quantity cuts moderated for all quintiles.

These dynamics point to different margins of adjustment by income. Smaller, more frequent purchases are the standard response to a binding short-run budget constraint: when the cost of a full tank rises, cash-constrained households spread fuel outlays across more visits rather than reduce travel, consistent with evidence on liquidity constraints and price search at the pump (Byrne and de Roos, 2017; Nishida and Remer, 2018). Higher-income households can shift toward remote work, substitute to more fuel-efficient vehicles, or defer discretionary travel, with transaction counts declining as these adjustments accumulate. Lower-income households, whose driving is concentrated in commuting and other non-discretionary travel, show the smallest reduction in trip counts even three months after the shock. Because we observe transactions at the store rather than the household level, we cannot fully separate the same households transacting more often from a shift in the customer mix. For example,

while we cannot rule out increased cross-neighborhood price search may bring additional buyers to a store, we show in Appendix E Fig. E2, that the price response at stations was essentially uniform across tracts. We nevertheless interpret the transaction-count evidence as consistent with, rather than proof of, a liquidity-constraint mechanism.

Lower-income households thus pay a larger share of income for the shock (Fig. 2b) and adjust primarily by reducing purchase size rather than travel (Fig. 2c). We turn next to whether the policy response intended to relieve the burden, a gasoline tax holiday, reaches the households most affected.

4 The distributional incidence of relief

4.1 State tax holidays

Two states enacted gasoline-tax holidays during our panel window, providing three quasi-experiments for measuring how much statutory relief reaches retail prices and how it is distributed across the income gradient. Georgia suspended its \$0.333/gal per-gallon excise from 20 March 2026 to 2 June 2026 (74 days; HB 1199 and Executive Order 05.15.26.02 ([Office of the Governor of Georgia, 2026a,b](#))). Indiana suspended its 7% gasoline sales tax beginning 8 April 2026 (Executive Order 26-09 ([Office of the Governor of Indiana, 2026a](#))), then added a suspension of its \$0.360/gal per-gallon excise tax on 6 May 2026 (Executive Order 26-11 ([Office of the Governor of Indiana, 2026b](#))). Both Indiana suspensions remained in effect through the end of the panel window.

We estimate pass-through with a station-level difference-in-differences design comparing stores in each treated state to stores in its immediate PADD-region neighbors before and after the start of the conflict. For each treatment τ , we estimate

$$\log p_{ist} = \sum_{k \in [-8, 12] \setminus \{-2, -1\}} \beta_k^\tau \cdot \mathbf{1}\{t - t_\tau^* = k\} \cdot \mathbf{1}\{s = s_\tau^*\} + \alpha_i + \gamma_{r(s),t} + u_{ist} \quad (4)$$

on the union of treated and control-state stores (NC and SC act as controls for Georgia while IL, MI, OH, and TN are controls for Indiana), where t_τ^* is the start week of treatment τ , s_τ^* is the treated state, α_i are store fixed effects, and $\gamma_{r(s),t}$ are date-by-PADD fixed effects that absorb common regional cross-state dynamics. The reference period is $k \in \{-2, -1\}$ and

standard errors are clustered at the store level. Pre-holiday lead coefficients are bounded near zero in both states (Fig. 3b), supporting the parallel-trends assumption required for causal interpretation of the coefficients.⁵

We summarize each treatment’s effect on retail prices with a pooled difference-in-differences specification that replaces the relative-week indicators in (4) with a single treatment indicator, yielding the average treatment effect on the treated, $\hat{\beta}_\tau^{\text{pooled}}$. The corresponding pass-through rate is

$$\rho_\tau = \frac{-\hat{\beta}_\tau^{\text{pooled}}}{\Delta \log p_\tau^{\text{theoretical}}}, \quad (5)$$

where $\Delta \log p_\tau^{\text{theoretical}}$ is the log-price change implied by full statutory pass-through.⁶

In Georgia, the tax holiday reduced retail prices for gasoline by 64% of the full statutory tax cut (Table 1, Panel C), implying incomplete but substantial pass-through broadly consistent with prior US state-holiday studies (Doyle and Samphantharak, 2008; Kilian and Zhou, 2024; Tsvetanov, 2024). Georgia’s reinstatement of the excise on 3 June 2026 provides a within-state test of whether pass-through is symmetric with the direction of the tax change. In a joint specification that estimates both the effect of the holiday and the reinstatement against a common pre-holiday reference, retail gasoline prices absorbed 63.9% of the tax cut but only 50.3% of the reinstatement. We can reject symmetric adjustment for both fuels (Appendix E, Fig. E8). Through the end of the panel, prices remained below the pre-holiday baseline by roughly 14% of the statutory tax for gasoline and 29% for diesel, so the reinstatement did not claw back all of the relief the holiday had delivered. Because the panel ends roughly four weeks after reinstatement, this comparison mixes the short-run dynamics of price adjustment with any steady-state asymmetry.

The canonical rocket-and-feather literature finds that retail prices rise quickly with a cost shock but fall more slowly when costs decline, reflecting competitive pressure or collusion (Borenstein et al., 1997) and differential customer search (Tappata, 2009). Our setting reverses the usual order (feathers then rockets): the holiday’s price reduction was immediate, but prices took time to “rocket” back when the tax was reimposed. This asymmetry can

⁵We report sensitivity of the pass-through results to broader control-state definitions (Appendix E, Fig. E6) and to excluding stores within a 30-minute drive of the treated-state border, where cross-border shopping could transmit the holiday’s price effects into control stores (Doyle and Samphantharak, 2008); pass-through moves by no more than 7 percentage points across border-exclusion bands (Fig. E6).

⁶Per-treatment formulas are shown in Appendix C.

in part be explained by tax-collection timing as the tax is collected when wholesale fuel is purchased, so the reimposition acted through the fuel remaining in station tanks on June 3rd, creating asymmetric input costs across stations. To the extent that stores not yet holding new, taxed fuel exerted competitive pressure on neighboring stores, prices rose more slowly. Customers exposed to months of higher prices may also have become more price-sensitive, with increased search behavior further restricting price increases.

In Indiana, retail prices fell by only 40% of the statutory cut during the initial sales-tax-only window, well below the per-gallon-excise estimates in this literature. Indiana’s sales tax on gasoline is posted monthly by the state’s Department of Revenue as a fixed per-gallon dollar amount anchored to pre-shock benchmark prices, so the amount suspended in this window had been set before the spike and released little relief per gallon.⁷ The low 40% measures how little of even that modest cut reached the pump. When Indiana added the per-gallon excise suspension on top of the existing sales-tax suspension on 6 May 2026, combined pass-through rose to 76%, a 36 percentage-point jump within the same state, retailers, and consumers. Holding statutory form and retailer composition fixed, this jump identifies the size of the per-gallon dollar relief, not statutory design, as the margin that governs pass-through (Marion and Muehlegger, 2011).

Focusing on Georgia’s per-gallon design, we compute a per-quintile offset rate as the ratio of per-gallon dollar relief to per-gallon dollar shock burden:

$$o_q = \frac{\tau_{GA} \cdot \rho_{GA}}{\bar{p}_{GA}^{\text{pre}} \cdot \hat{e}_q}, \quad (6)$$

where τ_{GA} is the statutory per-gallon excise tax rate, ρ_{GA} is the pooled pass-through estimate from (5), $\bar{p}_{GA}^{\text{pre}}$ is Georgia’s volume-weighted average retail gasoline price over the pre-event window (\$3.10/gal), and $\hat{e}_q = \exp(\hat{\beta}_q) - 1$ is the per-quintile percent change in gasoline expenditure recovered from the national pre/post log-expenditure regression of §3.2. Both terms are in dollars per gallon, so the offset rate is scale-invariant in gallons consumed and in household income with the only quintile-dependent input being $\hat{e}_q$. Confidence intervals use store-clustered standard errors via the delta method (Appendix C). Figure 3c and Table 1 (Panel D) show the offset rate is broadly similar across the income distribution but rises modestly with income, 26.9% at Q1 to 30.3% at Q5, with a Q5–Q1 difference of 3.5 percentage

⁷We provide further details on the tax structure in Appendix C.

points (SE 1.07) that rejects perfect uniformity ($\chi^2_4 = 13.96$, $p = 0.007$). This residual gradient reflects Q5 households’ larger total quantity decline (§3.2) and is small relative to the offset itself (28% on average) and to the $2.9\times$ regressivity of the shock. Because the formula imposes a single pass-through rate, we also re-estimate the pass through regression separately by tract-income quintile within Georgia, with quintile-matched control stores (Appendix C). Pass-through rises from 61% in the lowest-income quintile to 70% in the highest ($\chi^2_4 = 13.83$, $p = 0.008$). Because this runs in the same direction as the offset gradient, imposing a common ρ_{GA} if anything understates the tilt of delivered relief toward higher-income households.

In this setting, the holiday is approximately burden-proportional: the per-gallon dollar relief is uniform across consumers, while the offset rate varies only modestly because Q5’s larger total quantity decline yields a smaller per-gallon burden. The shock’s burden, however, is governed by income shares, so the post-relief burden inherits the regressive cross-quintile gradient essentially unchanged (§3.2; $2.9\times$ regressivity) (Borenstein and Davis, 2016; Stolper, 2024).

4.2 Counterfactual federal holiday

A suspension of the federal motor-fuel tax is under consideration by the 119th Congress (Lohman, 2026; Penn Wharton Budget Model, 2026). Using the pass-through and per-quintile offset parameters estimated above, we model the proposed four-month holiday, suspending the federal gasoline (\$0.184/gal) and diesel (\$0.244/gal) excise taxes simultaneously over a 1 June–1 October 2026 window.

First, we must extrapolate the burden of the shock to US household totals by applying our estimated store-week price and quantity responses to national on-highway consumption baselines from FHWA Highway Statistics.⁸ The annualized national household burden is approximately \$106 billion from gasoline and \$40 billion from diesel, for a total of \$146 billion per year, or \$1,111 per household, assuming the observed response persists (Appendix E1 reports 3-/6-month alternatives). The cumulative burden through the week of 28 June 2026 is \$50 billion. Figure 4a expresses this aggregate burden as a share of median US household income, and Table 1 (Panel E) describes the national burden and the different relief policies.

⁸Inputs and arithmetic are detailed in Appendix D.

We anchor pass-through to our Georgia state estimates (64% for gasoline and 80% for diesel). A federal holiday was debated but not enacted in 2022 (Penn Wharton Budget Model, 2022), so state suspensions provide the only realized US evidence on motor-fuel-tax pass-through. Under these parameters, the holiday would deliver effective retail price reductions of approximately \$0.12/gal per gallon of gasoline and \$0.20/gal per gallon of diesel, aggregating to approximately \$8 billion in household relief over the four-month window, or about 16% of the shock’s burden over the same period. The per-quintile offset documented in §4.1 carries over to the federal scale across both fuels. As before, uniform per-gallon relief offsets a similar fraction of every household’s per-gallon shock cost, meaning the post-relief share-of-income burden retains the regressive cross-quintile gradient, attenuated only slightly (Q1 from 1.72% to 1.62%; Q5 from 0.50% to 0.47%; Fig. 4b).

This net-of-relief burden varies widely across states (Fig. 4c,d). Heavy-driving, lower-income states in the South and Interior West (WV, MS, NV, WY, AL) bear the highest burden as a share of income, while densely populated and high-income states (DC, WA, RI, MA, MD) see the lowest burden. Decomposing each state’s burden into pre-shock *exposure* and PDI-estimated *response* relative to the US average, exposure alone accounts for 68% of the cross-state variance in log burden and response for another 40%.⁹ The two measures enter multiplicatively, so states on the same iso-burden contour line reach a similar share of income through different mixes. For example, WY has high exposure with a near-average response, while WV, NV, and MS experience higher price responses.

We next benchmark the holiday against a same-outlay means-tested cash rebate, modeled after COVID-19 pandemic relief, with a payment phased linearly across the income distribution and total outlay matched to the holiday’s four-month envelope.¹⁰ Per-household amounts range from \$100 for Q1 down to \$20 for Q5. The rebate attenuates the gradient more substantially (Q1 falls to 1.46%, Q5 to 0.49%), narrowing the cross-quintile burden ratio. A flat per-household dividend from the same envelope pays \$60 to every household but leaves that ratio essentially unchanged from the no-policy baseline (Q1 1.57%, Q5 0.45%). Our results highlight the distributional gradient that per-gallon relief leaves in place, showing that decoupling relief from consumption volume is not sufficient to close the income gradient,

⁹State burden is the product of exposure and response, so the additive log decomposition of its cross-state variance is exact: exposure (67.8%), response (39.6%), and a negative exposure–response covariance (−7.4%) sum to 100%.

¹⁰Further details on the construction of the rebate are in Appendix D.

requiring direct income targeting to do so. Our findings complement the work of [Bonnet et al. \(2025\)](#) who find targeted transfers deliver incomplete compensation, but are preferred to price subsidies during a 2022 price spike in France.

5 Discussion

The 2026 Iran oil shock raised US retail gasoline prices by 30% on average through the week of 28 June 2026, while quantity responses were substantially smaller at 6%. Combining the price and quantity results, we estimate a short-run price elasticity of demand of -0.23 , deepening to -0.36 by the end of the window. The estimated cumulative burden of \$50 billion as of the end of June 2026 was distributed $2.9\times$ more heavily on the lowest-income census tracts than on the highest. Since our panel observes revenue and volume for every transaction, we measure the demand response directly and decompose it into intensive and extensive margins (Fig. 2c).

Several states implemented fuel-tax holidays in an attempt to temporarily offset the burden for households. Georgia’s gasoline-tax holiday passed 64% of its statutory cut to retail prices, offsetting approximately 28% of every household’s per-gallon shock burden, with slight heterogeneity in pass-through by tract quintile. Adding Indiana’s per-gallon excise suspension on top of its existing sales-tax suspension raised combined pass-through from 40% to 76%, isolating the size of the per-gallon relief, rather than statutory design, as the margin that governs pass-through (§4.1). Our pass-through estimates (64%) are broadly consistent with prior aggregate literature for per-gallon excise holidays ([Doyle and Samphantharak, 2008](#); [Kilian and Zhou, 2024](#); [Marion and Muehlegger, 2011](#); [Tsvetanov, 2024](#)). Our results show that the tax holiday is burden-proportional in this setting as it attenuates the per-gallon magnitude of the shock approximately uniformly. The cross-income gradient is unchanged because the burden is set by income shares while the relief is set by consumption volume.

For the federal motor-fuel-tax holiday under consideration in the 119th Congress ([Lohman, 2026](#); [Penn Wharton Budget Model, 2026](#)), the projected aggregate relief at the Georgia per-gallon pass-through rates is approximately \$24 billion per year, or 16% of the shock’s annualized burden. National pass-through could depart from the state-scale estimate in either direction given differences in cost-shock structure and marginal supply conditions ([Boren-](#)

stein and Kellogg, 2014; Muehlegger and Sweeney, 2022). Irrespective of the magnitude, the per-quintile offset rate remains in a narrow band across the income distribution (26.9%–30.3% in our state-level estimate), and the regressive cross-quintile gradient of the burden remains intact in the post-relief residual. Alternative designs that decouple relief from consumption volume (means-tested rebates, capped per-household payments, or income-targeted refundable credits) can narrow the gradient while delivering the same aggregate relief.

Three caveats bound the interpretation of these findings. First, as our post-event window spans 18 weeks, long-run substitution may differ as households adjust vehicle stocks, work arrangements, and residential location (Bernard and Kichian, 2019). Second, the PDI panel is drawn from the convenience-store channel, which accounts for roughly 80% of US retail gasoline sales, but a modest share of national on-highway volume which is tilted toward lower-income neighborhoods (Appendix A.3). The national-burden extrapolation assumes the elasticity estimated within the panel generalizes to the households and channels we do not directly observe. Because the highest-income and least-burdened households are the most under-represented, the measured cross-quintile regressivity is, if anything, a conservative estimate of the national gradient. Lastly, the income gradient is identified across census tracts rather than individual households. Tract median income is a neighborhood proxy that cannot recover within-tract heterogeneity (Douenne, 2020). However, our within-tract, over-time design is robust to level differences, and contemporaneous cross-neighborhood shopping shifts are unlikely to explain our findings.

Our results speak directly to active policy debates over how to deliver relief during periods of high energy costs. Though popular for its salience, simplicity, and speed, a per-unit tax holiday lowers the level of a regressive fuel-price shock but not its incidence. Lowering the price at the pump is not equivalent to easing the burden where it falls hardest.

References

- Barkley, Aaron, David P. Byrne, and Xiaosong Wu (2022) “Price effects of calling out market power: A study of the COVID-19 oil price shock,” *Journal of Economics & Management Strategy*, Vol. 31, No. 4, pp. 923–941.
- Baumeister, Christiane and Lutz Kilian (2016) “Lower Oil Prices and the U.S. Economy: Is This Time Different?” *Brookings Papers on Economic Activity*, Vol. 2016, No. 2, pp. 287–357.
- Bento, Antonio M., Lawrence H. Goulder, Mark R. Jacobsen, and Roger H. von Haefen

- (2009) “Distributional and Efficiency Impacts of Increased US Gasoline Taxes,” *American Economic Review*, Vol. 99, No. 3, pp. 667–699.
- Bernard, Jean-Thomas and Maral Kichian (2019) “The long and short run effects of British Columbia’s carbon tax on diesel demand,” *Energy Policy*, Vol. 131, pp. 380–389.
- Bonnet, Odran, Étienne Fize, Tristan Loisel, and Lionel Wilner (2025) “Compensating against fuel price inflation: Price subsidies or transfers?” *Journal of Environmental Economics and Management*, Vol. 129, p. 103079.
- Borenstein, Severin, A. Colin Cameron, and Richard Gilbert (1997) “Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?” *The Quarterly Journal of Economics*, Vol. 112, No. 1, pp. 305–339.
- Borenstein, Severin and Lucas W. Davis (2016) “The Distributional Effects of US Clean Energy Tax Credits,” *Tax Policy and the Economy*, Vol. 30, No. 1, pp. 191–234.
- Borenstein, Severin and Ryan Kellogg (2014) “The Incidence of an Oil Glut: Who Benefits from Cheap Crude Oil in the Midwest?” *The Energy Journal*, Vol. 35, No. 1, pp. 15–34.
- Byrne, David P. and Nicolas de Roos (2017) “Consumer Search in Retail Gasoline Markets,” *The Journal of Industrial Economics*, Vol. 65, No. 1, pp. 183–193.
- Chouinard, Hayley and Jeffrey M. Perloff (2004) “Incidence of federal and state gasoline taxes,” *Economics Letters*, Vol. 83, No. 1, pp. 55–60.
- Coglianesi, John, Lucas W. Davis, Lutz Kilian, and James H. Stock (2017) “Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand,” *Journal of Applied Econometrics*, Vol. 32, No. 1, pp. 1–15.
- Davis, Lucas W. and Lutz Kilian (2011) “Estimating the effect of a gasoline tax on carbon emissions,” *Journal of Applied Econometrics*, Vol. 26, No. 7, pp. 1187–1214.
- Douenne, Thomas (2020) “The Vertical and Horizontal Distributive Effects of Energy Taxes: A Case Study of a French Policy,” *The Energy Journal*, Vol. 41, No. 3, pp. 231–254.
- Doyle, Joseph J. and Krislert Samphantharak (2008) “\$2.00 Gas! Studying the effects of a gas tax moratorium,” *Journal of Public Economics*, Vol. 92, No. 3, pp. 869–884.
- Fabra, Natalia and Mar Reguant (2014) “Pass-Through of Emissions Costs in Electricity Markets,” *American Economic Review*, Vol. 104, No. 9, pp. 2872–2899.
- Ganapati, Sharat, Joseph S. Shapiro, and Reed Walker (2020) “Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes,” *American Economic Journal: Applied Economics*, Vol. 12, No. 2, pp. 303–342.
- Georgia Department of Revenue (2026) “2026 Suspension of Georgia Motor Fuel Taxes – FAQs,” Georgia Department of Revenue.
- Indiana Department of Revenue (2026) “Departmental Notice #2: Gasoline Use Tax Rate,” Indiana Department of Revenue.
- Kilian, Lutz and Xiaoqing Zhou (2024) “Heterogeneity in the pass-through from oil to gasoline prices: A new instrument for estimating the price elasticity of gasoline demand,” *Journal of Public Economics*, Vol. 232, p. 105099.
- Knittel, Christopher R., Ben S. Meiselman, and James H. Stock (2017) “The Pass-Through

- of RIN Prices to Wholesale and Retail Fuels under the Renewable Fuel Standard,” *Journal of the Association of Environmental and Resource Economists*, Vol. 4, No. 4, pp. 1081–1119.
- Kopczuk, Wojciech, Justin Marion, Erich Muehlegger, and Joel Slemrod (2016) “Does Tax-Collection Invariance Hold? Evasion and the Pass-Through of State Diesel Taxes,” *American Economic Journal: Economic Policy*, Vol. 8, No. 2, pp. 251–286.
- Levin, Laurence, Matthew S. Lewis, and Frank A. Wolak (2017) “High Frequency Evidence on the Demand for Gasoline,” *American Economic Journal: Economic Policy*, Vol. 9, No. 3, pp. 314–347.
- Li, Shanjun, Joshua Linn, and Erich Muehlegger (2014) “Gasoline Taxes and Consumer Behavior,” *American Economic Journal: Economic Policy*, Vol. 6, No. 4, pp. 302–342.
- Lohman, Ali E. (2026) “Suspension of the Federal Gas Tax: In Brief,” Technical Report R48948, Congressional Research Service.
- Marion, Justin and Erich Muehlegger (2011) “Fuel tax incidence and supply conditions,” *Journal of Public Economics*, Vol. 95, No. 9, pp. 1202–1212.
- Muehlegger, Erich and Richard L. Sweeney (2022) “Pass-Through of Own and Rival Cost Shocks: Evidence from the U.S. Fracking Boom,” *The Review of Economics and Statistics*, Vol. 104, No. 6, pp. 1361–1369.
- National Association of Convenience Stores (2025) “Fuels Resource Center: The Convenience Store Fuels Report,” NACS Retail Fuels Resource Center, Convenience stores sell approximately 80% of the fuel purchased in the United States.
- Nishida, Mitsukuni and Marc Remer (2018) “The Determinants and Consequences of Search Cost Heterogeneity: Evidence from Local Gasoline Markets,” *Journal of Marketing Research*, Vol. 55, No. 3, pp. 305–320.
- Office of the Governor of Georgia (2026a) “Gov. Kemp Signs Major Tax Relief Bills for Hardworking Georgians,” Press release, State of Georgia, March.
- (2026b) “Gov. Kemp Suspends Gas Tax for Two Additional Weeks,” Press release, State of Georgia, May.
- Office of the Governor of Indiana (2026a) “Executive Order 26-09: Declaration of Energy Emergency and Waiver of Gas Sales Tax,” State of Indiana, Executive Department, April.
- (2026b) “Executive Order 26-11: Extension of Energy Emergency and Waiver of Gas Sales and Excise Taxes,” State of Indiana, Executive Department, May.
- Penn Wharton Budget Model (2022) “Effects of a Federal Gas Tax Holiday,” Penn Wharton Budget Model, March.
- (2026) “Federal Gas Tax Holiday: June 1, 2026 – October 1, 2026,” Penn Wharton Budget Model, May.
- Pizzi, Francisco (2026) “What Prices Can (and Can’t) Tell Us About Cross-Border Shopping? Evidence from Argentina’s Gasoline Market,” May, Working paper.
- Stolper, Samuel (2024) “Income and Energy Tax Pass-through: Evidence from Gas Stations,” Working paper.

- Tappata, Mariano (2009) “Rockets and feathers: Understanding asymmetric pricing,” *The RAND Journal of Economics*, Vol. 40, No. 4, pp. 673–687.
- Tsvetanov, Tsvetan (2024) “Tax holidays and the heterogeneous pass-through of gasoline taxes,” *Energy Economics*, Vol. 136, p. 107682.

6 Tables and figures




Figure 1. Price, quantity, and expenditure response to the 2026 Iran oil shock.
Notes: Event-study coefficients $\hat{\beta}_k$ from (2) on seasonally-adjusted log price (a), log quantity (b), and log expenditure (c) for aggregate gasoline (blue) and diesel (red); reference period $k \in \{-4, -3, -2, -1\}$. Faded points are weekly estimates with 95% confidence intervals from store-clustered standard errors; solid lines are 3-week centered rolling means. The vertical dashed line marks the event week (28 February 2026); horizontal reference lines mark each series' post-event average and most-recent-week estimate. Sample: $\sim 13,200$ stores observed weekly, January 2024–June 2026.




Figure 2. Per-household gasoline expenditure share by census-tract income quintile. *Notes:* Quintile assignment is cross-sectional and time-invariant; Q1 is the lowest-income quintile, Q5 the highest. **a**, Weekly PDI expenditures per-household as a share of weekly household income; solid lines are 3-week centered rolling means computed within the pre- and post-event windows; the vertical dashed line marks the event week. **b**, Pre-to-post change in expenditure share by quintile (percentage points). **c**, Within-quintile decomposition of the quantity response into intensive (gallons per transaction) and extensive (transaction count) margins, at two post-shock windows: weeks 1–4 (hollow markers) and weeks 13–18 (filled markers). Arrows trace each quintile’s shift from the early to the late window. In weeks 1–4, low-income households (Q1) increase transaction count by +7.6% while high-income households (Q5) hold roughly flat at +0.4%; by weeks 13–18 the extensive-margin gradient inverts, with Q1 down by -6.1% and Q5 down by -9.6%. The Q1 weeks-1–4-to-weeks-13–18 shift is -13.7 percentage points; the corresponding Q5 shift is -10.0 pp.; panels **a** and **b** use tract-clustered standard errors panel **a** clusters at the store level; 95% confidence intervals are shown.




Figure 3. State gasoline-tax holiday pass-through (a, b) and per-quintile offset within Georgia (c). *Notes:* **a**, Weekly seasonally-adjusted log retail price for Georgia (blue), Indiana (green), and other states (grey). Bottom arrows mark the three holiday windows; the vertical dashed line marks 28 February 2026. **b**, Event-study DiD coefficients from (4) on log retail price by weeks relative to each holiday's start, estimated separately for Georgia and Indiana against tight-neighbor controls (GA: NC, SC; IN: IL, MI, OH, TN); reference period $k \in \{-2, -1\}$. Dotted arrows mark the theoretical log-price drop under full statutory pass-through for each treatment; the legend reports observed pooled pass-through. **c**, Per-gallon offset rate of the shock burden by the Georgia holiday, by tract income quintile, from (6). Confidence intervals are 95% from store-clustered standard errors; Standard errors in panel **c** use the delta method.




Figure 4. National burden and counterfactual federal motor-fuel-excise holiday. *Notes:* **a**, US annual household burden (red) and full-year dual-fuel federal-holiday relief (blue) as share of median HH income, decomposed by fuel (light: gasoline; dark: diesel). **b**, Net burden as share of HH income by tract income quintile under three policies: no policy (grey), the dual-fuel federal holiday (red), and a same-outlay linearly means-tested cash rebate (blue). **c**, State-level net-of-relief burden as share of median HH income; each state uses its own PDI-estimated pre-event retail prices and pre/post elasticities scaled to FHWA MF-21 on-highway volumes, with thin-PDI states falling back to their EIA PADD pool. **d**, Exposure-response decomposition of **c**: horizontal axis is each state's pre-shock motor-fuel budget share of income (exposure); vertical axis is the state's PDI-estimated fuel-cost response minus the US average, in percentage points (response). Point color in **d** uses the same net-of-relief burden scale as the map in **c** (RdYlBu, 5th–95th-percentile limits). Dashed grey contours mark iso-burden curves at 1%, 2%, and 3% of HH income. States on the same contour reach similar burden via different exposure/response mixes.

Table 1. Headline empirical estimates.

Panel A. Demand response to the 2026 Iran oil shock					
	Gasoline		Diesel		
	Avg.	Recent	Avg.	Recent	
Price $\Delta \log p$	+30%	+21%	+36%	+23%	
Quantity $\Delta \log q$	-6%	-7%	-7%	-17%	
Expenditure	+22%	+12%	+27%	+3%	
Implied elasticity	-0.23	-0.39	-0.23	-0.87	
Panel B. Distributional incidence by census-tract income quintile					
	Q1	Q2	Q3	Q4	Q5
Pre-event share	3.4%	2.1%	1.8%	1.6%	1.3%
Post-event share	4.3%	2.7%	2.3%	2.0%	1.6%
Pre-to-post change (pp)	0.87 (0.03)	0.53 (0.02)	0.45 (0.02)	0.39 (0.02)	0.30 (0.01)
Panel C. State motor-fuel-tax holiday pass-through					
	Georgia		Indiana		
	Gasoline	Diesel	Sales tax	+ Excise	
Statutory relief (\$/gal)	\$0.333	\$0.373	\$0.187	\$0.609	
Pass-through ρ_τ	64%	80%	40%	76%	
Panel D. Per-quintile offset rate o_q within Georgia					
	Q1	Q2	Q3	Q4	Q5
Offset rate o_q	26.9% (0.72)	27.0% (0.73)	27.6% (0.68)	28.5% (0.71)	30.3% (0.79)
Panel E. National welfare aggregates and federal counterfactuals					
Annualized burden	\$106 billion gas + \$40 billion diesel = \$146B (\$1,111/HH)				
Federal holiday (4-mo)	\$8B (16% of burden)				
Means-tested rebate (Q1 max)	\$100/HH				
Flat per-HH dividend	\$60/HH				

Notes: Sample: $\sim 13,200$ US retail fuel sites, weekly, January 2024–June 2026; event date 28 February 2026. Panels A–D use SA log residuals from Eq. (1); standard errors (store-clustered in Panels B and D and tract-clustered in Panel B), are in parentheses. *Panel A* reports event-study coefficients (Eq. (2)): “Avg.” averages the 18-week post-event window; “Recent” is the panel’s most recent week (28 June 2026); implied short-run elasticity is $\Delta \log q / \Delta \log p$. *Panel B* reports the per-household gasoline expenditure in the PDI sample as a share of tract median income by income quintile; post-event Q1/Q5 ratio is $2.9\times$. *Panel C* reports the per-gallon dollar value of the suspended tax and pass-through from Eq. (5) against state-specific pre-Iran-shock retail-price baselines. *Panel D* reports the Georgia holiday’s per-gallon-burden offset by tract-income quintile (Eq. (6), delta-method SEs). *Panel E* aggregates the shock’s national burden and three same-cost relief instruments (four-month federal motor-fuel excise holiday, means-tested cash rebate, flat per-HH dividend), scaled to FHWA MF-21 on-highway consumption baselines. See §4.1, §4.2, and Appendices C, D for construction details.

Online Appendix for “The Incidence of Fuel-Price Shocks and Tax Holidays: Evidence from the 2026 Oil Shock”

Jacob Bradt and Reid B. Taylor

The following appendices are **for online publication only**.

Appendix A Data

Appendix B Empirical strategy: full specifications and inference

Appendix C Pass-through and per-quintile offset: derivations

Appendix D Welfare aggregation and counterfactual constructions

Appendix E Supplemental figures and tables

A Data

This appendix supplements §2 with the full data pipeline: sources, the sample-construction filters, ACS and EIA variable definitions with source URLs, and the cross-validation of our store-level price series against the US Energy Information Administration’s national reference series.

A.1 Data sources

Retail fuel transactions The primary data source is daily store-level transaction summaries from PDI Technologies, a point-of-sale processor whose data feeds aggregate records from US convenience-store retail fuel sites covering approximately 80% of US retail gasoline volume by site count. The raw data extend from January 2023 through early July 2026 at daily frequency. The empirical analysis uses the 2024-01-01 through 2026-06-28 window, where the end date is the start of the last complete delivery week (trailing partial delivery-cycle weeks are excluded by construction; §A.2). We aggregate PDI transaction records to the store-week-grade level at the start of the pipeline, anchored on PDI’s calendar-week assignment.

Store geography Store latitude/longitude are taken from PDI’s monthly store-status snapshots. We deduplicate the data to one canonical record per store, retaining the most recent snapshot with valid coordinates and carrying forward each store’s earliest `first_seen` and latest `last_seen` months. We restrict the analysis to the contiguous US and spatially join each store in the Albers equal-area projection (ESRI:102005) to state, core-based statistical area, county, and 2023 census tract using TIGER/Line boundary files from the US Census

Bureau. The 30 stores with missing or non-finite coordinates are excluded from quintile-stratified analyses as we cannot assign a census tract-based median income.

Tract-level demographics Tract-level economic and demographic estimates are pulled from the most recent five-year American Community Survey (2020–2024) via the Census Data API. Tract median household income is the primary stratification variable. The multi-variate decomposition reported in Appendix E additionally uses tract-level controls for racial and ethnic composition, housing tenure, and mean commute time, derived from the variables listed in §A.4.

External gasoline-price validation series To validate our store-level prices we use the EIA’s weekly US Regular All Formulations retail gasoline price¹, a stratified sample of approximately 1,000 retail outlets published in the EIA Petroleum Status Report. We build a comparable PDI national series as the volume-weighted average of regular-grade store prices, align it to the EIA’s Monday-anchored weeks, and compare the two over all 129 weeks in which both report data (2023-01-01 through panel end). The two series and their pre/post-event scatter are shown in Fig. E3 (Appendix E).

A.2 Sample construction

The estimation panel is built from the raw daily fuel-transaction records in seven sequential restrictions, summarized in Table A1 and detailed below.

Table A1. Sample-construction filters and parameters.

Step	Threshold / parameter	Rationale
1. Fuel-grade restriction	Six motor-vehicle grades	Drop non-fuel SKUs
2. Store-week aggregation	Sum revenue, volume, transactions	PDI calendar-week grain
3. Analysis-window restriction	2024-01-01 to 2026-06-28	Exclude delivery-cycle stub weeks
4. Per-day completeness	≥ 5 of 7 days per store-week	Drop materially incomplete weeks
5. Balanced-panel filter	$\geq 80\%$ of weeks per store-grade	Consistent operating set
6. Calendar housekeeping	Drop ISO week 54 and low-coverage weeks	Partial year-end weeks; vendor gaps
7. Tract-quintile assignment	Cross-sectional, time-invariant	Movement across quintiles ruled out

¹EIA series PET.EMM_EPMR_PTE_NUS.DPG.W.

Fuel-grade restriction and store-week aggregation (steps 1–2). Beyond the grade filter and store-week aggregation summarized in Table A1 and §2.1, rows reporting non-positive revenue or quantity are dropped.

Analysis-window restriction (step 3). The end of the analysis window is the start of the last complete week in the most recent PDI data delivery. This truncation excludes trailing partial weeks generated by the vendor’s delivery cycle, which can carry several reporting days per store while under-reporting weekly volume.

Per-day completeness (step 4). Store-weeks reporting fewer than five days of transactions (out of seven) are dropped. The threshold removes materially incomplete weeks due to short-term store closures, POS outages, or data-delivery interruptions without imposing a strict seven-day completeness requirement that would discard otherwise usable weeks containing a single missing day.

Balanced-panel filter (step 5). A store-grade combination failing the $\geq 80\%$ weekly-coverage threshold (Table A1, step 5) is dropped from all weeks, not only those it is missing, so within-store comparisons rest on a fixed operating set rather than a changing panel composition. Sensitivity to alternative coverage thresholds is reported in §2.2.

Calendar dates (step 6). ISO-8601 week numbering produces a partial four-day week (28–31 December in non-leap years) at year end. This overflow week is dropped because its shorter duration is not comparable to the surrounding seven-day weeks. We additionally drop any calendar week whose store coverage falls below 50% of the panel-wide median.

Tract-quintile assignment (step 7). Quintile breakpoints are computed once from the cross-section of stores in the final panel and are time-invariant by construction (Table A1, step 7; §2.2). Stores in tracts that lack an ACS income observation — e.g., tracts created or boundary-revised after the most recent ACS vintage, or with insufficient sample for B19013 — are retained in pooled analyses but excluded from quintile-stratified analyses.

A.3 Sample coverage and income representativeness

The PDI panel is a large but non-exhaustive sample of the US on-highway motor-fuel market. In 2025, the raw PDI data contains fuel sales at 22,219 stores transacting 11.7 billion gallons of gasoline, or 9.2% of US on-highway gasoline volume as reported in the Federal Highway Administration’s Highway Statistics Series MF-21 baseline. The balanced-panel estimation

sample of $\sim 13,200$ stores transacts 8.3 billion gallons, or 6.5% of the national on-highway total. Diesel coverage is lower (3.5% raw, 2.5% panel) because convenience stores capture a smaller share of on-highway diesel, much of which refuels at dedicated truck stops.

PDI stores span the full national income distribution but are tilted toward lower-income neighborhoods. The national household-weighted median tract income is \$80,375, whereas the median estimation-sample store sits in a tract with median income \$69,000, and 66.9% of sample stores are in below-national-median tracts. 55.7% of PDI stores are in the lowest two income quintiles nationwide and 8.4% in the highest quintile. Every quintile nonetheless retains ample support, the smallest containing over one thousand sample stores. The sample-selection filters drop a store set with essentially the same income composition as the retained sample (67.6% of dropped stores in below-median tracts, versus 66.9% of retained stores). Because the highest-income and least-burdened households are the most under-represented, the cross-quintile regressivity we document in §3.2 is, if anything, a conservative estimate of the national gradient. Our measure reflects the income of the neighborhoods where stores are located, not the exact income of the households that fuel at each store. Linking transactions to customer residence would require external foot-traffic or card-spend data and is left to future work.

A.4 ACS variable definitions

The income-quintile assignment in §2.2 uses tract median household income from ACS variable B19013_001. The multivariate decomposition in Appendix E additionally uses tract population by race and Hispanic origin (ACS table B03002, with non-Hispanic white B03002_003, non-Hispanic Black B03002_004, and Hispanic B03002_012 subsets), occupied housing tenure (B25003_001 total, B25003_002 owner-occupied), and aggregate worker travel time to work in minutes (B08013_001) normalized by total workers (B08303_001). Total household and total population denominators come from B11001_001 and B03002_001, respectively. All ACS variables are five-year estimates from the 2020–2024 ACS, obtained via the Census Data API through the `tidycensus` R package.

B Empirical strategy: full specifications and inference

This appendix presents the full specifications of the seasonal-adjustment, event-study, distributional, and difference-in-differences regressions summarized in the main text, the identifying assumptions of each design, and the cluster-robust inference protocol referenced throughout.

B.1 Identifying assumptions

Event-study leads-and-lags (§3.1). The identifying assumption is that the seasonal-adjustment model captures pre-event pricing dynamics and would have continued along that path in the post-event period absent the war. Pre-event lead coefficients clustered near zero (Fig. 1) provide support for the assumption, while the placebo event study at non-event dates (Appendix E, Fig. E1) bounds the magnitude of $|\hat{\beta}_k|$ expected under the null. Cross-validation of the seasonally-adjusted price series against the EIA weekly national retail gasoline price (Appendix E, Fig. E3; $r = 0.97$) supports the conclusion that the observed results are not a consequence of a contemporaneous shift in the relationship between sample prices and actual prices.

Pre/post by income quintile (§3.2). The identifying assumption is that the counterfactual within-quintile expenditure share absent the shock would have continued along its pre-event trajectory. The visible downward trajectory in Q1’s pre-event line in Fig. 2a reflects falling retail prices through 2025, suggesting within-quintile shares would have continued shrinking absent the shock. The simple pre-to-post comparison is thus conservative against finding a regressive incidence. The multivariate decomposition controlling jointly for tract race composition, homeownership rate, mean commute time, and urbanicity (Appendix E, Fig. E4) confirms the income gradient survives joint controls.

Customer catchment versus tract of residence (§3.2). The share in Eq. (3) divides station revenue in a tract by that tract’s resident household count and expresses it relative to that tract’s resident median income. It therefore recovers a household budget share only to the extent that a station’s customers both live in the tract and earn its resident income. The main threat to identification is that commuters and through-traffic contribute revenue in a tract that is then attributed to residents’ income. Because the incidence gradient compares low- to high-income tracts, the operative concern is that commuter spending in low-income tracts near job centers and highways is divided by low resident income, inflating the measured low-income share.

To test this, we merge each tract to its 2022 Longitudinal Employer–Household Dynamics (LEHD) LODES commuter-inflow ratio which is calculated as workplace-area jobs divided by the number of resident employed workers. Low-income tracts import more workers than high-income tracts (median inflow ratio of 1.06 in Q1 versus 0.63 in Q5). We then re-estimate Eq. (3) on the residential subsample of tracts whose inflow ratio does not exceed 1, where a station’s customers approximate its residents. For this subset of 6,151 tracts the Q1–Q5 share-change ratio is 3.0, essentially unchanged from the full-sample 2.9, and

stable across commuter-inflow terciles (Appendix E, Fig. E5). Removing the tracts where income-misattribution is likely the most severe therefore leaves the regressive gradient intact.

The inflow adjustment is static, so it does not address sorting induced by the shock itself: if sampled stations offer systematically lower prices and price sensitivity rises as income falls, the price shock could shift traffic toward sampled stations disproportionately in low-income tracts, inflating the measured Q1 share change through customer composition rather than household budgets. Two results bound this channel. The sampled price series tracks the EIA national retail series with no break at the event (Appendix E, Fig. E3), so sampled stations did not become differentially cheap relative to the broader market as the shock unfolded. And the per-quintile price event study (Appendix E, Fig. E2a) shows that stations in low-income tracts developed no widening price advantage over stations in high-income tracts that would attract disproportionate inflows. Residual reshuffling between sampled and unsampled stations within a tract cannot be excluded, but absent a differential price signal it lacks a mechanism to move low- and high-income tracts differently. The transaction-count evidence in §3.2 is read with the same caution.

DiD for state holidays (§4.1). The main identifying assumption is that gas prices in both treatment and control states were evolving along parallel trends, conditional on store and PADD-by-date fixed effects, and that this would have continued in the absence of the tax holidays. The pre-holiday lead coefficients $\hat{\beta}_k$ for $k < -2$ in Fig. 3b provide support for the pre-event period; however, a direct test of the post-period counterfactual is not possible. We restrict the control states to nearby neighbors that also share a PADD, for better compositional comparability and to absorb correlated regional shocks. Sensitivity to broader control definitions (full PADD, full panel) is reported in Appendix E, Fig. E6. Using nearby neighbors as controls raises a second threat: cross-border shopping transmits the holiday’s price effects across the state line (Doyle and Samphantharak, 2008), so control stores near the treated border may cut prices in response to competition from tax-suspended treated stores while treated border stores are disciplined by full-tax competitors. Both channels compress the treated–control gap and attenuate measured pass-through. We therefore re-estimate the pooled specification after excluding all stores within 10 to 50 straight-line miles of the treated-state border (the central 25-mile band corresponds to a 30-minute drive at 50 mph) — and, for the Indiana treatments, of the Kentucky border, given Kentucky’s own 11 May 2026 tax cut. Pass-through moves by no more than 7 percentage points across exclusion bands, and where it moves it moves modestly upward, consistent with the attenuation direction (Appendix E, Fig. E6).

The DiD estimator also requires no anticipation, which is plausible given the quick,

unexpected onset of the war and the resulting tax holidays. The event-study coefficients do not show evidence of anticipation in the weeks preceding the tax holiday legislation. Lastly, identification requires no other changes contemporaneous to the tax holiday that could also affect the outcome variable. While there were evolving economic and political changes taking place during the war, we are unaware of any other shocks that are perfectly correlated with the timing of both tax holidays that could plausibly explain the observed shifts in prices in the treated locations.

C Pass-through and per-quintile offset: derivations

This appendix presents the per-treatment theoretical pass-through formulas referenced in (5), including the Indiana sales-tax mechanics that determine the per-gallon dollar relief implied by the headline 7% statutory rate, and the derivation of the per-quintile offset rate (6).

C.1 Theoretical full-pass-through log-price change

The denominator of (5) is the log-price drop implied by 100% pass-through of the statutory tax cut. For Georgia,

$$\Delta \log p_{\text{GA}}^{\text{theoretical}} = \log(1 - \tau_{\text{GA}}/\bar{p}_{\text{GA}}) = -0.119,$$

where $\tau_{\text{GA}} = \$0.333/\text{gal}$ is the suspended per-gallon excise and $\bar{p}_{\text{GA}} = \$2.98/\text{gal}$ is the pre-holiday volume-weighted average regular-grade retail price in Georgia.

For Indiana’s initial sales-tax-only window,

$$\Delta \log p_{\text{IN}}^{\text{theoretical}} = \log(1 - \bar{g}/\bar{p}_{\text{IN}}) = -0.062,$$

where $\bar{p}_{\text{IN}} = \$3.13/\text{gal}$ is the analogous pre-holiday Indiana retail price and $\bar{g}$ is the volume-weighted-average posted Gasoline Use Tax (GUT) rate during the window (\$0.187/gal). Indiana’s “7% gasoline sales tax” is in practice a per-gallon dollar amount that the state Department of Revenue publishes monthly as 7% of the prior reference period’s average tax-exclusive retail price ([Indiana Department of Revenue, 2026](#)). The dollar amount suspended during the holiday equals the posted rate for the month of sale, not 7% of the contemporaneous retail price. Because the April 2026 GUT was anchored to February data (pre-Iran-shock) and the May 2026 GUT to March data (early-shock), the per-gallon relief during the holiday is materially smaller than a naive 7%-of- $\bar{p}_{\text{IN}}$ calculation would imply.

For the Indiana combined-tax window beginning 6 May 2026, the per-gallon relief is the

sum of the posted GUT and the suspended excise:

$$\Delta \log p_{\text{IN}+}^{\text{theoretical}} = \log(1 - (\bar{g}_+ + \tau_{\text{IN,ex}})/\bar{p}_{\text{IN}}) = -0.216,$$

where $\bar{g}_+ = \$0.249/\text{gal}$ is the average posted GUT during the combined-tax window and $\tau_{\text{IN,ex}} = \$0.360/\text{gal}$ is the suspended per-gallon excise. The combined per-gallon relief is $\$0.609/\text{gal}$.

C.2 Per-quintile offset derivation

The per-quintile offset rate (6) is the ratio of per-gallon dollar relief to per-gallon dollar shock burden. The numerator $\tau_{\text{GA}} \cdot \rho_{\text{GA}}$ is the dollar relief per gallon, uniform across consumers under any per-gallon mechanism. The denominator $\bar{p}_{\text{GA}}^{\text{pre}} \cdot \hat{e}_q$ is the dollar shock burden per gallon; $\bar{p}_{\text{GA}}^{\text{pre}}$ is Georgia’s volume-weighted average retail price of aggregate gasoline over the pre-event window, computed from the PDI panel ($\$3.10/\text{gal}$), and $\hat{e}_q = \exp(\hat{\beta}_q) - 1$ is the per-quintile percent change in expenditure recovered from a pre/post log-expenditure regression run separately by quintile q on the full national panel (cf. §3.2). Anchoring the denominator to Georgia’s own pre-shock price matches the state whose statutory relief the numerator prices, so o_q measures the share of a quintile- q household’s per-gallon shock burden that Georgia-priced relief covers. Both terms are in dollars per gallon, so the offset is dimensionless and scale-invariant: it depends neither on gallons consumed nor on household income, eliminating the conflation between in-sample PDI measurement and national-equivalent measurement that an empirical ratio of two regression coefficients would carry. The only quintile-dependent input is $\hat{e}_q$. Confidence intervals on o_q are delta-method approximations propagating sampling uncertainty in $\hat{\beta}_q$:

$$\text{SE}(o_q) \approx |o_q| \cdot \text{SE}(\hat{\beta}_q)/\hat{e}_q.$$

The pass-through estimate ρ_{GA} is treated as a fixed input for these intervals, since its sampling uncertainty is much smaller than that of $\hat{\beta}_q$ at the relevant scale.

C.3 Pass-through by tract-income quintile within Georgia

The offset rate (6) applies a single pass-through rate ρ_{GA} to every quintile. To test whether relief delivery itself varies with income, we re-estimate the pooled holiday specification of §4.1 separately by tract-income quintile. For each quintile q , the sample comprises Georgia stores in tracts assigned to quintile q and control-state stores (North and South Carolina) in the same quintile, so each quintile’s counterfactual price path is estimated from control

stores serving comparable neighborhoods. The specification retains store and date-by-PADD fixed effects, the joint holiday-plus-reinstatement treatment definition, and store-clustered standard errors; the per-quintile pass-through rate divides the holiday coefficient by the common full-pass-through log drop $\Delta \log p_{\text{GA}}^{\text{theoretical}}$, so cross-quintile variation reflects differential retail price response rather than differences in the statutory relief or the price base. Because each store belongs to exactly one quintile, the five subsamples are disjoint and the estimates are independent, and we test equality with the same inverse-variance-weighted Wald construction used for the offset rates. Pass-through ranges from 60% to 70% across quintiles and rises with income, from 61% in Q1 (SE 2.1 pp) to 70% in Q5 (SE 2.1 pp). Equality across quintiles is rejected ($\chi_4^2 = 13.83$, $p = 0.008$), with a Q5–Q1 difference of 9.9 percentage points (SE 3.0).

D Welfare aggregation and counterfactual constructions

This appendix presents the full national-extrapolation accounting summarized in §4.2, the federal counterfactual’s pass-through-anchor justification (in particular the diesel choice), and the construction of the means-tested cash-rebate benchmark.

D.1 National extrapolation

The aggregate household burden of the 2026 Iran shock and the counterfactual federal-holiday incidence reported in Fig. 4 are constructed by applying the PDI-measured pre-to-post log price and log quantity changes ($\widehat{\Delta \log p}$, $\widehat{\Delta \log q}$ from the headline regression on the aggregate-gasoline and diesel series) to publicly available national on-highway consumption baselines. The state-level burden and relief numbers displayed in the map of Fig. 4c and the exposure-response scatter of Fig. 4d use *state-specific* PDI inputs rather than the national elasticity applied uniformly to state consumption: for each state s and fuel $f \in \{\text{gas, dies}\}$, we re-estimate the same pre/post specification restricted to the store-week panel in s and take that state’s own volume-weighted pre-event retail price, so the burden $\text{burden}_{sf} = \text{gal}_{sf}^{\text{FHWA}} \cdot p_{sf}^{\text{PDI}} \cdot (\exp(\widehat{\Delta \log p}_{sf}) \exp(\widehat{\Delta \log q}_{sf}) - 1)$ reflects within-state incidence heterogeneity, not just pre-existing exposure differences. Two states have thin PDI coverage (16 and 3 gasoline stores in NJ and DC, respectively); four states are thin for diesel (17, 12, 10, and 1 for VT, NJ, RI, and DC). These states fall back to their EIA PADD’s pooled elasticity and pre-event price, so no state rests on insufficient support. The pre-event national baselines — total US households, state-taxed on-highway gasoline and diesel volumes, and pre-event US average retail prices — together with their public-data sources and the full arithmetic are tabulated in Table E1 (Appendix E). We anchor the volume baselines to FHWA Highway

Statistics Table MF-21 rather than EIA STEO or SEDS because MF-21 measures the state-taxed on-highway motor-fuel volumes that the federal excise applies to, excluding home heating oil, off-road agricultural / construction diesel, and dyed marine fuel that EIA’s total-distillate figure includes. This matches (i) the on-highway retail domain that the PDI panel observes and (ii) the excise base of the federal counterfactual, so the burden extrapolation and the federal-relief numerator are both anchored to the same reporting convention. Cumulative burden through the observed post-event window is reported alongside the annualized figure.

The per-household dollar figures in §4.2 and Fig. 4 are national-scale aggregates of total household gasoline expenditure from this extrapolation, distinct from the in-tract PDI revenue that the per-quintile share regression of Fig. 2 identifies; the share-pp gradient, a ratio of within-quintile shares, is unaffected by the scaling. The pooled response used for the burden is referenced to the full pre-event year, the appropriate counterfactual for an annualized figure, whereas the short-run elasticity in §3.1 is referenced to the four weeks immediately preceding the shock. The two imply window-average gasoline elasticities of -0.18 and -0.23, respectively (diesel -0.23 and -0.23), both within the range of prior short-run estimates.

D.2 Counterfactual federal holiday

The counterfactual incidence reported in Fig. 4b simulates a dual-fuel federal motor-fuel excise holiday where the existing federal \$0.184/gal gasoline excise and \$0.244/gal diesel excise are simultaneously suspended for a four-month window beginning 1 June 2026. We anchor pass-through to our own Georgia HB 1199 estimates: 64% for the gasoline excise (matched to the federal gasoline component; Fig. 3b) and 80% for the diesel excise (matched to the federal diesel component; Appendix E, Fig. E6). For the federal diesel anchor we use the full-PADD control-set point estimate rather than the tight-neighbor headline of 80% as the tight-neighbor specification for the diesel subsample relies on only two control states and is therefore noisier, while the full-PADD value sits at the central tendency of the three control sets reported for GA HB 1199 diesel. To our knowledge, GA HB 1199 is the only contemporaneous large-statutory-diesel-tax suspension in the US, making it the natural empirical anchor for federal diesel pass-through. The value sits within the range of the broader per-unit commodity-tax pass-through literature (Fabra and Reguant, 2014; Ganapati et al., 2020; Kopczuk et al., 2016; Marion and Muehlegger, 2011) and is materially higher than the lower-end pass-through values used in some external policy analyses.

The per-quintile dollar offset is $\$0.184/\text{gal} \cdot 64\% \cdot g_q^{\text{gas}} + \$0.244/\text{gal} \cdot 80\% \cdot g_q^{\text{dies}}$, where g_q^{gas} is per-household annual gasoline gallons in quintile q from the PDI panel (scaled to the EIA-anchored national gasoline baseline). The per-household diesel-gallons baseline is $g_q^{\text{dies}} = m_q \cdot \bar{g}^{\text{dies}}$, anchored in level to the FHWA MF-21 HH-weighted national on-highway

mean $\bar{g}^{\text{dies}}$ but scaled by a per-quintile multiplier m_q implied by per-tract-household diesel revenue observed in the PDI panel. We combine the PDI gradient with the FHWA level because convenience-store retail captures only a fraction of on-highway diesel throughput. Rescaling PDI diesel revenue directly to per-household burden would systematically underweight household exposure to commercial diesel throughput, while the combined construction retains both the cross-quintile pattern (a mild Q1>Q5 gradient consistent with low-income tracts’ proximity to freight corridors) and the correct aggregate per-household magnitude.

D.3 Counterfactual means-tested rebate

The same-outlay means-tested cash rebate benchmark introduced in §4.2 is parameterized at the income-quintile level: per-household rebates are $r_q = R \cdot w_q$ for $q = 1, \dots, 5$, with weights $w = (1.0, 0.8, 0.6, 0.4, 0.2)$ stepping down by twenty percentage points of the maximum at each successive quintile, and the maximum payment R chosen so that the aggregate fiscal outlay $\sum_q (N/5) \cdot R \cdot w_q$ equals the four-month dual-fuel federal-holiday relief envelope (\$8B). The resulting payment schedule ranges from \$100/HH at Q1 to \$20/HH at Q5. The schedule abstracts from per-child supplemental payments, filing-status heterogeneity, and continuous within-quintile income phase-out present in actual programs. It is intended as a representative progressive-relief benchmark rather than a faithful reproduction of any specific statute, and its role in Fig. 4b is to isolate the distributional consequences of decoupling relief from gasoline consumption volume while holding total federal cost constant.

D.4 Persistence assumption

The annualized aggregates in Fig. 4 assume the observed post-event level persists for a full calendar year. This is conservative if retail prices retreat as the underlying oil-market shock dissipates, and aggressive if the shock persists or amplifies through further conflict escalation. Sensitivity to alternative persistence horizons (the 18-week observed window, 3-month, 6-month, full-year) is reported alongside the headline figure in Table E1.

E Supplemental figures and tables




Figure E1. Placebo event study at non-event dates. The Fig. 1 event-study specification is re-estimated on pre-actual-event data only ($t < 28\text{February}2026$) for the headline aggregate-gasoline series, with the event week artificially anchored at two placebo dates: 31 August 2024 (blue; mid-pre-period) and 28 February 2025 (green; one calendar year before the actual conflict). Reference week is $k = -1$. **a**, Log price residual. **b**, Log quantity residual. Faded points are weekly estimates with 95% confidence intervals (store-clustered SE; bars are visually compressed by the very large within-store sample size). Bold lines are 3-week centered rolling means computed separately within the pre-placebo and post-placebo periods, with the line break at $k = 0$. Maximum absolute placebo coefficients across the two placebo dates are 5%–5% for log price (panel **a**) and 5%–6% for log quantity (panel **b**). For reference, the actual-event peak in Fig. 1 is 43% for log price and 13% for log quantity (red italicized annotation). The placebo maxima are roughly $6.7\times$ smaller than the actual-event price peak, consistent with the headline price and quantity responses being attributable to the 2026 Iran shock rather than to baseline residual variation. Some pre-placebo dynamic structure visible in the 31 Aug 2024 price-residual series ($k \in [-28, -12]$) reflects unsmoothed seasonal-adjustment residuals in the early panel weeks, where the SA model has limited identifying variation; this residual structure is bounded by the maximum-magnitude statistics reported above.




Figure E2. Per-quintile event-study dynamics at the actual Iran conflict event and at two non-event placebo dates. The Fig. 1 event-study specification is re-estimated separately by census-tract median-household-income quintile (Q1 lowest income, red, through Q5 highest, blue) for the headline aggregate-gasoline series. Columns are the actual event ($t^* = 28\text{February}2026$; grey shading marks post-event weeks; reference window $k \in \{-4, \dots, -1\}$) and two placebo anchors estimated on pre-actual-event data only, 31 August 2024 and 28 February 2025 (reference $k = -1$). Rows are log price residual (a–c) and log quantity residual (d–f). Faded points are weekly estimates with 95% confidence intervals (store-clustered SE, compressed by the large within-store sample); bold lines are 3-week centered rolling means computed separately within the pre- and post-anchor periods. At the actual event (a, d), post-event dynamics are similar across quintiles — price peaks of +34 to +36% and quantity declines within a ~ 5 -percentage-point band — so the per-quintile expenditure-share gradient of Fig. 2 emerges from the interaction of common within-quintile dynamics with pre-event exposure heterogeneity, not from a differential within-quintile shock response. At both placebo dates (b, c, e, f) no systematic per-quintile pre/post shift appears ($|\hat{\beta}_k| \leq 0.05$), supporting the parallel-trends assumption of the within-quintile pre/post design used in §3.2.




Figure E3. Cross-validation of the PDI-derived national retail-gasoline price series against the EIA weekly U.S. retail price. **a**, Weekly volume-weighted average price from this paper’s PDI panel (blue, $\sim 13,200$ stores) overlaid with the US Energy Information Administration’s weekly U.S. Regular All Formulations retail gasoline price (red, series PET.EMM.EPMR.PTE.NUS.DPG.W, compiled from a stratified sample of approximately 1,000 retail outlets independent of the PDI panel). The vertical dashed line marks the start of the Iran conflict (28 February 2026); the shaded region marks post-event weeks. The two series track each other closely throughout the 2023–2026 panel window, including the rapid post-event price increase. **b**, Week-by-week scatter of the PDI volume-weighted price against the EIA national price across all 129 weeks of overlapping data; the dashed line is the 45-degree reference. Pre-event weeks are plotted in grey and post-event weeks in blue. The two series agree at $r = 0.97$ (Pearson) with a mean absolute deviation of $\$0.092/\text{gal}$. Pre-event PDI prices sit modestly below EIA on average, reflecting the convenience-store retail channel covered by PDI versus the broader station mix surveyed by EIA; post-event the series converge and PDI tracks the EIA-recorded price spike one-for-one. The post-event agreement is the relevant validation for the headline event-study residual reported in Fig. 1, since that residual is anchored to pre-event seasonality and trend.




Figure E4. Multivariate decomposition of the post-event expenditure-share gradient. Coefficients on the post-event interaction of seven tract-level margins, from a tract-week fixed-effects regression of weekly per-household gasoline expenditure share on the indicated covariates (Appendix B). Hollow grey points are univariate estimates (one margin at a time, with tract and date fixed effects); filled blue points are joint estimates (all seven margins entered simultaneously, with the same fixed effects). Continuous covariates are standardized to mean zero and unit standard deviation so coefficients are directly comparable across margins; urbanicity dummies (*Metropolitan*, *Micropolitan*, with non-CBSA tracts as the omitted category) are unscaled. Horizontal bars are 95% confidence intervals from tract-clustered standard errors. The income gradient survives joint controls (univariate -0.16 pp per SD; joint -0.14 pp per SD), remaining the strongest single predictor of the post-event share change. Race composition retains independent significance (Black-share joint $+0.03$ pp per SD; Hispanic-share joint $+0.05$ pp per SD), though both are attenuated relative to their univariate point estimates by roughly half. Homeownership is largely absorbed by income (univariate -0.12 pp per SD; joint -0.03 pp per SD), consistent with homeownership's high cross-sectional correlation with median household income. Mean commute time and both urbanicity indicators (*Metropolitan*, *Micropolitan*) are statistically indistinguishable from zero in the univariate specification but become positive and significant in the joint specification, indicating that once tract-level income is conditioned on, longer commutes and CBSA membership are associated with a moderately larger post-event share increase. Sample: $n = 708,205$ tract-week observations across $n = 10,517$ tracts; pre-window starts 2025-02-28.




Figure E5. Customer-catchment robustness of the regressive expenditure-share gradient. The per-household gasoline expenditure share (Eq. (3)) attributes station revenue to a tract’s resident households and resident median income, so commuter and through-traffic spending in worker-importing tracts could bias the measured share and, because the incidence gradient compares low- to high-income tracts, its cross-quintile slope. Using Longitudinal Employer–Household Dynamics (LEHD) LODES commuter-inflow ratios (workplace-area jobs divided by resident employed workers, 2022 vintage), we restrict to residential tracts where a station’s customers approximate its residents in both number and income. **a**, Pre-to-post change in per-household gasoline expenditure share (percentage points) by tract income quintile, for the full sample (grey) and the residential subsample with inflow ratio at most 1 (blue); bars are 95% confidence intervals from tract-clustered standard errors. **b**, Q1/Q5 share-change ratio across catchment definitions: all tracts, the three commuter-inflow terciles (high to low inflow), and the residential subsample. The gradient is 3.0 on the 6,151 residential tracts, essentially unchanged from the full-sample 2.9, and stable across inflow terciles. Because low-income tracts import more workers than high-income tracts (median inflow 1.06 in Q1 versus 0.63 in Q5), the restriction removes precisely the tracts where income-misattribution would be most severe; the regressive incidence therefore reflects household budget shares rather than station siting relative to commuter flows.




Figure E6. Robustness of the Fig. 3 pooled pass-through estimates to control-set and border-exclusion choices. Both panels re-estimate the Fig. 3b pooled difference-in-differences (Appendix B) for four treatments: Georgia gasoline per-gallon excise, Indiana sales-tax-only, Indiana sales tax + per-gallon excise, and a parallel Georgia \$0.373/gal diesel excise suspension (HB 1199; 415 GA diesel stores) (Georgia Department of Revenue, 2026). Points are pooled pass-through rates (observed over theoretical full-pass-through log-price drop); bars are 95% store-clustered confidence intervals; the dashed line marks full statutory pass-through. **a**, Control-state set: tight neighbor (Fig. 3 headline), full PADD, and full panel. Estimates are stable; the Georgia-gasoline versus Indiana-sales-tax gap of about 24 percentage points persists, and Georgia diesel exceeds Georgia gasoline (80% tight-neighbor; 72% broader-control). **b**, Excluding stores within 10, 25, or 50 mi of the treated-state border (and the Kentucky border for Indiana) (Doyle and Samphantharak, 2008); the 25-mile band matches the 30-minute drive of Doyle and Samphantharak (2008) at 50 mph. Pass-through moves by at most 7 percentage points, mildly upward where it moves. Neither margin overturns the headline estimates.




Figure E7. Robustness of the Fig. 1 event-study trajectory to alternative seasonal-adjustment specifications. The headline aggregate-gasoline event study from Fig. 1 is re-estimated under five alternative seasonal-adjustment (SA) specifications and overlaid against the published SA trajectory (blue, bold) on the same window $k \in [-52, 17]$ and the same four-week immediate-pre reference $k \in \{-4, -3, -2, -1\}$. Each line is the 3-week centered rolling mean of the weekly $\hat{\beta}_k$ coefficients (store-clustered standard errors; 95% confidence intervals omitted here for visual clarity, with bands visually compressed by the very large within-store sample size). **a**, Log price residual. **b**, Log quantity residual. Five alternative SA variants are shown as light-to-dark grey lines: *No trend* (store + week-of-year FE only); *Calendar-year FE* (store + week-of-year + calendar-year FE, no continuous trend); *Linear trend* (store + week-of-year FE plus a linear time trend; the canonical event-study SA); *Linear trend, 2025+ pre-window only* (same as Linear but with the SA fit window restricted to the most recent pre-event year); and *Quadratic trend* (store + week-of-year FE plus linear and quadratic time trends). The bold blue *Published SA (Fig. 1)* trajectory adds moving-holiday dummies and, for quantity, a year-end saturation fixed effect on top of the quadratic-trend spec (Appendix B). Across all six lines, pre-event leads cluster within a narrow band and post-event trajectories trend together in direction and rough magnitude; the published SA sits near the centre of the alt-spec bundle in both panels. The visible bump around $k = -7$ that appears in all five alternative-spec quantity trajectories but is absent from the published-SA line reflects year-end inventory-reconciliation weeks (weeks 52–53 and 1–4) that the production SA’s quantity-only year-end saturation fixed effect absorbs — visual evidence of why this component of the production specification matters for the quantity series despite leaving the price series essentially unchanged.




Figure E8. Georgia motor-fuel-excise reinstatement (2026-06-03) and up-vs.-down pass-through symmetry test. Georgia’s HB 1199 excise suspension ended 2 June 2026; the pre-holiday statutory rates on gasoline (\$0.333/gal) and diesel (\$0.373/gal) were reinstated 3 June 2026. **a**, Weekly seasonally-adjusted log price for Georgia (blue) and the Fig. 3 tight-neighbor controls (NC, SC; grey), January 2026 through panel end. Blue shading marks the holiday window; red shading marks the post-reinstatement window. Solid lines are 3-week centered rolling means computed within each regime so the smoother does not average across the 3/20 and 6/3 breaks. **b**, DiD event-study coefficients around the reinstatement (reference $k \in \{-2, -1\}$), store + date $\times$ PADD FE, 95% CIs from store-clustered SE. Pre-reversal leads cluster near zero. A joint holiday-plus-reversal DiD gives a down-direction gasoline pass-through of 63.9% versus an up-direction rate of 50.3%; the corresponding diesel rates are 80.2% down and 51.4% up. A single-coefficient symmetry test $H_0 : \beta_{\text{rev}} = 0$ rejects at $p < 0.001$ (gasoline) and $p < 0.001$ (diesel), with post-reinstatement price gaps equal to -13.6% and -28.8% of statutory tax. The rejection runs in the *reverse* direction of classical rocket-and-feathers behaviour — retail prices declined more completely during the holiday than they rebounded after — though the 25-day post-reinstatement window mixes transition dynamics (panel **b**, $k \in \{0, 1\}$) with any longer-run gap.

Quantity	Value	Source / calculation
National pre-event baselines		
US households	131 million	Census CPS (2024)
On-highway gasoline consumption	127.3B gal/yr	FHWA MF-21 (2024)
On-highway diesel consumption	44.5B gal/yr	FHWA MF-21 (2024)
Pre-event gasoline price	\$3.00/gal	EIA weekly (Q4 2025 avg)
Pre-event diesel price	\$3.50/gal	EIA weekly (Q4 2025 avg)
PDI-measured pre/post change, gasoline (regular grade)		
$\Delta \log$ retail price	+0.277	Result 1
$\Delta \log$ quantity sold	−0.032	Result 1
Implied % Δ expenditure	+27.8%	$\exp(\Delta p) \exp(\Delta q) - 1$
PDI-measured pre/post change, diesel		
$\Delta \log$ retail price	+0.293	Result 1
$\Delta \log$ quantity sold	−0.066	Result 1
Implied % Δ expenditure	+25.4%	$\exp(\Delta p) \exp(\Delta q) - 1$
Implied national burden (annualized; full-year persistence)		
Gasoline, aggregate	\$106B/yr	baseline $\times$ % $\Delta \exp$
Diesel, aggregate	\$40B/yr	baseline $\times$ % $\Delta \exp$
Combined, aggregate	\$146B/yr	gasoline + diesel rows
Combined, per household	\$1,111/yr per HH	aggregate $\div$ HH count
Persistence sensitivity (combined burden per HH)		
18 weeks (observed)	\$385 per HH	annualized $\times$ 18/52
3 months	\$278 per HH	annualized $\times$ 13/52
6 months	\$556 per HH	annualized $\times$ 26/52
Full year (headline)	\$1,111 per HH	annualized $\times$ 52/52

Table E1. Methodology for the national-burden extrapolation (Appendix D; Fig. 4). Sections, top to bottom: (i) National pre-event baselines and their public-data sources; (ii) PDI-measured headline pre-to-post log-price and log-quantity changes used as elasticities, separately for the aggregate-gasoline series and the diesel series; (iii) implied annualized national burden, constructed by applying the PDI-measured implied percent change in nominal expenditure to the pre-event national expenditure baseline for each fuel, then summing; and (iv) sensitivity of the combined burden to alternative persistence assumptions. The headline figures reported in Fig. 4 use the full-year persistence assumption (the observed 18-week post-event level is assumed to persist for a calendar year). Three alternative persistence horizons (18-week observed window, 3 months, 6 months) are tabulated for reference, scaling the annualized aggregate by the indicated fraction of a 52-week year. The PDI–EIA price-path cross-validation reported in Fig. E3 confirms that the elasticities reported in the second section are anchored to a national price trajectory consistent with the independent EIA series ($r = 0.97$).