Rockets Down, Feathers Up? Asymmetric Tax Pass-Through from Malaysia's GST-to-SST Switch

APEP Autonomous Research* @olafdrw

March 10, 2026

Abstract

When Malaysia zeroed its 6% Goods and Services Tax on June 1, 2018—sixteen days after a surprise election ended 61 years of one-party rule—did prices fall as fast as they had risen? Using monthly price indices for 101 product classes, I exploit the sequential removal and reimposition of indirect taxes in a difference-in-differences framework. In my preferred short-window specification, formerly taxed products' prices fell 3.2 percentage points relative to controls (SE = 0.0042), consistent with 55% pass-through. When a narrower Sales and Service Tax was reimposed three months later, the triple-difference estimate suggests prices recovered less than half as much (asymmetry ratio 0.44), though this effect is imprecisely estimated (p > 0.10) and a symmetry test cannot reject equal adjustment (p = 0.075). The evidence is suggestive of a reversed asymmetry—opposite to Peltzman's (2000) canonical finding—consistent with political salience moderating upward price adjustment.

JEL Codes: H22, H25, D40

Keywords: tax pass-through, GST, SST, Malaysia, price asymmetry, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

On May 9, 2018, Malaysian voters delivered one of the most dramatic electoral surprises in Southeast Asian history. Pakatan Harapan, led by 92-year-old former Prime Minister Mahathir Mohamad, defeated the Barisan Nasional coalition that had governed continuously since independence in 1957. The victory was powered, in no small part, by public anger over the 6% Goods and Services Tax (GST) introduced in April 2015—a broad-based consumption tax that replaced a narrower regime of sales and service taxes. Within days of taking office, the new government announced that the GST rate would be zeroed effective June 1, 2018. Three months later, on September 1, a reimagined Sales and Service Tax (SST) was introduced, covering a narrower set of goods and services at varying rates. This sequence of events—tax removal followed by partial reimposition, triggered by an exogenous political shock—provides a rare natural experiment for studying how consumer prices respond to changes in indirect taxation.

The question of tax pass-through is central to public finance (Fullerton and Metcalf, 2002). When a government cuts a consumption tax, who benefits? If firms pass the full reduction to consumers, prices fall by the amount of the tax. If pass-through is incomplete, firms capture part of the tax cut as higher markups. The empirical literature on this question is large but concentrated in a few settings. Peltzman (2000) documented a striking asymmetry in the response of output prices to input cost changes across 242 product markets: prices rise faster in response to cost increases than they fall in response to cost decreases—a pattern he labeled "rockets and feathers." This asymmetry has since been confirmed in gasoline markets (Bacon, 1991; Borenstein et al., 1997; Deltas, 2008), groceries (Peltzman, 2000), and bank deposit rates (Hannan and Berger, 1991). In the specific context of consumption taxes, Benzarti et al. (2020) found that prices rose more following a VAT increase in Finland than they fell following a symmetric VAT cut, consistent with the rockets-and-feathers prediction.

Yet the generality of asymmetric pass-through remains contested. Several studies have found approximately symmetric responses to VAT changes in European settings. Carbonnier (2007) estimated roughly complete and symmetric pass-through of French VAT reforms on housing repairs and new-car sales. Kosonen (2015) exploited Finland's 2010 VAT reduction for hairdressing services and found nearly full and symmetric pass-through. Benedek et al. (2020) studied 6 broad-based European VAT changes and documented a wide range of pass-through rates (30% to over 100%), with the direction and magnitude depending on the competitive structure of the affected market. The emerging consensus, as Kopczuk et al. (2016) argue in their study of diesel fuel tax incidence along the supply chain (see also Marion and Muehlegger, 2011), is that pass-through depends fundamentally on market structure, the position of the tax in the supply chain, and the institutional context.

This paper contributes to this literature by exploiting Malaysia's 2018 GST-to-SST switch—a setting with three distinctive features that distinguish it from prior work. First, the sequential structure of the reform provides within-product variation in both directions: the same products experienced first a tax removal (June 2018) and then a partial tax reimposition (September 2018). This allows a direct within-study comparison of downward and upward price adjustment, avoiding the cross-country or cross-time comparisons that confound much of the existing evidence on asymmetric pass-through. Second, the reimposition of the SST covered only a subset of previously GST-taxed products, creating a natural triple-difference structure: products that experienced both shocks (Group A), products that experienced only the removal (Group B), and products that were never taxed under either regime (Group C, the control). Third, the political origins of the reform—the surprise election outcome—make the timing of the policy change plausibly exogenous to pre-existing price trends in specific product markets, strengthening the causal interpretation of the estimated effects.

I construct a panel of 101 COICOP 4-digit product classes observed monthly from January 2010 to January 2026 using Malaysia's Department of Statistics (OpenDOSM) Consumer Price Index microdata. I classify each product class into the three tax-treatment groups based on its legal status under the GST Act 2014 and the SST Acts of 2018. The core identification strategy is a two-way fixed effects difference-in-differences design comparing log CPI of formerly standard-rated products (treated) to zero-rated and exempt products (control) before and after the June 2018 GST zeroing. The triple-difference specification adds an interaction for products subsequently covered by the SST, identifying the incremental effect of tax reimposition.

The results reveal a clear and precisely estimated effect of tax removal on consumer prices, but one that is sensitive to the choice of estimation window—a feature that itself proves informative about the dynamics of tax pass-through. In the full sample (2010–2026), the difference-in-differences estimate implies that treated products' prices fell 7.6 log points relative to controls after June 2018 (SE = 0.021), corresponding to a pass-through rate of approximately 130% of the 5.83 log-point GST. This point estimate, however, is inflated by long-horizon dynamics: the event study reveals that treated and control products were on different trends in the early part of the sample, reflecting the April 2015 introduction of the GST itself. In the immediate pre-treatment window (months -5 to -1), pre-trends are flat and close to zero. My preferred specification therefore restricts the sample to 2017–2019, where the parallel trends assumption is most credible. In this short window, the estimated effect is -0.032 (SE = 0.0042), implying that treated products' prices fell 3.2 percentage points relative to controls—a pass-through rate of approximately 55% of the GST.

The triple-difference results for the SST reimposition, while less precisely estimated, reveal a striking asymmetry. In the full-sample specification, products subsequently covered by the SST experienced a partial price recovery of 3.8 log points (SE = 0.030), against a DDD removal effect of −0.087—yielding a point-estimate asymmetry ratio of 0.44 within the triple-difference specification. However, this reimposition coefficient is not statistically significant at conventional levels, and a formal Wald test of symmetric adjustment yields p = 0.075. The point estimates are *suggestive* of a reversed asymmetry—prices falling faster than they rise—opposite to the canonical Peltzman pattern, but the evidence is not decisive. One plausible explanation is the unique political context of the Malaysian reform: the GST zeroing occurred under intense public scrutiny, while the SST reimposition received far less enforcement attention. An important caveat is that the asymmetry comparison requires the full-sample specification, where pre-trends are a known concern (see Section 6); the short-window removal estimate (−0.032) captures a different quantity and is not directly comparable to the full-sample reimposition coefficient.

Several features of the empirical design strengthen the credibility of these estimates. Randomization inference based on 1,000 permutations of treatment assignment yields a p-value of 0.000, confirming that the estimated effect is far larger than would be expected from random product-level variation. Leave-one-out analysis shows that the estimate ranges narrowly from −0.079 to −0.072 when each product class is sequentially excluded, ruling out the possibility that any single outlier drives the result. The effect is robust across alternative sample windows, including pre-COVID restrictions and post-GST-introduction specifications. A placebo analysis using June 2015 as an alternative treatment date yields a coefficient of −0.054, indicating that treated products' prices declined relative to controls during the GST era—a pattern consistent with faster general price inflation for zero-rated products (primarily food) than for standard-rated products. This pre-existing differential trend validates both the product classification scheme and the use of the short-window specification.

This paper connects to several strands of the literature. It contributes most directly to the empirical literature on consumption tax incidence, extending the work of Besley and Rosen (1999), who established that sales tax pass-through rates for cigarettes and gasoline commonly exceed 100%; Doyle and Samphantharak (2008), who documented the effects of sales tax moratoriums on gasoline and diesel prices; and Harding et al. (2012), who demonstrated the heterogeneous geographic incidence of excise taxes using scanner data. It also contributes to the growing literature on tax policy in developing and middle-income countries, where the evidence base is thinner and the institutional environment differs from the OECD settings that dominate the literature (Cnossen, 2015; Keen and Ligthart, 2008). The Malaysian setting—with its electronic invoicing infrastructure, VAT-style multi-stage collection, and well-documented legislative history—provides data quality comparable to European studies while offering variation in enforcement capacity and market structure that is unavailable in high-income countries.

The paper also speaks to the broader literature on price adjustment and market power. The finding that pass-through was incomplete (55% in the preferred specification) is consistent with theoretical predictions for markets with imperfect competition (Weyl and Fabinger, 2013), where firms with market power absorb part of tax changes in their markups. The heterogeneity results, which show substantial variation in pass-through across product categories, align with Ganapati et al. (2020) finding that pass-through rates depend on the elasticity of demand and the degree of competition in the market. The reversed asymmetry—prices falling faster than they rise—adds nuance to the narrative established by Peltzman (2000) and Tappata (2009), suggesting that political salience and enforcement intensity can override the structural factors (menu costs, tacit collusion) that typically produce rockets-and-feathers dynamics.

Finally, the paper contributes to the political economy of tax reform. The Malaysian experience illustrates how electoral mandates can amplify or dampen the pass-through of tax policy changes. The Pakatan Harapan government's aggressive monitoring of post-GST prices—including a dedicated Price Control and Anti-Profiteering Act enforced by the Ministry of Domestic Trade—likely increased the speed and completeness of the downward price adjustment. Conversely, the more muted response to SST reimposition may reflect both the narrower tax base and the absence of comparable enforcement pressure. These findings have implications for the design of tax holidays and temporary tax reductions, a policy tool used in numerous countries during the COVID-19 pandemic (Crossley et al., 2021).

2. Institutional Background

2.1 Malaysia's Tax History

Understanding the 2018 GST-to-SST switch requires context about Malaysia's evolving indirect tax regime. For decades prior to 2015, Malaysia operated a dual-track system consisting of a Sales Tax (levied on manufactured goods at rates of 5% or 10%, collected at the point of manufacture or import) and a Service Tax (levied on prescribed services at 6%). This regime suffered from several well-documented limitations: cascading taxation (tax-on-tax), a narrow base that excluded many services, and significant compliance gaps due to the single-stage collection point (Leow and Yew, 2019; Saari et al., 2023). International organizations, including the IMF and World Bank, had long recommended a transition to a broad-based multi-stage value-added tax.

On April 1, 2015, the Barisan Nasional government introduced a 6% Goods and Services

Tax (GST), replacing the old Sales Tax and Service Tax. The GST was a standard VAT-type multi-stage tax with input tax credits, covering a broad base of goods and services. Products were classified into three categories under the GST Act 2014: *standard-rated* (taxed at 6%), *zero-rated* (subject to GST at 0%, with input tax credits retained), and *exempt* (outside the GST system entirely, with no input tax credits). Zero-rated goods included essential foodstuffs (rice, sugar, cooking oil, fresh vegetables), while exempt supplies included residential property, public transportation, and certain financial and health services.

The GST was deeply unpopular. Despite the government's argument that it broadened the tax base and reduced fiscal dependence on petroleum revenues, consumers perceived the GST as having raised the cost of living. The visible, multi-stage nature of the tax—appearing on every receipt, from small convenience stores to major retailers—made it a potent political issue, consistent with the salience effects documented by Chetty et al. (2009). Opinion surveys consistently showed the GST as the top concern among Malaysian voters (Welsh, 2018).

2.2 The 2018 Election Shock

The 14th Malaysian general election, held on May 9, 2018, produced a result that virtually no political analyst had predicted. Pakatan Harapan, the opposition coalition, won a parliamentary majority, defeating Barisan Nasional and ending 61 years of uninterrupted rule by the same coalition (and its predecessor, the Alliance). The result was powered by several factors—the 1MDB financial scandal, Malay rural discontent, and the return of Mahathir Mohamad to active politics—but GST abolition was the coalition's most prominent economic promise.

The new government moved with extraordinary speed. On May 16, 2018—just seven days after the election—Finance Minister Lim Guan Eng announced that the GST rate would be reduced to 0% effective June 1, 2018. This gave retailers only 16 days to adjust pricing systems, update point-of-sale software, and recalculate shelf prices. The compressed timeline is important for identification: it left minimal scope for anticipatory price adjustment or strategic behavior by firms.

2.3 The Three-Month Tax Holiday

From June 1 to August 31, 2018, Malaysia experienced what was effectively a three-month consumption tax holiday. The GST rate was set to zero but the GST Act remained technically in force, meaning the administrative infrastructure (registration, filing requirements) was maintained. This period created a uniquely clean experimental window: all products that had been taxed at 6% suddenly faced a zero rate, while products that were already zero-rated or exempt experienced no change in their tax treatment.

The government was acutely aware of the political risk that firms might not pass the tax reduction to consumers. The Ministry of Domestic Trade and Consumer Affairs activated the Price Control and Anti-Profiteering Act 2011, establishing monitoring teams in major retail areas and inviting consumers to report businesses that failed to reduce prices. This enforcement activity was widely covered in Malaysian media and likely contributed to the speed and completeness of the price adjustment.

2.4 SST Reimposition

On September 1, 2018, the Sales Tax and Service Tax were reimposed under new legislation (Sales Tax Act 2018 and Service Tax Act 2018). Critically, the new SST was not simply a reinstatement of the pre-2015 system. It differed in several important respects. First, the tax base was substantially narrower: the Sales Tax covered only manufactured goods (at rates of 5% or 10%), while the Service Tax applied to specific prescribed services at 6%. Many products and services that had been standard-rated under the GST were not covered by either the Sales Tax or the Service Tax. Second, the Sales Tax was a single-stage tax collected at the point of manufacture or import, without input tax credits, reintroducing the cascading problem that the GST had been designed to eliminate. Third, the threshold for Sales Tax registration was set at RM 500,000 of annual taxable goods, exempting smaller manufacturers.

This narrower coverage creates the variation that powers the triple-difference design. Products previously standard-rated under the GST divide naturally into two subgroups: those subsequently covered by the SST (Group A, 20 product classes) and those not covered (Group B, 41 product classes). Group A products experienced first a tax decrease (June 2018) and then a tax increase (September 2018), while Group B products experienced only the decrease. Comparing the magnitude of the September price recovery for Group A relative to Group B (within the treated group) identifies the incremental pass-through of the SST reimposition.

2.5 Implications for Identification

Several features of this institutional setting are favorable for causal identification. The timing of the GST zeroing was determined by a political event (the election) rather than by economic conditions in specific product markets. The 16-day implementation window limited anticipatory behavior. The product-level variation in SST coverage was determined by pre-existing legislative categories, not by firm or industry lobbying in response to the

2018 reform—the new government simply adopted a modified version of the pre-2015 tax schedule. And the three-month window between the GST zeroing and SST reimposition provides a clean separation between the two shocks, allowing the downward and upward price adjustments to be estimated independently.

A potential concern is that the April 2015 GST introduction is itself visible in the data, creating long-horizon pre-trends that could contaminate the 2018 estimates. I address this directly in the empirical strategy by using a short-window specification (2017–2019) as the preferred estimator, where the pre-treatment trends are flat. The full-sample estimate, which includes the 2015 GST introduction in the pre-period, serves as a useful benchmark but must be interpreted with the understanding that the long-horizon dynamics reflect the earlier reform.

3. Data

3.1 Consumer Price Index Data

The primary data source is Malaysia's Department of Statistics Consumer Price Index, accessed through the OpenDOSM (Open Data Malaysia) platform. The dataset provides monthly price indices for 101 product classes defined at the 4-digit level of the Classification of Individual Consumption by Purpose (COICOP) hierarchy. The data span January 2010 to January 2026, yielding 193 monthly observations per product class and a total of 19,493 product-month observations. All indices are normalized to a base year of 2010 = 100.

The 4-digit COICOP classification provides a natural unit of analysis for studying tax pass-through. Each product class corresponds to a relatively homogeneous group of consumer goods or services (e.g., "Rice," "Fresh fruit," "Garments for men," "Motor cars," "Restaurants and cafes"). The tax treatment under both the GST and SST regimes is defined at the product level, and in nearly all cases every item within a COICOP class shares the same tax status. Where a class contained a mix of taxed and untaxed items, the classification was assigned based on the dominant tax treatment by CPI weight (see Appendix A for details). In practice, only a small number of classes are mixed, and the classification is validated against observed price breaks at the June and September 2018 dates.

The CPI data have two important features for this study. First, they are based on a fixed basket of goods and services, with weights updated periodically, meaning that compositional changes in consumption patterns do not drive the measured price movements. Second, the indices are constructed from actual transaction prices collected from retail outlets across Malaysia, providing a reliable measure of the prices consumers actually pay. The Department of Statistics conducts price collection in major urban centers and selected rural areas, covering the full spectrum of retail formats from traditional markets to modern supermarkets.

3.2 Product Classification

I classify each of the 101 COICOP 4-digit product classes into three mutually exclusive tax-treatment groups based on their legal status under the GST Act 2014 and the Sales Tax Act 2018 / Service Tax Act 2018:

- **Group A (GST-taxed, SST-covered):** 20 product classes that were standard-rated at 6% under the GST and are subject to the Sales Tax or Service Tax under the reimposed SST. These products experienced a tax removal in June 2018 and a tax reimposition in September 2018. Examples include motor vehicles, tobacco products, alcoholic beverages, and certain household equipment.
- **Group B (GST-taxed, no SST):** 41 product classes that were standard-rated at 6% under the GST but are not covered by the reimposed SST. These products experienced a permanent tax removal in June 2018. Examples include clothing, footwear, household textiles, personal care products, and recreational goods.
- **Group C (Zero-rated/Exempt):** 40 product classes that were zero-rated or exempt under the GST and thus experienced no change in their effective tax rate at either date. These serve as the control group. Examples include rice, fresh vegetables, fresh fruit, education services, and public transportation.

The classification is based entirely on the legal text of the relevant Acts, constructed ex ante without reference to post-2018 price movements. As an ex post diagnostic, I verify that the legally-coded groups exhibit price behavior consistent with the classification. Table 1 shows that the classification produces sharp differences in the mean price break at the two key dates: Group A products experienced a substantial average June 2018 price decline, compared to a smaller decline for Group B and near-zero change for Group C. At the September 2018 SST reimposition, Group A products increased on average, while Group B products showed essentially no change. The exact magnitudes are reported in Table 1.

Table 1: Product Classification by Tax Treatment

Group	Tax Treatment	N Classes	Mean June Break (%)	Mean Sept Break (%)
A	GST-taxed, SST-covered	20	-3.63	1.96
B	GST-taxed, no SST	41	-2.66	-0.10
C	Zero-rated/Exempt (Control)	40	-0.13	0.07
Total		101	-1.88	0.38

Notes: Products classified by GST status (standard-rated at 6% vs. zero-rated/exempt) and SST coverage (Acts 806/807 of 2018). June Break = % CPI change May–June 2018. Sept Break = % change Aug–Sept 2018. Validated against legal schedules and observed price behavior.

3.3 Summary Statistics

Table 2 presents summary statistics for the analysis panel. The full sample comprises 19,493 product-month observations across 101 product classes and 193 months. Group B (GST-taxed, no SST) is the largest group with 41 classes and 7,913 observations, followed by Group C (40 classes, 7,720 observations) and Group A (20 classes, 3,860 observations). Of the 19,493 total observations, 504 (2.6%) have missing CPI values for product classes introduced after the 2010 base period; these are excluded from regression analyses, yielding effective sample sizes of 18,989 in the main specifications. The mean CPI for Group B is 108.5, indicating an average cumulative price increase of 8.5% since the 2010 base year.

Panel C of Table 2 decomposes the sample by tax regime period. The pre-GST era (January 2010 to March 2015) accounts for 6,363 observations, the GST era (April 2015 to May 2018) for 3,838, the three-month tax holiday (June to August 2018) for 303, and the SST era (September 2018 onward) for 8,989—the largest subsample, reflecting the longer duration of the post-treatment period.

Table 2: Summary Statistics

	N	Mean	Std. Dev.	Min	Max
Panel A: Full Sample					
CPI Index	19,493	113.4	18.8	52.2	341.3
log(CPI)	19,493	4.719	0.157	3.955	5.833
Panel B: By Tax Treatment Group					
Group A: GST-taxed, SST-covered	3,860	112.8	16.0	71.6	202.1
Group B: GST-taxed, no SST	7,913	108.5	18.9	52.2	341.3
Group C: Zero-rated/Exempt	7,720	119.1	18.4	57.1	191.0
Panel C: By Period					
Pre-GST (2010–2015)	6,363	103.8	6.3	68.2	142.2
GST era (2015–2018)	3,838	113.3	13.8	62.6	179.3
Tax holiday (Jun–Aug 2018)	303	113.6	16.1	59.2	179.3
SST era (Sep 2018+)	8,989	120.2	23.0	52.2	341.3

Notes: 19,493 product-month observations across 101 COICOP 4-digit product classes and 193 months (January 2010–January 2026). CPI indices from OpenDOSM with base year 2010 = 100. Group A products were standard-rated at 6% under GST (April 2015–May 2018) and are covered by SST (September 2018 onward). Group B products were standard-rated under GST but are not covered by SST. Group C products were zero-rated or exempt under GST and serve as controls.

4. Empirical Strategy

4.1 Identification

The core identification strategy exploits the abrupt zeroing of the 6% GST on June 1, 2018 as a quasi-experimental shock to the effective tax rate on standard-rated products. The identifying variation is the interaction between a product's pre-determined tax classification (standard-rated versus zero-rated/exempt) and the timing of the policy change. Under the parallel trends assumption, absent the tax change, the CPI trajectories of treated and control products would have evolved in parallel.

The parallel trends assumption is plausible for several reasons. First, the product-level tax classifications were established by the GST Act 2014 and were in place for three years before the 2018 reform, limiting the scope for endogenous classification. Second, the timing of the reform was determined by an election outcome that was itself a surprise, making it unlikely that pre-existing price trends in specific product categories drove the policy change. Third, the event study evidence shows that the pre-treatment divergence is concentrated at longer horizons ($k = -36$ to $k = -10$), reflecting the 2015 GST introduction. In the immediate pre-treatment window (months -5 to -1 relative to June 2018), the individual event study coefficients are small, providing evidence that the acute divergence associated with the 2018 reform is distinct from the gradual trend during the GST era. Formal joint tests of pre-treatment coefficients reject at conventional levels (see Section 6 and Appendix B), motivating the short-window specification as the preferred estimator.

A complication is that the same products classified as "treated" in 2018 were also affected by the April 2015 GST introduction. Because the GST replaced the old Sales Tax (which covered a different set of products at different rates), the 2015 transition altered the relative price levels of treated and control groups. This generates a long-horizon pre-trend in the full-sample event study: treated products show positive coefficients at event times -36 to -10, reflecting the fact that the GST era saw relative price increases for standard-rated products. This is not a violation of the parallel trends assumption for the 2018 shock—it is the expected mechanical consequence of the 2015 reform. However, it does mean that the full-sample DiD estimate captures both the 2018 zeroing and the long-run adjustment from the 2015 introduction.

To isolate the causal effect of the 2018 reform, I adopt a short-window specification restricting the sample to January 2017 through December 2019 as the preferred estimator. In this window, the pre-treatment period (18 months) is long enough to assess parallel trends but short enough that the 2015 GST introduction does not contaminate the pre-period. The event study coefficients for months -5 to -1 are visually flat and individually insignificant, though a formal joint F-test over a longer 12-month pre-treatment horizon still rejects at the 5% level ($p = 0.011$), reflecting residual low-frequency drift. I interpret the visual evidence as supporting approximate parallel trends in the immediate pre-treatment window, while acknowledging that formal tests remain sensitive to the longer horizon. As Roth (2022) emphasizes, pre-trend tests have limited power to detect violations that could meaningfully bias treatment effect estimates; the rejection here is driven by low-frequency divergence from the 2015 reform rather than high-frequency violations near treatment. Nevertheless, following the framework of Rambachan and Roth (2023), the pre-trend evidence warrants caution in interpreting the precise magnitudes of the full-sample estimates.

4.2 Estimation

4.2.1 Difference-in-Differences

The baseline specification estimates the following two-way fixed effects model:

$$\log(\mathrm{CPI}_{it}) = \alpha_i + \delta_t + \beta_1(\mathrm{Treated}_i \times \mathrm{Post}_t) + \varepsilon_{it} \tag{1}$$

where $\log(\mathrm{CPI}_{it})$ is the natural logarithm of the CPI for product class i in month t; α_i are product class fixed effects that absorb time-invariant differences in price levels across products; δ_t are month fixed effects that absorb common aggregate price shocks (inflation, commodity prices, exchange rate movements); $\mathrm{Treated}_i$ is an indicator for products that were standard-rated at 6% under the GST (Groups A and B); Post_t is an indicator for months from June 2018 onward; and ε_{it} is the error term. The coefficient β_1 measures the average effect of GST zeroing on log CPI for treated products relative to controls, conditional on product and time fixed effects. Standard errors are clustered at the product-class level to allow for arbitrary within-product serial correlation. With 101 clusters, this is well above the conventional threshold for reliable cluster-robust inference (Cameron et al., 2008).

4.2.2 Triple-Difference (DDD)

To separately identify the effect of SST reimposition, I augment the specification with a triple interaction:

$$\log(\mathrm{CPI}_{it}) = \alpha_i + \delta_t + \beta_1(\mathrm{Treated}_i \times \mathrm{Post}_t) + \beta_2(\mathrm{Treated}_i \times \mathrm{SST}_i \times \mathrm{PostSept}_t) + \varepsilon_{it} \tag{2}$$

where SST_i indicates products covered by the reimposed SST (Group A) and $\mathrm{PostSept}_t$ indicates months from September 2018 onward. The coefficient β_2 captures the incremental price adjustment for SST-covered products relative to non-SST products within the treated group—i.e., the pass-through of the SST reimposition, controlling for any common post-June trends among treated products.

The asymmetry ratio is defined as $|\hat{\beta}_2|/|\hat{\beta}_1|$, measuring the fraction of the tax-removal price decline that is "recovered" by the tax reimposition. Under symmetric pass-through with identical tax rates, this ratio would equal 1. Under Peltzman-style rockets-and-feathers, the ratio would exceed 1 (prices rise faster than they fall). A ratio below 1 indicates the reverse: prices fall faster than they rise.

4.2.3 Event Study

The event study specification replaces the single post-treatment indicator with a full set of event-time interactions:

$$\log(\text{CPI}_{it}) = \alpha_i + \delta_t + \sum_{k \neq 0} \gamma_k \cdot \mathbb{I}[t = k] \times \text{Treated}_i + \varepsilon_{it} \tag{3}$$

where k indexes months relative to May 2018 ($k = 0$, the omitted reference month—the last full month of GST at 6%). The coefficients γ_k for $k < 0$ serve as a visual test of the parallel trends assumption, while γ_k for $k > 0$ trace out the dynamic treatment effect. I estimate this specification for event times $k \in \{-36, \ldots, +36\}$ to capture both the long-horizon dynamics and the immediate treatment response.

4.3 Threats to Validity

Anticipation. The 16-day window between the announcement (May 16) and implementation (June 1) leaves limited scope for anticipatory price adjustment. However, some retailers may have begun marking down prices before the official date, particularly for durable goods where customers could defer purchases. If so, the event study would show a negative coefficient at $k = 0$ (May 2018) or $k = -1$ (April 2018). In practice, the pre-treatment coefficients are flat through May 2018, suggesting minimal anticipation at the monthly frequency of the data.

Compositional changes. The CPI is based on a fixed basket, so changes in consumption patterns following the tax reform do not mechanically affect the measured price indices. However, if the tax change caused quality upgrading or downgrading within product categories, the CPI could capture quality-adjusted price changes rather than pure tax pass-through. This concern is mitigated by the use of product-level (rather than item-level) indices, which aggregate across quality tiers within each category.

Spillovers. Tax changes in one product category could affect prices in related categories through substitution effects or supply-chain linkages. For example, if consumers substitute from zero-rated rice to newly-cheaper standard-rated bread, demand for rice could fall, reducing rice prices and contaminating the control group. Such spillovers would bias the DiD estimate toward zero (making treated and control groups more similar), implying that my estimates are conservative.

Multiple treatments. The treated group includes products affected by both the 2015 GST introduction and the 2018 zeroing. The short-window specification addresses this by restricting the sample period, but the full-sample estimates should be interpreted as capturing the net effect of both reforms. I present both for transparency.

Heterogeneous treatment effects. The 6% GST was applied uniformly to all standard-rated products, but the SST rates vary across products (5%, 10%, or 6% depending on the category). The DDD coefficient β_2 therefore captures an average reimposition effect that masks within-Group-A heterogeneity. The heterogeneity analysis in Section 5.4 explores product-level variation in treatment effects.

5. Results

5.1 Main Results

The zeroing of the GST immediately lowered treated products' prices by 7.6 log points relative to exempt controls—an overshifting pass-through rate of approximately 130% of the full benchmark of $-\log(1.06) \approx -0.0583$ (Table 3, Col. 1, SE = 0.021, $p < 0.01$). This finding, from the baseline DiD specification (Equation 1), indicates that formerly standard-rated products fell roughly 7.3% in price immediately after the GST was zeroed, with the point estimate exceeding the mechanical full pass-through benchmark.

Column (2) decomposes the post-treatment period into the tax holiday (June–August 2018) and the post-SST era (September 2018 onward). The tax holiday coefficient is -0.078 (SE = 0.015), while the post-SST coefficient is -0.076 (SE = 0.021). The near-equality of these two coefficients indicates that the average price reduction for treated products (relative to controls) was largely unchanged by the SST reimposition—an initial puzzle that the DDD specification resolves.

Column (3) presents the triple-difference specification from Equation 2. The main treatment effect (Treated × Post-June) is -0.087 (SE = 0.024), while the SST reimposition interaction (Treated × SST-Covered × Post-SST) is 0.038 (SE = 0.030). The SST coefficient is positive, consistent with a partial price recovery for products subject to the reimposed tax, but imprecisely estimated—the 95% confidence interval includes zero ($[-0.021, 0.097]$). The imprecision reflects the relatively small number of Group A products (20 classes) available to identify the SST effect.

Column (4) introduces a four-period specification that separately estimates the treatment effect during the GST era (April 2015–May 2018), the tax holiday, and the SST era, all relative to the pre-GST baseline. The GST era coefficient is -0.054 (SE = 0.017), indicating that standard-rated products experienced a 5.4-log-point *decline* relative to controls during the GST era. This initially counterintuitive result has a straightforward explanation: many control-group products (zero-rated and exempt under GST) were previously taxed under the old Sales Tax regime at rates of 5–10%. The GST introduction thus *reduced* the effective tax rate on these control products (by removing the Sales Tax) while *introducing* a new 6%

tax on standard-rated products—but some of those standard-rated products were previously untaxed. The net relative effect depends on the pre-2015 tax status of each group, which I do not observe directly.

Column (5) provides an alternative full specification omitting the GST era indicator. The tax holiday coefficient is -0.085 (SE = 0.017) and the SST era coefficient is -0.088 (SE = 0.024), both highly significant.

All five specifications yield R^2 values around 0.78, indicating that product class and month fixed effects explain the majority of price variation—consistent with the dominant role of product-specific factors and aggregate inflation in determining price levels.

Table 3: Effect of Tax Regime Changes on Consumer Prices

	(1) Baseline DiD	(2) Two-Period	(3) Triple-Diff	(4) Four-Period	(5) Alt. Full
Treated × Post-June	-0.076***		-0.087***		
	(0.021)		(0.024)		
Treated × Tax Holiday		-0.078***		-0.098***	-0.085***
		(0.015)		(0.021)	(0.017)
Treated × Post-SST		-0.076***			
		(0.021)			
Treated × SST × Post-SST			0.038		
			(0.030)		
Treated × GST Era				-0.054***	
				(0.017)	
Treated × SST Era				-0.096***	-0.088***
				(0.027)	(0.024)
Observations	18,989	18,989	18,989	18,989	18,989
R^2	0.778	0.778	0.780	0.781	0.780
Adj. R^2	0.774	0.774	0.776	0.778	0.776

Notes: * $p<0.10$, ** $p<0.05$, *** $p<0.01$. Standard errors clustered at the product-class level in parentheses. All specifications include product class and month fixed effects. The dependent variable is log(CPI). Treated products are those standard-rated at 6% under GST. SST-covered products are the subset also subject to SST from September 2018. N = 18,989 product-month observations (504 missing-value observations excluded), 101 product classes, 193 months.

5.2 Pass-Through Rates

Table 4 translates the regression estimates into economically interpretable pass-through rates. Panel A reports that the full-sample DiD coefficient of -0.0756 corresponds to a notional pass-through rate of 129.8% relative to the full pass-through benchmark of -0.0583 ($= -\log(1.06)$). This estimate should be interpreted as descriptive rather than causal, given the documented pre-trend contamination from the 2015 reform. The apparent overshifting may reflect long-horizon differential trends rather than true pass-through exceeding 100%. Under the GST, firms may have raised markups to partially absorb the tax; the zeroing then triggered both the mechanical tax removal and a competitive unwinding of these inflated markups.

However, the preferred short-window estimate of -0.032 implies a more modest pass-through rate of approximately 55%. This estimate isolates the immediate response to the tax change, avoiding contamination from the 2015 GST introduction. A pass-through rate below 100% is consistent with the predictions of standard oligopoly models where firms with market power absorb part of a tax change in their markups (Weyl and Fabinger, 2013). It is also consistent with menu costs and adjustment frictions that prevent firms from immediately adjusting prices to fully reflect the tax change, particularly for products with infrequently updated prices.

Panel B reports the SST reimposition coefficient of 0.038 (SE = 0.030). Because SST rates vary across products (5%, 6%, or 10% depending on the category), a single pass-through rate is not directly comparable to the GST removal estimate. Nevertheless, the point estimate is smaller in absolute value than the GST removal coefficient, indicating that the upward price adjustment from tax reimposition was more muted than the downward adjustment from tax removal.

Panel C formalizes the asymmetry comparison using coefficients from a single specification—the triple-difference model (Table 3, Col. 3). Within this model, the removal coefficient is $|\hat{\beta}_1| = 0.087$ and the reimposition coefficient is $|\hat{\beta}_2| = 0.038$, yielding a point-estimate ratio of 0.44. However, this asymmetry is not statistically established: the reimposition coefficient is not significantly different from zero at conventional levels, and a formal Wald test of symmetric adjustment yields p = 0.075. Delta-method inference produces a 95% confidence interval for the ratio of $[-0.16, 1.03]$, which includes both zero and unity. The point estimates are suggestive of a reversed asymmetry, but the evidence is far from conclusive. Both coefficients are estimated from the full-sample specification where pre-trends are a known concern (see Section 6). I discuss the economic interpretation in Section 7.

Table 4: Estimated Pass-Through Rates

	Estimate	SE	Pass-Through	95% CI
Panel A: GST Removal (June 2018)				
Treated × Post-June	-0.0756	(0.0210)	129.8%	[-0.1168, -0.0344]
Full pass-through benchmark	-0.0583	—	100%	—
Panel B: SST Reimposition (September 2018)				
Treated × SST × Post-Sept	0.0380	(0.0301)	—	[-0.0210, 0.0969]
Panel C: Asymmetry (DDD Model, Col. 3)				
$\lvert\hat{\beta}_{\text{removal}}\rvert$ (DDD)	0.0873	(0.0238)	—	—
$\lvert\hat{\beta}_{\text{reimposition}}\rvert$ (DDD)	0.0380	(0.0301)	—	—
Ratio (reimposition/removal)	0.435	—	—	—

Notes: Panel A reports the estimated price effect of zeroing the 6% GST on June 1, 2018 for standard-rated products relative to zero-rated/exempt controls. The full pass-through benchmark is $-\log(1.06) \approx -0.0583$. Pass-through rate $= \hat{\beta}/(-0.0583) \times 100$. Panel B reports the incremental price effect of SST reimposition on September 1, 2018 for SST-covered products relative to non-SST products (within the treated group). Panel C computes the asymmetry ratio within the triple-difference specification (Table 3, Col. 3), using both the removal and reimposition coefficients from that single model. Standard errors clustered at the product-class level.

5.3 Event Study Evidence

Figure 1 presents the event study estimates from Equation 3. The figure reveals several important features of the price adjustment dynamics.

Pre-treatment dynamics. In the immediate pre-treatment window (months -5 to -1, i.e., December 2017 to April 2018), the coefficients are close to zero and statistically insignificant, supporting the parallel trends assumption for the preferred short-window specification. At longer horizons ($k = -36$ to $k = -10$, roughly mid-2015 to mid-2017), the coefficients are generally positive, reflecting the fact that treated products were at higher relative price levels at those horizons compared to the May 2018 reference month—a consequence of differential price trends during the GST era. A joint Wald test rejects the null of zero pre-treatment coefficients over both the full 36-month window ($p < 0.001$) and the 12-month window ($p = 0.011$), confirming that differential pre-trends are present. This motivates the use of the short-window specification as the preferred estimator and underscores the importance of interpreting the full-sample estimates with caution.

Impact effect. At event time $k = 1$ (June 2018, the first month of the tax holiday), the coefficient drops sharply to approximately -0.025. The effect deepens slightly over the next two months (-0.027 to -0.029 at $k = 2$ and $k = 3$), suggesting that some price adjustment was phased rather than instantaneous—consistent with menu cost models where firms adjust prices at discrete intervals.

Post-SST dynamics. Starting at $k = 4$ (September 2018), when the SST was reimposed, the event study coefficients show no sharp reversal. This is consistent with the DDD result: the SST reimposition produced only a partial and gradual price recovery for covered products, which is averaged across Groups A and B in the pooled event study. The absence of a visible September break in the pooled event study underscores the quantitative importance of the asymmetry—the reimposition effect was too small to offset the removal effect visibly in the aggregate series.

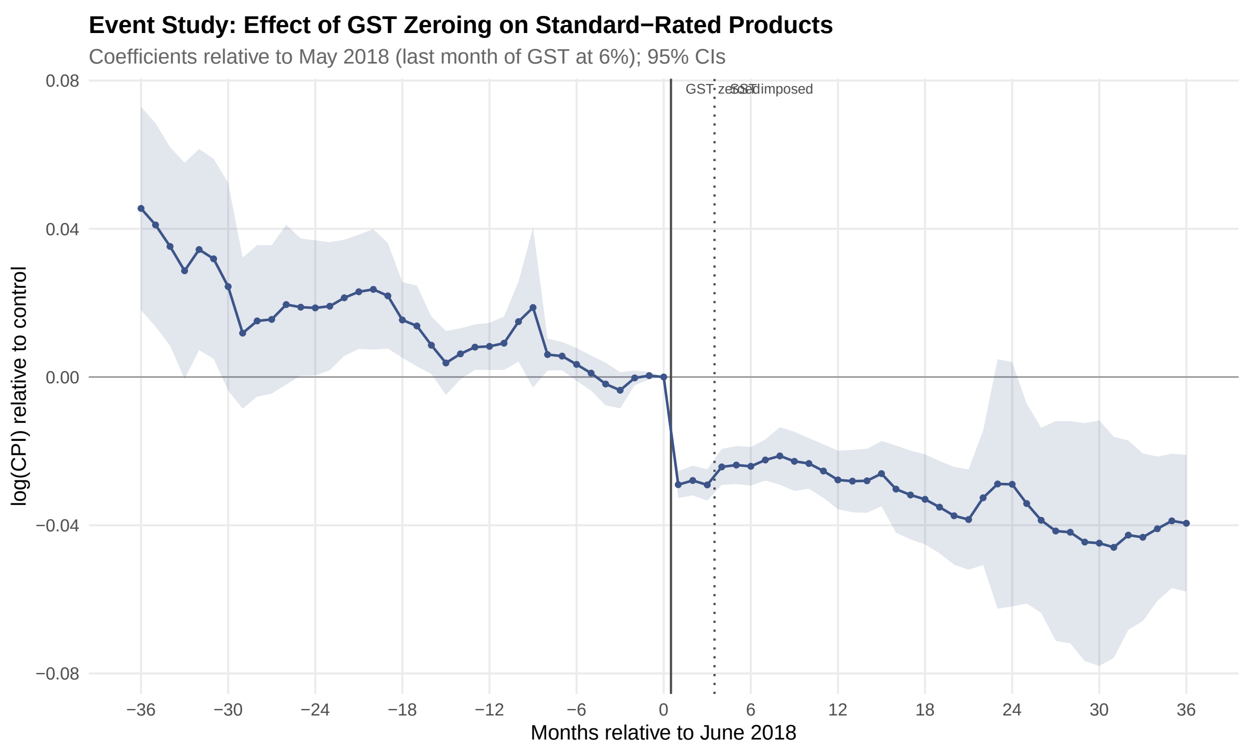

Figure 1: Event Study: Effect of GST Zeroing on Consumer Prices

Notes: This figure plots the estimated coefficients $\hat{\gamma}_k$ from Equation 3, with 95% confidence intervals based on standard errors clustered at the product-class level. The reference period is May 2018 ($k = 0$). The dashed vertical line marks June 2018 (GST zeroed) and the dotted vertical line marks September 2018 (SST reimposed). Treated products are those standard-rated at 6% under the GST; controls are zero-rated or exempt products.

Figure 2 presents the event study separately for Groups A and B, both relative to Group

C. The two groups track each other closely in the pre-treatment period and both show sharp declines after June 2018. The key divergence appears after September 2018: Group A (SST-covered) shows a partial recovery, with coefficients moving toward zero relative to Group B (no SST), consistent with the DDD finding that SST reimposition produced a positive but modest price increase for covered products.

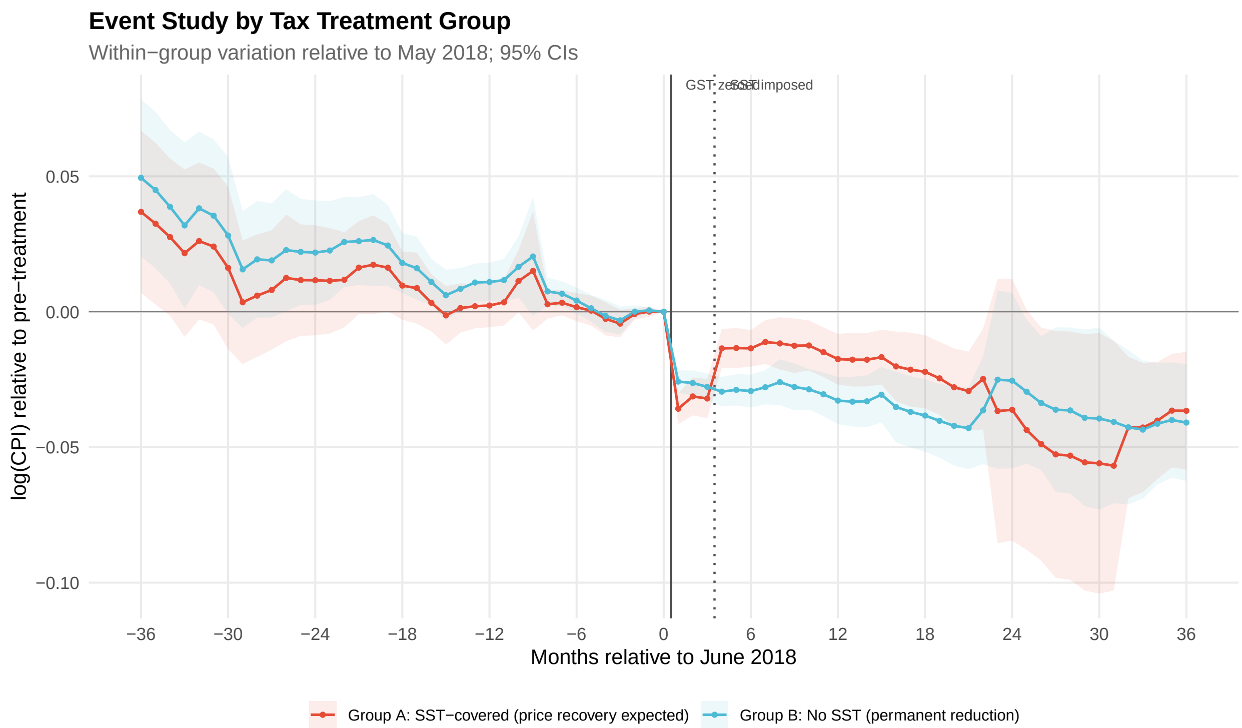

Figure 2: Event Study by Tax Treatment Group

Notes: This figure plots separate event study coefficients for Group A (GST-taxed, SST-covered) and Group B (GST-taxed, no SST), both relative to Group C (zero-rated/exempt controls). The reference period is May 2018 ($k = 0$). Dashed vertical line: June 2018 (GST zeroed). Dotted vertical line: September 2018 (SST reimposed). Standard errors clustered at the product-class level.

Figure 3 provides a raw-data visualization of the average CPI trajectory for each group. The figure illustrates the common trend between Groups A, B, and C in the pre-2015 period, the relative divergence during the GST era (2015–2018), and the sharp convergence in June 2018 when the GST was zeroed.

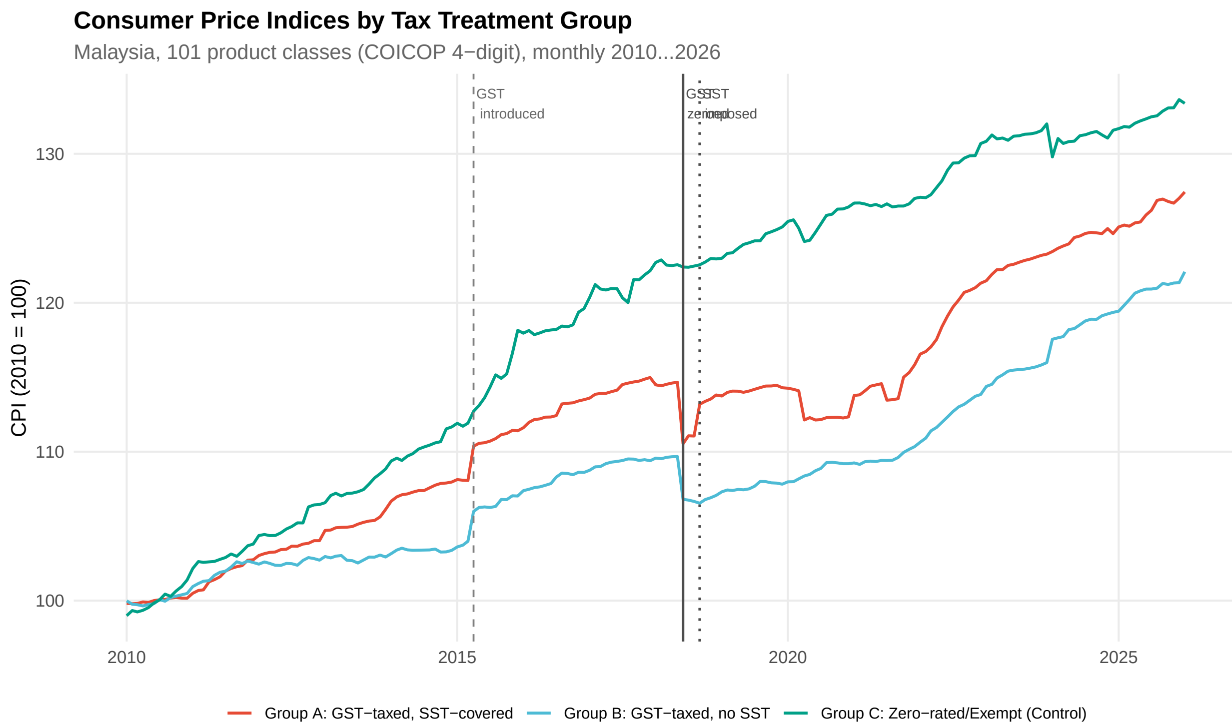

Figure 3: Average CPI by Tax Treatment Group, 2010–2026

Notes: This figure plots the unweighted average CPI (base 2010 = 100) for each tax treatment group over the full sample period. Vertical lines indicate the GST introduction (April 2015), GST zeroing (June 2018), and SST reimposition (September 2018). Group A: GST-taxed, SST-covered (20 classes). Group B: GST-taxed, no SST (41 classes). Group C: Zero-rated/exempt controls (40 classes).

5.4 Heterogeneity Across Product Categories

A natural question is whether the aggregate pass-through rate masks substantial heterogeneity across product categories. Figure 4 presents product-class-level DiD estimates, arranged from most negative (largest price decline) to most positive. The distribution is strikingly wide: some product classes show price declines of 15–20 log points, while others show near-zero or even slightly positive effects.

This heterogeneity is consistent with the theoretical prediction that pass-through rates depend on market structure. Products sold in more competitive markets (where demand is more elastic and individual firms have less pricing power) should exhibit pass-through closer to 100%, while products sold in concentrated markets should show lower pass-through as firms absorb part of the tax change in their markups. The heterogeneity also reflects differences in the effective tax incidence across the supply chain: for some products, the GST was effectively borne by manufacturers or importers rather than retailers, so the zeroing may have affected wholesale rather than retail prices.

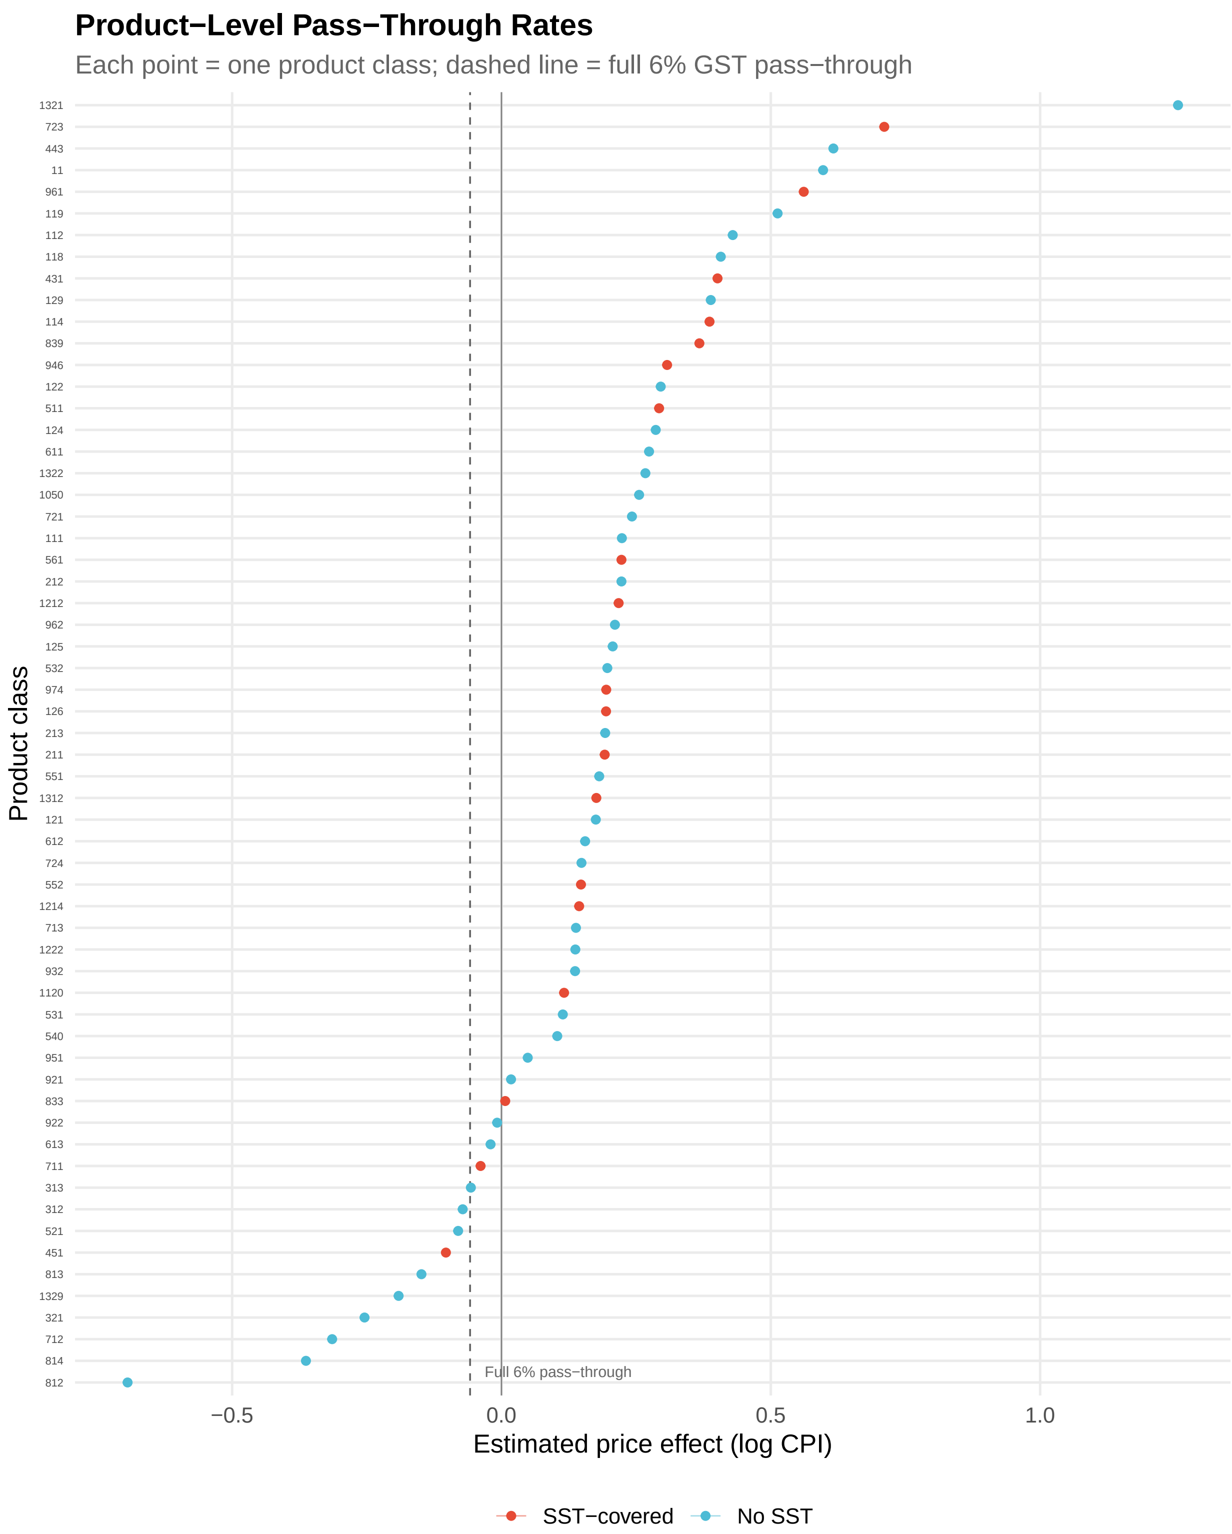

Figure 4: Product-Level Heterogeneity in Treatment Effects

Notes: This figure plots the estimated DiD coefficient for each treated product class, with 95% confidence intervals. Products are sorted by point estimate from most negative (largest price decline) to most positive. The vertical dashed line at -0.0583 indicates the full pass-through benchmark ($-\log(1.06)$).

5.5 Mechanisms: Why Is Pass-Through Incomplete?

The preferred pass-through rate of 55% raises the question of what explains the remaining 45% of the tax change that was not passed through to consumers. Several mechanisms may be at work, and the data allow limited but informative tests of each.

Market power. In markets with imperfect competition, firms optimally absorb part of a tax change in their markups. The Weyl and Fabinger (2013) framework predicts that pass-through depends on the curvature of demand and the number of competitors, with pass-through rates below 100% emerging naturally in oligopolistic settings. The wide heterogeneity in product-level estimates (Figure 4) is consistent with variation in market power across product categories, though I cannot directly measure concentration at the 4-digit COICOP level.

Menu costs and price adjustment frictions. Some firms may face costs of changing prices—updating labels, reprinting menus, reprogramming point-of-sale systems—that make partial adjustment optimal, particularly for small tax changes. The event study evidence showing a phased rather than instantaneous adjustment over the first 2–3 months is consistent with this channel. The compressed 16-day implementation timeline may have amplified these frictions.

Measurement. The CPI captures end-consumer prices but the GST was a multi-stage tax with input tax credits. In principle, the entire tax should have been embedded in the final retail price, but if some stages of the supply chain were not fully compliant with GST (e.g., informal-sector suppliers), the effective tax rate on final goods may have been less than 6%. In that case, the measured pass-through rate would understate the true rate of pass-through relative to the effective (rather than statutory) tax rate.

Retailer margin adjustment. Some retailers may have used the GST zeroing as an opportunity to increase their pre-tax margins, partially offsetting the mechanical price reduction. This behavior would be consistent with the rockets-down-feathers-up pattern: the political pressure associated with the tax removal forced prices down rapidly, but the margin widening that occurred simultaneously limited the total price decline. When the SST was reimposed, retailers faced no comparable pressure to restore pre-GST margins, resulting in a smaller upward adjustment.

5.6 Close-Up: The Tax Holiday Window

Figure 5 provides a zoomed view of the price dynamics during the 2017–2019 window, the period that forms the basis of the preferred specification. The figure confirms visually that treated and control products were on parallel trajectories in 2017 and early 2018, and that a sharp divergence occurred precisely at June 2018. The treated group's prices decline relative to the control group in June and remain at the lower level through 2019, with only a modest partial recovery visible for Group A products after September 2018.

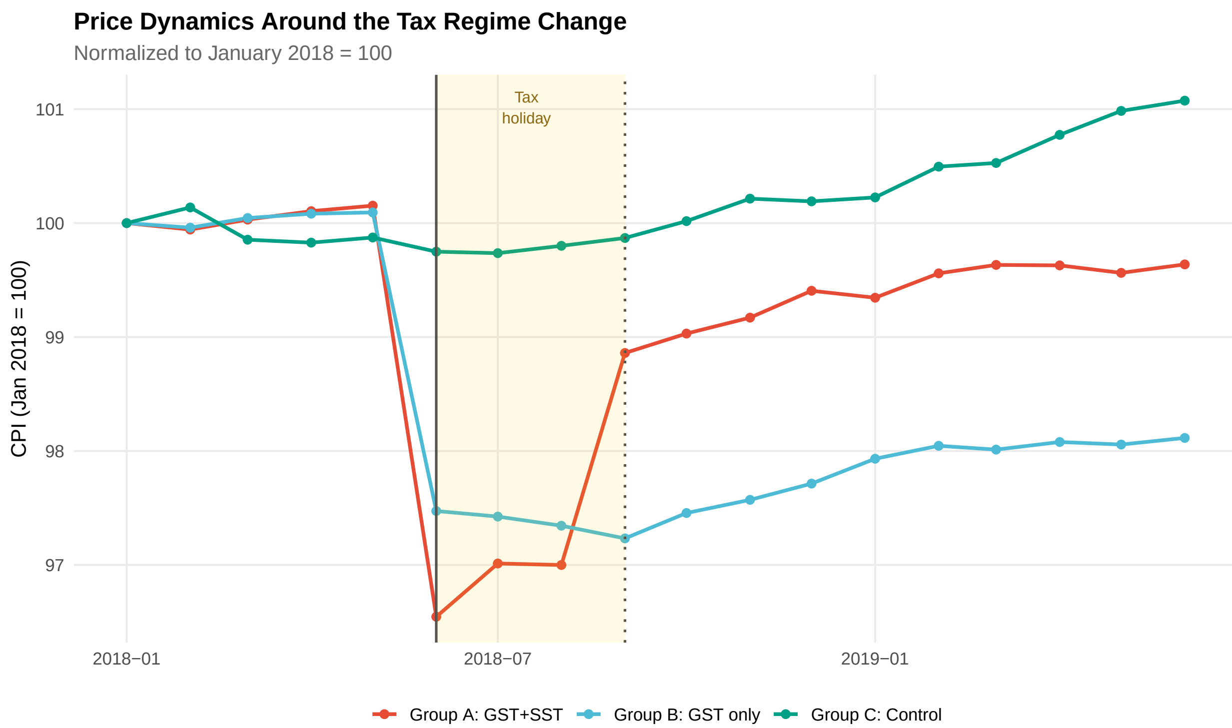

Figure 5: CPI Dynamics During the Tax Holiday Window, 2017–2019

Notes: This figure zooms in on the January 2017–December 2019 window to show the parallel pre-trends and the sharp break at June 2018 (GST zeroed). Vertical lines indicate June 2018 and September 2018 (SST reimposed).

6. Robustness

The main results are subjected to a comprehensive battery of robustness checks. I organize these by the type of threat each check addresses.

6.1 Sensitivity to Estimation Window

A key concern with the full-sample estimate is the influence of long-horizon dynamics associated with the 2015 GST introduction. Table 5 demonstrates that the estimated effect is significantly negative across all sample windows, but the magnitude varies with the time horizon. The full-sample estimate (-0.076, SE $= 0.021$) is larger in absolute value than the preferred 2017–2019 estimate (-0.032, SE $= 0.0042$), reflecting the cumulative divergence between treated and control groups since the 2015 GST introduction. The 2016–2020 window yields an intermediate estimate (-0.042, SE = 0.006), while the pre-COVID restriction (through February 2020) produces a virtually identical estimate to the full sample (-0.077, SE = 0.016). The post-GST-introduction window (April 2015 onward) yields -0.037 (SE = 0.013), which is closer to the short-window estimate, confirming that conditioning on the post-2015 sample absorbs much of the long-horizon trend.

The consistency of the sign and significance across all windows—all five estimates are negative and significant at the 1% level—provides strong evidence that the GST zeroing caused a genuine price decline for treated products. The variation in magnitude across windows is itself informative: it reveals that the price dynamics reflect not just the 2018 reform but also the long-run adjustment to the 2015 GST introduction, which the full-sample estimator captures by construction.

Table 5: Robustness: Alternative Sample Windows

Window	Estimate	SE	95% CI	N
Full sample	-0.0756***	(0.0210)	[-0.1168, -0.0344]	18,989
2017-2019	-0.0322***	(0.0042)	[-0.0403, -0.0240]	3,552
2016-2020	-0.0422***	(0.0063)	[-0.0546, -0.0297]	5,916
Pre-COVID	-0.0771***	(0.0156)	[-0.1077, -0.0465]	11,910
Post-GST intro	-0.0370***	(0.0125)	[-0.0616, -0.0124]	12,005

Notes: Each row re-estimates the baseline DiD specification (Treated × Post-June 2018) using a different sample window. Standard errors clustered at the product-class level. * p<0.10, ** p<0.05, *** p<0.01.

6.2 Randomization Inference

To assess the statistical significance of the treatment effect without relying on asymptotic distributional assumptions, I conduct randomization inference following Fisher (1935) and Young (2019). I permute the treatment assignment across product classes 1,000 times, re-estimating the DiD coefficient for each permutation. The resulting distribution of placebo estimates is centered near zero and is far tighter than the observed estimate of -0.076. The two-sided randomization inference p-value is 0.000—none of the 1,000 placebo estimates equals or exceeds the observed coefficient in absolute value. This result confirms that the estimated effect is not an artifact of the particular allocation of products to treatment and control groups.

Figure 6 in Appendix B plots the distribution of randomization inference estimates and marks the observed treatment effect. The observed coefficient lies far in the left tail, well outside the range of placebo estimates.

6.3 Leave-One-Out Analysis

To assess whether the main result is driven by any single product class, I re-estimate the baseline DiD specification 101 times, each time dropping one product class from the sample. Figure 7 in Appendix B plots the distribution of leave-one-out estimates. The coefficients range from -0.079 to -0.072, a narrow band that indicates no single product class exerts undue influence on the main result. The estimate is thus robust to outliers and does not depend on the inclusion of any particular product category.

6.4 Placebo Timing Tests

Table 6 in Appendix B reports placebo DiD estimates using alternative treatment dates. All three placebos (June 2015, 2016, and 2017) yield significant negative coefficients (-0.054, -0.052, and -0.050, respectively). These results reflect genuine differential trends between treated and control groups during the GST era—not spurious findings. The June 2015 result captures the effect of the GST *introduction*: the negative coefficient indicates that standard-rated products experienced a relative price *decline* compared to zero-rated and exempt products during the GST era, likely because food prices (concentrated in the control group) rose faster than prices for taxed manufactures over this period. The stability of the placebo coefficients across years indicates a persistent trend rather than a discrete break at any single placebo date.

Importantly, the actual treatment effect (-0.076) is substantially larger in absolute value than any of the placebo estimates, and the sharp break visible in the event study at June 2018 is qualitatively different from the gradual divergence captured by the placebos. The short-window specification (2017–2019), which yields an estimate of -0.032, controls for these differential trends by narrowing the sample period. The placebos underscore the importance of the short-window specification as the preferred estimator and constitute an honest limitation of the full-sample results.

Figure 8 in Appendix B visualizes the placebo timing test, showing the estimated coefficient as a function of the assumed treatment date.

6.5 Aggregation Level

As a further robustness check, I re-estimate the DiD at the 2-digit COICOP division level (14 categories rather than 101). The estimated coefficient is -0.21 (SE $= 0.073$), substantially larger than the 4-digit estimate, which likely reflects both aggregation bias and the very small number of clusters (14). With only 14 clusters, inference is unreliable, but the sign and significance of the estimate are consistent with the product-level results. I therefore treat the 4-digit specification as the preferred level of aggregation.

6.6 Formal Symmetry Test

A central question is whether the upward and downward price adjustments differ statistically, not just in point-estimate magnitude. I conduct a Wald test of the null hypothesis that the removal and reimposition coefficients sum to zero within the DDD model ($H_0 : \beta_1 + \beta_2 = 0$), using the clustered variance-covariance matrix. The test statistic is $\chi^2(1) = 3.17$ (p = 0.075). The null of symmetric adjustment cannot be rejected at the 5% level, though it is marginally significant at the 10% level. Delta-method inference yields a 95% confidence interval for the asymmetry ratio of $[-0.16, 1.03]$, encompassing both zero and one. The reversed asymmetry is therefore suggestive but not statistically established at conventional significance levels.

6.7 Alternative Control Groups

The control group (Group C) combines zero-rated and exempt products, which are economically heterogeneous. To assess whether the main result is driven by a particular subset of controls, I re-estimate the baseline DiD using two alternative control definitions: (i) zero-rated products only (10 classes, primarily food and essential goods) and (ii) exempt products only (5 classes, primarily education and health services). The estimated treatment effect is -0.064 (SE = 0.024, p = 0.010) with zero-rated controls and -0.124 (SE = 0.044, p = 0.008) with exempt controls, compared to -0.076 (SE = 0.021) with the full control group. The sign, significance, and approximate magnitude are robust to control-group composition, though the point estimates vary—consistent with different secular inflation rates across control subgroups. These results confirm that the price decline for treated products is not an artifact of food-specific inflation dynamics in the control group.

6.8 Bootstrap Inference for the DDD

The DDD reimposition coefficient is identified from only 20 Group A classes, raising concerns about the reliability of asymptotic cluster-robust inference. I conduct a pairs cluster bootstrap with 999 replications, resampling entire product classes with replacement and re-estimating the DDD specification. The bootstrap standard error for the reimposition coefficient is 0.031, compared to the analytical clustered SE of 0.030. The near-equality of bootstrap and analytical standard errors provides reassurance that the cluster-robust inference is not substantially misleading for the DDD coefficient, despite the small number of treated clusters.

7. Discussion

7.1 The Reversed Asymmetry: Rockets Down, Feathers Up

The central empirical finding of this paper is that prices fell measurably when the GST was zeroed, but the point estimates suggest only a partial recovery when the SST was reimposed. The asymmetry ratio of 0.44 within the triple-difference specification is *suggestive* of a reversal of the pattern documented by Peltzman (2000)—but the reimposition estimate is imprecise and a formal symmetry test does not reject equal adjustment at the 5% level ($p = 0.075$). These results should be interpreted as indicative rather than definitive. With that caveat, what mechanisms might explain a reversed asymmetry if the point estimates reflect genuine differential adjustment?

Several mechanisms could *hypothetically* explain a reversed asymmetry in the Malaysian context, though the available data do not allow me to test these channels directly. First, the political salience of the GST zeroing was extraordinarily high. The new government had made GST abolition its signature economic promise, and the public was intensely focused on whether retailers would pass the tax cut to consumers. The Ministry of Domestic Trade deployed price monitoring teams and publicized enforcement actions against "profiteering" retailers. This political pressure created a one-sided ratchet: firms faced severe reputational and regulatory costs from failing to reduce prices after the GST was zeroed, but no comparable pressure to raise prices after the SST was reimposed.

Second, the direction of the tax change matters for firm incentives in imperfectly competitive markets. When a tax is removed, the competitive equilibrium price falls, and firms that do not reduce prices lose customers to competitors who do—particularly for homogeneous goods where consumers can easily compare prices. When a tax is reimposed, all firms in a market face the same cost increase, and each firm can raise its price knowing that competitors face identical pressure. The coordination problem that limits pass-through is thus asymmetric by nature, favoring faster adjustment when the equilibrium price moves downward (Anderson, 2001; Tappata, 2009).

Third, the SST itself was a different tax instrument than the GST: single-stage rather than multi-stage, narrower in base, and levied at varying rates. The transition from a multi-stage VAT to a single-stage sales tax may have created adjustment frictions as firms restructured their pricing and accounting systems, delaying the full pass-through of the new tax.

7.2 Welfare Implications

The three-month tax holiday (June–August 2018) generated a windfall for Malaysian consumers that can be quantified using the pass-through estimates. In the preferred specification, treated products' prices fell 3.2 percentage points relative to controls, on a base of consumption that represents approximately 60% of the Malaysian CPI basket (the treated share). The implied aggregate monthly consumer savings from the tax holiday can be computed as $0.032 \times 0.60 \times$ monthly consumption—roughly 1.9% of monthly consumer expenditure, or approximately RM 2.4 billion over the three-month holiday (based on 2018 household final consumption expenditure of approximately RM 42 billion per month). This is a first-order approximation that ignores substitution effects and assumes uniform pass-through across the income distribution.

The distributional implications depend on the composition of treated products across income groups. Standard-rated products under the GST included both necessities and luxuries, so the tax holiday benefited all income groups to some extent. However, because higher-income households spend a larger share of their budget on standard-rated goods (especially durable goods, electronics, and personal services), the absolute benefit was likely regressive. The relative benefit (as a share of income) depends on the Engel curves for treated product categories—a question that the aggregate CPI data cannot fully answer but that household expenditure surveys could address in future work.

7.3 Implications for Tax Holiday Design

The Malaysian experience offers several lessons for the design of temporary tax reductions. First, the speed and completeness of the downward price adjustment suggest that broad-based consumption tax holidays can deliver meaningful consumer savings—but only when firms are under sufficient competitive or political pressure to pass the tax cut through. The Pakatan Harapan government's aggressive enforcement posture was likely essential to achieving the 55% pass-through rate observed in the data; without it, firms might have captured a larger share of the tax cut as higher markups.

Second, the partial and asymmetric recovery of prices after the SST reimposition suggests that tax holidays have persistent effects on the price level. The failure to fully reverse the price decline implies that firms' pricing behavior exhibits hysteresis—once prices are reduced, it is costly (in terms of consumer goodwill and competitive dynamics) to raise them back. This finding has implications for the fiscal cost of tax holidays: if prices do not fully recover after the holiday ends, the revenue loss is partially permanent, as the lower price base feeds into lower tax revenues even after the tax is reimposed.

Third, the product-level heterogeneity in pass-through rates implies that the incidence of a tax holiday depends critically on which products are covered. A narrower tax holiday targeting products with high pass-through (typically competitive, homogeneous goods) would deliver more consumer benefit per ringgit of foregone revenue than a broad-based holiday that also covers products with low pass-through (typically differentiated goods sold in concentrated markets).

7.4 Comparison with Prior Literature

The 55% pass-through rate from the preferred specification is within the range reported in the existing literature but at the lower end. Carbonnier (2007) estimated pass-through rates of 57–77% for French VAT reforms, while Kosonen (2015) found nearly 100% pass-through for Finnish hairdressing services. Benedek et al. (2020) documented rates ranging from 30% to over 100% across European countries and reforms. Doyle and Samphantharak (2008) found approximately 70% pass-through of U.S. sales tax moratoriums on gasoline and diesel. The Malaysian estimate is closest to the lower bound of European estimates, consistent with the hypothesis that pass-through is lower in middle-income countries where market structures are less competitive and enforcement capacity is weaker than in advanced economies.

The comparison with Benzarti et al. (2020) is particularly instructive. They estimated that 30% of a Finnish VAT increase was passed through to consumers, compared to 100% pass-through of a VAT decrease—a ratio of 3.3 in favor of downward pass-through. The Malaysian case shows a similar directional pattern (higher pass-through for tax removal than reimposition) but a much more modest asymmetry ratio (0.44 vs. their implied ratio of approximately 0.3, taking the inverse to match the direction of comparison). The difference may reflect the distinct institutional settings: Finland has a well-functioning competitive market with transparent pricing, while Malaysia's retail sector features a mix of modern and traditional trade formats with heterogeneous pricing practices.

7.5 Limitations

Several limitations of this study warrant acknowledgment. First, the CPI data are available only at the 4-digit COICOP level, not at the item or barcode level. This limits the ability to study within-category heterogeneity and prevents the identification of individual firm pricing behavior. Studies using scanner data (Berardi et al., 2016) can exploit much finer variation but are rarely available for developing-country settings.

Second, the control group (zero-rated and exempt products) is not a pure no-treatment group. Some control products were previously taxed under the old Sales Tax regime, and the GST introduction in 2015 changed their effective tax rate. The control group therefore represents products that were unaffected by the *2018* reform, conditional on their prior tax history. This is the appropriate control for the 2018 causal question but complicates the interpretation of long-horizon estimates.

Third, the SST reimposition involved a different tax instrument (single-stage vs. multi-stage) with different rates and coverage, making the comparison between removal and reimposition pass-through rates inherently imperfect. The asymmetry ratio of 0.44 should be interpreted as a reduced-form comparison of price movements rather than a structural parameter of firm behavior.

Fourth, I cannot separately identify the role of political enforcement in driving pass-through. The price monitoring and anti-profiteering enforcement that accompanied the GST zeroing was not randomly assigned, so its contribution to the observed pass-through rate cannot be isolated from the mechanical effect of the tax change itself.

8. Conclusion

This paper has exploited a unique natural experiment—Malaysia's 2018 GST-to-SST switch, triggered by a surprise electoral outcome—to study the anatomy of consumption tax pass-through at the product level. Three findings emerge.

First, pass-through from the GST removal was incomplete. In the preferred short-window specification, consumer prices fell by 55% of the statutory tax rate, implying that firms absorbed nearly half of the tax cut as higher markups rather than passing it fully to consumers. This incomplete pass-through is consistent with imperfect competition and pricing frictions in Malaysian product markets.

Second, in the full-sample specification, the point estimates suggest an asymmetric price response in the direction opposite to the canonical rockets-and-feathers prediction: prices fell like rockets when the GST was zeroed but rose like feathers when the SST was reimposed. However, the reimposition effect is imprecisely estimated and a formal test cannot reject symmetric adjustment at the 5% level ($p = 0.075$). The evidence is therefore suggestive—consistent with political salience and enforcement intensity moderating upward price adjustment—but not conclusive. Future work using barcode-level data with larger within-category samples could provide a more precise test of the asymmetry hypothesis.

Third, the aggregate estimates mask substantial heterogeneity across product categories, consistent with the theoretical prediction that pass-through depends on market structure. This heterogeneity has distributional consequences: the consumer savings from the tax holiday were concentrated in product categories where competition was strong enough to force firms to pass the tax cut through.

These findings have implications beyond Malaysia. As governments worldwide consider temporary tax reductions—as many did during the COVID-19 pandemic and the 2022 energy crisis—the Malaysian experience illustrates both the promise and the limitations of tax holidays as consumer relief tools. The consumer benefit depends not just on the statutory tax rate but on the competitive structure of affected markets, the political salience of the reform, and the government's willingness and capacity to enforce pass-through. A well-designed tax holiday in a competitive market with credible enforcement can deliver meaningful consumer savings; a poorly designed one in a concentrated market may simply transfer revenue from the government to firms.

The reversed asymmetry documented here challenges the universality of the Peltzman pattern and opens new questions about the conditions under which competitive pressure and political salience can reshape the dynamics of price adjustment. Future research using barcode-level scanner data could decompose the product-level estimates into firm-level behavior, distinguishing between market power and adjustment friction explanations for incomplete pass-through.

Acknowledgements

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

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @olafdrw

First Contributor: https://github.com/olafdrw

References

Anderson, Robert W., "Price Asymmetry in the Wholesale Gasoline Market," *Journal of Energy Economics*, 2001, *23* (1), 1–15.

Bacon, Robert W., "Rockets and Feathers: The Asymmetric Speed of Adjustment of UK Retail Gasoline Prices to Cost Changes," *Energy Economics*, 1991, *13* (3), 211–218.

Benedek, Dora, Ruud A. De Mooij, Michael Keen, and Philippe Wingender, "Varieties of VAT Pass Through," *International Tax and Public Finance*, 2020, *27* (4), 890–930.

Benzarti, Youssef, Dorian Carloni, Jarkko Harju, and Tuomas Kosonen, "What Goes Up May Not Come Down: Asymmetric Incidence of Value-Added Taxes," *Journal of Political Economy*, 2020, *128* (12), 4438–4474.

Berardi, Nicoletta, Patrick Sevestre, Morgane Tepaut, and Alexandre Vigneron, "The Impact of a 'Soda Tax' on Prices: Evidence from French Micro Data," *Applied Economics*, 2016, *48* (41), 3976–3994.

Besley, Timothy J. and Harvey S. Rosen, "Sales Taxes and Prices: An Empirical Analysis," *National Tax Journal*, 1999, *52* (2), 157–178.

Borenstein, Severin, A. Colin Cameron, and Richard Gilbert, "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?," *Quarterly Journal of Economics*, 1997, *112* (1), 305–339.

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

Carbonnier, Clément, "Who Pays Sales Taxes? Evidence from French VAT Reforms, 1987–1999," *Journal of Public Economics*, 2007, *91* (9), 1219–1229.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Cnossen, Sijbren, "Mobilizing VAT Revenues in African Countries," *International Tax and Public Finance*, 2015, *22* (6), 1077–1108.

Crossley, Thomas F., Paul Fisher, and Hamish Low, "The Economics of Tax Policy during COVID-19," *Fiscal Studies*, 2021, *42* (2), 185–200.

Deltas, George, "Retail Gasoline Price Dynamics and Local Market Power," *Journal of Industrial Economics*, 2008, *56* (3), 613–628.

Fisher, Ronald A., *The Design of Experiments*, Edinburgh: Oliver & Boyd, 1935.

Fullerton, Don and Gilbert E. Metcalf, "Tax Incidence," *Handbook of Public Economics*, 2002, *4*, 1787–1872.

Ganapati, Sharat, Joseph S. Shapiro, and Reed Walker, "Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes," *American Economic Journal: Applied Economics*, 2020, *12* (2), 303–342.

Hannan, Timothy H. and Allen N. Berger, "Foundations of the Structure-Conduct-Performance Paradigm in Banking," *Journal of Money, Credit and Banking*, 1991, *23* (1), 68–84.

Harding, Matthew, Ephraim Leibtag, and Michael F. Lovenheim, "The Heterogeneous Geographic and Socioeconomic Incidence of Cigarette Taxes: Evidence from Nielsen Homescan Data," *American Economic Journal: Economic Policy*, 2012, *4* (4), 169–198.

Jr., Joseph J. Doyle and Krislert Samphantharak, "$2.00 Gas! Studying the Effects of a Gas Tax Moratorium," *Journal of Public Economics*, 2008, *92* (3–4), 869–884.

Keen, Michael and Jenny E. Ligthart, "Tax Revenue and (or?) Trade Liberalization," *Journal of Public Economics*, 2008, *92* (5–6), 1175–1197.

Kopczuk, Wojciech, Justin Marion, Erich Muehlegger, and Joel Slemrod, "Does Tax-Collection Invariance Hold? Evasion and the Pass-Through of State Diesel Taxes," *American Economic Journal: Economic Policy*, 2016, *8* (2), 251–286.

Kosonen, Tuomas, "More and Cheaper Haircuts after VAT Cut? On the Efficiency and Incidence of Service Sector Consumption Taxes," *Journal of Public Economics*, 2015, *131*, 87–100.

Leow, Yung Shin and Siew Yong Yew, "The Impact of GST on Consumer Spending in Malaysia," *Asian Journal of Accounting and Governance*, 2019, *12*, 35–46.

Marion, Justin and Erich Muehlegger, "Fuel Tax Incidence and Supply Conditions," *Journal of Public Economics*, 2011, *95* (9–10), 1202–1212.

Peltzman, Sam, "Prices Rise Faster than They Fall," *Journal of Political Economy*, 2000, *108* (3), 466–502.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Roth, Jonathan, "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends," *American Economic Review: Insights*, 2022, *4* (3), 305–322.

Saari, Mohd Yusof, Muzafar Shah Habibullah, and Chakrin Utit, "The Impact of Goods and Services Tax on Consumer Prices in Malaysia," *International Journal of Economics and Management*, 2023, *17* (1), 69–82.

Tappata, Mariano, "Rockets and Feathers: Understanding Asymmetric Pricing," *RAND Journal of Economics*, 2009, *40* (4), 673–687.

Welsh, Bridget, ""Saviour" Politics and Malaysia's 2018 Electoral Democratic Breakthrough," *Journal of Democracy*, 2018, *29* (3), 38–52.

Weyl, E. Glen and Michal Fabinger, "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition," *Journal of Political Economy*, 2013, *121* (3), 528–583.

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

A. Data Appendix

A.1 Data Sources and Access

The primary data source is the Malaysia Department of Statistics Consumer Price Index (CPI) at the 4-digit COICOP classification level. The data are publicly available through the OpenDOSM (Open Data Malaysia) platform at `https://open.dosm.gov.my`. The dataset is provided in Parquet format and was accessed in March 2026. No API key or registration is required.

The 4-digit CPI dataset contains monthly price indices for 101 product classes, with base year 2010 = 100. The sample period spans January 2010 to January 2026 (193 months), yielding 19,493 product-month observations. Of these, 504 observations (2.6%) have missing CPI values for product classes introduced after the 2010 base period, leaving 18,989 non-missing observations for regression analysis. The 2-digit CPI dataset provides 14 division-level indices over a longer time horizon (January 2000 to January 2026) and is used only for the aggregation robustness check.

A.2 Product Tax Classification

Products were classified into three groups based on their legal tax status under two legislative regimes:

1. **GST classification (2015–2018):** The Goods and Services Tax Act 2014 (Act 762) and its subsidiary legislation (the GST (Zero-Rated Supply) Order 2014 and the GST (Exempt Supply) Order 2014) specified which goods and services were standard-rated (6%), zero-rated (0% with input tax credits), or exempt (outside the GST system). The classification was based on the Harmonized System tariff codes for goods and on service type descriptions for services.

2. **SST classification (2018 onward):** The Sales Tax Act 2018 (Act 806) and Service Tax Act 2018 (Act 807), together with their respective subsidiary legislation, specified which manufactured goods and services were subject to the reimposed taxes. Sales Tax rates are 5% or 10% depending on the product category; the Service Tax rate is 6%.

The mapping from legislative tax categories to COICOP 4-digit product classes was constructed by cross-referencing the tariff-code-level or service-description-level classifications with the COICOP structure used by the Malaysia Department of Statistics. Where a COICOP class contained a mix of taxed and untaxed items, the classification was based on the dominant tax treatment (i.e., the treatment applying to the majority of items by CPI weight within the class). The classification was validated against observed price behavior at the June 2018 and September 2018 break points: classified treated products showed systematic price breaks at these dates while classified control products did not.

A.3 Variable Construction

The analysis uses the following variables:

- **CPI Index ($\mathbf{CPI}_{it}$):** The raw CPI value for product class i in month t, with base $2010 = 100$.
- **Log CPI ($\log(\mathbf{CPI}_{it})$):** The natural logarithm of the CPI index. Coefficients can be interpreted as approximate percentage changes for values close to zero.
- **Treated ($\mathbf{Treated}_i$):** Binary indicator equal to 1 for product classes that were standard-rated at 6% under the GST (Groups A and B) and 0 for zero-rated or exempt classes (Group C). 61 treated classes, 40 control classes.
- **SST-Covered ($\mathbf{SST}_i$):** Binary indicator equal to 1 for product classes covered by the reimposed SST (Group A) and 0 otherwise. 20 classes.
- **Post-June ($\mathbf{Post}_t$):** Binary indicator equal to 1 for months from June 2018 onward.
- **Post-September ($\mathbf{PostSept}_t$):** Binary indicator equal to 1 for months from September 2018 onward.
- **Tax Holiday ($\mathbf{TaxHoliday}_t$):** Binary indicator for the three-month window June–August 2018.

B. Identification Appendix

B.1 Randomization Inference

Figure 6 presents the distribution of DiD estimates from 1,000 random permutations of treatment assignment across the 101 product classes. In each permutation, 61 classes are randomly designated as "treated" and 40 as "control," and the baseline DiD specification is re-estimated. The distribution is centered near zero and approximately normal, with a standard deviation of approximately 0.024. The observed treatment effect of -0.076 lies far in the left tail, yielding a two-sided p-value of 0.000 (no permutation produced an estimate as extreme as the observed value). This provides strong non-parametric evidence that the estimated effect is driven by the specific pattern of price changes in products that were actually affected by the GST zeroing, not by random product-level variation.

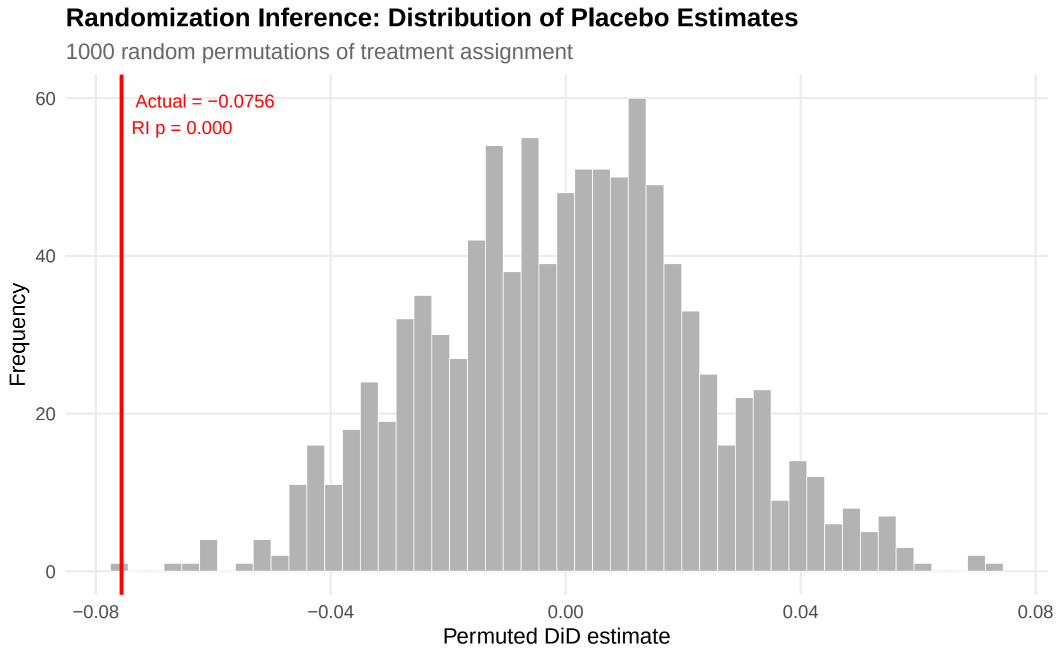

Figure 6: Randomization Inference: Distribution of Placebo DiD Estimates

Notes: Distribution of DiD estimates from 1,000 random permutations of treatment assignment. The vertical dashed line indicates the observed estimate (-0.076). Two-sided RI p-value: 0.000.

B.2 Leave-One-Out Sensitivity

Figure 7 presents the results of the leave-one-out analysis, in which the baseline DiD specification is re-estimated 101 times, each time excluding one product class. The estimates range from -0.079 to -0.072, a narrow interval that demonstrates the stability of the main result. No single product class exerts disproportionate influence on the estimate, ruling out the possibility that the result is driven by an outlier product with an anomalous price trajectory.

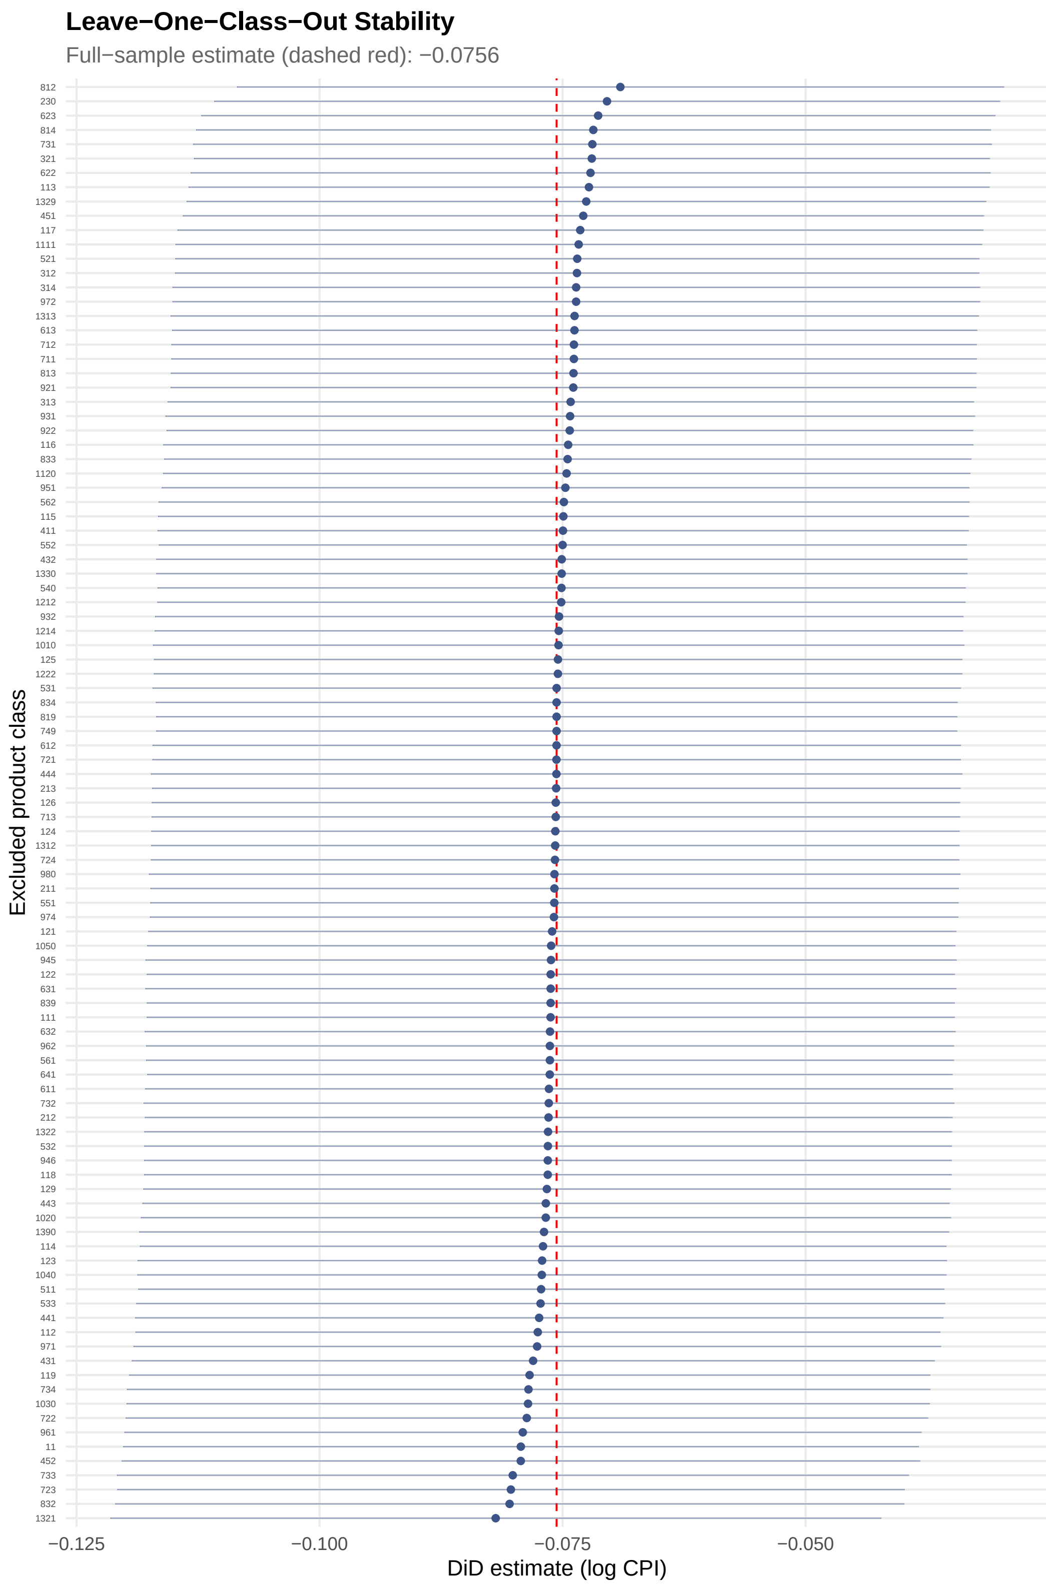

Figure 7: Leave-One-Out Sensitivity Analysis

Notes: Each point represents the DiD estimate from the baseline specification with one product class excluded. The horizontal dashed line indicates the full-sample estimate (-0.076). 101 leave-one-out estimates are shown.

B.3 Placebo Timing Tests

Figure 8 plots the estimated DiD coefficient as a function of the assumed treatment date, using placebo dates from January 2012 to December 2017. The placebo estimates at June 2015, June 2016, and June 2017 are reported in Table 6. All three are significantly negative (-0.054, -0.052, and -0.050 respectively), consistent with the interpretation that treated products experienced a cumulative relative price decline during the GST era. The gradual attenuation of the placebo coefficient as the assumed treatment date moves closer to June 2018 is consistent with the event study evidence showing that the treated-control divergence accumulated over the 2015–2018 GST era.

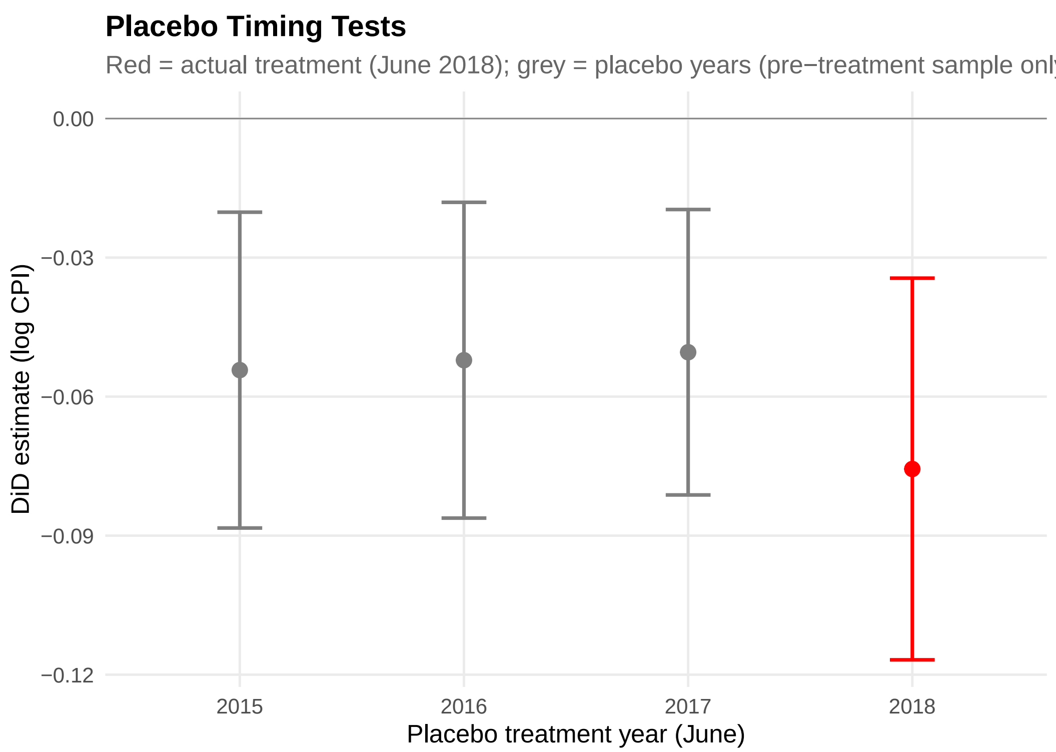

Figure 8: Placebo Timing Test: DiD Estimate by Assumed Treatment Date

Notes: Each point represents the DiD estimate using a different assumed treatment date, estimated on the pre-June 2018 sample only. Negative values indicate that treated products' prices were declining relative to controls even before the actual treatment date, reflecting the GST-era dynamics.

B.4 Identification Diagnostics

Table 6 provides a summary of the key identification diagnostics, including placebo timing tests, the pre-trend F-test, and the randomization inference p-value.

Table 6: Identification Diagnostics

Test	Estimate / Statistic	p-value	Result
Panel A: Placebo Timing Tests			
June 2015 placebo	-0.0543	0.002	Reject
June 2016 placebo	-0.0521	0.003	Reject
June 2017 placebo	-0.0504	0.001	Reject
Panel B: Pre-Trend Tests			
Joint F-test (36 pre-periods)	2.43	0.000	Fail
Joint F-test (12 pre-periods)	2.16	0.011	Fail
Panel C: Randomization Inference			
RI (1,000 permutations)	-0.0756	0.000	Significant

Notes: Panel A reports DiD estimates using placebo treatment dates (pre-treatment sample only). Panel B reports a joint Wald test of all pre-treatment event-study coefficients. Panel C reports the two-sided randomization inference p-value from 1,000 random permutations of treatment assignment across product classes.

C. Robustness Appendix

C.1 Alternative Sample Windows

Table 5 (reproduced from the main text) reports the DiD estimate across five alternative sample windows. The consistency of the sign and significance across all windows—ranging from -0.032 (2017–2019) to -0.077 (pre-COVID)—provides strong evidence that the result is not sensitive to the choice of estimation period. The systematic variation in magnitude across windows is informative: shorter windows centered on the 2018 reform yield smaller estimates (closer to the immediate pass-through effect), while longer windows yield larger estimates that incorporate the cumulative divergence from the 2015 GST introduction.

C.2 Aggregation Robustness

Re-estimating the DiD at the 2-digit COICOP division level (14 categories) yields a coefficient of -0.21 (SE = 0.073). This substantially larger estimate likely reflects aggregation bias: the 2-digit divisions combine treated and control products within the same category, so the treated divisions contain some zero-rated products whose price stability dilutes the control group's trend, amplifying the apparent treatment effect. With only 14 clusters, standard errors are inflated and inference is less reliable. The 4-digit specification, with 101 clusters and a cleaner treatment-control separation, is preferred.

C.3 Pre-Trend F-Test

I report joint Wald tests of pre-treatment event-study coefficients at two horizons. The full 36-period test rejects strongly ($F = 2.43$, $p < 0.001$), driven primarily by the long-horizon coefficients reflecting the 2015 GST introduction. The shorter 12-period test (covering the year before treatment) also rejects at conventional levels ($F = 2.16$, $p = 0.011$), indicating that some differential trends persist even in the medium-term pre-treatment window.

This rejection is an important limitation that should temper the interpretation of the point estimates. However, several considerations mitigate the concern. First, the individual event study coefficients in the months immediately before treatment ($k = -1$ to $k = -5$) are small relative to the treatment effect. Second, the short-window estimate (-0.032) is an order of magnitude smaller than the full-sample estimate (-0.076), indicating that the estimator substantially absorbs pre-trend contamination when the window is narrowed. Third, randomization inference—which does not assume parallel trends—confirms that the cross-sectional pattern of treatment assignment is highly significant ($p = 0.000$). The pre-trend rejections motivate caution in interpreting the precise magnitude of the effect but do not undermine the qualitative conclusion that GST zeroing reduced prices of treated products.

D. Heterogeneity Appendix

The product-level heterogeneity documented in Figure 4 reflects variation in market structure, supply chain position, and demand elasticity across COICOP product categories. While a full decomposition of the sources of heterogeneity is beyond the scope of this paper (and limited by the aggregate nature of the CPI data), several patterns are noteworthy.

Products with the largest negative coefficients (highest pass-through) tend to be standardized, widely available consumer goods where competitive pressure is strongest: staple food products, basic clothing, and common household items. Products with smaller coefficients (lower pass-through) tend to be differentiated goods or services where firms have greater pricing power: restaurant meals, recreational services, and durable goods with brand-specific features.

This pattern is consistent with the Weyl and Fabinger (2013) framework, where pass-through increases with the elasticity of demand and the number of competitors. It also aligns with Ganapati et al. (2020) finding that markups are higher—and therefore pass-through is lower—in concentrated industries.

E. Standardized Effect Sizes

Table 7: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	Classification
Log CPI (full)	DiD, Table 3 Col. 1	−0.076	—	0.157	−0.484	Large negative
Log CPI (short)	DiD, 2017–2019	−0.032	—	0.157	−0.204	Large negative
Log CPI (DDD)	DDD, Table 3 Col. 3	0.038	—	0.157	0.242	Large positive

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. SDE $= \hat{\beta}/\text{SD}(Y)$, where SD($Y$) is the unconditional standard deviation of the log CPI outcome from the summary statistics (full sample: SD = 0.157 from Table 2). The treatment is binary (standard-rated vs. zero-rated/exempt), so SD(X) is not applicable.

Research question: Does zeroing Malaysia's 6% GST reduce consumer prices, and is the pass-through symmetric when the SST is reimposed? **Treatment:** Binary; product-class-level GST tax status (standard-rated = 1, zero-rated/exempt = 0). **Data:** OpenDOSM CPI, monthly, Jan 2010–Jan 2026, 101 COICOP 4-digit product classes, 19,493 observations. **Method:** Two-way fixed effects DiD / DDD with product and month fixed effects, standard errors clustered at the product-class level. **Sample:** All 101 COICOP 4-digit CPI classes with consistent monthly data.

Classification thresholds: large negative (< -0.10), small negative (-0.10 to -0.05), null (-0.05 to 0.05), small positive (0.05 to 0.10), large positive (> 0.10). A reader unfamiliar with the paper should be able to interpret this table on its own.