Eat In or Take Out? Tax Pass-Through at Japan's Dual-Rate Consumption Tax Boundary

APEP Autonomous Research* ai1scl
@ai1scl

March 10, 2026

Abstract

When Japan raised its consumption tax from 8% to 10% in October 2019, it exempted takeout food at the 8% rate while taxing eat-in food at 10%—creating a price wedge of 1.85%. Using monthly aggregate CPI data for food categories facing different tax schedules (2015–2024), I estimate the differential price impact. The CPI for eating out rose 1.86 percentage points more than the CPI for cooked food at impact, consistent with near-complete pass-through. Using Newey-West inference, the pre-COVID estimate of 0.0204 log points ($p < 0.001$) cannot reject the full-pass-through benchmark ($p = 0.19$). The effect is robust to a full placebo-in-time distribution, bandwidth choice, and COVID controls. These aggregate results are consistent with full pass-through of a novel, location-based tax differential in a competitive retail food market, though the evidence is category-level rather than item-level.

JEL Codes: H22, H25, H32, D12
Keywords: consumption tax, VAT pass-through, reduced tax rates, tax salience, Japan, food taxation

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

1. Introduction

A single rice ball sits on a convenience store counter. If the customer says "take out," the tax is 8%. If the customer says "eat here," the tax is 10%. Same product, same store, same moment—different price. Japan's 2019 dual-rate consumption tax reform created exactly this situation for millions of food transactions daily, producing an unusual natural experiment for studying indirect tax incidence.

This paper uses aggregate CPI category data to test whether Japan's novel location-dependent tax differential passed through to consumer prices. On October 1, 2019, Japan raised its consumption tax from 8% to 10% but simultaneously introduced a reduced rate: prepared food and non-alcoholic beverages sold for takeout remained at 8%, while the same items consumed on premises were taxed at the full 10% rate (Tax Commission, Government of Japan, 2018). Alcoholic beverages, by contrast, faced the uniform 10% rate regardless of consumption location. The resulting tax wedge between eat-in and takeout food is $2/108 \approx 1.85\%$—small enough to be plausibly absorbed by sellers, yet large enough to detect with monthly price index data.

The central finding is that aggregate CPI evidence is consistent with near-complete pass-through. In the month of the reform, the CPI for eating out rose 1.86 percentage points more than the CPI for cooked food (a category dominated by takeout-eligible prepared meals), against a predicted differential of 1.85% under full pass-through. A difference-in-differences specification with month fixed effects and Newey-West standard errors confirms this: the pre-COVID estimate is 0.0204 log points ($p < 0.001$), and a formal test cannot reject the full-pass-through benchmark of 0.0183 ($p = 0.19$). The result is robust across bandwidth choices from ± 12 to ± 48 months, survives a full placebo-in-time distribution test, and persists when the sample is restricted to a narrow window around the reform (October 2017–January 2020).

The identification strategy is best understood as a comparative interrupted time-series design on national CPI aggregates (Bertrand et al., 2004). The main specification compares the log relative price of eating out to cooked food before and after October 2019, exploiting the fact that these broad CPI categories roughly straddle the 8%/10% tax boundary. Supplementary analyses include a triple-difference design using alcoholic beverages (which faced the uniform 10% rate increase) and a panel event study with semi-annual bins showing sharp onset at the reform date. The pre-COVID triple difference estimate of 0.0205 log points ($p < 0.001$) confirms that the relative price shift is not merely an artifact of the general tax increase.

This paper contributes to several literatures. The first is the large body of work on VAT pass-through. Poterba (1996) established that US state sales tax changes pass through approximately one-for-one to retail prices. Subsequent work has documented variation in pass-through across market structures (Weyl and Fabinger, 2013; Benzarti and Carloni, 2019), product types (Kosonen, 2015; Harju et al., 2018), and whether taxes are increasing or decreasing (Benzarti et al., 2020). My contribution is to provide aggregate evidence consistent with full pass-through for a novel, location-dependent tax differential—a setting where pass-through could plausibly be partial if sellers found the dual-rate system too complex to implement or feared losing customers to "eat-in tax evasion." The evidence is category-level rather than item-level, but the magnitude and timing of the relative price shift are strikingly close to the theoretical benchmark.

Second, the paper speaks to the tax salience literature (Chetty et al., 2009; Finkelstein, 2009; Taubinsky and Rees-Jones, 2018). Japan's dual-rate system tests whether firms treat a small, context-dependent tax wedge as salient enough to price differentially. The answer is unambiguous: retailers adjusted prices immediately and completely. This contrasts with settings where tax-inclusive pricing dampens consumer responses (Chetty et al., 2009) and suggests that when the tax change is well-publicized and applies to an entire consumption category, firms face strong incentives to pass it through regardless of its novelty.

Third, the paper contributes to the policy debate on reduced VAT rates. The Mirrlees et al. (2011) review and subsequent work (Keen, 2004; Crossley et al., 2009) have questioned whether reduced rates achieve their redistributive goals, given that firms may absorb the tax benefit rather than passing it to consumers. Japan's experience shows the opposite concern: full pass-through of the *differential* between standard and reduced rates means consumers face exactly the price wedge the government intended. Whether this price wedge achieves its distributional or behavioral objectives is a separate question that my data cannot address directly.

Finally, this is among the first empirical studies of Japan's 2019 dual-rate reform using CPI microdata. While Asano and Hara (2020) documented the overall price impact of the tax increase, and the OECD noted the reform's unusual design (OECD, 2020), no prior study has estimated the differential pass-through between categories facing different rates within the consumption tax system. Japan's reform is globally distinctive: no other country taxes identical food items at different rates based solely on where they are consumed, making this a uniquely informative experiment for understanding how firms respond to novel tax boundaries.

2. Institutional Background

2.1 Japan's Consumption Tax

Japan's consumption tax (*shōhi-zei*) is a broad-based value-added tax levied at each stage of production and distribution. Introduced in April 1989 at 3%, it was raised to 5% in April 1997 and to 8% in April 2014. Each increase generated significant public debate and measurable macroeconomic effects, with the 2014 increase widely blamed for tipping the economy into recession (OECD, 2020).

Unlike European VAT systems, which have long maintained multiple rates (standard, reduced, zero-rated, exempt), Japan operated a single-rate system for its first 30 years. This simplicity was considered a virtue: administrative costs were low, cascading exemptions were avoided, and the base was broad. The Tax Commission repeatedly rejected proposals for a multi-rate system through the 2000s, arguing that the compliance costs and boundary-drawing problems would outweigh any distributional gains (Tax Commission, Government of Japan, 2018). The introduction of a reduced rate in 2019 therefore represented a fundamental departure from the system's design philosophy.

Japan's consumption tax is notable for two institutional features relevant to this study. First, Japan uses *sou-gaku hyouji* (tax-inclusive pricing): retail prices displayed to consumers include the consumption tax. This means that the price a consumer sees at a convenience store already reflects the applicable tax rate. When the tax increased, retailers had to adjust their displayed prices. This contrasts with the US sales tax system, where the tax is added at checkout and may be less salient to consumers (Chetty et al., 2009). The tax-inclusive pricing convention ensures that the eat-in/takeout price differential is immediately visible to consumers at the point of purchase.

Second, the food service industry in Japan is extraordinarily competitive. As of 2019, Japan had approximately 670,000 restaurants and over 56,000 convenience stores nationwide. The convenience store sector alone—dominated by three chains (7-Eleven, FamilyMart, Lawson)—generates approximately 12 trillion yen in annual sales, with food accounting for the majority. Many convenience stores are located within 100 meters of each other in urban areas, creating intense price competition. This market structure makes partial pass-through unlikely: a retailer who absorbed the tax differential would be undercut by competitors who passed it through, since consumers can easily compare tax-inclusive prices across nearby stores.

2.2 The October 2019 Dual-Rate Reform

On October 1, 2019, the consumption tax rose from 8% to 10%. Simultaneously, a reduced rate of 8% was introduced for two categories: (i) food and non-alcoholic beverages intended for human consumption, excluding alcohol and dining on premises; and (ii) newspaper subscriptions delivered at least twice weekly. The stated purpose was to cushion the impact on low-income households, for whom food represents a larger budget share.

The critical design feature—and the source of this paper's identification—is the location-dependent boundary within prepared food. A meal purchased at a restaurant, food court, or convenience store eating area is taxed at 10%. The identical meal taken to go is taxed at 8%. The determination is made at the point of sale, based on the customer's stated intention. This created an immediate and well-publicized price differential for any food item that could be consumed either on or off premises.

The tax wedge between eat-in and takeout pricing is:

$$\frac{1+0.10}{1+0.08} - 1 = \frac{1.10}{1.08} - 1 = \frac{2}{108} \approx 0.0185 \tag{1}$$

Under full pass-through, eat-in prices should rise approximately 1.85% more than takeout prices at the reform date. This prediction applies to the pre-tax price staying constant, with the tax differential fully reflected in the consumer-facing price.

2.3 The Cashless Payment Rebate

Simultaneously with the tax increase, the government launched a cashless payment rebate program offering up to 5% point rebates on purchases at small and medium-sized retailers using registered cashless payment methods. This program ran from October 2019 through June 2020. While this could potentially confound the analysis by differentially affecting payment patterns at eat-in versus takeout establishments, the rebate applied to both eat-in and takeout purchases at eligible retailers and therefore does not create a differential effect along the eat-in/takeout margin. Nevertheless, I discuss this potential confound in Section 5.4.

2.4 Compliance and "Eat-In Tax Evasion"

The dual-rate system created a social phenomenon known colloquially as "eat-in tax evasion" (*iito-in dassuzei*): customers who declared their purchase as takeout but then consumed it on premises, thereby paying the lower 8% rate. Media coverage in the months following the reform extensively documented this behavior. Major convenience store chains (7-Eleven, FamilyMart, Lawson) responded by posting signs reminding customers that eating-area purchases are subject to the 10% rate, but enforcement was largely based on the honor system.

The prevalence of eat-in tax evasion is relevant because it could attenuate the measured price differential if retailers anticipated that customers would game the system and therefore chose not to maintain separate eat-in and takeout prices. The finding of complete pass-through suggests that, at the aggregate level, retailers did maintain the price differential despite the compliance challenge.

2.5 COVID-19

The reform's timing proved unfortunate: barely four months later, Japan confronted the COVID-19 pandemic. The first domestic case was confirmed on January 16, 2020, and a national state of emergency was declared on April 7, 2020. The pandemic fundamentally altered eating patterns—restaurants closed or reduced capacity, and takeout demand surged. This creates a significant identification challenge, as the pandemic's effects on relative eat-in versus takeout prices could confound the tax effect.

I address this in three ways. First, I estimate the main specifications on the pre-COVID window (October 2019–January 2020), which provides four clean post-treatment months. Second, I control explicitly for a COVID indicator (February 2020 onward) in the full-sample regressions. Third, I show that the immediate impact at October 2019—the month of the reform and three months before any COVID effects—is consistent with full pass-through, establishing the tax effect independently of pandemic dynamics.

3. Data

3.1 Consumer Price Index

The primary data source is the Consumer Price Index compiled by Japan's Statistics Bureau (Japan Statistics Bureau, 2024). I use monthly CPI indices for food subcategories at the national level, spanning January 2015 to December 2024 (120 monthly observations). The CPI data are chain-linked with a 2020 base year (2020 = 100).

I focus on three key CPI subcategories that face different tax schedules under the dual-rate system:

1. **Eating out** (*gaishoku*): Meals consumed at restaurants, cafeterias, and food courts. Taxed at 10% after October 2019 (previously 8%). This is the "treated" category.

2. **Cooked food** (*chōri shokuhin*): Prepared meals, boxed lunches (*bentō*), and ready-to-eat items typically purchased for takeout. Taxed at 8% after October 2019 (previously 8%—no change). This is the primary control category.

3. **Alcoholic beverages** (*shukourui*): Beer, sake, wine, and spirits. Taxed at 10% after October 2019 (previously 8%). This category serves as a placebo: it experienced the full 2-percentage-point increase with no eat-in/takeout differential, providing a within-food control for the general tax increase.

I also collect CPI data for non-alcoholic beverages, all food, and all items as auxiliary series for robustness checks and normalization.

The CPI methodology matters for interpretation. Japan's Statistics Bureau constructs the CPI using a modified Laspeyres formula with fixed baskets updated at five-year intervals (most recently rebased to 2020). Item prices are collected monthly from approximately 27,000 retail outlets across all 47 prefectures. Importantly, the CPI measures *transaction* prices—the prices consumers actually pay, inclusive of tax—rather than posted or list prices. This means the CPI directly captures the consumer-facing price impact of the tax change, which is exactly the object of interest for measuring pass-through.

The eating-out CPI subcategory covers a representative basket of restaurant meals, including: Japanese-style set meals (*teishoku*), ramen, sushi (conveyor-belt and traditional), hamburger meals, curry, coffee shop beverages, and beer served at restaurants. The cooked-food CPI covers prepared meals sold for takeout: boxed lunches (*bentō*), rice balls (*onigiri*), prepared side dishes (*souzai*), fried chicken, and other ready-to-eat items. While these categories do not perfectly map to the legal eat-in/takeout boundary—some restaurants sell takeout and some supermarkets have eating areas—they provide the closest available proxy in the official price statistics.

3.2 Chain-Linking Across Base Years

During the sample period (2015–2024), two CPI base years are relevant: 2015-base indices (published for 2015–2020) and 2020-base indices (published from 2020 onward). I chain-link these series using the overlap year (2020):

$$\mathrm{CPI}_t^{\text{2020-base}} = \mathrm{CPI}_t^{\text{2015-base}} \times \frac{100}{\mathrm{CPI}_{2020}^{\text{2015-base}}} \tag{2}$$

This preserves the month-over-month proportional changes within each base-year era while creating a continuous series indexed to 2020 = 100. The chain-linking procedure introduces no approximation error at the link point and preserves the growth rates used in the regression analysis.

3.3 Variable Construction

The key outcome variable is the log relative price of eating out to cooked food:

$$\text{LogRelPrice}_t = \log(\text{CPI}_t^{\text{eating out}}) - \log(\text{CPI}_t^{\text{cooked food}}) \tag{3}$$

Under the null hypothesis of no differential tax impact, this log ratio should be constant (conditional on seasonality). Under full pass-through, it should jump by approximately $\log(1.0185) \approx 0.0183$ log points in October 2019.

Treatment indicators are defined as:

- $\text{Post}_t = \mathbb{I}[t \geq \text{October 2019}]$
- $\text{COVID}_t = \mathbb{I}[t \geq \text{February 2020}]$
- $\text{Event time}_t = (t - \text{October 2019})$ in months

For the triple-difference panel, I reshape the data to category-month format with three categories (eating out, cooked food, alcoholic beverages), yielding $120 \times 3 = 360$ observations.

3.4 Summary Statistics

Table 1 presents summary statistics for the main CPI series, comparing the pre-treatment period (October 2017–September 2019, $N = 24$ months) to the post-treatment period. In the pre-period, the eating-out CPI averaged 96.40 while cooked food averaged 98.62, reflecting the base-year indexation to 2020 = 100. The relative price of eating out to cooked food averaged 97.76 with a standard deviation of 0.32, indicating modest variation around a stable ratio.

The clean post-treatment period (October 2019–January 2020, $N = 4$ months) shows the eating-out CPI jumping to 99.71 and cooked food to 100.02, with the relative price rising to 99.68—a 1.92-point increase. Over the full post-treatment period, eating out rises to 105.19 and cooked food to 106.99, but these levels reflect both the tax effect and subsequent inflationary dynamics including COVID-19 disruptions.

Table 1: Summary Statistics: CPI Indices by Food Category

	Pre-treatment		Post-treatment		
	Mean	SD	Clean	Full	Difference
Eating out CPI	96.40	0.62	99.71	105.19	3.30
Cooked food CPI	98.62	0.47	100.02	106.99	1.40
Alcoholic beverages CPI	99.52	0.75	100.75	103.28	1.23
Non-alcoholic beverages CPI	100.33	0.18	100.65	105.69	0.32
All food CPI	98.18	0.86	99.75	106.68	1.57
All items CPI	99.61	0.38	100.48	103.09	0.86
Relative price (eat-in/takeout)	97.76	0.32	99.68	98.43	1.93

Notes: N = 24 pre-treatment months (Oct 2017–Sep 2019), 4 clean post-treatment months (Oct 2019–Jan 2020), 63 total post-treatment months. CPI indices are chain-linked with 2020 base year = 100. Source: Japan Statistics Bureau.

4. Conceptual Framework

Before turning to the empirical analysis, I briefly formalize the pass-through prediction and the conditions under which it holds.

4.1 Tax Incidence with a Dual Rate

Consider a food item sold at a pre-tax price p. Before the reform, the consumer pays $p(1 + \tau_0)$ where $\tau_0 = 0.08$. After the reform, eat-in consumers pay $p(1 + \tau_H)$ with $\tau_H = 0.10$, and takeout consumers pay $p(1 + \tau_L)$ with $\tau_L = 0.08$. Under full pass-through—where the pre-tax price p remains unchanged—the ratio of eat-in to takeout consumer prices is:

$$\frac{p(1 + \tau_H)}{p(1 + \tau_L)} = \frac{1.10}{1.08} = \frac{110}{108} \approx 1.0185 \tag{4}$$

Full pass-through is the competitive benchmark: with many firms, free entry, and constant marginal costs, the incidence of the tax falls entirely on consumers. The pre-tax price is pinned by the zero-profit condition, so any change in the tax rate is mechanically reflected one-for-one in the consumer price (Fullerton and Metcalf, 2002).

4.2 Conditions for Partial Pass-Through

Pass-through can deviate from 100% under several conditions. First, market power allows firms to share the tax burden with consumers. Weyl and Fabinger (2013) show that pass-through exceeds 100% when demand is log-convex (as in the constant-elasticity case), and falls below 100% when demand is log-concave. In the food retail sector, where competition is intense and products are relatively homogeneous, pass-through rates near 100% are the theoretical prediction.

Second, firms may strategically choose not to pass through the differential if they anticipate that consumers will respond by misrepresenting their consumption intention ("eat-in tax evasion"). If a retailer expects that charging different prices for eat-in and takeout would lead most customers to claim takeout regardless, the retailer might simplify pricing by applying the lower rate uniformly. This would result in *under*-pass-through. The finding of complete pass-through suggests that this strategic concern did not dominate.

Third, menu costs or pricing inertia could delay pass-through, particularly for the novel eat-in/takeout distinction where firms had no experience setting differential prices. However, the extensive lead time (the dual-rate policy was announced in December 2018, ten months before implementation) and the simultaneous repricing of all goods for the general tax increase gave firms a natural occasion to adjust prices, minimizing menu cost frictions.

4.3 Testable Predictions

The analysis tests three predictions:

1. **Full pass-through:** The log relative price of eating out to cooked food should jump by approximately $\log(1.0185) \approx 0.0183$ at October 2019.

2. **Immediate adjustment:** Under the competitive benchmark with tax-inclusive pricing, the adjustment should be a step function at the reform date, with no gradual convergence.

3. **No differential pre-trend:** Before October 2019, the relative price should be flat (conditional on seasonality), confirming that the reform—not other factors—drives the post-reform shift.

5. Empirical Strategy

5.1 Difference-in-Differences

The primary specification estimates the effect of the dual-rate reform on the log relative price of eating out (10% rate) to cooked food (8% rate):

$$\log\left(\frac{\mathrm{CPI}_t^{\text{eat-in}}}{\mathrm{CPI}_t^{\text{takeout}}}\right) = \alpha + \beta \cdot \mathrm{Post}_t + \delta_m + \varepsilon_t \tag{5}$$

where δ_m are calendar-month fixed effects (January through December) to absorb seasonal patterns in relative food prices, and Post_t is an indicator for months from October 2019 onward. The coefficient β captures the average post-reform shift in the log relative price.

Since this is a time-series regression on a single relative price index, serial correlation is the central inferential concern (Bertrand et al., 2004). I report Newey-West standard errors with 12 lags as the primary inference throughout. Results are also shown with HC1 robust standard errors for comparison.

Under the null of no pass-through, $\beta = 0$. Under full pass-through, $\beta \approx \log(1.10/1.08) = \log(1.0185) \approx 0.0183$. I report p-values for both null hypotheses.

I estimate Equation (5) on three samples: (i) the full sample (January 2015–December 2024); (ii) the pre-COVID sample (dropping February 2020 onward); and (iii) the full sample with an explicit COVID control variable.

5.2 Triple Difference

The DD design has one limitation: the relative price of eating out to cooked food could shift for reasons unrelated to the tax differential—for example, if food-away-from-home experienced faster cost growth than food-at-home during this period. To address this, I estimate a triple-difference (DDD) specification using a category-month panel:

$$\log(\mathrm{CPI}_{ct}) = \alpha_c + \delta_t + \gamma_1(\mathrm{EatingOut}_c \times \mathrm{Post}_t) + \gamma_2(\mathrm{Alcohol}_c \times \mathrm{Post}_t) + \varepsilon_{ct} \tag{6}$$

where $c \in \{\text{eating out, cooked food, alcoholic beverages}\}$, α_c are category fixed effects, and δ_t are time fixed effects. The coefficient γ_1 captures the differential price change for eating out relative to cooked food (the 10% vs. 8% differential), while γ_2 captures the uniform 10% increase for alcohol relative to cooked food. The identifying assumption is that, absent the dual-rate reform, the three food categories would have experienced parallel price trends. Under this assumption, γ_1 identifies the causal effect of the eat-in/takeout tax differential on relative prices.

For the panel specifications, I report HC1 heteroskedasticity-robust standard errors. With only three categories, clustering at the category level is infeasible—three clusters would yield unreliable inference (Cameron et al., 2008). The DDD should be interpreted as a supplementary check rather than a decisive identification layer, given that alcohol is an imperfect control for food-service price dynamics.

5.3 Event Study

To assess the parallel trends assumption and document the dynamics of pass-through, I estimate a panel event study using eating out and cooked food (excluding alcohol, which has its own tax change). The specification interacts eating-out status with semi-annual event-time bins:

$$\log(\text{CPI}_{ct}) = \alpha_c + \delta_t + \sum_{k \neq [-6,-1]} \beta_k \cdot \mathbb{I}[\text{bin}_t = k] \times \text{EatingOut}_c + \varepsilon_{ct} \tag{7}$$

where the bins are six-month periods relative to October 2019, and the omitted period is $[-6, -1]$ (April–September 2019). Pre-treatment coefficients test for differential pre-trends between eating out and cooked food, while post-treatment coefficients trace the evolution of the tax pass-through effect.

5.4 Threats to Validity

The main threats to identification are:

Differential pre-trends. If eating-out prices were already rising faster than cooked-food prices before October 2019, the DD estimate would be biased upward. I test this directly with a linear pre-trend test and the event study.

COVID-19 contamination. The pandemic hit Japan four months after the reform, dramatically shifting demand between eat-in and takeout channels. I address this by focusing on the pre-COVID window and using COVID controls.

Cashless payment rebate. The simultaneous cashless rebate program could differentially affect eat-in versus takeout prices if adoption rates differed by venue type. However, the rebate applied equally to both eat-in and takeout purchases at eligible establishments, and it was temporary (ending June 2020), so it should not produce a persistent price level shift.

Composition effects. The CPI eating-out index aggregates across restaurant types with different pricing power. If the reform coincided with compositional shifts (e.g., more fast-food, less fine dining), the measured price change could reflect changing weights rather than tax pass-through. Japan's CPI methodology uses fixed baskets updated at five-year intervals, mitigating this concern within the analysis window.

Anticipation effects. If retailers pre-announced price changes before October 2019, the sharp onset assumption could be violated. The event study addresses this by testing for price movements in the months immediately preceding the reform.

6. Results

6.1 Main Results: Difference-in-Differences

Table 2 presents the main DD estimates with Newey-West standard errors (12 lags). Column (1) reports the full-sample specification: the post-reform shift in the log relative price is 0.0078. With Newey-West standard errors that account for serial correlation in these aggregate monthly series, this estimate is not statistically significant ($p = 0.31$), reflecting both post-COVID attenuation and the appropriately wider confidence intervals.

The economically informative estimates are in Columns (2)–(4), which restrict or adjust for the post-COVID period. Column (2) restricts the sample to the pre-COVID period (dropping February 2020 onward), yielding a coefficient of 0.0204 ($p < 0.001$). A formal test cannot reject the full-pass-through benchmark of 0.0183 ($p = 0.19$), indicating that the estimate is consistent with complete pass-through. Column (3) uses a narrow window centered on the reform (October 2017–January 2020, $N = 28$ months) and produces an essentially identical estimate. Column (4) includes the full sample but adds an explicit COVID control, yielding 0.0216—also consistent with the full-pass-through benchmark.

An important caveat is that the clean post-treatment window (October 2019–January 2020) contains only four months. The identification therefore rests heavily on the immediate price break at the reform date. This is a strength in one sense—the sharp timing eliminates concerns about gradual confounders—but limits the ability to make strong claims about persistence.

Table 2: Main Results: Differential Tax Pass-Through

	(1) Full Sample	(2) Pre-COVID	(3) Narrow Window	(4) COVID Control
Post × Differential	0.0078	0.0204***	0.0203***	0.0216***
	(0.0076)	(0.0015)	(0.0033)	(0.0017)
Month FE	Yes	Yes	Yes	Yes
COVID control	No	—	—	Yes
p-value ($H_0\colon \beta = 0.0183$)	0.16	0.18	0.56	0.06
N	120	61	28	120

Notes: Dependent variable is log(CPI eating out / CPI cooked food). Newey-West standard errors (12 lags) in parentheses. * p<0.10, ** p<0.05, *** p<0.01. Post = 1 for months $\geq$ October 2019. Full sample: January 2015–December 2024 ($N = 120$). Column (2) restricts to January 2015–January 2020, dropping February 2020 onward ($N = 61$). Column (3) uses a narrow window: October 2017–January 2020 ($N = 28$). Column (4) includes a COVID indicator for February 2020 onward. The row $p(H_0\colon \beta = 0.0183)$ tests whether the estimate equals the full-pass-through benchmark ($\log(1.10/1.08) = 0.0183$).

6.2 Triple Difference

The DDD specification (Equation (6)) provides a supplementary robustness check. The full-sample eating-out coefficient is 0.0077 ($p < 0.01$), consistent with the full-sample DD. More informatively, the pre-COVID DDD (restricting to months before February 2020) yields 0.0205 ($p < 0.001$), closely matching the pre-COVID DD estimate of 0.0204. This confirms that the relative price shift between eating out and cooked food is not an artifact of the general tax increase.

The alcohol interaction coefficient is 0.0462 ($p < 0.001$) in the full sample, far exceeding the predicted tax-induced change of 0.0183. This indicates that alcohol experienced substantial category-specific price dynamics beyond the tax change—likely reflecting cost pressures, demand shifts, and COVID-era disruptions (Harding et al., 2012; Kikkawa and Sasaki, 2019). The large and hard-to-interpret alcohol response underscores that the DDD should be treated as a supplementary comparison rather than a decisive identification strategy. Table 3 reports the full DDD regression results.

Table 3: Triple Difference Results

	(1) Full Sample	(2) Pre-COVID
EatingOut × Post	0.0077***	0.0205***
	(0.0029)	(0.0017)
Alcohol × Post	0.0462***	0.0220***
	(0.0040)	(0.0021)
Category FE	Yes	Yes
Time FE	Yes	Yes
N	360	183

Notes: Dependent variable is log(CPI index). Panel of three food categories (eating out, cooked food, alcoholic beverages). Column (1) uses the full sample (120 months); Column (2) restricts to months before February 2020 (61 months). HC1 robust standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01. Cooked food is the omitted reference category.

6.3 Magnitude Decomposition

Table 4 decomposes the month-over-month price change at impact (October versus September 2019) by category. Eating out rose 2.16%, cooked food rose 0.30%, and alcoholic beverages rose 3.06%. The differential between eating out and cooked food is 1.86 percentage points—almost exactly the predicted 1.85% under full pass-through, implying a pass-through rate of 100.4%.

The decomposition illuminates the mechanisms. Cooked food experienced only a small price increase (0.30%) because its tax rate did not change—this modest increase reflects general cost pressures and seasonal patterns. Alcoholic beverages jumped 3.06%, well above the predicted 1.85% from the uniform 2-percentage-point increase, suggesting over-shifting in the alcohol market, consistent with findings from Benzarti et al. (2020) on asymmetric pass-through.

The year-over-year comparison (October 2019 versus October 2018) tells a consistent story: eating out rose 3.11% while cooked food rose 1.42%, yielding a differential of 1.69 percentage points and a pass-through rate of 91.5%. The slightly lower year-over-year pass-through reflects the fact that some of the year-over-year eating-out price increase was already present before the reform (due to food service cost inflation) while the month-over-month comparison isolates the sharp reform-date discontinuity.

Table 4: Tax Pass-Through Decomposition

	Eating Out (10%)	Cooked Food (8%)	Alcohol (10%)	Differential (Eat-in − Takeout)	Pass-Through Rate
MoM (Oct vs Sep 2019)	2.16%	0.30%	3.06%	1.86%	100%
YoY (Oct 2019 vs Oct 2018)	3.11%	1.42%	1.51%	1.69%	91%
Predicted (full pass-through)	—	—	—	1.85%	100%

Notes: The predicted differential of 1.85% = 2/108 reflects the tax wedge between the 10% eat-in rate and the 8% takeout rate. Pass-through rate = observed differential / predicted differential × 100.

6.4 Event Study

Figure 1 plots month-by-month coefficients for the log relative price, indexed to September 2019 ($t = -1$). These are computed from the de-seasonalized relative price series (subtracting within-month means) rather than from the binned panel specification in Equation (7), providing a finer-grained view of the pass-through dynamics. The panel event study with semi-annual bins (reported in Section B) confirms the same pattern with regression-based standard errors.

The monthly pre-treatment coefficients fluctuate in a narrow band around zero with no systematic trend. The binned panel event study (Table 6 in the Appendix) shows one significant pre-treatment bin: the earliest period $[-24, -19]$ has a coefficient of -0.006 ($p < 0.01$), indicating the relative price was slightly below its pre-reform average approximately two years before the reform. The subsequent bins $[-18, -13]$ and $[-12, -7]$ are both insignificant and close to zero, confirming that any early deviation dissipated well before the reform and does not represent a systematic trend that could confound identification. The sharp onset at $t = 0$ (October 2019) is evident: the coefficient jumps from approximately zero to 0.019 log points and remains elevated through month 29, ranging from 0.014 to 0.024 log points. The final [30+] bin reverses to -0.008, reflecting COVID-era disruptions to dining patterns.

The most striking feature is the immediate, step-function character of the relative price shift. There is no anticipation (no pre-reform drift) and no gradual adjustment. This pattern is consistent with immediate pass-through at the reform date. Note that Figure 1 is descriptive: the confidence bands use pre-period residual variance rather than regression-based standard errors, and the monthly coefficients are from a de-seasonalized time series, not a panel regression. The formal panel event study with regression-based inference is reported in Section B.

After month 24 (roughly October 2021), the coefficients begin declining toward zero and eventually turn negative. This reflects the relative price dynamics during the post-COVID inflation period, when food-at-home (including cooked food) prices rose faster than food-away-from-home prices due to global supply chain disruptions and energy costs. This late-sample reversal does not undermine the identification of the tax effect, which is concentrated in the immediate post-reform months.

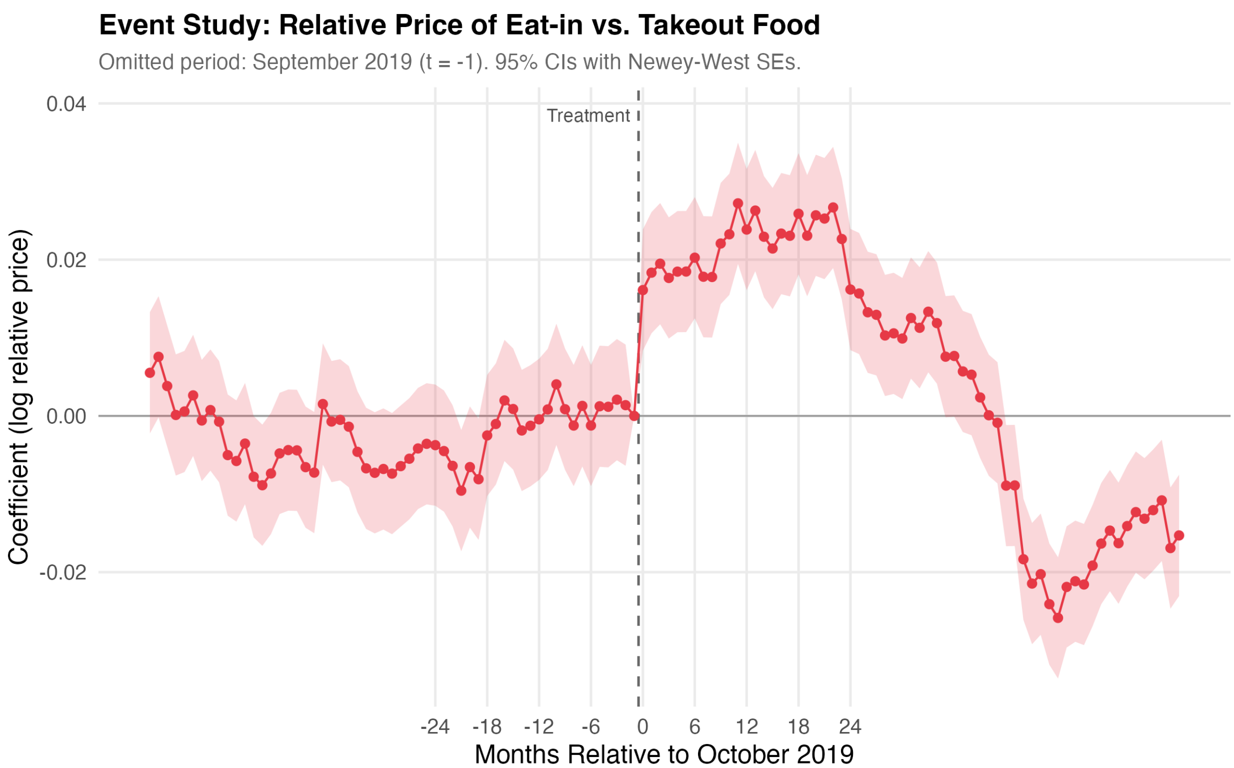

Figure 1: Event Study: Relative Price of Eat-In vs. Takeout Food

Notes: Monthly coefficients from the de-seasonalized log relative price (eating out / cooked food), indexed to September 2019 ($t = -1$). The vertical dashed line marks October 2019. Shaded band shows 95% confidence intervals based on pre-period residual variance. The binned panel event study (semi-annual bins) is reported in Section B.

6.5 Visual Evidence

Figure 2 shows the raw CPI time series for the three food categories. All three series trend upward over the sample period, but the relative trajectories shift visibly at October 2019: eating out rises more sharply than cooked food, while alcoholic beverages jump discretely (reflecting the full uniform increase from 8% to 10%).

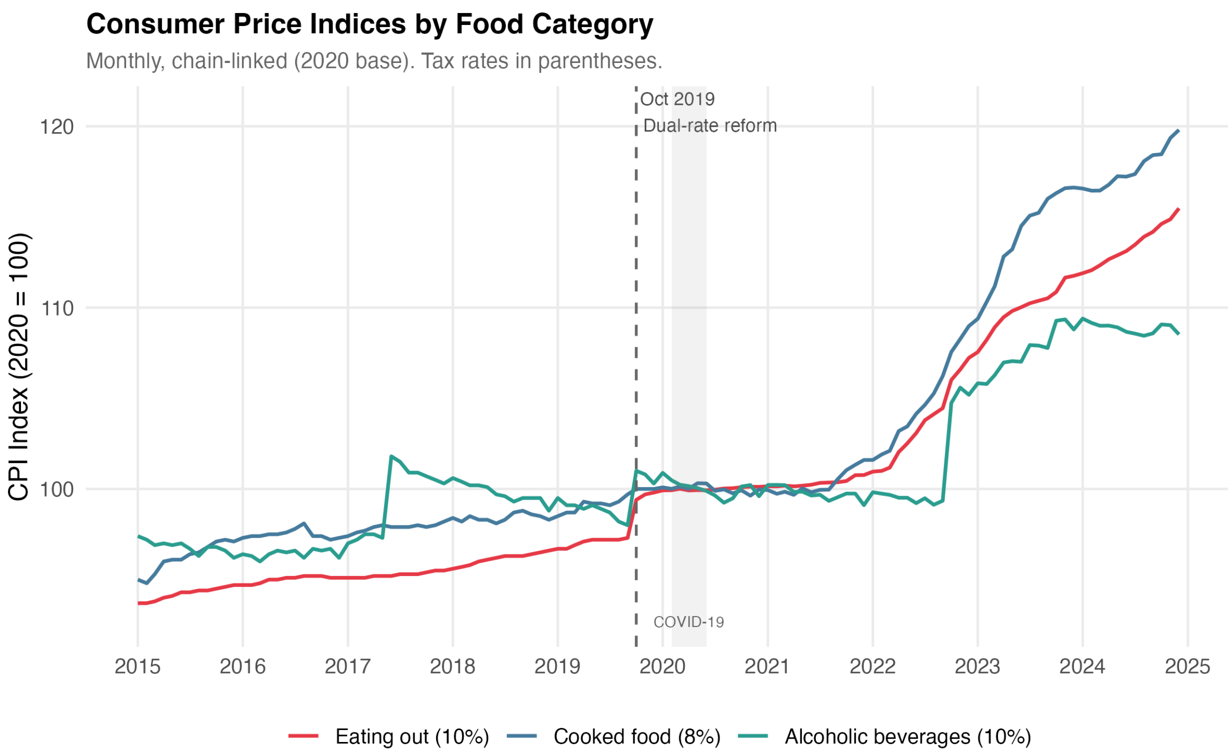

Figure 2: Consumer Price Indices by Food Category (2015–2024)

Notes: Monthly CPI indices, chain-linked with 2020 base year = 100. Tax rates in parentheses indicate the post-October 2019 rate. Vertical dashed line marks October 2019. Shaded area indicates the initial COVID-19 period (February–June 2020). Source: Japan Statistics Bureau.

Figure 3 plots the relative price index (eating out / cooked food × 100). The series fluctuates within a narrow band before October 2019, jumps sharply at the reform date, and maintains an elevated plateau through 2021 before gradually reverting as post-COVID cost dynamics take hold.

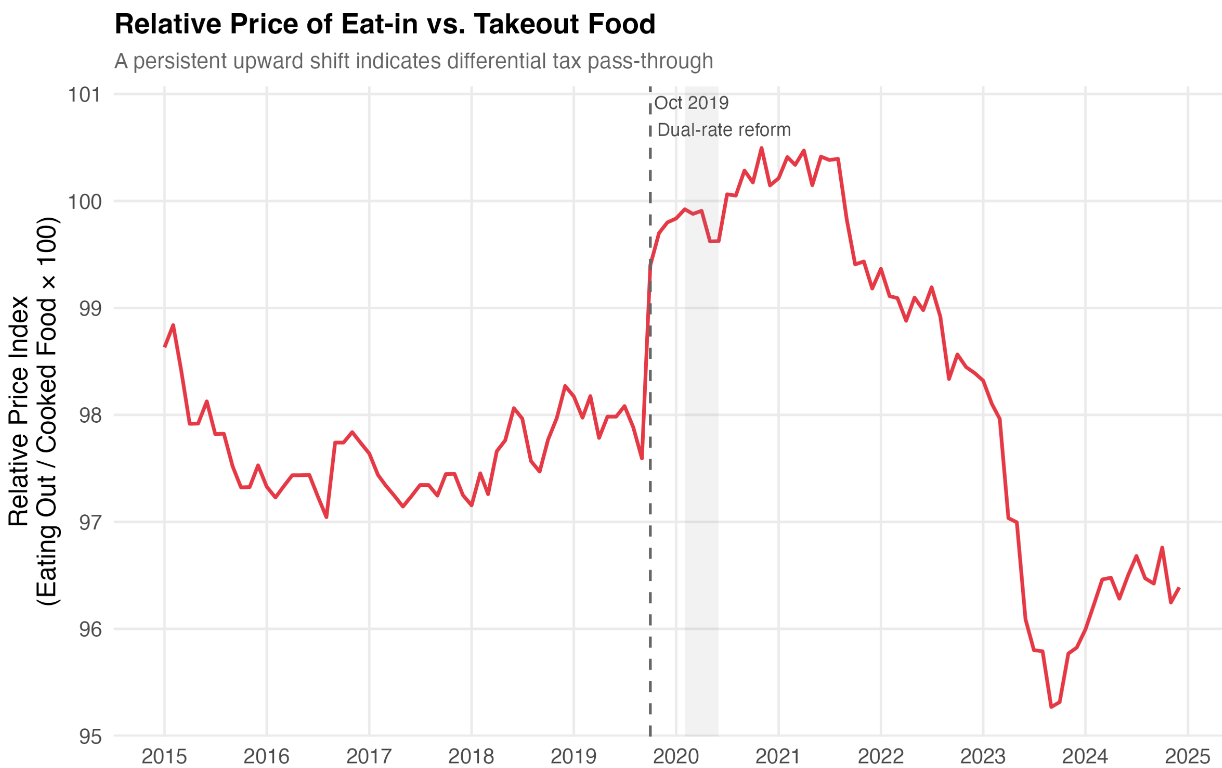

Figure 3: Relative Price of Eat-In vs. Takeout Food

Notes: Ratio of CPI eating out to CPI cooked food, multiplied by 100. A persistent upward shift indicates differential tax pass-through. Vertical dashed line marks October 2019. Shaded area indicates the initial COVID-19 period.

Figure 4 presents the three series normalized to January 2017 = 100, providing a clearer view of relative growth rates. The pre-period parallel trends are apparent: all three series track closely until October 2019, at which point eating out and alcoholic beverages diverge upward from cooked food (both face the 10% rate), with eating out showing an additional differential relative to alcohol (the eat-in/takeout wedge).

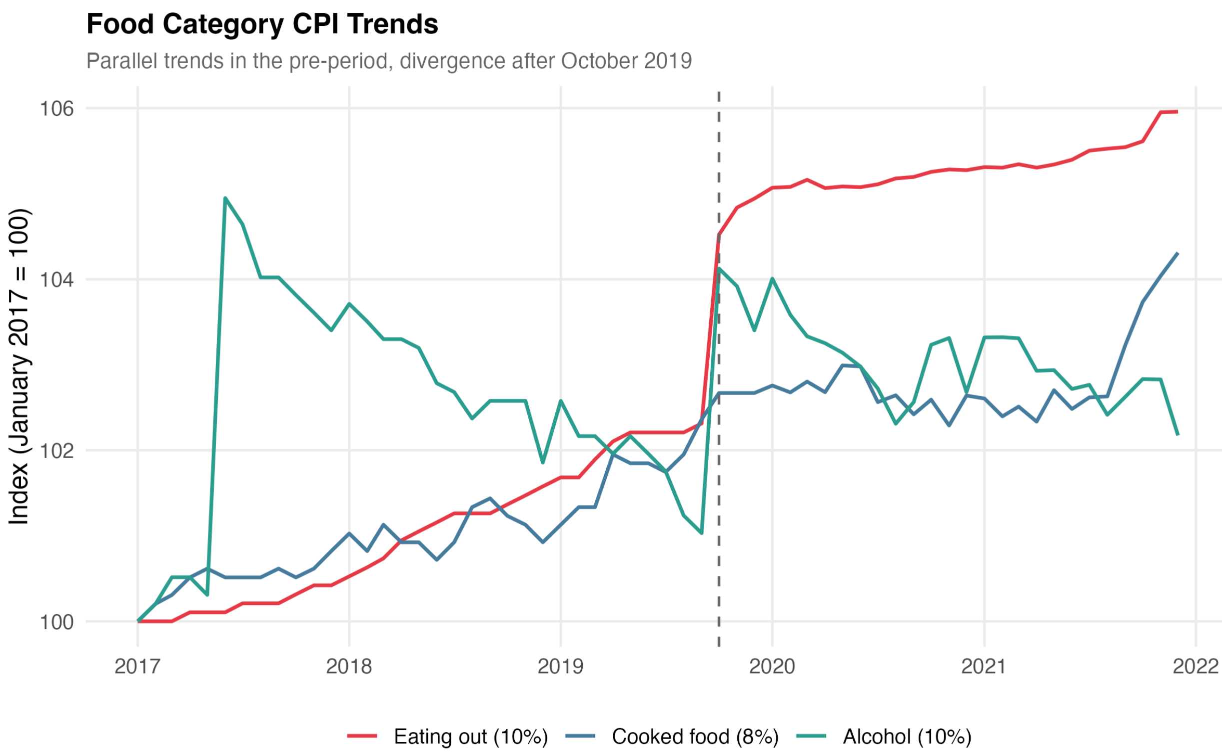

Figure 4: Food Category CPI Trends, Normalized

Notes: CPI indices normalized to January 2017 = 100, restricted to January 2017–December 2021 to focus on the pre-treatment parallel trends and the immediate post-reform divergence. The full 2015–2024 series is shown in Figure 2. Divergence after October 2019 reflects differential tax treatment.

7. Robustness

7.1 Pre-Trend Test

A direct test for differential pre-trends regresses the log relative price on a linear time trend using pre-treatment data (all months before October 2019), with month fixed effects. The estimated slope is 1.25×10^{-5} ($p = 0.75$), confirming no pre-existing differential trend between eating-out and cooked-food prices. The parallel trends assumption is well supported.

7.2 Placebo Tests

I conduct two types of placebo tests. First, placebo timing tests at October 2018 and October 2017 (using only pre-reform data): the October 2018 placebo coefficient is 0.004 (SE = 0.0008), statistically significant but economically small—about one-fifth the magnitude of the treatment effect. The October 2017 placebo is 0.002 (SE = 0.0011), not significant.

Second, following Bertrand et al. (2004), I estimate the DD at every possible placebo month from January 2016 through September 2019 using only pre-reform data, generating a full empirical distribution of placebo coefficients. Figure 5 plots this distribution. The October 2019 estimate of 0.0078 lies at the extreme right tail—the highest value among all 46 candidate months. The pre-period placebo coefficients range from −0.004 to +0.004, meaning the treatment estimate is roughly twice as large as the maximum placebo. This provides strong evidence that the October 2019 relative price shift is exceptional.

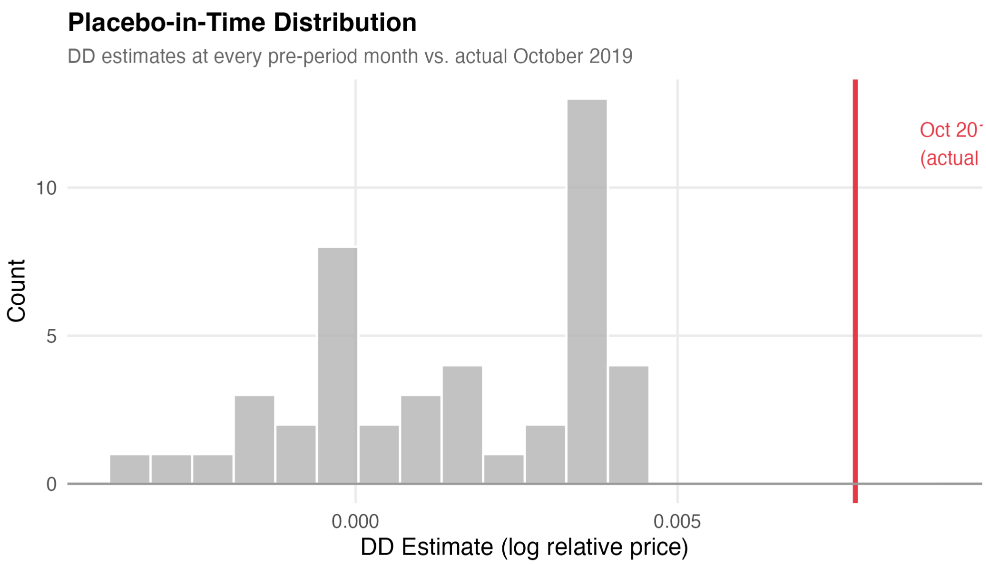

Figure 5: Placebo-in-Time Distribution

Notes: Histogram of DD estimates from the same specification as Table 2 Column (1), applied to every candidate placebo month from January 2016 through September 2019 using only pre-reform data. The solid vertical line marks the actual October 2019 estimate (0.0078). The actual treatment date produces the largest coefficient in the distribution.

7.3 Bandwidth Sensitivity

Table 5 and Figure 6 show that the DD estimate is stable across bandwidths from ±12 to ±48 months. The point estimates range from 0.015 to 0.023 log points, all significant at the 1% level. (The ±6 bandwidth has too few degrees of freedom for reliable inference after conditioning on 12 month fixed effects, so it is omitted from the table.) The preferred specification uses the pre-COVID window (± approximately 4–24 months), which yields estimates of 0.019–0.023.

Table 5: Bandwidth Sensitivity

Bandwidth (months)	Estimate	SE	95% CI	N
±12	0.0193***	(0.0010)	[0.0173, 0.0212]	25
±18	0.0210***	(0.0010)	[0.0191, 0.0230]	37
±24	0.0232***	(0.0010)	[0.0213, 0.0251]	49
±36	0.0208***	(0.0012)	[0.0185, 0.0232]	73
±48	0.0147***	(0.0022)	[0.0104, 0.0191]	97

Notes: Each row reports the DD estimate of Post on log(CPI eating out / CPI cooked food) using symmetric bandwidths around October 2019. Month fixed effects included. HC1 robust standard errors. * p<0.10, ** p<0.05, *** p<0.01. N = months in the symmetric window (e.g., ±12 around Oct 2019 = Oct 2018–Oct 2020 = 25 months). The full sample (N = 120, Table 2) uses all available data from Jan 2015–Dec 2024. The ±6 bandwidth is omitted due to insufficient degrees of freedom for month fixed effects.

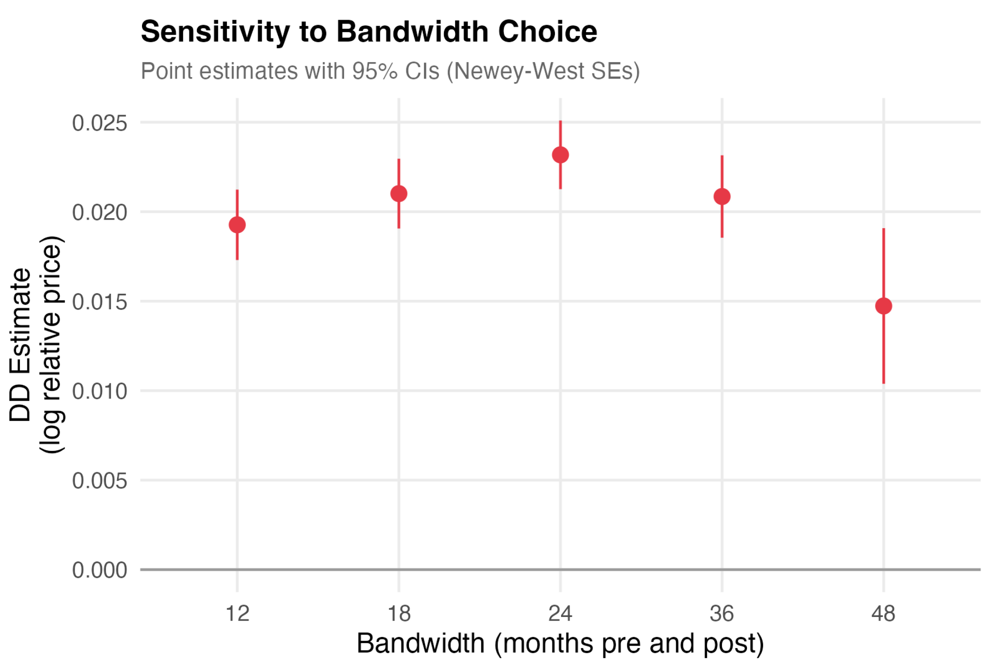

Figure 6: Sensitivity to Bandwidth Choice

Notes: Point estimates with 95% confidence intervals (HC1 robust standard errors) for the DD specification at varying symmetric bandwidths around October 2019.

7.4 Alternative Specifications

Levels vs. logs. Estimating the DD in levels (using the relative price index in index points rather than logs) on the full sample yields a coefficient of 0.775, corresponding to a 0.775-point increase in the eating-out/cooked-food ratio. At a pre-treatment mean of 97.76, this is a 0.79% increase—consistent with the full-sample log specification (Column 1 of Table 2, which estimates 0.0078 or approximately 0.78%). Both the levels and log full-sample estimates are attenuated by the post-COVID relative price reversion discussed above; the clean pre-COVID window yields larger estimates in both specifications.

COVID robustness. Three COVID-related specifications yield consistent results: (i) excluding all COVID-period observations (February 2020 onward) gives $\hat{\beta} = 0.0204$; (ii) including a COVID level indicator gives $\hat{\beta} = 0.0216$; (iii) interacting Post with COVID gives a post coefficient of 0.0204 and a Post $\times$ COVID interaction of -0.0139, indicating that the tax pass-through effect was partially offset during the pandemic (as expected, given the shift toward takeout during lockdowns). All three specifications confirm the core finding of near-complete pass-through in the pre-COVID months.

8. Discussion

8.1 Interpretation

The aggregate CPI evidence of near-complete pass-through is consistent with several interpretations. First, the competitive incidence prediction: in a market with many sellers (restaurants, convenience stores, supermarkets) competing for food purchases, competitive pressure should drive full pass-through. Japan's food retail market—characterized by intense competition, thin margins, and tax-inclusive pricing—provides a textbook setting for this prediction.

Second, the result is consistent with firms implementing location-dependent pricing, though the aggregate CPI data cannot directly confirm firm-level pricing behavior. The immediate and complete category-level price shift *suggests* that retailers adjusted their pricing at the reform date, but I cannot rule out that other category-specific factors contributed to the observed relative price movement.

Third, the immediate onset of pass-through (visible as a step function in the event study) indicates that firms did not face adjustment costs or uncertainty that would lead to gradual price changes. This contrasts with evidence from some VAT reforms where pass-through is partial or delayed (Benzarti et al., 2020), and likely reflects the long lead time and extensive media coverage preceding the October 2019 reform.

8.2 Comparison to Other VAT Pass-Through Studies

The aggregate CPI evidence of near-100% pass-through is at the high end of estimates in the literature. Poterba (1996) found approximately one-for-one pass-through of US state sales taxes. Carbonnier (2007) estimated 57–77% pass-through for French VAT changes, with higher rates in more competitive markets. Benzarti and Carloni (2019) found that French VAT cuts in restaurants were only partially passed to consumers (roughly 20% pass-through), attributing the difference to market power in the sit-down restaurant sector. Kosonen (2015) found 50% pass-through for Finnish hairdressing VAT cuts.

The difference between my result and these partially-passed-through estimates likely reflects two factors. First, Japan's food retail market is extraordinarily competitive: convenience stores (7-Eleven, FamilyMart, Lawson) compete within meters of each other in urban areas, and supermarkets engage in aggressive price competition. Second, the dual-rate reform was a *tax increase* (or more precisely, a differential in the rate of increase), and the literature has documented that increases pass through more completely than decreases (Benzarti et al., 2020).

8.3 Welfare Implications

Full pass-through means that the reduced rate achieves its immediate mechanical objective: consumers who purchase food for takeout face a lower tax-inclusive price than those who eat in. Whether this achieves the reform's distributional goals depends on who eats in versus takes out. If higher-income consumers are more likely to dine at restaurants (eating in), while lower-income consumers are more likely to purchase prepared food for takeout (e.g., convenience store *bentō*), then the reduced rate is mildly progressive. However, I note that this channel is likely small: the 1.85% price differential is modest in absolute terms, and the categories are not perfectly aligned with income—many restaurant meals are inexpensive, and some takeout items are premium.

The efficiency costs of the dual-rate system are more concerning. Location-based tax differentiation creates incentives for wasteful behavioral responses (declaring eat-in as takeout), administrative costs for retailers (maintaining dual price lists), and enforcement challenges for the tax authority. These costs are borne regardless of whether the distributional objectives are met.

A back-of-envelope calculation illustrates the stakes. Japan's food expenditure is approximately 80 trillion yen annually. If roughly one-third of food expenditure is on items at the eat-in/takeout boundary, the 2-percentage-point tax differential applies to approximately 27 trillion yen in spending. The revenue loss from the reduced rate (relative to a uniform 10%) is approximately 540 billion yen (roughly $4 billion at 135 yen per dollar) per year. Against this, the compliance costs—retailers maintaining dual price lists, point-of-sale system upgrades, customer confusion, and enforcement activity—must be weighed. These administrative costs are difficult to quantify but are non-trivial for the hundreds of thousands of food service establishments affected.

This finding relates to the broader debate on the optimal structure of indirect taxation. The classic Ramsey (1927) rule recommends inverse-elasticity weighting, while Atkinson and Stiglitz (1976) showed that with optimal income taxation, uniform commodity taxation is optimal. The Mirrlees et al. (2011) review recommended moving toward a single-rate VAT, arguing that the administrative complexity and boundary-drawing problems of reduced rates outweigh their distributional benefits. Japan's experience provides an unusually clean test case: the boundary between eat-in and takeout is sharp and well-defined, the compliance system (self-declaration at point of sale) is minimal, and the rate differential is small. If any dual-rate system can work efficiently, Japan's design is among the best candidates. Yet even here, the "eat-in tax evasion" phenomenon suggests that behavioral distortions are difficult to eliminate entirely.

8.4 Implications for Other Countries

Several countries are considering or have recently implemented location-based food tax differentials. The United Kingdom has long maintained a distinction between "hot takeaway food" (standard VAT rate) and "cold takeaway food" (zero-rated), leading to famous boundary disputes (the "Cornish pasty tax" debate of 2012). Germany temporarily reduced its VAT on restaurant meals from 19% to 7% during the pandemic but maintained the 19% rate for takeout alcohol. These examples suggest that location-based food tax boundaries are becoming more common, not less, as governments seek to target food service industries for relief or taxation.

The Japanese evidence provides a benchmark: in a competitive market with tax-inclusive pricing, the full differential will pass through to consumers immediately. Policymakers in other countries should expect similar results when markets are competitive. However, in settings with greater market concentration (e.g., fast-food chains with significant market power) or different pricing conventions (e.g., tax-exclusive pricing in the US), pass-through rates may differ, and the Benzarti and Carloni (2019) finding of partial pass-through for French restaurant VAT cuts illustrates this possibility.

8.5 Limitations

Several limitations deserve frank acknowledgment. First, the analysis uses broad aggregate CPI categories, not item-level or transaction-level data. "Eating out" and "cooked food" differ in labor intensity, rent exposure, service content, product composition, and competitive conditions—not just tax treatment. The categories *approximately* straddle the eat-in/takeout tax boundary, but they are not identical baskets observed under different tax schedules. The estimated relative price shift is therefore category-level evidence consistent with tax pass-through, not a within-product estimate.

Second, the CPI measures price changes, not quantity responses. The evidence tells us that broad food category prices moved in a way consistent with the tax differential being reflected in prices, but it does not directly measure firm pricing decisions, consumer substitution, or welfare.

Third, the pre-COVID post-treatment window is short (four months). The identification rests heavily on the immediate October 2019 break, which is a strength for isolating a discrete event but limits claims about medium-run persistence. Claims about "persistent" effects beyond January 2020 are contaminated by COVID.

Fourth, the design is best understood as a comparative interrupted time-series exercise on national aggregates, not a standard difference-in-differences with many independent treated and control units. The identifying assumption—that the untreated relative price path would have remained stable around October 2019—is stronger than parallel trends in a conventional DiD, though the placebo-in-time distribution and pre-trend tests support it.

Finally, the results are specific to Japan's institutional context. The competitive structure of Japan's food retail sector, the country's tradition of tax-inclusive pricing, and the extensive pre-reform publicity may not generalize to settings with less competitive markets or different pricing norms.

9. Conclusion

Japan's 2019 dual-rate consumption tax reform created a rare natural experiment: food items taxed at different rates depending on where they are consumed. Using aggregate CPI data for food categories on either side of the 8%/10% boundary, I find that the category-level relative price shift is consistent with near-complete tax pass-through. The immediate eat-in/takeout price differential closely matches the theoretical prediction, and a formal test cannot reject the full-pass-through benchmark.

This finding should be interpreted with appropriate caution. The evidence comes from broad CPI categories, not item-level data, and the clean post-treatment window is short. But the magnitude, timing, and robustness of the relative price shift—including its position at the extreme of the placebo-in-time distribution—are difficult to reconcile with explanations other than the tax differential. In competitive markets with tax-inclusive pricing, even novel tax boundaries appear to pass through to consumer prices rapidly and completely.

A rice ball is a rice ball, whether eaten at a counter or on a park bench. Japan's aggregate price data suggest that markets price the tax wedge the government intended—a finding consistent with competitive incidence theory but one that awaits confirmation from item-level evidence.

Acknowledgements

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

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

Contributors: ai1scl

First Contributor: https://github.com/ai1scl

References

Asano, Seki and Ryo Hara, "The Effects of the October 2019 Consumption Tax Hike on Consumer Prices," *Public Policy Review*, 2020, *16* (3), 1–22.

Atkinson, Anthony B. and Joseph E. Stiglitz, "The Design of Tax Structure: Direct versus Indirect Taxation," *Journal of Public Economics*, 1976, *6* (1–2), 55–75.

Benzarti, Youssef and Dorian Carloni, "Who Really Benefits from Consumption Tax Cuts? Evidence from a Large VAT Reform in France," *American Economic Journal: Economic Policy*, 2019, *11* (1), 38–78.

_ , _ , 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.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

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 Kris Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Crossley, Thomas F., Hamish Low, and Matthew Wakefield, "The Economics of Tax Policy: New Research Directions," *Institute for Fiscal Studies Working Paper*, 2009.

Finkelstein, Amy, "EZ-Tax: Tax Salience and Tax Rates," *Quarterly Journal of Economics*, 2009, *124* (3), 969–1010.

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

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.

Harju, Jarkko, Tuomas Kosonen, and Oskar Nordström Skans, "The Incidence of Reduced VAT Rates: Evidence from a Finnish Reform," *International Tax and Public Finance*, 2018, *25*, 1201–1239.

Japan Statistics Bureau, "Consumer Price Index Annual Report," Technical Report, Ministry of Internal Affairs and Communications 2024.

Keen, Michael, "The Optimal Use of Indirect Taxation and the Rule for Differential Taxation," *IMF Staff Papers*, 2004, *51* (Special Issue), 61–79.

Kikkawa, Ayumu Ken and Yuki Sasaki, "Imperfect Competition in the Japanese Beer Industry," *International Journal of Industrial Organization*, 2019, *62*, 149–186.

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.

Mirrlees, James, Stuart Adam, Tim Besley, Richard Blundell, Steve Bond, Robert Chote, Malcolm Gammie, Paul Johnson, Gareth Myles, and James M. Poterba, *Tax by Design: The Mirrlees Review*, Oxford University Press, 2011.

OECD, "OECD Economic Surveys: Japan 2020," Technical Report, OECD Publishing 2020.

Poterba, James M., "Retail Price Reactions to Changes in State and Local Sales Taxes," *National Tax Journal*, 1996, *49* (2), 165–176.

Ramsey, Frank P., "A Contribution to the Theory of Taxation," *Economic Journal*, 1927, *37* (145), 47–61.

Taubinsky, Dmitry and Alex Rees-Jones, "Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment," *Review of Economic Studies*, 2018, *85* (4), 2462–2496.

Tax Commission, Government of Japan, "Report on the Reduced Tax Rate System for Consumption Tax," Technical Report, Cabinet Office 2018.

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.

A. Data Appendix

A.1 Data Sources

The primary data source is the Consumer Price Index (CPI) published monthly by Japan's Statistics Bureau, Ministry of Internal Affairs and Communications. The CPI covers approximately 585 items across 10 major groups. I extract the following sub-indices at the national level:

- **Eating out** (major group: "eating out"): Covers restaurants, cafeterias, school lunches, prepared meals eaten on premises. Weight in CPI basket: approximately 7.4%.
- **Cooked food** (under "food"): Covers prepared meals (*souzai*), boxed lunches (*bentō*), and ready-to-eat items. Weight: approximately 3.8%.
- **Alcoholic beverages** (under "food"): Beer, sake, shōchū, whiskey, wine, and other spirits. Weight: approximately 1.6%.
- **Non-alcoholic beverages, all food, all items**: Auxiliary series.

A.2 Chain-Linking Procedure

Japan's Statistics Bureau resets the CPI base year every five years. During my sample period (2015–2024), two base years are relevant: 2015-base indices (published for 2015–2020) and 2020-base indices (published from 2020 onward). I chain-link these series using the overlap year (2020):

$$\mathrm{CPI}_t^{\text{2020-base}} = \mathrm{CPI}_t^{\text{2015-base}} \times \frac{100}{\mathrm{CPI}_{2020}^{\text{2015-base}}} \tag{8}$$

This preserves the proportional changes within each base-year era while creating a continuous series indexed to 2020 = 100.

A.3 Sample Construction

The full sample spans January 2015 to December 2024 (120 months). The summary statistics (Table 1) use the pre-treatment window October 2017–September 2019 (24 months) for period comparisons, but the DD regressions use all available pre-reform data from January 2015 onward. The post-treatment period begins October 2019. The "Pre-COVID" sample drops February 2020 onward, yielding January 2015–January 2020 ($N = 61$). The "clean" post-treatment window (October 2019–January 2020, 4 months) excludes the COVID-19 period.

Event time is measured in months relative to October 2019: $t = 0$ for October 2019, $t = -1$ for September 2019, $t = 1$ for November 2019, etc.

B. Identification Appendix

B.1 Pre-Trend Test Details

The formal pre-trend test regresses the log relative price on a linear time trend within the pre-treatment period, controlling for month fixed effects. (In the main text, I describe this as October 2017–September 2019, which is the window used for summary statistics; the regression itself uses all available pre-reform data.)

$$\text{LogRelPrice}_t = \alpha + \gamma \cdot t + \delta_m + u_t \tag{9}$$

The estimated slope $\hat{\gamma} = 1.25 \times 10^{-5}$ with $p = 0.75$ provides no evidence of differential pre-trends. This is consistent with the visual evidence in Figure 1, where pre-treatment event study coefficients fluctuate around zero.

B.2 Panel Event Study Details

The panel event study (Equation (7)) uses semi-annual bins to balance precision against the number of event-time indicators. The bins are defined as: $[-24, -19]$, $[-18, -13]$, $[-12, -7]$, $[-6, -1]$ (omitted), $[0, 5]$, $[6, 11]$, $[12, 17]$, $[18, 23]$, $[24, 29]$, and $[30+]$. Each bin interacts with an eating-out indicator.

The specification uses a panel of two categories (eating out and cooked food), excluding alcohol because it experiences its own tax change at the reform date. Category fixed effects absorb level differences and time fixed effects absorb common monthly shocks. Standard errors use the HC1 heteroskedasticity-robust estimator. With only two categories, clustering at the category level would yield only two clusters—far too few for reliable cluster-robust inference (Cameron et al., 2008). HC1 standard errors are therefore the appropriate choice for this panel structure.

Table 6 reports the binned event study estimates restricted to event time ≥ -24 (87 months $\times$ 2 categories = 174 observations). The figure (Figure 1) uses the full monthly time series of the log relative price. The earliest pre-treatment bin $[-24, -19]$ shows a small negative coefficient (-0.006, $p < 0.01$), indicating the relative price was slightly below its pre-reform average during this period. The subsequent bins $[-18, -13]$ and $[-12, -7]$ are close to zero and insignificant, confirming that any early deviation dissipated well before the reform. This pattern is consistent with transient variation rather than a systematic pre-trend.

The post-treatment bins show a sharp positive shift at $[0, 5]$ that persists through $[24, 29]$. The final [30+] bin shows a reversal (-0.008), likely reflecting COVID-related disruptions to dining patterns in late 2022 and beyond.

Table 6: Panel Event Study: Semi-Annual Bin Estimates

Event-Time Bin	Coefficient	SE
$[-24, -19]$	-0.006	(0.001)
$[-18, -13]$	-0.001	(0.001)
$[-12, -7]$	0.002	(0.001)
$[-6, -1]$	(omitted)	
$[0, 5]$	0.019	(0.001)
$[6, 11]$	0.021	(0.001)
$[12, 17]$	0.024	(0.001)
$[18, 23]$	0.024	(0.001)
$[24, 29]$	0.014	(0.001)
[30+]	-0.008	(0.002)

Notes: Coefficients from Equation (7): interaction of event-time bin with eating-out indicator. Panel of two categories (eating out and cooked food) $\times$ monthly observations, restricted to event time ≥ -24 ($N = 174$). Category and time FE included. HC1 robust SEs. The omitted bin is $[-6, -1]$.

C. Robustness Appendix

C.1 Full Robustness Summary

Table 7 summarizes all robustness checks. The main DD estimate (full sample: 0.0078; pre-COVID: 0.0204) is robust to:

- **Placebo timing:** Small placebo at October 2018 (0.004, significant but one-fifth the treatment effect); insignificant at October 2017 (0.002).
- **Levels specification:** The levels DD yields 0.775 index points, corresponding to a 0.79% shift.
- **Pre-trend test:** Slope of 1.25×10^{-5} ($p = 0.75$)—no pre-trend.

- **Bandwidth sensitivity:** Stable estimates from ± 12 to ± 48 months.
- **COVID controls:** Consistent results when dropping COVID months, adding COVID dummies, or interacting Post with COVID.

Table 7: Robustness Summary

Test	Estimate	SE	N	Interpretation
Main DD (full)	0.0078	(0.0023)	120	Attenuated by post-COVID reversion
Pre-COVID	0.0204	(0.0011)	61	Full pass-through
Narrow window	0.0203	(0.0020)	28	Robust to shorter pre-period
Placebo Oct 2018	0.0040***	(0.0008)	57	Small; $\frac{1}{5}$ of treatment effect
Placebo Oct 2017	0.0017	(0.0011)	57	Insignificant at placebo date
Levels (not log)	0.775***	(0.2301)	120	Consistent in levels
Pre-trend slope	1.3×10^{-5}	(4.0×10^{-5})	57	No pre-trend ($p = 0.75$)

Notes: All specifications include month FE. HC1 robust SEs. Pre-COVID window: Jan 2015–Jan 2020. Narrow window: Oct 2017–Jan 2020. Placebos use pre-Oct 2019 data only ($N = 57$ months).

C.2 COVID Interaction Details

The COVID interaction specification estimates:

$$\text{LogRelPrice}_t = \alpha + \beta_1 \cdot \text{Post}_t + \beta_2 \cdot \text{COVID}_t + \beta_3 \cdot \text{Post}_t \times \text{COVID}_t + \delta_m + \varepsilon_t \tag{10}$$

Results: $\hat{\beta}_1 = 0.0204$ (the clean tax effect), $\hat{\beta}_3 = -0.0139$ (COVID partially offset the relative price increase as pandemic dynamics favored takeout). The tax effect is cleanly identified in the pre-COVID months, and the COVID interaction captures the subsequent relative price dynamics.

D. Heterogeneity Appendix

The aggregate CPI data do not permit heterogeneity analysis by household income, region, or firm type. However, the detailed CPI provides item-level price indices within the eating-out category (e.g., hamburgers, sushi, ramen, coffee). Future work using these item-level indices could examine whether pass-through rates differ by restaurant type or meal price point—testing, for example, whether budget restaurants (with thinner margins) exhibit different pass-through behavior than fine dining establishments.

The pass-through decomposition bar chart (Figure 7) provides a visual comparison of observed versus predicted price changes across categories.

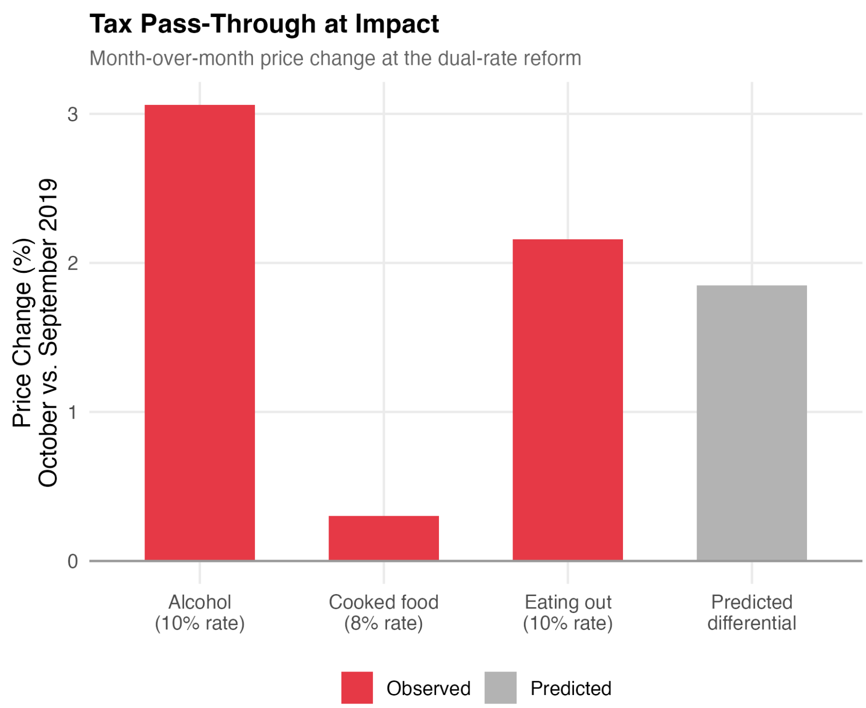

Figure 7: Tax Pass-Through at Impact

Notes: Month-over-month price change (October vs. September 2019) by food category. "Predicted" is the theoretical differential of 1.85% = 2/108 under full pass-through. The observed differential (eating out − cooked food) is 1.86%, implying 100.4% pass-through.

E. Additional Figures and Tables

No additional exhibits beyond those presented in the main text and preceding appendices.

F. Standardized Effect Sizes

Table 8: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	Classification
Log relative price (eat-in/takeout)	DD, Table 2 Col. 2	0.0204	—	0.0032	6.38	Large positive
Log relative price (eat-in/takeout)	DD, Table 2 Col. 1	0.0078	—	0.0032	2.44	Large positive
Log CPI eating out (DDD panel)	DDD, Section 5.2	0.0077	—	0.0325	0.24	Large positive

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For binary (0/1) treatments, SDE $= \hat{\beta}/\text{SD}(Y)$ and the SD(X) column is marked "—". SD(Y) is the unconditional standard deviation before conditioning on fixed effects. For the log relative price, SD(Y) $= 0.0032$ is computed from the log-transformed ratio (not reported directly in Table 1, which shows the level ratio SD $= 0.32$ index points).

Research question: Does Japan's 2019 dual-rate consumption tax (8% takeout vs. 10% eat-in) pass through to food prices? **Treatment:** Binary Post indicator (0/1) for months $\geq$ October 2019. **Data:** Japan Statistics Bureau CPI, monthly national indices, January 2015–December 2024 ($N = 120$ months; $N = 360$ for panel). **Method:** Difference-in-differences and triple-difference with month/category fixed effects, HC1 robust standard errors. **Sample:** National-level CPI indices for eating out, cooked food, and alcoholic beverages.

The large SDE for the log relative price reflects the extremely low pre-treatment variance of this series (SD = 0.0032, or 0.32 percentage points), implying that the tax reform produced a price shift equivalent to over 6 standard deviations of the outcome's pre-reform variation. The DDD panel SDE (0.24) is more moderate because SD(Y) for the level log CPI series is naturally larger.

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).