Equity and Efficiency of Consumption Taxes: Evidence from India's GST

Pierre Bachas, Davi Bhering, Pulak Ghosh*

Abstract

When income taxes and transfers are constrained, should governments redistribute through higher taxes on luxury goods? To quantify this equity-efficiency tradeoff, we exploit a large VAT reform in India that reduced tax rates on a broad set of non-essential goods. Rate cuts are partially passed through to consumer prices, generate moderate quantity responses, and induce little firm-product entry and exit or product misclassification. Together, these findings point to limited efficiency costs from differentiated commodity taxation. Combined with evidence that luxury consumption is a powerful tag for household income, our estimates imply that luxury goods should face substantially higher tax rates than standard goods under India's current tax and transfer system. We then show what it would take for uniform commodity taxation to be preferred. Achieving this benchmark requires simultaneously broadening the income tax base, increasing its progressivity, and sharply improving transfer targeting, well beyond current practice in most developing countries.

JEL codes: O23, H23, H21

*Pierre Bachas: Senior Economist, World Bank Research and International Tax Observatory pbachas@worldbank.org; Davi Bhering: PhD Candidate, Paris School of Economics, davi.bhering@psemail.eu; Pulak Ghosh: Professor, Indian Institute of Management Bangalore, pulak.ghosh@iimb.ac. This work was supported by the World Bank Knowledge for Change Program (KCP IV). We are indebted to Sam Asher and Lucie Gadenne who made this project possible. We thank Gaston Bernheim, Santiago Cesteros and Radhika Kaul for excellent research assistance. We benefited from comments by Amory Gethin, Felipe Lobel, Dario Tortarolo, Liam Wren-Lewis, Farah Zahir and Gabriel Zucman. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not represent the views of the World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent. An abbreviated paper was presented at India's National Institute of Public Finance and Policy (NIPFP) 2025 State Forum, and published as part of the conference's compendium.

1 Introduction

Most developing countries face tight fiscal space alongside persistently high income inequality (Chancel et al. 2025). Redistribution is conventionally assigned to income taxes and transfers, while consumption taxes are expected to remain broad-based and uniform (Atkinson and Stiglitz 1976). This benchmark becomes less compelling when informational, administrative, or political constraints limit the reach or progressivity of income taxes and transfers (Jensen 2022). In such settings, theory predicts that differentiated commodity taxation should complement direct redistribution (Deaton 1977). Yet the quantitative importance of this insight is unclear, as its redistributive benefits and efficiency costs are difficult to measure.

This paper quantifies the extent to which differentiated commodity taxation can complement constrained income taxes and transfers. While value-added taxes (VATs) often rely on reduced rates for essential goods, luxury consumption is far more strongly associated with income and provides a powerful tag for redistribution. Figure 1 illustrates this contrast using India's household expenditure survey. Under full pass-through, reducing tax rates on necessities implies that for every dollar received by the poorest households, the richest receive fifty cents. By contrast, almost all of the burden from a luxury tax falls on the top decile.

A central question is whether the redistributive gains from taxing luxury goods outweigh the efficiency costs of differentiated commodity taxation. A common objection to multi-rate VAT systems is that they distort consumption while encouraging product misclassification and changes in firms' product mix. These concerns may be particularly important in developing countries, where limited administrative capacity facilitates evasion and informality pressures formal firms. Evaluating this trade-off requires estimating incidence and behavioral responses to large tax changes across many goods, yet most VAT reforms are small or narrowly targeted. We study India's 10-percentage-point rate cut on a broad set of non-essential goods using administrative tax data. Our product-level difference-in-differences estimates imply moderate efficiency costs but substantial equity gains, supporting a sizable optimal luxury-tax premium under India's current tax and transfer system. Eliminating this premium requires

large, simultaneous expansions in income-tax coverage and progressivity together with major improvements in transfer-targeting capacity.

Our empirical setting is India's Goods and Services Tax (GST), a value-added tax that applies different rates across goods partly to achieve redistributive objectives. The standard rate is 18 percent, while a 28 percent luxury rate applies to a broad set of non-essential goods, including cosmetics, furniture, and appliances. In November 2017, roughly three-quarters of these goods were reassigned to the 18 percent bracket, generating large, plausibly exogenous variation in indirect tax rates across goods. The reform was intended to build political support for the recently introduced GST and was facilitated by a federal guarantee that compensated states for lost revenue.

In a first part, we study how registered firms respond to this sharp rate cut using GST returns from Karnataka—the fifth-largest state by GDP—covering monthly filings from July 2017 to March 2019. Firms report quantities, sales, and taxes at the product level. Our core sample comprises 28,000 firms and 844 products (HSN-4 level), totaling 2.4 million firm-product-month observations. We estimate behavioral responses using a difference-in-differences design that compares, within firms, products affected by the rate cut to those that were not. The setting features a single treatment period, no treatment reversal, and a uniform 10-percentage-point tax change. Multi-product firms account for the vast majority of sales, and allow us to isolate product-level responses within firms. We assess potential spillovers using variation in treatment intensity and firms that never report treated products.

Our empirical analysis yields four main findings on the incidence and efficiency effects of differentiated commodity taxation. First, the 10-percentage-point statutory rate cut reduced effective tax rates by 8.1 percentage points on average. The remaining gap reflects the fact that some firms already applied rates below 28 percent to treated products before the reform, either strategically or due to system complexity.

Second, we find incomplete pass-through of the tax cut to consumer prices. Prices of treated products fall by 3.4 percent, implying only 51 percent pass-through of the effective tax

reduction, despite anti-profiteering provisions. This estimate is within the range of evidence from other VAT cuts globally (Benzarti et al. 2020), suggesting that commodity taxes are borne jointly by consumers and producers.[1]

Third, quantity responses imply moderate efficiency costs of maintaining higher rates on these goods. At the extensive margin, we find no evidence of increased entry of treated products by either incumbent or new firms. At the intensive margin, quantities rise by 3.5 percent in response to a 3.4 percent price decline, increasing reported sales by 6.7 percent. Behavioral responses are too small to offset the mechanical revenue loss from the rate cut, resulting in a substantial decline in tax revenue. Firm-level value added responds slightly more. Elasticities are larger among smaller firms filing simplified returns, while price pass-through is lower among the largest firms.

Fourth, we study whether differentiated commodity taxation induces spillovers to related but untreated products through product misclassification or substitution. While products such as make-up and watches were reclassified, related goods such as essential oils and clocks remained taxed at 18 percent. Using within-firm over-time comparisons, we find no evidence of spillovers to related products relative to more distant alternatives.

Two complementary analyses provide longer pre-trends and strengthen the credibility of our findings. Re-estimating the effects of the November 2017 reform using scanner data from a large hypermarket chain yields similar estimates, with slightly lower price pass-through and larger quantity responses, consistent with observed firm-size differences in the tax data. Using tax data, we also examine the July 2018 reform, which cut rates on remaining 28 percent goods, finding comparable effects.

The results are robust to a range of alternative samples and specifications. Because identification relies on multi-product firms, we test for SUTVA violations by examining whether control products sold by firms that also sell treated products respond differently from the same products sold by other firms, finding no evidence of spillovers. Although firms

[1]Producer prices increase significantly, consistent with higher profits for firm owners; however, we lack balance-sheet or corporate tax data to directly test this hypothesis.

report only their total consumer-sales share, our analysis focuses on final-consumer sales and is robust to stricter sample restrictions and excluding high-upstreamness products. Finally, our conclusions are unchanged when related products are defined using machine-learning measures instead of HSN classifications.

In a second part, we examine the equity–efficiency trade-off of differentiated commodity taxation and the role of luxury taxes as a complement to income taxes and transfers. We calibrate the analysis using the behavioral parameters estimated from the near-elimination of the 28 percent GST bracket.[2] The estimated quantity responses imply moderate efficiency costs. Using a sufficient-statistic approach, we estimate a deadweight loss of 14 cents per tax dollar. Likewise, the tax cuts are far from self-financing: each mechanically lost dollar of revenue is offset by just 23 cents from higher sales.

To assess consumer incidence and targeting, we combine household expenditure data with statutory tax rates and our pass-through estimates. The GST was initially mildly progressive, but the reform weakened this progressivity by disproportionately reducing tax burdens on higher-income households. Moreover, the luxury tax bracket was poorly targeted: goods in the top bracket were no more concentrated among high-income households than those in the standard bracket.

We then combine the estimated efficiency costs of taxation with household consumption patterns in a multi-person model of optimal commodity taxation following Diamond (1975). Using the estimated quantity elasticities, pass-through rates, and inverse-income social welfare weights, we calibrate the optimal commodity tax schedule. Defining luxury goods as those with the steepest Engel curves to match the baseline consumption share of the initial 28 percent GST bracket, we find that, abstracting from income taxation and targeted transfers, the optimal tax rate on luxury goods exceeds that on standard goods by 26 percentage points.

We next introduce income taxes and transfers to examine how limitations in their coverage, progressivity, and targeting shape the optimal luxury tax. At the time of the reform, India's

[2]Combined, the November 2017 and July 2018 reforms reduced the rate on 81 percent of top-bracket goods to 18 percent, leaving only a small set—such as "sin goods"—taxed at 28 percent.

income tax covered only the top 4 percent of households, the top statutory rate was 30 percent, and transfers remained imperfectly targeted. Under this system, the optimal luxury-tax premium is 24 percentage points and falls to near zero only under large, simultaneous expansions of income tax coverage and progressivity together with sharply improved transfer targeting, consistent with the Atkinson–Stiglitz benchmark.

Across a range of parameter values and calibration assumptions, the luxury-tax premium under the current PIT typically exceeds 20 percentage points. It falls below this level only when society places substantially less weight on redistribution or under assumptions addressing external validity. Specifically, our parameters are identified from a reform that compressed rate differentiation, whereas our counterfactual raises luxury tax rates. Relative to the observed reform, higher luxury tax rates could induce greater product misclassification and higher pass-through. Accounting for these differences conservatively yields a lower-bound estimate of 12 percentage points.

The implied welfare gains are economically meaningful. A revenue-neutral reform of optimal commodity taxes raises welfare by over 1 percent of expenditure for each of the bottom four deciles, with two-thirds of these gains achievable through optimal luxury taxation alone. These gains are notable given that they arise within an already differentiated GST whose average effective consumption tax rate is only 7 percent.

Taken together, our results indicate that differentiated commodity taxation can improve redistribution at moderate efficiency cost. Optimal tax rates on luxury goods remain substantially above the standard rate when structural or political constraints limit redistribution via income taxes and transfers. Our analysis abstracts from other considerations relevant to differentiated commodity taxation, including paternalistic motives, ethical considerations, and labor–leisure distortions (Saez 2002). Even so, our findings suggest that uniform commodity taxation is optimal only under more restrictive conditions than commonly assumed. This perspective may become increasingly relevant as technological change, globalization, and wealth concentration reshape labor and capital taxation (Korinek and Lockwood 2026).

The remainder of the paper is organized as follows. Section 2 describes the institutional context. Section 3 presents the data and empirical strategy. Section 4 reports the administrative-data results. Section 5 analyzes scanner data. Section 6 develops and calibrates the optimal luxury tax model. Section 7 concludes.

1.1 Related Literature and Contributions

Our paper contributes to the literature in three main ways. First, we estimate the key parameters governing the efficiency costs of differentiated commodity taxation, including quantity elasticities, tax pass-through, extensive-margin responses, and product misreporting. While a seminal theoretical literature has identified these parameters as central inputs into optimal commodity taxation (Belan and Gauthier 2006; Boadway et al. 1994; Diamond 1975; Ramsey 1927), empirical work has not, to our knowledge, estimated this full set of parameters across a broad range of goods.[3] This complements a growing empirical literature on indirect taxation that instead studies specific features of VAT systems, including exemption design (Gadenne et al. forthcoming; Liu et al. 2021), withholding (Brockmeyer and Hernandez 2022; Garriga and Tortarolo 2024; Waseem 2022), and tax evasion (Almunia et al. 2023; Gupta et al. 2026; Naritomi 2019; Okunogbe and Pouliquen 2022; Pomeranz 2015; Waseem 2023).

Our second contribution is to show that differentiated luxury taxation can play an important redistributive role and to quantify the implications for optimal tax design. While the poor targeting of reduced rates on necessities is established and reflected in policy advice (Bachas et al. 2024b; Crawford et al. 2010; International Monetary Fund and World Bank 2018), we show that the case for luxury rates is fundamentally different. Luxury consumption provides a much stronger signal of household income than consumption of necessities, making differentiated luxury taxation a more effective tool for improving equity. This is particularly relevant in contexts where informational, implementation and political constraints limit both

[3]Related work on optimal commodity taxation estimated demand parameters using parametric demand systems (Deaton 1977; Gaarder 2019), while another strand examines the demand and welfare effects of taxes on "sin goods" (Allcott et al. 2019; Dubois et al. 2020).

the effectiveness of income taxation and the ability to accurately target transfers (Banerjee et al. 2024; Hanna and Olken 2018; Jensen 2022; Muralidharan et al. 2016). We quantify these implications by calibrating the optimal luxury-tax premium under alternative tax-and-transfer systems, contributing to the literature on second-best tax and transfer design (Bergeron et al. 2024; Best et al. 2015; Gadenne et al. 2024; Huang and Rios 2016; Niehaus et al. 2013). Our results show that the uniform commodity tax benchmark becomes relevant in practice only after large, simultaneous expansions in income tax coverage and progressivity, together with substantially improved transfer targeting.

Finally, we contribute to the literature on the incidence of commodity taxation. Prior work documents substantial heterogeneity in tax pass-through across products, markets, and institutional settings (Agarwal et al. 2025; Benedek et al. 2020; Benzarti 2024; Benzarti et al. 2024; Brusco and Velayudhan 2025; Harju et al. 2018). By estimating consumer incidence across a broad set of non-essential goods—including luxury goods—we provide new evidence on how tax incidence varies with luxury status. We find no evidence that more luxurious goods exhibit different pass-through rates or quantity responses, suggesting that, in our setting, taxing luxury goods does not entail systematically larger efficiency costs than taxing other non-essential goods.

2 Context

India's Goods and Services Tax (GST) The GST, introduced in July 2017, is India's main indirect tax and represented one of the most significant fiscal reforms in the country's recent history. Before its introduction, India's indirect tax landscape was fragmented: each state levied its own VAT while the central government imposed additional indirect taxes. The GST replaced 17 such taxes, creating a unified system intended to reduce cascading taxation on interstate trade and to harmonize tax rules across the country, as captured in the reform's motto, *"One Nation, One Tax."*

The GST is structured as a standard value-added tax: firms remit tax on their sales and can claim input tax credits on eligible purchases. Almost all goods and services are covered, including imports, with exclusions such as liquor, petroleum, electricity, and land. Labor costs are not creditable. GST revenue follows the destination principle. For intra-state transactions, the tax is split evenly between the central and state governments. For inter-state transactions, the tax is levied centrally and then apportioned between the destination state and the center. A portion of the central government's GST revenue (41%) is then redistributed to states according to rules based on population, income, and area. Overall, states retain over 70 percent of total GST proceeds (PRS Legislative Research 2024). The GST Council governs the system by setting tax rates and assigning products to tax brackets. States hold two-thirds of the voting power, with the central government holding the remaining one-third.

Tax Rate Setting at the Onset of GST The GST Council assigned products to five tax brackets: exempt, 5 percent, 12 percent, 18 percent, and 28 percent. The allocation of products to these rates followed two guiding principles: preserving progressivity and maintaining continuity with pre-GST tax rates. Progressivity was a central design criterion. Essential goods were placed in lower-rate brackets, while non-essential and luxury items were assigned higher rates. As noted in the minutes of the 4$^{\text{th}}$ GST Council meeting, *"The rate structure should not be regressive in nature."*

Continuity with the pre-GST tax structure was also a central design priority. The Council sought to keep GST rates close to the combined federal and average state tax rates that prevailed before the reform. Given the political and technical challenges of the rollout, this continuity was viewed as a pragmatic way to facilitate implementation. As noted in the minutes prepared ahead of the GST roll-out, *"the Committee had recommended the rates after taking into account the present tax incidence on account of Central Excise, Service Tax and VAT."* The minutes further map each GST bracket to a range of pre-GST combined VAT and excise rates. For example, goods taxed at 15–21 percent before the reform were

assigned to the 18 percent GST bracket, with analogous mappings across all brackets. The initial GST schedule was thus calibrated to replicate the pre-GST tax structure on average.

Product Classification under GST The GST uses the Harmonized System of Nomenclature (HSN) established by the World Customs Organization to classify goods and define tax rates. The HSN has three main levels: 96 HSN-2 chapters, 1,228 HSN-4 headings, and 5,612 HSN-6 subheadings. India further differentiates some HSN-6 subheadings by adding two digits, creating HSN-8 codes. Most GST rates are defined at the HSN-4 level, though some are specified at finer levels of detail.[4]

The November 2017 Rate Cut We exploit a 10 percentage point tax cut applied to most products initially taxed at 28 percent. This reduction took effect on November 15, 2017, four and a half months after the GST was introduced. Several political and institutional factors contributed to the decision. The central government had a strong incentive to boost the popularity of its flagship economic reform through a visible rate cut. The introduction of the GST drew criticism for its extensive rate differentiation, despite being framed as a tax-simplification reform. These concerns threatened to undermine the reform's primary objective of unifying the tax system rather than changing overall tax burdens. Criticism focused particularly on the 28 percent bracket, which many viewed as excessively high for a broad set of goods. As we show later, it was only imperfectly targeted at luxury consumption.

From a revenue perspective, states—holding two-thirds of the votes in the GST Council—had little incentive to oppose the rate cut. The central government had guaranteed 14 percent annual GST revenue growth during the reform's first five years, effectively shielding states from the fiscal consequences of lower tax rates.

The combination of political pressure, public dissatisfaction, and the revenue guarantee

[4]Table A.1 illustrates the HSN hierarchy. For example, powder face is an HSN-8 code nested within the HSN-6 category make-up or skin care powders, the HSN-4 category make-up or skin care preparations, and the HSN-2 category essential oils and resinoids; perfumery, cosmetic or toilet preparations. Level of detail varies across HSNs, with some branches more granular than others.

for states created a favorable environment for the rate cuts. Yet the short interval since the GST's introduction and the highly compressed decision process made the rate-cut reform largely unexpected, leaving little scope for anticipatory responses. The GST Council met on November 10, announced the decision the next day, and implemented it on November 15.

Karnataka The data is sourced from Karnataka, India's fifth-largest state economy, with a population of 65 million and a GDP per capita of INR 185,000 (USD 2,890) in 2018. Although the data is from a single state, the institutional features of the GST apply nationally. States primarily play an operative and enforcement role. Karnataka is considered a relatively high-capacity state within India (Ministry of Panchayati Raj and Indian Institute of Public Administration 2024). However, its own-source tax revenue relative to state GDP is close to the national average (PRS Legislative Research 2024), indicating that it is not an outlier in terms of fiscal capacity.

3 Data and Empirical Strategy

3.1 Tax Rates by Product

We construct a product-level panel of GST rates from July 2017 to March 2020 by combining two sources. First, we use the Ministry of Finance rate booklets for July 2017 and March 2020, which served as the official reference for firms. These booklets list HSN-4 codes and short product descriptions by tax bracket. Most HSN-4 codes appear only once, but some appear multiple times because narrower categories, typically at the HSN-6 level, face different tax rates. Although the booklets report HSN-4 codes, the accompanying descriptions usually identify the corresponding HSN-6 products, allowing us to match each entry to the HSN-6 classification. Comparing the 2017 and 2020 booklets allows us to identify each product's initial and final tax rates.

Second, we use GST Council meeting minutes, which record rate changes in tables that

mirror the booklet format. We systematically extract these tables to date each change and validate the changes inferred from the booklets.

The resulting dataset contains monthly GST product rates from July 2017 to March 2020. Most products correspond to HSN-4 codes (91%); the remainder are groups of HSN-6 codes within an HSN-4 sharing the same rate (9%) (Table A.2). We refer to these HSN-based tax-rate categories as *products* throughout the paper.

3.2 Data

Our main dataset consists of monthly GST returns for registered firms in Karnataka from July 2017 to March 2019. We complement them with a nationally representative household expenditure survey and scanner data from a large hypermarket chain. We use the official end-2017 exchange rate (64 INR/USD) for all currency conversions.

Firm-Level Tax Returns Firms with annual sales exceeding INR 2 million (USD 31,000) must register and file regular GST returns. Firms below the threshold can either register for GST voluntarily or file the composition scheme.[5] Firms that choose the composition scheme do not appear in our data as they file a separate return (GSTR-4). All GST-registered firms must file a monthly tax return (GSTR-3B), containing total sales, tax liabilities, and input tax credits. They must also file a detailed tax return (GSTR-1), with frequency depending on firm size. Firms with annual sales above INR 15 million (USD 234,000) must file it monthly, and firms below quarterly. This return includes a breakdown between business-to-business sales (GSTR-1 Table 4) and business-to-consumer sales (GSTR-1 Tables 5 and 7).

Firm–HSN Tax Returns Firms with annual sales above INR 15 million are required to file an HSN-level sales return (GSTR-1 Table 12), reporting monthly sales, taxes, quantities, and units by HSN code. Half of firm–HSN–month observations are reported at the HSN-4

[5]The composition scheme levies a turnover-based tax at a 1 percent rate and does not allow input tax deductions. Compliance requirements are also simplified compared to the standard GST regime.

level, where most rate changes occur, and half at the HSN-8 level. Reporting granularity is a firm-level characteristic and remains stable over time and across HSN codes. We therefore retain firms' reporting level rather than aggregate HSN codes, as aggregation would make quantities and implied prices less meaningful.

These data allow us to construct unit values that proxy for pre-tax and tax-inclusive (consumer) prices. Our main analysis exploits firm–HSN tax returns, which are rarely available in standard VAT data because they typically neither distinguish goods nor report quantities. Although only 26 percent of firms in Karnataka report HSN-level sales, they account for three-quarters of sales, allowing our product-level analysis to cover most economic activity.

HSN Registration File To complement our analysis of smaller firms that do not file the HSN-level return, we use the GST registration file (REG-01). At registration, firms list the top five HSN codes they expect to sell. Although non-binding, these codes provide indicative information for firms that file only the firm-level return.

Hypermarket Barcode-Level Data The data comprise barcode-level transactions from one of the largest retail chains in Karnataka, covering April 2016 to July 2019. Barcodes identify individual goods by brand and package size (e.g., "Heinz Ketchup 200g Bottle" and "Nivea Men's Deo Sport 150ml"). For each barcode, we observe quantities, pre-tax and tax-inclusive sales, and discounts. We compute taxes paid as the difference between tax-inclusive and pre-tax sales.

India National Sample Household Survey This nationally representative survey contains household-level expenditure data for 345 survey item codes, which we match to HSN codes (typically at the four-digit level). We use the 2011–2012 wave, which is the most recent round available prior to the 2017 tax reform.

3.3 Sample Construction and Data Cleaning

Final Consumer Sales Given the credit–invoice structure of the VAT, taxes on sales to registered firms and final consumers have different implications. Registered firms recover VAT through input tax credits, whereas final consumers cannot. To assess the equity and efficiency implications of consumption taxes, we therefore focus on the burden borne by final consumers and restrict the analysis to non-creditable transactions. Final consumers include households and entities unable to claim input credits, such as unregistered firms, firms under the composition scheme, and firms selling only exempt products.

An empirical challenge is that firms do not report the breakdown between sales to businesses and final consumers at the product level. However, monthly firm-level VAT returns report the overall share of sales to each client type. Our main analysis focuses on firms selling primarily to final consumers, defined as those with a baseline (pre–November 2017) Business to consumer (B2C) sales share above 50 percent. This firm level restriction excludes half of the observations in the HSN-level data.[6] To assess the sensitivity of our results to this restriction, we perform several robustness checks: varying the B2C threshold defining final-consumer-facing firms; restricting the sample to products covered by the household survey or the hypermarket data, or classified as downstream by Antràs et al. (2012); and redefining the B2C share using the full sample period rather than only the baseline period.

Exclusion of Services We exclude services for three reasons. First, the Services Accounting Code (SAC) is much coarser than the HSN system for goods, with all services grouped under HSN-2 Chapter 99, limiting within-category tax variation. Second, GST rates for some services depend on characteristics omitted from SAC codes, such as hotel room tariffs or restaurant air-conditioning. Third, broad SAC categories aggregate heterogeneous activities, making quantities difficult to define.

[6]Forty-five percent are excluded because their baseline B2C share is below 50 percent, while the remaining 5 percent first appear after November 2017 and are assigned a missing baseline share. We prefer this approach because the reform itself may affect the B2C share, but relax this restriction when studying firm entry.

Data Cleaning To ensure data quality and sample consistency, we drop a total of 9.6 percent of observations, including those with invalid or missing HSN codes (because the GST filing portal does not provide a drop-down menu for HSN selection) and those with reported sales but missing quantities (see Table A.3).

Moreover, because the GST portal does not automatically map HSN codes to statutory rates, firms can report infeasible rates—that is, rates outside the statutory range applicable to a given HSN-4 code.[7] These errors are broadly symmetric and only weakly correlated with firm characteristics such as size, location, consumer base, and months observed (Appendix B). To abstract from this reporting issue, we exclude pre-reform observations with rates outside a one-percentage-point band around the statutory range for the corresponding HSN-4 code. This removes 3.7 percent of observations and ensures that differences between effective and statutory rates reflect meaningful variation rather than reporting error.[8]

Finally, because quantities may be reported inconsistently and HSN codes may encompass heterogeneous product bundles, we exclude firm–HSNs whose maximum-to-minimum price ratio exceeds 100, affecting 1.4 percent of observations.

3.4 Treatment Groups: Definition and Description

Products with a Rate Cut The main analysis focuses on products experiencing a rate reduction from 28 percent to 18 percent in November 2017. Of the 1,367 products, 164 experienced this reduction (Table A.4). An additional 108 products experienced rate changes in other periods, most notably in July 2018, which we analyze separately, or changes across different rate brackets in November 2017.

Table A.5 tabulates the share of products in each tax bracket at the onset of the GST (July 2017) and at the end of our sample (March 2020). In July 2017, 17 percent of products

[7] Most HSN-4 codes correspond to a single rate; a minority map to multiple HSN-6 codes with different rates (e.g., 18 and 28%), in which case the full range is considered feasible.

[8] Results are robust to including these observations (Appendix B). Although doing so weakens the relationship between statutory and effective tax rates, elasticities with respect to effective tax rates are unchanged, indicating stable behavioral responses to taxes paid. We view the restriction as improving external validity, though unrestricted statutory-rate elasticities may be more relevant for the November 2017 reform.

were taxed at 28 percent. By March 2020, 88 percent of these had moved to the 18 percent bracket, leaving only 2 percent of all products taxed at 28 percent. Most of this adjustment occurred in November 2017. The near-elimination of the 28 percent bracket is consistent with limited "cherry-picking" and reflects the government's objective of simplifying the GST. Products that remained at 28 percent were largely those taxed for externality or internality reasons, such as tobacco and soft drinks, akin to excise taxes. The main exception was large cars. Rate changes outside the shift from the 28 to the 18 percent bracket were more targeted and appear to have been driven in part by product-specific considerations.

Table A.6 lists the 30 most common products experiencing a rate reduction, ranked by their number of firm–product–month observations before November 2017. The table underscores the breadth of the reform: treated items span food products (e.g., chocolate, pastries), hygienic goods (e.g., make-up, shaving), house goods (lamps, mirrors), and construction materials (e.g., wires, ceramics). No single product dominates the sample: none accounts for more than 10 percent of pre-reform observations and the top 30 products collectively represent 76 percent of treated-product observations.

Products in the Control Group To construct a comparable control group, our main specification focuses on products always taxed at 18 percent or 12 percent. Using the household survey, we compare the distribution of consumption across tax brackets by income. Figure A.1 shows that products taxed at 12 percent and 18 percent have consumption profiles similar to those initially taxed at 28 percent, reflecting the limited targeting of the 28 percent bracket toward higher-income households. By contrast, exempt and 5 percent products are consumed disproportionately by poorer households and are therefore less comparable.

As with the treated products, control products are heterogeneous (Table A.7), although "Medicaments in Measured Doses" accounts for 22% of observations and is excluded in a robustness check. The top 30 products account for 85% of the control sample, including medical items (e.g., drugs, medical instruments), hygiene products (e.g., soap, insecticides),

household goods (e.g., brooms, computers, pens), construction materials (e.g., pipes, taps, glue), and food products (e.g., jams, pastries, butter).

Descriptive Statistics Table 1 summarizes the main sample used to analyze the November 2017 rate cut. It includes 844 products: 164 treated (28% to 18%) and 680 control products (always taxed at 18% or 12%). Since firms often sell multiple HSN codes, the data contain 445,040 unique firm–HSN pairs. The resulting panel is unbalanced because firms enter and exit and their product mix changes over time. On average, a firm–HSN appears in 5.3 of the 21 months, yielding 2,355,873 firm–HSN–month observations. Sales are highly skewed: before November 2017, median monthly sales per firm–HSN equal INR 7,440 (USD 116) in months with positive sales, while the mean is INR 364,000 (USD 5,690). Treated and control firm–HSN observations are similar across most pre-reform characteristics.

Most firms sell multiple products. Table 2 shows that, among the 42,899 tax returns filed by the 17,783 firms before November 2017, 66 percent reported sales of more than one product, a share that rises to 89 percent when sales weighted. Firms tend to sell both treated and control products within a given month: 42 percent of tax returns (57% sales-weighted) contained at least one treated product, and among these returns, 81 percent (88% sales-weighted) also reported sales of a control product. On average, firms reported sales of 7.6 different products per month during the four pre-reform months, of which 1.4 were treated, 3.3 were control, and 2.9 were excluded.

3.5 Empirical Strategy

To estimate the causal impact of the rate reduction, we use a standard difference-in-differences (DiD) design with a single treatment date in November 2017, no treatment reversals, and a binary treatment indicator capturing exposure to the 10-percentage-point tax cut. Our specification compares products that experienced a rate cut in November 2017, which constitute the treated group, with products taxed at 18 percent or 12 percent that did not

experience a rate change, which form the control group:

$$y_{iht} = \sum_{t=Jul17,\neq Oct17}^{Mar19} \beta_t \cdot \text{treat}_h \times \text{month}_t + \alpha_{ih} + \gamma_{it} + \epsilon_{iht} \tag{1}$$

where y_{iht} denotes the outcome of interest for firm i, HSN h, in month t. The indicator treat_h equals 1 if HSN h was subject to the November 2017 rate cut (28% to 18%) and 0 otherwise. α_{ih} and γ_{it} are firm-HSN and firm-time fixed effects, respectively, which control for unobserved heterogeneity across HSNs within firms and for firm-level shocks over time. The coefficient of interest, β_t, captures the difference in outcomes between treated and control HSNs relative to the omitted period, October 2017 ($\beta_{Oct17} = 0$). Standard errors are clustered at the firm and HSN-4 level.

Identification relies on the parallel trends assumption that, absent the reform, outcomes for treated and control products within firms would have followed similar trends, which we assess by examining pre-treatment dynamics. We also report estimates from a before–after DiD specification comparing the pre-reform period (July–October 2017) with the post-reform period (April 2018–March 2019). This specification allows for a transition period (November 2017–March 2018) and captures a full year after the reform, providing a transparent summary measure of the treatment effect.

The availability of firm–HSN information allows us to control for time-invariant differences across firms in the items reported under a given HSN. Although firms may not sell identical items within the same HSN, the composition of items reported under a given code is unlikely to change over time, mitigating concerns about within-HSN composition or quality shifts. Moreover, the granularity of the firm–HSN–month panel, together with the fact that most firms sell multiple products, allows us to exploit within-firm variation across products while controlling for firm-level shocks that could otherwise bias the estimates. Since most firms are multi-product firms, this within-firm variation is central to our identification.

One concern is that the rate cut may also affect control products sold by treated firms through spillovers (e.g. if lower prices on treated products attract new consumers). Such

spillovers would violate the Stable Unit Treatment Value Assumption (SUTVA). To address it, we bound the plausible magnitude of spillover effects on control products.

Tax and Price Elasticities Building on this identification strategy, we estimate quantity and sales elasticities with respect to effective tax rates and consumer prices by instrumenting these variables with interactions between treatment status and period indicators (pre, transition, and post). Because the statutory tax cut is common across treated products, the resulting elasticities closely approximate the difference-in-differences coefficients rescaled by the statutory rate change and the corresponding change in effective tax rates.[9] Section 6.2 discusses how these estimates map into the elasticities used for optimal taxation.

4 Administrative Tax Data Analysis

4.1 Main Results

From Statutory to Effective Tax Rates Statutory tax rates (STRs)—the rates set by law—need not translate one-for-one into effective tax rates (ETRs), the rates firms actually remit. Because some firms already reported ETRs of 12 or 18 percent on treated products before the reform, they should respond little, if at all, to the statutory rate cut. These lower pre-reform ETRs may reflect reporting errors due to the complexity of the tax system or strategic under-reporting.[10] We therefore first estimate the reform's effect on ETRs before turning to prices and quantities.

Figure 2A shows the average STR and ETR for treated and control products. By construction, the STR of treated products fell from 28 percent to 18 percent. Prior to the reform, their average ETR was 26.9 percent and converged to 18 percent between November 2017 and February 2018. For control products, a mix of goods taxed at 12 percent and 18

[9] With only the post-period instrument, the 2SLS estimate is exactly equal to this ratio. Adding the transition-period instrument yields a weighted average of the period-specific ratios.

[10] In Section 6.3, we treat the initial STR–ETR gap as a reduced-form measure of evasion to calibrate the evasion response to rate dispersion when evaluating the optimal luxury tax rate.

percent, the STR remained at about 16 percent, while the average ETR slightly declined from 15.7 percent to 15.3 percent.

Figure 2B reports event-study estimates from Equation 1.[11] The reform reduced ETRs by 8.1 percentage points, relative to a 10-point statutory cut, implying an 81 percent pass-through of the statutory reduction to effective tax rates. This incomplete pass-through points to frictions in multi-rate VAT systems. In Section 6, we incorporate this baseline STR–ETR gap in an extension that accounts for misreporting costs associated with rate differentiation.

Price Incidence The extent of tax pass-through determines whether the gains from the tax cut accrue to consumers or producers. Figure 3A shows that consumer prices decline gradually post reform, stabilizing after three months at about 0.034 log points below baseline. This decline is substantially smaller than the 0.066 log-point reduction in effective tax rates. Relative to effective rates, pass-through is 51 percent (s.e. = 9), allowing us to reject both zero and full pass-through.[12] Half of the tax cut is passed through to consumers despite anti-profiteering provisions in GST.[13]

Prices respond to realized changes in effective tax liabilities rather than statutory rates. Figure A.3 shows that firm–products with initially misaligned ETRs experience only small declines in ETRs and no detectable price response. By contrast, among firm–products initially applying the correct rates, prices fall by 0.039 log points alongside a 0.074 log-point decline in ETRs, implying 53 percent pass-through. These results suggest that pass-through estimates based on CPI data and statutory rate changes may understate true pass-through when effective tax adjustments are incomplete.

[11] Figure A.2 reports the event-study estimates for log ETR.

[12] Relative to the 0.081 log-point (10 percentage point) statutory rate cut, pass-through is 42 percent.

[13] GST law requires that "any reduction in the rate of tax (...) shall be passed on (...) by way of a commensurate reduction in prices" (Government of India 2017). Despite a dedicated enforcement framework, implementation was limited. By September 2018, 1,289 applications had been filed but only 31 investigations completed, mostly for very large firms, with recoveries of INR 150 crore (USD 23.4 million) (GST Council 2018). The provision was also legally challenged as the notion of a "commensurate" price reduction gave authorities broad discretion in measuring profiteering (Mishra and Phadke 2020).

Quantities We next examine how quantities respond to the consumer price reduction induced by the rate cut. Figure 3B presents estimates of Equation 1 for log quantities. In the first three months following the reform, quantities mirror the price dynamics, rising as prices gradually decline. Thereafter, treated products exhibit a persistent increase of 0.035 log points, marginally significant at the 5 percent level.

Relative to the statutory rate reduction, the estimates imply a quantity elasticity of -0.43. Instrumenting effective tax rates with interactions between treatment status and the pre-, transition-, and post-reform periods yields an ETR elasticity of -0.50. Instrumenting consumer prices with the same interactions yields a price elasticity of demand of -1.13.

Sales Pre-tax sales define the VAT base absent non-proportional input adjustments.[14] Their responsiveness to tax changes therefore provides a sufficient statistic for the revenue consequences of the reform. By construction, pre-tax sales equal quantities times pre-tax producer prices. Following the incomplete pass-through documented above, Figure 4A shows that producer prices increase after the rate reduction. Together with the increase in quantities, this raises pre-tax sales by 0.067 log points (Figure 4B). The implied elasticity of the tax base with respect to the effective tax rate is -0.99.

Extensive Margin The rate reduction may also affect the tax base along the extensive margin if firms begin selling, or reporting sales of, treated HSNs after the reform. To study this channel, we construct a balanced firm–HSN panel with a binary indicator for positive monthly sales. The reporting rate is the probability that a firm reports positive sales of an HSN in a given month, conditional on reporting it at least once during the sample period.[15]

[14] Firms' final tax liability depends on value added—sales net of input tax credits—which we analyze in Section 4.5, as it is observed only at the firm level rather than at the product level.

[15] We distinguish three extensive-margin responses: (i) firm–HSN entry or exit, when an incumbent firm starts or stops reporting a given HSN; (ii) firm exit, when a firm stops reporting all HSNs; and (iii) firm entry, when a new firm begins filing firm–HSN tax returns. The main sample, defined using pre-November-2017 B2C shares, mechanically excludes margin (iii) for firms not reporting before the reform and therefore captures only margins (i) and (ii). To incorporate all three margins, we redefine the sample using average B2C shares over the full period. Results are similar using the original sample, reflecting that most firms already reported before the reform (Figure A.4).

Figure 5A plots the reporting rate for treated and control HSNs. The probability that a firm reports a given HSN increases steadily over the sample period, but treated and control HSNs follow nearly identical trajectories. This common upward trend is consistent with the expansion of the GST system and overall economic growth.

We next estimate Equation 1 using a binary reporting indicator to capture within-firm extensive-margin responses. Figure 5B presents the estimated treatment effects. HSN subject to the rate cut exhibit no meaningful change in their likelihood of being sold. The point estimate is close to zero and precisely estimated: we reject increases above 0.1 percentage points, less than 1 percent of the initial 20 percent reporting rate.

4.2 Product Misclassification and Substitution

Misreporting and Substitution Mechanisms A first concern in multi-rate commodity tax systems is that their complexity creates opportunities for firms to misreport higher-taxed goods as lower-taxed ones, eroding the intended fiscal and distributional effects of rate differentiation. In our setting, firms may have reported products taxed at 28 percent as goods subject to lower rates. To the extent that misreporting entails detection risk, the November 2017 reform reduced incentives to misclassify and may have induced firms to stop doing so.[16]

A second concern is that differentiated tax rates distort relative prices. In our setting, the reform may have shifted demand toward products whose rates were reduced and away from related products. Unlike misclassification, such substitution reflects a genuine reallocation of consumption and therefore has welfare implications beyond its revenue effects. Although substitution and misclassification have distinct welfare consequences, they can generate similar empirical patterns when products are close substitutes or plausible candidates for misclassification. In both cases, we would expect reported quantities and sales of untreated related products to decline following the rate cut.

[16]This prediction follows from the standard tax-evasion framework with continuous and convex misreporting costs (Allingham and Sandmo 1972). It need not hold under fixed evasion costs or if changing behavior risks revealing past evasion to the tax authority.

Defining Related Products To make progress, we assume misreporting occurs toward related products, for which relabeling is easier to justify during an audit and more plausibly presented as an honest mistake. Because the GST relies on the HSN, with most rate differentiation at the HSN-4 level, we exploit its hierarchical structure to define relatedness. We partition untreated products into: (i) unrelated HSN-4 products, whose HSN-2 chapter contains no treated products; and (ii) related HSN-4 products, which share an HSN-2 code with at least one treated product. This classification isolates the products most susceptible to misclassification while providing a plausible control group (see Table A.8). Substitution is likewise more likely within HSN-2 categories, although economic substitutability varies across products. We therefore later restrict the set of potential substitutes to construct a tighter comparison group.

Table A.9 illustrates the classification by listing, for each HSN-2 chapter containing a treated product, the most common treated product and the most common related control product prior to the tax cut. While the degree of similarity within HSN-2 chapters varies, many treated–control pairs correspond to clearly related products (e.g., washing preparations and bar soap; sugar confectionery and other sugars), but others are more distinct (e.g., lighters and brooms; malt extract and bakery products).

Related Products Results We estimate Equation 1 at the firm–HSN level, comparing related and unrelated HSNs within firms. Under a misreporting mechanism, quantities and sales of related HSNs should decline relative to unrelated HSNs. Figure 6 presents event-study estimates for quantities, sales, consumer prices, and reporting rates. We find no evidence of misreporting. The quantity effect is instead positive (0.016 log points; s.e. $= 0.021$), allowing us to reject declines exceeding 2.5 percent at the 95 percent confidence level (Panel A). Sales move similarly (Panel B), consistent with the absence of price responses (Panel C). Finally, we find no evidence of extensive-margin adjustments: changes in reporting rates are economically small and precisely estimated around zero after the transition period (Panel D).

We further probe the absence of relabeling (and substitution) through three robustness exercises in Table A.10. First, mislabeling may arise only in HSN-2 chapters where treated products account for a large share of activity. Restricting the sample to chapters in which treated goods represent at least the median pre-reform sales share (≈15 percent) yields estimates slightly higher than the baseline. Second, responses may be concentrated among firms selling both treated and related products. Restricting the sample to such firms again yields similar results. Finally, allowing for substitution toward goods initially taxed at 5 percent—excluded from the main control group—leaves the main findings largely unchanged, although the estimated effect on consumer prices becomes statistically significant, consistent with negative spillovers to the prices of related goods.

LLM Matching of Related Products Our estimates rule out economically meaningful misreporting or substitution within the HSN-4 classification used to implement GST rate differentiation in India. This is consistent with the literature, which generally finds substitution only among very closely related goods.[17] We therefore use a large language model to identify the control products most closely related to each treated product, including at the HSN-6 level where available, following the procedure described in Appendix C. Using Equation (1), we estimate spillover effects separately for products classified as substitutes or complements to treated products, relative to neutral control products.

The results are reported in Table C.1. Estimated coefficients generally have the expected sign—negative for substitutes and positive for complements—but are small and statistically insignificant. Allowing each treated product to match multiple substitutes and complements yields no evidence of spillovers. Restricting the analysis to the closest substitute and complement produces a significant positive effect for complements but not substitutes.

[17] Across settings, substitution effects are typically confined to narrowly defined product categories. For example, Asatryan and Gomtsyan (2023) find substitution from taxed fresh beef and pork to poultry, but not to other food products. Hoenink et al. (2021) document modest within-category substitution for dairy products. Following Mexico's soda tax, substitution was limited to bottled water (Colchero et al. 2016), while Philadelphia's soda tax generated no clear substitution to other beverages (Cawley et al. 2019). More broadly, Benzarti et al. (2024) find no spillover effects across goods.

However, this effect disappears under the higher-confidence matching approach. Overall, these findings suggest that any household substitution or firm misclassification is confined to relationships finer than the HSN-4 classification used to implement India's GST.

4.3 Within-Firm Spillovers

Our identification relies on within-firm variation in treatment status across products, absorbing firm–time shocks. A potential concern is that the rate cut generates firm-level spillovers, affecting untreated products and violating SUTVA. To assess this possibility, we exploit heterogeneity in firms' baseline exposure to treated products.[18] If spillovers are quantitatively important, control products sold by firms with greater exposure should exhibit larger post-reform changes: increasing under positive spillovers, such as higher firm-level demand, and decreasing if treated products crowd out sales of untreated products. We therefore estimate a regression restricted to control products, using the firm's baseline treated-sales share, rescaled to the mean, as the explanatory variable:

$$\ln y_{ipt} = \gamma_t \left(\text{TreatedShare}_i \times \text{Month}_t\right) + \lambda_{ip} + \delta_{pt} + \varepsilon_{ipt}, \tag{2}$$

where y_{ipt} denotes the outcome for firm i, HSN p, at time t; TreatedShare_i is the base share of sales coming from treated products, rescaled by the average baseline share of treated sales among exposed firms (14%); λ_{ip} are firm–HSN fixed effects and δ_{pt} are product–time fixed effects. Identification comes from comparing, within periods and HSN codes, firm–product observations that were more exposed to treated products at baseline with those that were not exposed or relatively less so.

The estimates for the impact on control products of the average firm for consumer prices, quantities, and sales are -0.001, -0.006, and -0.007 log points, respectively (Figure A.5).

[18] Before November 2017, the median control product was sold by a firm for which treated products accounted for 3.5 percent of sales (mean: 11.7%). Among firms selling at least one treated product, the median exposure rises to 5.0 percent (mean: 14.0%), indicating substantial cross-firm variation in exposure.

These effects are substantially smaller than the direct effects of the rate reduction (-0.034, 0.035, and 0.067), amounting to roughly 3 percent, 17 percent, and 10 percent of their respective magnitudes. This suggests that within-firm spillovers are limited and unlikely to drive the estimated treatment effects.

4.4 Additional samples and reforms

Figure 7 summarizes robustness checks for pass-through, the sales and quantity elasticities, and the reporting rate (see Table A.11 for the full results). Estimates remain close to the baseline across specifications. The main exceptions are household-survey products and sales-weighted regressions, which yield somewhat larger quantity and sales responses. Because these specifications either focus on household consumption or place greater weight on high-sales products, we use them to provide conservative bounds on the implied demand elasticities in our calibration scenarios (Section 6). We then consider five sets of robustness checks.

The first set varies the control group, which in the baseline comprises products always taxed at 12 or 18 percent, by: (i) adding products taxed at 5 percent; (ii) retaining only products taxed at 18 percent; (iii) excluding *medicament*, a large regulated control product; and (iv) restricting the comparison to HSN-2 chapters containing a treated product or, conversely, excluding those chapters to limit indirect exposure to the reform.

Second, we examine robustness to the definition of final-consumer exposure. Because registered suppliers can claim tax credits on their inputs, our focus is on the incidence of rate changes on final consumers. In the baseline, firms are classified as final-consumer facing if most pre-reform sales are to final consumers. We then consider several alternatives: (i) restricting the sample to downstream products as defined by Antràs et al. (2012); (ii) products appearing in the household survey; (iii) products sold in the hypermarket data; (iv) firms with over 80 percent of their pre-reform sales to final consumers; and (v) classifying firms based on their share of sales to final consumers over the entire sample period rather than only before November 2017.

The third set relaxes data consistency filters. We reintroduce (i) firm-products with extreme price variation and (ii) firm-products applying infeasible HSN-4 rates.

The fourth set tests robustness to sample balance. While our baseline imposes no balance requirement, we consider: (i) a weak balance condition requiring firm–product observations to appear both before and after November 2017; (ii) a strong balance condition requiring products to be sold every month up to March 2018.

The fifth set tests weighting and aggregation. In the baseline, observations are kept at the firm–HSN level because quantities and prices cannot be coherently aggregated across HSN codes. Sales can. We therefore consider three alternatives: (i) weighting observations by pre–November 2017 sales; (ii) aggregating sales to the firm–*product* level at which tax rates are defined (HSN-4), even when firms report more narrowly; and (iii) aggregating all treated and all control products within each firm.

Finally, the results are not driven by a small number of products: product-level regression estimates for changes in consumer prices, quantities, and sales are centered around the average effects and are close to normally distributed (Figure A.6).

Elasticities by Luxury Intensity A concern for differentiated commodity taxation is that luxury goods may have more elastic tax bases than other products. The broad set of products affected by the reform allows us to test this by interacting the treatment effect with each product's Engel-curve steepness, our measure of luxury intensity (Appendix D). We find no evidence that tax-base elasticities increase with luxury intensity: the interaction coefficients are small, statistically insignificant, and opposite in sign to the hypothesis.

Second Reform: July 2018 Rate Cut Twenty-five products that remained in the 28 percent tax bracket after the November 2017 reform were cut to 18 percent in July 2018. Although the reform affected fewer products, it complements the earlier reform by occurring

well after GST implementation and targeting higher-ticket items.[19] We estimate its effects using the same specification as for the November 2017 reform, with a control group comprising products taxed at 12 or 18 percent throughout the sample period. Summary statistics are reported in Table A.13.

The July 2018 estimates show the same qualitative patterns as the November 2017 reform, although the implied elasticities are somewhat larger (Figure A.7). Consumer prices fall by 0.038 log points, implying 63 percent pass-through of the 0.061 log-point tax cut. Quantities increase by 0.057 log points and pre-tax sales by 0.080 log points, implying quantity and sales elasticities of -0.95 and -1.32, compared with -0.50 and -0.99 in the November 2017 reform. Overall, these results show that our findings generalize to a distinct set of higher-ticket goods and a later stage of GST implementation.

4.5 Firm-Level Results Covering All Registered Firms

Our firm–product design exploits within-firm variation to absorb time-varying firm-level shocks while matching the product-level policy variation. However, value added—the GST tax base—is reported only at the firm level, and product-level GST returns are filed only by medium and large firms. Consequently, the analysis excludes most small registered firms (74% of firms but only 26% of sales; Table A.14).

The GST also includes two firm-level data sources: (i) a consolidated monthly return reporting value added (GSTR-3B), and (ii) a registration form listing up to five main products. We use these data to classify firms as treated if they list at least one product affected by the November 2017 reform, and as controls otherwise. Because value added is observed only at the firm level, this design allows us to estimate the effects of tax changes on value added while including all registered firms, regardless of size. We first estimate the reduced-form (intent-to-treat) effects of reform exposure. We then instrument firm-level tax rate changes

[19] The reform affected electrical appliances (e.g., vacuum cleaners and refrigerators), home renovation items (e.g., paints and varnishes), and household goods (e.g., suitcases and mirrors) (see Table A.12). These products accounted for 2.6 percent of sales (see Table A.4).

with this binary exposure to estimate the causal effects of tax changes on firm outcomes.

Figure A.8 reports firm-level event studies. Treated firms see ETRs fall by 3 pp, versus 0.7 pp for controls, confirming that the registration form is informative (Panel A).[20] This implies a first-stage reduction in ETRs of 2.3 percentage points (Panel B). Relative to control firms, treated firms raise sales by 0.036 log points (Panel C) and value added by 0.054 log points (Panel D), implying elasticities with respect to ETR of -1.69 and -2.49, respectively.[21] The larger response of value added may reflect fixed input costs, although the difference is not statistically significant.

Comparing Firm-Product and Firm-Level Estimates The firm-level strategy yields larger elasticities than the baseline but covers all registered firms rather than only the larger firms filing detailed product-level returns. Table A.16 reveals a clear decline in sales elasticities with firm size: -2.61 among small firms filing simplified returns (Column 1), compared with -1.89 among firms filing detailed firm-level returns (Column 2) and -1.16 among firms filing product-level returns (Column 3).

Among firms filing product-level returns, the binary firm-level specification yields a sales elasticity of -1.16 and a value-added elasticity of -1.67, consistent with the baseline product-level sales elasticity of -0.99. Replacing the binary registration-based treatment with firms' actual pre-reform shares of sales in treated products, constructed from product-level tax returns, produces a stronger first stage and elasticities of -0.62 for sales and -1.40 for value added (Column 4). Because this measure captures the intensity of each firm's exposure to the reform rather than simply whether it sold any treated products, we view these estimates as our preferred firm-level specification.[22]

[20] At the firm level, we define STRs as the average statutory rate across listed products (e.g., one product taxed at 18% and another at 28% implies an STR of 23%). Before the November 2017 reform, treated firms averaged an STR of 23.1 percent, versus 13.7 percent for controls.

[21] We define value added as monthly sales minus intra-state purchases, with purchases inferred from firm-to-firm transactions. We do not observe inter-state input purchases. The analysis is restricted to firm-months with positive value added.

[22] The smaller elasticities suggest that the binary treatment overstates firms' responses by assigning the same treatment intensity to firms with different exposure to the reform.

Informality Our estimated tax-base elasticity captures all margins of adjustment that affect taxable sales, including real economic responses, substitution toward the informal sector, and tax evasion. We cannot, however, separately identify the contributions of these channels. The firm-level results show larger sales elasticities among smaller registered firms, consistent with greater exposure to informal competition and scope for evasion. By contrast, responses among medium to large firms remain moderate, and we find no evidence of extensive-margin adjustment through exit from the formal sector, despite longstanding concerns that higher tax rates may induce greater informality (e.g., Emran and Stiglitz 2005; Gordon and Li 2009). Informality nevertheless remains important for the distributional consequences of tax reform and is explicitly incorporated into the optimal-tax and welfare analysis in Section 6.

5 Scanner-Level Data Analysis

Our analysis faces three main limitations: a short pre-treatment window close to the GST's introduction and coinciding with India's main festive season; imperfect restriction to final-consumer sales; and potentially coarse product definitions at the HSN-4 level. To address these issues, we complement the tax data with a proprietary barcode-level dataset covering all transactions of a large hypermarket chain in Karnataka. Although limited to a single firm, this dataset overcomes these limitations: it provides a long pre-treatment period;[23] transactions are to final consumers; and barcodes identify identical products over time (same quality, volume, and brand).

We replicate the main analysis, adapting sample construction as needed. Because HSN codes are unavailable, treatment is defined using observed effective tax rates. Treated barcodes are those with an effective tax rate of 28 percent between August and October 2017 and 18 percent between December 2017 and March 2019; control barcodes have effective tax rates of 12 percent or 18 percent throughout. We restrict the sample to barcodes sold in all periods,

[23]The dataset runs from April 2016 to July 2019. We begin the analysis in July 2016, after the system's introduction phase, and end in March 2019 for comparability with the main analysis.

as many appear intermittently because of seasonality, product updates, or discontinuation.[24] To account for seasonality, we include department-month fixed effects. Identification therefore compares treated and control barcodes within store departments. The specification otherwise mirrors the firm–product analysis, with department–month fixed effects replacing firm–month fixed effects and barcode fixed effects replacing HSN fixed effects. We estimate the following regression:

$$y_{bdt} = \sum_{t=\text{Jul16},\, t\neq\text{Oct17}}^{\text{Mar19}} \beta_t \cdot (\text{treat}_b \times \text{month}_t) + \alpha_b + \gamma_{td} + \epsilon_{bdt} \tag{3}$$

where y_{bdt} denotes the outcome for barcode b, of store department d, in month t. The coefficients β_t capture the dynamic effects of the reform, with the pre-reform month of October 2017 omitted as the reference period ($\beta_{\text{Oct17}} = 0$). The model includes barcode fixed effects (α_b) and store department–month fixed effects (γ_{td}). Standard errors are clustered at the level of the hypermarket classification variable used to identify treated and control barcodes (946 clusters).

Table A.17 reports the number of treated and control barcodes by store department. In total, there are 1,508 treated barcodes across 10 departments, most of which belong to *Personal Care & Cosmetics*, followed by *Home Care* and *Breakfast*. The control group includes 2,491 barcodes, more evenly distributed across the 10 departments.[25]

Figure 8A shows the event-study estimates from equation 3 for consumer prices.[26] Treated and control barcodes exhibit parallel price trends throughout the long pre-period from July 2016 to November 2017, with no visible break at the GST's introduction in July 2017. Following the rate cut, consumer prices fall by 0.024 log points after the transition period,

[24] Barcodes must appear in at least five transactions during a month to be included.

[25] Table A.18 reports the number of barcodes at the hypermarket *Family* level, analogous to the HSN-4 classification in the tax data (Table A.6).

[26] Before the GST's introduction in July 2017, only state VAT is observed, as central duties (Central Value Added Tax, CENVAT) were embedded in retailers' pre-tax prices and remitted upstream. As a result, pre-tax prices and effective tax rates cannot be recovered for the full set of taxes later unified under the GST. This limitation does not affect quantities or consumer prices, which include all tax components and are invariant to tax reporting.

while effective tax rates fall by 0.081 log points, implying 30 percent pass-through.[27] The estimates are precise enough to reject both zero and full pass-through at the 1 percent level. The hypermarket ranks among the largest one percent of firms in the administrative data, and its lower estimated pass-through is consistent with the sales-weighted administrative estimates showing lower pass-through among larger firms (Figure 7).

Consumer price adjustment following the reform is non-linear: prices initially fall, then rise, and fall again. This pattern reflects the hypermarket's strategic use of temporary discounts during the transition. Consumer prices incorporate promotional discounts, and Figure 8B plots their contribution to price changes. The initial price decline is entirely explained by deeper discounts on treated products, indicating that short-run pass-through occurs through temporary promotions rather than permanent price cuts. After two months, discounts unwind and consumer prices rise, allowing the retailer to retain more of the tax cut, before declining again as lower prices become embedded. This behavior is consistent with greater pricing sophistication—and potentially greater market power—than the average firm in the tax data.

Figure 8C reports the quantity results. Pre-trends are stable in the year preceding the rate cut. Following the reform, quantities sold of treated barcodes slightly rise, with a point estimate of 0.042 log points, significant at the 10 percent level. Given the size and formality of the hypermarket, quantity responses likely reflect genuine increases in sales rather than newly reported informal activity. This quantity response is slightly larger, though not statistically different from that in the tax data. Combined with the lower price pass-through, it however implies a higher price elasticity of -1.75.

Overall, the hypermarket analysis, based on an independent data source, is qualitatively consistent with our main findings.

[27] In the hypermarket data, rate cuts are defined using the retailer's self-reported tax brackets. The chain's size, centralized structure, and standardized reporting make tax reporting highly reliable.

6 Optimal Commodity Rate Differentiation

This section examines the implications of our empirical results for optimal commodity taxation. We first analyze equity and efficiency separately, then combine them to derive the optimal rate differentiation between luxury and standard goods.

6.1 Equity

GST Progressivity at Baseline Matching the household expenditure survey to commodity-level GST rates, we find that the average effective consumption tax rate increased monotonically across the income distribution before the reform, from 5.9 percent in the bottom decile to 8.4 percent in the top (Figure A.9A). The near-elimination of the 28 percent tax bracket through the November 2017 and July 2018 rate cuts was mechanically regressive, reducing the average tax rate by 0.7 percentage points for the bottom decile but 1.3 percentage points for the top (Figure A.9B).

Retargeting the Top Tax Bracket Despite the relative progressivity of the initial GST rate assignment, substantial scope remains for better targeting. To illustrate, we construct an alternative schedule based on goods' Engel-curve slopes. Before the reform, products in the 28 percent bracket accounted for 11 percent of goods consumption (seven percent of total consumption, including services). Holding the size of the top bracket fixed at 11 percent of goods consumption, we reassign products to the top bracket in descending order of their Engel-curve slopes. Products removed from the top bracket receive the consumption-weighted average tax rate of the newly included products, leaving the average tax rate unchanged.

Under the original assignment, top-decile households devote only 1.8 times more of their expenditure to 28 percent-rated goods than bottom-decile households. Under the progressive reassignment, this ratio rises to 28 (Figure A.10A). Holding the top bracket's share fixed at 11 percent of goods consumption, the reassignment raises the top decile's average tax rate by 1.4 percentage points while lowering it for all households below the eighth decile (Figure A.10B).

6.2 Efficiency: Tax Revenues and the MCPF

Elasticities for the Calibration To assess the efficiency implications of the reform, we first describe how the empirical elasticities are incorporated into the calibration. We primarily use the elasticities estimated for products matched to the household survey, ensuring consistency with the set of products used in the calibration. This choice is also conservative, as survey-matched products exhibit more elastic quantity and sales responses than the full sample. We also consider value-added elasticities in an extension. Although value added is the relevant VAT tax base, our preferred empirical design identifies responses at the firm-product level, whereas value added is observed only at the firm level, making the corresponding estimates less precise.

For the calibration, we treat the estimated elasticities as compensated. Empirically, adjusting for the income effect changes the estimated elasticity by only 0.06 in absolute value. This small difference reflects the limited budget share of treated goods and incomplete pass-through to consumer prices. Moreover, because treated and control products have similar income elasticities, reform-induced income effects should generate similar demand shifts across product groups within firms. These common shifts are largely absorbed by the firm-time fixed effects, so the remaining variation primarily reflects substitution responses. We apply the correction whenever uncompensated elasticities are required.[28]

We use elasticities with respect to effective tax rates and model the initial compliance gap between effective and statutory rates in an extension. This approach generalizes more readily to settings with different levels of compliance and yields more conservative estimates than using elasticities with respect to statutory tax changes. We remain agnostic about the extent to which the sales elasticity reflects evasion, since both real and evasion responses affect the tax base. Any evasion component in the quantity response would instead imply lower welfare costs. Appendix E provides the formal derivations underlying these mappings.

[28] By the Slutsky equation, $\varepsilon^u s = \varepsilon^c s - \varepsilon_p \sum_i a^i \theta^i \eta$, where a^i is household i's share of treated-goods consumption, θ^i is its budget share of treated goods, and η is their income elasticity.

Tax Revenues The revenue impact of a rate change depends on the elasticity of the tax base, since behavioral responses may offset part of the mechanical revenue loss:

$$dR = \left(1 + \frac{t}{1+t}\varepsilon_s\right) dM,$$

where dR is total revenue change, dM the mechanical change holding the tax base fixed, t the initial tax rate, and ε_s the uncompensated sales elasticity. Using the full-sample elasticity, each rupee of mechanical revenue loss translates into an actual loss of 0.77 rupees, implying that behavioral responses offset 0.23 rupees. Using the larger elasticity for survey-matched goods, the corresponding loss falls to 0.63 rupees. Thus, the rate cut generated a substantial revenue loss, only partly offset by higher reported sales.

Marginal Cost of Public Funds (MCPF) The MCPF, defined as the total private cost per dollar of revenue raised, equals one plus the marginal excess burden per dollar of marginal revenue, and can be written as (Dahlby 2008):

$$\text{MCPF} = 1 + \frac{-\frac{t}{1+t}\,\varepsilon_x}{1 + \frac{t}{1+t}\,\varepsilon_s},$$

where ε_x is the compensated quantity elasticity governing the marginal excess burden, and ε_s the uncompensated tax-base elasticity governing marginal revenue. Evaluated at the pre-reform tax rate, the MCPF is 1.14 using the full-sample elasticities and 1.40 using survey-matched goods. Thus, raising one rupee through indirect taxation costs the private sector 1.14–1.40 rupees. These estimates are upper bounds if part of the observed quantity response reflects changes in evasion rather than consumption.

6.3 Optimal Luxury Taxation

Optimal Tax Rule We derive a general expression for optimal indirect tax rates using the many-person Ramsey rule (Diamond 1975). We consider a continuous mass of households

$i \in [0,1]$, with exogenous income y^i. Producer prices are denoted by q_j, and ad-valorem taxes t_j yield consumer prices $p_j = (1+t_j)q_j$. The indirect utility function of household i is given by $v(p, y^i)$. The government maximizes a social welfare function $G(\cdot)$, subject to a revenue requirement. The Lagrange multiplier μ represents the marginal value of public funds (MVPF); the social cost of raising one additional unit of government revenue. We assume zero cross-price elasticities and do not account for producer surplus, both of which we address in an extension. The baseline model and its extensions are developed in Appendix F.[29]

Under these assumptions, the optimal indirect tax on good j is:

$$\frac{t_j}{1+t_j} = \frac{1 - \varepsilon_{p,j} \int_i \frac{g^i}{\mu} a_j^i \, di}{-\varepsilon_{s,j}} \tag{4}$$

where g^i denotes social marginal welfare weights and $a_j^i = \frac{x_j^i}{x_j}$, household's i share of total consumption of good j, with x_j^i the consumption of good j by household i and $x_j = \int_i x_j^i \, di$ the aggregate consumption of good j. The term $\varepsilon_{p,j}$ is the elasticity of consumer prices with respect to the tax and $\varepsilon_{s,j}$ the uncompensated elasticity of pre-tax sales with respect to the tax (i.e., the tax base elasticity).

The optimal tax on good j is higher when (i) demand is less elastic ($\varepsilon_{s,j}$ is closer to zero); (ii) households with lower welfare weights consume relatively more of the good; (iii) tax pass-through is lower, so consumers bear a smaller share of the burden; and (iv) the MVPF is higher.

The equity–efficiency trade-off arises along two dimensions. First, the government trades off tax-base distortions ($\varepsilon_{s,j}$) against revenue needs, captured by the MVPF (μ). This determines the overall level of indirect taxation, even if goods share identical elasticities and social values. Second, the relationship between a good's elasticity and social weight also matters. If goods with higher social weights are more elastic, efficiency and equity pull in opposite directions, making the case for differentiated rates an empirical question.

[29] The setup assumes fixed labor supply, so the efficiency cost we measure reflects distortions in consumption choices across goods rather than labor-leisure distortions.

Income Taxes and Transfers As in many low- and middle-income countries, India's personal income tax (PIT) applied to only a small share of the population. In 2017, 4 percent of the adult population paid PIT (Ministry of Finance, Government of India 2019). The statutory schedule for fiscal year 2017–2018 was progressive above an exemption threshold of INR 250,000 (USD 3,900), with brackets at INR 500,000 and INR 1 million and marginal rates of 5, 20, and 30 percent. Taxable income was further reduced by deductions and exemptions, effectively raising the income level at which households began paying PIT. To capture these provisions, we calibrate an effective exemption threshold so that the share of taxpaying households in the survey equals the observed 4 percent coverage rate.

We consider three schemes for redistributing PIT revenues.[30] The baseline redistributes revenues as universal lump-sum transfers. At the other extreme, a precisely targeted scheme allocates all transfers to the bottom 50 percent of the income distribution. Finally, we consider a more policy-relevant scenario based on the National Rural Employment Guarantee Act (NREGA), one of India's largest social-protection programs. We assume 75 percent of transfers accrue to the bottom half of households and 25 percent to the top half (Dutta et al. 2014; Imbert and Papp 2015). Accounting for 10 percent administrative costs and 17 percent leakage leaves net transfers equal to 73 percent of PIT revenues (Banerjee et al. 2020; Muralidharan et al. 2016).

Calibration We observe consumption shares, a_j^i, in the household survey, and take elasticities of quantity demanded from the empirical estimates for survey products.[31] We define social marginal welfare weights, g^i, as the inverse of household per capita post-tax-and-transfer expenditure, our proxy for disposable income, corresponding to log-utility preferences (CRRA coefficient equal to one). To discipline the model, we calibrate μ, the marginal value of public funds, so that the optimal tax rate for goods initially taxed at 18 percent equals 18 percent.

[30] The baseline PIT schedule applied to the survey yields less revenue than 2017 PIT collections. We therefore scale PIT revenues to match aggregate collections and avoid understating the PIT's role.

[31] As quantity and sales elasticities are unrelated to product luxuriousness (Appendix D), we use the same elasticities for all products.

This normalization allows us to focus on the luxury tax premium—the difference between the optimal tax rate on luxury goods and the standard rate—rather than the overall level of indirect taxation.

We treat PIT schedule changes as exogenous and model PIT coverage expansions by proportionally lowering all tax brackets, including the exemption threshold. PIT and its redistribution alter consumption shares, a_j^i, and social marginal welfare weights, g^i. We capture these effects by estimating flexible Engel curves by decile, which recover how demand for each good varies with income (Appendix G).

Results Figure 9 shows the resulting optimal tax rates for each household-survey product, highlighting selected major products along the income-elasticity distribution. Optimal rates range from close to zero for products such as potatoes to roughly 60 percent for gold ornaments. For the discussion that follows, we refer to the products with the steepest Engel curves accounting for 11 percent of goods consumption—matching the expenditure share of goods initially taxed at 28 percent—as luxury products. These products have income elasticities above 1.6, averaging 2.2.

Figure 10A shows the optimal luxury-tax premium: the difference between the luxury and standard rates (calibrated at 18%). Without a personal income tax (PIT), the optimal luxury-tax premium is 26 pp. Under India's current PIT, it falls only slightly to 24 pp, reflecting the tax's limited coverage and progressivity. We next relax constraints on PIT rates and coverage separately. Raising all PIT rates by 20 pp (increasing the top marginal rate to 50%), while holding the tax base fixed, lowers the premium to 21 pp. Expanding PIT coverage to the top income quintile, while holding rates fixed, also lowers the premium to 21 pp. Combining both reforms reduces the premium to 15 pp. Thus, even sizable expansions in PIT progressivity and coverage reduce the optimal premium by just over one-third, leaving substantial scope for luxury taxation.

We next add targeted transfers. Combining both PIT reforms with NREGA-like transfers

further reduces the premium to 13 pp. Only when the expanded PIT finances transfers that are perfectly targeted to the bottom 50 percent, with no leakage, does the premium fall sharply, to 1 pp.[32] These results provide an empirical counterpart to the Atkinson–Stiglitz benchmark. Uniform commodity taxation becomes approximately optimal only once the conditions underlying the theorem—a broad, progressive income tax combined with well-targeted transfers—are satisfied. Under India's current institutions, these conditions are far from met, leaving substantial scope for luxury taxation.[33]

How General Is the Case for Luxury Taxation? In Figure 10B, we assess the robustness of the calibrated 24 percentage point luxury-tax premium under the current PIT system to a range of alternative assumptions, including tax-base elasticities, the incidence of tax increases, producer surplus, survey undercoverage of top incomes, informal consumption, and welfare weights.

First, we consider alternative elasticity estimates. Using firm-sales-weighted elasticities, which are closer to aggregate estimates, raises the optimal premium to 34 pp, while using the value-added elasticity reduces it modestly, to 20 pp.

Second, we assess the robustness of extrapolating from the observed GST rate cut, which compressed tax-rate differences across products, to a counterfactual reform that (i) raises luxury tax rates and (ii) increases tax-rate differentiation. Because tax increases often exhibit higher pass-through than tax cuts (Benzarti et al. 2020), we assume conservatively full pass-through. This reduces the optimal luxury-tax premium to 17 pp, as the efficiency cost of the reform is enhanced.[34] Greater rate differentiation may also create stronger incentives for

[32] Precisely targeting transfers to the bottom 50 percent, without expanding the PIT base and rates, reduces the optimal tax premium only to 19 pp, as revenue under the current PIT remains limited.

[33] Figure A.11 reinforces this point by focusing on the PIT exemption threshold, a key distinction between tax systems in high-income and developing countries (Jensen 2022). Even with a 50 percent top PIT rate and NREGA-like transfers financed by the additional revenue, the luxury-tax premium declines gradually, reaching negligible levels only when the PIT extends to roughly the top 40 percent of the income distribution, well beyond its coverage in most developing countries.

[34] Higher pass-through amplifies the tax-base response by increasing the consumer price response. This follows from the identity $\varepsilon_s = \varepsilon_p(1 + \varepsilon_{xp}) - 1$, where ε_{xp} is the price elasticity of demand.

product relabeling and evasion than rate compression.[35] To capture this possibility, we extend the model in Appendix F to allow for evasion, calibrating it so that relabeling of high-tax goods into lower-tax categories reproduces the observed pre-reform gap between statutory and effective tax rates, and the resulting tax under-payment accrues fully to firms. Under this assumption, the optimal luxury-tax premium falls to 17 pp. Combining full pass-through with the assumption that the entire statutory–effective tax rate gap reflects evasion lowers the premium to 12 pp. Because both assumptions are deliberately conservative, this estimate should be interpreted as a lower bound.

Third, we consider extensions incorporating producer surplus, survey under-coverage of top incomes, and informal consumption (see Appendices F and G). We first extend the model to account for firm profits, attributing them in a simplified way to the top 1 percent. This reduces the optimal luxury-tax premium to 17 pp, as higher luxury taxes lower firm profits, which now enter the social welfare calculation. We next correct top incomes for the survey's understatement of inequality (Bharti et al. 2026). This increases the premium to 32 pp by increasing the concentration of consumption among high-income households. Finally, we allow for informal consumption by assuming that goods taxed at 18 percent are consumed informally in the proportions estimated by Bachas et al. (2024a), while luxury goods remain fully formal. This biases against luxury taxation because untaxed informal consumption makes standard-rate goods relatively more progressive. Even so, the optimal premium remains 21 pp. Combining all three extensions yields an optimal premium of 22 pp.

Finally, we examine the sensitivity of our results to the assumed degree of inequality aversion. Our baseline uses log-utility social welfare weights ($\gamma = 1$). As expected, lower inequality aversion ($\gamma = 0.5$) reduces the optimal luxury-tax premium to 13 percentage points, while greater inequality aversion ($\gamma = 2$) increases it to 42 percentage points. Despite these differences, the comparative statics are remarkably stable as show in Figure A.12:

[35] Under standard models, changes in rate differentiation affect incentives for relabeling and evasion symmetrically. However, adjustment frictions, enforcement dynamics, or salience effects may generate larger responses when rate differentials increase than when they decrease.

expanding the coverage, progressivity, and targeting of income taxes and transfers reduces the premium by similar proportions across values of γ. Thus, although the preferred level of luxury taxation depends on redistributive preferences, its role within the broader fiscal system remains largely unchanged.

Consumer Welfare We use a second-order approximation to indirect utility to estimate the expenditure-equivalent welfare effects of implementing optimal commodity taxation in a revenue-neutral manner. The approximation accounts for both tax-induced price changes and behavioral consumption responses (Chetty 2009):

$$\Delta W_{ij} \;=\; -\theta_j^i \, \varepsilon_p \left[\Delta \ln(1+t_j) + \tfrac{1}{2}\big(1+\varepsilon_s\big)\big(\Delta \ln(1+t_j)\big)^2 \right]. \tag{5}$$

Where θ_j^i denotes consumption by household i on good j, ε_p the pass-through rate, ε_s the sales elasticity, and $\Delta \ln(1+t_j)$ the change in the log tax rate.[36] Tax reductions increase welfare to the extent that they are passed through to consumer prices, with consumers' reoptimization contributing a second-order adjustment.

Figure A.13 illustrates the consumer-welfare consequences of revenue-neutral reforms to India's baseline commodity tax schedule. It compares three scenarios: unrestricted optimal commodity taxation (corresponding to the rates in Figure 9), optimal luxury taxation (as previously defined), and a simpler reclassification that assigns the goods with the steepest Engel curves to the 28% tax bracket. In our preferred specification which models behavioral responses (Panel B), optimal commodity taxation generates welfare gains for the bottom eight deciles while concentrating welfare losses on the top decile. The bottom four deciles gain more than 1 percent of expenditure, with gains reaching over 1.5 percent for the poorest households, while the richest households lose a similar amount. Restricting differentiation to luxury goods captures about two-thirds of the welfare gains under the fully optimal commodity-tax schedule. These gains are sizable given that the reform is revenue neutral, operates within

[36] The derivation is in Appendix F. An isoelastic-demand approximation yields nearly identical consumer-surplus estimates.

a tax system whose average effective consumption tax rate is only 7 percent, and improves upon an already differentiated rate schedule.

Panels A and C bracket our preferred specification. Relative to Panel B, Panel A reports welfare effects considering only tax-induced price changes, while Panel C additionally accounts for informal consumption. Together, they show that the qualitative conclusions are robust, although informality reduces the welfare gains.

7 Conclusion

Combining causal estimates from a large VAT reform in India with evidence on the distributional targeting of taxed goods, we quantify how expansions of income taxes and transfers reduce the optimal luxury-tax premium, identifying when uniform commodity taxation becomes optimal in practice. Our estimates imply that luxury taxation remains desirable under a wide range of plausible reforms and converges to zero only after large simultaneous expansions of both income taxes and transfers.

These findings help explain why differentiated taxation of luxury and non-essential consumption remains common across developing countries. As shown in Table A.19, several large emerging economies—including China, Brazil, Türkiye, and South Africa—maintain higher tax burdens on a broad set of non-essential goods through differentiated VAT rates, excises, and selective consumption taxes.[37]

The case for luxury taxation should nevertheless be weighed against several potential drawbacks. Differentiated VAT structures increase administrative complexity, generate classification disputes, and require periodic updating as consumption patterns evolve and formerly luxurious goods diffuse through the income distribution. They may also raise horizontal equity concerns when households with similar incomes consume different bundles

[37]The scope of differentiated commodity taxation varies: some countries tax a broad set of non-essential goods consumed disproportionately by higher-income households, while others focus on luxuries such as yachts, private aircraft, and high-end vehicles. Such taxes often pursue multiple objectives, including redistribution, externality correction, and industrial or trade policy.

(Saez 2002). Political economy cuts both ways: luxury taxation may be easier to enact than income tax reform, but differentiated rates are also more vulnerable to lobbying. Conversely, some luxury goods may generate status-based consumption externalities, providing an additional rationale for higher taxes beyond redistribution (Bursztyn et al. 2018).

These results also have implications for the evolution of the tax mix. The conventional view is that, as administrative capacity improves, redistribution should increasingly rely on direct taxes and transfers rather than commodity taxation. But this presumes direct tax bases become easier to tax than consumption. That need not be the case. Luxury consumption is often highly visible, whereas income, wealth, and profits may become increasingly mobile and easier to conceal as globalization, technological change, and rising wealth concentration reshape economies. In such environments, differentiated commodity taxation may remain an important complement to direct redistribution (Korinek and Lockwood 2026).

References

Agarwal, S., Ghosh, P., & Zhang, J. (2025). Tax policy transmission and household expenditures. *The Review of Economics and Statistics*, 1–46. https://doi.org/10.1162/rest_a_01584

Allcott, H., Lockwood, B. B., & Taubinsky, D. (2019). Regressive sin taxes, with an application to the optimal soda tax. *The Quarterly Journal of Economics*, *134*(3), 1557–1626. https://doi.org/10.1093/qje/qjz017

Allingham, M. G., & Sandmo, A. (1972). Income tax evasion: A theoretical analysis. *Journal of Public Economics*, *1*(3), 323–338. https://doi.org/https://doi.org/10.1016/0047-2727(72)90010-2

Almunia, M., Hjort, J., Knebelmann, J., & Tian, L. (2023). Strategic or confused firms? evidence from "missing" transactions in uganda. *The Review of Economics and Statistics*, *105*(5), 1119–1130. https://doi.org/10.1162/rest_a_01180

Antràs, P., Chor, D., Fally, T., & Hillberry, R. (2012). Measuring the upstreamness of production and trade flows. *American Economic Review*, *102*(3), 412–16. https://doi.org/10.1257/aer.102.3.412

Asatryan, Z., & Gomtsyan, D. (2023). *The incidence of vat enforcement: Evidence from agricultural markets* (tech. rep.) [SSRN Working Paper No. 4591073]. ZEW – Leibniz Centre for European Economic Research.

Atkinson, A., & Stiglitz, J. (1976). The design of tax structure: Direct versus indirect taxation. *Journal of Public Economics*, *6*(1), 55–75. https://doi.org/https://doi.org/10.1016/0047-2727(76)90041-4

Bachas, P., Gadenne, L., & Jensen, A. (2024a). Informality, Consumption Taxes, and Redistribution. *The Review of Economic Studies*, *91*(5), 2604–2634. https://doi.org/10.1093/restud/rdad095

Bachas, P., Jensen, A., & Gadenne, L. (2024b). Tax equity in low- and middle-income countries. *Journal of Economic Perspectives*, *38*(1), 55–80. https://doi.org/10.1257/jep.38.1.55

Banerjee, A., Duflo, E., Imbert, C., Mathew, S., & Pande, R. (2020). E-governance, accountability, and leakage in public programs: Experimental evidence from a financial management reform in india. *American Economic Journal: Applied Economics*, *12*(4), 39–72. https://doi.org/10.1257/app.20180302

Banerjee, A., Hanna, R., Olken, B. A., & Sverdlin Lisker, D. (2024). Social protection in the developing world. *Journal of Economic Literature*, *62*(4), 1349–1421. https://doi.org/10.1257/jel.20241646

Belan, P., & Gauthier, S. (2006). Optimal indirect taxation with a restricted number of tax rates. *Journal of Public Economics*, *90*(6-7), 1201–1213. https://doi.org/10.1016/j.jpubeco.2005.10.004

Benedek, D., De Mooij, R. A., Keen, M., & Wingender, P. (2020). Varieties of VAT pass through. *International Tax and Public Finance*, *27*(4), 890–930. https://doi.org/10.1007/s10797-019-09566-5

Benzarti, Y. (2024). *Tax incidence anomalies* (NBER Working Paper No. 32819). National Bureau of Economic Research. https://doi.org/10.3386/w32819

Benzarti, Y., Carloni, D., Harju, J., & Kosonen, T. (2020). What goes up may not come down: Asymmetric incidence of value-added taxes. *Journal of Political Economy*, *128*(12), 4438–4474. https://doi.org/10.1086/710558

Benzarti, Y., Garriga, S., & Tortarolo, D. (2024). *Can VAT cuts and anti-profiteering measures dampen the effects of food price inflation?* (NBER Working Paper No. 32241). National Bureau of Economic Research. https://doi.org/10.3386/w32241

Bergeron, A., Tourek, G., & Weigel, J. L. (2024). The state capacity ceiling on tax rates: Evidence from randomized tax abatements in the drc. *Econometrica*, *92*(4), 1163–1193. https://doi.org/https://doi.org/10.3982/ECTA19959

Best, M. C., Brockmeyer, A., Kleven, H. J., Spinnewijn, J., & Waseem, M. (2015). Production versus revenue efficiency with limited tax capacity: Theory and evidence from pakistan. *Journal of Political Economy*, *123*(6), 1311–1355. https://doi.org/10.1086/683849

Bharti, N. K., Chancel, L., Piketty, T., & Somanchi, A. (2026). Income and wealth inequality in india, 1922–2023: The rise of the billionaire raj. *The World Bank Economic Review*, lhag013. https://doi.org/10.1093/wber/lhag013

Boadway, R., Marchand, M., & Pestieau, P. (1994). Towards a theory of the direct-indirect tax mix. *Journal of Public Economics*, *55*(1), 71–88. https://doi.org/10.1016/0047-2727(94)90081-7

Brockmeyer, A., & Hernandez, M. (2022). *Taxation, information, and withholding: Evidence from costa rica* (Discussion Paper No. 17716). Centre for Economic Policy Research.

Brusco, G., & Velayudhan, T. (2025). Vat incidence in real vat systems. *Journal of Public Economics*, *249*, 105447. https://doi.org/https://doi.org/10.1016/j.jpubeco.2025.105447

Bursztyn, L., Ferman, B., Fiorin, S., Kanz, M., & Rao, G. (2018). Status goods: Experimental evidence from platinum credit cards. *The Quarterly Journal of Economics*, *133*(3), 1561–1595. https://doi.org/10.1093/qje/qjy006

Cawley, J., Frisvold, D., Hill, A., & Jones, D. (2019). The impact of the philadelphia beverage tax on purchases and consumption by adults and children. *Journal of Health Economics*, *67*, 102225.

Central Board of Indirect Taxes and Customs. (2019). Circular no. 113/32/2019-gst: Clarification regarding gst rates and classification of certain goods [Government of India, Ministry of Finance].

Chancel, L., Gómez-Carrera, R., Moshrif, R., & Piketty, T. (Eds.). (2025). *World inequality report 2026* [Accessed: 2025-12-16]. World Inequality Lab.

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, *1*(1), 451–488. https://doi.org/10.1146/annurev.economics.050708.142910

Colchero, M. A., Popkin, B. M., Rivera, J. A., & Ng, S. W. (2016). Beverage purchases from stores in mexico under the excise tax on sugar sweetened beverages: Observational study. *bmj*, *352*.

Cotti, C., Courtemanche, C., Maclean, J. C., Nesson, E., Pesko, M. F., & Tefft, N. W. (2022). The effects of e-cigarette taxes on e-cigarette prices and tobacco product sales: Evidence from retail panel data. *Journal of Health Economics*, *86*, 102676.

Crawford, I., Keen, M., & Smith, S. (2010). Value added tax and excises. In J. Mirrlees, S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles, & J. Poterba (Eds.), *Dimensions of tax design: The Mirrlees review* (pp. 275–422). Oxford University Press.

Dahlby, B. (2008). *The marginal cost of public funds: Theory and applications* (1st ed., Vol. 1). The MIT Press.

Deaton, A. (1977). Equity, efficiency and the structure of indirect taxation. *Journal of Public Economics*, *8*(3), 299–312. https://doi.org/10.1016/0047-2727(77)90032-2

Deaton, A. (2005). Measuring poverty in a growing world (or measuring growth in a poor world). *Review of Economics and Statistics*, *87*(1), 1–19.

Decker, S. L., & Schwartz, A. E. (2000). *Cigarettes and alcohol: Substitutes or complements?* (Working Paper No. 7535). National Bureau of Economic Research. https://doi.org/10.3386/w7535

Diamond, P. A. (1975). A Many-Person Ramsey Tax Rule. *Journal of Public Economics*, *4*(4), 335–342.

Dubois, P., Griffith, R., & O'Connell, M. (2020). How well targeted are soda taxes? *American Economic Review*, *110*(11), 3661–3704. https://doi.org/10.1257/aer.20171898

Dutta, P., Murgai, R., Ravallion, M., & van de Walle, D. (2014). *Right to work? assessing india's employment guarantee scheme in bihar*. World Bank. https://doi.org/10.1596/978-1-4648-0130-3

Emran, S., & Stiglitz, J. (2005). On selective indirect tax reform in developing countries. *Journal of Public Economics*, *89*(4), 599–623. https://doi.org/10.1016/j.jpubeco.2004.04.007

Gaarder, I. (2019). Incidence and distributional effects of value added taxes. *The Economic Journal*. https://doi.org/https://doi.org/10.1111/ecoj.12576

Gadenne, L., Nandi, T. K., & Rathelot, R. (forthcoming). Taxation and supplier networks: Evidence from india [Forthcoming]. *American Economic Journal: Applied Economics*.

Gadenne, L., Norris, S., Singhal, M., & Sukhtankar, S. (2024). In-kind transfers as insurance. *American Economic Review*, *114*(9), 2861–97. https://doi.org/10.1257/aer.20220822

Garriga, P., & Tortarolo, D. (2024). Firms as tax collectors. *Journal of Public Economics*, *233*, 105092. https://doi.org/https://doi.org/10.1016/j.jpubeco.2024.105092

Gordon, R., & Li, W. (2009). Tax structures in developing countries: Many puzzles and a possible explanation. *Journal of Public Economics*, *93*(7), 855–866. https://doi.org/https://doi.org/10.1016/j.jpubeco.2009.04.001

Government of India. (2017). The central goods and services tax act, 2017 [Act No. 12 of 2017, Section 171].

GST Council. (2018). Detailed agenda note, 30th GST Council meeting [New Delhi, 28 September 2018].

Gupta, B., Velayudhan, T., Vohra, V., & Garg, A. (2026). *Enhancing revenue collection capacity through automation: Evidence from a vat* [SSRN Working Paper]. https://doi.org/10.2139/ssrn.6433741

Hanna, R., & Olken, B. A. (2018). Universal basic incomes versus targeted transfers: Anti-Poverty programs in developing countries. *Journal of Economic Perspectives*, *32*(4), 201–226. https://doi.org/10.1257/jep.32.4.201

Harju, J., Kosonen, T., & Skans, O. N. (2018). Firm types, price-setting strategies, and consumption-tax incidence. *Journal of Public Economics*, *165*, 48–72. https://doi.org/10.1016/j.jpubeco.2018.06.006

Hoenink, J. C., Waterlander, W. E., Mackenbach, J. D., Mhurchu, C. N., Wilson, N., Beulens, J. W., & Nghiem, N. (2021). Impact of taxes on purchases of close substitute foods: Analysis of cross-price elasticities using data from a randomized experiment. *Nutrition Journal*, *20*(1), 75.

Huang, J., & Rios, J. (2016). Optimal tax mix with income tax non-compliance. *Journal of Public Economics*, *144*, 52–63. https://doi.org/https://doi.org/10.1016/j.jpubeco.2016.10.001

Imbert, C., & Papp, J. (2015). Labor market effects of social programs: Evidence from india's employment guarantee. *American Economic Journal: Applied Economics*, *7*(2), 233–263. https://doi.org/10.1257/app.20130401

International Monetary Fund & World Bank. (2018). Tax policy assessment framework (TPAF): Value added tax module.

Jensen, A. (2022). Employment structure and the rise of the modern tax system. *American Economic Review*, *112*(1), 213–34. https://doi.org/10.1257/aer.20191528

Korinek, A., & Lockwood, L. (2026). *Public finance in the age of ai: A primer* (Working Paper No. 34873). National Bureau of Economic Research. https://doi.org/10.3386/w34873

Liu, L., Lockwood, B., Almunia, M., & Tam, E. H. F. (2021). VAT Notches, Voluntary Registration, and Bunching: Theory and U.K. Evidence. *The Review of Economics and Statistics*, *103*(1), 151–164. https://doi.org/10.1162/rest_a_00884

Ministry of Finance, Government of India. (2019). Income Tax Return Statistics: Assessment Year 2018–19 [Version 1.1].

Ministry of Panchayati Raj & Indian Institute of Public Administration. (2024). *Status of devolution to panchayats in states: An indicative evidence based ranking* (Report) [Press Information Bureau]. Government of India.

Mishra, H., & Phadke, A. (2020). Evaluating the constitutionality of a deficient transition: The anti-profiteering law in India. *NLIU Law Review*, *10*(1), 237–269.

Muralidharan, K., Niehaus, P., & Sukhtankar, S. (2016). Building state capacity: Evidence from biometric smartcards in india. *American Economic Review*, *106*(10), 2895–2929. https://doi.org/10.1257/aer.20141346

Naritomi, J. (2019). Consumers as tax auditors. *American Economic Review*, *109*(9), 3031–72. https://doi.org/10.1257/aer.20160658

Niehaus, P., Atanassova, A., Bertrand, M., & Mullainathan, S. (2013). Targeting with agents. *American Economic Journal: Applied Economics*, *5*(1), 206–238. https://doi.org/10.1257/app.5.1.206

Okunogbe, O., & Pouliquen, V. (2022). Technology, taxation, and corruption: Evidence from the introduction of electronic tax filing. *American Economic Journal: Economic Policy*, *14*(1), 341–72. https://doi.org/10.1257/pol.20200123

Pomeranz, D. (2015). No taxation without information: Deterrence and self-enforcement in the value added tax. *American Economic Review*, *105*(8), 2539–69. https://doi.org/10.1257/aer.20130393

PRS Legislative Research. (2024). *State of State Finances* (tech. rep.) [FY 2024–25]. PRS Legislative Research.

Ramsey, F. P. (1927). A Contribution to the Theory of Taxation. *The Economic Journal*, *37*(145), 47–61. https://doi.org/10.2307/2222721

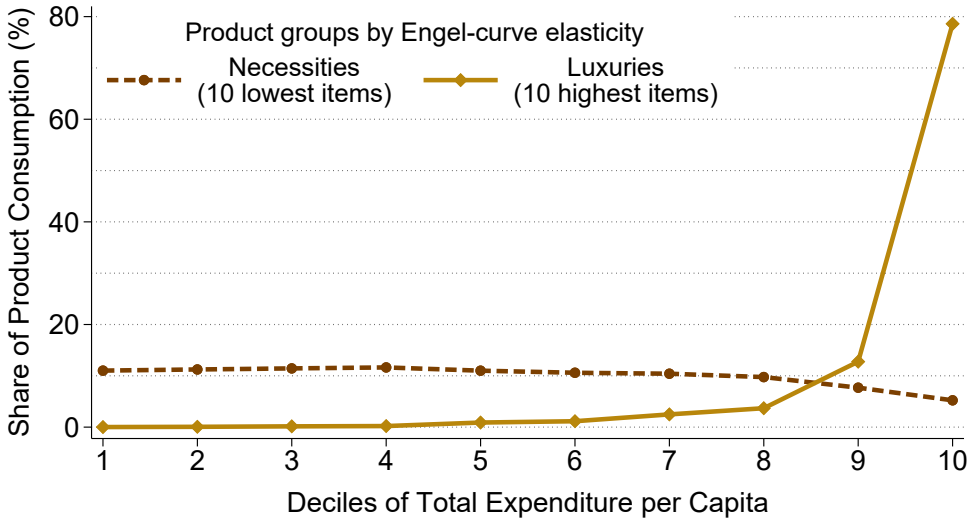
Saez, E. (2002). The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes. *Journal of Public Economics*, *83*(2), 217–230.

Waseem, M. (2022). The Role of Withholding in the Self-Enforcement of a Value-Added Tax: Evidence from Pakistan. *The Review of Economics and Statistics*, *104*(2), 336–354. https://doi.org/10.1162/rest_a_00959

Waseem, M. (2023). Overclaimed refunds, undeclared sales, and invoice mills: Nature and extent of noncompliance in a value-added tax. *Journal of Public Economics*, *218*, 104783. https://doi.org/https://doi.org/10.1016/j.jpubeco.2022.104783

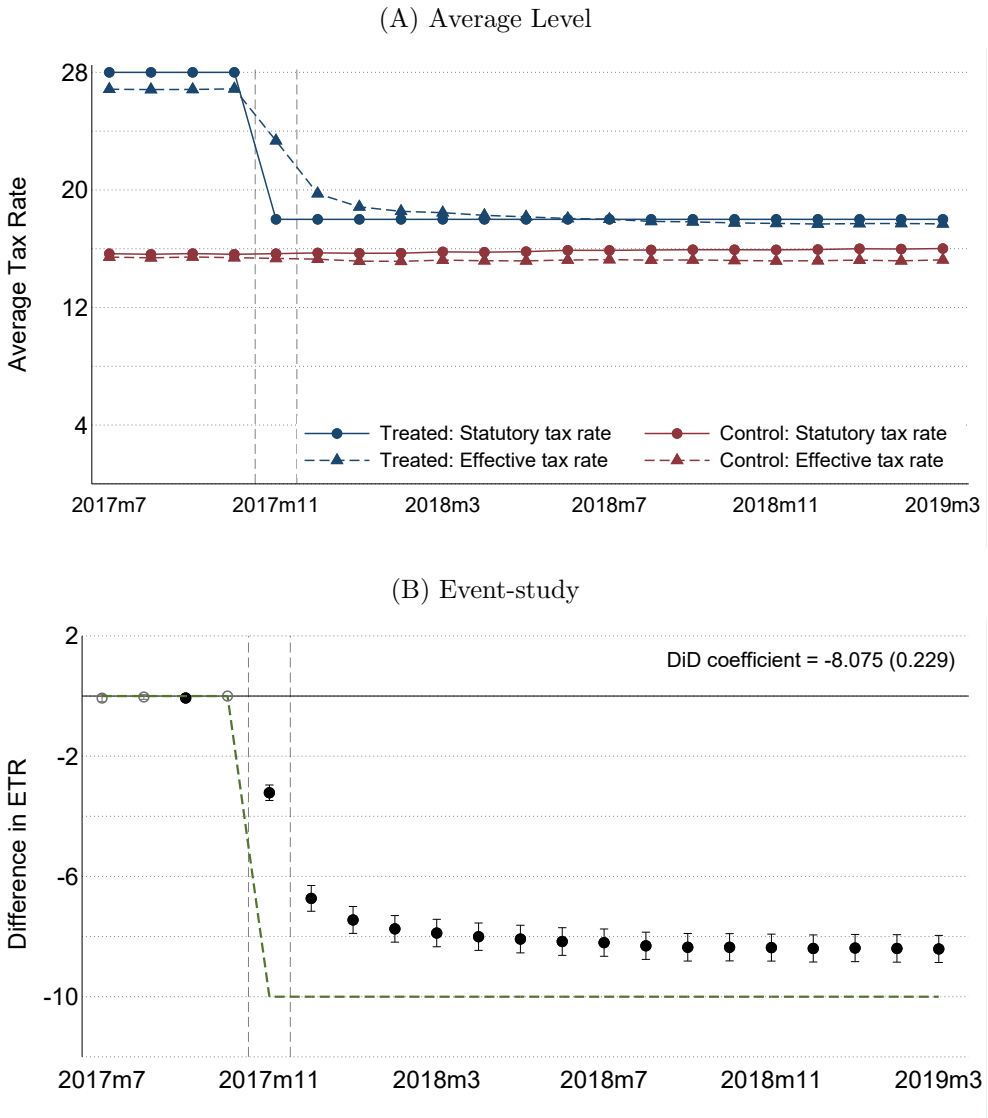
Figures and Tables

Figure 1: Consumption Patterns of Necessities and Luxury Goods



Note: This figure shows the share of consumption of selected goods across deciles of household total expenditure per capita using the 2011-12 household expenditure survey. The luxury group (gold solid line) consists of the 10 goods with the highest Engel curve elasticities among all survey goods, and the necessity group (brown dashed line) the 10 goods with the lowest. Each of these product groups account for about 1.5% of total consumption. This figure is discussed in Section 1.

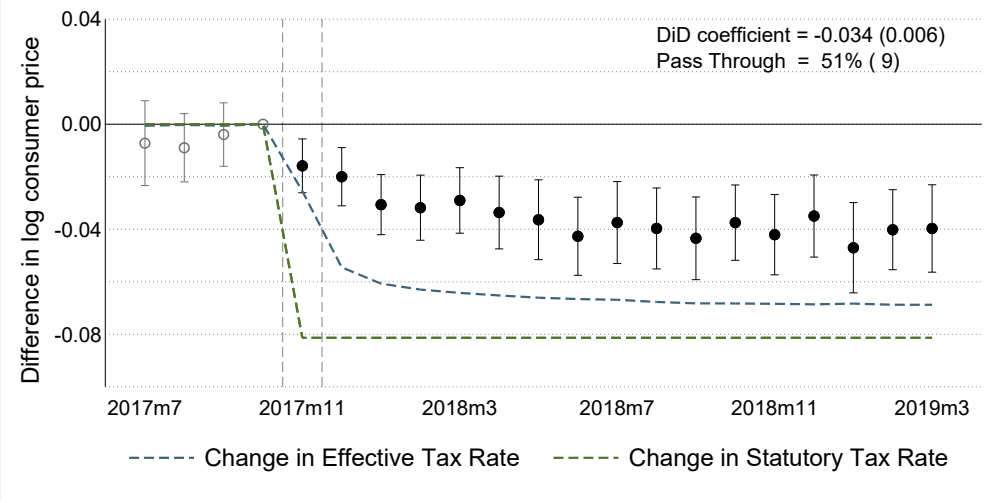
Figure 2: Statutory & Effective Tax Rates Around the Nov. 2017 Reform



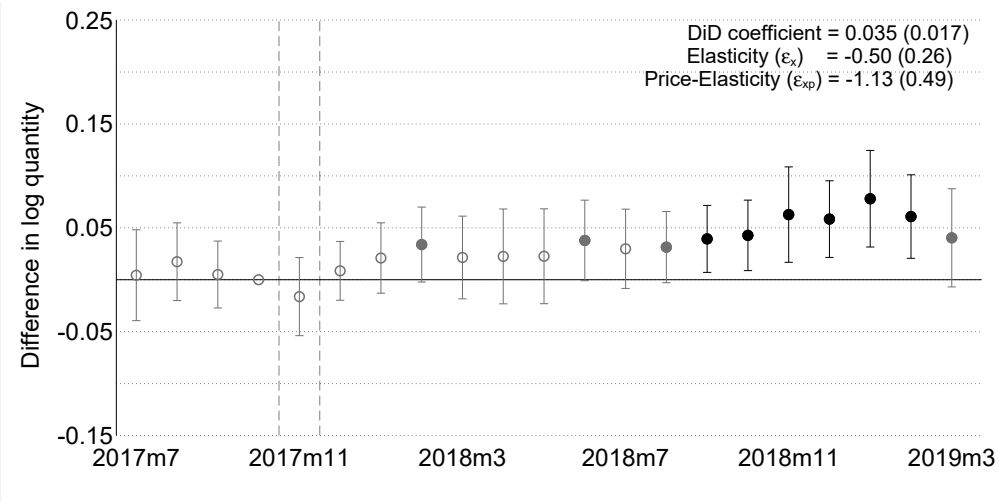
Note: Panel A presents the average Statutory Tax Rate (STR) and the average Effective Tax Rate (ETR) by treatment group. Treated HSNs are those that experienced a reduction in STR from 28% to 18%. Control HSNs are those taxed at 12% or 18% that did not experience a rate change during the period. Panel B presents the event-study coefficients for changes in ETRs. The dashed green line represents the change in STR. ETRs are winsorized at the 1st and 99th percentiles. Hollow markers denote coefficients that are not statistically significant at the 10% level; grey markers denote significance at the 10% level and black markers at the 5% level. This figure is discussed in Section 4.1.

Figure 3: Effects of the November 2017 GST Rate Cut: Consumer Prices & Quantities

(A) Consumer Price and Tax Pass-Through



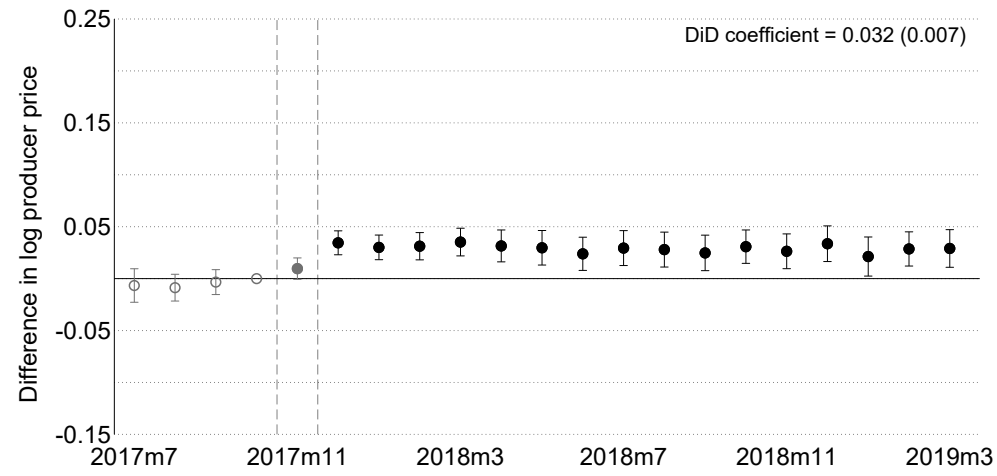
(B) Quantities



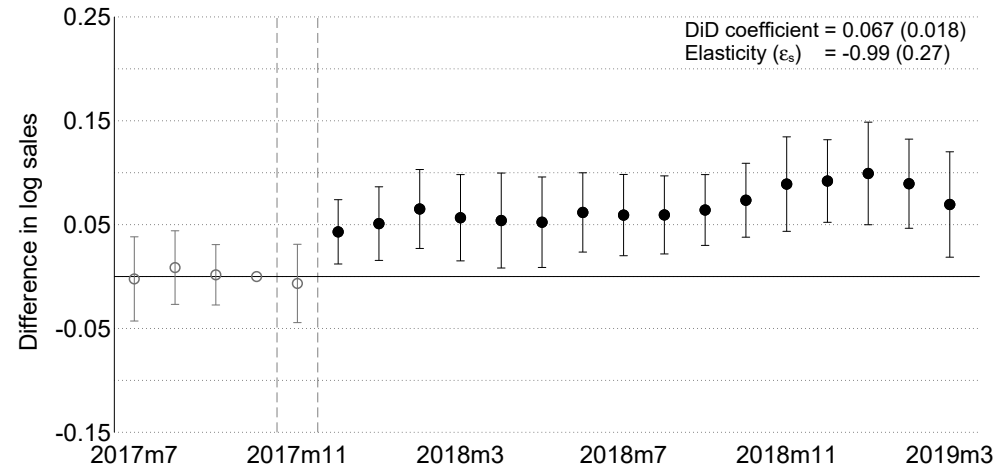
Note: Panel A presents the event-study coefficients for log consumer price. Treated HSNs are those that experienced a reduction in Statutory Tax Rates (STRs) from 28% to 18%. Control HSNs are those taxed at 12% or 18% that did not experience a rate change during the period. The green and the blue dashed lines indicate respectively the difference in log STR and log Effective Tax Rates (ETR) between the two groups. Pass-through is defined as the ratio between the reduction in consumer price and the reduction in ETR. Panel B presents the event-study coefficients for log quantity. The *DiD coefficient* is the difference-in-differences estimate for the months after March 2018. The *Elasticity* (ε_x) is the elasticity of quantity to ETR, instrumented with the treatment group and post periods. Hollow markers denote coefficients that are not statistically significant at the 10% level; grey markers denote significance at the 10% level and black markers at the 5% level. This figure is discussed in Section 4.1.

Figure 4: Effects of the November 2017 GST Rate Cut: Pre-Tax Price and Sales

(A) Pre-Tax Price (producer price)



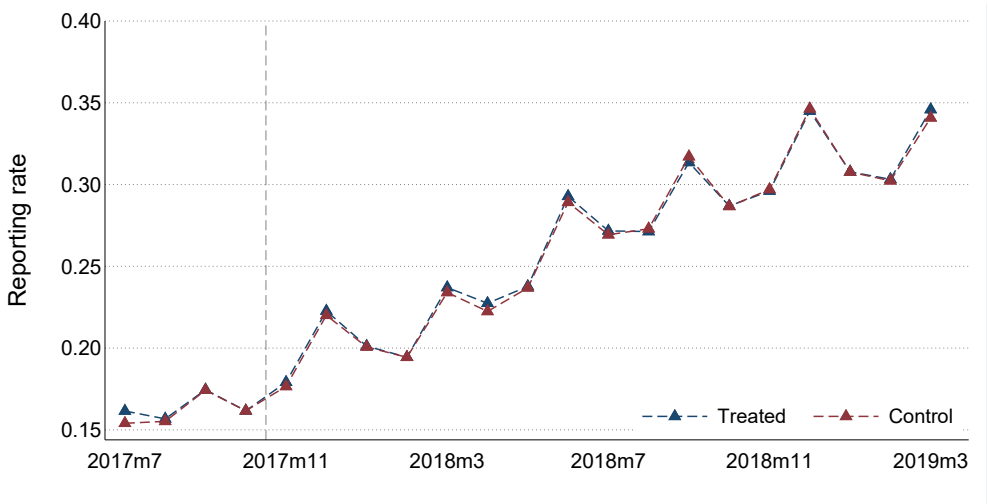
(B) Sales Value



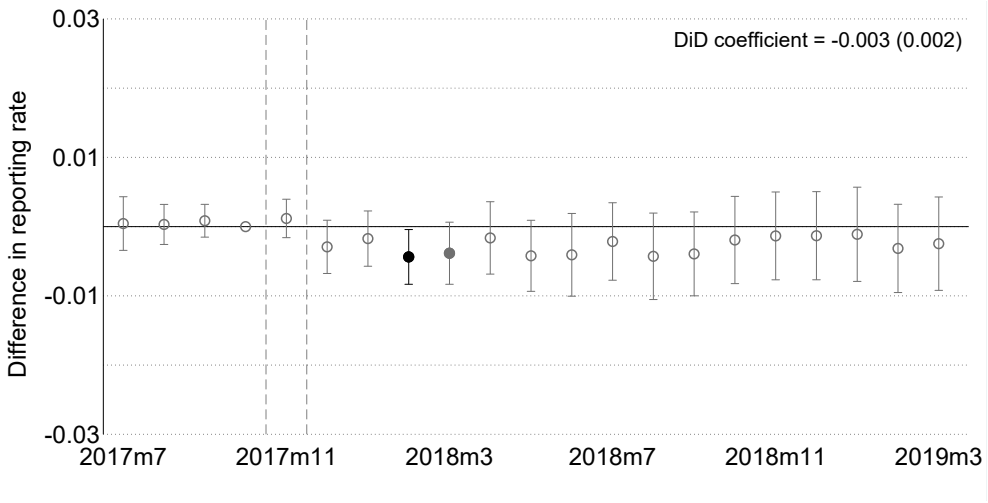
Note: This figure presents the event-study coefficients for log pre-tax prices and log sales. Sales amounts are pre-tax. Treated HSNs are those that experienced a reduction in Statutory Tax Rates from 28% to 18%. Control HSNs are those taxed at 12% or 18% that did not experience a rate change during the period. The *DiD coefficient* is the difference-in-differences estimate for the months after March 2018. The *Elasticity* (ε_s) is the elasticity of sales to Effective Tax Rates, instrumented with the treatment group and post periods. Hollow markers denote coefficients that are not statistically significant at the 10% level; grey markers denote significance at the 10% level and black markers at the 5% level. This figure is discussed in Section 4.1.

Figure 5: Effects of the Nov. 2017 GST Rate Cut: Reporting Rate (Extensive Margin)

(A) Probability of reporting a product: levels

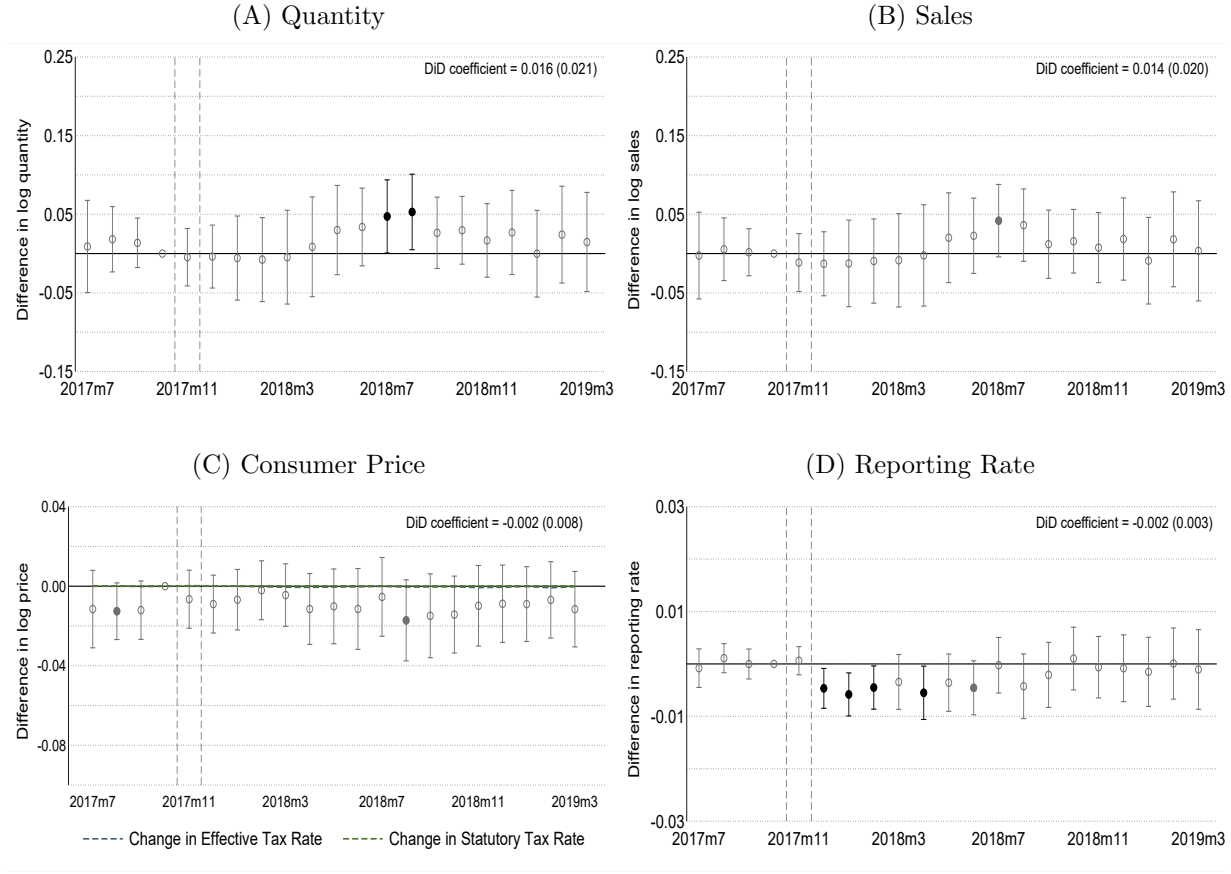


(B) Probability of reporting a product: Event-study coefficients



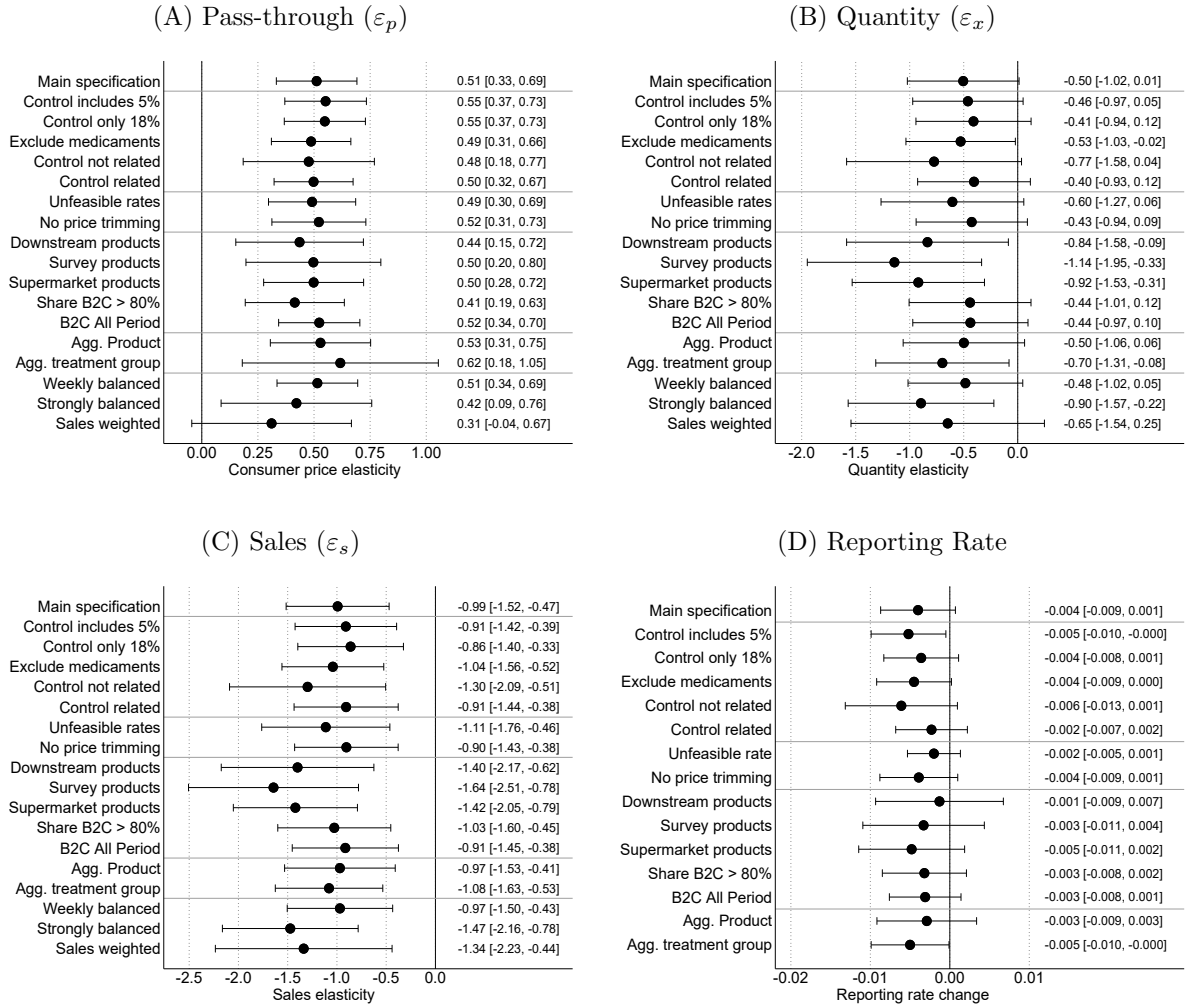
Note: This figure presents the results for the reporting rate. The reporting rate is defined as the probability of reporting sales of an HSN in a given month, conditional on reporting sales of that HSN in any month during the sample period. Treated HSNs are those whose statutory tax rate was reduced from 28% to 18%. Control HSNs are those taxed at 12% or 18% that did not experience a rate change during the period. The DiD coefficient is the difference-in-differences estimate for the months after March 2018. Hollow markers denote coefficients that are not statistically significant at the 10% level; grey markers denote significance at the 10% level and black markers at the 5% level. This figure is discussed in Section 4.1.

Figure 6: Results for Related but Untreated Products



Note: This figure presents the results for log quantity, log sales, log price, and the reporting rate, comparing related and non-related control products. Control products are those taxed at 12% or 18% that did not experience a rate change during the period. Related products are control products in the same HSN-2 category as a treated product. Non-related products are control products without any treated product in the same HSN-2 category. The *DiD coefficient* is the difference-in-differences estimate for the months after March 2018. Hollow markers denote coefficients that are not statistically significant at the 10% level; grey markers denote significance at the 10% level and black markers at the 5% level. This figure is discussed in Section 4.2.

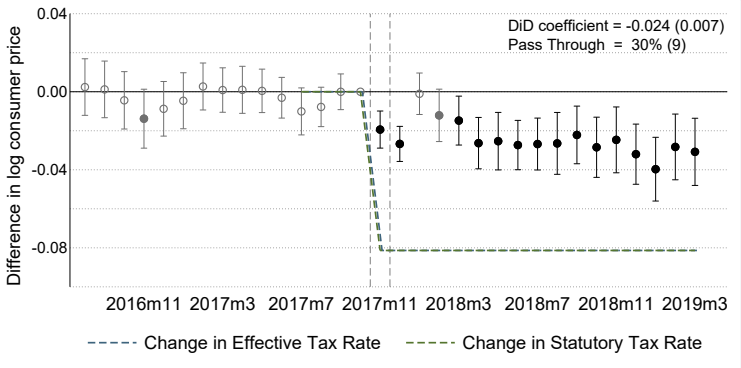
Figure 7: Alternative Samples and Specifications



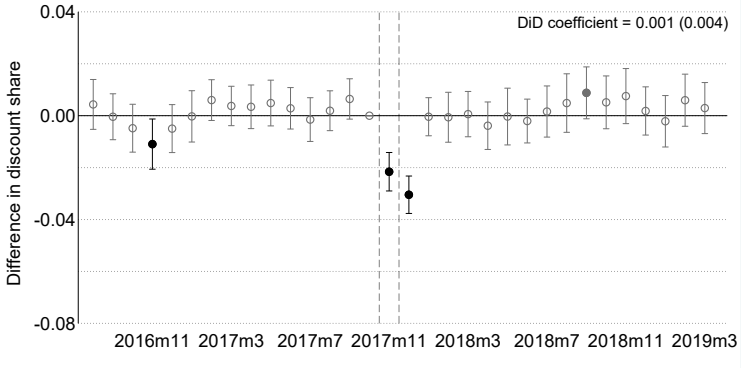
Note: The figure reports estimates from a series of alternative specifications. Each point estimate in Panels A–C comes from a regression of the outcome on the Effective Tax Rates, instrumented with the treatment group and post periods. Panel A reports the pass-through of consumer prices, and Panels B and C report the elasticities of quantities and sales, respectively. Panel D reports the OLS difference-in-differences coefficient for the reporting rate, defined as the probability that a firm reports sales of a product in a given month, conditional on reporting sales of that product at least once during the period analyzed. The coefficient estimates are reported to the right of each panel, with their 95% confidence intervals in brackets. The figure is discussed in Sections 4.4 and 5.

Figure 8: Hypermarket Scanner Data: Consumer Price, Discount, and Quantity Response

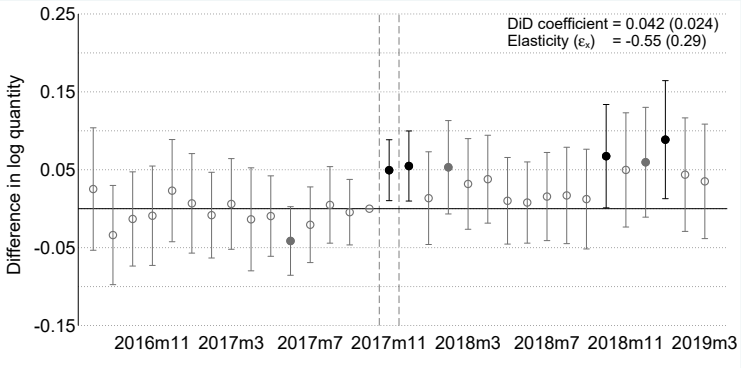
(A) Consumer Price (Inclusive of discounts)



(B) Discounts

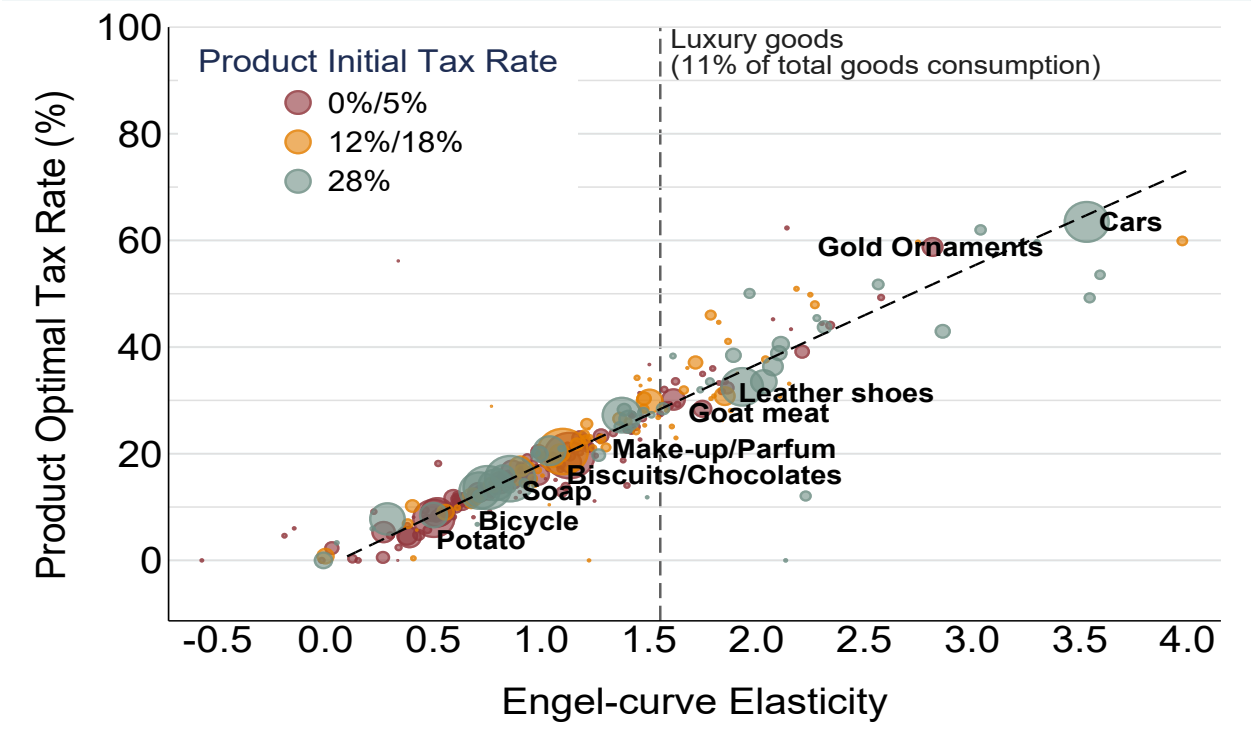


(C) Quantity



Note: This figure presents event-study coefficients for consumer prices, discount shares, and quantities sold, comparing barcodes that experienced a rate reduction from 28% to 18% in November 2017 with barcodes continuously taxed at 18% or 12% in the hypermarket dataset. The share of the price reduction attributable to discounts is defined as Discount Contribution $= -\frac{\text{Discount}}{\text{Consumer Price+Discount}} = -\frac{\text{Discount}}{\text{Pre-discount Price}}$. Effective and Statutory Tax Rates are equalized in the figure. Hollow markers denote coefficients that are not statistically significant at the 10% level; grey markers denote significance at the 10% level and black markers at the 5% level. This figure is discussed in Section 5.

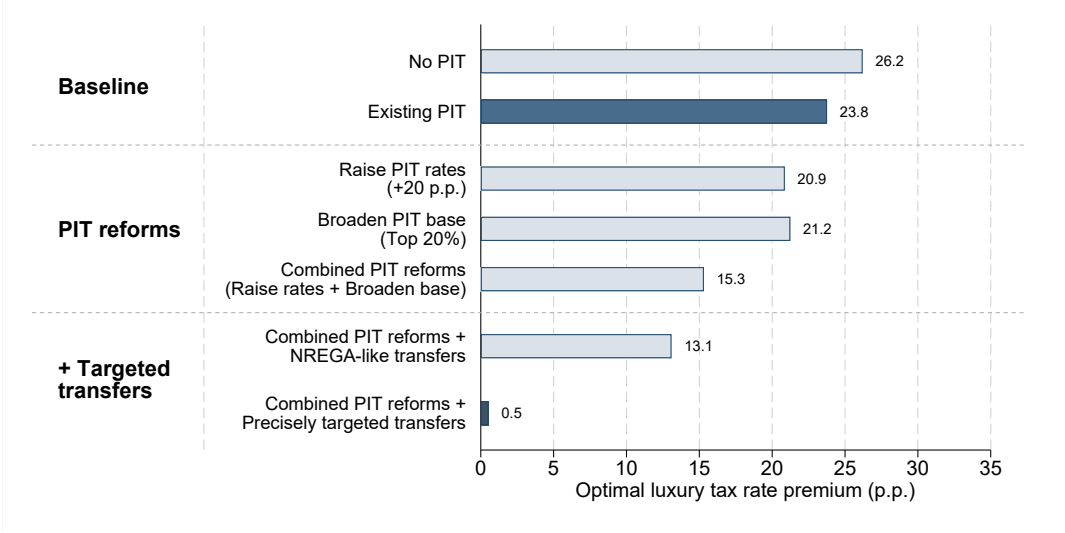
Figure 9: Optimal Commodity Tax Rates by Survey Product



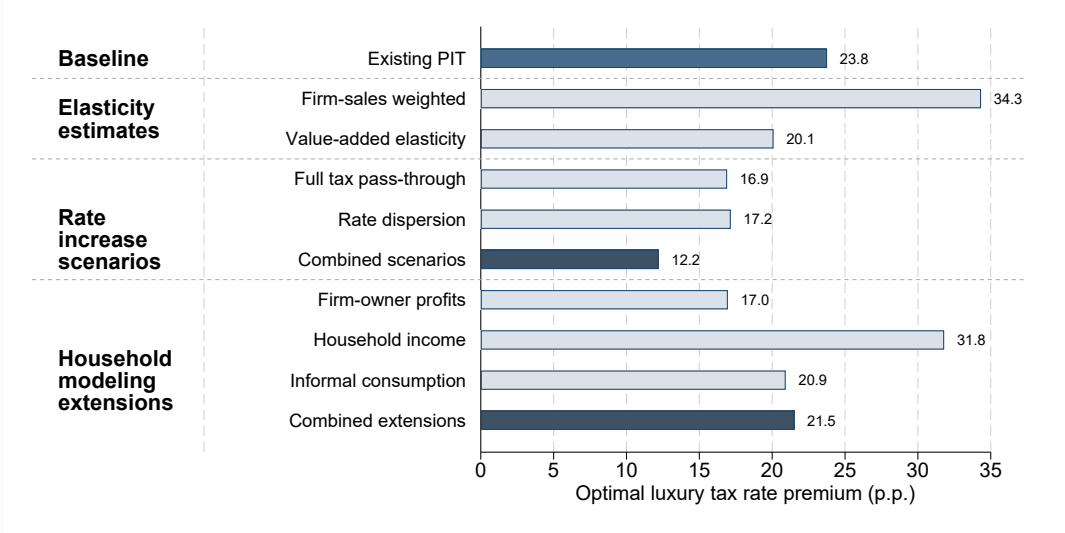
Note: This figure presents the estimated optimal tax rate for each survey good as a function of its Engel-curve elasticity. The dashed line shows the fitted relationship. Optimal tax rates below zero (i.e., subsidies) are winsorized at zero. Each circle represents a survey good, with circle size proportional to total consumption. Product labels illustrate selected high-consumption goods at different points along the optimal tax schedule. The vertical line labeled "Luxury goods" marks the threshold above which goods account for 11% of total goods consumption, matching the consumption share of goods in the pre-reform GST 28% tax bracket. The figure is discussed in Section 6.3.

Figure 10: Optimal Luxury Tax Premium under Alternative Tax Systems

(A) Relaxing Income Tax and Transfers Constraints



(B) Robustness and Extensions



Note: This figure reports the optimal luxury-tax premium, defined as the difference between the optimal indirect tax rate on luxury goods and the optimal rate on goods in the standard 18% bracket. Luxuries are the basket of goods with the highest Engel-curve elasticities, with a consumption share matched to that of goods initially taxed at 28%. Unless stated otherwise, PIT revenue is rebated as universal lump-sum transfers. Panel A varies the personal income tax and transfer system: starting from either no PIT or the existing PIT, it considers a 20 p.p. increase in PIT rates, a broadening of the PIT base by lowering PIT brackets to cover the top 20% of the income distribution, the combination of both reforms, and the combination of these reforms with NREGA-like transfers or transfers targeted to the bottom 50% of the income distribution. Panel B reports robustness results under alternative assumptions on tax-base elasticities, tax incidence, and household income, each relative to the baseline specification, together with specifications combining adjustments. Further details on the model are provided in Appendix F, and details on the calibration are provided in Appendix G. This figure is discussed in Section 6.3.

Table 1: Summary Statistics for Treated and Control Products

	Treated	Control	All
Products	164	680	844
Firm-HSNs	128,927	316,113	445,040
Of which before Nov 17	44,417	107,886	152,303
Before November 2017:			
Months reported	2.05 (1.17) [2.00]	2.04 (1.16) [2.00]	2.04 (1.16) [2.00]
HSN level reported	6.23 (1.97) [8.00]	6.28 (1.97) [8.00]	6.27 (1.97) [8.00]
Final consumers sales (Firm share)	0.81 (0.15) [0.83]	0.79 (0.16) [0.80]	0.80 (0.16) [0.81]
Sales ('000 rupees)	208 (2146) [8.11]	429 (15553) [7.17]	364 (13135) [7.44]
Sales (Log value)	9.09 (2.41) [9.00]	9.00 (2.56) [8.88]	9.03 (2.52) [8.92]
Quantity (Log units)	3.77 (2.62) [3.53]	4.16 (2.74) [3.95]	4.05 (2.71) [3.83]
Price pre-tax (Log value per unit)	5.32 (2.02) [5.08]	4.84 (2.12) [4.57]	4.98 (2.10) [4.74]
Observations	682,960	1,672,913	2,355,873
Of which before Nov 17	90,935	220,107	311,042

Note: This table presents the sample used in the firm–HSN-level analysis for the November 2017 rate reduction from 28% to 18%. The *Treated* column includes products that experienced this rate reduction in November 2017. The *Control* column includes products that did not undergo any rate change between July 2017 and March 2020 and were taxed at 12% or 18% throughout. The *All* column combines both groups. The table is divided into three groups of rows: Products, Firm–HSN, and Observations. Standard deviations are in parentheses and the median in brackets. Statistics are measured before November 2017. This table is discussed in Section 3.4.

Table 2: Firm-Level Summary Statistics

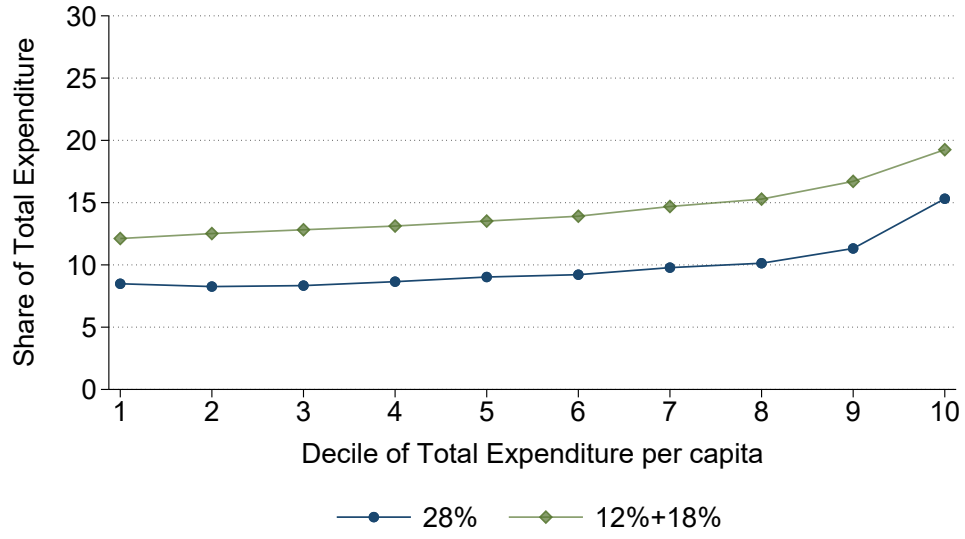
Number of Firms	27,081	
before Nov 17	17,783	
Number of Tax returns	248,379	
before Nov 17	42,899	
For tax returns before Nov 17:	Mean	Median
Sales ('000 rupees)	6,770	881
Treated	442	0
Control	2,200	44
Other	4,128	140
Number of products	7.60	3
Treated	1.45	0
Control	3.27	1
Other	2.90	1
Sells $>$ 1 product	0.66	1
Sells $>$ 1 product (Sales weighted)	0.89	1
Sells treated product	0.42	0
Sells treated product (Sales weighted)	0.57	1
Conditional on selling treated product:		
Sells $>$ 1 product	0.91	1
Sells $>$ 1 product (Sales weighted)	0.99	1
Sells control product	0.81	1
Sells control product (Sales weighted)	0.88	1

Note: This table reports descriptive statistics for the firms and tax returns (firm–months) included in the main sample, as defined in Section 3. Products refer to the number of products reported in a firm-product tax return. Rows labeled (Sales weighted) are weighted by firm sales. This table is discussed in Section 3.4.

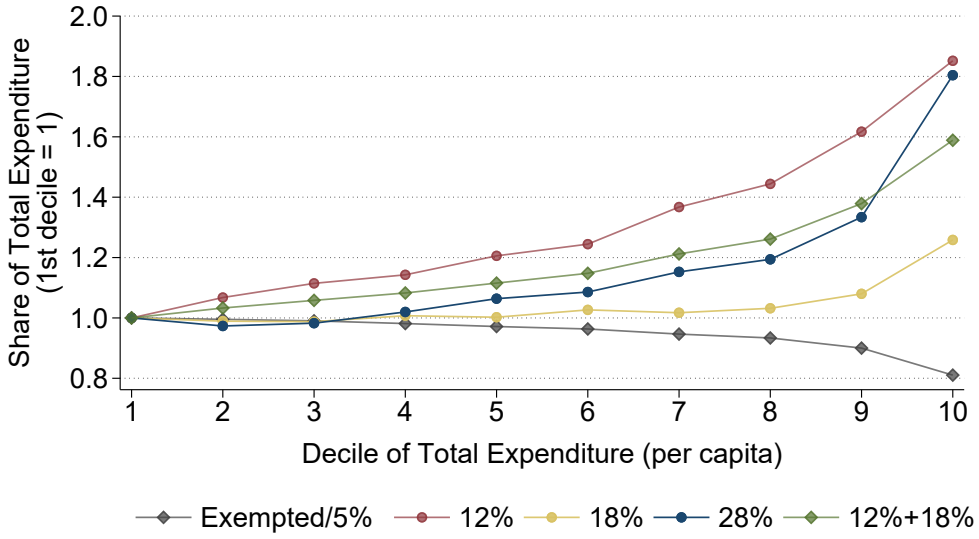
Appendix A: Additional Figures and Tables

Figure A.1: Expenditure Shares by GST Rate Across Expenditure Deciles

(A) Baseline shares

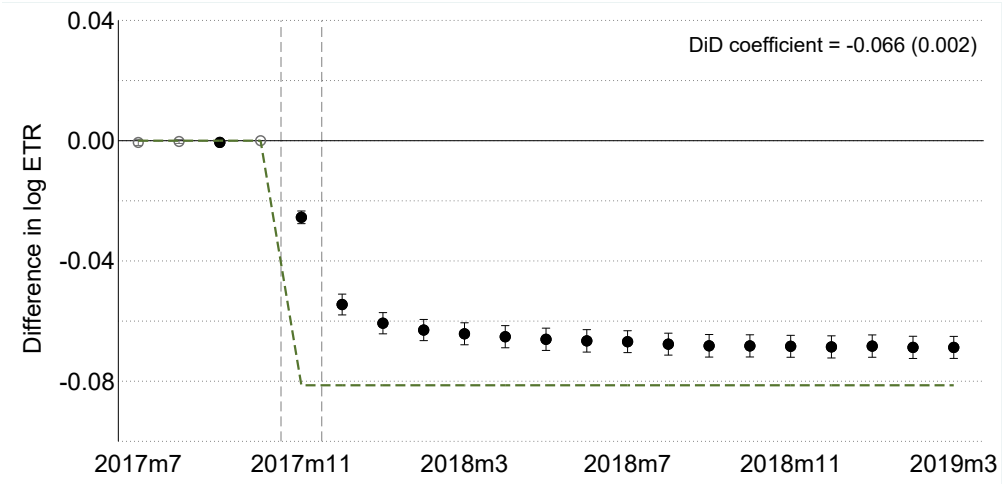


(B) Normalized by bottom decile



Note: Panel A presents the share of total consumption by product GST rate in July 2017 across deciles of household total expenditure per capita, based on the 2011–12 household expenditure survey. Panel B normalizes each tax-rate share by the corresponding share of the first decile. Services, goods received as gifts or free of charge, and petrol and diesel for vehicles are excluded. This figure is discussed in Section 3.4.

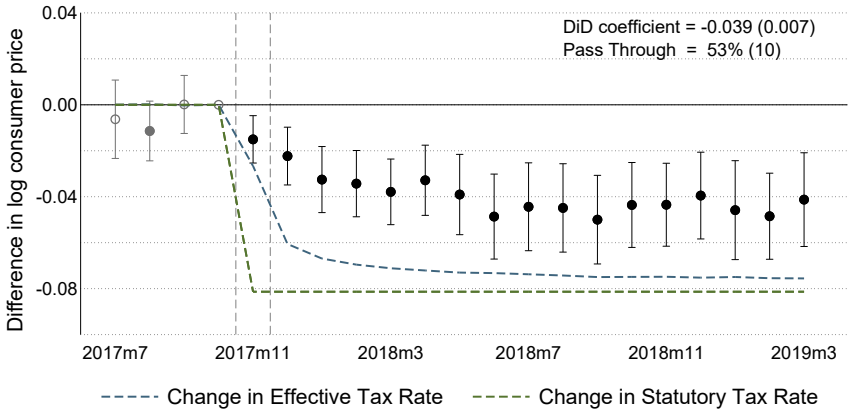
Figure A.2: Statutory & Effective Tax Rates Around the Nov. 2017 Reform (logs)



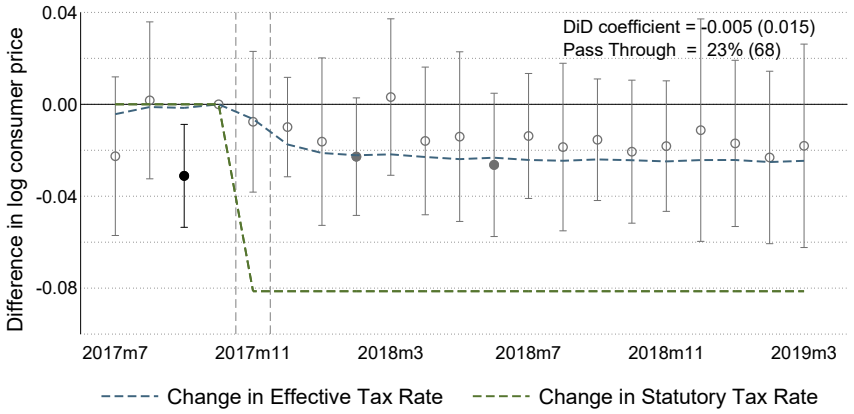
Note: The figure presents the event-study coefficients for changes in log Effective Tax Rates. The dashed green line represents the change in log Statutory Tax Rates (STRs). Treated HSNs are those that experienced a reduction in STR from 28% to 18%. Control HSNs are those taxed at 12% or 18% that did not experience a rate change during the period. Hollow markers denote coefficients that are not statistically significant at the 10% level; grey markers denote significance at the 10% level and black markers at the 5% level. This figure is discussed in Section 4.1.

Figure A.3: Consumer Price and Pass-Through by STR–ETR Match Status

(A) ETR matches STR

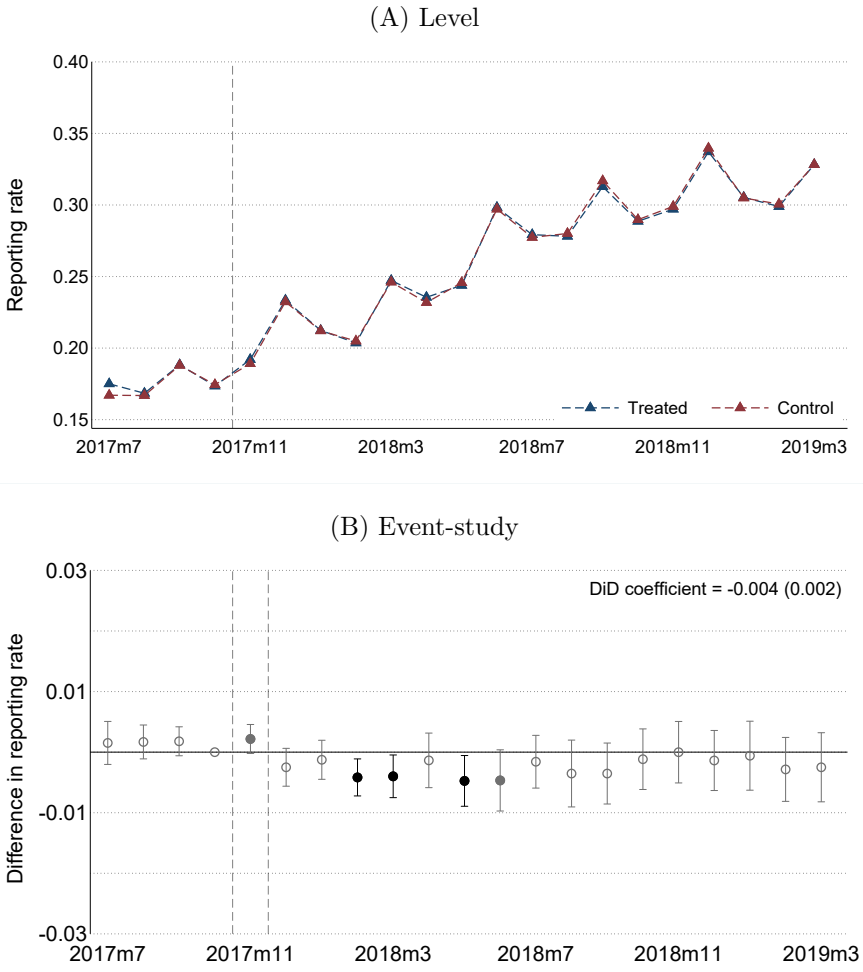


(B) ETR does not match STR



Note: This figure presents the event-study coefficients for the difference in log consumer price. Treated HSNs are those that experienced a reduction in Statutory Tax Rates (STRs) from 28% to 18%. Control HSNs are those taxed at 12% or 18% that did not experience a rate change during the period. The green and the blue dashed lines indicate respectively the difference in log STR and log Effective Tax Rates (ETR) between the two groups. Pass-through is defined as the ratio between the reduction in consumer price and the reduction in ETR. The group *ETR matches STR* includes treated firm-HSNs with a pre-November 2017 ETR no lower than 27% and no higher than 29%, that is, those with an ETR error lower than 1%. The group *ETR does not match STR* includes treated firm-HSNs with a pre-November 2017 ETR lower than 27% or higher than 29%, that is, those with an ETR error greater than 1%. The *DiD coefficient* is the difference-in-differences estimate for the months after March 2018. Hollow markers denote coefficients that are not statistically significant at the 10% level; grey markers denote significance at the 10% level and black markers at the 5% level. This figure is discussed in Section 4.1.

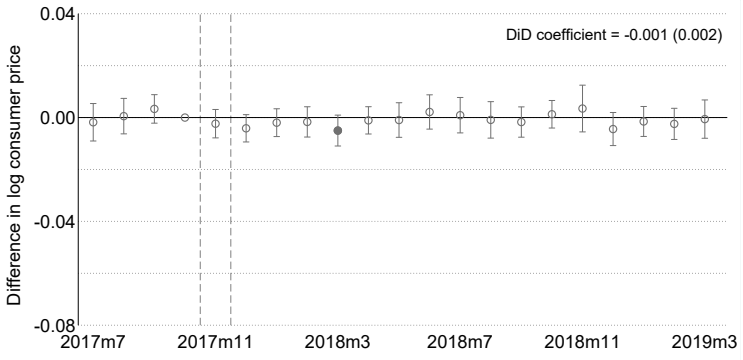
Figure A.4: Effects of the Nov. 2017 GST Rate Cut: Reporting Rate, Alternative Definition



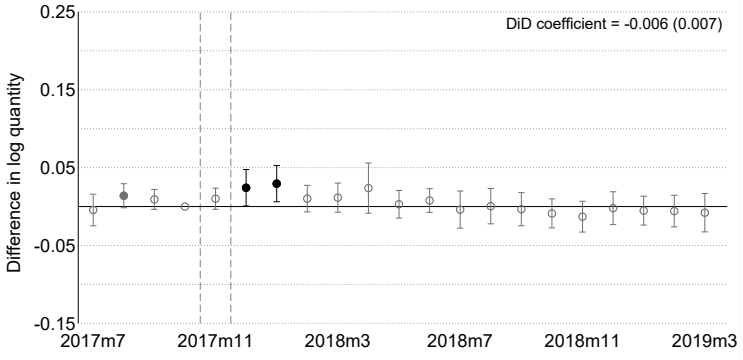
Note: This figure presents the results for the reporting rate. The reporting rate is defined as the probability of reporting selling an HSN in a given month conditional on having sold it at any point in the period analyzed (contrarily to our baseline definition, shown in Figure 5 which conditions the sample to HSN sold in the pre-reform period). Treated HSNs are those that experienced a reduction in Statutory Tax Rates (STRs) from 28% to 18%. Control HSNs are those taxed at 12% or 18% that did not experience a rate change during the period. The *DiD coefficient* is the difference-in-differences estimate for the months after March 2018. Hollow markers denote coefficients that are not statistically significant at the 10% level; grey markers denote significance at the 10% level and black markers at the 5% level. This figure is discussed in Section 4.1.

Figure A.5: Within-Firm Spillover Effects to Untreated Products

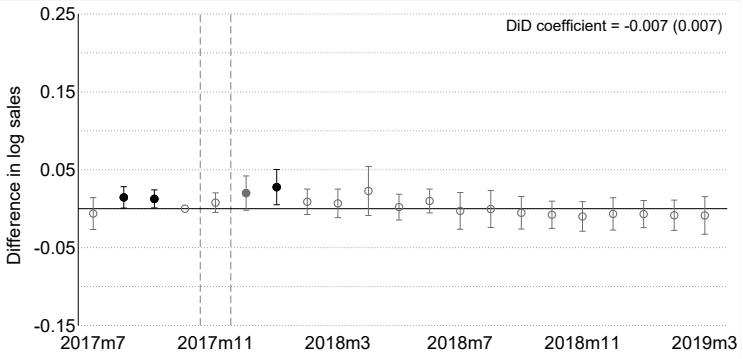
(A) Consumer Price



(B) Quantity

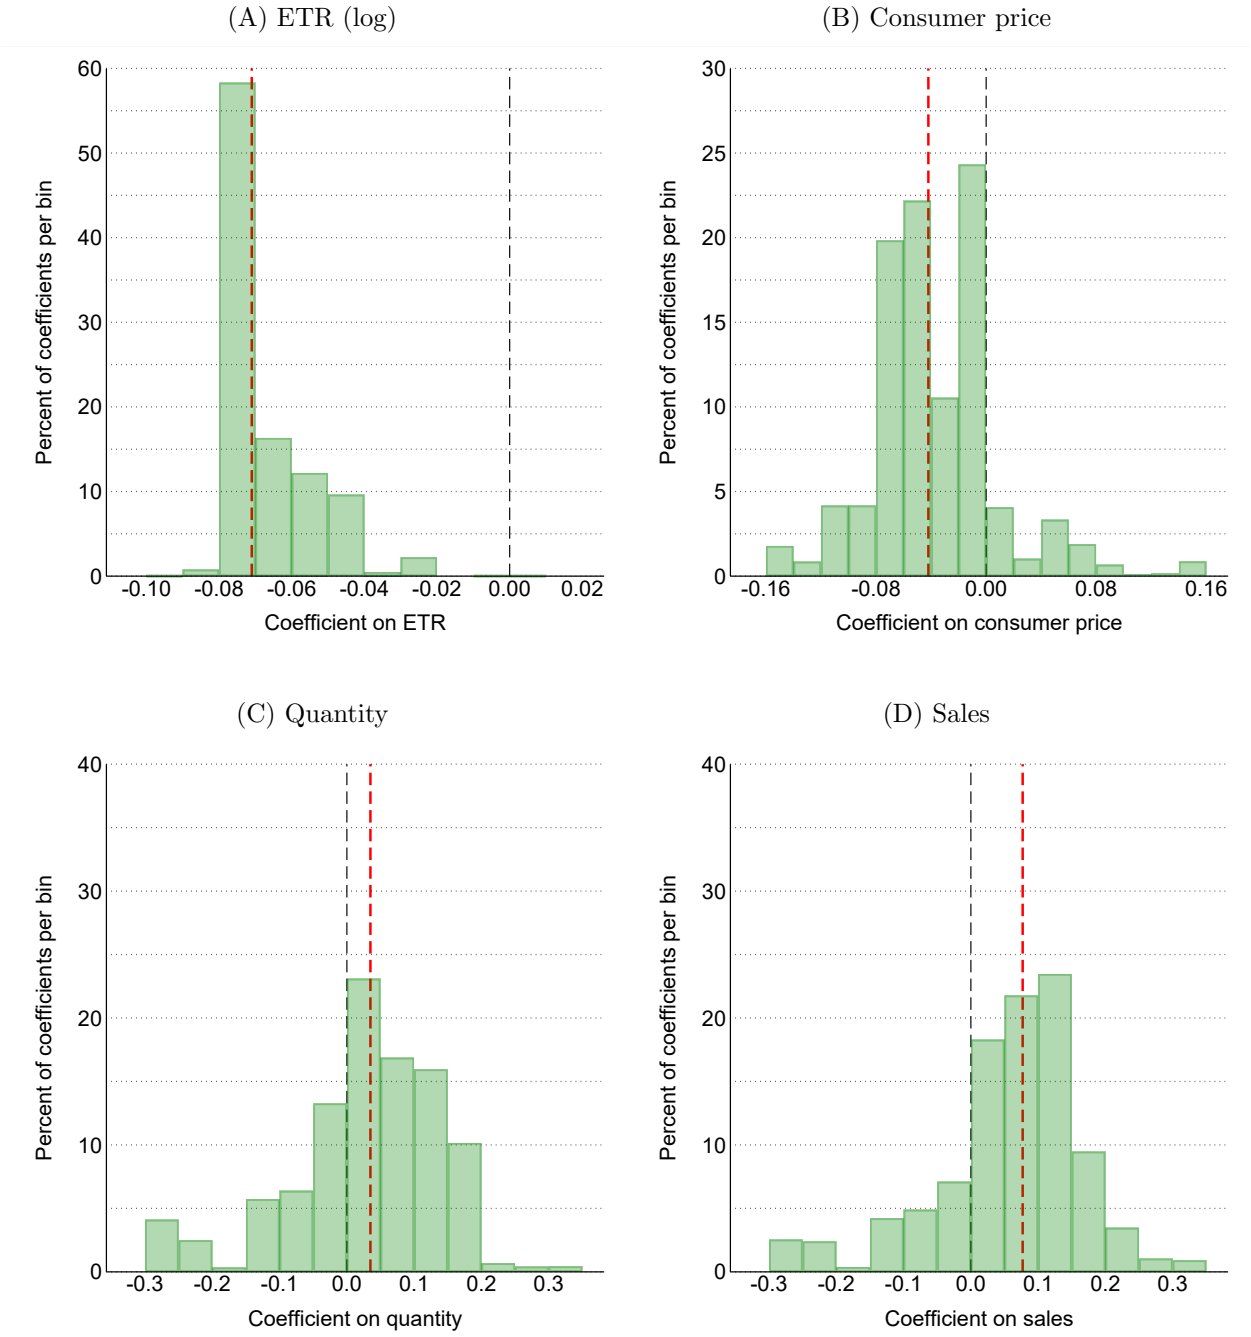


(C) Sales



Note: This figure plots the coefficients γ_t from equation (2) for consumer prices, quantities, and sales, estimated on control HSNs only, where the baseline firm-level treated share is interacted with monthly indicators. The baseline firm-level treated share is rescaled by the average baseline share of treated sales among exposed firms, equal to 14%. All regressions include firm–HSN fixed effects (λ_{ip}) and HSN–time fixed effects (δ_{pt}), with standard errors clustered at the firm and HSN-4 levels. Hollow markers denote coefficients that are not statistically significant at the 10% level; grey markers denote significance at the 10% level and black markers at the 5% level. This figure is discussed in Section 4.3.

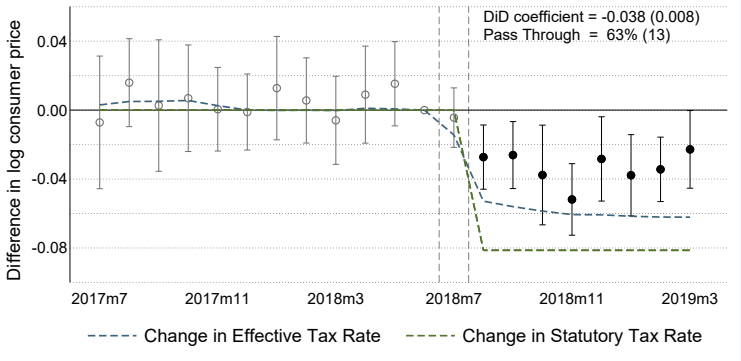
Figure A.6: Distribution of Product-Level Treatment Effects



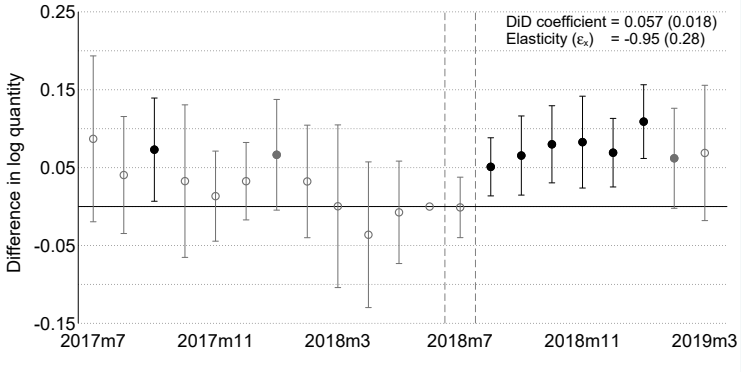
Note: This figure displays coefficients of the difference-in-differences for each treated product for log sales, quantity, Effective Tax Rates and consumer prices. Coefficients are weighted by the number of observations before November 2017 and winsorized at each histogram's maximum and minimum values. The red line indicates the average. This figure is discussed in Section 4.4.

Figure A.7: Effects of the July 2018 GST Rate Cuts

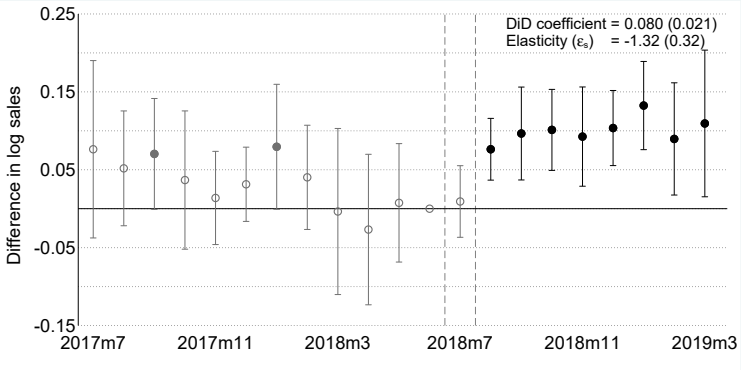
(A) Consumer Price



(B) Quantity

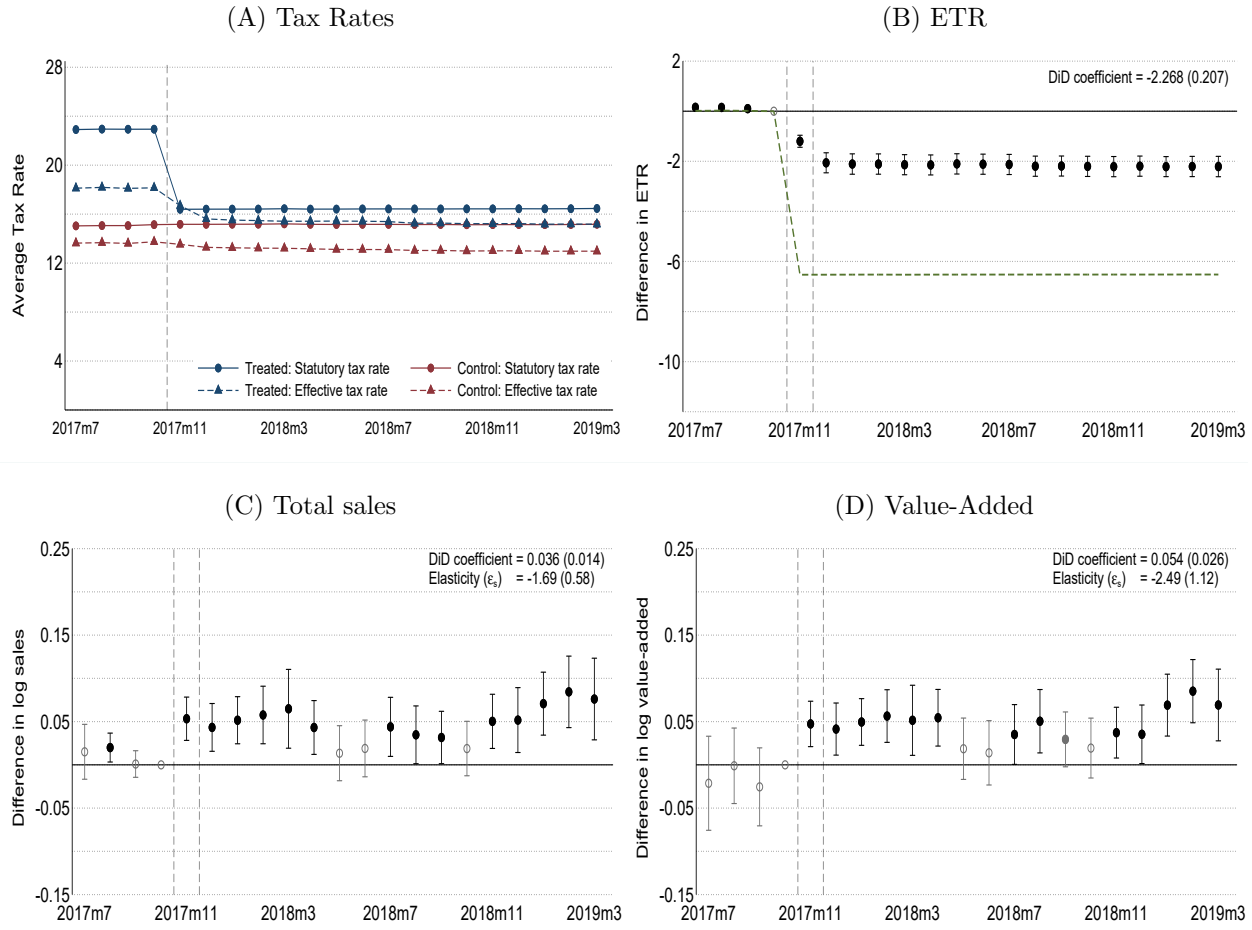


(C) Sales



Note: This figure presents the results for log quantity, log sales and log price. Treated HSNs are those that experienced a reduction in Statutory Tax Rates (STRs) from 28% to 18% in July 2018. Control HSNs are those taxed at 12% or 18% that did not experience a rate change between July 2017 and March 2019. The *DiD coefficient* is the difference-in-differences estimate for the months after August 2018 compared to the months before July 2018. Hollow markers denote coefficients that are not statistically significant at the 10% level; grey markers denote significance at the 10% level and black markers at the 5% level. This figure is discussed in Section 4.4.

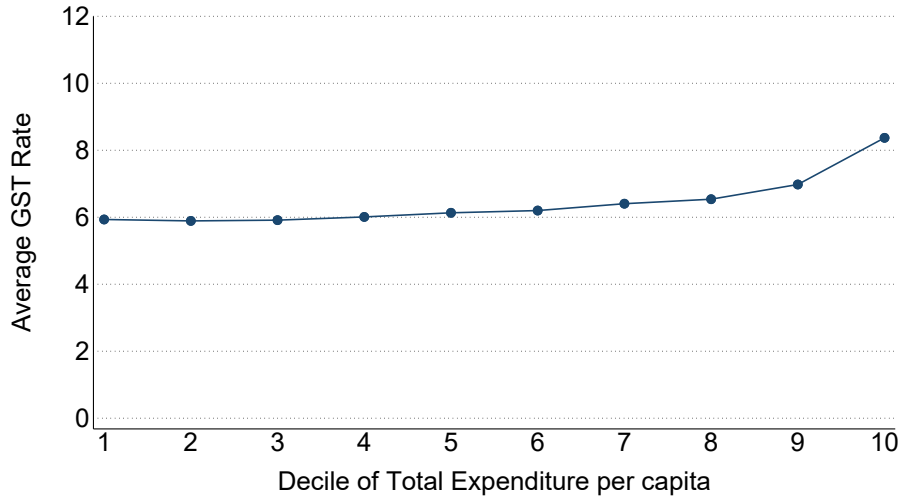
Figure A.8: Firm-Level Results (Nov. 2017 GST Rate Cut), All Registered Firms



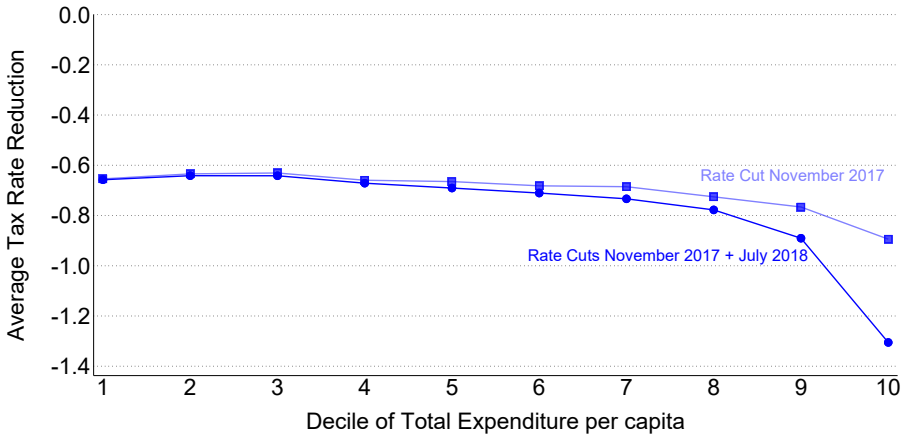
Note: Panel A presents the average statutory tax rate and the average effective tax rates by treatment status using firm-level tax returns (GSTR-3B). Panel B presents the event-study coefficients comparing treated and control firms for the change in ETR and STR, Panel C for the log total sales, and Panel D for the log value added. Treated firms are those that, in the registration form, report selling at least one treated product. Effective tax rates are winsorized at the 1st and 99th percentile. Value added is defined as the monthly difference between sales and intra-state purchases, calculated from firm-to-firm transactions. The analysis is restricted to firm-months with positive value added. Hollow markers denote coefficients that are not statistically significant at the 10% level; grey markers denote significance at the 10% level and black markers at the 5% level. This figure is discussed in Section 4.5.

Figure A.9: Average GST Rates Across Expenditure Deciles

(A) Baseline average rate



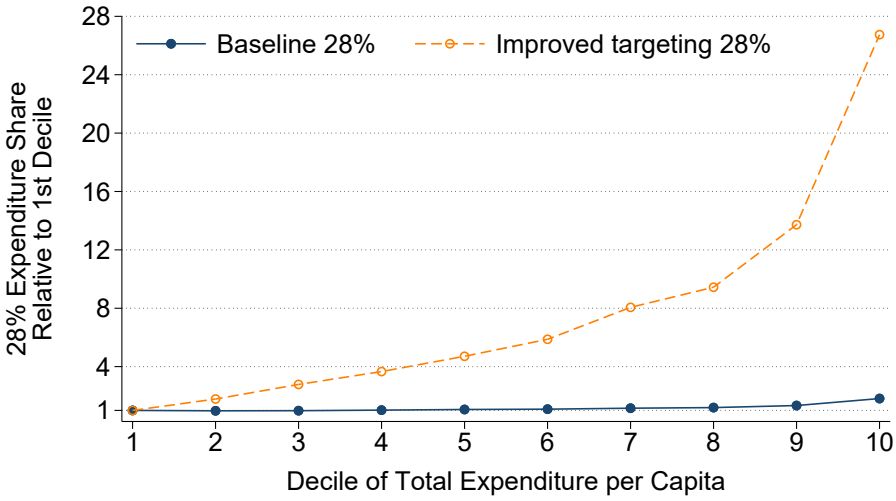
(B) Average rate reduction



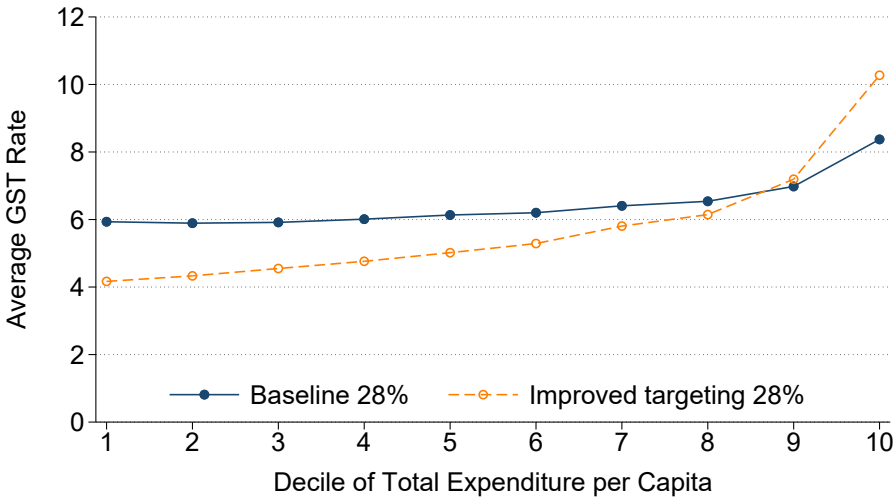
Note: Panel A presents the average GST rate in July 2017 across deciles of household total expenditure per capita, defined as total taxes divided by total consumption in each decile. Panel B presents the corresponding average rate reduction following the November 2017 and July 2018 rate cuts, defined as total taxes before minus total taxes after the reduction divided by total consumption in each decile. Services, goods received as gifts or free of charge, and petrol and diesel for vehicles are excluded. This figure is discussed in Section 6.1.

Figure A.10: Improving the Targeting of the 28% GST Bracket

(A) Normalized 28% share

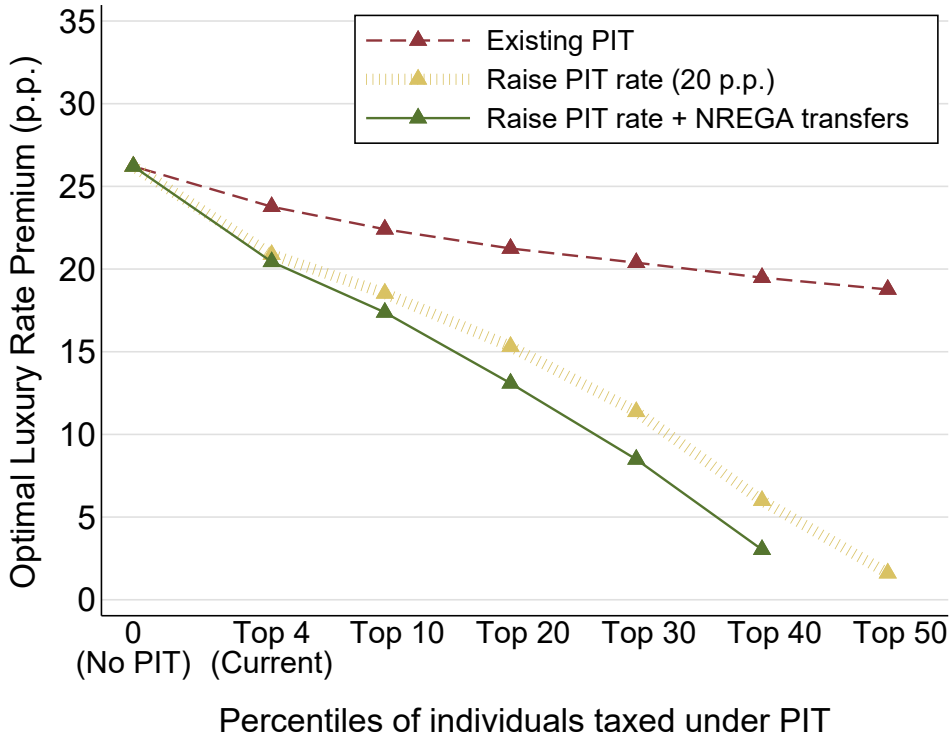


(B) Average GST rate



Note: Panel A presents the normalized expenditure share of goods taxed at 28% across deciles of household total expenditure per capita, normalized by the share of the first decile. Panel B presents the corresponding average GST rate. Improved targeting reassigns to the 28% bracket the goods with the highest ratio between top and bottom deciles, keeping the total amount taxed at 28% and the average tax rate constant. Services, goods received as gifts or free of charge, and petrol and diesel for vehicles are excluded. This figure is discussed in Section 6.1.

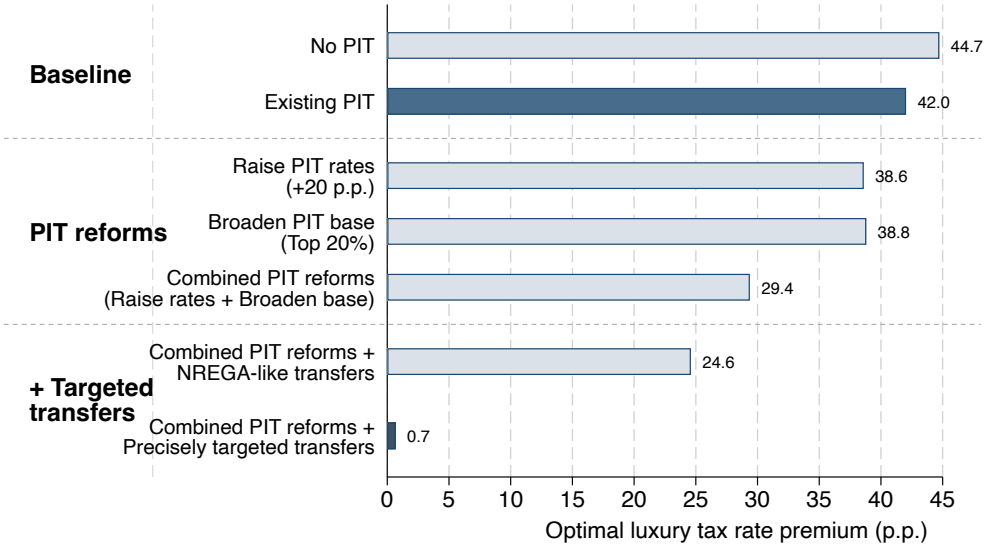
Figure A.11: Optimal Luxury Taxation: Expanding PIT Coverage



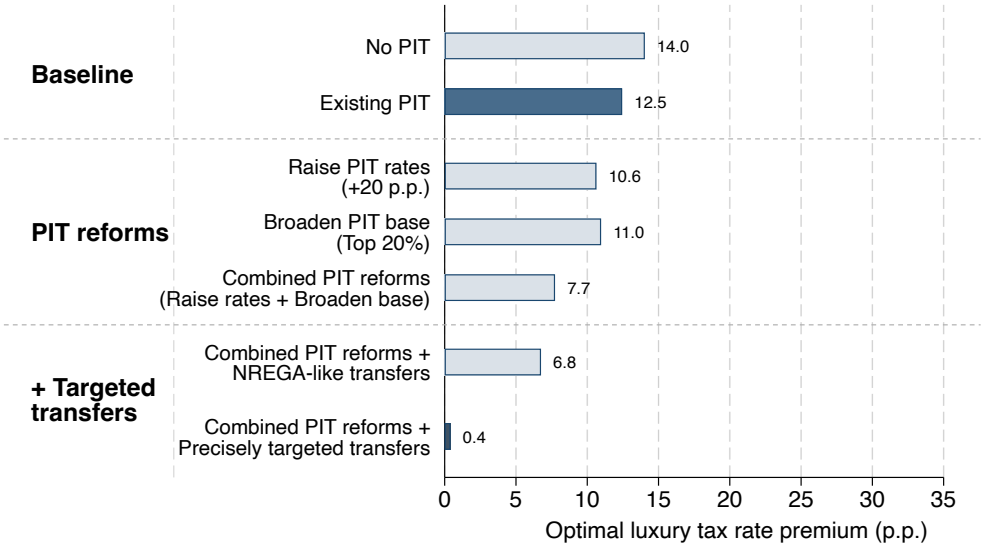
Note: This figure reports the optimal luxury-tax premium, defined as the difference between the optimal indirect tax rate on luxury goods and the optimal rate on goods in the standard 18% bracket. Luxuries are the basket of goods with the highest Engel-curve elasticities, with a consumption share matched to that of goods initially taxed at 28%. Unless stated otherwise, PIT revenue is rebated as universal lump-sum transfers. The x-axis varies the personal income tax (PIT) threshold, expressed as the percentiles taxed at least at the lowest PIT rate. It ranges from a scenario without PIT (0—No PIT) to increasingly broad PIT bases. The lines represent various personal income tax and transfer systems: the existing PIT rate structure, a 20 p.p. increase in all PIT rates, and the combination of a 20 p.p. increase in PIT rates with NREGA-like transfers. Further details on the model are provided in Appendix F, and details on the calibration are provided in Appendix G. This figure is discussed in Section 6.3.

Figure A.12: Optimal Luxury Tax Premium, Varying Inequality Aversion

(A) High Inequality Aversion (γ=2)



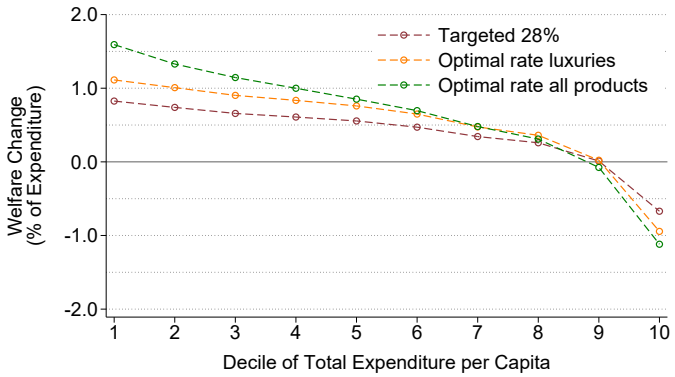
(B) Low Inequality Aversion (γ=0.5)



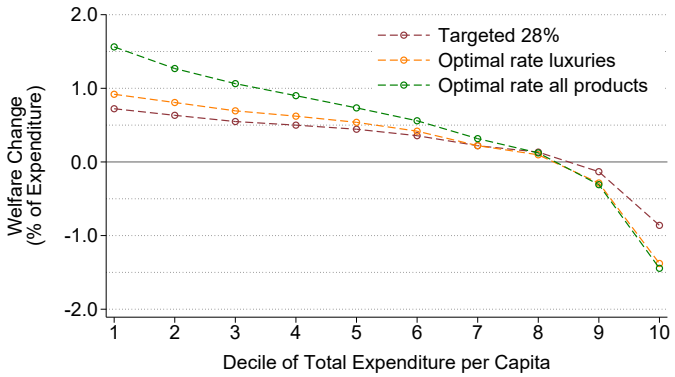
Note: This figure reports the optimal luxury-tax premium, defined as the difference between the optimal indirect tax rate on luxury goods and the optimal rate on goods in the standard 18% bracket. Luxuries are the basket of goods with the highest Engel-curve elasticities, with a consumption share matched to that of goods initially taxed at 28%. We define social marginal welfare weights as $g^i = (x^i)^{-\gamma}$, where x^i denotes household per capita expenditure and γ is the coefficient of relative inequality aversion. Our baseline sets $\gamma = 1$ (log utility), and we examine alternative values of γ in this figure. This figure is discussed in Section 6.3.

Figure A.13: Revenue-Neutral Welfare Effects of Commodity Tax Reform

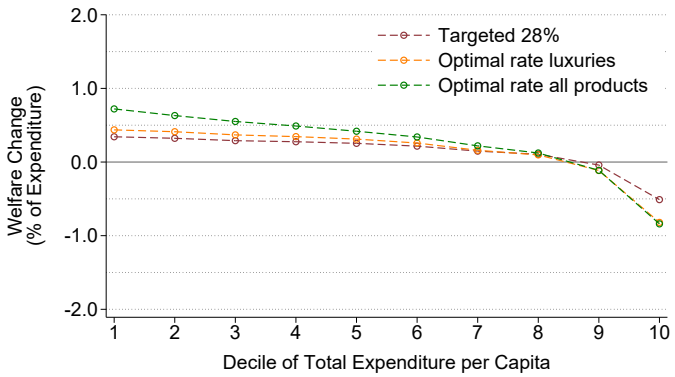
(A) Tax Induced Price Change



(B) Including Consumption Responses



(C) Including Consumption Responses & Informality



Note: This figure shows the average welfare change by decile of household total expenditure per capita under alternative assignments of goods to tax brackets, computed from a second-order approximation derived in Appendix G. Panel A reports the welfare change from the price change using the estimated tax pass-through, Panel B adds the substitution correction using the estimated sales elasticities for survey products, and Panel C accounts for informal consumption by restricting welfare effects to the formal share of spending. Tax rates are linearly adjusted to keep the overall tax burden across rate-assignment alternatives constant. This figure is discussed in Section 6.3.

Table A.1: Example of the HSN Structure

Level description	Item code	Level	Level-digit
Essential oils; cosmetic and toilet preparations; soap and washing; lubricating preparations	6	Group	-
Essential oils and resinoids; perfumery, cosmetic or toilet preparations	33	Chapter	2
Beauty or make-up preparations and preparations for the care of the skin	3304	Heading	4
Make-up or skin care powders, incl. baby powders, whether or not compressed	330491	Subheading	6
Powder face	33049110	HSN code	8

Note: This table presents one example of the HSN structure for an HSN-8 code. "Powder face" is an HSN-8 code, included in the HSN-6 code for "make-up or skin care powders". This HSN-6 code is part of the HSN-4 code for "beauty, make-up or skin care preparations". The corresponding HSN-2 code is "essential oils and resinoids, perfumery, cosmetic or toilet preparations". The proximity among products in a level varies from case to case, so HSN levels can be more or less specific than this example. This table is discussed in Section 2.

Table A.2: Number of Products per HSN-4

Products Per HSN-4	Number of HSN-4s	Percent of HSN-4s
1	1,133	91.30
2	92	7.41
3	14	1.13
4	2	0.16
Total	1,241	100.00

Note: Products are defined as unique combinations of an HSN-4 code and rate path, characterized by its initial rate, final rate, and any rate-change date. When an HSN-4 contains multiple rates, products are defined at the HSN-6 level, grouping HSN-6 codes that share the same rate path. The table shows that 91% of HSN-4s are a product and 9% of HSN-4s have more than one product. This table is discussed in Section 3.1.

Table A.3: Observation-Level Sample Cleaning and Restriction Steps

Cleaning Step	Number of observations	Share (%)
Initial sample	12,815,721	100
Services	832,835	6.5
Incomplete Tax Returns		
Missing product code	361,296	2.8
Non-existent code	384,889	3.0
HSN-2 code	105,158	0.8
Multiple rates	154,946	1.2
No quantity	236,195	1.8
Final consumers focus		
Final consumer sales share < 0.5	5,837,308	45.5
Final consumer sales share missing	642,063	5.0
Consistency filters		
Infeasible rate	468,690	3.7
Extreme price variation	177,652	1.4
Final sample	3,614,689	28.2
Before November 2017	469,671	-

Note: This table reports the number of observations removed at each sample cleaning and restriction step. The top row gives the initial sample of firm–product–months with sales, and the bottom row shows the final number of observations after data cleaning. The share column reports each value as a percentage of the initial sample. The number of observations concerned by each cleaning step can vary depending on the elimination order, since some observations may present more than one of the attributes listed above. This table is discussed in Section 3.3.

Table A.4: Product-Level Sample Restrictions

Treatment Group	Products	Observations (N)	Observations share	Sales share
Totals before November 2017	1367	469,671	100	100
Treated products	164	90,935	19.4	6.6
Control products	680	220,107	46.8	32.9
Not related	332	91,287	19.4	10.0
Related	348	128,820	27.4	22.9
Other treated	108	39,022	8.3	5.7
July 2018	32	16,797	3.6	2.6
Other periods	29	12,797	2.7	2.3
Other rates	47	9,428	2.0	0.8
Other control	415	119,607	25.4	54.7
Rate 0% (exempted)	140	25,405	5.4	29.6
Rate 5%	246	74,867	15.9	15.4
Rate 28%	29	19,335	4.1	9.7

Note: This table presents the number of products in each group. A product is defined as an HSN-4 level code and a rate path. There are 1,241 HSN-4 level codes and 1,367 HSN-4 level code–rate paths in the sample. The *Treated* group includes products that experienced a rate reduction from 28% to 18% in November 2017. The *Control* group includes products taxed at 18% or 12% throughout the entire period. The Control subgroup labeled *Not Related* consists of control products that are not in the same HSN-2 code as a treated product, while the subgroup *Related* includes control products that are in the same HSN-2 code as a treated product. The *Other Treated* group includes all products that experienced a rate change between July 2017 and March 2020 but are not in the main Treated group. Subgroup *July 2018* refers to products that had a rate reduction in that month, out of which 25 had a rate reduction from 28 to 18 percent. Subgroup *Other periods* include products with rate reductions in months other than November 2017 or July 2018. Subgroup *Other rates* include products that had a rate reduction in November 2017 but not from 28% to 18%. The *Other control* category includes all products without a rate change that are not part of the main control group. Subgroup *Rate 0% (exempted)* refers to exempted products; subgroup *Rate 5%* to products taxed at 5%; and subgroup *Rate 28%* to products taxed at 28%. This table is discussed in Section 3.4.

Table A.5: Product-Level GST Rates in July 2017 and March 2020

Initial \ Final	0%	5%	12%	18%	28%	Total initial (%)	Treated (%)
0%	122	0	0	0	0	9	0
5%	11	264	0	0	0	20	4
12%	4	15	167	0	0	14	10
18%	1	4	33	513	0	40	7
28%	0	0	6	198	29	17	88
Total final (%)	10	21	15	52	2		

Note: This table presents the number of products in each combination of initial (July 2017) and final (March 2020) GST rates. Products are defined by an HSN-4 code and their combination of initial and final tax rates. This table is discussed in Section 3.4.

Table A.6: Top Products with a Rate Reduction in November 2017

Product name	Product code (HSN-4)	Treated products (%)	Cumulative share (%)
Make-Up	3304	7.6	7.6
Shaving Preparations	3307	7.2	14.8
Washing, Cleaning Or Degreasing Preparations	3402	5.9	20.7
Hair Preparations	3305	4.9	25.5
Chocolate	1806	3.9	29.4
Electrical Circuits (Switchers, Lamp-Holders, Fuses)	8536	3.0	32.4
Malt Extract	1901	2.9	35.3
Polishes And Creams	3405	2.9	38.2
Electrical Transformers	8504	2.9	41.1
Insulated Electric Conductors (Wires)	8544	2.9	44.1
Air Vacuum Pumps	8414	2.5	46.5
Liquid Or Cream Soap	3401	2.4	48.9
Mountings	8302	2.2	51.2
Lamps	8539	2.2	53.4
Ceramic Flags And Paving	6907	2.0	55.4
Lighting Or Signalling Equipment	8512	2.0	57.3
Perfumes And Toilet Waters	3303	1.9	59.3
Furniture	9403	1.9	61.2
Rubber Articles (Gaskets, Seals, Bands)	4016	1.8	63.0
Extracts Of Coffee	2101	1.3	64.3
Sanitary Fixtures	6910	1.3	65.7
Plastic Articles (Office And School Supplies)	3926	1.3	67.0
Inner Tubes	4013	1.3	68.2
Bread, Pastry, Cakes, Biscuits	1905	1.2	69.4
Suitcases, Vanity Cases, Briefcases	4202	1.2	70.6
Baths And Sinks	3922	1.2	71.8
Batteries	8506	1.1	72.9
Mirrors	7009	1.0	73.9
Painters' Colours	3213	1.0	75.0
Razors	8212	1.0	75.9
Other Products	.	24.1	100.0

Note: This table presents the names and HSN codes of the main products that had a rate reduction in November 2017. The name is a simplified version of the official HSN name. Additionally, the table displays the share and cumulative share that each product represented among treated products, considering the number of observations before November 2017. For example, Shaving Preparations is the second most important category, with its HSN-4 code being 3307. It accounted for 7.2% of the firm–products with a rate change. Shaving Preparations and Make-Up accounted together for 14.8% of firm–products with a rate change. This table is discussed in Section 3.4.

Table A.7: Top Control Products (12 or 18% GST Rate Throughout)

Product name	Product code (HSN-4)	Control products (%)	Cumulative share (%)
Medicaments In Measured Doses	3004	21.9	21.9
Insecticides, Rodenticides, Herbicides	3808	6.9	28.8
Bar Soap	3401	3.4	32.2
Tubes, Pipes And Hoses Of Plastic	3917	3.4	35.6
Medical Instruments	9018	3.0	38.6
Brooms And Brushes	9603	3.0	41.5
Screws, Bolts, Nuts Of Iron/Steel	7318	2.9	44.4
Petroleum And Bituminous Oil	2710	2.6	47.0
Telephone Sets	8517	2.5	49.6
Prepared Glues And Adhesives	3506	2.5	52.1
Computers	8471	2.3	54.3
Medicaments (Not In Measured Doses)	3003	2.1	56.4
Packing And Closure Goods Of Plastic	3923	2.0	58.4
Centrifuges And Filtering Apparatus	8421	1.9	60.3
Wadding, Gauze, Bandages	3005	1.9	62.2
Household And Hygienic Articles Of Iron Or Steel	7323	1.8	64.0
Butter And Other Milk Fats And Oils	0405	1.8	65.8
Household And Hygienic Articles Of Plastic	3924	1.7	67.5
Pens	9608	1.7	69.2
Prepared Or Preserved Fruit And Nuts	2008	1.7	70.9
Taps, Cocks, Valves For Pipes	8481	1.6	72.4
Bread, Pastry, Cakes, Biscuits	1905	1.5	74.0
Concretes, Cements And Additives	3824	1.5	75.5
Self-Adhesive Items Of Plastics	3919	1.5	76.9
Plates, Sheets, Film Or Strip Of Plastic	3920	1.4	78.3
Padlocks And Locks	8301	1.4	79.6
Jams And Fruit Or Nut Pastes	2007	1.4	81.0
Ball Or Roller Bearings	8482	1.3	82.3
Therapeutic Blood	3002	1.3	83.6
Discs, Tapes And Storage Devices	8523	1.3	84.9
Other Products	.	15.1	100.0

Note: This table presents the names and HSN codes of the main control products. The name is a simplified version of the official HSN name. Additionally, the table displays the share and cumulative share that each product represented among control products, considering the number of observations before November 2017. This table is discussed in Section 3.4.

Table A.8: Description of Related and Unrelated Groups

	Related	Not Related	Control
Products	348	332	680
Firm–HSNs	203,138	112,975	316,113
Of which before Nov 17	65,096	42,790	107,886
Before November 2017:			
Months reported	1.98 (1.14) [1.00]	2.13 (1.18) [2.00]	2.04 (1.16) [2.00]
HSN level reported	6.10 (1.99) [8.00]	6.55 (1.91) [8.00]	6.28 (1.97) [8.00]
Final consumers sales (Firm share)	0.80 (0.16) [0.82]	0.78 (0.17) [0.78]	0.79 (0.16) [0.80]
Sales ('000 INR)	510 (19891) [6.78]	314 (4989) [7.75]	429 (15553) [7.17]
Sales (Log value)	8.95 (2.53) [8.82]	9.08 (2.59) [8.96]	9.00 (2.56) [8.88]
Quantity (Log units)	3.79 (2.65) [3.53]	4.69 (2.78) [4.61]	4.16 (2.74) [3.95]
Price pre-tax (Log value per unit)	5.16 (2.18) [5.01]	4.39 (1.93) [4.10]	4.84 (2.12) [4.57]
Observations	1,015,273	657,640	1,672,913
Of which before Nov 17	128,820	91,287	220,107

Note: This table presents the sample used in the substitution analysis. The table includes only control group products. *Control* products were taxed at 18% or 12% throughout July 2017 to March 2020. The *Related* column includes control products that are in the same HSN-2 code as a treated product. The *Not Related* column includes control products that are not in the same HSN-2 code as a treated product. The table is divided in three groups of rows: Products, Firm–HSN, and Observations. Standard deviations are in parentheses and the median in brackets. Statistics are measured for months before November 2017. This table is discussed in Section 4.2.

Table A.9: Main Treated and Related Product within an HSN-2

HSN-2	Treated product	HSN-4	Related product	HSN-4	Share of treated sales HSN-2	Share of related group
38	Organic Solvents	3814	Insecticides, Rodenticides, Herbicides	3808	5.5	7.2
34	Washing, Cleaning Or Degreasing Preparations	3402	Liquid Or Cream Soap	3401	46.1	3.6
90	Equipment For Photographic	9010	Medical Instruments	9018	8.0	3.4
39	Plastic Articles (Office And School Supplies)	3926	Tubes, Pipes And Hoses Of Plastic	3917	13.7	3.0
96	Lighters	9613	Brooms And Brushes	9603	1.7	2.7
73	Sanitary Ware Of Iron/Steel	7324	Screws, Bolts, Nuts Of Iron/Steel	7318	0.7	2.5
84	Air Vacuum Pumps	8414	Computers	8471	15.5	2.4
85	Electrical Circuits (Switchers, Lamp-Holders, Fuses)	8536	Telephone Sets	8517	17.6	2.0
19	Malt Extract	1901	Bread, Pastry, Cakes, Biscuits	1905	18.5	1.6
40	Rubber Articles (Gaskets, Seals, Bands)	4016	Tubes, Pipes And Hoses Of Rubber	4009	10.6	1.2
83	Mountings	8302	Padlocks And Locks	8301	59.3	1.1
33	Make-Up	3304	Dental Hygiene Preparations	3306	83.1	1.1
48	Wallpaper	4814	Paper And Paperboard Cut	4823	1.0	1.1
68	Worked Monumental Or Building Stone	6802	Abrasive Powder Or Grain	6805	94.7	1.0
17	Sugar Confectionery	1704	Sugars, Sugar Syrups, Artificial Honey, Caramel	1702	0.4	0.9
21	Extracts Of Coffee	2101	Sauces And Mixed Condiments	2103	5.8	0.8
76	Roofs, Doors And Windows	7610	Table, Kitchen Or Household Articles Of Aluminium	7615	0.4	0.8
82	Razors	8212	Hand Tools (Blow Lamps, Clamps, Anvils)	8205	21.6	0.7
32	Painters' Colours	3213	Synthetic Organic Colouring Matter	3204	4.0	0.6
95	Festive Articles	9505	Wheeled Toys And Dolls	9503	0.6	0.5
70	Mirrors	7009	Glassware For Kitchen, Toilet, Office, Decoration	7013	32.8	0.5
74	Household Items Of Copper	7418	Copper Tube Or Pipe Fittings	7412	61.9	0.3
44	Plywood	4412	Wood Sawn Or Chipped Lengthwise	4407	27.6	0.3
91	Watches	9102	Alarm Clocks, Wall Clocks	9105	95.3	0.3
94	Furniture	9403	Medical Furniture	9402	78.8	0.1
36	Ferro-Cerium And Spark Alloys	3606	Safety Fuses, Igniters, Detonators	3603	22.6	0.1
42	Suitcases, Vanity Cases, Briefcases	4202	Clothing Accessories Of Leather	4203	92.3	0.0
69	Ceramic Flags And Paving	6907	Bricks, Blocks, Tiles	6901	81.3	0.0
67	Artificial Flowers	6702	Bird Skins And Feather	6701	21.0	0.0
43	Articles Of Furskin	4303	Raw Furskins	4301	0.0	0.0
25	Granite, Basalt, Sandstone (Building Stone)	2516	Marble, Travertine, Ecaussine (Building Stone)	2515	1.6	0.0

Note: This table presents the names and HSN-4 of the largest treated and related product in each HSN-2 chapter containing a treated product. The largest products are based on the number of observations (firm–HSN–months) before November 2017. The names shown are simplified versions of the HSN names. Column 1 lists the HSN-2 code. Column 2 displays the name of the largest treated product, and column 3 its corresponding HSN code. Columns 4 and 5 present the name and HSN-4 of the largest related product. Column 6 indicates the percentage of pre-November 2017 sales with a rate reduction in the HSN-2. Column 7 presents the percentage of the related group that this related product represents. For example, the main product with a rate reduction in HSN-2 34 was *Washing, Cleaning Or Degreasing Preparations* and the main related product in HSN-2 34 was *Liquid Or Cream Soap*. Column 6 indicates that products with a rate reduction accounted for 46.1% of observations in HSN-2 34 before November 2017, while column 7 indicates that *Liquid Or Cream Soap* accounted for 3.6% of observations in the entire related group before November 2017. This table is discussed in Section 4.2.

Table A.10: Indirect Effects of the Tax Cut

	(1) Main specification	(2) Abv. median exposure	(3) Include 5%	(4) Within firm
ETR	-0.000	-0.001	-0.001	-0.001
	(0.000)	(0.001)	(0.000)	(0.001)
Consumer Price	-0.002	-0.006	-0.014	0.004
	(0.008)	(0.010)	(0.007)	(0.010)
Quantity	0.016	0.035	-0.002	0.038
	(0.021)	(0.025)	(0.018)	(0.026)
Sales	0.014	0.030	-0.015	0.043
	(0.020)	(0.023)	(0.017)	(0.023)
Reporting Rate	-0.002	-0.002	-0.003	-0.009
	(0.002)	(0.003)	(0.002)	(0.005)
N	1,517,273	1,036,135	2,036,713	900,925
N (extensive)	6,551,685	4,490,388	8,548,113	3,501,876

Note: Each column-row cell reports the point estimate from a separate difference-in-differences regression comparing *Related* and *Not Related* control products. *Control* products were taxed at 18% or 12% throughout July 2017 to March 2020. *Related* products includes control products that are in the same HSN-2 code as a treated product. *Not Related* products includes control products that are not in the same HSN-2 code as a treated product. Estimates compare the pre period (July–October 2017) to the post period (April 2018–March 2019). All outcome variables are in logs. Column (1) presents the main specification with no threshold restriction on the treatment group. Column (2) restricts the treatment group to products with an above-median share of treated products in their HSN group. Column (3) uses a broader sample that includes 5% rate products. Column (4) restricts the treatment group to products sold by firms that also sell directly treated products. Standard errors, clustered at the firm and HSN-4 level, are reported in parentheses. This table is discussed in Section 4.2.

Table A.11: Robustness of Baseline Estimates and Elasticities

		Control group					Cleaning Steps		Sales to final consumers					Balance			Aggregation level	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)
	Main specification	Include 5%	Only 18%	Exclude medicaments	Only not related	Only related	No price trimming	Include sales only	Downstream products	Survey products	Supermarket products	B2C > 80%	B2C whole period	Weakly balanced	Strongly balanced	Sales weighted	Firm-Product	Firm-Treatment
ETR	-0.066	-0.066	-0.066	-0.066	-0.065	-0.066	-0.066	-0.066	-0.065	-0.062	-0.066	-0.064	-0.066	-0.067	-0.070	-0.072	-0.067	-0.065
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)
Consumer Price	-0.034	-0.037	-0.036	-0.032	-0.032	-0.033	-0.035	-0.034	-0.029	-0.032	-0.034	-0.026	-0.035	-0.036	-0.030	-0.024	-0.035	-0.041
	(0.006)	(0.006)	(0.006)	(0.006)	(0.009)	(0.006)	(0.007)	(0.006)	(0.009)	(0.009)	(0.007)	(0.007)	(0.006)	(0.006)	(0.012)	(0.014)	(0.007)	(0.015)
Quantity	0.035	0.032	0.028	0.036	0.052	0.027	0.030	0.035	0.056	0.074	0.064	0.029	0.030	0.036	0.064	0.051	0.034	0.046
	(0.017)	(0.017)	(0.018)	(0.017)	(0.027)	(0.017)	(0.017)	(0.017)	(0.025)	(0.025)	(0.021)	(0.018)	(0.018)	(0.018)	(0.024)	(0.035)	(0.019)	(0.021)
Sales	0.067	0.061	0.057	0.070	0.086	0.060	0.061	0.066	0.092	0.104	0.096	0.066	0.061	0.067	0.104	0.100	0.066	0.071
	(0.018)	(0.018)	(0.018)	(0.018)	(0.027)	(0.018)	(0.018)	(0.018)	(0.026)	(0.028)	(0.022)	(0.019)	(0.018)	(0.018)	(0.025)	(0.035)	(0.019)	(0.018)
Reporting Rate	-0.004	-0.005	-0.004	-0.004	-0.006	-0.002	-0.004	-0.004	-0.001	-0.003	-0.005	-0.003	-0.003	-0.013	–	-0.009	-0.003	-0.005
	(0.002)	(0.002)	(0.002)	(0.002)	(0.004)	(0.002)	(0.002)	(0.002)	(0.004)	(0.004)	(0.003)	(0.003)	(0.002)	(0.003)		(0.004)	(0.003)	(0.002)
Consumer Price Elasticity (ε_p)	0.512	0.552	0.548	0.487	0.477	0.498	0.522	0.512	0.436	0.497	0.498	0.414	0.523	0.515	0.422	0.311	0.529	0.617
	(0.092)	(0.093)	(0.092)	(0.090)	(0.149)	(0.090)	(0.107)	(0.092)	(0.145)	(0.153)	(0.114)	(0.113)	(0.092)	(0.092)	(0.171)	(0.181)	(0.114)	(0.223)
Quantity Elasticity (ε_x)	-0.505	-0.461	-0.409	-0.528	-0.774	-0.405	-0.426	-0.505	-0.835	-1.141	-0.920	-0.441	-0.438	-0.484	-0.895	-0.648	-0.498	-0.697
	(0.264)	(0.261)	(0.272)	(0.258)	(0.413)	(0.266)	(0.263)	(0.264)	(0.382)	(0.412)	(0.313)	(0.288)	(0.272)	(0.271)	(0.344)	(0.457)	(0.287)	(0.315)
Sales Elasticity (ε_s)	-0.993	-0.910	-0.861	-1.041	-1.298	-0.907	-0.904	-0.993	-1.400	-1.644	-1.422	-1.027	-0.915	-0.969	-1.474	-1.337	-0.970	-1.080
	(0.267)	(0.263)	(0.273)	(0.264)	(0.404)	(0.270)	(0.268)	(0.268)	(0.395)	(0.440)	(0.321)	(0.293)	(0.275)	(0.273)	(0.352)	(0.458)	(0.287)	(0.278)
Quantity-Price Elasticity (ε_{xp})	-1.131	-0.989	-0.805	-1.207	-2.031	-0.853	-1.001	-1.131	-1.955	-2.390	-2.025	-1.108	-1.008	-1.062	-2.122	-2.223	-1.027	-1.131
	(0.489)	(0.415)	(0.477)	(0.529)	(0.820)	(0.518)	(0.477)	(0.489)	(0.840)	(1.002)	(0.657)	(0.717)	(0.468)	(0.494)	(1.049)	(1.735)	(0.548)	(0.466)
N	2,155,746	2,676,356	1,588,322	1,898,883	1,215,454	1,543,463	2,281,117	2,162,245	1,019,763	727,922	1,902,447	1,165,411	2,080,779	1,195,002	199,689	1,332,849	1,400,701	185,936

Note: Each column-row cell reports the point estimates from a separate outcome and robustness specification. Estimates come from difference-in-differences regressions comparing treated and control products across three periods: pre (July–October 2017, omitted), transition (November 2017–March 2018), and post (April 2018–March 2019, reported). Estimates compare the pre period to the post period. All outcome variables are in logs. Consumer Price, Quantity, and Sales elasticities are estimated with log ETRs instrumented with treatment-by-period interactions. The Quantity–Price elasticity is estimated with log Consumer Price instrumented with treatment-by-period interactions. Column (1) presents the results for the main specification. The Control group columns vary the definition of the control products by: (2) including 5% rate products; (3) restricting to products taxed at 18%; (4) excluding medicaments; (5) restricting to products not in the same HSN-2 of treated products; and (6) restricting to products in the same HSN-2 of treated products. The Cleaning Step columns add back observations removed in cleaning steps: (7) adds observations removed because of extreme price variation; and (8) adds observations that only reported sales and not quantities. The Sales to final consumers columns test alternative definitions of sales to final consumers, by: (9) restricting to downstream products (Antràs et al. 2012); (10) selecting products appearing in the National Sample Survey; (11) restricting to treated products in the supermarket database; (12) restricting to firms for which more than 80% of pre-November 2017 sales were made to final consumers; and (13) restricting to firms for which more than 50% of the whole period sales were made to final consumers. The Balance columns implement sample-balance restrictions, including: (14) weak balance, defined as the firm–product appearing before and after November 2017 at least once; (15) strong balance across the entire period, defined as the firm–product appearing all periods; and (16) weighting by firm–product pre-November 2017 sales with weights winsorized at the 97.5 percentile. The Aggregation level columns aggregate observations at broader definitions: (17) at the firm–product–month level; and (18) at the firm–treatment–month level (treated or control). Standard errors are reported in parentheses. This table is discussed in Section 4.4.

Table A.12: Top Products with a Rate Reduction in July 2018

Product name	Product code (HSN-4)	Treated products (%)	Cumulative share (%)
Glaziers Putty, Caulking, Surface Preparations	3214	16.8	16.8
Heating And Thermic Appliances, Hair-Dressers	8516	16.6	33.5
Paints And Varnishes (Non-Aqueous Dissolved)	3208	11.8	45.3
Electric Accumulators	8507	11.4	56.7
Paints And Varnishes (Aqueous Dissolved)	3209	8.2	64.8
Refrigerators And Freezers	8418	7.1	71.9
Suitcases, Vanity Cases, Briefcases	4202	5.8	77.7
Laundry Washing Machines	8450	5.6	83.3
Food Grinders And Mixers, Juice Extractors	8509	5.6	88.9
Shavers	8510	2.6	91.5
Vacuum Cleaners	8508	2.0	93.5
Paints And Varnishes (Enamels, Lacquers, Distempers)	3210	1.5	95.1
Scent And Toilet Sprays	9616	1.1	96.2
Trailers And Semi-Trailers	8716	0.8	97.0
Framed Mirrors	7009	0.6	97.6
Chains Of Copper	7419	0.5	98.1
Motor Cars For The Transport Of Persons	8703	0.5	98.6
Globes For Lamps And Lanterns	7020	0.4	99.0
Works Trucks	8709	0.3	99.3
Festive Articles	9505	0.2	99.6
Other Articles Of Wood	4421	0.1	99.7
Wooden Frames For Art Or Mirrors	4414	0.1	99.8
Special Purpose Vehicles	8705	0.1	99.9
Carving And Moulding Material And Articles	9602	0.1	100.0
Worked Monumental Or Building Stone	6802	0.0	100.0

Note: This table presents the names and HSN codes of the main products that had a rate reduction in July 2018. The name is a simplified version of the official HSN name. Additionally, the table displays the share and cumulative share that each product represented among treated products, considering the number of observations before July 2018. This table is discussed in Section 4.4.

Table A.13: Summary Statistics for Treated and Control Products (July 2018 Reform)

	Treated	Control	All
Products	25	680	705
Firm–HSNs	19,103	320,005	339,108
Of which before Jul 18	11,070	194,592	205,662
Before July 2018:			
Months reported	3.54 (3.39) [2.00]	3.81 (3.52) [2.00]	3.80 (3.52) [2.00]
HSN level reported	5.84 (1.98) [4.00]	6.20 (1.98) [8.00]	6.18 (1.98) [8.00]
Final consumers sales (Firm share)	0.82 (0.16) [0.87]	0.80 (0.17) [0.83]	0.80 (0.17) [0.83]
Sales ('000 INR)	418 (3310) [17.56]	345 (11994) [6.00]	348 (11713) [6.32]
Sales (Log value)	9.82 (2.48) [9.77]	8.84 (2.54) [8.70]	8.89 (2.54) [8.75]
Quantity (Log units)	2.95 (2.22) [2.71]	4.03 (2.68) [3.81]	3.98 (2.67) [3.74]
Price pre-tax (Log value per unit)	6.86 (1.95) [6.85]	4.80 (2.06) [4.57]	4.91 (2.10) [4.66]
Observations	90,430	1,692,014	1,782,444
Of which before Jul 18	39,197	742,242	781,439

Note: This table presents the sample used in the firm–HSN-level analysis for the July 2018 rate reduction from 28% to 18%. The *Treated* column includes products that experienced this rate reduction in July 2018. The *Control* column includes products that did not undergo any rate change between July 2017 and March 2020 and were taxed at 12% or 18% throughout. The *All* column combines both groups. The table is divided into three groups of rows: Products, Firm–HSN, and Observations. Standard deviations are in parentheses and the median in brackets. Statistics are measured before July 2018. This table is discussed in Section 4.4.

Table A.14: Description of Firm-Level Tax Returns

	Firm Return	+ Detailed Firm Return	+ Firm–Product Return	All
Firms	156,716	243,203	88,335	334,970
Tax returns	313,888	607,071	213,870	1,134,829
	27.7%	53.5%	18.8%	100%
Sales	1,030	2,094	20,951	5,353
	(16,097)	(70,427)	(353,000)	(162,066)
	[213]	[342]	[1,428]	[368]
	5.3%	20.9%	73.8%	100%
Value-Added	767	1,309	15,731	3,877
	(10,456)	(68,392)	(320,450)	(148,046)
	[78]	[100]	[520]	[120]
	5.5%	18.1%	76.5%	100%
Taxes	58	217	2,468	597
	(1,990)	(2,850)	(53,414)	(23,323)
	[13]	[37]	[152]	[34]
	2.7%	19.4%	77.9%	100%
Products registered	2.74	2.77	3.18	2.84
	(1.56)	(1.58)	(1.60)	(1.59)
	[2.00]	[2.00]	[3.00]	[3.00]
	26.7%	52.2%	21.1%	100%

Note: This table divides firm-level tax returns into three mutually exclusive groups: (i) simplified tax return (ii) detailed firm-level tax return, and (iii) firm–product tax return. The All column pools all tax returns. Firms may file more than one tax return type over the entire period. For each variable, the table reports the mean, standard deviation in parentheses, and median in brackets. The last line of each block reports the group shares for that variable considering all periods. Sales and tax amounts are reported in thousand INR. Value added is defined as total sales minus intra-state inputs. This table is discussed in Section 4.5.

Table A.15: Description of Firm-Level Treatment and Control Group

	Treated	Control	All
Firms	74,902	146,615	221,517
Tax returns	1,007,651	2,088,209	3,095,860
with details	715,791	1,472,637	2,188,428
with firm–product	207,716	388,932	596,648
before Nov 2017	172,991	378,208	551,199
Before November 2017:			
Sales	3,641 (61,917) [294]	4,345 (147,676) [321]	4,124 (127,150) [312]
Value-Added	2,608 (57,361) [82]	3,134 (125,001) [113]	2,969 (108,416) [103]
Taxes	491 (5,629) [50]	453 (16,430) [34]	465 (13,970) [38]
Products registered	3.01 (1.56) [3.00]	2.26 (1.49) [2.00]	2.49 (1.55) [2.00]

Note: This table presents the treated and control groups used in the firm-level administrative tax data results. Treated firms are those that, in the registration form, report selling at least one treated product. The sample excludes firms that report in the registration form selling exempt products, products taxed at 5%, non-treated products taxed at 28%, services, or products treated in later periods. For each variable, the table reports the mean, standard deviation in parentheses and median in brackets. Value added is defined as total sales minus intra-state inputs. Sales, value-added, and tax amounts are reported in thousand INR. This table is discussed in Section 4.5.

Table A.16: Comparison of Firm and Firm–Product Level Results

	(1) Firm Return Simplified	(2) Firm Return Detailed	(3) Firm–Product Return	(4) Firm–Product Share Treated
ETR	-0.017	-0.018	-0.023	-0.064
	(0.002)	(0.002)	(0.002)	(0.002)
Sales	0.050	0.037	0.027	0.043
	(0.016)	(0.014)	(0.018)	(0.020)
Sales Elasticity (ϵ_s)	-2.609	-1.892	-1.157	-0.618
	(0.702)	(0.610)	(0.600)	(0.249)
Value-Added	0.056	0.060	0.045	0.112
	(0.025)	(0.029)	(0.024)	(0.028)
Value-Added Elasticity (ϵ_{va})	-2.983	-2.971	-1.668	-1.403
	(1.184)	(1.334)	(0.814)	(0.338)
N	883,217	1,568,696	582,565	1,014,684
N value-added	694,448	1,206,335	491,811	854,015

Note: This table compares firm-level estimates across reporting forms: (i) simplified tax return (ii) detailed firm-level tax return, and (iii) firm–product tax return. Columns (1)–(3) define treatment using the registration form, where firms are considered treated if they list at least one treated product. Column (4) uses a continuous treatment defined as the share of treated sales observed in the firm–product tax returns before November 2017. Value added is defined as the monthly difference between sales and intra-state purchases. Standard errors are reported in parentheses. This table is discussed in Section 4.5.

Table A.17: Barcodes by Department & Treatment Status (Hypermarket Data)

Department	Treated barcodes	Control barcodes	All
Personal Care & Cosmetics	1,021	528	1,549
Home Care	225	169	394
Breakfast	65	118	183
DIY & Auto Accessories	47	31	78
Leisure	36	392	428
Processed Food	87	836	923
Home	12	10	22
Luggage	5	0	5
Home Needs & Home Upkeep	9	405	414
Small White Appliance	1	2	3
Total	1,508	2,491	3,999

Note: This table reports the number of treated and control barcodes in the hypermarket sample, grouped by store department as defined by the retailer. Treated barcodes correspond to barcodes with a rate reduction from 28% to 18% in November 2017, while control barcodes belong to unaffected categories taxed at 12% or 18%. This table is discussed in Section 5.

Table A.18: Treated Barcodes by Barcode "Family" (Hypermarket Data)

Barcode "Family"	Barcodes (Count)	Of Treated (%)	Cumulative (%)
Hair Care	291	19.3	19.3
Deodorants and Talcs	256	17.0	36.3
Skin Care	237	15.7	52.0
Housekeeping	136	9.0	61.0
Laundry	89	5.9	66.9
Shaving	89	5.9	72.8
Confectionery	64	4.2	77.1
Personal Wash	54	3.6	80.6
Baby	45	3.0	83.6
Health Mix	36	2.4	86.0
Electricity	26	1.7	87.7
Coffee	24	1.6	89.3
Other families	162	10.7	100.0
Total	1,508	100.00	100.0

Note: This table lists the barcode "Families" with treated items in the hypermarket sample. Barcode Families follow the retailer's internal product classification. "Barcodes (Count)" reports the number of distinct treated barcodes within each Family. "Of Treated (%)" gives the share of all treated barcodes accounted for by that Family, and the cumulative share is reported in the last column. This table is discussed in Section 5.

Table A.19: Taxes of Luxury & Non-Essential Goods, 10 Largest Emerging Economies

Country	Tax Instruments	Illustrative Examples	Coverage
China	VAT differentiation	Luxury vehicles, yachts, private aircraft, jewelry, watches, cosmetics	Broad non-essential
India	GST differentiation + Cess	SUVs, air conditioners, consumer durables, premium motorcycles, aerated beverages	Broad non-essential
Brazil	Excises (IPI) and selective taxes	Automobiles, perfumes, cosmetics, selected electronics	Broad non-essential
Mexico	VAT differentiation and IEPS excises	-	Limited luxury taxation
Indonesia	Luxury Goods Sales Tax (PPnBM)	Luxury vehicles, yachts, private aircraft, luxury housing and apartments	Narrow luxury
Türkiye	Special Consumption Tax (ÖTV)	Passenger vehicles, luxury vehicles, selected luxury imports	Broad non-essential
Argentina	Internal Taxes (*Impuestos Internos*)	Luxury vehicles, yachts, private aircraft	Narrow luxury
South Africa	Ad valorem excises	Luxury vehicles, cosmetics, perfumes, electronics, watches	Broad non-essential
Vietnam	Special Consumption Tax	Automobiles, motorcycles, yachts, private aircraft	Narrow luxury
Thailand	Excise Tax System	Luxury vehicles, yachts, selected luxury imports	Narrow luxury

Note: This table describes the taxation of luxury and non-essential consumption based on tax guides in the 10 largest emerging economies, ranked by GDP. Country-level tax systems were analyzed using Claude Sonnet 4.6 (Anthropic, 2026), informed by PwC Worldwide Tax Summaries, country chapters for China, India, Brazil, Mexico, Indonesia, Türkiye, Argentina, South Africa, Vietnam, and Thailand; the Deloitte International Tax Source; KPMG Tax News Flash; and country-specific official sources, including the GST Council of India (56th Meeting, September 2025), the Indonesia Directorate General of Taxes (2025), and the South African Revenue Service Ad Valorem Excise Duty Schedule. Examples are illustrative rather than exhaustive. Coverage classifications are based on the types of goods covered rather than tax rates. *Narrow luxury* systems primarily target goods consumed by a small share of households (e.g., yachts, private aircraft, luxury housing, or high-end vehicles). *Broad non-essential* systems extend to a wider set of goods disproportionately consumed by higher-income households, including consumer durables and selected services. *Limited luxury taxation* refers to systems in which differentiation primarily occurs through VAT relief for necessities or excises on goods associated with externalities rather than through explicit taxation of luxury consumption. Tax systems and coverage reflect information current as of mid-2026. This table is discussed in Section 7.

Appendix B: Baseline STR–ETR Gap

Under the GST reporting system firms manually report their tax liabilities, and HSN codes are not automatically mapped to statutory rates, as discussed in Section 3. This leaves scope for a wedge between the Statutory Tax Rate (STR) and the Effective Tax Rate (ETR) of firm-products. By definition, the STR of treated products fell from 28 percent to 18 percent. Without imposing any restriction on feasible rates, the average ETR of treated products prior to the reform was 25.3 percent, implying a 2.7 percentage point gap relative to the STR (Figure B.2A). After the reform, the ETR converged to 18 percent. We consider three non-mutually exclusive explanations for the baseline wedge between STRs and ETRs among treated products: limited observability in the data, inattention, and strategic misreporting.

Inattention Inattention may explain a third of the baseline wedge. We estimate its contribution using errors that are unfavorable to firms. Figure B.1A shows that, without restricting to feasible rates, 4 percent of treated firm-products continued to remit the pre-reform 28 percent rate after April 2018, five months after the rate cut, consistent with VAT reporting mistakes documented in other countries (e.g., Almunia et al. 2023). Conversely, about half of the lower-rate reporting reflects firms applying 0, 5, or 12 percent, and these mistakes persist with similar frequency after the reform.

Panels B and C of Figure B.1 show that these mistakes also occur among control firm-products, particularly before November 2017, when more goods faced the 28 percent rate. Moreover, errors are symmetric between favorable and unfavorable mistakes for 12 percent control products, which lie in the middle of the tax schedule. Because 28 percent is the top rate, however, unfavorable errors for treated products before the reform cannot be offset by equally large errors in the opposite direction. Assuming the treated group has the same share of inattentive firms as the control group, inattention alone explains a 0.9 pp wedge if inattentive firms were equally likely to select any existing tax rate (0%, 5%, 12%, or 18%).

Limited Observability Limited observability may explain another third of the initial gap. About 20 percent of treated HSN-4 codes contain lower-rated HSN-6 subcodes. Because some firms report only at the HSN-4 level, the econometrician cannot observe the applicable statutory rate within these heterogeneous codes, creating measurement error. To quantify its importance, we use firm-products reporting at the HSN-8 level to recover the within-HSN-4 sales distribution across subcodes and reweight the statutory rate for HSN-4 reporters accordingly. We apply this adjustment to the 34 treated products with lower-rated HSN-6 subcodes and to eight products—mostly medical instruments in Chapter 90—whose applicable rate depends on end use according to GST notifications (Central Board of Indirect Taxes and Customs 2019). The adjustment reduces the gap by 0.9 percentage points (Figure B.2B).

Consistent with both inattention and limited observability, firms selling only to other firms—where strategic misreporting is less attractive as buyers claim input tax credits—still exhibit a 2.4 percentage point gap (Figure B.2C). Together, the evidence suggests that non-strategic reasons account for about two-thirds of the gap.

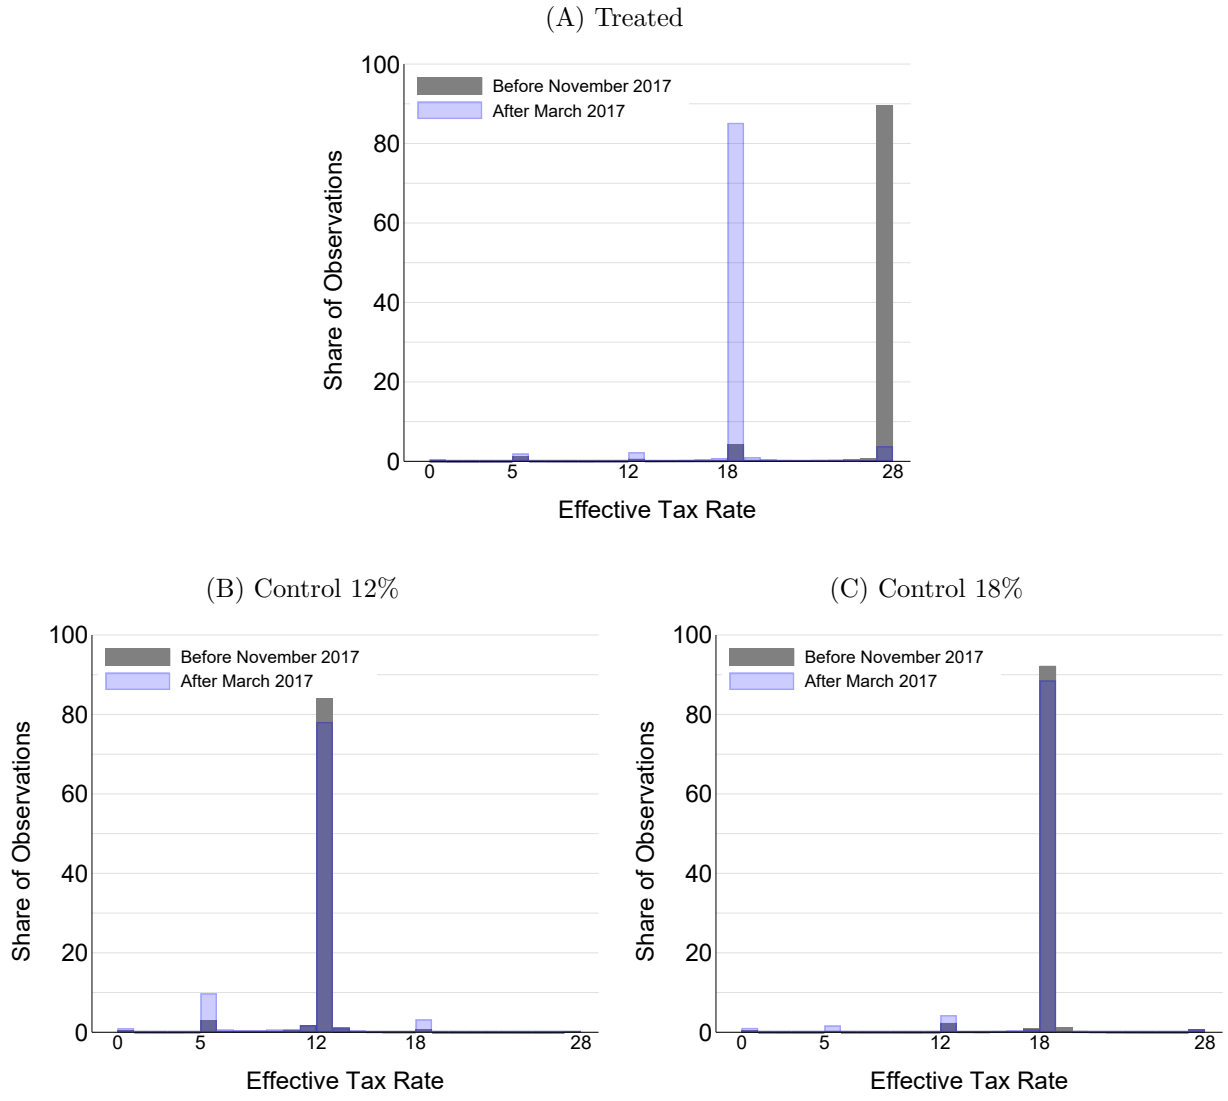
Strategic Misreporting The residual gap may reflect strategic behavior intended to minimize tax liabilities. Table B.1 compares treated firm-products with an initial ETR of at least 27% (compliers) to those with a lower initial ETR (non-compliers). The two groups are broadly similar in terms of sales, reporting duration, location, and the share of sales to final consumers. The main difference is reporting precision, consistent with more accurate product classification among compliers and with the limited observability analysis.

To further investigate the strategic misreporting, Table B.2 regresses the STR–ETR gap on proxies for strategic behavior. Column 1 shows that larger firm-products exhibit smaller gaps, with a 10 percent increase in sales associated with a 0.013pp reduction, while a 10pp higher share of sales to final consumers is associated with a 0.213pp larger gap. More granular reporting and longer reporting histories are also associated with smaller gaps, whereas location is insignificant. The signs of these coefficients are consistent with strategic behavior, although

they could also reflect inattention. These proxies nonetheless explain little of the overall variation, with an R^2 of just 0.059. Column 2 adds HSN fixed effects. The coefficients remain statistically significant, while the R^2 rises to 0.29, indicating that a substantial share of the variation in the STR–ETR gap reflects differences across HSNs. Column 3 includes firm fixed effects and product fixed effects. The R^2 rises further to 0.49, indicating that differences across firms explain an additional 20 percentage points of the variation in the STR–ETR gap beyond differences across products. Even under the conservative upper-bound assumption that all of this incremental explanatory power reflects strategic behavior, firm-level strategic reporting accounts for at most 20 percent of the overall variation in the STR–ETR gap. Overall, the evidence suggests that strategic behavior may contribute to the STR–ETR gap, but its quantitative importance appears limited.

Combining the evidence on inattention, observability, B2B firms, and STR-ETR gap predictors, we conclude that the gap is mainly non-strategic. Tax motivated strategic behavior appears to account for no more than a third.

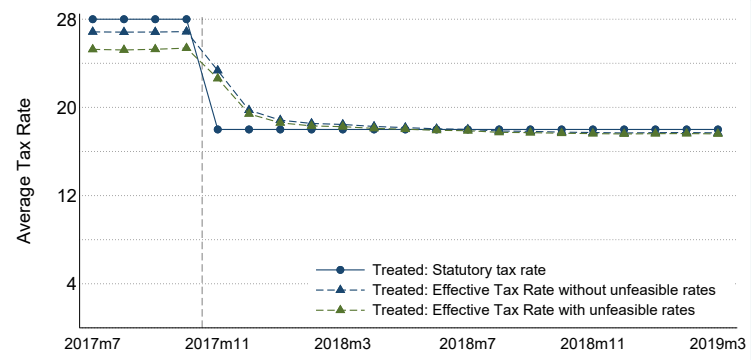
Figure B.1: Product Effective Tax Rates Before and After the November 2017 Reform



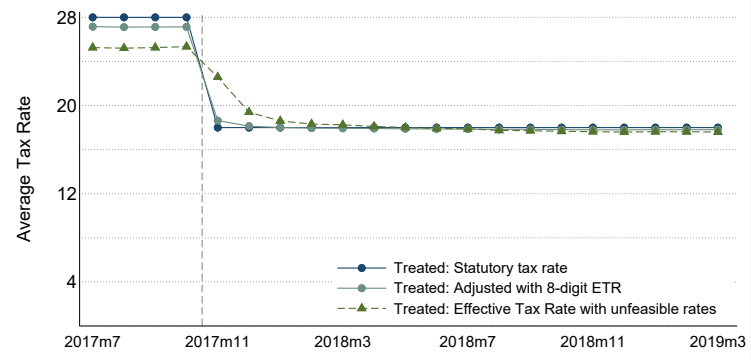
Note: This figure presents the distribution of Effective Tax Rates (ETR) before and after November 2017 rate reduction. Treated products are those that experienced a reduction in STR from 28% to 18% in November 2017. ETRs are winsorized at the 1st and 99th percentiles and rounded, such that the 18% bin includes values between 17.5 and 18.5. An observation is a unique firm-HSN-month. This figure is discussed in Appendix B.

Figure B.2: Average Corrected Statutory Tax Rates and Effective Tax Rates

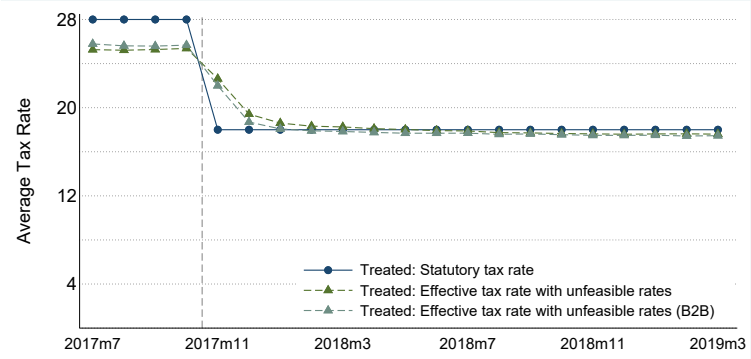
(A) Infeasible effective tax rates



(B) Corrected Statutory Tax Rates



(C) B2B firms



Note: This figure presents the average Statutory Tax Rate (STR) and the average Effective Tax Rate (ETR) for treated products. This figure is discussed in Appendix B.

Table B.1: Statistics of ETR Compliers and Non-Compliers

	Initial ETR above 27%	Initial ETR below 27%	All
Products	161	158	163
Firm–HSNs	40,891	16,071	56,962
Before November 2017:			
Months reported	2.04 (1.17) [2.00]	2.13 (1.17) [2.00]	2.06 (1.17) [2.00]
HSN level reported	6.38 (1.95) [8.00]	5.70 (1.97) [4.00]	6.18 (1.98) [8.00]
Sales ('000 INR)	197 (1933) [7.82]	269 (3110) [5.95]	218 (2338) [7.22]
Sales (Log value)	9.05 (2.42) [8.96]	8.80 (2.47) [8.69]	8.98 (2.44) [8.88]
Quantity (Log units)	3.89 (2.47) [3.53]	3.81 (2.34) [3.47]	3.87 (2.43) [3.53]
Price pre-tax (Log value per unit)	5.32 (1.97) [5.08]	5.15 (1.93) [4.86]	5.27 (1.96) [5.01]
Firms	6,934	4,975	8,342
Firm Sales ('000 INR)	4081 (19053) [332]	4495 (22463) [227]	4202 (20106) [294]
Final consumers sales (Firm share)	0.81 (0.15) [0.83]	0.83 (0.16) [0.88]	0.81 (0.16) [0.84]
Bangalore Urban (Firm location)	0.48 (0.50) [0.00]	0.50 (0.50) [1.00]	0.48 (0.50) [0.00]
Observations	83,352	34,227	117,579

Note: This table presents descriptive statistics for treated products based on the match between Statutory Tax Rates (STRs) and Effective Tax Rates (ETRs) before November 2017. All statistics are based on tax returns before November 2017. Treated products are those that experienced a reduction in STR from 28% to 18%. The column *ETR matches STR* includes treated firm–HSNs with a pre-November 2017 average ETR no lower than 27%. The column *ETR below 28%* includes treated firm–HSN with a pre-November 2017 ETR lower than 27%. Standard deviations are in parentheses and the median in brackets. This table is discussed in Appendix B.

Table B.2: Determinants of the STR–ETR Initial Gap

	STR–ETR Initial Gap		
	(1)	(2)	(3)
Sales (log)	-0.137	-0.108	-0.041
	(0.052)	(0.040)	(0.033)
No. periods reported	-0.160	-0.132	-0.109
	(0.065)	(0.045)	(0.037)
Share of non-ITC sales	2.133	1.847	-
	(0.615)	(0.484)	
Bangalore urban	0.237	0.088	-
	(0.199)	(0.168)	
Lower-rated product	2.017	-	-
	(0.434)		
Reported at HSN-8	-1.678	-	-
	(0.436)		
Product FE	No	Yes	Yes
Firm FE	No	No	Yes
Observations	117,336	116,817	115,180
R^2	0.059	0.286	0.491

Note: The table reports coefficients from regressions of the STR–ETR gap in percentage points on firm–HSN and firm characteristics. Column (1) includes no fixed effects, Column (2) includes HSN FE, and Column (3) includes Firm and HSN FE. The sample is treated products in the months before November 2017 and an observation is a firm–HSN–month. *Sales (log)* measure monthly sales at the firm–HSN level. *No. periods reported* is the number of months the firm–HSN was reported before November 2017. *Share of non-ITC sales* is the share of firm sales before November 2017 to buyers not claiming input tax credits. *Bangalore urban* equals 1 if the firm is located in Bangalore Urban pincode. *Lower-rated product* indicates if the HSN-4 category contains rates lower than 28% at GST onset and HSN-4s with classification disputes in Chapter 90. *Reported at HSN-8* indicates reporting at HSN-8 granularity. Standard errors are reported in parentheses. This table is discussed in Appendix B.

Appendix C: ML Matching of Related Products

The analysis in Section 4.2 finds no evidence of economically meaningful misreporting or substitution at the HSN-4 level—the level at which GST rate differentiation is implemented in India. However, because it compares treated products with all untreated products within an HSN-2 chapter, it may miss spillovers occurring only among the closest substitutes and complements within narrower product classifications (e.g., HSN-6). This appendix therefore probes these finer margins using a large language model (GPT-4) to match each treated product to its closest control substitutes and complements, including at the HSN-6 level for firms reporting at that level. We then estimate spillover effects separately for control products classified as substitutes and complements by comparing their post-reform outcomes with a fixed set of control products that are classified as neutral. We implement this procedure on the administrative tax panel and then replicate it using the supermarket scanner data.

Candidate Pairs We collapse the tax panel to one record per HSN code (HSN-6 when firms report it, otherwise HSN-4) and attach HSN descriptions. The resulting sample contains 4,953 unique HSNs, of which 811 are originally treated and 4,142 are controls. Crossing each treated HSN with every control HSN within the same HSN-2 chapter yields 231,030 candidate pairs for classification.

Matching Procedures We consider three increasingly stringent matching approaches—*broad matching*, *best-pair matching*, and *best-pair high-confidence-threshold matching*—that trade sample size for greater confidence that matched HSNs are genuinely related. This restricts candidate substitutes and complements to increasingly comparable goods while allowing matching within relatively detailed categories. The number of unique HSNs in each specification is reported in the corresponding columns of Table C.1.

Under Broad matching, we classify each candidate pair using a GPT-4 family model accessed through OpenAI's Chat Completions API. The model assigns one of three labels—*substitute*, *complement*, or *neutral*—according to whether the HSN serve similar uses, are commonly consumed together, or have no obvious economic relationship. We refer to this as the *broad matching* approach because each treated HSN can be matched to multiple substitutes and complements.

Under Best-pair matching, the model is shown a treated HSN together with all eligible control HSNs within the same HSN-2 chapter and asked to identify the single most likely substitute and complement. For each match, it also reports a confidence score reflecting the strength of the relationship. If the score falls below a pre-specified threshold, no substitute or complement is assigned.

The Best-pair, high-confidence-threshold matching approach further restricts the sample to substitute and complement assignments with model-reported confidence scores of at least 85 (on a 0–100 scale).

Estimation To estimate spillover effects, we first remove treated HSNs from each matched dataset. We then estimate separate regressions for control HSNs classified as substitutes and complements, comparing their post-reform outcomes with those of the control HSNs classified as neutral. Each regression uses the same event-study difference-in-differences specification as Equation (1).

We then turn to the hypermarket scanner panel, following the same procedure. We generate candidate pairs within each store department and classify them using the same LLM pipeline. Because barcode product descriptions are often short and abbreviated, we augment the matching approach by enriching each barcode description with information retrieved via the Serper API before classification (this is applied only to the hypermarket data, as the HSN descriptions in the tax panel are complete).

Results Figure C.1 reports event-study estimates of quantity responses in the tax panel under the three matching approaches, separately for products identified as substitutes and complements. Table C.1 summarizes the corresponding quantity and price estimates for both data sources. While the estimated coefficients generally have the expected sign (negative for substitutes and positive for complements), most are small and statistically insignificant. The only coefficient significant at conventional levels is the positive quantity response for complements under best-pair matching in the tax data. However, this result is not robust: the estimate shrinks and becomes statistically insignificant under the high-confidence-threshold matching approach. Taken together, the three approaches provide little evidence of economically meaningful substitution or complementarity in households' responses to the GST rate reductions, a conclusion further supported by the supermarket panel.

Interpretation and Relation to the Literature Taken together, these results provide little evidence of economically meaningful substitution by households or product relabeling by firms among closely related products.

This conclusion is consistent with the broader empirical literature, which generally finds that cross-product spillovers following targeted tax reforms are modest and largely confined to narrowly defined product clusters. In food and beverage markets, taxes primarily reduce purchases of the taxed product with limited substitution toward untaxed alternatives: the Mexican tax on sugar-sweetened beverages shifts demand toward water and other

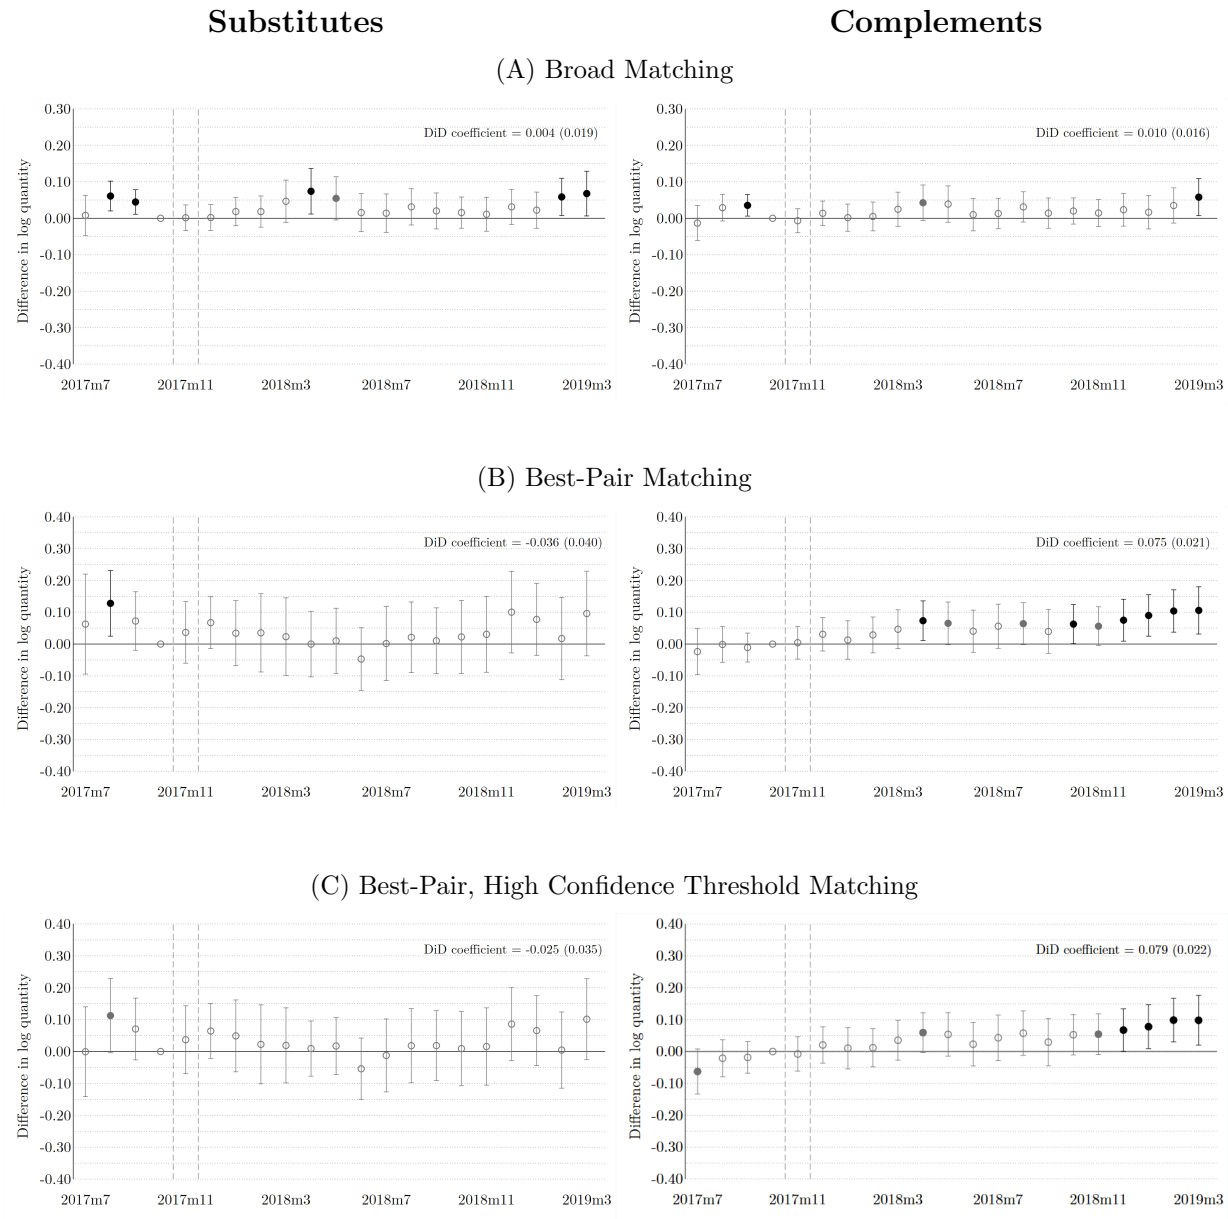
untaxed drinks (Colchero et al. 2016); the Philadelphia beverage tax produces mainly spatial substitution with limited cross-beverage substitution (Cawley et al. 2019); and a virtual-supermarket experiment in New Zealand finds modest within-category substitution for dairy, meats, and snacks (Hoenink et al. 2021). Cleaner substitution is documented in narrow nicotine markets (Cotti et al. 2022) and between alcohol and tobacco products (Decker and Schwartz 2000). Similarly, Asatryan and Gomtsyan (2023) find within-meat substitution following a VAT enforcement reform in Armenia, while Benzarti et al. (2024) detect no meaningful substitution following a temporary VAT cut in Argentina.

Table C.1: Treatment Effects in Tax and Supermarket Data

	Substitution Effects				Complementarity Effects			
			# HSNs				# Barcodes	
Approach	Quantity	Price	Indirectly Treated	Control	Quantity	Price	Indirectly Treated	Control
Panel A: Tax Data								
Broad Matching	0.004	-0.004	1780	2607	0.010	-0.005	2536	2607
	(0.019)	(0.008)			(0.016)	(0.007)		
Best-Pair Matching	-0.036	-0.019	116	2607	0.075***	0.000	474	2607
	(0.040)	(0.023)			(0.021)	(0.010)		
Best-Pair, High Confidence Threshold Matching	-0.025	-0.012	113	2607	0.079***	-0.003	383	2607
	(0.035)	(0.025)			(0.022)	(0.010)		
Panel B: Supermarket Data								
Broad Matching	-0.045	-0.004	1215	530	-0.028	-0.001	1560	530
	(0.036)	(0.008)			(0.036)	(0.008)		
Best-Pair Matching	-0.015	0.006	19	530	0.120	-0.002	25	530
	(0.087)	(0.021)			(0.094)	(0.020)		
Best-Pair, High Confidence Threshold Matching	-0.027	0.033	27	530	-0.015	0.001	24	530
	(0.122)	(0.021)			(0.075)	(0.024)		

Note: This table reports pooled difference-in-differences treatment-effect estimates comparing control products classified as substitutes or complements to control products classified as neutral. Panel A uses the administrative tax (HSN) panel; Panel B uses the hypermarket scanner panel. Under *# HSN* and *# Barcode* headings the two sub-columns report, respectively, the number of unique indirectly treated HSNs (HSN-6 where valid, otherwise HSN-4) entering the regression – that is, HSNs classified as substitutes (for the substitution columns) or complements (for the complementarity columns) of treated products – and the number of unique neutral HSNs and barcodes. The control group is held fixed across the three approaches in each panel and consists of control products classified exclusively as neutrals under the broad-matching classification; it is the same set of products on the substitution and complementarity sides within a given approach. The regressions in Panel A follows Equation (1) The regressions in Panel B follows Equation (3). The three approaches in Panel A are broad matching, best-pair matching, and best-pair, high confidence threshold matching; in Panel B the best-pair matching step is improved with the Serper description-enrichment procedure, as described in the text. Quantity and price are expressed in logs. Significance levels: $^{***}\,p < 0.01$, $^{**}\,p < 0.05$, $^{*}\,p < 0.10$.

Figure C.1: Tax Data: Quantity Response by Classification Approach



Notes: This figure plots event-study coefficients of the quantity response in the administrative tax (HSN) panel, estimated separately under each of the three classification approaches discussed in the text. Panel A uses broad matching, Panel B best-pair matching, and Panel C best-pair, high confidence threshold matching. Within each panel, the left subplot reports estimates for products identified as substitutes of the originally treated products, and the right subplot reports estimates for products identified as complements. All specifications use the flexible treatment definition and an unrestricted control group of products classified as neutrals. The regressions include firm × product × quantity-unit fixed effects, firm × month fixed effects, and standard errors two-way clustered at the firm level and at the HSN-4 level. Coefficients are reported in event time around the GST rate reduction; vertical bands are 95% confidence intervals.

Appendix D: Heterogeneity by Engel Curves

An important determinant of the equity–efficiency trade-off in indirect taxation is the interaction between demand elasticities and the income distribution of consumption. Taxing goods disproportionately consumed by higher-income households is more attractive when demand for those goods is relatively inelastic, whereas high elasticities increase the associated efficiency costs.

Interaction with Income Elasticity. Using the sample of products matched to the household expenditure survey, we test whether treatment effects vary with products' income elasticities by interacting the treatment indicator with each product's income elasticity in the elasticity specification. Formally, we estimate:

$$y_{iht} = \beta_1 \widehat{\text{ETR}}_{iht} + \beta_2 \left(\widehat{\text{ETR}}_{iht} \times \hat{\beta}_j^{\text{Engel}} \right) + \alpha_{ih} + \gamma_{it} + \epsilon_{iht} \tag{6}$$

Where y_{iht} denotes the log outcome for firm i, HSN h in period t. α_{ih} and γ_{it} are firm-HSN and firm-time fixed effects.

ETR_{iht} is the log effective tax rate, and $\hat{\beta}_j^{\text{Engel}}$ is the product-level income elasticity. Because ETR_{iht} and its interaction with income elasticity are endogenous, they are instrumented using the triple interaction between product income elasticity, the treatment group indicator (treat_h), and the post-period indicator (post_t). We cluster standard errors at the firm and HSN-4 product level.[38]

The interaction coefficient captures how tax elasticities vary with product income elasticity. Table D.1 shows that it is consistently small and statistically insignificant for the elasticities of sales, quantities, and consumer prices with respect to effective tax rates, as well as for the elasticity of quantities with respect to consumer prices.

Meta-Regression The second approach uses a meta-regression. We first estimate Engel curve slopes for each survey product. Treated products span a wide range of income elasticities, with an average of 1.4, a 25th percentile below one, and a 75th percentile above three (Figure D.1). We then estimate product-specific quantity and price elasticities with respect to effective tax rates in the administrative data and regress these estimates on the corresponding Engel curve slopes.

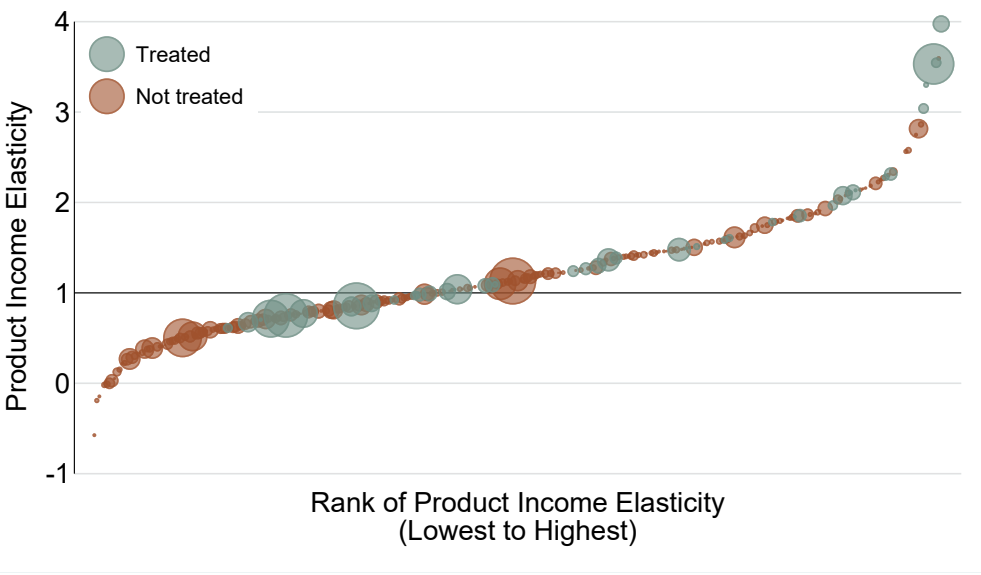
[38] For four-digit HSN codes that can be directly matched to survey products, we use the corresponding income elasticity. For codes without a direct match, we assign the income elasticity of the closest survey product within the same two-digit HSN category.

$$\hat{\beta}_j^{\text{elasticity}} = \alpha + \delta \hat{\beta}_j^{\text{Engel}} + \epsilon_j \tag{7}$$

We plot estimated quantity and price elasticities with respect to the ETR against each product's Engel-curve slope (Figure D.2). Panel A reports quantity elasticities and Panel B price elasticities. Both relationships are essentially flat, suggesting that behavioral responses vary little with product luxuriousness.

Table D.2 confirms these findings under alternative weighting schemes (survey consumption shares, equal product weights, tax-data observations, and tax-data sales shares). None of the interaction coefficients is statistically significant, except in the survey-weighted specification, where the coefficient has the opposite sign to the theoretical prediction: more luxurious products exhibit significantly less elastic quantity responses. Point estimates range from -0.20 to 0.58 for quantity elasticities and from -0.10 to 0.09 for the corresponding implied price elasticities.

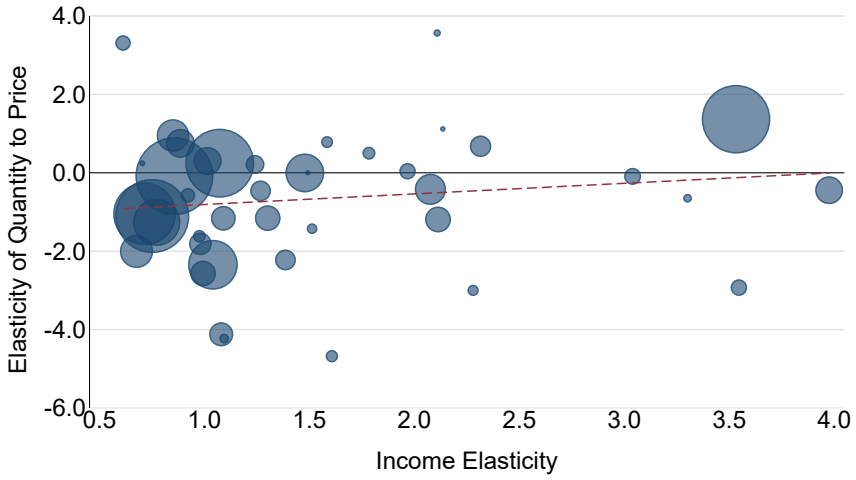
Figure D.1: Product Income Elasticity by Treatment Status



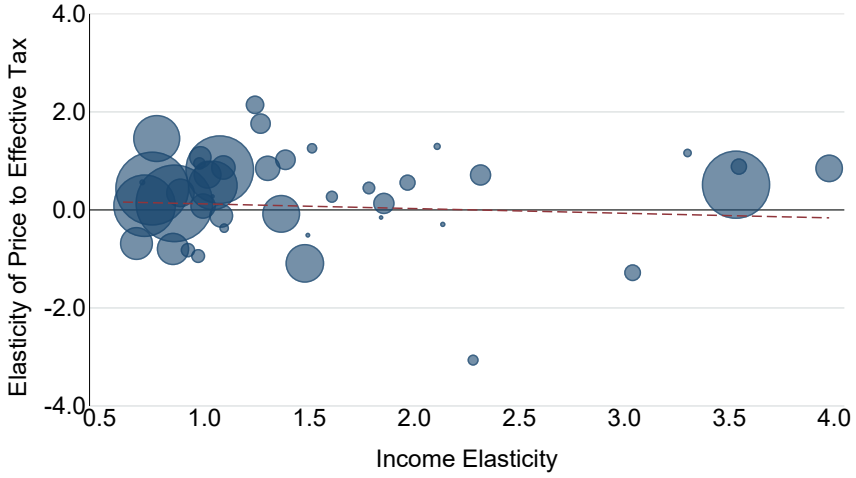
Note: This figure presents Engel beta coefficient of survey products. Treated products are those with a rate reduction from 28% to 18% in November 2017. The circle size increases with product consumption in the survey. This figure is discussed in Appendix D.

Figure D.2: Correlation Between Product Elasticity and Engel-Curve Slope

(A) Quantity Elasticity



(B) Consumer Price Elasticity



Note: This figure plots the correlation between product quantity and consumer price elasticities with respect to effective tax rates and the Engel-curve slope. Each circle represents a survey product, and circle size increases with the product's total consumption in the survey. The red dashed line reproduces the linear fit from the survey equal-weight regression as reported in Table D.2. Each panel excludes the two survey products with the lowest and the two survey products with the highest values of the elasticity for readability. This figure is discussed in Appendix D.

Table D.1: Heterogeneous Treatment Effects by Product Income Elasticity

	(1) Quantity on ETR	(2) Sales on ETR	(3) Consumer Price on ETR	(4) Quantity on Consumer Price
ETR/Consumer Price Elasticity	-0.538	-1.165	0.373	-1.484
	(0.469)	(0.442)	(0.154)	(0.833)
Interaction with income elasticity	0.051	0.163	0.111	0.217
	(0.267)	(0.238)	(0.095)	(0.353)
N	2,177,469	2,177,469	2,177,469	2,177,469
Mean income elasticity (treated)	1.434	1.434	1.434	1.434
Mean income elasticity (control)	1.411	1.411	1.411	1.411

Note: This table reports IV estimates from the continuum specification interacting the treatment with product-level income elasticity. For columns (1)–(3), the effective tax rate (ETR) and its interaction with income elasticity are instrumented by the treatment interactions. For column (4), the consumer price (inclusive of taxes) and its interaction with income elasticity are instrumented. Income elasticity is measured from Engel curves, using the HSN-2 classification when the HSN-4 measure is unavailable. All regressions include firm–product–unit and firm–period fixed effects, with standard errors clustered by firm and HSN-4 product code. This table is discussed in Appendix D.

Table D.2: Regression of Quantity and Price Elasticities on Product Income Elasticity

Weights:	Equal weight (Survey)	Consumption (Survey)	N Observations (Tax Returns)	Product sales (Tax Returns)
Dependent variable: Elasticity of quantity to consumer price				
Income elasticity	0.270	0.584	0.446	-0.200
	(0.355)	(0.179)	(0.316)	(0.298)
Dependent variable: Elasticity of consumer price to effective taxes				
Income elasticity	-0.096	0.033	0.019	0.087
	(0.375)	(0.145)	(0.123)	(0.110)
Observations (Survey Products)	44	44	44	44

Note: This table reports estimates from regressions of survey-product elasticities on the corresponding Engel-curve coefficients, as specified in equation (7). Each column corresponds to a different weighting scheme: equal weighting of survey products, weighting by survey consumption shares, weighting by the number of matched HSN-level observations in firm–product tax returns prior to November 2017, and weighting by pre-tax sales shares of the corresponding HSN codes over the same period. For each elasticity measure, we exclude the survey products with the highest and lowest estimated elasticities from the corresponding regression. Standard errors are reported in parentheses. The sample includes 44 survey products treated in November 2017. This table is discussed in Appendix D.

Appendix E: Elasticity Estimation

To link our reduced-form event-study estimates to the underlying demand elasticities, we set up an equilibrium model of consumer demand. Let x_{kjs} denote the quantity of product j sold by firm k in period s. The tax-inclusive consumer price, p_{kjs}, is a function of the effective tax rate, t^{etr}_{kjs}, and the pre-tax producer price, q_{kjs}, such that $p_{kjs} = (1 + t^{etr}_{kjs})q_{kjs}$. We specify a general demand function $f(\cdot)$ that depends on this tax-inclusive consumer price, accompanied by a demand shifter A_{kjs}:

$$x_{kjs} = A_{kjs} f\left((1 + t^{etr}_{kjs})q_{kjs}\right)$$

Taking the natural logarithm yields:

$$\ln(x_{kjs}) = \ln(A_{kjs}) + \ln f\left((1 + t^{etr}_{kjs})q_{kjs}\right)$$

In our difference-in-differences framework, we model the unobserved demand shifter $\ln(A_{kjs})$ using a combination of firm-product and firm-time fixed effects plus an error term, such that $\ln(A_{kjs}) = \alpha_{kj} + \gamma_{ks} + \epsilon_{kjs}$.

$$\ln(x_{kjs}) = \alpha_{kj} + \gamma_{ks} + \ln f\left((1 + t^{etr}_{kjs})q_{kjs}\right) + \epsilon_{kjs}$$

In the empirical estimation, we use a shock to the statutory tax rate.[39] Our estimated difference-in-differences coefficient, β^x, recovers this derivative under the identification assumption that the unobserved error term (ϵ_{kjs}) is exogenous to the statutory tax rate change, and that there are no spillovers across products and firm-level channels (satisfying SUTVA), implying the treatment does not alter the firm-time fixed effect, expressed as $\frac{\partial \gamma_{ks}}{\partial \ln(1+t^{str}_{js})} = 0$.

We decompose the observed coefficient in three parameters governing this response. The price elasticity of demand for product j ($\varepsilon_{xp,j}$), the tax pass-through elasticity to the consumer price ($\varepsilon_{p,j}$), and the elasticity of the effective tax rate with respect to the statutory tax rate ($\varepsilon_{etr,j}$)[40]:

$$\beta^x \equiv \frac{d\ln(x_{kjs})}{d\ln(1 + t^{str}_{js})} = \frac{\partial \ln f(\cdot)}{\partial \ln\left((1 + t^{etr}_{js})q_{kjs}\right)} \times \frac{d\ln\left((1 + t^{etr}_{js})q_{kjs}\right)}{d\ln(1 + t^{etr}_{js})} \times \frac{d\ln(1 + t^{etr}_{kjs})}{d\ln(1 + t^{str}_{js})}$$

[39] The statutory tax rate change is the same for all firm-products so using a dummy is the same as the statutory tax rate change up to a rescale.

[40] Given the discrete policy change, we estimate arc elasticities over the reform interval.

$$\varepsilon_{xp,j} \equiv \frac{\partial \ln f(\cdot)}{\partial \ln\left((1+t_{js}^{etr})q_{kjs}\right)}$$
$$\varepsilon_{p,j} \equiv \frac{d \ln\left((1+t_{js}^{etr})q_{kjs}\right)}{d \ln(1+t_{js}^{etr})}$$
$$\varepsilon_{etr,j} \equiv \frac{d \ln(1+t_{kjs}^{etr})}{d \ln(1+t_{js}^{str})}$$

Substituting our defined parameters, the total equilibrium quantity response simplifies to:

$$\beta^{x} \equiv \frac{d \ln(x_{kjs})}{d \ln(1+t_{js}^{str})} = \varepsilon_{xp,j}\varepsilon_{p,j}\varepsilon_{etr,j}$$

We also estimate each elasticity directly. Under the same identification assumptions, regressing the log of the effective tax factor on the statutory tax change, alongside our set of fixed effects, the reduced-form coefficient directly identifies the elasticity of the effective tax rate with respect to the statutory rate: $\beta^{etr} = \varepsilon_{etr,j}$.

$$\ln(1+t_{kjs}^{etr}) = \beta^{etr} \ln(1+t_{js}^{str}) + \alpha_{kj} + \gamma_{ks} + \nu_{kjs}$$

Furthermore, this specification is the first-stage regression for a two-stage least squares (2SLS) estimation. By instrumenting the endogenous effective tax rate with the exogenous statutory policy shock, the second stage replaces the observed $\ln(1+t_{js}^{etr})$ with its fitted values, $\ln(\widehat{1+t_{js}^{etr}})$.

As we observe prices, we can apply the identical empirical strategy to the tax-inclusive consumer price. This coefficient isolates the equilibrium consumer price response, which perfectly captures the product of the effective pass-through elasticity ($\varepsilon_{p,j}$) and the tax compliance elasticity ($\varepsilon_{etr,j}$):

$$\ln(p_{kjs}) = \beta^{p} \ln(1+t_{js}^{str}) + \alpha_{kj} + \gamma_{ks} + \omega_{kjs}$$

For tax-exclusive sales (producer revenue), defined as $s_{kjs} = q_{kjs}x_{kjs}$, the total derivative is the sum of the producer price response and the quantity response. Expanding this using the chain rule, and using the relation between the post- and pre-tax price elasticities $\varepsilon_{p,j} = (1+\varepsilon_{q,j})$, where $\varepsilon_{q,j}$ denotes the elasticity of the producer price with respect to the

tax (producer-price pass-through):

$$\begin{aligned}\beta^s &\equiv \frac{d\ln(s_{kjs})}{d\ln(1+t_{js}^{str})} \\ &= \frac{d\ln(x_{kjs})}{d\ln(1+t_{js}^{str})} + \frac{d\ln(q_{kjs})}{d\ln(1+t_{js}^{str})} \\ &= \varepsilon_{etr,j}\left[\varepsilon_{p,j}(1+\varepsilon_{xp,j})-1\right]\end{aligned}$$

E.1 Tax Evasion

We can extend the framework to a setting with tax evasion. Let the parameter $\lambda \in (0,1]$ capture the share of actual quantities sold that are formally reported. In this case, the observed (reported) quantity, x_{kjs}, may differ from the actual consumer demand, which we denote as x_{kjs}^*. Assuming the reported share is a function of the statutory tax rate, the reported sales equation becomes:

$$x_{kjs} = A_{kjs} f\left((1+t_{kjs}^{etr})q_{kjs}\right)\lambda(1+t_{js}^{str})$$

Taking the natural logarithm yields:

$$\ln(x_{kjs}) = \ln(A_{kjs}) + \ln f\left((1+t_{kjs}^{etr})q_{kjs}\right) + \ln\lambda(1+t_{js}^{str})$$

Applying our difference-in-differences framework, the total derivative of the observed log quantity with respect to the statutory tax factor now includes an additional term reflecting the reporting response:

$$\beta_{rep}^x \equiv \frac{d\ln(x_{kjs})}{d\ln(1+t_{js}^{str})} = \frac{d\ln(x_{kjs}^*)}{d\ln(1+t_{js}^{str})} + \frac{d\ln\lambda(1+t_{js}^{str})}{d\ln(1+t_{js}^{str})}$$

We can define this new parameter as the reporting elasticity, $\varepsilon_\lambda \equiv \frac{d\ln\lambda(\cdot)}{d\ln(1+t_{js}^{str})}$. Substituting our previous derivation for the actual demand response, the estimated coefficient for reported sales is:

$$\beta_{rep}^x = \beta^x + \varepsilon_\lambda$$

This composite parameter highlights the identification challenge introduced by evasion. If a reduction in the statutory tax rate lowers the evasion share—meaning the reported share λ

increases when the tax falls—then the reporting elasticity is negative ($\varepsilon_{\lambda,j} < 0$). Consequently, the estimated β^x_{rep} coefficient will overestimate the true equilibrium quantity response in absolute value, as the observed increase in sales is driven by both demand expansion and formalization.

Note that baseline differences in evasion across firm-products and aggregate shocks to reporting behavior across firm-periods are absorbed by the fixed effects (α_{kj} and γ_{ks}). Thus, bias in the estimated coefficient β^x_{rep} arises only if evasion responds directly to the product's tax rate. Moreover, the tax-base response including evasion, captured by the reported-sales coefficient (β^s_{rep}), is the relevant parameter for optimal taxation when evasion is incorporated (Section F.5).

E.2 Income Effects

We can model income effects in the demand function as:

$$x_{kjs} = A_{kjs}\, f\big((1 + t^{etr}_{js})q_{kjs},\, y_{js}\big)$$

Taking logs and totally differentiating with respect to the statutory tax factor decomposes the reduced-form response into a price channel and an income channel:

$$\beta^x_{inc} \equiv \frac{d\ln(x_{kjs})}{d\ln(1 + t^{str}_{js})} = \underbrace{\frac{\partial \ln f(\cdot)}{\partial \ln\big((1 + t^{etr}_{js})q_{kjs}\big)} \cdot \frac{d\ln\big((1 + t^{etr}_{js})q_{kjs}\big)}{d\ln(1 + t^{str}_{js})}}_{\beta^x} + \frac{\partial \ln f(\cdot)}{\partial \ln y_{js}} \cdot \frac{d\ln y_{js}}{d\ln(1 + t^{str}_{js})}$$

The first term captures the elasticities defined earlier, so that $\beta^x = \varepsilon_{xp,j}\varepsilon_{p,j}\varepsilon_{etr,j}$. For the income channel, define the income elasticity of demand for product j:

$$\eta_{js} \equiv \frac{\partial \ln f(\cdot)}{\partial \ln y_{js}}$$

By Slutsky's identity, the tax-induced price reduction expands real income proportionally to the treated product's expenditure share scaled by pass-through and compliance, $s_{js} = \theta_{js}\,\varepsilon_{p,j}\varepsilon_{etr,j}$. The income channel thus contributes $\eta_{js}\, s_{js}$ to the treated-product response.

This macroeconomic shock symmetrically affects control products j', contributing $\eta_{j's}\, s_{js}$ to their response. The firm-time fixed effect γ_{ks} absorbs the common component, leaving

only the differential income elasticity in the difference-in-differences estimator:

$$\beta^x_{inc} = \beta^x + (\eta_{js} - \eta_{j's})\, s_{js}$$

The estimated β^x_{inc} captures the compensated demand coefficient (β^x) plus a residual term corresponding to the part of the income effect not netted out by the control group. In our setting, η_{js} and $\eta_{j's}$ are nearly equal given the poor initial targeting of products, and the residual income shock is further mitigated by imperfect pass-through, so that $\beta^x_{inc} \approx \beta^x$.

Appendix F: Optimal Commodity Taxation

F.1 Model Setup under Perfect Pass-Through

We derive a general expression for optimal commodity tax rates in the spirit of the many-person Ramsey rule (Diamond 1975). We start with the simplest case in which we assume perfect pass-through, zero cross-price elasticities, and no firm surplus. We consider a continuum of households $i \in [0, 1]$, each with exogenous income y^i. Producer prices are denoted by q_j, and ad valorem taxes t_j imply consumer prices $p_j = (1 + t_j)q_j$. The indirect utility function of household i, denoted by $v(p, y^i)$, is defined as the maximum attainable utility given the price vector and the household's income:

$$v(p, y^i) = \max_{x^i} \; u^i(x^i) \quad \text{s.t.} \quad p \cdot x^i \leq y^i$$

where x^i denotes the chosen consumption vector and $u^i(\cdot)$ the household-specific direct utility function. Firms operate under constant returns to scale and perfect competition, implying zero profits. These assumptions allow us to abstract from firm behavior and ownership in the baseline. The government maximizes a social welfare function $G(\cdot)$ subject to a revenue requirement. The associated Lagrange multiplier, μ, is the shadow value of public funds, capturing the marginal social value of relaxing the government budget constraint by one unit.

The government social welfare function is given by

$$W = \int_i G\big(v(p, y^i)\big)\; di \;+\; \mu \sum_j t_j q_j x_j,$$

where aggregate demand for good j is

$$x_j = \int_i x_j^i(p, y^i)\, di,$$

and consumer prices satisfy

$$p_j = (1 + t_j)q_j.$$

The government chooses ad-valorem taxes $\{t_j\}$ to maximize welfare. Differentiating welfare with respect to t_j yields

$$\frac{dW}{dt_j} = \int_i G'\big(v(p, y^i)\big)\, \frac{\partial v(p, y^i)}{\partial p_j} \frac{\partial p_j}{\partial t_j}\, di + \mu \left(q_j x_j + \sum_k t_k q_k \frac{\partial x_k}{\partial p_j} \frac{\partial p_j}{\partial t_j} \right).$$

Assuming perfect pass-through ($\partial p_j / \partial t_j = q_j$), using Roy's identity

$$\frac{\partial v(p, y^i)}{\partial p_j} = -\frac{\partial v(p, y^i)}{\partial y^i} x_j^i(p, y^i),$$

and with no cross-price effects, ruling out substitution and fiscal externalities from other products, the first-order condition simplifies to

$$-q_j \int_i G'\big(v(p, y^i)\big)\, \frac{\partial v(p, y^i)}{\partial y^i} x_j^i\, di + \mu q_j \left(x_j + t_j q_j \frac{\partial x_j}{\partial p_j} \right) = 0.$$

Define the social marginal welfare weight

$$g^i \equiv G'\big(v(p, y^i)\big)\, \frac{\partial v(p, y^i)}{\partial y^i},$$

and individual consumption shares

$$a_j^i \equiv \frac{x_j^i}{x_j}.$$

The optimal tax condition can then be written as

$$\int_i g^i a_j^i\, di = \mu \left(1 + t_j q_j \frac{1}{x_j} \frac{\partial x_j}{\partial p_j} \right).$$

Define the (own-price) Marshallian demand elasticity

$$\varepsilon_{xp,j} \equiv \frac{\partial x_j}{\partial p_j} \frac{p_j}{x_j}.$$

Using $p_j = (1 + t_j)q_j$, the Ramsey optimal tax rule is

$$\boxed{\tau_j \equiv \frac{t_j}{1 + t_j} = \frac{1 - \int_i \frac{g^i}{\mu} a_j^i \, di}{-\varepsilon_{xp,j}}}$$

Interpretation This expression makes explicit the trade-off faced by the government. The denominator, $-\varepsilon_{xp,j}$, captures the responsiveness of the tax base: under full pass-through, taxes affect consumer prices one-for-one, such that the elasticity of the tax base coincides with the elasticity of quantities to prices. Goods with more elastic demand therefore generate larger distortions and face lower optimal tax rates. The numerator reflects distributional considerations, aggregating individuals' social marginal welfare weights according to their consumption shares of good j. Goods that are disproportionately consumed by individuals with higher social weights should face lower optimal tax rates. The optimal tax rate of good j reflects the trade-off between these efficiency and distributional concerns.

Equal Social Marginal Welfare Weights In the special case where the government assigns an equal social marginal welfare weight to all individuals ($g^i = \bar{g}$ for all i), the optimal tax rule reduces to

$$\frac{t_j}{1 + t_j} = \frac{1 - \bar{g}/\mu}{-\varepsilon_{xp,j}}.$$

In this case, the optimal indirect tax rate depends only on the relative valuation of public funds, μ, compared to the private marginal utility of income, $\bar{g}$, and on the efficiency cost of taxation. When $\mu = \bar{g}$, so that the government values public and private consumption equally at the margin, the optimal indirect tax rate is zero.

Compensated Elasticity To express the optimal tax rule in terms of compensated elasticities, we apply the Slutsky equation, $\frac{\partial x_j^i}{\partial p_j} = \left(\frac{\partial x_j^i}{\partial p_j}\right)^c - x_j^i \frac{\partial x_j^i}{\partial y^i}$, to the government's first-order condition. This isolates the substitution effect by absorbing the income effect of the price change into a net social marginal welfare weight:

$$\tilde{g}^i \equiv g^i + \mu t_j \frac{q_j x_j^i}{y^i} \eta_j^i$$

Where η_j^i is the income elasticity of household i for good j. This modified weight, $\tilde{g}^i$, credits households for the "fiscal externality" from the share of any marginal income returned to the government as tax revenue. The optimal tax rule can then be rewritten in terms of the compensated elasticity. ($\varepsilon^c_{xp,j}$):

$$\tau_j = \frac{1 - \int_i \frac{\tilde{g}^i}{\mu} a_j^i \, di}{-\varepsilon^c_{xp,j}}$$

F.2 Imperfect Pass-Through

In this section, we relax the assumption that producer prices are independent of taxes. Firms are allowed to adjust producer prices in response to taxation, so that pass-through of taxes into consumer prices may be incomplete. We continue to abstract from profit income in this section.[41]

Using Roy's identity,

$$\frac{\partial v(p, y^i)}{\partial p_j} = -\frac{\partial v(p, y^i)}{\partial y^i} x_j^i(p, y^i),$$

and allowing for imperfect pass-through so that consumer prices satisfy

$$p_j = p_j(t_j), \qquad \frac{\partial p_j}{\partial t_j} \neq q_j,$$

the derivative of welfare with respect to t_j is

$$\frac{dW}{dt_j} = -\int_i G'\big(v(p, y^i)\big) \frac{\partial v(p, y^i)}{\partial y^i} x_j^i \frac{\partial p_j}{\partial t_j} \, di \; + \; \mu \left(q_j x_j + t_j \frac{d(q_j x_j)}{dt_j} \right).$$

Define social marginal welfare weights and consumption shares as

$$g^i \equiv G'\big(v(p, y^i)\big) \frac{\partial v(p, y^i)}{\partial y^i}, \qquad a_j^i \equiv \frac{x_j^i}{x_j}.$$

Dividing by $q_j x_j$ and rearranging yields

$$\int_i g^i a_j^i \frac{\partial p_j}{\partial t_j} \frac{1}{q_j} \, di = \mu \left(1 + t_j \frac{1}{q_j x_j} \frac{d(q_j x_j)}{dt_j} \right).$$

[41] Incomplete pass-through can arise from competitive supply with decreasing returns, among other mechanisms. The optimal tax rule derived here assigns zero welfare weight to producer surplus, whether accruing as profits to firm owners or returns to other factors. Section F.5 relaxes this assumption.

Define the pass-through elasticity and the elasticity of pre-tax sales (tax base),

$$\varepsilon_{p,j} \equiv \frac{\partial p_j}{\partial t_j} \frac{1+t_j}{p_j}, \qquad \varepsilon_{s,j} \equiv \frac{d(q_j x_j)}{dt_j} \frac{1+t_j}{q_j x_j}.$$

Using $p_j = (1+t_j)q_j$, the first-order condition can be written as

$$\int_i g^i a_j^i \, \varepsilon_{p,j} \, di = \mu \left(1 + \frac{t_j}{1+t_j} \varepsilon_{s,j}\right),$$

which implies the optimal tax rule under imperfect pass-through:

$$\boxed{\tau_j \equiv \frac{t_j}{1+t_j} = \frac{1 - \int_i \frac{g^i}{\mu} a_j^i \varepsilon_{p,j} \, di}{-\varepsilon_{s,j}}}.$$

Interpretation This expression generalizes the baseline Ramsey rule by allowing for imperfect pass-through. Distributional considerations are now mediated by the extent to which taxes are transmitted to consumer prices: the numerator weights social marginal welfare weights by both consumption shares and the pass-through elasticity $\varepsilon_{p,j}$, so that lower pass-through reduces the welfare impact of a given tax change on consumers and weakens distributional concerns. With incomplete pass-through, part of the tax burden is absorbed through producer prices. Implicitly, this formulation assigns zero social weight to changes in pre-tax prices accruing to firms. This may be justified either by the social welfare structure, which assigns zero social weight to firm owners, or by supply-side assumptions under which changes in pre-tax prices do not generate welfare-relevant surplus.

Relative to the baseline with full pass-through, the denominator now captures the elasticity of the tax base rather than the elasticity of quantities. Under full pass-through, consumer prices move one-for-one with the tax, so the elasticity of sales coincides with the elasticity of quantities. With imperfect pass-through, however, taxes affect both prices and quantities, and the relevant efficiency cost is governed by the elasticity of sales, which determines how strongly the tax distorts revenue. As a result, imperfect pass-through creates a wedge between consumer incidence and tax-base distortions, and optimal taxation must account separately for how taxes affect prices paid by consumers and the evolution of the tax base.

F.3 Firm Ownership

We now extend the analysis by allowing individuals to own firm. Individual income includes profits, so that

$$y^i = y_0^i + \pi^i, \qquad \pi^i = \sum_j \sigma_j^i \pi_j,$$

where σ_j^i is the share of firm j owned by individual i, and $\pi_j = \gamma q_j x_j$ are aggregate profits, where $\gamma \in (0, 1]$ captures the profit margin. The individual value function is therefore

$$v(p, y^i) = \max_x \; u(x) \quad \text{s.t.} \quad p \cdot x \leq y^i.$$

Differentiating welfare with respect to t_j, and assuming no cross-profit effects ($\partial \pi_k / \partial t_j = 0$ for all $k \neq j$), yields

$$\frac{dW}{dt_j} = -\int_i G'\big(v(p, y^i)\big) \frac{\partial v(p, y^i)}{\partial y^i} \left(x_j^i \frac{\partial p_j}{\partial t_j} - \sigma_j^i \frac{d\pi_j}{dt_j} \right) di \; + \; \mu \left(q_j x_j + t_j \frac{d(q_j x_j)}{dt_j} \right).$$

Define social marginal welfare weights, consumption shares, and profit-to-expenditure shares as

$$g^i \equiv G'\big(v(p, y^i)\big) \frac{\partial v(p, y^i)}{\partial y^i}, \qquad a_j^i \equiv \frac{x_j^i}{x_j}, \qquad s_j^i \equiv \frac{\pi_j^i}{p_j x_j},$$

where π_j^i is the profit income of individual i derived from good j, and s_j^i represents this income as a share of total consumer expenditure on good j.

Dividing by $q_j x_j$ and rearranging, the first-order condition becomes

$$\int_i g^i \left(a_j^i \varepsilon_{p,j} - s_j^i \varepsilon_{\pi,j} \right) di = \mu \left(1 + \frac{t_j}{1 + t_j} \varepsilon_{s,j} \right),$$

where

$$\varepsilon_{p,j} \equiv \frac{\partial p_j}{\partial t_j} \frac{1 + t_j}{p_j}, \qquad \varepsilon_{s,j} \equiv \frac{d(q_j x_j)}{dt_j} \frac{1 + t_j}{q_j x_j}, \qquad \varepsilon_{\pi,j} \equiv \frac{d\pi_j}{dt_j} \frac{1 + t_j}{\pi_j}.$$

Since profits are a constant fraction of sales revenue, the proportional response of profits to taxation is equal to the response of sales revenue: $\varepsilon_{\pi,j} = \varepsilon_{s,j}$.

Substituting this relationship into the first-order condition yields the optimal tax rule:

$$\boxed{\frac{t_j}{1 + t_j} = \frac{1 - \int_i \frac{g^i}{\mu} \left(a_j^i \varepsilon_{p,j} - s_j^i \varepsilon_{s,j} \right) di}{-\varepsilon_{s,j}}}.$$

Interpretation Allowing for firm ownership introduces an additional welfare channel through firms' profits. In addition to consumer-price effects, the numerator now captures how taxation affects individuals' incomes through changes in profits, weighted by the distribution of profits of good j across individuals and the social marginal welfare weights of firm owners. Under the constant-margin assumption, profits are proportional to sales, so that profit responses to taxation are governed by the same elasticity as sales. The denominator continues to reflect the efficiency cost of taxation through distortions of the tax base. Optimal taxation therefore trades off consumer incidence, profit incidence, and efficiency considerations. The benchmark formulation without firm ownership is recovered as the special case in which profits are zero (e.g., due to perfectly offsetting costs) or firm owners have zero social weight.

F.4 Cross-Price Effects

We now allow for (i) imperfect pass-through and (ii) cross-price effects.

Consumer prices are functions of the full vector of taxes,

$$p_m = p_m(\boldsymbol{t}), \qquad \frac{\partial p_m}{\partial t_j} \text{ need not be zero for } m \neq j.$$

Differentiating welfare with respect to t_j yields

$$\frac{dW}{dt_j} = -\int_i G'\big(v(p,y^i)\big)\,\frac{\partial v(p,y^i)}{\partial y^i}\sum_m x_m^i(p,y^i)\frac{\partial p_m}{\partial t_j}\,di \;+\; \mu\left(q_j x_j + \sum_m t_m \frac{d(q_m x_m)}{dt_j}\right).$$

Define social marginal welfare weights and consumption shares as

$$g^i \equiv G'\big(v(p,y^i)\big)\,\frac{\partial v(p,y^i)}{\partial y^i}, \qquad a_m^i \equiv \frac{x_m^i}{x_m}.$$

Dividing the first-order condition by $q_j x_j$ and rearranging yields

$$\int_i \frac{g^i}{\mu}\left[\sum_m a_m^i \frac{x_m}{x_j}\frac{\partial p_m}{\partial t_j}\frac{1}{q_j}\right] di = 1 + \sum_m t_m \frac{1}{q_j x_j}\frac{d(q_m x_m)}{dt_j}.$$

Introduce cross pass-through elasticities and tax-base (sales) elasticities:

$$\varepsilon_{p,mj} \equiv \frac{\partial p_m}{\partial t_j}\frac{1+t_j}{p_m}, \qquad \varepsilon_{s,mj} \equiv \frac{d(q_m x_m)}{dt_j}\frac{1+t_j}{q_m x_m}.$$

Using $q_j = p_j/(1+t_j)$ and $d(q_m x_m)/dt_j = \varepsilon_{s,mj}\,(q_m x_m)/(1+t_j)$, the first-order condition can be written as

$$\int_i \frac{g^i}{\mu}\left[\sum_m a_m^i \frac{x_m p_m}{x_j p_j}\,\varepsilon_{p,mj}\right] di = 1 + \sum_m \frac{t_m}{1+t_j}\frac{q_m x_m}{q_j x_j}\,\varepsilon_{s,mj}.$$

Let

$$\omega_{mj} \equiv \frac{p_m x_m}{p_j x_j}$$

denote the ratio of aggregate expenditure on good m to aggregate expenditure on good j (i.e., a relative expenditure weight). Rearranging yields the optimal tax rule under imperfect pass-through and cross-price effects:

$$\boxed{\tau_j \equiv \frac{t_j}{1+t_j} = \frac{1 - \int_i \frac{g^i}{\mu}\left[\sum_m a_m^i \omega_{mj}\,\varepsilon_{p,mj}\right] di + \sum_{m\neq j} \frac{t_m}{1+t_m}\omega_{mj}\,\varepsilon_{s,mj}}{-\varepsilon_{s,jj}}}.$$

Interpretation The generalized optimal tax formula highlights two channels through which cross-price effects matter. First, the numerator captures indirect consumer-price effects. When a tax on good j affects the prices of other goods, consumer welfare depends not only on the price of good j, but also on induced price changes for goods $m \neq j$. These effects are weighted by consumption shares, social marginal welfare weights, and the relative aggregate importance of good m to good j. As a result, goods whose taxation indirectly raises the prices of goods disproportionately consumed by individuals with higher social weights face lower optimal tax rates. Second, cross-price effects operate through the tax base. A tax on good j changes the quantities sold of other goods through substitution, affecting government revenue collected on those goods. These effects depend on tax-base cross-elasticities and are scaled by existing tax wedges and the relative size of each tax base. When taxing good j reduces the tax base of goods that are already highly taxed or account for a large share of tax revenue, the marginal efficiency cost of taxation increases.

F.5 Evasion

In this section, we introduce evasion, allowing the reported quantity (x_j^{rep}) to differ from actual consumer demand (x_j), informal consumption to vary across households, and firm and consumer evasion to differ. Let $\lambda_j^{c,i} \in [\lambda_j^i, 1]$ denote the share of household i's formal consumption (pays commodity tax), and let $\lambda_j^i(t_j) \in (0,1]$ denote the share of household i's

consumption that firms report to the tax authority, such that $x_j^{rep,i} = x_j^i \lambda_j^i(t_j)$.[42] Household i pays the tax-inclusive price $p_j = (1+t_j)q_j$ on the taxed share ($\lambda_j^{c,i}$) and the informal price q_j^{inf} on the informal share ($1-\lambda_j^{c,i}$). The effective price for good j faced by household i is a weighted average: $p_j^{e,i} = \lambda_j^{c,i}(1+t_j)q_j + (1-\lambda_j^{c,i})q_j^{inf}$, and for simplicity, we treat the informal price as fixed with respect to the tax ($\partial q_j^{inf}/\partial t_j = 0$).

The indirect utility function for household i incorporates its own effective price vector, $p^{e,i}$:

$$v(p^{e,i}, y^i) = \max_{x^i} \; u^i(x^i) \quad \text{s.t.} \quad p^{e,i} \cdot x^i \leq y^i$$

The government's social welfare function evaluates consumer welfare over individual effective prices, while tax revenues depend on aggregate reported sales ($\int_i x_j^i \lambda_j^i \, di$):

$$W = \int_i G\big(v(p^{e,i}, y^i)\big) \; di \; + \; \mu \sum_j t_j q_j \int_i x_j^i \lambda_j^i \, di$$

Defining the social marginal utility of income for individual i as $g^i \equiv G'(v)\frac{\partial v}{\partial y^i}$ as before and incorporating the modified, household-specific consumer incidence, the first-order condition of welfare with respect to the tax rate t_j:

$$\begin{aligned}
\frac{dW}{dt_j} &= -\int_i G'\big(v(p^{e,i}, y^i)\big) \frac{\partial v(p^{e,i}, y^i)}{\partial y^i} x_j^i \frac{\partial p_j^{e,i}}{\partial t_j} \, di \; + \; \mu \left(q_j \int_i x_j^i \lambda_j^i \, di + t_j \frac{d}{dt_j} \int_i q_j x_j^i \lambda_j^i \, di \right) = 0 \\
&= -\int_i g^i x_j^i \frac{\partial p_j^{e,i}}{\partial t_j} \, di \; + \; \mu \left(q_j \int_i x_j^i \lambda_j^i \, di + t_j \frac{d}{dt_j} \int_i q_j x_j^i \lambda_j^i \, di \right) = 0
\end{aligned}$$

To evaluate this condition, we formalize the relevant elasticities. The effective-price elasticity ($\varepsilon^i_{p^e,j}$) is now household-specific, while total sales ($\varepsilon_{s,j}$) and the reported formal share ($\varepsilon_{\lambda,j}$) are aggregate objects:

$$\varepsilon^i_{p^e,j} \equiv \frac{\partial p_j^{e,i}}{\partial t_j} \frac{1+t_j}{p_j^{e,i}}, \qquad \varepsilon_{s,j} \equiv \frac{d(q_j x_j)}{dt_j} \frac{1+t_j}{q_j x_j}, \qquad \varepsilon_{\lambda,j} \equiv \frac{d\bar{\lambda}_j}{dt_j} \frac{1+t_j}{\bar{\lambda}_j}$$

where $\bar{\lambda}_j \equiv \int_i a_j^i \lambda_j^i \, di$ is the consumption-weighted average reporting share, with $a_j^i \equiv x_j^i / x_j$ and $\int_i a_j^i \, di = 1$, so that $\int_i x_j^i \lambda_j^i \, di = x_j \bar{\lambda}_j$.

The household-specific effective-price elasticity is a function of the pass-through elasticity

[42] We define the lower bound of $\lambda_j^{c,i}$ to equal λ_j^i as the consumer cannot evade the commodity tax more than the firm selling the product (e.g. taxes are remitted by firms). This bound holds household by household.

and the household's formal consumption share. Because the informal price is fixed, the tax transmits to the effective price only through household i's formal share, such that $\varepsilon^i_{p^e,j} = \omega^i \varepsilon_{p,j}$ with $\omega^i \equiv \frac{\lambda_j^{c,i}(1+t_j)q_j}{p_j^{e,i}} \in (0,1]$ the formal (taxed) share of household i's effective expenditure on good j. Sales, reported sales and evasion elasticities relate at the aggregate level as: $\varepsilon^{rep}_{s,j} = \varepsilon_{s,j} + \varepsilon_{\lambda,j}$. Substituting these elasticities and solving for the optimal tax rate $\tau_j \equiv \frac{t_j}{1+t_j}$ provides the optimal tax rule with heterogeneous evasion:

$$\tau_j = \frac{1 - \int_i \frac{g^i}{\mu} a^i_j \frac{\lambda_j^{c,i}}{\bar{\lambda}_j} \varepsilon_{p,j} \, di}{-\varepsilon^{rep}_{s,j}}$$

Firm Evasion To focus on firm evasion, we consider the case in which informal consumption and firm evasion shares are homogeneous across households ($\lambda_j^{c,i} = \lambda_j^c$ and $\lambda_j^i = \lambda_j$). The optimal tax rule simplifies to:

$$\boxed{\tau_j = \frac{1 - \frac{\lambda_j^c}{\lambda_j} \int_i \frac{g^i}{\mu} a^i_j \varepsilon_{p,j} \, di}{-\varepsilon^{rep}_{s,j}}}$$

The modified optimal tax rule introduces two new forces. In the numerator, the marginal welfare cost of taxation is scaled by the ratio of the consumer's effective tax burden to the firm's reported tax base ($\lambda_j^c/\lambda_j \geq 1$). This term captures the allocation of the "evasion surplus" between consumers and firms. When $\lambda_j^c = \lambda_j$, the consumer fully captures the surplus. In this scenario, the consumption affected by the tax perfectly matches the reported tax base, and the evasion terms in the numerator cancel out. When the firm captures part of the surplus ($\lambda_j^c > \lambda_j$), it charges the tax-inclusive price on a larger share of units than it actually remits taxes for to the government. The consumer suffers the price distortion on a higher share of their consumption than the tax base. In the extreme case where the firm captures all the surplus ($\lambda_j^c = 1$), the welfare lost is scaled by the evasion share ($1/\lambda_j$). This magnifies the social cost per dollar of public revenue, constraining the optimal rate.

In the denominator, the relevant parameter becomes the elasticity of reported sales ($\varepsilon_{s,j} + \varepsilon_{\lambda,j}$). Because tax increases lead to higher evasion ($\varepsilon_{\lambda,j} < 0$), the reported tax base is more elastic than actual consumption. Knowing that higher taxes will lead to revenue losses from both depressed real sales and increased evasion, the optimal tax rate is lower than it would be in a perfectly compliant economy.

Informal Consumption To isolate the effect of informal markets, we consider the case where $\lambda_j^{c,i} = \lambda_j^i$ for all i. Here, consumers fully capture the evasion surplus, but informality varies across households. Defining the share of formal consumption of good j attributable to household i as:

$$\tilde{a}_j^i \equiv \frac{x_j^i \lambda_j^i}{\int_i x_j^i \lambda_j^i \, di} = a_j^i \, \frac{\lambda_j^i}{\bar{\lambda}_j}, \qquad \int_i \tilde{a}_j^i \, di = 1,$$

the optimal tax rule simplifies to:

$$\boxed{\tau_j = \frac{1 - \displaystyle\int_i \frac{g^i}{\mu} \, \tilde{a}_j^i \, \varepsilon_{p,j} \, di}{-\varepsilon_{s,j}^{rep}}}$$

Relative to the firm evasion case, the key difference lies in the numerator, where actual consumption shares (a_j^i) are replaced by formal consumption shares ($\tilde{a}_j^i$). The progressivity of the tax depends on who buys the good formally, rather than who consumes it overall, because only formal consumption is taxed.

F.6 Welfare

The welfare change for household i from a change in the price of good j is the area under its demand curve, which we express in log prices:

$$\Delta W_{ij} \;=\; -\int_{\ln p_j^0}^{\ln p_j^1} p \, x_{ij}(p) \, d\ln p,$$

where the integrand $V(\ln p) = p \, x_{ij}(p)$ is spending on good j. Along the demand curve, spending responds to the price with elasticity $\frac{d \ln V}{d \ln p} = \frac{d \ln(p\,x)}{d \ln p} = 1 + \varepsilon_{xp}$. A second-order approximation of the welfare change around observed initial spending $v_{ij} = p_j^0 x_{ij}^0$ therefore gives:

$$\Delta W_{ij} \;\approx\; -v_{ij}\Big[\Delta \ln p_j + \tfrac{1}{2}\big(1 + \varepsilon_{xp}\big)\,(\Delta \ln p_j)^2\Big].$$

The first term is the mechanical burden implied by the envelope theorem, and the second term is the substitution correction. Consumer prices respond to the statutory tax factor with the estimated pass-through elasticity, $\Delta \ln p_j = \varepsilon_p \, \Delta \ln(1 + t_j)$. Using the identities linking elasticities, $\varepsilon_s = \varepsilon_q + \varepsilon_{xp}\varepsilon_p$ and $\varepsilon_q = \varepsilon_p - 1$, yields the expression:

$$\Delta W_{ij} \;\approx\; -v_{ij} \, \varepsilon_p\Big[\Delta \ln(1 + t_j) + \tfrac{1}{2}\big(1 + \varepsilon_s\big)\,(\Delta \ln(1 + t_j))^2\Big].$$

This equation provides a second-order approximation in the tax change for any demand system, assuming negligible income effects and zero cross-price substitution. Computing instead the exact consumer-surplus change under isoelastic demand gives $\Delta W_{ij} = -v_{ij}\,[f_j^{\,1+\varepsilon_{xp}} - 1]/(1+\varepsilon_{xp})$ with $f_j = \big((1+t_j^1)/(1+t_j^0)\big)^{\varepsilon_p}$.

Appendix G: Model Estimation

This section explains how we take the formula derived in Section F.2 to the data.

Household Consumption Survey Surveys typically capture only a fraction of aggregate household consumption (Deaton 2005). To address this issue and express all values in 2017 terms, we rescale total consumption in the survey to match aggregate household consumption from the national accounts in 2017, mapping the consumption categories reported in the survey to the corresponding national accounts items.

We perform the analysis at the percentile level, with each percentile containing the same number of individuals. Aggregating to percentiles eliminates zero consumption observations for some goods, which would otherwise complicate the estimation of income elasticities and the adjustment of consumption after redistribution. Households are ranked by per capita total consumption, excluding non-market items, gifts and free collections, and vehicle purchases. Vehicle purchases are excluded because of their lumpy nature but added back when constructing total percentile consumption. Non-market items include cooked meals received as assistance or free at the workplace, and taxes paid directly by households, such as road cess and house tax.

Income Since the survey reports consumption but not income, we define household baseline post-tax income as total household consumption, $y^i_{\text{post, baseline}} = c^i$, and apply the 2017 PIT schedule to recover baseline pre-tax income, $y^i_{\text{pre, baseline}} = y^i_{\text{post, baseline}} + t^i_{2017}$. We use household income, rather than individual income, as the tax unit because the survey does not provide sufficient information on intra-household allocation or on the income of individual household members. This approach treats all household income as being earned by a single household member.

Consumption Shares The counterfactual luxury bracket is constructed by selecting the goods with the highest income elasticities until their combined consumption reaches the same share of total consumption as goods initially taxed at 28 percent, namely 11 percent of total goods consumption. To compute consumption in the 18 percent bracket and in the

counterfactual luxury bracket, we consider only goods covered by the GST. We therefore exclude services, non-market items, gifts and free collections, and petrol and diesel for vehicles, which are outside the GST base.

Personal Income Tax For the baseline personal income tax (PIT), we use the standard statutory schedule in place in 2017, with marginal rates of 5 percent, 20 percent, and 30 percent above thresholds of INR 250,000, INR 500,000, and INR 1,000,000, respectively. The Indian PIT system includes several deductions and exemptions that cannot be measured in the survey. For example, taxable income can be reduced by up to INR 150,000 in specified savings and investments, a further INR 50,000 in pension contributions, INR 200,000 in home-loan interest, and INR 25,000 in health-insurance premiums. The exemption threshold was also higher for taxpayers above 60 and 80 years old.

To account for these features, we calibrate the effective exemption level so that 4 percent of households in the survey are liable for PIT, matching the share of the adult population that actually pays PIT (Ministry of Finance, Government of India 2019). Specifically, we set the exemption adjustment equal to the difference between income at the 96th percentile and the statutory exemption threshold of INR 250,000. We then shift all income tax brackets by this same exemption amount. When considering counterfactual expansions on the PIT coverage up to a given percentile, we reduce proportionally all tax brackets including this exemption component.

Transfers In the baseline, PIT revenues are redistributed lump-sum, with no revenue loss. Targeted transfers are lump-sum transfers to the bottom 50 percent, with no revenue loss. For NREGA transfers we consider that 75 percent goes to the bottom 50 and 25 percent to the top 50 following evidence from Dutta et al. (2014) and Imbert and Papp (2015), with 10 percent administrative costs (Dutta et al. 2014) and 17 percent leakage (Banerjee et al. 2020; Muralidharan et al. 2016), resulting in net transfers equal to 73 percent of revenues. Because the baseline PIT schedule applied to the survey yields lower PIT revenues than the amount actually collected in 2017, we apply a scaling factor to PIT revenues throughout, calibrated to match this aggregate in the baseline, to avoid understating the role of the PIT.[43]

[43] According to (Ministry of Finance, Government of India 2019), personal income tax collected from individuals amounted to INR 3,178,450 millions in the 2017 fiscal year (April 2017 to March 2018). To focus on personal income only, we remove 27.2 percent of aggregate individual PIT revenues, corresponding to the share of business income taxed under the PIT.

Social Weights Social marginal welfare weights are defined as the inverse of post-tax income:

$$g^i = \frac{1}{y^i_{\text{post}}}.$$

Elasticities For the elasticities, we use the parameters estimated in the specification for survey products: $\varepsilon_p = 0.497$ and $\varepsilon_{xp} = -2.390$. We then recover the sales elasticity using the relationship

$$\varepsilon_{s,j} = \varepsilon_{xp,j}\varepsilon_{p,j} + (\varepsilon_{p,j} - 1).$$

Deriving the elasticity of sales ($\varepsilon_{s,j}$) from the primitive parameters allows us to remain consistent when recomputing sales elasticities in robustness exercises. For example, simulating full pass-through must lead to $\varepsilon_{s,j}$ becoming more negative (a larger magnitude) to keep the identity valid, and treating the sales elasticity independently would implicitly reduce the quantity-to-price elasticity.

Marginal Value of Public Funds To discipline the model, we calibrate the marginal value of public funds (μ) so that the optimal tax rate for goods initially taxed at 18 percent equals 18 percent. Interpreting the 18 percent rate as optimal allows us to focus on optimal rate differentiation rather than the overall level of taxation.

Personal Income Tax Expansion We treat changes in the PIT schedule as exogenously given. For a tax schedule f: $y^i_{\text{post, f}} = y^i_{\text{pre, baseline}} - t^i_{\text{f}}$. PIT and its redistribution affect the level and composition of consumption across goods ($c^i_j(y^i_{\text{post, f}})$) and social marginal welfare weights ($g^i = \frac{1}{y^i_{\text{post, f}}}$). To capture the change in household consumption induced by taxation and redistribution, we estimate how consumption of each good responds to changes in income along the distribution. For each group of goods j, we estimate separate Engel-curve regressions by decile of the income distribution. Within each decile d, we allow for nonlinear income effects through a quadratic specification in log income:

$$\log c^i_j = \alpha_{jd} + \beta_{1jd} \log y^i + \beta_{2jd} (\log y^i)^2 + u^i_j,$$

where i indexes percentiles, d denotes the resource decile to which percentile i belongs, c^i_j is consumption of good j, and y^i denotes percentile-level total consumption observed in the survey ($y^i_{\text{post, baseline}}$). Estimating one regression per decile, together with the quadratic specification, allows the income elasticity of consumption for each good to vary flexibly along the distribution and fits observed consumption patterns better than a single pooled

Engel curve. Then we update observed consumption according to the change in income after personal income tax and transfers.

Extension Parameters This paragraph presents the remaining parameters used to implement the extension formulas derived in Section F. In the evasion extension, we set the firm reporting parameter $\lambda = 0.89$, based on the pre-reform gap between the statutory tax rate and the effective tax rate among treated products, and set the consumer reporting parameter $\lambda_c = 1$. This assumption attributes all evasion gains to firm owners, with no sharing with consumers, reducing the optimal rate differentiation. In the firm-owners extension, we allocate all profits to the top 1 percent and set the profit margin equal to the aggregate value-added margin observed pre-reform, $\gamma = 0.72$. In the extension correcting top incomes, we add savings based on income-to-consumption levels of Bharti et al. (2026), constructing an augmented income measure, $y = c + t + s$.

Welfare effects We evaluate three revenue-neutral counterfactuals. Targeted 28% reassigns the goods with the steepest Engel curves to the 28 percent bracket (Section 6.1). Optimal rate luxuries assigns this same basket the optimal luxury rate of 42 percent instead of 28 percent. Optimal rate all products assigns each product its optimal rate from Equation (4), with μ calibrated to preserve the initial average effective tax rate. Welfare effects follow Equation (5). The mechanical analysis includes only the direct price effect of the tax change, while the behavioral analysis additionally incorporates the resulting quantity response. The informality analysis further accounts for the fact that taxation affects only the formal share of expenditure. To ensure revenue neutrality, we adjust all tax rates within each counterfactual by a common scaling factor. Let t_j^0 and $\tilde{t}_j$ be the initial GST and counterfactual rates, $b_j^0 = q_j x_j^0$ the initial tax base, and $R^0 = \sum_j t_j^0 b_j^0$ the initial total tax revenue. All rates in a given counterfactual are scaled by a common factor κ to match initial tax revenue. In the mechanical scenario the base is fixed, so $R^0 = \kappa_M \sum_j \tilde{t}_j b_j^0$. In the behavioral scenario the base responds with the estimated sales elasticity, $b_j(\kappa) = b_j^0 \left[(1 + \kappa\tilde{t}_j)/(1 + t_j^0)\right]^{\varepsilon_s}$, so that $R^0 = \kappa_B \sum_j \tilde{t}_j b_j(\kappa_B)$, which we obtain by iterating from κ_M until convergence.