Equity versus Efficiency of Indirect Taxes: Evidence from a Large Reform in India

Pierre Bachas, Davi Bhering, Pulak Ghosh*

April 7, 2026

Abstract

We study the redistributive and efficiency effects of differentiated commodity taxation in a constrained tax system. Using administrative data and a large VAT cut in India, we show that tax reductions are partially passed through to prices, generate moderate quantity responses, and induce little product mislabeling. As a result, well-targeted differentiated commodity taxes achieve meaningful redistribution at limited efficiency cost. Our calibrations yield an optimal luxury tax premium of 20 percentage points. Given the features of India's income tax and transfer system—high exemption thresholds, moderate rates, and imperfect targeting—we show that overturning this conclusion would require large and simultaneous relaxation along all margins.

JEL codes: E26, H21, H23, 023

*Pierre Bachas: Senior Economist, World Bank Research, 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 Santiago Cesteros and Radhika Kaul for excellent research assistance. We are also grateful for suggestions from Liam Wren-Lewis 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.

1 Introduction

Most developing countries combine tight fiscal space with persistently high income inequality (Chancel et al. 2025). Fiscal policy must therefore reconcile efficient revenue collection with redistribution. When income is observable and can be taxed nonlinearly, and when transfers can be precisely targeted, differentiated commodity taxation has no role in optimal redistribution (Atkinson and Stiglitz 1976). In many developing countries, however, informational, administrative, and political constraints sharply limit the coverage and progressivity of income taxes and transfers (Jensen 2022). As a result, redistribution cannot rely solely on direct taxation, creating scope for indirect consumption taxes and for their rate structure.

This paper shows that, under these constraints, differentiated commodity taxation can complement direct taxes by targeting consumption patterns correlated with income. However, this result comes with a twist relative to prevailing practice. Differentiation is typically implemented through reduced rates on essential goods, while luxury goods are more rarely taxed above the standard rate. We show that this logic may be reversed: reduced rates on necessity goods deliver limited redistribution, whereas higher taxes on luxury goods offer greater redistributive scope and can offset limited taxation at the top of the income distribution.

To fix ideas, Figure 1 illustrates the asymmetric tagging potential of commodity taxation using India's household expenditure survey. The ratio of consumption by the top income decile to that of the bottom decile is effectively unbounded for the most luxurious goods, but just below one-half for the most essential goods. Consequently, when the tax rate on necessities is reduced, for every dollar saved by the poorest households, roughly fifty cents accrue to the richest.[1] By contrast, a tax increase on luxuries is highly targeted: almost all of the burden falls on the top decile, with virtually none on the bottom decile.

A common objection to higher tax rates on luxury goods is that they generate large efficiency costs, through reduced consumption, increased evasion, and product misclassification, yet this claim has received limited empirical support. To assess this equity–efficiency trade-off, we study a large change in VAT rates in India for a broad set of goods, and apply a difference-in-differences design using granular

[1]Moreover, as the poor disproportionately consume in the informal sector, most of the benefits from rate reduction on necessities de-facto accrues to richer households (Bachas et al. 2023).

administrative data on product-level prices, quantities, and sales. The results point to moderate efficiency costs along both the intensive and extensive margin, but substantial distributional gains. Combining these dimensions, our calibrations yields an optimal luxury tax premium of around 20 percentage points. Given current features of India's income tax and transfer system—high exemption thresholds, moderate rates, and incomplete targeting—overturning this conclusion would require large and simultaneous relaxations along all margins.

India's main consumption tax is a value-added tax—the Goods and Services Tax (GST)—with five rate brackets: exempt, 5%, 12%, 18%, and 28%. A central rationale for this multi-rate structure is equity. The standard rate is 18%, while the 28% rate applies to a broad set of goods, including cosmetics, furniture, and household appliances. The income targeting underlying this classification dates back to the 1990s and, while progressive, is imperfect: on average, the slopes of the Engel curves for goods taxed at 28% are similar to those for goods taxed at 18%. In November 2017, a reform abruptly reclassified three-quarters of goods previously taxed at 28% to the 18% rate, generating large and plausibly exogenous variation in tax rates across goods. The tax cut reflected the central government's incentive to bolster support for the newly introduced federal GST and was supported by states that were guaranteed compensation for any resulting revenue losses.

In a first part, we study how registered firms respond to a sharp cut in consumption tax rates, using all GST returns from Karnataka—the fifth-largest Indian state by GDP—between July 2017 and March 2019. The data are rich: firms report monthly quantities, sales, and taxes at the product level. Our core sample comprises 105,000 firms and 842 products (HSN-4 level), totaling five million firm–product–month observations. To estimate behavioral responses, we implement a differences-in-differences design that compares, within firms, the prices, quantities, and sales of products that experienced a rate cut with those that did not. The empirical setting consists of a single treatment period, no treatment reversal, and a large and uniform treatment intensity—a ten-percentage-point tax cut. We analyze multi-product firms, which account for the vast majority of sales and are thus the relevant unit of analysis. This permits to include firm-by-time and firm-by-product fixed effects, isolating product-specific responses within firms. We assess potential within-firm spillovers using variation in treatment intensity and firms that never sell treated products.

Our empirical analysis yields four main findings on effective tax rates, price

pass-through, and direct and indirect efficiency costs. 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% to treated products before the reform, either strategically or due to system complexity. Among baseline compliers, the decline in effective rates matches the statutory change.

Second, we find substantial but incomplete pass-through of the tax cut to consumer prices. On average, prices of treated products decline by 3.5%: slightly over half of the effective rate reduction is passed through as lower prices, despite a contemporaneous anti-profiteering campaign.[2] This magnitude matches existing evidence from VAT rate reductions in other contexts (Benzarti et al. 2020).

Third, the efficiency costs of the tax cut are moderate. At the extensive margin, we find no evidence of increased entry of treated products among either incumbent or new firms. At the intensive margin, quantities rise by 3.2% in response to a 3.5% price decline, raising sales by 6.3%. These behavioral responses are too small to offset the mechanical revenue loss from the rate cut, implying a sizable drop in tax revenues. Firm-level value-added responds similarly. Sales elasticities are larger among smaller firms filing simplified returns, while price pass-through is lower among largest firms.

Fourth, we study indirect efficiency effects. While items such as make-up and watches were reclassified, related products such as essential oils and clocks remained taxed at 18%. A common concern with differentiated commodity taxation is that it may induce avoidance through product misclassification or shift demand toward close substitutes. To assess these channels, we compare the behavior of unaffected products that are closely related to treated goods with that of more distant alternatives. Using within-firm over-time comparisons, we find no evidence of even modest declines in quantities or sales for related products relative to unrelated ones.

Two complementary analyses reinforce our main findings. First, we revisit the November 2017 reform using scanner data from a large hypermarket chain. Although less representative, these data span a longer period (including pre-GST), report products at the barcode level, and capture only final-consumer sales. Results are qualitatively similar, with somewhat lower price pass-through and slightly larger quantity responses, consistent with firm-size heterogeneity in the tax data. Second, we study the July 2018 reform, which reduced rates on the smaller set of goods still

[2]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.

taxed at 28%, and find similar effects.

The results are robust to a range of alternative samples and specifications. Since identification relies on multi-product firms, we assess within-firm spillovers (potential SUTVA violations) by examining the behavior of control products sold by firms that also sell treated products, relative to the same products in other firms, finding no effects. Our analysis aims to focus on products sold to final consumers, yet firms only report their total share of sales to final consumers. While our main sample already excludes firms selling mainly to other firms, the results are robust to more stringent restrictions, and to excluding products with high upstreamness indices. Finally, related products are defined using HS classifications, yet machine-learning measures of relatedness yield similar evidence of limited mislabeling. The findings are robust to weighting products by sales and estimation at the firm level.

In a second part, we examine the equity–efficiency trade-off of differentiated commodity taxation and the role of luxury taxes in complementing income taxes and transfers. We discipline this analysis with parameters estimated from the quasi-elimination of the 28% tax bracket induced by the combined November 2017 and July 2018 reforms.[3] Efficiency effects are modest: applying a sufficient-statistic approach based on our quantity elasticity with respect to the effective tax rate, we estimate a deadweight loss of 13 cents per tax dollar, implying a low marginal cost of public funds. In revenue terms, for every dollar of tax revenue lost from the rate cut, roughly 20 cents is recouped through behavioral increases in sales.

To assess distributional effects on consumers, we combine household expenditure data with our pass-through estimates. At baseline, the GST was progressive: top-decile households faced an effective tax rate on consumption of 9.5% vs. 6% for the bottom decile, and consumed twice the share of goods taxed at either 18 or 28%. The reform reduced GST progressivity, lowering taxes by 1.2% of annual consumption for the top decile versus 0.6% for the bottom. However, the original assignment of goods to the standard and luxury brackets was poorly targeted: consumption baskets taxed at higher rates were no more progressive than those taxed at standard rates. Reassigning to the top tax bracket the goods with the steepest Engel-curve slopes—while holding its consumption share at 11%—would lower effective GST rates for all deciles from the eighth and below (from 6.0% to 4.8% for the bottom decile) and raise them only

[3]The November 2017 and July 2018 reforms reduced rates on 75% and 20% of top-bracket goods, respectively; only a small set—such as "sin goods"—remained subject to the top 28% rate.

for the top decile (from 9.5% to 10.7%).

We then set up a multi-person model of optimal commodity taxation following (Diamond 1975). The economy consists of a continuum of households with exogenous incomes, each consuming a standard and a luxury good. The government sets commodity tax rates to balance equity and efficiency, initially abstracting from income taxation. Calibrating the model using our estimated elasticities, pass-through rates, and distributional consumption patterns, we find that the optimal tax rate on luxury goods is more than 20 percentage points higher than that on standard goods.

We next introduce income taxes and transfers, and examine how limits to their coverage and progressivity affect the optimal structure of luxury taxes. We consider India's system at the time of the reform: income taxation featured a high exemption threshold (covering only the top 4 percent of households); statutory rates rose from 5 to 30 percent only at the very top of the distribution; and a key transfer program, NREGA, targeted the bottom half of the population with partial accuracy (Y. Liu et al. 2018). Starting from this baseline, luxury tax premia meaningfully decline only under large and simultaneous relaxations of all three constraints. Consistent with the Atkinson and Stiglitz (1976) benchmark, the luxury tax premium converges to zero when the exemption threshold is halved, income tax rates are doubled and, and revenue collected is redistributed to the bottom 50 percent with perfect accuracy.

Several considerations are relevant for external validity. First, the reform equalized rates across non-essential goods, which may generate asymmetric effects relative to reforms that introduce differentiation. Second, price pass-through has been systematically found to be lower for tax cuts than for rises (Benzarti et al. 2020). Finally, our analysis abstracts from other dimensions of luxury taxation, including paternalistic objectives, ethical considerations, and labor–leisure distortions (Saez 2002).

Bearing these caveats in mind, our results suggest that, in Karnataka, the VAT raises revenue at a moderate efficiency cost even under differentiated rates. Higher rates on selected luxury goods can increase tax progressivity in settings where personal income taxation is constrained by structural or political factors. Reflecting these considerations, India reintroduced a 40% luxury ("de-merit") tax bracket in 2025, underscoring the continued policy relevance of differentiated consumption taxes.

The paper is organized as follows. Section 2 discusses the context. Section 3 presents the data and empirical strategy. Section 4 shows the main results. Section 5 analyzes scanner-level data. Section 6 calibrates optimal luxury tax premia.

1.1 Related Literature and Contributions

This paper contributes to three strands of the literature. First, we contribute to work on indirect tax pass-through (see recent studies by Benedek et al. 2020; Benzarti et al. 2020; Carbonnier 2007; Harju et al. 2018; Velayudhan and Brusco 2022). Using rich administrative and scanner data, we measure how prices respond to large tax cuts across a broad set of goods. Our estimates fall within the existing range: pass-through of tax cuts to consumer prices is incomplete and lower among larger firms. The results also highlight the need to account for changes in effective, not just statutory, tax rates.

Second, we build on research on the efficiency costs of indirect taxation (Agarwal et al. 2025; Belan and Gauthier 2006; Gaarder 2022; Garg et al. 2023; Harris et al. 2018) and how tax design affects revenue collection (Pomeranz 2015; Waseem 2023). We exploit a large, sudden reform affecting a broad set of goods and combine detailed product-level tax returns from multi-product firms with granular scanner data. This allows us to directly estimate efficiency costs through quantity responses and extensive-margin adjustments (Benzarti and Carloni 2019; L. Liu et al. 2021).

Third, we contribute to the literature on optimal tax design in developing countries. A key strand emphasizes second-best designs in the presence of administrative and informational constraints (Best et al. 2015; Huang and Rios 2016; Londoño-Vélez and Ávila-Mahecha 2024). We study the role of higher rates on luxury goods where income tax and transfers are limited. By examining mislabeling—concerns often raised in policy debates but rarely tested empirically—we provide insights on the merits of multi-rate VAT systems. The targeting capacity of luxury taxation also connects our analysis to the broader literature on tax and transfer targeting, including work on India's main social assistance programs.

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 (ITC) on eligible purchases. Almost all goods and services are covered, 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 subsequently redistributed to states according to rules based on population, income, and area. Overall, states retain over 70% of total GST proceeds (State of State Finances 2024). The GST Council governs the system, including setting rates and assigning products to tax brackets. States hold two-thirds of the voting power in the GST Council and the central government the remaining one-third.

Tax rates setting at the onset of GST The GST Council assigned products to five tax brackets: exempt, 5%, 12%, 18%, and 28%.[4] 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 4th GST Council meeting, *"The rate structure should not be regressive in nature; items of mass consumption should not be taxed at higher rate."*

Continuity with the pre-GST tax structure was also a central design priority. The Council aimed to keep GST rates close to the combined federal and state tax rate that prevailed before the reform. Given the political and technical challenges of the GST roll-out, preserving continuity at the introduction stage was viewed as pragmatic. As noted in the minutes prepared ahead of the GST roll-out (see Appendix), *"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 specific band of pre-GST combined VAT and excise rates. For instance, goods with a pre-GST rate between 15% and 21% were assigned to the

[4]A few products were taxed at special rates outside the main GST brackets. Precious stones were taxed at 0.25%, while gold and other precious metals, including jewelry, were taxed at 3%. In addition, selected products were subject to a GST Compensation Cess on top of the 28% rate, mainly aerated waters, pan masala, tobacco products, and some motor vehicles.

18% bracket, with similar rules applied across categories. Thus, the initial GST rate schedule was calibrated to reproduce the effective pre-GST tax structure on average.

Product classification under GST The GST uses the Harmonized System Nomenclature (HSN) established by the World Customs Organization to define the tax rates. The HSN consists of three main levels: 96 2-digit codes known as chapters, 1,228 4-digit codes called headings, and 5,612 6-digit codes referred to as subheadings. India further differentiates some products within a 6-digit code by adding two additional digits, thus creating 8-digit codes. The large majority of GST rates for goods are set at the 4-digit level, though a few products have rates specified at a more detailed level.

Table A.1 presents an example of the HSN structure: "powder face" is an 8-digit code within the 6-digit code for "make-up or skin care powders", part of the 4-digit code for "make-up or skin care preparations". The corresponding 2-digit code is "essential oils and residuals, perfumery, cosmetic or toilet preparations". The granularity of the breakdown varies and HSN levels can be more or less specific than this example.

Reasons for the November 2017 Rate Cut We exploit a 10 percentage point tax cut applied to most products initially taxed at 28%. This reduction took effect on November 15, 2017, four and a half months after GST was introduced. Several political and institutional factors contributed to the decision. For the central government, there was a strong incentive to boost the popularity of the GST—its flagship economic reform—through a visible rate cut. The government faced criticism over the extent of rate differentiation in a reform framed as simplifying the tax system (see Appendix). These critiques threatened to erode the success of a reform whose core objective was the unification of state-level tax systems rather than a change in overall tax burdens. Public criticism specific to the 28% bracket had intensified: the rate was viewed as excessively high for this broad set of goods. As we show later, this perception is consistent with the distributional consumption patterns of these products.

From a revenue perspective, states, holding two-thirds of the votes in the GST Council, had limited incentives to oppose the cut. The central government had guaranteed them a minimum annual revenue growth of 14% for the first five years of GST, securing their support at the reform's launch. This compensation mechanism muted the fiscal costs of rate reductions for the states.

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

guarantee for states created a favorable environment for rate cuts. Yet, the fast timeline, relative to the broader GST introduction, and highly compressed decision process, made the reform's timing largely unexpected and left little scope for anticipatory responses. The GST Council met on November 10, announced the decision on November 11, and implemented the change 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 185,000 rupees (USD 2,870) in 2018. Although the data come from a single state, institutional features of the GST—such as the statutory framework, reporting requirements, IT infrastructure, and rate schedule—apply nationally. States primarily play an operative and enforcement role. Karnataka is considered as 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 (State of State Finances 2024), indicating that it is not an outlier in terms of fiscal capacity.

3 Data and Empirical Strategy

3.1 Tax Rates by Products

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 serve as the official reference for firms. These booklets list four-digit HSN codes and short product descriptions within each tax bracket. Most codes appear only once, but some appear multiple times when narrower subsets—usually six-digit HSNs—face different rates. Because these descriptions typically correspond to six-digit HSN product names, we link them to the six-digit classification using fuzzy matching. Comparing the 2017 and 2020 booklets yields the initial and final GST rates for each product.

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. When the minutes refer to categories narrower than the four-digit level, the descriptions allow us to identify the relevant six-digit HSN codes within each category.

The resulting dataset contains monthly GST product rates from July 2017 to March 2020. The vast majority of products correspond to an HSN-4 level (91%); the remaining products correspond to groups of HSN-6 codes (9%) (Table A.2).[5] Thereafter, anytime in the paper we refer to this specific taxonomy of goods defined for our purposes, we will use the italicized *products*.

3.2 Data

Our main dataset consists of GST tax returns for all registered firms in Karnataka from July 2017 to March 2019. We complement these data with a nationally representative household expenditure survey for India and barcode-level transaction data from a large retail chain.

Firm Level Tax Returns In our analysis period, firms with annual sales exceeding 2 million rupees (approximately USD 31,000) must register and file regular GST returns.[6] Firms below the threshold can either register for GST voluntarily or participate in the turnover tax composition scheme.[7] 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 more detailed tax return (GSTR-1), with frequency depending on firm size. Firms with annual sales above 15 million rupees (approximately USD 234,000) must fill 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–Product Tax Returns Firms with annual sales above 15 million rupees are required to file a product-level sales return (GSTR-1 Table 12), reporting monthly sales, taxes, quantities, and units by HSN code.[8] These data allow us to construct

[5]Formally, we define a product as a unique combination of an HSN-4 code and a GST rate path, characterized by its initial rate, final rate, and any rate-change date. Most HSN-4 codes map to a single rate path. When an HSN-4 contains multiple rate paths, we instead define products as groups of HSN-6 codes within that HSN-4 that share the same rate path.

[6]Using the official end-2017 exchange rate of 64 INR/USD for all currency conversion

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

[8]About half of firm–HSN–month observations are reported at the HSN-4 level, where most rate changes occur, and half at the finer HSN-8 level. We retain the firm-reported level of aggregation,

unit values as to proxy for pre-tax and tax-inclusive (consumer) prices. Our main analysis uses firm–product tax returns, which rarely exist in standard VAT data that typically do not distinguish products or report quantities. While only one third of firms in Karnataka list products, they account for three quarters of sales and tax payments, allowing product-level analysis to cover most economic activity.

Product Registration File To complement our analysis for smaller firms that do not file the product-level report, we draw on the GST registration file (REG-01). At registration, firms list the top five HSN codes they expect to sell. The list is non-binding and firms may ultimately sell other products, but it provides indicative information for those that file only the firm-level return.

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

India National Sample Household Survey The database contains household-level expenditure data for 345 survey product codes, which we match to HSN codes. The survey is nationally representative. We use the 2011–2012 wave, which is the most recent round available prior to the 2017 tax reform.[9] Since the survey does not collect income information, we rank households based on total per capita expenditure.[10]

3.3 Sample Construction and Data Cleaning

Sales to Final Consumers Given the credit–invoice structure of the VAT, taxes on sales to registered businesses and to final consumers have different implications. Registered firms recover VAT through input tax credits, while final consumers cannot.

as aggregating HSN codes reported separately would make quantities and implied prices harder to interpret, even within the same HSN-4 category.

[9]The subsequent available India National Sample Household Survey was conducted in 2022–2023, which is equally distant in time from the period analyzed.

[10]When ranking households, we exclude vehicles (personal transport equipment), which represent highly lumpy consumption items, free collections, gifts, and charitable transfers.

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 as well as 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, this information is available in monthly firm-level VAT returns, which report the overall share of sales to each client type. Our main analysis therefore focuses on firms selling primarily to final consumers, defined as those with a baseline (pre–November 2017) B2C sales share above 50%. This restriction excludes 50.5% of observations.[11]

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, the hypermarket data, or those classified as downstream by Antràs et al. (2012); and redefining the B2C share using the full sample period rather than just the baseline period.

Exclusion of Services We exclude services for three reasons. First, services are classified under the SAC system, which is much coarser than the HSN classification for goods: at the two-digit level, all services fall under Chapter 99, making it difficult to identify rate changes within subcategories. Second, GST rates for some services depend on attributes not captured by SAC codes—for example, hotel rates vary by room tariff and restaurant rates by whether the establishment is air-conditioned. Third, quantities are conceptually difficult to define for many services, and the aggregation of heterogeneous activities within broad SAC codes makes it challenging to construct consistent service-level quantity measures.

Data Cleaning Procedures To ensure data quality and a consistent sample, we drop the 9.6% of observations with invalid or missing HSN codes (the GST portal does not provide a drop-down menu) and those with reported sales missing quantities (see Table A.3 for details).

[11] Among excluded observations, 45.5% have a baseline B2C share below 50%, and an additional 5.0% first appear after November 2017. Because the baseline share is defined using pre-reform data, firms entering after this date are assigned a missing value. We prefer this approach because the reform may itself affect the B2C share, but relax this restriction when studying firm entry.

Moreover, as the GST portal does not automatically map HSN codes to statutory rates, firms can report infeasible rates—rates outside the set of statutory rates applicable to products listed under a given four-digit HSN code.[12] These reporting errors are broadly symmetric up and down, and only weakly predicted by firm characteristics such as size, location, consumer base, and experience (Appendix B). To abstract from this implementation issue, we exclude firm–product observations with rates falling outside a one–percentage-point band around the statutory rate range within the HSN-4 code in the pre-reform period. This restriction removes 3.7% of observations and ensures that differences between effective and statutory rates reflect meaningful variation rather than reporting errors.[13]

Finally, because quantities may be reported inconsistently and HSN codes may cover heterogeneous product bundles, we exclude firm–products with implausibly large price dispersion over time (price max/min > 100). This affects 1.4% of observations. To conservatively limit measurement error in unit values, we drop the entire firm–product rather than only the outlier month.

3.4 Treatment groups definition and description

Products with a rate cut The main analysis focuses on products experiencing a rate reduction from 28% to 18% in November 2017. 164 out of the 1363 products fall in this category (Table A.4). We analyze the 104 products that faced other rate changes separately, including those concerned by later reforms (notably the July 2018 changes), and products with changes from different brackets in November 2017.

Table A.5 tabulates the share of products in each tax bracket at the onset of GST (July 2017) and at the end of our sample (March 2020). In July 2017, 17% of products were taxed at 28%. By March 2020, 87% of these had moved to the 18% bracket, leaving only 2% of all products still taxed at 28%. Most of this adjustment occurred in November 2017. The near-elimination of the 28% bracket is consistent with minimal "cherry-picking" and reflects the government's objective of tax simplification.

[12] 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 of rates is considered feasible.

[13] Results are robust to including these observations (Appendix B). While their inclusion weakens the relation between statutory and effective rates, elasticities with respect to effective tax rates remain unchanged, indicating stable behavioral responses to taxes paid. We consider that our restrictions offer greater external validity, yet for evaluating the November 2017 reform specifically, statutory-rate elasticities estimated from the unrestricted sample are most relevant (see conclusion).

Products that remained at 28% were typically those taxed for externality or internality reasons, such as tobacco and soft drinks, akin to an excise. An exception are large cars which also stayed at 28%. By contrast, rate changes from other brackets seem partly motivated by equity and efficiency considerations. The table also shows that 18% is the most common rate: its product share rises from 40% to 52% over the period. The values in parenthesis show similar patterns when weighting by pre-reform sales.

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, pasta), hygienic goods (e.g., make-up, shaving), house goods (lamps, office supplies), and construction materials (e.g., wires, ceramics). No single product dominates the sample: none accounts for more than 10% of pre-reform observations and the top 30 products collectively represent 76% of treated-product observations

Products in the control group To construct a comparable control group to the treated products, our main specification focuses on products that were always taxed at 18% or 12%. Using the household survey, we analyze the consumption profile by tax bracket along the income distribution. Figure 2 reveals that products taxed at 12% and 18% exhibit a similar distributional consumption pattern to those initially taxed at 28%, given that in practice the assignment of goods to the 28% bracket was not very progressive. Products exempted and taxed at 5% are, on average, consumed proportionally more by poorer households, making them less comparable. Later, we perform several robustness checks by varying the definition of the control group.

As with the treated group, the control products are heterogeneous (Table A.7), although the top item"Medicaments in Measured Doses" accounts for a substantial share of control observations (22%), we will drop it in robustness tests. The top 30 products together represent 85% of the control sample. They span medical items (e.g., drugs, medical instruments), hygiene (e.g. soap, insecticides), household goods (e.g., brooms, computers, pens), construction materials (e.g., pipes, taps, glues), and food products (e.g., jams, bread and pastries, butter).

Descriptive statistics Table 1 summarizes the main sample used to analyze the November 2017 rate cut. It includes 844 products: 164 treated (from 28% to 18% rate) and 680 control products (that remained at 18% or 12% rate). Since firms

often sell multiple HSN codes, the data contain 449,368 unique firm–HSN pairs. The resulting panel is unbalanced: firms enter and exit, and their product mix varies over time. On average, a firm–HSN appears in 5.3 of the 21 months, yielding 2,379,017 firm–HSN–month observations, of which 314,077 before November 2017. Sales are skewed: before November 2017, the median monthly sales per firm–HSN are roughly 7,360 rupees (USD 115) in months when the product is sold, while the mean reaches 362,000 rupees (USD 5,660). Treated and control firm–HSNs appear similar along most pre-reform dimensions.

Most firms sell multiple products: Table 2 shows that out of the 19,800 firms present before November 2017, 66% reported selling more than one product in these four months, a share that rises to 88% weighted by sales. Given that most firms in the data sell multiple products, firms tend to sell both treated and control products: 42% (65% sales weighted) of firms sold at least one treated product in these months, and among these firms 83% (93% sales weighted) also sold a control product in the same month. Firms report selling 13.3 different HSNs per month on average in the four pre-reform months, out of which 2.5 are treated, 6.0 are control and 4.8 are 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% or 12% 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 faced the 10 percentage point rate cut (from 28% to 18%) in November 2017 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 for each period t the difference in the outcome between treated and control products

relative to the omitted period, October 2017 ($\beta_{Oct17} = 0$). We cluster standard errors at the firm and HSN-4 product 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 a post-reform period (April 2018–March 2019). This approach allows for a transition period (November 2017–March 2018), and captures a full year after the reform, thus providing a transparent summary measure of the treatment effect.

The availability of firm–HSN information allows us to control for differences in HSN sold across firms. Although different firms might not sell identical items within the same HSN code, the composition of items that a given firm reports under a given code is less likely to change over time. This mitigates concerns about within-HSN code composition or quality shifts. Moreover, the granularity of the firm–HSN–month panel, together with the fact that most firms sell multiple products, lets us exploit within-firm variation across products, thereby controlling for firm-level unobservables that could bias the estimated effects. Since the majority of firms are multi-product firms, incorporating this dimension is essential for the relevance of the analysis.

One concern is that the rate reduction may also affect control products sold by treated firms through spillovers—for example, if lower prices on treated products attract additional consumers. Such spillovers would violate the Stable Unit Treatment Value Assumption (SUTVA). To address this concern, we bound the plausible magnitude of spillover effects on control products.

Elasticities Building on this identification strategy, we estimate elasticities of quantities and sales with respect to the effective tax rate and consumer prices by instrumenting these endogenous variables with the interaction between treatment status and period indicators (pre, transition, and post), which captures the exogenous variation generated by the statutory rate reform. This strategy allows us to estimate the uncompensated elasticities of quantities, sales, and consumer prices with respect to the effective tax rate, as well as the uncompensated price elasticity of demand (the response of quantities to consumer prices). Appendix C presents a simple theoretical framework mapping our empirical estimates to the underlying structural parameters which we rely on for the optimal tax model.

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 fully map into effective tax rates (ETRs), the rates firms actually remit. This is because some firms already reported ETRs of 12 or 18% on treated products pre-reform and thus should not respond to the rate cut.[14] These patterns can reflect reporting errors induced by the complexity of the tax system or strategic misreporting to lower taxes. We thus first examine how the reform changed effective tax rates.

Figure 3.A shows the average STR and ETR for treated and control products. By construction, the STR of treated products fell from 28% to 18%. Prior to the reform, their average ETR was 26.9% and converged to 18% between November 2017 and February 2018. For control products, a mix of goods taxed at 12% and 18%, the STR remained at about 16%, while the average ETR slightly declined from 15.7% to 15.3%.

Figure 3.B reports event-study estimates with the full set of fixed effects (Equation 1).[15] The reform reduced ETRs by 8.1 percentage points, relative to a 10-point statutory cut, implying 81% pass-through to effective rates. This incomplete pass-through points to frictions in multi-rate VAT systems: in our optimal tax simulations (Section 6), we treat this baseline misreporting as a fiscal cost of rate differentiation.

Price Incidence The extent of tax pass-through determines whether the gains from the tax cut accrue to consumers or producers. Figure 4.A shows that consumer prices decline gradually after the reform, stabilizing after three months at about 0.035 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 53 percent (95 percent CI: 33–73 percent), allowing us to reject both zero and full pass-through.[16] Overall, just over half of the tax cut is passed through to consumers.

Prices respond to realized changes in effective tax liabilities rather than statutory

[14] Given the sample restrictions in Section 3.3, this arises among legally feasible rates within an HSN-4 code. We define an "error" as a case in which multiple HSN-6 categories within the same HSN-4 code carry different rates, and the firm reports the lower rate despite listing an HSN code legally subject to the higher rate.

[15] Figure A.1 reports the event-study estimation of log ETR.

[16] Relative to the 0.081 log-point (10 percentage point) statutory rate cut, pass-through is 43 percent.

changes per se. Figure A.2 shows that firm–products with initially misaligned ETRs experience only small ETR declines and no detectable price response. By contrast, when firms initially applied the correct rates, prices fall by 0.040 log points alongside a 0.069 log-point ETR decline, implying 55 percent pass-through. This also implies that pass-through estimates based on CPI data and statutory rate changes may understate true pass-through when effective tax adjustments are incomplete, introducing attenuation bias in reduced-form estimates.[17]

Quantities We next examine how quantities respond to the consumer price reduction induced by the rate cut. Direct observation of quantities is uncommon, as in many countries VAT returns require firms to report only sales values. Figure 4.B presents estimates of Equation 1 for log quantities. In the first three months following the reform, quantities mirror the price dynamics: as prices gradually decline, quantities increase. After this adjustment period, treated products display a persistent increase of 0.032 log points on average, marginally significant at the 5 percent level.

Relative to the 10 percentage point statutory rate reduction (0.081 log points), the estimates imply a quantity elasticity with respect to the statutory tax rate of -0.39. Because statutory changes translate only partially into effective tax reductions, this understates behavioral responses to realized tax changes. Instrumenting ETR changes using treatment status interacted with pre, transition, and post reform periods yields an ETR elasticity of -0.46. Instrumenting consumer prices with the same interaction implies a demand elasticity of -1.01.

Sales Pre-tax sales define the VAT base absent non-proportional input adjustments.[18] 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 5.A shows that producer prices increase after the rate reduction. Together with the increase in quantities, this raises pre-tax sales by 0.063 log points (Figure 5.B). The implied elasticity of the tax base with respect to the effective tax rate is -0.94.

[17] This concern may be particularly relevant in developing countries, where gaps between statutory and effective tax rates can be substantial due to compliance frictions and reporting errors.

[18] Firms' final tax liability depends on value added—sales net of input tax credits—which we analyze in Section 4.6, as it is observed only at the firm level rather than at the product level.

Extensive Margin The rate reduction may also affect the tax base along the extensive margin if firms begin selling—or start reporting sales of—treated products after the reform. To study this channel, we construct a balanced firm–product panel with a binary indicator for positive monthly sales. The reporting rate is defined as the probability that a firm reports positive sales of a product in a given month, conditional on having sold that product at least once during the sample period.[19] Figure 6.A plots the reporting rate for treated and control products. The probability that a firm sells a given product increases steadily over the sample period, but treated and control products follow nearly identical trajectories. This shared upward trend is consistent with the expansion of the GST system and overall economic growth during this period.

We next estimate Equation 1 using a binary reporting indicator to capture within-firm extensive-margin responses. Figure 6.B presents the estimated treatment effects. Products subject to the rate cut exhibit no meaningful change in their likelihood of being sold. The point estimate is close to zero (it is actually slightly negative) and precisely estimated: we can reject increases exceeding 0.1 percentage points—less than a 1 percent change relative to the 15 to 20 percent baseline reporting rate.

4.2 Product Misclassification and Substitution

Defining misreporting and substitution A first concern in multi-rate commodity tax systems is that their complexity create opportunities for firms to misreport higher-taxed goods as lower-taxed ones, thereby eroding the intended fiscal and distributional effects of rate differentiation. In our setting, firms may have reported goods taxed at 28% 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 discontinue such behavior.[20]

A second concern is that differentiated tax rates distort relative prices: in our

[19] We distinguish three extensive-margin responses: (i) firm–product entry or exit, when an incumbent firm starts or stops reporting a given HSN code; (ii) firm exit, when a firm stops reporting all HSN codes; and (iii) firm entry, when new firms begin filing firm–product tax returns. The main sample, constructed using pre-November-2017 B2C shares, mechanically excludes margin (iii) for firms not reporting pre-reform, and only captures margins (i) and (ii). To allow for all three margins, we re-define the sample using average B2C shares over the full period. Results using the original sample are similar, reflecting that most firms were already reporting pre-reform (Figure A.3.B).

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

setting, the reform may shift demand toward goods whose rates were reduced and away from related products. Unlike mislabeling, such substitution reflects a genuine reallocation of consumption and therefore carries welfare implications beyond its revenue effects. Although substitution and mislabeling 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 a decline in reported quantities and sales of untreated related goods following the rate cut.

Defining related products To make progress, we assume misreporting occurs toward related products in the taxonomy, for which relabeling is easier to justify during an audit and may be framed as an honest mistake. The GST relies on the Harmonized System of Nomenclature (HSN), with most rate differentiation at the four-digit level, which we exploit to define relatedness. We partition untreated products into: (i) unrelated HSN-4 products, whose HSN-2 chapter contains no treated goods; and (ii) related HSN-4 products, which share an HSN-2 code with at least one treated good. This approach isolates the products most prone to misclassification while providing a plausible control group (see Table A.8). Substitution is also more likely within HSN-2 categories, although economic substitutability varies widely even within these groups. We therefore later restrict the set of potential substitutes to construct a tighter comparison group of products for which substitution appears most plausible.

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

Related products results We estimate Equation 1 at the firm–product level, comparing related to unrelated products within firms. Under a misreporting mechanism, quantities and sales of related products should fall relative to unrelated products. Figure 7 presents event-study estimates for quantities, sales, consumer prices, and reporting rates. We find no evidence of mislabeling: the coefficient on quantities sold of related products is actually positive and statistically insignificant (0.016 log points, s.e. = 0.021), such that we reject declines larger than 2.5 percent at the 95 percent

confidence level (Panel A). Sales move in parallel (Panel B), in line with the absence of price responses (Panel C). Finally, we detect no adjustment along the extensive margin: changes in reporting rates are economically small and precisely estimated around zero (Panel D).

We further probe the absence of relabeling (and substitution) through three robustness exercises presented in Table A.10. First, mislabeling may arise only in HSN–2 chapters where treated products account for a substantial share of activity. Restricting the sample to chapters in which treated goods represent at least the median pre-reform sales share ($\approx$ 15 percent) yields estimates nearly identical to the baseline. Second, responses may be concentrated among firms selling both treated and related products. Limiting the sample to such firms again produces similar results. Finally, allowing for substitution toward goods initially taxed at 5 percent—excluded from the main control group—does not alter the conclusions.

Machine Learning Matching of Related Products Our estimates rule out economically meaningful misreporting or substitution at the level at which GST rate differentiation is implemented in India—namely, the four-digit HSN classification—although they cannot rule out responses within narrower product categories. This is in line with the literature, which only finds substitution among narrowly defined, highly substitutable goods rather than across broader product classes.

To probe these finer margins, we use a machine-learning procedure to identify, for each treated product, its closest substitute and complement. Under the most restrictive matching criteria, we find small positive and marginally significant effects for complements, and small negative but statistically insignificant effects for substitutes. When the algorithm is able to assign to a treated products a substitutes or a complement, these generally lie at a level of disaggregation closer to the 8-digit HSN classification, suggesting that any cross-price responses occur at a finer level of product similarity than the one targeted by the reform and are unlikely to be first-order in India's multi-rate GST system, which typically differentiates goods at the HSN-4 level.

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 may generate firm-level spillovers, affecting untreated products and violating SUTVA. To assess this

possibility, we exploit heterogeneity in firms' pre-reform exposure to treated products. If spillovers are quantitatively important, control products sold by firms with greater baseline exposure should exhibit larger post-reform increases in quantities or sales. Before November 2017, the median control product was sold by a firm for which treated products accounted for 3.5% of pre-reform sales (mean: 11.7%). Among firms selling at least one treated product, the median exposure rises to 5.0% (mean: 14%). 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 mean share, which for interpretation we rescale 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 in 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.4). These are an order of magnitude smaller than the corresponding direct effects of the rate reduction (−0.035, 0.032, and 0.063). As a result SUTVA violation due to within-firm spillovers could account at most for 10–20% of the estimated treatment effects.

4.4 Robustness to Alternative Samples and Specifications

Figure 8 summarizes key robustness checks (see Table A.11 for the full set). The first set varies the composition of the control group—initially comprising products always taxed at 12 or 18%—by broadening or tightening the pool of control products: (i) including products taxed at 5%; (ii) restricting the sample to products taxed at 18%; (iii) excluding *medicament*, a large regulated control product; and (iv) either restricting comparisons to products within the same HSN-2 chapters or excluding those chapters to limit indirect exposure to the reform.

The second set examines robustness to our definition of final-consumer exposure.

As registered suppliers can claim tax credits on their inputs, our focus is on the incidence of rate changes on final consumers. We test several alternatives: (i) restricting the sample to downstream products as defined by Antràs et al. (2012); (ii) to products appearing in the household survey; (iii) to products sold in the hypermarket data; and (iv) to firms with over 80% of their pre–reform sales to final consumers.[21]

The third set examines the data consistency filters. We reintroduce in the sample (i) firm-products with extreme price variation and (ii) firm-products applying unfeasible rates for the HSN-4 code.

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 February 2018.

The fifth set considers robustness to weights and aggregation. In the baseline, we keep observations at the firm-reported HSN level, since quantities (and prices) cannot be coherently aggregated across HSN codes. By contrast, sales can be aggregated, which makes this robustness check particularly informative for that outcome. We consider three alternatives: (i) weighting observations by pre–November 2017 sales; (ii) aggregating firm-level sales at firm-*product* level, which tax rates are defined (HSN-4) even when firms report at narrower levels; and (iii) aggregating all treated and all control products within each firm.

Across all sets of robustness checks, the estimates remain close to those in the main specification, reinforcing the robustness of these results. Moreover, the findings 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 close to normally distributed (Figure A.5).

4.5 Second Reform: Rate reduction in July 2018

Thirty-three products that remained in the 28% tax bracket after the November 2017 reform experienced a rate cut to 18% in July 2018. Although fewer products were impacted, the reform complements the earlier reform: it occurs well after GST implementation and applies to higher-ticket items representing a non-trivial share of

[21] In the baseline, firms are classified as final-consumer facing if a majority of their pre-reform sales are to final consumers. Restriction (v) instead uses the share of sales to final consumers computed over the entire sample period, rather than only pre–November 2017.

sales (See Table A.4).[22] We estimate effects using the same specification as for the November 2017 reform, where the control group consists of products that are taxed at 12 or 18% during our complete time-frame. (See Table A.12 for summary statistics.)

The results are similar despite differences in timing and composition (Figure A.6). Consumer prices decline by 0.028 log points, implying a 55% pass-through relative to the 0.052 log-point tax change. Quantities increase by 0.037 log points and pre-tax sales rise by 0.062 log points. The elasticities of quantities and sales with respect to the ETR are -0.73 and -1.22, compared with -0.58 and -1.07 in the first reform. Overall, the July 2018 episode reinforces the November 2017 findings and extends them to a distinct set of higher-value goods.

4.6 Firm Level Results and Heterogeneity by Firm Size

Our main firm–product design exploits within-firm variation to absorb time-varying firm-level shocks and closely match the policy margin, as the reform applies at the product level. However, this strategy has two limitations. First, value added-the GST tax base-is only observed at the firm level. Second, only medium and large firms are required to report product-level GST returns, excluding most small firms (82% of registered firms, representing 26% of total sales; Table A.14).

The GST system also provides two firm-level data sources: (i) a consolidated monthly return filed by all firms reporting value added (GSTR-3), and (ii) a registration form listing up to five main products sold by the firm. Using this information, we implement a complementary firm-level design in which firms are classified as treated if they list at least one product affected by the November 2017 rate cut, and as controls otherwise. This approach allows us to include the full population of registered firms regardless of size. We first estimate reduced-form effects of reform exposure on firm outcomes (an intent-to-treat specification). We then instrument firm-level tax rate changes using this predetermined exposure to estimate the causal effect of tax changes on firm outcomes. [23]

We report the firm-level event study results in Figure A.7. The registration form contains meaningful information: treated firms experience a post-reform decline in

[22] This later reform affected electrical appliances (e.g. vacuum cleaners, refrigerators), house renovation items (e.g. paints, varnishes) and household goods (e.g. suitcases, mirrors) (See Table A.13).

[23] To maintain comparability with the firm–product analysis, we exclude firms listing products treated in other periods, exempt goods, and services. The resulting sample includes 174,000 firms reporting on average 2.5 products at registration (Table A.15). Overall, 28% of firms are classified as treated.

ETRs of 3 percentage points, compared to 0.7 percentage points for control firms (Panel A).[24] This results in a first stage redution ETR reduction of 2.3pp (Panel B). Relative to control firms, treated firms increase reported sales by 0.036 log points (Panel C); value-added increases by 0.054 log points (Panel D). These yield elasticities of reported sales and value-added with respectively to the tax rate of respectively of -1.69 and -2.49. [25] The higher elasticity of value added compared to sales is consistent with the presence of fixed costs in inputs.

Comparing firm-product and firm-level estimates The firm-level strategy yields somewhat larger elasticities, but it covers all registered firms rather than only those filing detailed product-level returns. To facilitate comparison, Table A.16, column 3, applies the same firm-level specification to the subset of firms filing product-level returns. The resulting sales elasticity (-1.16) and value-added elasticity (-1.67) are similar to the product-level sales elasticity.[26]

Firm-Size Heterogeneity These results imply that sales of smaller firms—those filing only the firm-level return—are more tax elastic. This is consistent with the idea that their customers can switch to informal stores, or that their own reporting and evasion behavior is more sensitive to tax rates. Table A.16 confirms this pattern: firms filing the (simplifying) firm-level return exhibit a sales elasticity of -2.61, compared to -1.89 for firms filling the detailed firm level return, and -1.16 for firms filing the product-level returns, revealing a clear gradient of sales elasticities by firm size segment.

5 Scanner Level Data Analysis

Our analysis faces three main limitations: a short pre-treatment window close to the GST's introduction; imperfect restriction to final-consumer sales; and potentially coarse

[24] We define STRs as the average statutory rate across listed products (e.g., a firm reporting one product taxed at 18% and another at 28% is assigned an STR of 23%). Before November 2017, treated firms had an average STR of 23.1%, compared to 13.7% for control firms.

[25] Value added is defined as the monthly difference between sales and intra-state purchases, inferred from firm-to-firm transactions. We do not observe inputs bought through inter-state transactions. The analysis is restricted to firm-months with positive value added.

[26] An alternative approach defines firm-level treatment using baseline exposure to treated products in the firm-product returns. This measure cannot be constructed for smaller firms and yields a less transparent intensity design. For directly comparable outcomes, however, the estimated elasticities are similar; see column 4 of Table A.16.

product definitions at the 4-digit HSN level. To address these issues, we complement the tax data with a proprietary barcode-level dataset covering all transactions of a large supermarket chain in Karnataka. Although limited to a single firm, this dataset directly overcomes these limitations: it provides a long pre-treatment period;[27] all 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 effective tax rates. Treated barcodes are taxed at 28% between August and October 2017 and at 18% between December 2017 and March 2020; control barcodes are taxed at 12% or 18% throughout. We restrict the sample to barcodes sold in all periods, as many appear intermittently due to seasonality, product updates, or discontinuation. To address seasonality, we include department-by-month fixed effects. Identification therefore compares barcodes experiencing a rate cut to barcodes continuously taxed at 12% or 18% within store departments. The specification mirrors the firm–product analysis, with department–time fixed effects replacing firm–period fixed effects and barcode fixed effects capturing the product dimension. We estimate the following regression:

$$y_{bdt} = \sum_{t=\text{Aug16},\, t\neq\text{Nov17}}^{\text{Mar19}} \beta_t \cdot (\text{treat}_b \times \text{month}_t) + \alpha_b + \gamma_{td} + \epsilon_{bdt} \tag{3}$$

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 tax category level.

Table A.17 reports the number of treated and control barcodes by store department (defined by the hypermarket). In total, there are 1,508 treated barcodes across 10 departments, most of which belong to the *Personal Care & Cosmetics* department, followed by *Home Care* and *Breakfast* departments. The control group includes 2,491 barcodes, more evenly distributed across the 10 departments.[28]

[27] 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.

[28] Table A.18 reports the number of barcodes at the hypermarket's *Family* level. The leading *Families* mirror the HSN-4 classification (Table A.6). For instance, the *Personal Care & Cosmetics* department contains the *Hair Care* and *Skin Care* families, corresponding to HSN-4 *Hair Preparations*

Figure 9.A shows the event-study estimates from equation 3 for consumer prices.[29] 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. After the November 2017 rate cut, consumer prices fall, by an estimated 0.024 log points past the initial transition period. This decline is smaller than the corresponding reduction in effective tax rates—nearly equal to the 0.081 log-point statutory cut—implying a pass-through rate of 30%. The estimates are precise enough to reject both zero and full pass-through at the 1% level.

In the hypermarket data, effective and statutory tax rate reductions coincide, as treatment is defined using the retailer's self-reported tax brackets. Given the chain's size, centralized structure, and standardization, we are confident in the accuracy of its product classification and tax reporting. The hypermarket's sales place it within the top one percent of firms in the administrative data, and its lower estimated pass-through is consistent with the sales-weighted result using the administrative tax data, which indicates that larger firms exhibit lower pass-through (Figure 8).

Consumer price adjustment in the months following the reform is non-linear: prices initially fall, then rise, and fall again. This pattern reflects the hypermarket's strategic use of discounts during the transition. Consumer prices include temporary promotional discounts, and Figure 9.B plots their contribution to price changes. Treated products exhibit a reduction in price due to discounts of the same magnitude as the initial price decline, indicating that short-run pass-through is entirely driven by increased discounting. Rather than implementing an immediate permanent price cut, the hypermarket initially passes through the tax reduction via temporary, salient discounts. After two months, discounts unwind: consumer prices rise, allowing the retailer to capture the tax cut, before declining again as the price reduction becomes embedded. This behavior is consistent with greater pricing sophistication—and potentially higher market power—than that of the average firm in the tax data.

Figure 9.C reports the quantity results. Pre-trends are stable in the year preceding the rate cut. Following the reform, quantities sold of treated barcodes slightly increase, with a point estimate of 0.042 log points, which is not statistically significant. Given

and *Make-Up*, two of the key treated *products* in the tax data.

[29] Before the GST's introduction in July 2017, only state VAT is observed, as central duties (CENVAT) were embedded in prices and remitted upstream, preventing recovery of pre-tax prices or effective tax rates. This limitation does not affect quantities or consumer prices, which are invariant to tax reporting and include all tax components.

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 entirely independent data source, corroborates our main results of partial pass-through and moderate quantity responses.

6 Optimal Commodity Rate Differentiation

This section examines the implications of our results for optimal commodity tax design, asking whether luxury goods should be taxed more heavily, and by how much. We first analyze efficiency and equity considerations separately, then combine them to derive the optimal degree of rate differentiation between luxury and standard goods. We finally extend the analysis to account for income taxes and transfers.

6.1 Efficiency

Tax Revenues The tax revenue impact of a tax rate change depends on the elasticity of the tax base with respect to the statutory rate change. The change in revenue can be written as:

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

Behavioral responses partly offset the mechanical revenue loss from the rate reduction. Absent behavioral effects, revenue would fall one-for-one with the rate cut. Using our baseline sales elasticity (proxying for the tax base) with respect to the ETR, $\varepsilon_s = -0.94$, and $t = 0.28$, the revenue change is $-(1 - 0.94 \times \frac{0.28}{1+0.28}) \approx -0.79$. Thus, one dollar of mechanical revenue loss translates into \$0.79 of actual loss, implying a behavioral offset of \$0.21. At the upper bound of the 95% confidence interval, $\varepsilon_s = -1.45$, behavioral responses offset \$0.68 per mechanical dollar. Using instead the STR elasticity, which is one-third smaller ($\varepsilon_s^{STR} \approx -0.77$), the revenue loss is of -\$0.83. Thus, the rate cut led to a large revenue loss, with only modest compensation from higher reported sales.

Marginal Cost of Public Funds The deadweight loss of indirect taxes is given by the amount of transactions that cease to exist as a result of taxation. Unlike tax

revenues, which also depend on how pre-tax prices adjust and affect the tax base (proxied with pre-tax sales), efficiency costs depend on the quantity elasticity, which captures the number of transactions forgone independently of the distribution of the tax burden between consumers and firms. Following the sufficient-statistics approach, and under simplifying assumptions of quasi-linear utility, price-taking firms, and no cross-price effects, the welfare change from a small tax change is:

$$dW = dM \cdot \frac{t}{1+t} \cdot \varepsilon_q$$

where $\varepsilon_q \geq 0$ denotes the absolute value of the uncompensated quantity elasticity with respect to the tax rate.[30] The Marginal Cost of Public Funds (MCPF), defined as the total private cost per dollar of revenue raised, is:

$$\text{MCPF} = 1 + \frac{\frac{t}{1+t} \cdot \varepsilon_q}{1 + \frac{t}{1+t} \cdot \varepsilon_s}.$$

With $\varepsilon_q = 0.46$, $\varepsilon_s = -0.94$, and $t = 0.28$, the MCPF is 1.13. This means that raising one dollar of revenue through indirect taxation reduces private welfare by about \$1.13. Even at the upper bound of the 95% confidence interval for ε_q, the welfare cost per revenue dollar remains modest.

6.2 Equity

GST Progressivity at Baseline Goods taxed at 12%, 18%, and 28% are consumed disproportionately by richer households, while exempt and 5%-rated goods are consumed relatively more by poorer households. The share of total consumption spent on goods in the top three tax brackets rises from 21% in the bottom decile to 41% in the top decile (see Figure A.8). The average effective consumption tax rate therefore increases monotonically across the income distribution, from 6.0% in the bottom decile to 9.3% in the top decile (Figure A.9).

The quasi-elimination of the 28% tax bracket (with the combined November 2017 and July 2018 rate cuts) was mechanically regressive, reducing average tax burdens more for richer households (Figure A.10). The reduction amounts to 0.7 percentage

[30] Under quasi-linear utility there is no income effect, so the uncompensated and compensated elasticities coincide. With general preferences the compensated elasticity would be more appropriate for measuring deadweight loss.

points for the bottom decile but reaches 1.3 percentage points for the top.

Optimal Product Assignment Despite the relative progressivity of the initial GST rate assignment, it falls well short of the most progressive allocation. Figure A.11 normalizes expenditure shares in each tax bracket to those of the bottom decile and shows that goods taxed at 18% exhibit a steeper income gradient than those taxed at 28%: the top decile devotes 2.5 times more of its expenditure to 18%-rated goods, compared with 1.8 times more for 28%-rated goods.

This leaves substantial scope for a more progressive rate assignment. Some goods—such as PCs, cameras, and musical instruments—are consumed by top-decile households at levels more than 100 times those of the median, while other goods initially taxed at 28% (e.g., soap, shampoo, shaving products, and chocolate) display much flatter consumption profiles. To illustrate this scope, we construct an alternative rate assignment based on Engel-curve slopes. Before the reform, goods taxed at 28% accounted for 11% of total consumption. Holding this share fixed, we reassign products to the 28% rate starting with those exhibiting the steepest Engel slopes until cumulative consumption reaches 11%. Displaced goods are assigned the weighted-average tax rate of the replacing goods, keeping the overall average rate unchanged.

Under the original assignment, top-decile households consume 28%-rated goods 1.8 times more (as a share of total expenditure) than the bottom decile; under the progressive reassignment, this ratio rises to nearly 16 (Figure A.12). Implementing the optimal reassignment for just 11% of total consumption would raise the top-decile average tax rate by 1.4 percentage points, while reducing it for all households below the 8th decile (Figure A.13).

6.3 Equity-Efficiency Trade-off

Theory We derive a general expression for optimal indirect tax rates using the many-person Ramsey rule (Diamond 1975). Optimal rates rise with the inverse of demand elasticities and fall with a good's social value: more elastic goods should be taxed less, while goods with lower social value should be taxed more.

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, that is, the social cost of raising one additional unit of government revenue. Firms operate under constant returns to scale and perfect competition.

For simplicity of exposition and in light of the results on substitution, we initially assume zero cross-price elasticity and present the formula allowing for it in the Appendix D. Under these assumptions, the optimal indirect tax on good j is:

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

where g^i denotes the generalized social marginal welfare weights, a_j^i is the share of household i in the total consumption of good j , defined as $a_j^i = \frac{x_j^i}{x_j}$, with x_j^i denoting 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 ε_j^p is the elasticity of consumer prices with respect to taxes, and ε_j^s is the uncompensated elasticity of pre-tax sales with respect to the tax (i.e., the tax base elasticity).[31]

The optimal tax on good j is higher when (i) demand for the good is less elastic (that is, when ε_j^s is closer to zero), (ii) the covariance between g^i and s_j^i is higher, that is, when richer households consume relatively more of the good, and (iii) the elasticity of prices with respect to taxes is lower, implying that consumers bear a smaller fraction of the burden (iv) the marginal value of public funds is higher.

The equity–efficiency trade-off arises along two dimensions. First, the government trades off tax-base distortions (ε_j^s) against revenue needs, captured by the marginal utility of public funds (μ). This trade-off determines the overall level of indirect taxation, even when goods share identical elasticities and social values. Second, there is a trade-off between a good's elasticity and its social weight: if luxury goods are more elastic than essential goods, efficiency and equity considerations pull in opposite directions, making the desirability of higher rates on such goods ambiguous.[32]

[31] The derivation is provided in Appendix D. For a direct mapping with the empirical estimation, we keep the formula in terms of uncompensated elasticities. This could be rewritten in terms of compensated elasticities by applying the Slutsky equation as Diamond 1975.

[32] Given our objectives, the formula abstracts from Pigouvian motives for indirect taxation. For some goods disproportionately consumed by higher-income households—such as medicines or school tuition—governments may have other reasons to apply lower tax rates.

Calibration We observe s_j^i in the household consumption survey and set social marginal welfare weights g^i equal to the inverse of total expenditure. Demand elasticities are taken from the empirical estimates. Although we find no statistically significant relationship between quantity elasticities and Engel-curve elasticities (Appendix E), we adopt a conservative approach and use the point estimate of 0.217, implying that goods with steeper Engel curves are less elastic.[33]

To discipline the model, we calibrate μ so that the optimal tax rate for goods initially taxed at 18% equals 18%. Interpreting the 18% rate as optimal allows us to focus on optimal rate differentiation rather than the level of taxation.

Optimal Rate Differentiation We calibrate optimal tax rates for each survey product. Figure 10 reports these estimates, ranking products by their Engel-curve elasticity. Negative optimal rates (i.e., subsidies) are winsorized at zero. Optimal tax rates rise with the Engel-curve elasticity but nonlinearly, as quantity elasticities also increase with the Engel curve. The schedule peaks at an Engel elasticity of about 3, corresponding to a tax rate of 50%, beyond which higher demand elasticities dominate the gains from greater progressivity. Product names illustrate examples of important products along the distribution of optimal rates: potatoes (0%), tea leaves (5%), chicken (12%), cooked snacks (18%), goat meat (28%), and gold ornaments (50%).

We also consider the more realistic scenario with a limited number of tax brackets. Holding fixed the sales value of goods assigned at baseline to the standard 18% versus luxury 28% brackets (11% of sales value), we reshuffle goods across brackets optimally. This yields an optimal tax rate for the reclassified luxury goods of approximately 38%.

6.4 Optimal Luxury Taxes with Income Taxes and Transfers

Income Taxes and Transfers Canonical public-finance theory prescribes that redistribution should be achieved through incomes taxes and transfers, implying no redistributive role for other instruments such as differentiated consumption taxes (Atkinson and Stiglitz 1976). This result, however, relies on the ability of governments to implement broad and flexible income-tax and transfer schedules across the full income distribution—an assumption at odds with the reality of most developing

[33]The elasticity of good j is given by $\varepsilon_j = \varepsilon_0 + \delta\left(\hat{\beta}_j^{\text{Engel}} - \hat{\beta}_{\text{treated}}^{\text{Engel}}\right)$, where δ is the slope coefficient estimated in equation (20).

countries. Once constraints are considered, and statutory policy choices are viewed as endogenous to them, then the use of consumption taxes for both revenue and equity objectives can be a sound and second-best optimal policy (Huang and Rios 2016).

As in many low- and middle-income countries, India's personal income tax (PIT) covered only a small share of the population. In 2011–2012, the year of the household survey, roughly 4 percent of adults paid PIT. The PIT featured a progressive statutory schedule above an exemption threshold of INR 180,000 (USD 2,800), with marginal rates of 10, 20, and 30 percent applied to successive income brackets. The top 30% rate applied to income above INR 800,000 (USD 12,300). Taxable income was further narrowed by numerous exemptions and deductions.

Model extension We extend the analysis by considering counterfactual expansions of India's PIT system relative to the benchmark, examining both broader coverage and higher tax rates. Throughout, we compute optimal indirect tax rates conditional on the PIT system—that is, we do not solve for the joint optimal design of indirect and personal income taxes. Instead, we adjust the optimal commodity rate calibration to reflect post-PIT income. By reducing income at the top of the distribution, the PIT affects optimal indirect tax rates through two channels. First, lower disposable income reduces consumption, which we model using the Engel curves estimated from the household survey. Second, lower post-tax income increases the social marginal welfare weights assigned to top income groups, , as their post-tax income is reduced.

As discussed in Section 6.2, the initial structure of indirect tax rates was far from maximally progressive. To isolate the role of rate differentiation, we apply an optimal tax to the goods with the highest Engel-curve slopes (hereafter, luxuries), holding fixed the total consumption share subject to the top rate at 11% (the share of goods initially taxed at 28%). We then compare the resulting optimal rates to those applied to goods initially taxed at 18%, and define the difference as the luxury tax premium.

Results Figure 11 presents estimates of optimal luxury rate premiums under alternative personal income tax (PIT) scenarios. The horizontal axis varies the PIT exemption threshold, expressed as a fraction of the initial exemption thresholds and mapped to the corresponding income percentiles. Moving from right to left, the figure ranges from a scenario without a PIT, to progressively lower PIT thresholds, with all thresholds reduced proportionally. In the FY 2011 baseline scenario, we use the PIT

thresholds in place in 2011, adjusted to account for standard deductions. In the 80% scenario, all thresholds are reduced proportionally by 20%. Each line corresponds to a different PIT schedule. The blue line represents the PIT schedule implemented in India in FY 2011–2012; the point on the horizontal axis corresponding to the FY 2011 thresholds therefore reflects the actual policy environment in that year, with both thresholds and marginal tax rates unchanged. All other points on the blue line, as well as those on the other lines, represent counterfactual scenarios that vary thresholds, rates, or both. The red line illustrates a scenario in which all PIT marginal tax rates are increased by 20 percentage points. The green line excludes from the optimal indirect tax calculation the share of income subject to PIT, which is equivalent to assuming a 100% marginal tax rate on income above the PIT threshold.

The optimal premia are largely flat at around 20% as income tax thresholds are lowered (blue line), implying an optimal rate of 38% for luxury goods, and a luxury premium of 20 percentage points relative to the 18% rate. This flatness reflects the limited role of the PIT: even with lower thresholds, average PIT rates are low,[34] leading to modest reductions in consumption and, thus, little change in luxury premia.

The optimal luxury premium slightly lowers when PIT marginal rates are higher, at about 18% (red line). Yet, PIT rates have to increase much more for luxury premia to meaningfully fall. We consider a counterfactual in which income subject to PIT is excluded from the indirect tax base (green line). In this case, indirect taxes are optimally levied only on income not covered by PIT, resulting in a much larger decline in luxury premia—to about 15 percentage points at FY 2011–2012 thresholds. As PIT coverage expands, rate differentiation falls further, reaching roughly 7% when the PIT extends to the 73rd income percentile (a 50% threshold reduction).

7 Conclusion and Policy Implications

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

[34]The income weighted average rate, conditional on paying the PIT, is about 10%.

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

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. (2023). Informality, Consumption Taxes, and Redistribution. *The Review of Economic Studies*, rdad095. https://doi.org/10.1093/restud/rdad095

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., & Carloni, D. (2019). Who Really Benefits from Consumption Tax Cuts? Evidence from a Large VAT Reform in France. *American Economic Journal: Economic Policy*, *11*(1), 38–63. https://doi.org/10.1257/pol.20170504

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

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

Carbonnier, C. (2007). Who pays sales taxes? evidence from french vat reforms, 1987–1999. *Journal of Public Economics*, *91*(5-6), 1219–1229.

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.

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

Gaarder, I. (2022). Incidence and distributional effects of value added taxes. *The Economic Journal*. https://doi.org/https://doi.org/10.1111/ecoj.12576

Garg, S., Ghosh, P., & Tan, B. J. (2023). Within firm supply chains: Evidence from india. *Journal of International Economics*, *144*(100). https://doi.org/10.1016/j.jinteco.2023.103793

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

Harris, T., Phillips, D., Warwick, R., Goldman, M., Jellema, J., Goraus, K., & Inchauste, G. (2018). Redistribution via VAT and cash transfers: An assessment in four low and middle income countries.

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

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

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

Liu, Y., Barrett, C. B., Narayanan, S., Das, U., Chau, N., Soundararajan, V., Sheahan, M., & Deininger, K. (2018). *Impact evaluation of the national rural employment guarantee scheme in india* (3ie Grantee Final Report). International Initiative for Impact Evaluation (3ie). New Delhi.

Londoño-Vélez, J., & Ávila-Mahecha, J. (2024). Behavioural responses to wealth taxation: Evidence from colombia. *The Review of Economic Studies*, *92*(4), 2624–2655. https://doi.org/10.1093/restud/rdae076

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.

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

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

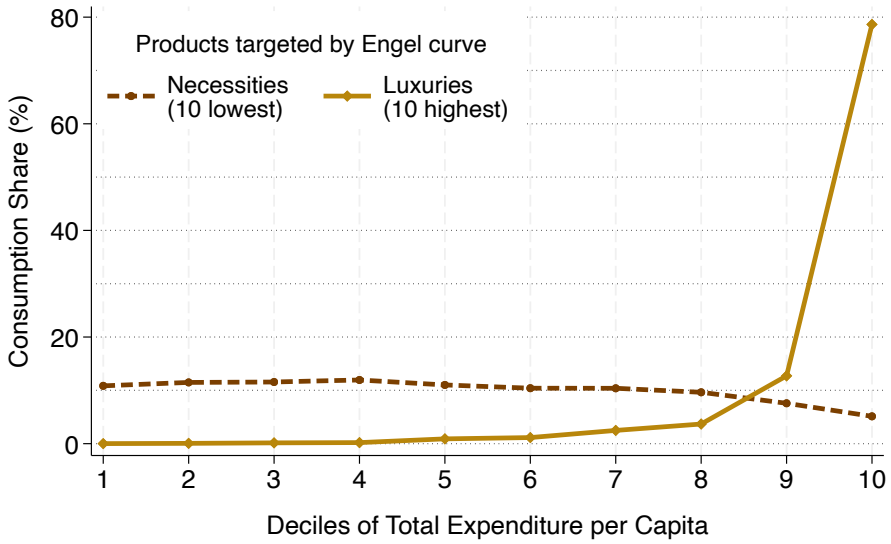
State of State Finances. (2024). *State of State Finances* (tech. rep.) [FY 2024–25]. PRS Legislative Research.

Velayudhan, T., & Brusco, G. (2022). *Does the Informal Sector Escape the VAT?* (Working Paper).

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

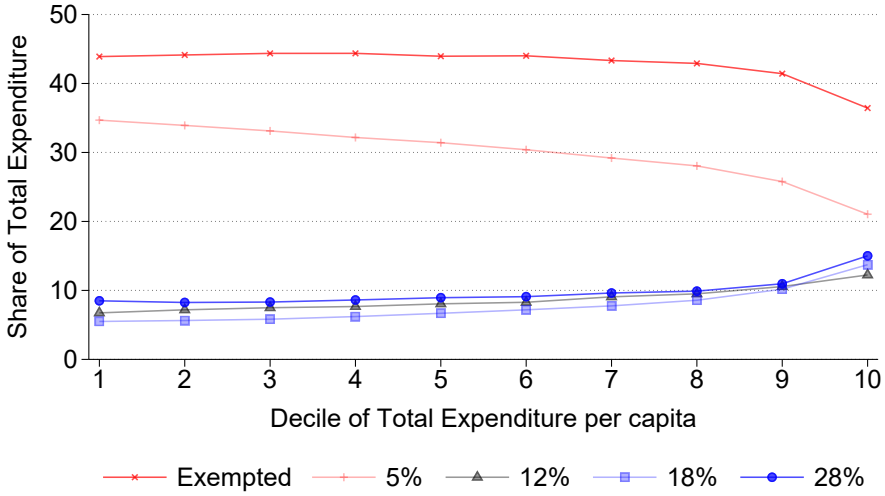
Figures

Figure 1: Distributional Consumption Shares: Necessities vs Luxuries



Note: This figure shows the share of consumption of selected products across deciles of household total expenditure per capita using the 2011-12 household expenditure survey. The luxury group (blue line) consist of the 10 products with the highest Engel curve elasticities among all survey products, and the necessity group (red line) the 10 products with the lowest. Each of these product groups account for about 1.5% of total consumption.

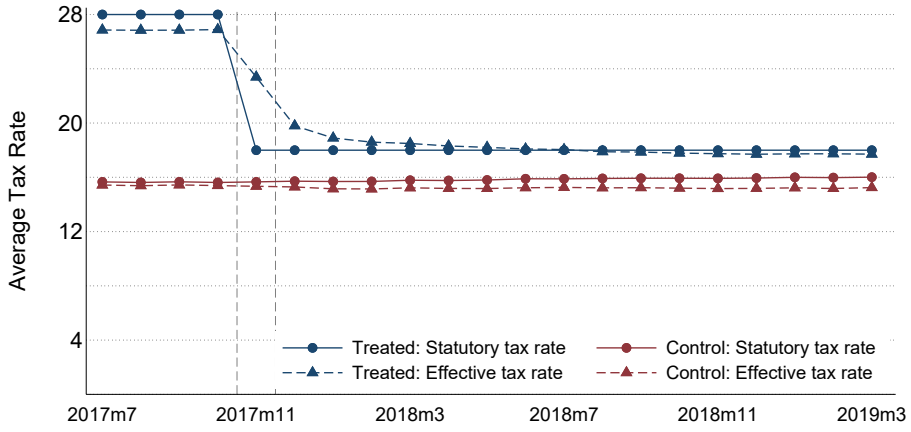
Figure 2: Distributional Expenditure Shares by Baseline GST rate



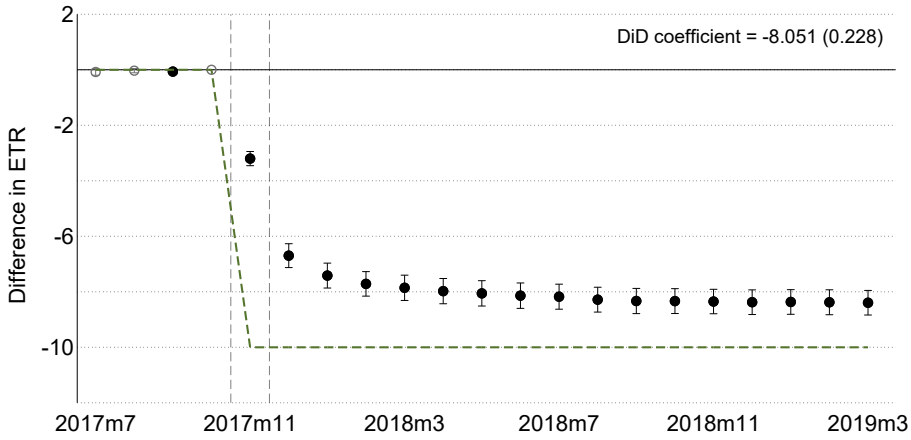
Note: This figure presents the share of total consumption by product GST rate in July 2017 across deciles of household total expenditure per capita, based on the 2011-12 household expenditure survey. Products received as gifts or social transfers and jewelry are excluded.

Figure 3: Average Statutory and Effective Tax Rates

(a) Average Level



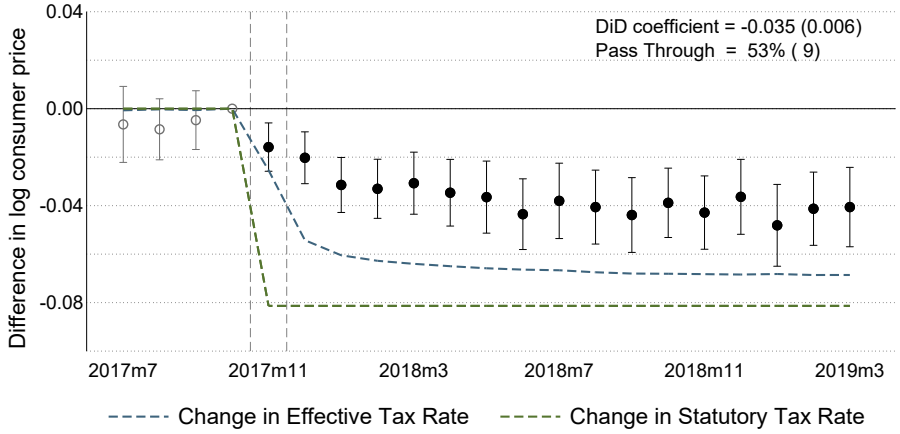
(b) Event-study



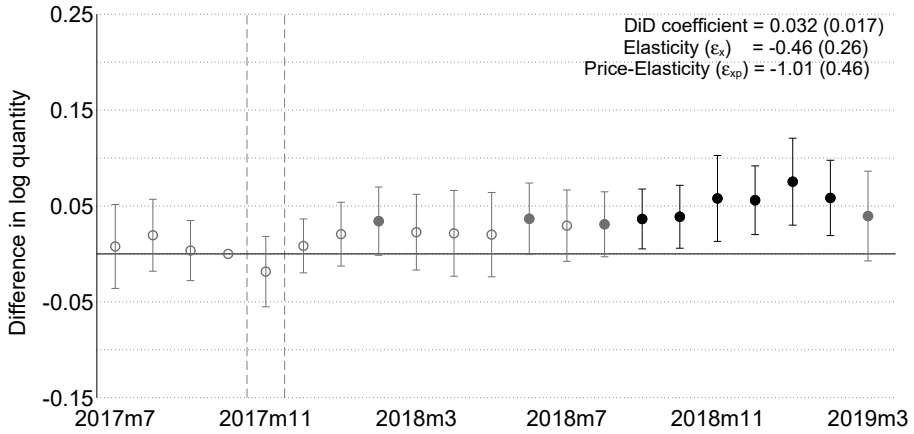
Note: Panel A presents the average Statutory Tax Rate (STR) and the average Effective Tax Rate (ETR) by treatment group. Panel B presents the event-study coefficients for changes in ETRs. The dashed green line represents the change in STR. Treated products are those that experienced a reduction in STR from 28% to 18%. Control products are those taxed at 12% or 18% that did not experience a rate change during the period. ETRs are winsorized at the 1^{st} and 99^{th} percentiles.

Figure 4: Consumer Price and Quantity

(a) Consumer Price and Tax Pass-Through



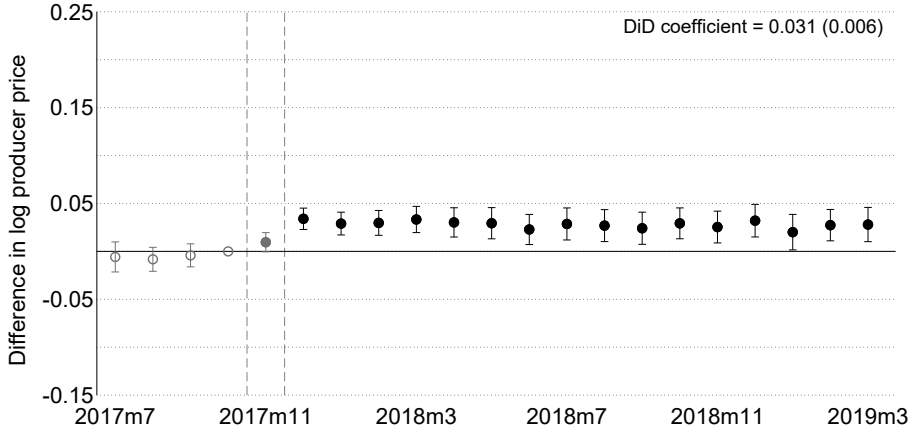
(b) Quantity



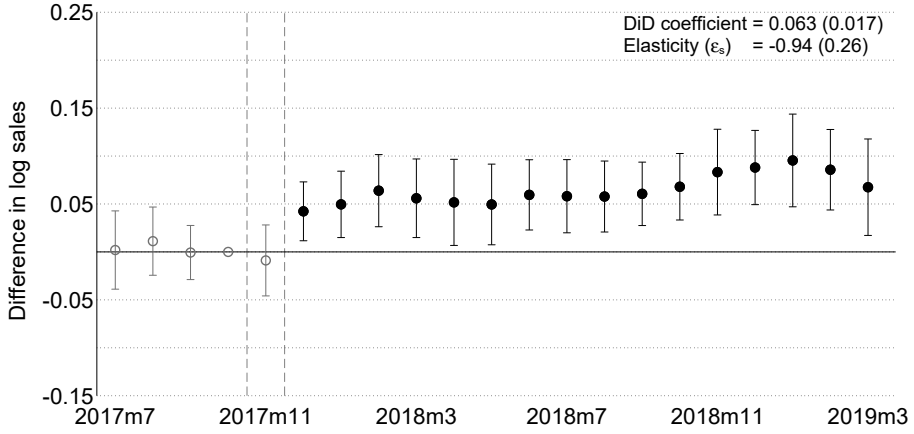
Note: Panel A presents the event study coefficients for log consumer price. The green and the blue dashed lines indicate respectively the difference in log STR and log Effective Tax Rates (ETR) between the two groups. Pass-Thought is defined as the ration between the reduction in consumer price and the reduction in ETR. This figure presents the event study coefficients for log quantity. The *Elasticity* (ϵ_x) is the elasticity of quantity to ETR, instrumented with the treatment group and post periods. The *DiD coefficient* are the difference-in-differences coefficient for the months after January 2018.

Figure 5: Pre-Tax Price and Sales

(a) Pre-Tax Price (producer price)



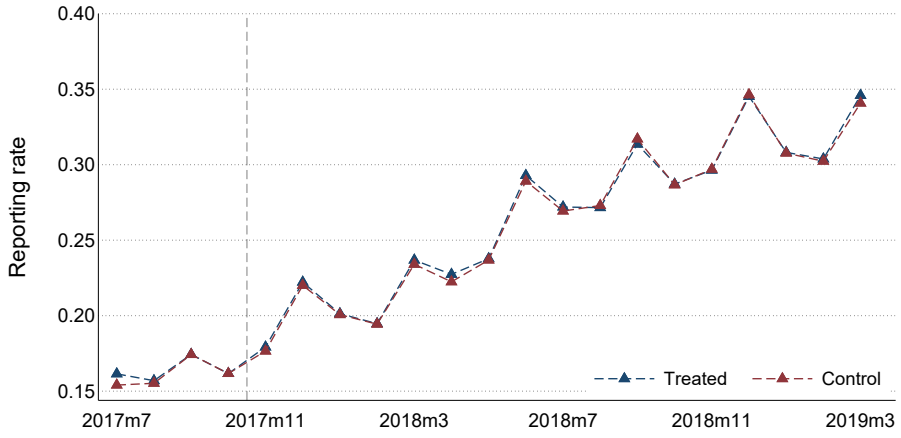
(b) Sales Value



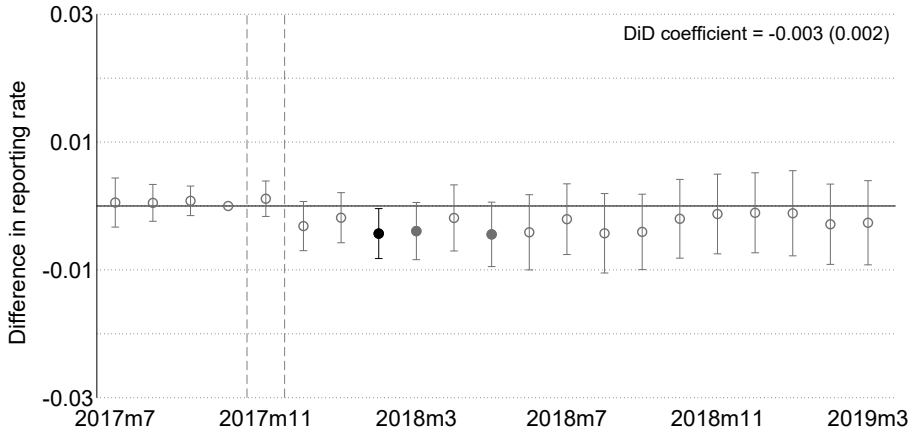
Note: This figure presents the event study coefficients for log pre-tax prices and log sales. Sales amounts are pre-tax. Treated products are those that experienced a reduction in Statutory Taxes Rates (STRs) from 28% to 18%. Control products are those taxed at 12% or 18% that did not experience a rate change during the period. The *DiD coefficient* are the difference-in-differences coefficient for the months after January 2018. The *Elasticity* (ϵ_s) is the elasticity of sales to ETR, instrumented with the treatment group and post periods.

Figure 6: Reporting Rate (Extensive Margin)

(a) Probability of reporting a product: levels

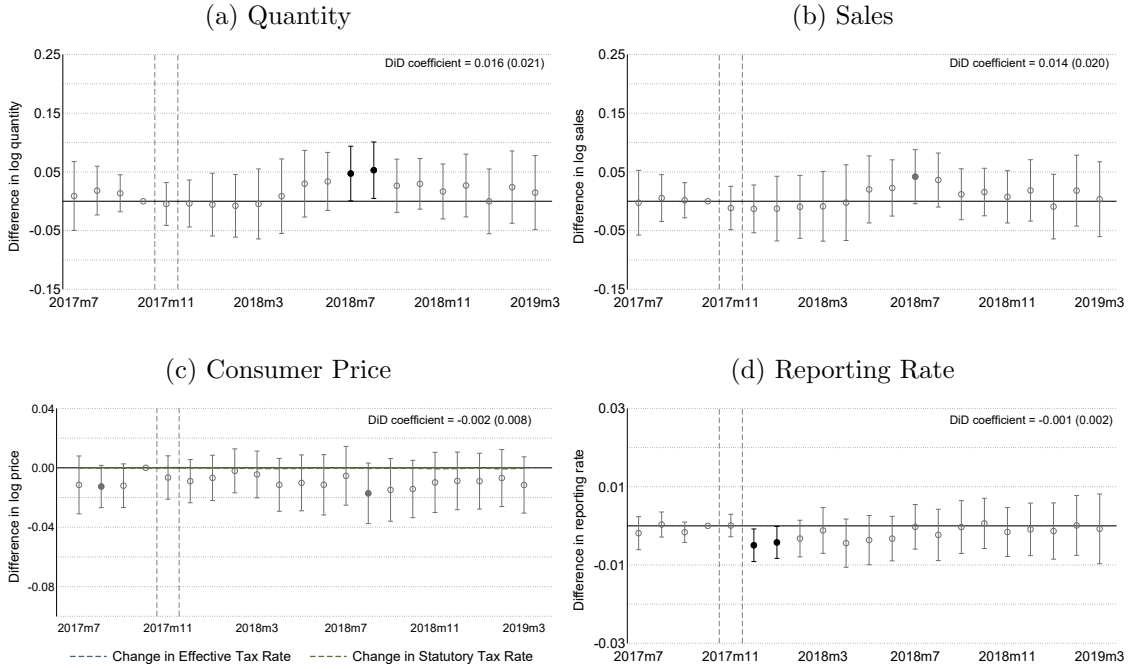


(b) Probability of reporting a product: Event-study coefficients



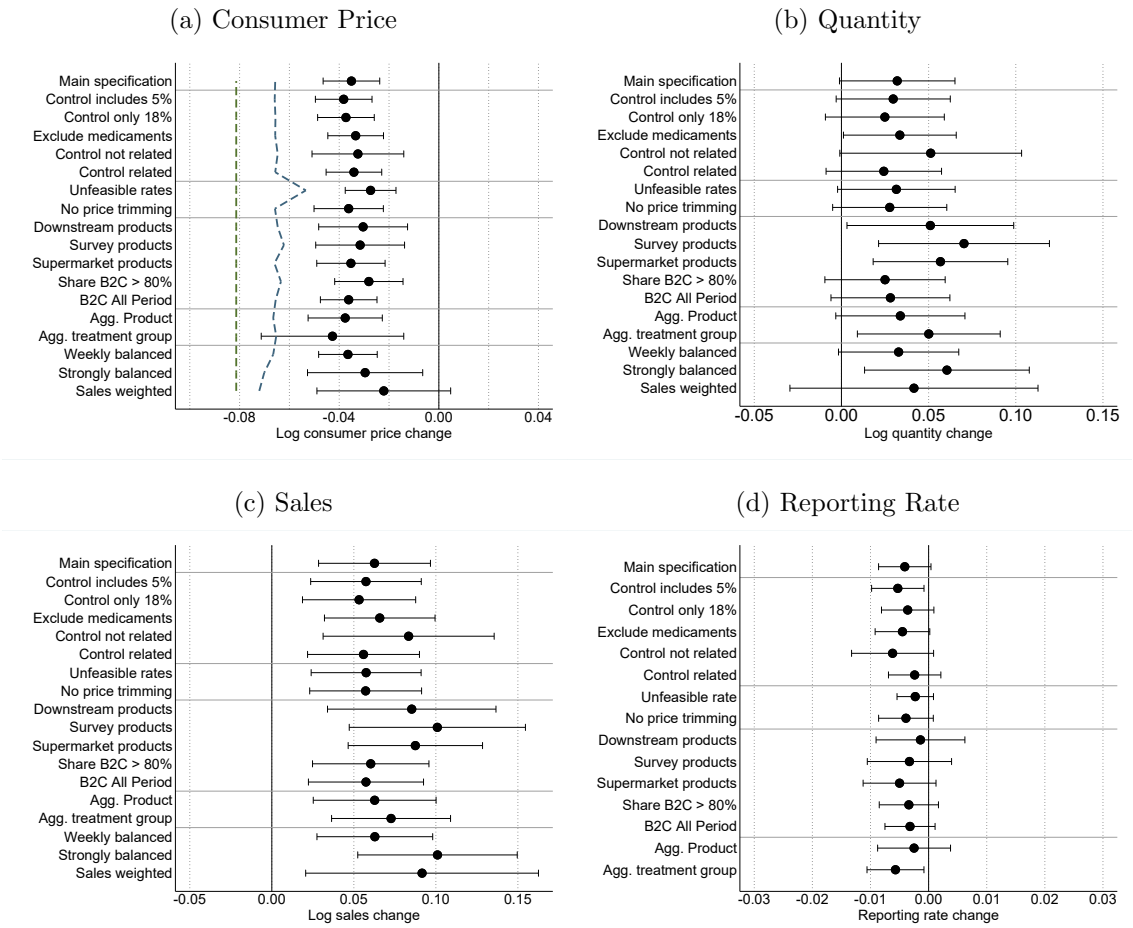
Note: This figure presents the results for the reporting rate. The reporting rate is defined as the probability of reporting selling a product in a month conditional on having sold it any month in the period analyzed. Treated products are those that experienced a reduction in Statutory Taxes Rates (STRs) from 28% to 18%. Control products are those taxed at 12% or 18% that did not experience a rate change during the period. The *DiD coefficient* are the difference-in-differences coefficient for the months after January 2018.

Figure 7: 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* refers to the difference-in-differences estimate for the months after January 2018.

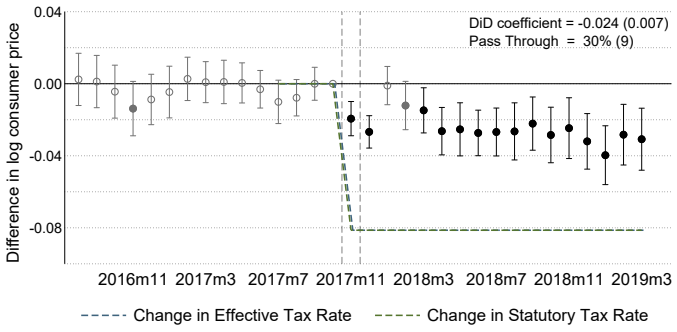
Figure 8: Robustness to Alternative specifications



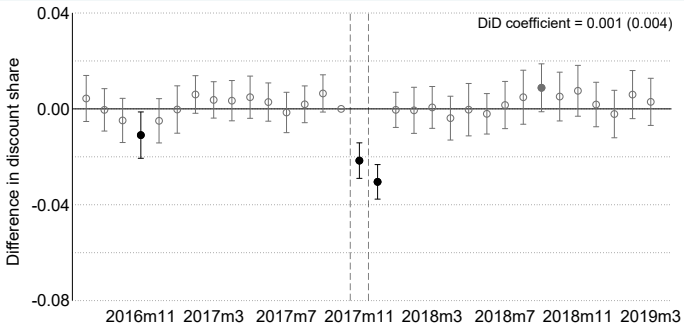
Note: The figure reports the differences-in-differences coefficients of a series of robustness tests. Each coefficient reports the point estimates from a separate robustness specification. Estimates come from difference-in-differences regressions comparing treated and control products across three periods: pre (July–October 2017, omitted), transition (Nov 2017–Jan 2018), and post (Feb 2018–Mar 2019, reported). Consumer price (Panel A), quantity (Panel B) and sales (Panel C) are are in logs. The reporting rate (Panel D) is defined as the probability of reporting selling a product in a month conditional on having sold it any month in the period analyzed.

Figure 9: Hypermarket consumer price, discount, and quantity response

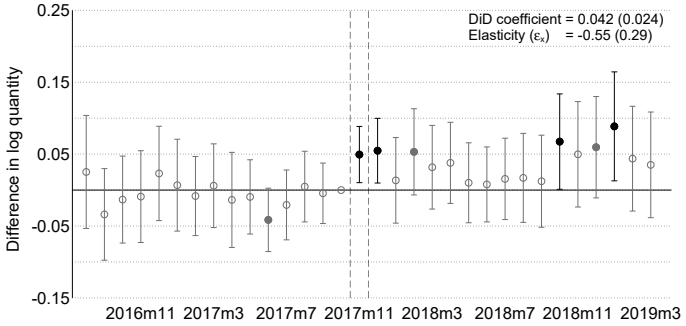
(a) Consumer Price (Inclusive of discounts)



(b) Discounts

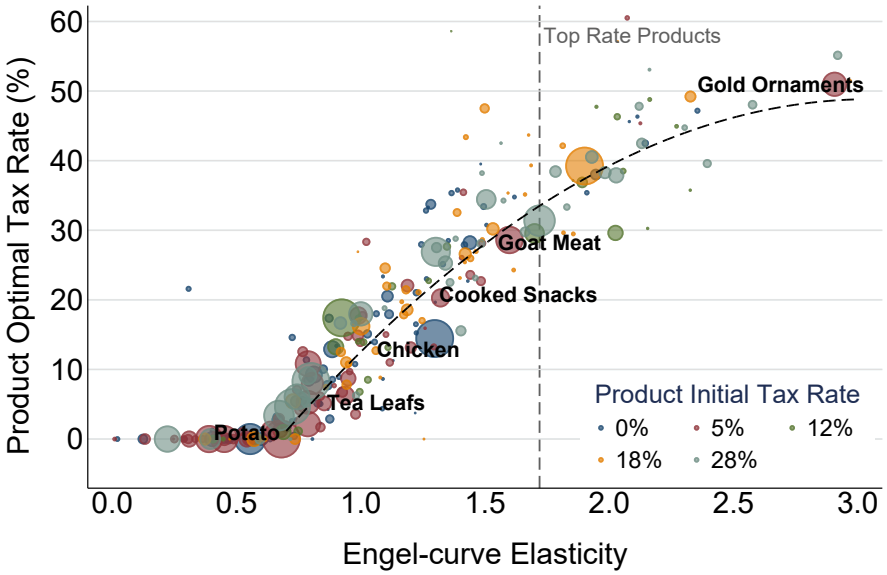


(c) Quantity



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 $= -\frac{\text{Discount}}{\text{Consumer Price+Discount}} = -\frac{\text{Discount}}{\text{Pre-discount Price}}$.

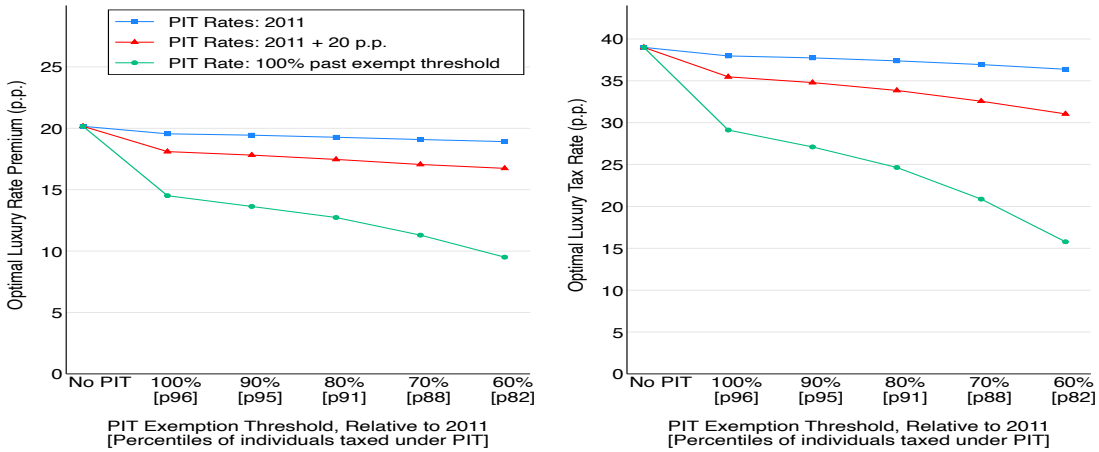
Figure 10: Optimal Tax Rate for Each Survey Product



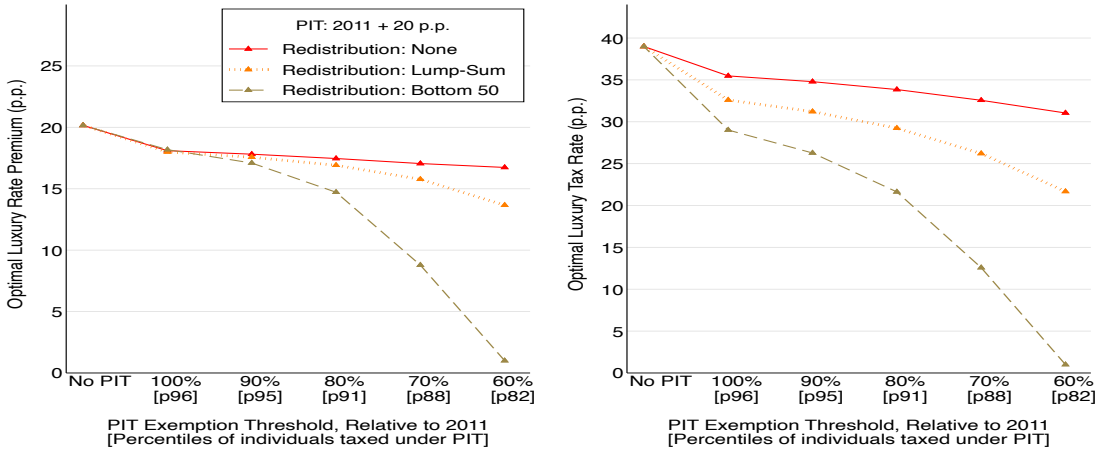
Note: This figure presents our estimates of optimal tax rates by survey product as a function of the Engel-curve elasticity parameter. Optimal tax rates below zero (i.e., subsidies) are winsorized at zero. Each circle represents one survey product, and circle sizes reflect total consumption. Product names illustrate examples of products with high total consumption at different levels of the optimal tax rate. The dashed line shows a quadratic fitted relationship.

Figure 11: Optimal Luxury Taxation with Income Taxes and Transfers

(a) Varying Income Exemption Thresholds and PIT Rates (No redistribution)



(b) Varying Income Exemption Thresholds and Redistribution of PIT Revenue



Note: This figure shows calibrated optimal rates for the basket of goods with the highest Engel-curve elasticity ("luxuries"), accounting for the same consumption share as goods initially taxed at 28% (11% of total consumption). Optimal rate differentiation is defined relative to that of a basket of standard goods. In the baseline scenarios, the standard rate is calibrated to 18% rates, but then allowed to vary. The x-axis varies the personal income tax (PIT) threshold, expressed as a share of its current value: it ranges from a scenario without PIT to increasingly lower PIT thresholds. The left side figures corresponds to the luxury rate *premium* (luxury rate minus standard rate) and the right side figures to the luxury tax *rate*. Panel (a) shows scenarios varying the tax rates of the PIT. The blue line corresponds to the PIT rates applied in India in FY 2011–2012 (top MTR 30%). The red line increases all rates by 20 percentage points (top MTR 50%). The green line assumes a 100% MTR on income above the exemption threshold (equivalent to excluding from optimal rate estimation the income subject to PIT). Panel (b) shows scenarios varying the redistribution of PIT revenues, under a rate schedule of with a top MTR 50% (same as red line in panel (a))). The orange line redistributes the PIT revenue in a lump-sum manner. The brown line redistributes the PIT revenue equally to the bottom 50% of households.

Tables

Table 1: Description of treatment and control group

	Treated	Control	All
Products	164	680	844
Firm-HSNs	133,255	316,113	449,368
Of which before Nov 17	45,831	107,886	153,717
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	6.29 (1.97) [8.00]	6.28 (1.97) [8.00]	6.29 (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)	204 (2115) [7.80]	429 (15553) [7.17]	362 (13071) [7.36]
Sales (Log value)	9.05 (2.42) [8.96]	9.00 (2.56) [8.88]	9.02 (2.52) [8.90]
Quantity (Log units)	3.75 (2.61) [3.50]	4.16 (2.74) [3.95]	4.04 (2.71) [3.83]
Price pre-tax (Log value per unit)	5.30 (2.01) [5.07]	4.84 (2.12) [4.57]	4.98 (2.10) [4.74]
Observations	706,104	1,672,913	2,379,017
Of which before Nov 17	93,970	220,107	314,077

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 in three groups of rows: Products, Firm-HSN, and Observations. Standard deviations are in parenthesis and the median in brackets. para Statistics are measured before November 2017.

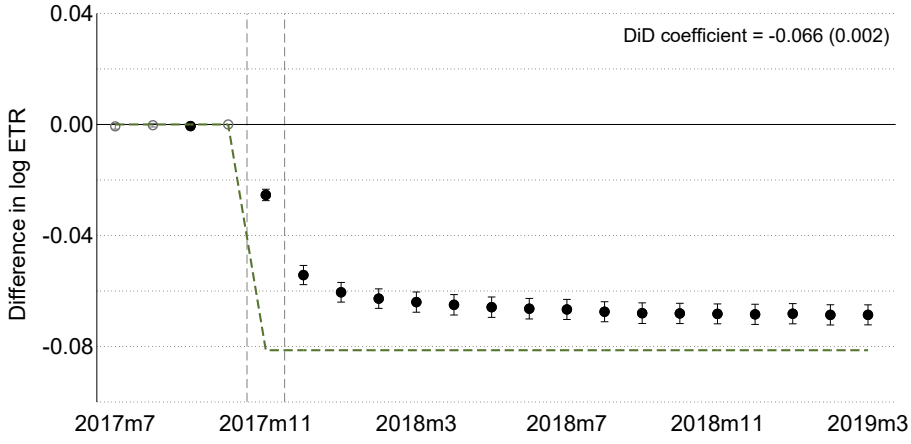
Table 2: Description of firms in the tax data

Firms	28,370	
Firms before Nov 17	19,793	
Tax returns	260,563	
Tax returns before Nov 17	47,641	
For tax returns before Nov 17:	Mean	Median
Sales ('000 rupees)	8,988	1,045
Sales treated	538	0
Sales control	2,891	56
Sales excluded	5,558	205
HSNs	13.30	3
HSNs treated	2.47	0
HSNs control	5.96	1
HSNs excluded	4.82	1
Has treated product	0.42	0
Has treated product (Sales)	0.65	1
Has >1 product	0.66	1
Has >1 product (Sales)	0.88	1
Conditional on selling treated product:		
Has >1 product	0.91	1
Has >1 product (Sales)	0.99	1
Has control product	0.83	1
Has control product (Sales)	0.93	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. "HSNs" refer to the number of HSNs reported in a firm-product tax return. Rows ended with (Sales) are weighted by firm sales.

Appendix A: Additional Figure and Tables

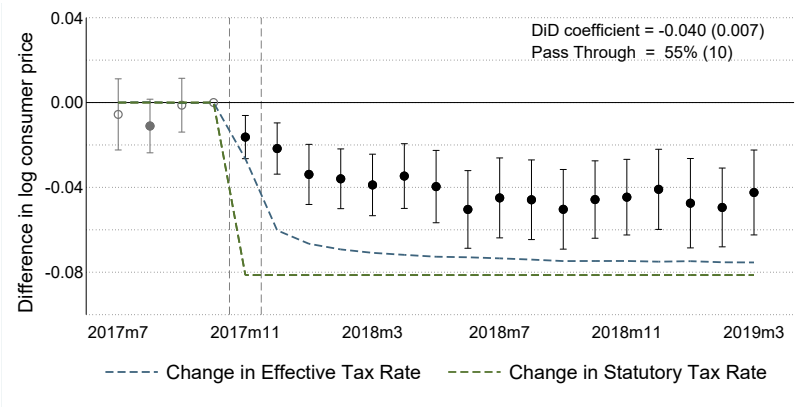
Figure A.1: Average Statutory and Effective Tax Rates in logs



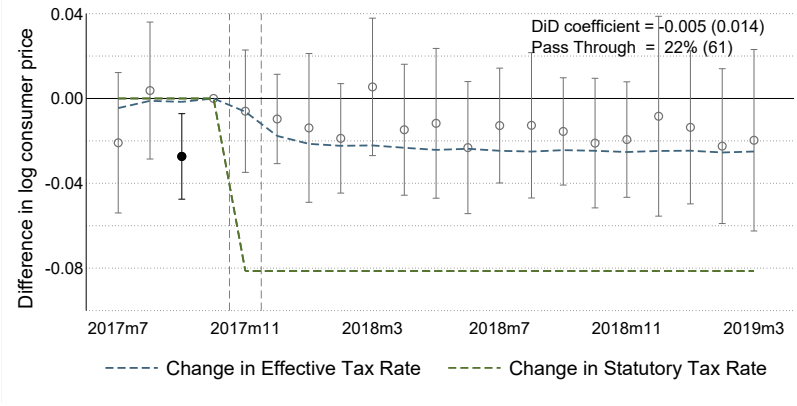
Note: The figure presents the event-study coefficients for changes in log ETRs. The dashed green line represents the change in log STR. Treated products are those that experienced a reduction in STR from 28% to 18%. Control products are those taxed at 12% or 18% that did not experience a rate change during the period.

Figure A.2: Consumer Price and Pass-Thought by STR-ETR match status

(a) ETR matches STR



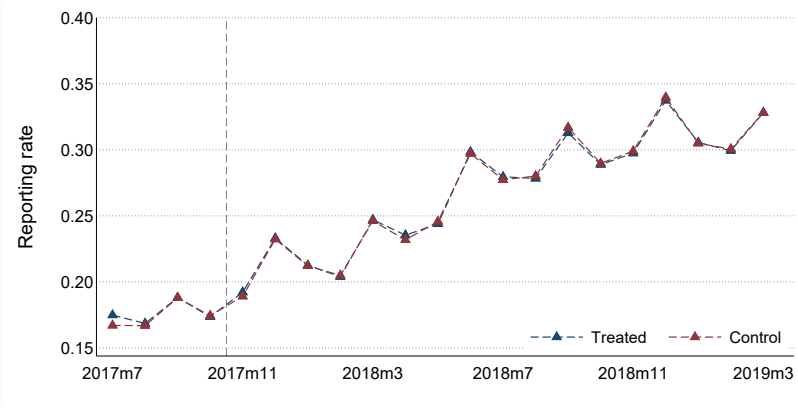
(b) ETR does not match STR



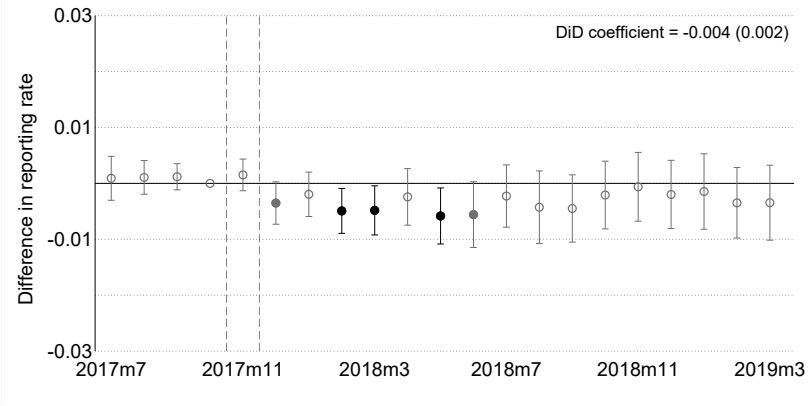
Note: This figure presents the event study coefficients for the difference in log consumer price. Treated products are those that experienced a reduction in Statutory Taxes Rates (STRs) from 28% to 18%. Control products are those taxed at 12% or 18% that did not experience a rate change during the period. 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-Thought is defined as the ration between the reduction in consumer price and the reduction in ETR. The pre-period for the difference-in-differences coefficients are the months before November 2017 months and the post period are the months after January 2018. The group *ETR matches STR* includes treated firm-HSNs with a pre-November 2017 ETR no lower than 27% and no higher 29%, that is, those with an ETR error lower than 1%. The group *ETR does not match STR* includes treated firm-HSN with a pre-November 2017 ETR lower than 27% or higher than 29%, that is, those with an ETR error greater than 1%.

Figure A.3: Reporting Rate

(a) Level



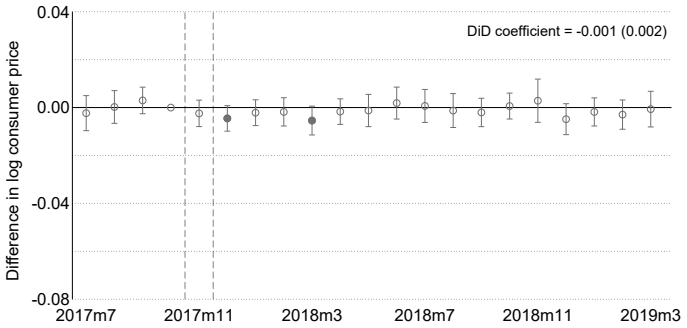
(b) Event-study



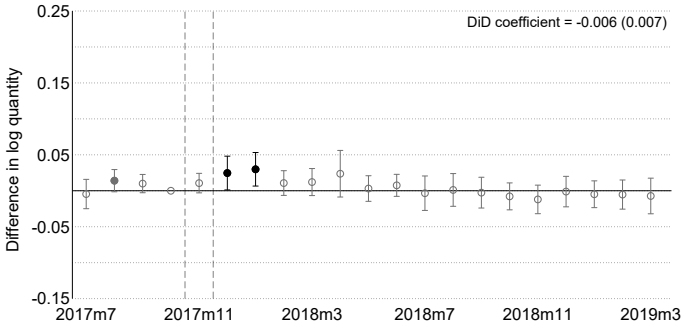
Note: This figure presents the results for the reporting rate. The reporting rate is defined as the probability of reporting selling a product in a month conditional on having sold it any month in the period analyzed. The sample definition uses the average share of final consumer sales before November 2017 to restrict to sales to final consumers. Treated products are those that experienced a reduction in Statutory Taxes Rates (STRs) from 28% to 18%. Control products are those taxed at 12% or 18% that did not experience a rate change during the period. The *DiD coefficient* are the difference-in-differences coefficient for the months after January 2018.

Figure A.4: Within-Firm Spillover

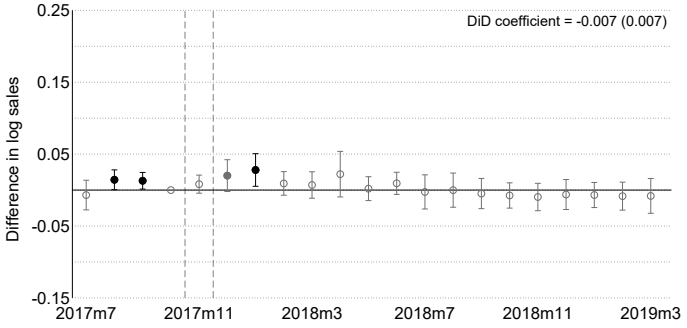
(a) Consumer Price



(b) Quantity

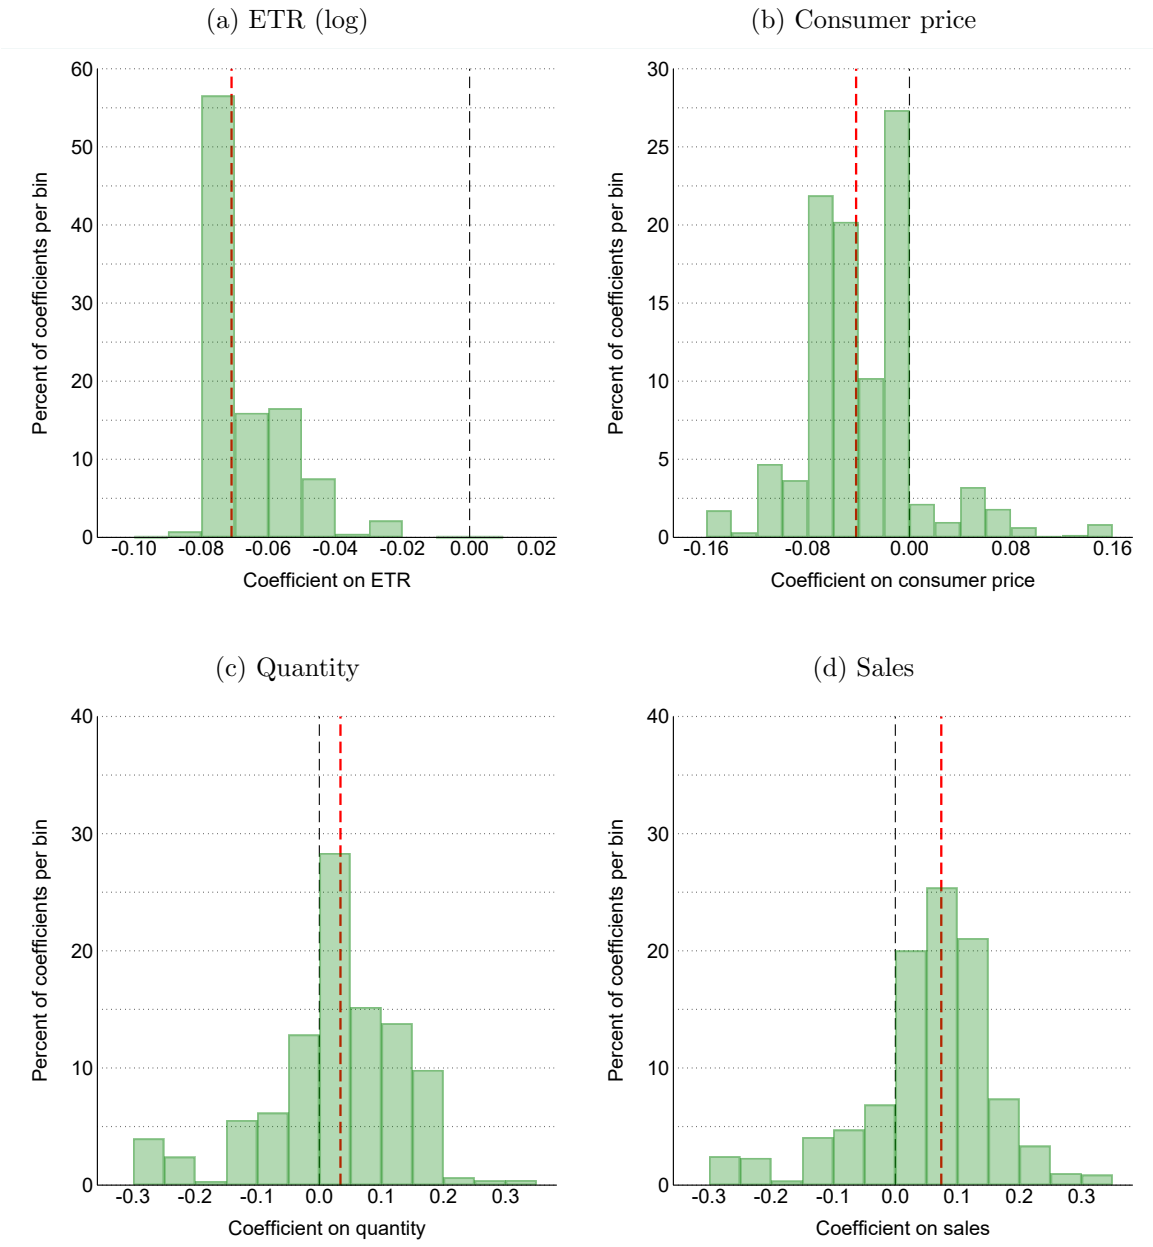


(c) Sales



Note: plots the coefficients γ_t from equation (2) for consumer prices, quantities, and sales, estimated on control products only, where the baseline firm-level treated share is interacted with monthly indicators; all regressions include firm–HSN fixed effects (λ_{ip}) and product–time fixed effects (δ_{pt}), with standard errors clustered at the firm and 4-digit product levels. The mean baseline firm-level treated share is 13.1% among firm-HSNs with a positive exposition.

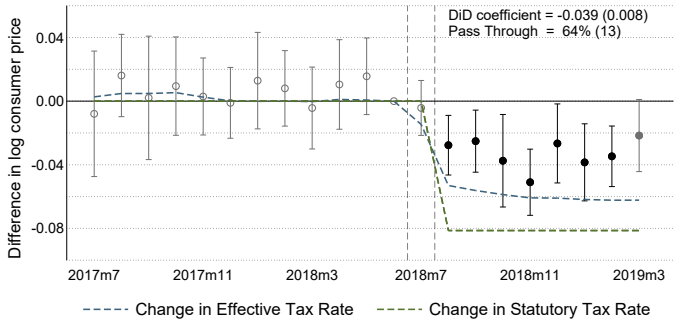
Figure A.5: Distribution of product coefficients



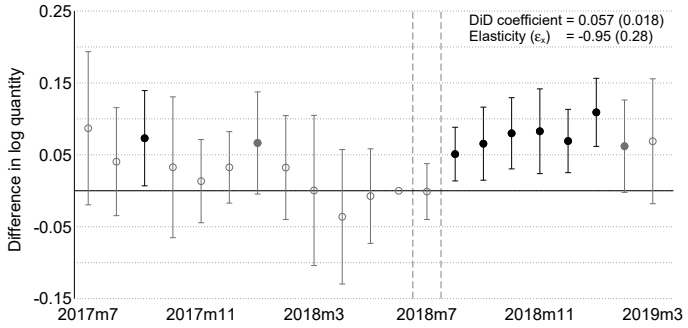
Note: This figure displays coefficients of the difference-in-differences for each treated product for log sales, quantity, effective tax rates and consumer prices. Coefficients are weighted by number of observations before November 2017 and winterized on each histogram maximum and minimum values. The red line indicates the aggregate coefficient.

Figure A.6: Treated in July 2018

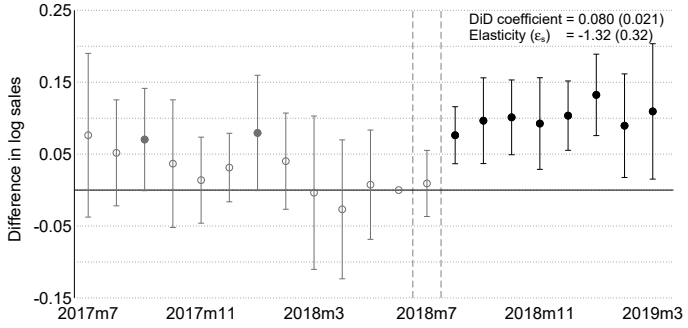
(a) Consumer Price



(b) Quantity

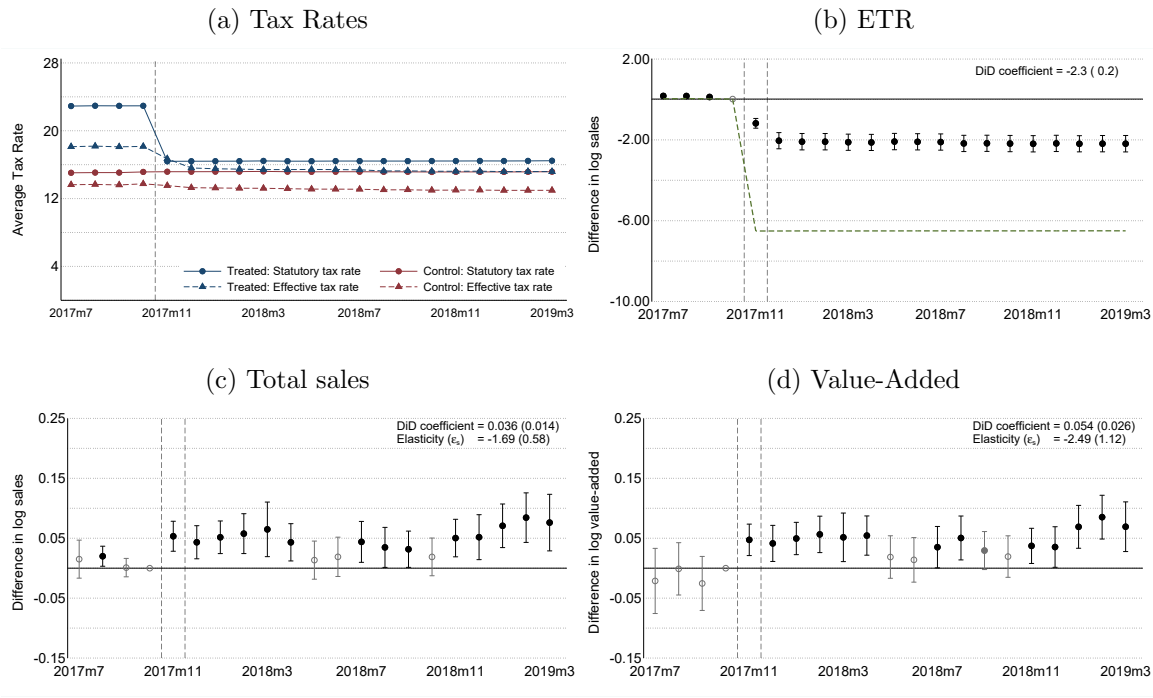


(c) Sales



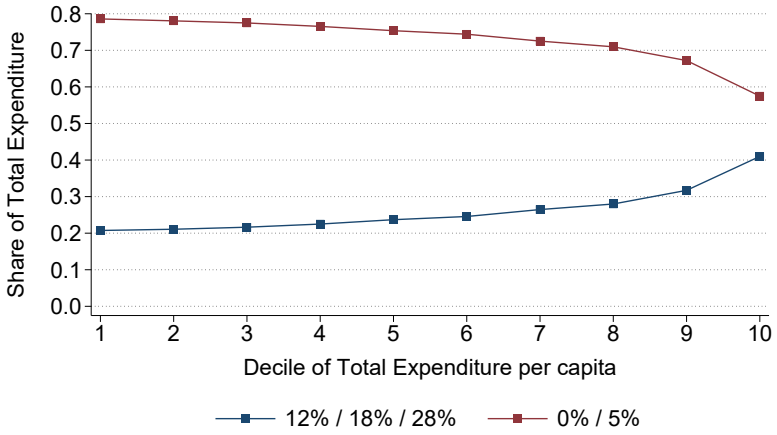
Note: This figure presents the results for log quantity, log sales and log price. reated products are those that experienced a reduction in Statutory Taxes Rates (STRs) from 28% to 18% in July 2018. Control products are those taxed at 12% or 18% that did not experience a rate change between July 2017 and March 2019. The *DiD coefficient* refers to the difference-in-differences estimate for the months after August 2018 compared to months before July 2017.

Figure A.7: Firm tax returns responses



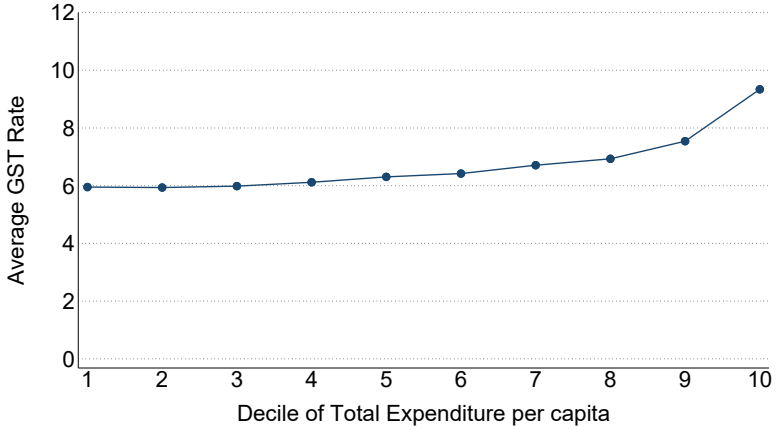
Note: Panel A presents the average statutory tax rate and the average effective tax rates by treatment status using firm-level tax returns (GSTR3). Panel B present the event-study coefficients comparing treated and control firms for the change in ETR and STR, Panel C for the log total sales, and Panel 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, inferred from firm-to-firm transactions. We do not observe inputs bought through inter-state transactions. The analysis is restricted to firm-months with positive value added.

Figure A.8: Expenditure share by total consumption decile and groups of baseline tax rate



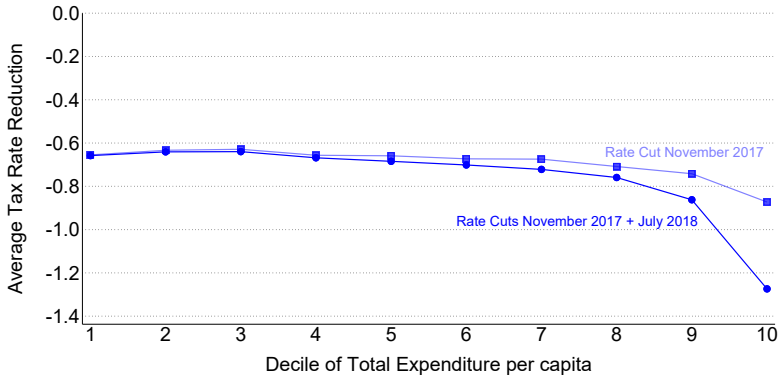
Note: This figure presents the share of total consumption by groups of GST rate in July 2017 across deciles of household total expenditure per capita. Services, goods received as gifts or social transfers and jewellery are excluded from the share of total expenditure.

Figure A.9: Average baseline GST tax rate by total consumption decile



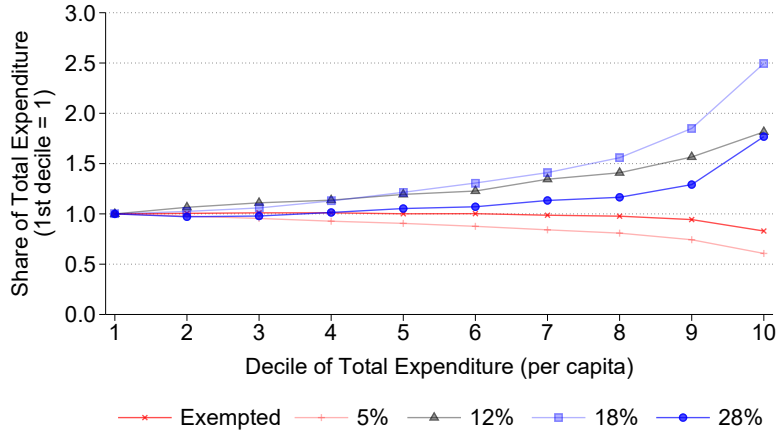
Note: This figure presents the average GST rate in July 2017 across deciles of household total expenditure per capita. The average rate is defined as total taxes divided by total consumption in each decile. Services, goods received as gifts or social transfers and jewellery are excluded from the share of total expenditure.

Figure A.10: Reduction in average tax rate after rate reductions



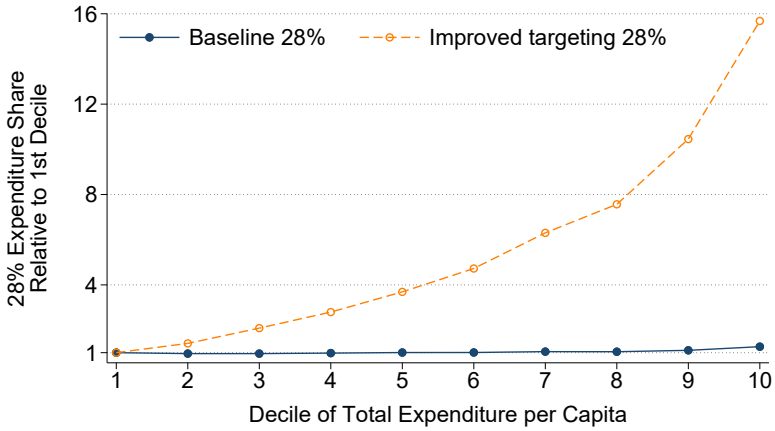
Note: This figure presents the reduction GST rate across deciles of household total expenditure per capita following the November 2017 and July 2018 rate reductions. The average rate reduction is defined as total taxes before minus total taxes after the reduction divided by total consumption in each decile. Services, goods received as gifts or social transfers and jewellery are excluded from the share of total expenditure.

Figure A.11: Normalized expenditure share by total consumption decile and baseline tax rate



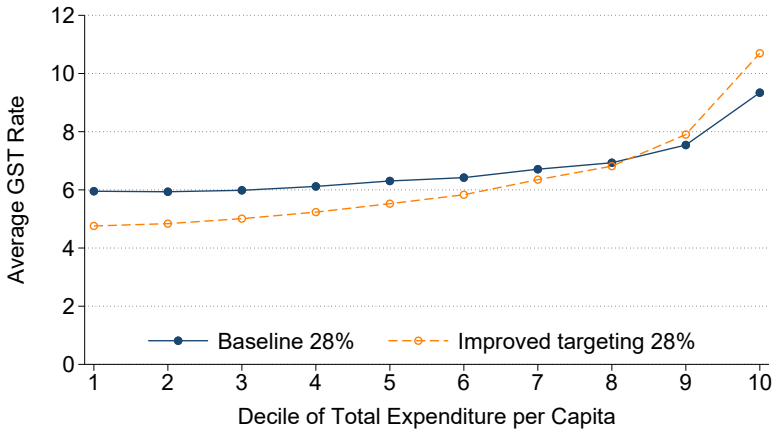
Note: This figure presents the share of total consumption by GST rate in July 2017 across deciles of household total expenditure per capita normalized by the consumption share of the first decile. Services, goods received as gifts or social transfers and jewellery are excluded from the share of total expenditure.

Figure A.12: Normalized 28% rate expenditure share by total consumption decile and target



Note: This figure presents the consumption share of goods taxed at 28% across deciles of household total expenditure per capita normalized by the consumption share of the first decile. Optimal targeting assign goods with the higher ratio between 5th and 10th percentile to the 28% keeping the total amount taxed at 28% and the average rate constant. Services, goods received as gifts or social transfers and jewellery are excluded from the share of total expenditure.

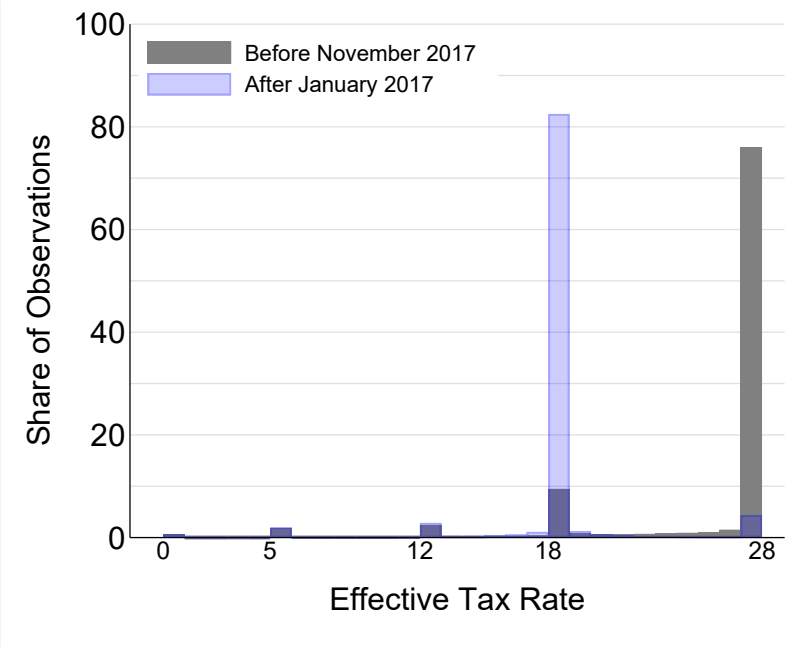
Figure A.13: Average rate with improved targeting by total consumption decile



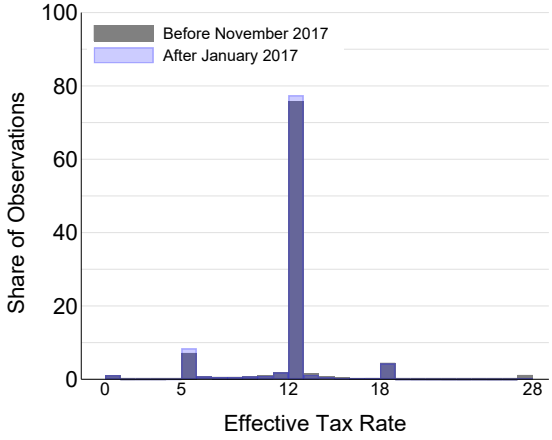
Note: This figure presents average GST tax rate across deciles of household total expenditure per capita normalized by the consumption share of the first decile. Optimal targeting assign goods with the higher ratio between 5th and 10th percentile to the 28% keeping the total amount taxed at 28% and the average rate constant. Services, goods received as gifts or social transfers and jewellery are excluded from the share of total expenditure.

Figure A.14: Histogram of products Effective Tax Rates before and after rate reduction

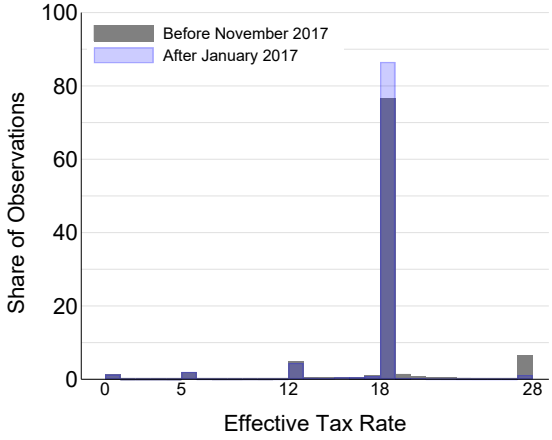
(a) Treated



(b) Control 12%



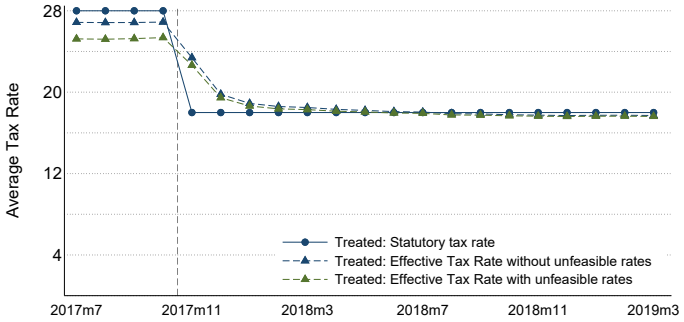
(c) Control 18%



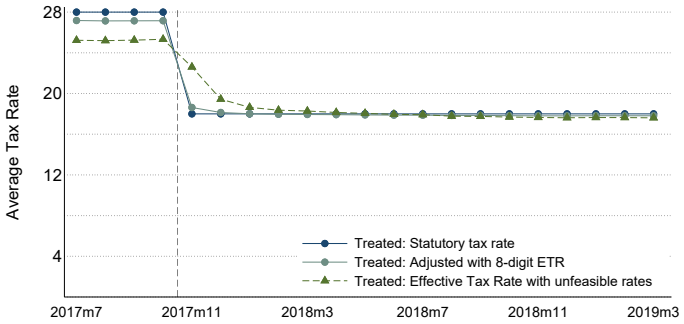
Note: This figure presents the distribution of Effective Tax Rates (ETR) before and after November 2017 rate reduction. Treated products are those that experienced a reduction in STR from 28% to 18% in November 2017. ETRs are winsorized at the 1st and 99th percentiles and rounded, such that the 18% bin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

Figure A.15: Average corrected Statutory Tax Rates and Effective Tax Rates

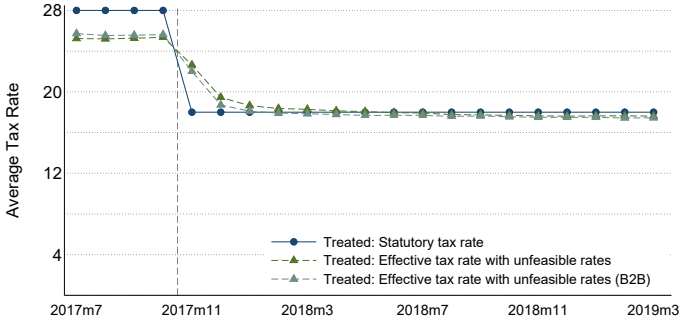
(a) Unfeasible effective tax rates



(b) Corrected Statutory Tax Rates

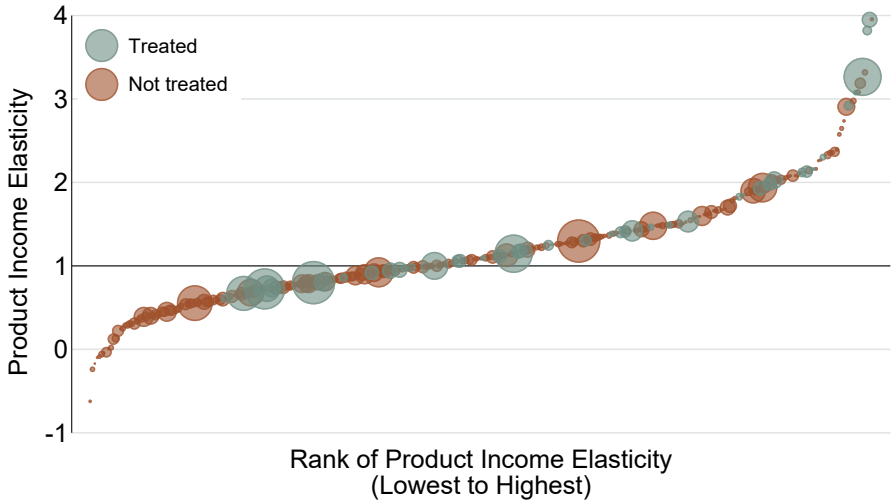


(c) B2B firms



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

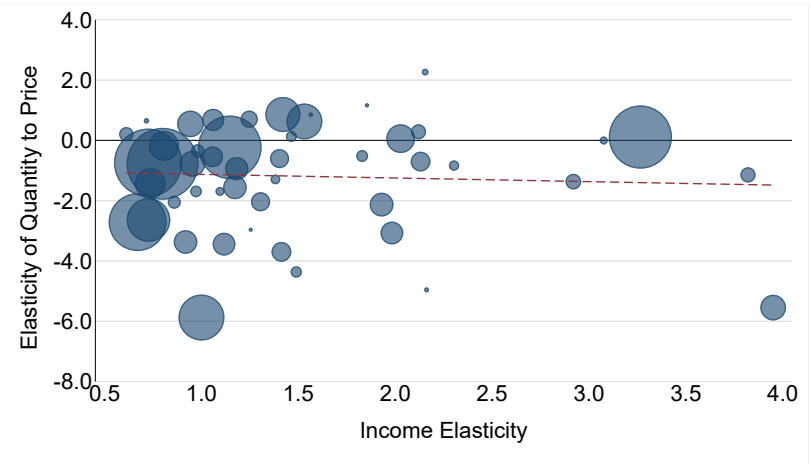
Figure A.16: Product Income Elasticity by Treatment Status



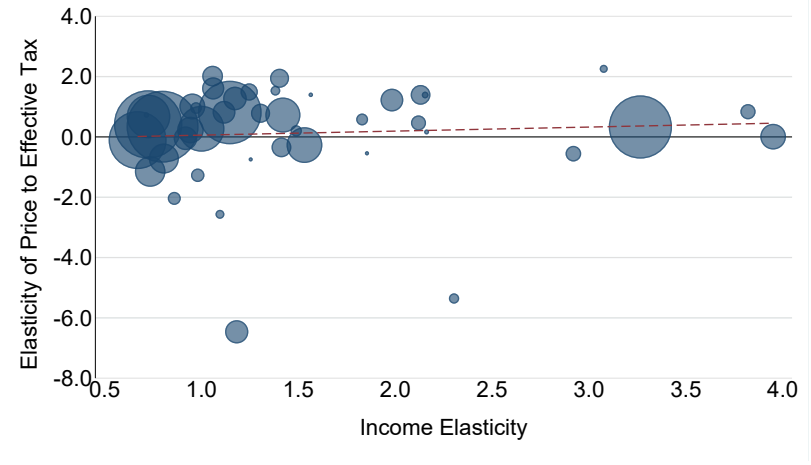
Note: This figure presents Engel beta coefficient of survey products. Treated products are those with a rate reduction from 28% to 18% in November 2017. The circle size increases with product consumption in the survey.

Figure A.17: 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 shows the linear fit from an unweighted regression. Panel B excludes the two products with the lowest and the two with the highest consumer price elasticity values for better visualization.

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 a 8-digit HSN code. "Powder face" is an 8-digit code, included in the 6-digit code for "make-up or skin care powders". This 6-digit code is part of the 4-digit code for "beauty, make-up or skin care preparations". The corresponding 2-digit code is "essential oils and residuals, 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.

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

Products Per HSN-4	Number of HSN-4s	Percent of HSN-4s
1	1,136	91.54
2	90	7.25
3	13	1.05
4	2	0.16
Total	1,241	100.00

Notes: Products are defined as unique combinations of an HSN-4 code and rate path. 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 91.4% of HSN-4s are a product and 8.6% of HSN-4s have more than one product.

Table A.3: Description of observation-level sample cleaning and restriction steps

Cleaning Step	Number of observations	Share (%)
Initial sample	12,815,721	100.0
Services	832,835	6.5
Incomplete Tax Returns		
Missing product code	361,296	2.8
Non-existant code	384,889	3.0
HSN 2-digit 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		
Unfeasible 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.

Table A.4: Description of product-level sample restrictions

Treatment Group	Products	Observations (N)	Observations share	Sales share
Totals before November 2017	1363	469,671	100	100
Treated products	164	93,970	20.0	6.7
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	104	35,987	7.7	5.7
Aug 2018	32	16,797	3.6	2.6
Other periods	18	8,582	1.8	2.2
Other rates	54	10,608	2.3	0.9
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 1,268 HSN-4 level codes and 1,423 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 2-digit HSN code as a treated product, while the subgroup *Related* includes control products that are in the same 2-digit HSN 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 *August 2018* refers to products that had a rate reduction in that month. Subgroup *Other periods* include products with rate reductions in months other than November 2017 or August 2018. Subgroup *Other rates* include products that had a rate reduction in November 2017 but not from 28% to 18%. The *Excluded 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%.

Table A.5: Tax Rates in July 2017 and March 2020

Before \ After	0%	5%	12%	18%	28%	Total before (%)	Treated before (%)
0%	122 (10)	0 (0)	0 (0)	0 (0)	0 (0)	9 (10)	0 (0)
5%	11 (0)	267 (29)	0 (0)	0 (0)	0 (0)	20 (30)	4 (1)
12%	4 (0)	16 (0)	167 (15)	0 (0)	0 (0)	14 (16)	11 (2)
18%	1 (0)	4 (0)	33 (1)	513 (17)	0 (0)	40 (19)	7 (8)
28%	0 (0)	0 (0)	6 (0)	194 (15)	29 (11)	17 (26)	87 (57)
Total after (%)	10 (10)	21 (30)	15 (17)	52 (32)	2 (11)		
Treated after (%)	12 (6)	7 (0)	19 (10)	27 (46)	0 (0)		

Note: This table presents the number of products and sales-weighted shares in parentheses in each combination of initial (July 2017) and final (March 2020) GST rates. Products are defined by a 4-digit HSN code and their combination of initial and final tax rates. Sales-weighted shares use pre-November 2017 sales.

Table A.6: Top products with a rate reduction in November 2017

Product name	Product code (4-digit HSN)	Treated products (%)	Cumulative share (%)
Make-Up	3304	7.4	7.4
Shaving Preparations	3307	6.9	14.3
Washing, Cleaning Or Degreasing Preparations	3402	5.7	20.0
Hair Preparations	3305	4.7	24.7
Rubber Articles (Gaskets, Seals, Bands)	4016	3.8	28.5
Chocolate	1806	3.8	32.3
Electrical Circuits (Switchers, Lamp-Holders, Fuses)	8536	2.9	35.1
Malt Extract	1901	2.9	38.0
Polishes And Creams	3405	2.8	40.8
Eletrical Transformers	8504	2.8	43.6
Insulated Electric Conductors (Wires)	8544	2.8	46.4
Air Vacuum Pumps	8414	2.4	48.9
Liquid Or Cream Soap	3401	2.3	51.2
Plastic Articles (Office And School Supplies)	3926	2.2	53.4
Mountings	8302	2.2	55.5
Lamps	8539	2.2	57.7
Ceramic Flags And Paving	6907	1.9	59.6
Lighting Or Signalling Equipment	8512	1.9	61.5
Perfumes And Toilet Waters	3303	1.9	63.3
Furniture	9403	1.8	65.2
Extraxts Of Coffee	2101	1.3	66.5
Sanitary Fixtures	6910	1.3	67.7
Inner Tubes	4013	1.2	69.0
Bread, Pastry, Cakes, Biscuits	1905	1.2	70.1
Suitcases, Vanity Cases, Briefcases	4202	1.1	71.3
Baths And Sinks	3922	1.1	72.4
Batteries	8506	1.1	73.5
Mirrors	7009	1.0	74.5
Painters' Colours	3213	1.0	75.5
Razors	8212	0.9	76.4
Other Products	.	23.6	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, Hair Preparations is the second most important category, with its 4-digit HSN product code being 3305. It accounted for 6.8% of the firm-products with a rate change. Hair Preparations and Make-Up accounted together for 15.3% of firm-products with a rate change.

Table A.7: Top control products

Product name	Product code (4-digit HSN)	Treated 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.

Table A.8: Description of Related and Not Related 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 rupees)	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 2-digit HSN codeas a treated product. The *Not Related* column control products that are not in the same 2-digit HSN code as a treated product. The table is divided in three groups of rows: Products, Firm-HSN, and Observations. Statistics are measured for months before November 2017.

Table A.9: Main Treated and Related Product per 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	24.5	3.0
96	Lighters	9613	Brooms And Brushes	9603	1.8	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	13.3	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	Extraxts Of Coffee	2101	Sauces And Mixed Condiments	2103	5.8	0.8
76	Roofs, Dors 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	79.1	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 Acessories 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 most relevant treated and control product in each HSN-2 with a treated and control products. Product relevance is based on the number of observations (firm-HSN-months) before November 2017. The names shown are simplified versions of the HSN names. Columns 1 lists the HSN-2 code. Column 2 displays the name of the treated product, and column 3 its corresponding HSN code. Columns 4 and 5 present the name and code, respectively, 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 indicates the percentage of related group that the related product represents. For example, the main product with a rate reduction in HSN-2 34 was and the main control product in HSN-2 34 was . Column 8 indicates products with a rate reduction accounted for 40% of pre-November 2017 sales in HSN-2 34, while column 9 indicates that accounted for 2.3% of observations in the related group before November 2017.

Table A.10: Indirect Effect 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 indirectly treated and control products. Estimates compare the pre period (July–October 2017) to the post period (February 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 the broader sample that includes 5% rate products. Column (4) restricts the treatment group to products sold by firms that also sell directly treated products (within-firm variation). Standard errors, clustered at the firm and 4-digit product level, are reported in parentheses.

Table A.11: Robustness Checks

		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)	(19)
	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	Weekly balanced	Strongly balanced	Sales weighted	Firm-Product	Firm-Treatment	Product
ETR	-0.066	-0.066	-0.066	-0.066	-0.065	-0.066	-0.066	-0.066	-0.065	-0.062	-0.066	-0.063	-0.065	-0.067	-0.070	-0.072	-0.067	-0.065	-0.073
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.003)
Consumer Price	-0.035	-0.038	-0.037	-0.033	-0.032	-0.034	-0.036	-0.035	-0.030	-0.032	-0.035	-0.028	-0.036	-0.036	-0.030	-0.022	-0.038	-0.043	-0.148
	(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)	(0.008)	(0.015)	(0.065)
Quantity	0.032	0.030	0.025	0.034	0.051	0.024	0.028	0.032	0.051	0.070	0.057	0.025	0.028	0.033	0.060	0.042	0.034	0.050	0.019
	(0.017)	(0.017)	(0.017)	(0.016)	(0.027)	(0.017)	(0.017)	(0.017)	(0.024)	(0.025)	(0.020)	(0.018)	(0.017)	(0.018)	(0.024)	(0.036)	(0.019)	(0.021)	(0.085)
Sales	0.063	0.057	0.053	0.066	0.083	0.056	0.057	0.063	0.085	0.101	0.087	0.060	0.057	0.063	0.101	0.092	0.063	0.073	-0.058
	(0.017)	(0.017)	(0.018)	(0.017)	(0.027)	(0.017)	(0.017)	(0.017)	(0.026)	(0.027)	(0.021)	(0.018)	(0.018)	(0.018)	(0.025)	(0.036)	(0.019)	(0.019)	(0.060)
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	-0.014	_	-0.010	-0.002	-0.006	_
	(-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)	(-0.014)		(-0.010)	(-0.002)	(-0.006)	
Consumer Price Elasticity (ϵ_p)	0.526	0.567	0.563	0.503	0.492	0.515	0.541	0.541	0.454	0.492	0.524	0.442	0.542	0.530	0.417	0.291	0.490	0.439	1.993
	(0.090)	(0.092)	(0.091)	(0.089)	(0.149)	(0.088)	(0.108)	(0.108)	(0.144)	(0.153)	(0.109)	(0.113)	(0.091)	(0.090)	(0.165)	(0.182)	(0.112)	(0.185)	(0.856)
Quantity Elasticity (ϵ_q)	-0.464	-0.426	-0.366	-0.487	-0.759	-0.359	-0.396	-0.396	-0.757	-1.093	-0.829	-0.387	-0.407	-0.447	-0.853	-0.524	-0.210	-0.394	-0.331
	(0.259)	(0.256)	(0.266)	(0.253)	(0.409)	(0.260)	(0.255)	(0.255)	(0.376)	(0.405)	(0.301)	(0.278)	(0.267)	(0.266)	(0.342)	(0.467)	(0.243)	(0.292)	(1.101)
Sales Elasticity (ϵ_s)	-0.938	-0.859	-0.803	-0.984	-1.266	-0.844	-0.855	-0.855	-1.302	-1.601	-1.304	-0.945	-0.865	-0.918	-1.436	-1.232	-0.720	-0.955	0.669
	(0.262)	(0.259)	(0.269)	(0.259)	(0.400)	(0.265)	(0.263)	(0.263)	(0.391)	(0.434)	(0.310)	(0.284)	(0.271)	(0.269)	(0.350)	(0.466)	(0.245)	(0.268)	(0.764)
Quantity-Price Elasticity (ϵ_q^p)	-1.012	-0.891	-0.699	-1.076	-1.951	-0.730	-0.898	-0.898	-1.737	-2.274	-1.734	-0.909	-0.910	-0.959	-2.045	-1.934	-0.498	-1.108	-0.138
	(0.464)	(0.396)	(0.455)	(0.498)	(0.782)	(0.490)	(0.446)	(0.446)	(0.769)	(0.979)	(0.576)	(0.639)	(0.444)	(0.472)	(1.031)	(1.822)	(0.493)	(0.597)	(0.523)
N	2,177,475	2,698,083	1,610,056	1,920,615	1,236,855	1,565,202	2,304,439	2,184,025	1,028,495	743,121	1,924,164	1,178,924	2,102,144	1,206,533	201,495	1,345,563	1,409,557	187,758	8,841

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 (Nov 2017–Jan 2018), and post (Feb 2018–Mar 2019, reported).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 adds back observations removed in cleaning steps: (7) adds observations removed because of extreme price variation; and (8) adds observations removed because of extreme price variation and 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 and Chor, 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 implements sample-balance restrictions, including: (14) weekly 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; (18) at the firm-treatment-month level (treated or control); and (19) at the product-month level, including only products sold in all periods by at least 10 firms.

Table A.12: Description of treatment and control group

	Treated	Control	All
Products	25	680	705
Firm-HSNs	19,103	320,005	339,108
Of which before Nov 17	11,070	194,592	205,662
Before November 2017:			
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 rupees)	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 Nov 17	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 in three groups of rows: Products, Firm-HSN, and Observations. Standard deviations are in parenthesis and the median in brackets. para Statistics are measured before November 2017.

Table A.13: Products with a rate reduction in July 2018

Product name	Product code (4-digit HSN)	Treated products (%)	Cumulative share (%)
Glaziers Putty, Calking, Surface Preparations	3214	16.8	16.8
Heating And Thermic Appliances, Hair-Dressers	8516	16.6	33.5
Paints And Vanishes (Non-Aqueous Disolved)	3208	11.8	45.3
Electric Accumulators	8507	11.4	56.7
Paints And Vanishes (Aqueous Disolved)	3209	8.2	64.8
Refrigerators And Freezers	8418	7.1	71.9
Suitcases, Vanity Cases, Briefcases	4202	5.8	77.7
Laudry 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 Vanishes (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.

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. The All column pools all groups. For each variable, the table reports the mean, standard deviation in parentheses, and median in brackets. The last line of each block reports shares relative to total in all periods. Sales and tax amounts are reported in thousand rupees. Value-added is defined as total sales minus intra-state inputs.

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 rupees.

Table A.16: Reconciliation of firm and firm-product level results

	(1) Firm Return	(2) + Detailed Firm Return	(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) firms filing only the firm-level tax return (GSTR-3), (ii) firms filing the detailed firm-level tax return with B2B/B2C information (GSTR-1) along with the firm-level return (GSTR-3), and (iii) firms filing the firm-product tax return (GSTR-1 Table 12) along with the two previous tax returns. 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, inferred from firm-to-firm transactions.

Table A.17: Balanced barcodes by department and treatment status

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 final balanced 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%. Counts reflect unique barcodes that appear in all required periods of the balanced panel.

Table A.18: Treated barcodes by barcode "Family"

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 balanced 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.

Table A.19: Description of treatment and control group

	ETR 28%	ETR below 28%	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 rupees)	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 rupees)	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%.

Table A.20: 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 Bangalure Urban pincode. *Lower-rated product* indicates if the 4-digit HSN 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 8-digit granularity.

Table A.21: 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 2-digit HSN classification when the 4-digit measure is unavailable. All regressions include firm-product-unit and firm-period fixed effects, with standard errors clustered by firm and 4-digit product code.

Table A.22: 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	0.765	0.601	0.385	-0.158
	(2.288)	(2.393)	(1.034)	(0.599)
Dependent variable: Elasticity of consumer price to effective taxes				
Income elasticity	-0.396	-0.151	0.051	0.095
	(0.725)	(0.476)	(0.135)	(0.109)
Observations (Survey Products)	46	46	46	46

Notes: This table reports estimates from regressions of survey-product elasticities on the corresponding Engel-curve coefficients, as specified in equation (21). 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. Standard errors are reported in parentheses. The sample includes 48 survey products treated in November 2017.

A GST context

A.1 Rate definition

The minutes of the 4th GST Council meeting explained the criteria for defining the GST rates based on the tax rates previously in place. As described below, it clearly states the goal to apply GST rates close to the combination of state and union-level rates previously in place.

> *(i) There shall be a category of goods which shall be exempt from GST and this would include items like food grains.*
> *(ii) There shall be a low band of tax rate of 5% and would generally cover goods which presently attract combined tax rate of Central Excise and VAT (including cascading on account of these two taxes) between 3% and less than 9%. Such goods are normally consumed by the vulnerable sections of the society or have high impact on inflation.*
> *(iii) There shall be a standard tax rate of 12% and would generally cover goods which presently attract combined tax rate of Central Excise and VAT (including cascading on account of these two taxes) between 9% and less than 15%.*
> *(iv) There shall be another standard tax rate of 18% and would generally cover goods which presently attract combined tax rate of Central Excise and VAT between 15% and less than 21% (including cascading on account of these two taxes).*
> *(v) There shall be a higher band of tax rate of 28% and would generally cover goods which presently attract combined tax rate of Central Excise and VAT equal to or more than 21% (including cascading on account of these two taxes).*
> *(vi) Supply of services shall generally be taxed at the rate of 18%.*

The initial GST rates were set in May 2017 during the 14th GST Council meeting. In this Meeting the minute reinforces the goal to stay close to existing rates, also indicating that in some cases it indicated a different rate.

> *The Committee had recommended the rates after taking into account the present tax incidence on account of Central Excise, Service Tax and VAT*

(...) In certain cases, the Committee had recommended lower or higher GST rates, vis-à-vis the present tax incidence (including embedded taxes) taking into consideration: (a) the current economic and social realities; (b) ensuring moderate tax incidence on items of common use; and (c) ensuring moderate GST rates for intermediates and capital goods.

More generally, given the high uncertainty in tax collections and the technical implementation challenges of GST, the Central Government and the States, via the GST Council, initially maintained existing rates before undertaking major adjustments. Since any change to rates or exemptions is proposed within the GST Council achieving a qualified majority often was hard. As a result, product-specific adjustments were difficult to implement.

A.2 GST criticis

The political opposition to the government actively criticized the GST reform, particularly the rate differentiation. For instance, Rahul Gandhi, a leading Member of Parliament from the opposition, famously dubbed GST as "Gabbar Singh Tax," referencing the Bollywood villain who terrorizes poor villagers, to suggest that GST was harming small businesses. In a speech, he stated: "GST means Gabbar Singh Tax. This is causing a loss to the country. Small shopkeepers are finished. Lakhs of youths have gone unemployed."[35] In contrast, he referred to his version of GST as the "Genuine Simple Tax."[36] Former Finance Minister P. Chidambaram also voiced strong criticisms of the GST implementation. He pointed to rate differentiation as an issue, referring to GST as "a very, very imperfect law. This is not the GST which we had envisaged... What has been implemented, however, is a GST with seven, or possibly more, rates. It is a mockery of GST... When we have rates like 0.25, 3, 5, 12, 18, 28 and 40, and possibly more because of the discretion vested with state governments, how can we call this a 'one nation, one tax' regime?"[37]

[35] Rahul Gandhi's remarks on GST: https://www.financialexpress.com/india-news/rahul-gandhi-dubs-gst-as-gabbar-singh-tax-mocks-pm-narendra-modi-over-note-ban/903655

[36] Rahul Gandhi's remarks on "Genuine Simple Tax": https://sambadenglish.com/rahul-gandhi-mocks-at-gst-again

[37] P. Chidambaram's remarks on GST: https://www.livemint.com/Politics/9ttJhbK99hgl9rT0kwCP8H/P-Chidambaram-slams-current-iteration-of-GST-calls-it-moc.html

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% to 18%. Without imposing any restriction on feasible rates, the average ETR of treated products prior to the reform was 25.3%, implying a 2.7 percentage point gap relative to the STR (Figure A.15.A). After the reform, the ETR gradually converged to 18% between November 2017 and February 2018. We consider three non-mutually exclusive explanations for the baseline wedge between the STR and the ETR among treated products: limited observability in the data, inattention, and strategic misreporting.

Inattention may account for 33% of the baseline wedge. To estimate it, we use errors that are unfavorable to firms to bound its contribution to the initial STR–ETR gap. Figure A.14.A shows that without restricting to feasible rates 4 percent of treated firm-products continued to remit the pre-reform 28% rate after April 2018, despite the rate reduction having occurred six months earlier, a pattern consistent with mistakes in VAT reporting documented in other countries (e.g., Almunia et al. 2023). Half of the lower rate is explained by firms applying 0%, 5%, or 12%, and these mistakes persist with similar magnitude after the rate reduction. Importantly, panel B and C of Figure A.14 show that these mistakes also occur among control firm-products, particularly before November 2017, when the 28% bracket applied to a broader set of goods. In addition, these errors are symmetric between favorable and unfavorable mistakes for the 12% control products, which occupy the middle of the tax schedule. Because 28% is the highest existing rate, however, unfavorable errors among treated products pre-rate reduction cannot be offset by equally large errors in the opposite direction. Assuming that a comparable share of firms mistakenly applied 18% prior to the reform, inattention alone explains 0.9 percentage points of the wedge if inattentive firms were equally likely to select any existing rate (0%, 5%, 12%, or 18%).

Observability constraints may account for an additional 33% of the initial gap. Specifically, 20% of treated HSN–4 codes include lower-rated HSN–6 subcodes. Since some firms report at the 4-digit level, this within-code rate heterogeneity creates measurement error in the applicable statutory rate. To quantify its importance, we use

the ETRs of firm-products reporting at the 8-digit level to recover the within-HSN–4 distribution of sales across subcodes, and we reweight the STR of 4-digit firm-products accordingly. There are 34 treated products with a lower-rated HSN–6 subcode within the same HSN–4 category. We additionally classify eight products—primarily in Chapter 90 (medical instruments)—as subject to classification disputes based on GST notifications, since their applicable rate depends on the use case (Central Board of Indirect Taxes and Customs 2019). This adjustment reduces the gap by 0.9 percentage points (Figure A.15.B).

Finally, consistent with both inattention and observability mechanisms, fully B2B firms—for which strategic misreporting is less attractive because buyers claim input tax credits—still exhibit a 2.4 percentage point gap (Figure A.15.C). In light of this evidence, non-strategic mechanisms can account for two-thirds of the initial gap. The residual gap may reflect strategic behavior, a possibility we examine directly in the analysis of misreporting (Section 4.2).

To evaluate the extend that the remaining gap may be strategic, Table A.19 compares baseline compliers to non-complier firm-products in the treatment group, defining complier as those with a STR-ETR gap below 1 percentage point. The table indicates that they are broadly similar in observables such as sales, reporting duration, location, and the share of sales to final consumers. The main difference is reporting precision, consistent with greater classification precision.

To quantify the explanatory power of these variables, Table A.20 presents regression results a regression of predictors of strategic behaviour on STR-ETR gap. The analysis reveals that variables plausibly related to strategic behavior are statistically significant. Column 1 shows that firm-product sales are negatively correlated with the gap, with a 10% increase in size associated with a 0.013pp reduction in the gap, while sales to final consumers are positively correlated with the gap, with a 10pp increase in this share associated with a 0.217pp increase in the gap. More granular reporting and a longer reporting history, capturing a "learning" channel, reduce the gap, while location is statistically insignificant. However, these strategic proxies jointly have very limited explanatory power of the overall gap, as indicated by an R^2 of 0.065. Column 2 adds HSN fixed effects, comparing firm-products within HSN codes. The R^2 rises to 0.29, reinforcing that a portion of the discrepancy is HSN-specific. Column 3 includes product and firm fixed effects, which account for firm-level strategic behavior and fully exploit the firm-product dimension of the data. The R^2 increases 0.49, implying

that even under the conservative assumption that all firm-level dispersion reflects strategic behavior, it could explain 20% of STR–ETR gap on top of the product-level dispersion.

Combining the evidence on inattention, observability, B2B firms, and gap predictors, we conclude that the gap is largely non-strategic, although strategic behavior likely accounts for up to 20% of it.

C Elasticity estimation

To formally link our reduced-form event-study estimates to underlying structural parameters, we set up an equilibrium framework 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_{js}^{etr}, and the pre-tax producer price, q_{kjs}, such that $p_{kjs} = (1 + t_{js}^{etr})q_{kjs}$. We specify a general demand function $f(\cdot)$ that depends on this tax-inclusive consumer price, accompanied by a demand shifter A_{kjs}:

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

Taking the natural logarithm yields:

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

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 $\ln(A_{kjs}) = \alpha_{kj} + \gamma_{ks} + \epsilon_{kjs}$.

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

In the empirical estimation, we use a shock to the statutory tax rate. 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 $\frac{\partial \gamma_{ks}}{\partial \ln(1+t_{js}^{str})} = 0$.[38]

[38]Because the policy change is discrete rather than marginal, we estimate arc elasticities over the

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

We define the three structural parameters governing this response.[39] The price elasticity of demand for product j (ε_j^{xp}), the tax pass-through elasticity to the consumer price (ε_j^{p}), and the elasticity of the effective tax rate with respect to the statutory tax rate (ε_j^{etr}):

$$\varepsilon_j^{xp} \equiv -\frac{\partial \ln f(\cdot)}{\partial \ln\left((1+t_{js}^{etr})q_{kjs}\right)} \tag{5}$$

$$\varepsilon_j^{p} \equiv \frac{d\ln\left((1+t_{js}^{etr})q_{kjs}\right)}{d\ln(1+t_{js}^{etr})} \tag{6}$$

$$\varepsilon_j^{etr} \equiv \frac{d\ln(1+t_{js}^{etr})}{d\ln(1+t_{js}^{str})} \tag{7}$$

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

$$\beta^{str} \equiv \frac{d\ln(x_{kjs})}{d\ln(1+t_{js}^{str})} = -\varepsilon_j^{xp}\varepsilon_j^{p}\varepsilon_j^{etr}$$

We also estimate the tax compliance parameter 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: $\beta^{etr} = \varepsilon_j^{etr}$.

$$\ln(1+t_{kjs}^{etr}) = \beta^{etr}\ln(1+t_{js}^{str}) + \alpha_{kj} + \gamma_{ks} + \nu_{kjs} \tag{8}$$

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 $\ln(1+t_{kjs}^{etr})$ with its fitted values, $\ln(\widehat{1+t_{kjs}^{etr}})$.

As we directly observe prices, we can apply the identical empirical strategy to the

reform interval.

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

tax-inclusive consumer price. This coefficient isolates the equilibrium consumer price response, which perfectly captures the product of the effective pass-through elasticity (ε_j^p) and the tax compliance elasticity (ε_j^{etr}):

$$\ln(p_{kjs}) = \beta^p \ln(1 + t_{js}^{str}) + \alpha_{kj} + \gamma_{ks} + \omega_{kjs} \tag{9}$$

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 gives and using the relation between post- and pre-tax price elasticity $\varepsilon_j^p = (1 + \varepsilon_j^q)$:

$$\begin{aligned}
\varepsilon_j^s &\equiv \frac{d\ln(s_{kjs})}{d\ln(1 + t_{js}^{str})} \\
&= \frac{d\ln(x_{kjs})}{d\ln(1 + t_{js}^{str})} + \frac{d\ln(q_{kjs})}{d\ln(1 + t_{js}^{str})} \\
&= \varepsilon_j^{etr}\left[\varepsilon_j^p(1 - \varepsilon_j^{xp}) - 1\right]
\end{aligned} \tag{10}$$

C.1 Evasion

We can extend the framework to a setting with 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:

$$x_{kjs} = A_{kjs} f\left((1 + t_{js}^{etr})q_{kjs}\right) \lambda(1 + t_{js}^{str}) \tag{11}$$

Taking the natural logarithm yields:

$$\ln(x_{kjs}) = \ln(A_{kjs}) + \ln f\left((1 + t_{js}^{etr})q_{kjs}\right) + \ln \lambda(1 + t_{js}^{str}) \tag{12}$$

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:

$$\beta^x_{rep} \equiv \frac{d\ln(x_{kjs})}{d\ln(1+t^{str}_{js})} = \frac{d\ln(x^*_{kjs})}{d\ln(1+t^{str}_{js})} + \frac{d\ln\lambda(1+t^{str}_{js})}{d\ln(1+t^{str}_{js})} \tag{13}$$

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

$$\beta^x_{rep} = \beta^x + \varepsilon^\lambda \tag{14}$$

This composite parameter highlights the identification challenge introduced by evasion. If a decrease in the statutory tax rate reduces the share of evasion—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 actual demand expansion and formalization. Notice that any baseline differences in evasion across firm-products, or any aggregate shocks to reporting behavior across firm-periods, are absorbed by our fixed effects (α_{kj} and γ_{ks}). Therefore, the bias in the estimated coefficient β^x_{rep} only arises if the level of evasion is directly responsive to the product statutory tax rate. In addition, the reported sales coefficient, β^s_{rep}, is the coefficient of interest for optimal taxation in the setting incorporating evasion in Section D.5.

D Optimal Commodity Taxation

D.1 Baseline Model

We derive a general expression for optimal indirect 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 ownership. We consider a continuum of households $i \in [0,1]$, each with exogenous income y^i in the baseline model. 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 exogenous 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 represents the chosen consumption vector and $u^i(\cdot)$ is the household-specific direct utility function. Firms operate under constant returns to scale and perfect competition, so that profits are zero and producer prices are independent of taxes. These assumptions allow us to abstract from firm behavior and ownership in the baseline case. The government maximizes a social welfare function $G(\cdot)$ subject to a revenue requirement. The associated Lagrange multiplier μ represents 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 assuming no cross-price effects, 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

$$\varepsilon_j^x \equiv -\frac{\partial x_j}{\partial p_j}\frac{p_j}{x_j}.$$

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

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

Interpretation This expression makes explicit the trade-off faced by the government. The denominator, ε_j^x, captures the responsiveness of the tax base: under full pass-through, taxes affect consumer prices one-for-one, so the elasticity of the tax base coincides with the elasticity of quantities demanded. 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 components.

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

$$\frac{t_j}{1+t_j} = \frac{1-\bar{g}/\mu}{\varepsilon_j^x}.$$

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.

D.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.

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

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

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

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

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

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

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

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

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 ε_j^p, 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.

D.3 Firm Ownership

We now extend the analysis by allowing individuals to be firm owners. Individual income includes profit income, 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 $\pi_j = q_j x_j$ are aggregate profits. 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

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

where

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

Assuming that profits are a constant proportion of sales revenue, $\pi_j = \gamma q_j x_j$, where $\gamma \in (0, 1]$ captures the profit margin. Under this assumption, the proportional response of profits to taxation is equal to the response of sales revenue: $\varepsilon_j^\pi = \varepsilon_j^s$.
Substituting this relationship into the first-order condition yields the optimal tax rule:

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

Interpretation Allowing for firm ownership introduces an additional welfare channel through profit income. In addition to consumer-price effects, the numerator now captures how taxation affects individuals' incomes through changes in profits, weighted by distribution of profit income of good j across individuals and the social marginal welfare weights of firm owners. Under costless production, profits are proportional to sales, so that profit responses to taxation are governed by the same elasticity as quantities sold. 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 profit income is zero (due to offsetting cost or not being considered) or firm owners have zero social weight.

D.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:

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

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

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

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:

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

Interpretation The generalized optimal tax formula highlights two distinct 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 responds not only to changes in the price of good j itself, but also to induced changes in the prices of goods $m \neq j$. These effects are weighted by individuals' consumption shares, social marginal welfare weights, and by the relative aggregate importance of good m compared to good j. As a result, goods whose taxation indirectly raises the prices of goods that are 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 alters the quantities sold of other goods through substitution effects, thereby affecting government revenue collected

on those goods. These effects depend on cross-elasticities of the tax base and are scaled by existing tax wedges and by the relative size of the tax base associated with each good. When taxing good j reduces the tax base of goods that are already highly taxed or that account for a large share of tax revenue, the marginal efficiency cost of taxation is amplified.

D.5 Optimal Tax Rule with Evasion

In this section, we introduce evasion, allowing the reported quantity (x_j^{rep}) to differ from actual consumer demand (x_j). Let $\lambda_j(t_j) \in (0, 1]$ denote the share of quantities sold that are formally reported to the tax authority, such that $x_j^{rep} = x_j\lambda_j(t_j)$. To formalize the consumer's problem under evasion, we must adjust the prices they effectively face. Let $\lambda_j^c \in [\lambda_j, 1]$ denote the share of consumption on which the consumer effectively bears the tax burden. The consumer pays the tax-inclusive price $p_j = (1+t_j)q_j$ on the taxed share (λ_j^c) and the pre-tax price q_j on the evaded share ($1 - \lambda_j^c$). The effective price for good j is a weighted average: $p_j^e = (1 + \lambda_j^c t_j)q_j$.
The indirect utility function for household i incorporates the effective price vector, p^e:

$$v(p^e, y^i) = \max_{x^i} \; u^i(x^i) \quad \text{s.t.} \quad p^e \cdot x^i \leq y^i \tag{15}$$

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

$$W = \int_i G\big(v(p^e, y^i)\big) \; di \; + \; \mu \sum_j t_j q_j x_j \lambda_j \tag{16}$$

Applying Roy's identity, the welfare effect of a tax change depends on the incidence of the evasion surplus. Defining the social marginal utility of income for individual i as $g^i \equiv G'(v)\frac{\partial v}{\partial y^i}$ as before, the first-order condition of welfare with respect to the tax rate t_j incorporates this modified consumer incidence:

$$\begin{aligned}
\frac{dW}{dt_j} &= -\int_i G'\big(v(p, y^i)\big) \, \frac{\partial v(p, y^i)}{\partial y^i} (\lambda_j^c x_j^i) \frac{\partial p_j}{\partial t_j} \, di \; + \; \mu \left(q_j x_j \lambda_j + t_j \frac{d(q_j x_j \lambda_j)}{dt_j} \right) = 0 \\
&= -\int_i g^i \lambda_j^c x_j^i \frac{\partial p_j}{\partial t_j} \, di \; + \; \mu \left(q_j x_j \lambda_j + t_j \frac{d(q_j x_j \lambda_j)}{dt_j} \right) = 0
\end{aligned} \tag{17}$$

Define:

$$\varepsilon_j^p \equiv \frac{\partial p_j}{\partial t_j}\frac{1+t_j}{p_j}, \qquad \varepsilon_j^s \equiv -\frac{d(q_j x_j)}{dt_j}\frac{1+t_j}{q_j x_j}, \qquad \varepsilon_j^\lambda \equiv \frac{d\lambda_j}{dt_j}\frac{1+t_j}{\lambda_j} \tag{18}$$

With sales, reported sales and evasion elasticities relating as: $\varepsilon_j^{s,rep} = \varepsilon_j^s - \varepsilon_j^\lambda$. 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 evasion:

$$\boxed{\tau_j = \frac{1 - \frac{\lambda_j^c}{\lambda_j}\int_i \frac{g^i}{\mu} a_j^i \varepsilon_j^p \, di}{\varepsilon_j^{s,rep}}} \tag{19}$$

Interpretation 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, which can also be interpreted as the indirect incidence of taxes on informal market prices. When $\lambda_j^c = \lambda_j$, the consumer fully captures the surplus (e.g., via lower prices in the informal sector). 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 efficiency leakage 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 ($\varepsilon_j^s - \varepsilon_j^\lambda$). Because tax increases leads to higher evasion ($\varepsilon_j^\lambda < 0$), 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 in a perfectly compliant economy.

E Heterogeneity by Products' Engel Curve Slopes

A key dimension of the equity–efficiency trade-off in indirect taxation is product-level demand elasticity. Governments may wish to tax goods disproportionately consumed by richer households, but such targeting is less attractive if those goods exhibit high elasticities and therefore entail larger efficiency costs.

Interaction term The first was to asses this trade-off is interacting the income elasticity of each product with the treatment status in the elasticity estimations. Formally, we run:

$$y_{iht} = \beta_1 \widehat{\text{ETR}}_{iht} + \beta 2 \left(\widehat{\text{ETR}_{iht} \times \hat{\beta}_j^{\text{Engel}}} \right) + \alpha_{ih} + \gamma_{it} + \epsilon_{iht} \tag{20}$$

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.

The interaction term gives us how the elasticity vary with the income elasticity. Table A.22 presents the results. It reveals that the interaction coefficient is not significant and very close to zero for the elasticity of sales, quantity and consumer prices to effective tax rates, and quantity to consumer prices.

Meta regression The second approach employes a meta regression. We first estimate Engel curve slopes for each survey product. Treated survey products span a range of Engel elasticities: the average is 1.8, but the 25th percentile is below one and the 75th above three (Figure A.16). We then correlate these Engel elasticities with our estimated quantity and price elasticities with respect to the ETR at the product level. We regress products specific-elasticities estimated in the administrative data-on products' Engel curve slopes.

$$\hat{\beta}_j^{\text{elasticity}} = \alpha + \delta \hat{\beta}_j^{\text{Engel}} + \varepsilon_j \tag{21}$$

.

We find essentially a flat relationship, suggesting that product "luxuriousness" is not a key determinant of behavioral responses in our setting (Figure A.17). Table A.22 reports the results under different weighting schemes (by product consumption share in the survey, equal product weighting, number of observations in tax data, and sales share in the tax data). None of the specifications yield statistically significant coefficients. The point estimates range from –0.12 to 0.45 for the quantity elasticity (and from –0.13 to 0.19 for the implied price elasticity). In the main specification using the firm-product tax returns, we have estimated a quantity elasticity of 1.12. Taken at face value, these coefficients imply that even a sizable difference of one in the Engel-curve slope between two goods would translate into only modest changes in the quantity elasticity on the order of –12% to +17%.

We then plot the relationship between our estimated quantity and price elasticities with respect to the ETR and each product's Engel-curve slope (Figure A.17). Panel (a) reports the correlation for quantity elasticities, and Panel (b) for consumer price elasticities. We find no systematic variation in elasticities along the Engel curve, suggesting that product "luxuriousness" is not a key determinant of behavioral responses in our setting.