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

Pierre Bachas, Davi Bhering, Pulak Ghosh*

January 13, 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 only partially passed through to prices, generate modest quantity responses, and induce little product mislabeling. As a result, well-targeted differentiated commodity taxes can achieve substantial redistribution at moderate 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 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 both direct and indirect efficiency costs. First, the ten–percentage-point statutory tax rate cut led to an average decline of 6.6 percentage points in effective tax rates. The remaining gap is due to some firms applying rates different from 28% to treated products, pre-reform. About half of this imperfect baseline compliance seems driven by non-strategic ambiguity or mistakes. Conditional on baseline compliance, the decline in effective rates matches the statutory cut.

Second, we find substantial but incomplete pass-through of the tax cut to consumer prices. On average, prices of treated products decline by 2.7%: only 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 cuts are moderate. At the extensive margin, we can rule out even small increases in the entry of treated products in both existing and new firms. At the intensive margin, quantities sold rise by 3.2% (relative to a 2.7% decline in prices), leading to an increase in pre-tax sales values of 5.8%. These behavioral responses are too small to compensate for the mechanical revenue loss from the rate cut, implying a substantial decline in tax revenues. Firm-level value-added mirrors the sales response. Sales elasticities are larger among firms filing simplified tax returns and among smaller firms, but show no systematic relationship with other proxy measures of exposure to informal-sector competition.

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, the data cover a longer period, extending to the pre-GST period, reports products at the barcode level, and only captures final-consumer sales. The

[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.

results are qualitatively similar, with somewhat lower pass-through to consumer prices and slightly larger quantity responses, consistent with heterogeneity by firm size in the tax data and some market power among large retailers. Second, we analyze the July 2018 reform, which reduced the tax rate on the smaller set of goods that remained taxed at 28%. Its effects are similar to those of the November 2017 reform.

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, aggregation to the product level, 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 14 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

[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.

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 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 the main 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 discusses equity. Section 6 calibrates optimal luxury tax premia.

1.1 Related Literature and Contributions

This paper contributes to three related strands of the literature. First, we contribute to the literature on the pass-through of indirect taxes (see recent studies by Benedek et al. 2020; Benzarti et al. 2020; Carbonnier 2007; Harju et al. 2018; Velayudhan and Brusco 2022). Leveraging a rich dataset, we precisely measure how firm prices respond to large statutory rate reductions across a broad set of goods. Our pass-through estimates fall squarely within the existing range: tax cuts are only partially passed through to consumer prices, with lower pass-through among larger firms. Our results underscore the importance of accounting for changes in effective tax rates following statutory reforms in developing countries.

Second, we expand on existing research on the efficiency costs associated with indirect taxes (Agarwal et al. 2025; Belan and Gauthier 2006; Gaarder 2022; Garg et al. 2023; Harris et al. 2018) and the impact of their design for revenue collection (Pomeranz 2015; Waseem 2023). Our study leverages a large, sudden reform applied to a wide range of normal goods, uniquely combining detailed product-level tax data from multi-product firms with granular scanner data. We provide direct estimates of efficiency costs by directly analyzing quantity responses and exploring effects at the extensive margin (Benzarti and Carloni 2019; L. Liu et al. 2021).

Third, we contribute to the literature on optimal tax and transfer design in developing countries. A strand emphasizes the role of second-best tax designs in settings with administrative and informational constraints (Best et al. 2015; Huang and Rios 2016; Londoño-Vélez and Ávila-Mahecha 2024). We examine the role of higher rates on luxuries in contexts with constrained income taxes. By investigating substitution and mislabeling behaviors—issues frequently raised in policy discussions but rarely assessed empirically—our study provides critical insights into the practical implications of multi-rate VAT systems. The targeting power of luxury taxation,

connects us to the much more extensive literature on transfer design; including concerning India's main social assistance program.

2 Context

India's Goods and Services Tax The Goods and Services Tax (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

[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.

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 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.[5]

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.

[5]The average annual growth rate of state taxes was 7.6% between fiscal years 2013 and 2016 (Gupta and Rajaraman 2020).

3 Data and Empirical Strategy

3.1 Tax Rates

We assemble a dataset of GST rates by products from July 2017 to March 2020 by merging two main sources. The first source is the publicly available rate booklets for July 2017 and March 2020 from the Ministry of Finance. These booklets are the official source available to firms to check tax rates. In each tax bracket, the booklets list the corresponding four-digit HSN codes along with a short product description. Most codes appear only once, but if a subset of the four-digit code, typically a six-digit code, is taxed differently, the four-digit code appears more than once with different descriptions. These descriptions usually correspond to names of six-digit HSNs, which we can thus link to the six-digit HSN database via fuzzy matching. Combining the 2017 and 2020 booklets yields the initial and final rates for every product.

The second source is the set of GST Council meeting minutes. These minutes record each rate decision in tables that mirror the booklet format, listing the affected four-digit codes and product descriptions. We extract these tables to date every rate change and cross-check them against the changes inferred from the 2017 and 2020 booklets. When a minute refers to a narrower product than the four-digit level, the accompanying description enables us to identify the relevant six-digit codes within that four-digit 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).[6] 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 comprises the tax returns of all GST-registered firms in Karnataka from July 2017 to March 2019. We complement this data with the nationally representative expenditure survey for India, and a dataset of transactions at the barcode

[6]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.

level 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.[7] Firms below the threshold can either register for GST voluntarily or participate in the turnover tax composition scheme.[8] 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). For B2B sales, the return records invoice-level information with sales, taxes and counterpart.

Firm by Product Tax Returns Firms with annual sales above 15 million rupees must also fill a complementary tax return which includes a summary of sales by product (GSTR-1 Table 12). This tax return requires firms to report monthly sales, taxes, quantities, quantity unit by HSN code.[9] The reporting of quantities, sales, and taxes by product, also allows us to calculate pre-tax and tax-inclusive prices (henceforth, consumer prices).

Our main analysis uses firm-product tax returns, which offer key advantages over standard VAT filings that seldom distinguish products and rarely report quantities or prices. In Karnataka, only about one third of returns include HSN-level details, but these firms account for three-quarters of total sales and tax payments. The resulting data provide granular insight into firm behavior at the product level while still capturing the vast majority of economic activity and revenues in the state.

[7]We use the official end-2017 exchange rate of 63.93 INR per USD as the reference rate for all currency conversions.

[8]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.

[9]Approximately half of the firm-HSN-month observations are reported at the HSN-4 level, which is the level of most rate changes, and half at the much finer HSN-8 level. We keep the data at the firm-reported level, since aggregating across HSN codes that firms report separately makes quantities and implied prices harder to interpret, even within the same HSN-4.

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 dataset The data contain barcode-level transactions from one of the largest retail chains in Karnataka, from April 2016 to July 2019. Barcodes correspond to detailed goods, distinguished by brand and size (e.g. "Heinz Ketchup 200g Bottle", "Nivea Men's Deo Sport 150ml"). For each barcode it reports quantity, pre-tax value, tax-inclusive value, and discounts. We compute the tax paid as the pre- to post-tax sales difference.

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.[10] Since the survey does not collect income information, we rank households based on total per capita expenditure.[11]

3.3 Data cleaning and sample restrictions

Sales to final consumers Given the credit-invoice structure of the VAT, taxes on sales to other VAT-registered businesses and taxes on sales to final consumers have different implications. Registered businesses can claim input credits and therefore recover the VAT paid on their purchases, while final consumers cannot. To assess the equity and efficiency implications of consumption taxes, we focus on the burden falling on final consumers. Our analysis therefore aims to be restricted to transactions where the VAT is not creditable. Final consumers include not only individuals but also entities that cannot claim input credits, which include, firms that are unregistered, under the composition scheme, or selling exempted products only.

An empirical challenge is that firms do not report sales to businesses versus final consumers at the product level. However, they report this breakdown in their monthly

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

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

firm-level VAT returns, which provide the total share of sales to each type of client. Our main analysis focuses on firms that primarily sell to final consumers, defined as those with a baseline (pre–November 2017) share of sales to final consumers exceeding 50%. Applying this restriction removes 53% of observations.[12]

Given the size of this restriction, we run extensive robustness checks: imposing alternative thresholds to define final-consumer-facing firms; restricting the sample to products covered in the household consumption survey, the hypermarket dataset, or those classified as downstream by Antràs and Chor (2012); and redefining the share of sales to final consumers using the full sample period rather than the baseline period.

Exclusion of Services We exclude services from the analysis. Services are classified under the SAC system, which is far coarser than the HSN codes for goods: at the two-digit level, all services fall under Chapter 99, making it hard to identify rate changes for specific services within a given subcategory. Moreover, GST rates for some services depend on attributes not captured by SAC codes—for example, hotel rates vary by room tariff, and restaurant rates on whether the establishment was air-conditioned. Finally, defining quantities is conceptually difficult for many services, and the aggregation of heterogeneous activities within broad SAC codes makes it challenging to construct meaningful and consistent service-level quantity measures.

Data quality issues We take several steps to ensure data quality. Because the GST portal does not provide a drop-down menu for HSN codes, firms can report inexistent or missing product codes; we drop these observations. We also remove observations where firms report sales but omit quantities, to maintain a stable sample across dependent variables. Table A.3 summarizes these cleaning steps: 2.9% of firm–HSN–month observations have invalid HSN-4 codes, 0.8% report coarse HSN-2 or incomplete HSN-4 codes, 1.5% have missing codes, and 5.7% do not report quantities.

One concern is that firms may inconsistently report quantities, as quantities do not affect tax liability, or use the same HSN code for a varying bundle of goods within a given code. We thus exclude firm–products with implausibly large price variation over time: specifically, those where the ratio of the highest to the lowest reported

[12] 49% of observations have a baseline B2C share below 50%, and an additional 4% first appear after November 2017. Because the baseline share is defined using pre–November 2017 data, firms entering after this date are assigned a missing value. We prefer this approach because the reform may affect the share of sales to final consumers, but relax this restriction when studying firm entry.

price exceeds 100. This leads us to drop 1.4% of observations; to be conservative, we remove the entire firm–product across all periods rather than only the outlying 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. 165 out of the 1368 products fall in this category (Table A.4). We analyze the 105 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, 85% 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. 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.[13]

Table A.7 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

[13] Table A.6 shows similar patterns when weighting by pre-reform sales.

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.8), although the top item"Medicaments in Measured Doses" accounts for a substantial share of control observations (22%), and will thus be dropped in some robustness tests. The top 30 products together represent 86.8% 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. The sample includes 846 products: 165 treated (from 28% to 18% rate) and 681 control products (that remained at 18% or 12% rate). Since firms often sell multiple HSN codes, the data contain 491,476 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 5.6 of the 21 months, yielding 2,728,208 firm–HSN–month observations, of which 401,715 before November 2017. Sales are skewed: before November 2017, the median monthly sales per firm–HSN are roughly 6,780 rupees (USD 106) in months when the product is sold, while the mean reaches 407,000 rupees (USD 6,360). Treated and control firm–HSNs look similar along most pre-reform dimensions.

Most firms sell multiple products: Table 2 shows that out of the 20,000 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 13.2 different HSNs per month on average in the four pre-reform months, out of which 2.5 are treated, 6.0 are control and 4.9 are excluded.

3.5 Empirical strategy

To analyze behavioral responses, we employ an event-study difference-in-differences (DiD) approach at the firm-HSN-month level. 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. The regression specification is given by:

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

In this specification, y_{iht} denotes the outcome of interest for firm i, HSN h, in month t. The indicator treatment_h equals 1 if HSN h faced the 10 percentage point rate cut (from 28% to 18%) in November 2017 and 0 otherwise. The coefficients 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 level to account for potential serial correlation and product-specific shocks.

This specification corresponds to the standard DiD framework, featuring a single treatment timing in November 2017, no reversal of treatment status, and a binary treatment variable indicating exposure to the 10 percentage point tax cut. Identification relies on the standard parallel trends assumption: in the absence of the reform, the outcomes of treated and control products within firms would have evolved similarly. The event-study specification also allows us to assess pre-treatment trends, providing a check for the parallel trends assumption. We also present results for a before and after DiD specification. In this case, we compare the months before the rate reduction—July to October 2017—with months after—April 2018 to March 2019—allowing for a transition period and capturing a full year post reform.

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 HSN 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 could also affect control products sold by treated firms through spillovers, for example, if lower prices on treated products draw more consumers to the store. Such spillovers would violate the Stable Unit Treatment Value Assumption (SUTVA). We provide an analysis that bounds the plausible magnitude of these effects on control products.

4 Main Results on Administrative Data

4.1 Effective Tax Rates

From Statutory to Effective Tax Rates In contexts with constraints to administrative capacity, it cannot be taken for granted that changes in statutory tax rates (STR)—the rates set by the government—translate into changes in effective tax rates (ETR), defined as the rates firms actually pay. With differentiated rates, it is also not guaranteed that firms apply the prescribed rate accurately, particularly since the tax declaration system did not automatically map products to their corresponding statutory rates. A key advantage of the administrative data is that it allows us to observe monthly the tax rates firms actually recorded and paid for each product sold.

Figure 3.A shows the average STR and ETR for treated and control products. By definition, the STR of treated products dropped from 28% to 18%. Yet, prior to the reform, the average ETR was only 25.3% and gradually converged to 18% between November 2017 and February 2018. For control products, the STR remained at 15.6%, which is consistent with the treatment group composition of products taxed at 12% and 18%. The average ETR was closer to 15.0% and remained stable after November 2017. Figure 3.B presents the results of the event-study in Equation 1 with ETR as the dependent variable. The ETR reduction was 6.5 percentage points on average, compared to an STR reduction of 10 percentage points. This numbers indicate that 65% of the reduction in STR translated in a reduction in ETR.

Baseline STR-ETR gap The 3.5 percentage point gap between STR and ETR for treated products at baseline warrants closer scrutiny. What explains this discrepancy, and does it reflect strategic behavior?

One source of this discrepancy arises from specific sub-products that are taxed at lower rates than the broader HSN code under which they fall. For instance, "bindis" are tax exempt, but are part of the treated "make-up" code (HSN 3304). We cannot identify such sub-products precisely since the differentiation occurs below the level at which firms report HSN codes: in this case, bindis correspond to a 10-digit product code (HSN 33049940). A related issue occurs when an HSN-4 category contains multiple HSN-6 codes subject to different rates. Firms may misreport the applicable rate or report only the HSN-4 code even when they sell an HSN-6 product taxed differently. To assess this channel, Figure A.2 compares the ETR of treated products with and without specific sub-products and HSN-6 codes taxed at a different rate. When restricting to treated products with no lower-rated subproducts, the pre-reform ETR rises to 26.5%, suggesting that lower-rated sub-products might explain up to 40% of the initial ETR–STR gap, but cant fully account for it.

Another source is inconsistent tax filing. Under the GST reporting, firms can freely enter tax amounts, including under rates that do not actually exist. To assess this, Figure A.3 examines the distribution of ETRs for treated products before and after the rate change. Before November 2017, 74% of treated firm–products applied an ETR of 28%, consistent with the STR. The remaining 26% of firm–products applied lower rates, generating the initial ETR–STR gap. Among these non-compliers, firms overwhelmingly reported rates that actually exist; only 12% fall outside these values (8% when allowing for values very close to an existing rate). After the rate reduction compliance improves: over 80% of treated firm–products now report the correct 18% rate, while the densities at 5% and 12% remain unchanged. Notably, 4% continue to pay 28% even after November 2017, despite the reduction. This suggests that filing errors are not fully strategic: firms also make mistakes that work against their own interest. Moreover, because 28% is the highest rate in the system, symmetric mistakes mechanically generate a downward bias in ETRs: firms can err only by applying lower existing rates when the correct STR is 28%.

To further investigate strategic behavior in the initial STR-ETR gap, Table A.9 divides firm-products into baseline compliers (STR-ETR match) and non-compliers (STR-ETR don't match). The table presents statistics pre-November 2017 for these

two groups. Non-complying observations are similar in value and, if anything, sold by larger firms than complying observations. Both groups are equally urban and similarly exposed to final consumers. The only notable difference is reporting precision: complying firm-products are reported with at a more detailed HSN level, consistent with these firms reporting more precisely the products they sell for tax classification.

To pin down the share of the mismatch that might be explained by these factors, Table A.10 presents the results of the regression of the STR-ETR gap on these variables, controlling for HSN code combined with month fixed effects and restricting the sample to treated products in the months preceding November 2017. The correlations reveal that the same HSN sold by a firm with a higher share of sales to final consumers is more likely to exhibit a larger gap (i.e., a lower ETR). The coefficient for the number of months the firm sold the HSN before November 2017 is negative and statistically significant, indicating that more recurrence is correlated with more accurate reporting. The magnitudes, however, are rather small. A 3 percentage point increase in the share of sales to final consumers, such as the difference between groups with and without STR-ETR matching, is correlated with only a 0.06 percentage point lower ETR. Even an order-of-magnitude larger difference of 30 percentage points could explain only about 25% of the initial STR-ETR gap of 2.5 percentage points observed in the treated group. The relatively low R^2 of 0.27 further confirms that the available variables and fixed effects can explain only a fraction of the observed variation.

Overall, the STR-ETR gap we observe can be partially explained by specific sub-products taxed at lower rates and by firms reporting incorrectly. Although we find evidence of strategic mismatches between statutory and effective tax rates, these factors cannot account for the bulk of the gap. It is also worth noting that the initial months following the introduction of the GST may have been a particularly challenging period for firms to assign the correct rates, which could have further contributed to the observed discrepancies. We return to this point when discussing the implications of having multiple tax rates.

4.2 Price Incidence

The extent of tax pass-through to consumer prices determines whether producers or consumers benefits from the tax rate change. Figure 4.A shows that consumer prices decline gradually after the reform, stabilizing after three months at an average reduction

of 0.027 log points. While economically meaningful, this decline is substantially smaller than the 10 percentage point (0.081 log-point) cut in statutory rates. Relative to statutory rates, pass-through is 33 percent; however relative to effective tax rates—the relevant margin for firms' pricing decisions—pass-through rises to 51 percent, with a 95 percent confidence interval of 33 to 69 percent. The estimates are precise enough to reject both zero and full pass-through.

To confirm that changes in effective tax rates (ETRs) drive consumer price responses, Figure A.4 splits the sample by whether STRs and ETRs coincided prior to November 2017. Treated firm–products with incorrect initial ETRs experience only a small decline in ETRs (0.022 log points) and show no change in consumer prices. By contrast, for firm–products with correct initial rates, ETRs closely track the statutory cut, and consumer prices fall by 0.040 log points. The implied pass-through is close to 60%, allowing us to reject both zero and full pass-through. Overall, this exercise confirms that price adjustments respond to changes in ETRs and that pass-through with respect to effective tax rates is stable at slightly above one half.

An implication of this result is that pass-through estimates based on Consumer Price Index (CPI) data may be misleading if changes in ETRs are not observed. Ignoring the mapping from statutory to effective rates can spuriously attribute incomplete pass-through to other mechanisms. This concern is relevant for developing countries, where gaps between STR and ETRs can be substantial.

4.3 Quantity, Sales and Extensive Margin Responses

Quantities Having established that the reform led to a reduction in consumer prices, we analyze how the quantities sold responded to the rate cut. Observing quantities directly is rather rare since, in most countries, the VAT only require firms to report sales. The elasticity of quantities is a key parameter as it determines the dead-weight loss of taxation. Figure 4.B presents the results of Equation 1 for quantities. They follow an inverse pattern to prices in the first three months post reform. While prices gradually decrease once the reform is implemented, quantities gradually increase. After this transition period, the quantities of treated products remain fairly stable, with an average increase of 0.032 log points, which is marginally significant statistically.

Compared to the 10 percentage point (0.081 log points) rate reduction, the coefficient implies an elasticity of quantities to the STR of 0.39. Given that the

STR reduction only partially translated in ETR reduction, the figure presents the elasticity of quantity with respect to ETR. The change in ETR is instrumented by the treatment status and the three broad periods (pre-reform, transition, post-reform "equilibrium"). When considering the change in ETR, the quantity elasticity increases to 0.58, but remains modest. We also estimate the implied elasticity of quantities to consumer prices, instrumenting the change in consumer prices with the treatment. The (endogenous) quantity elasticity with respect to prices is 1.2.

Sales Pre-tax sales determine the tax base over which the tax is collected and the extent to which the tax base responds to rate changes defines the tax revenue impact from the reform.[14] When the base is highly elastic, the increase in reported sales may partly offset the rate cut. Particularly in developing countries, there is concern that the elasticity of the tax base is high due to the relative ease of under-reporting sales, and shift activity to the informal sector. Higher elasticity of sales with respect to the tax rate would imply lower revenue-maximizing rates.

Pre-tax sales equal quantities sold times pre-tax producer prices. Figure 5.A shows that producer prices increase following the rate reduction. This rise is the mechanical counterpart of incomplete pass-through to consumer prices. Figure 5.B shows that the joint increase in producer prices and quantities raises pre-tax sales by 0.058 log points, nearly one-for-one with the decline in effective tax rates. Accordingly, the implied elasticity of sales with respect to the effective tax rate is precisely estimated at 1.0.

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.[15]

[14] Firms' final tax liability depends on value added, sales after deducting input tax credits. We analyze value-added in Section 4.8, as it is only observed at the firm level (and not product level).

[15] 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.5.B).

Figure 6.A plots the reporting rate for treated and control products. The probability that a firm sells a given product rises steadily over the sample period, but treated and control products follow nearly identical trajectories. This common upward trend is consistent with the expansion of the GST systems and the growth of the Indian economy during this period. We estimate equation 1 using the binary reporting indicator to capture the within–firm extensive-margin response. Figure 6.B plots the difference in reporting rates between treated and control products. Products subject to the rate cut show no meaningful change in their likelihood of being sold: the estimated effect is essentially zero, with a precisely estimated coefficient of –0.001 percentage points relative to a baseline mean of 0.25 percentage points. We can therefore reject, at the 95% confidence level, any increase in the reporting rate larger than 0.5%.

4.4 Product Misreporting and Substitution

Defining misreporting and substitution A key concern in multi-rate commodity tax systems is that they 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 is costly, the rate harmonization implemented in November 2017 removed incentives to mislabel and should have led firms to discontinue misreporting.[16] In practice, such mislabeling could attenuate the fiscal gains from the rate reduction to the extent that the observed increase in reported sales reflects a reclassification away from other products.

A second concern is that differential tax rates distort relative prices. These tax-induced wedges may alter consumer choices through substitution toward related, lower-taxed products. Because substitution requires a real reallocation in the consumption basket, it entails welfare consequences in addition to the revenue impact. In our setting, the rate reduction generated a price reduction that could induce consumers to reduce consumption of the previously lower-taxed good in favor of the good whose rate was reduced. Although mislabeling and substitution do not have identical welfare implications, they can generate similar patterns in the data for products that are close substitutes or plausible candidates for misreporting. In both cases, we would expect a

[16]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.

decline in the reported quantities and sales of lower-taxed goods that are related to the goods experiencing a rate reduction.

Defining related products To make progress, we assume that misreporting occurs toward products classified as related in the Indian tax system, for which misclassification is easier to justify during an audit and can plausibly be framed as an honest mistake. In our setting, the GST relies on the Harmonized System of Nomenclature (HSN), with most rate differentiation occurring at the HSN-4 level. We therefore use the HSN hierarchy to define related products, which provides a natural, policy-determined, and exogenous grouping. We divide untreated products into two categories: (i) unrelated HSN-4 products, whose HSN-2 chapter contains no treated products; and (ii) related HSN-4 products, defined as those that share an HSN-2 code with at least one treated product. This classification isolates the products for which misclassification is most plausible while providing a clean control group (see Table A.11). Substitution is also more likely within HSN-2 categories. However, as we discuss below, the plausibility of substitution varies considerably even within an HSN-2, and we later restrict the set of potential substitutes to obtain a tighter group of products that are more likely affected by substitution.

Table A.12 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).

Results on related products We estimate equation 1 at the firm–product level, comparing within-firm outcomes for related and unrelated control products. If mislabeling is pronounced, we would expect the quantities and sales of related products to decline relative to unrelated products. Figure 7 reports results for quantities, sales, consumer prices, and reporting rates of related products relative to unrelated products. Panel A shows no statistically significant difference in quantities sold. The point estimate is actually positive (0.012 log points, s.e. $= 0.020$), allowing us to rule out a decline in quantities of more than 1.6 percent at the 95 percent confidence level. Consistent with essentially flat prices (Panel D), the estimates for sales closely

mirror those for quantities (Panel B). Mislabeling could also occur via the extensive-margin, with firms ceasing to report related products. Panel C shows no differential change in reporting rates over time, ruling out economically meaningful mislabeling or substitution responses.

Mislabeling or substitution may be detectable only when treated products represent a meaningful share of activity within an HSN-2 chapter. To assess this, Figure A.6 restricts the sample of related products to chapters in which treated goods account for at least the median pre-reform sales share (14%). The results are nearly identical to those in the baseline specification. A related concern is that effects may occur only within firms that sell both treated and related products. We thus further restrict the sample to firm–product pairs in which the firm sold both a treated and a related product prior to November 2017. Figure A.7 shows that the results remain similar: quantities and sales exhibit small, positive, and statistically insignificant changes.[17] Last, substitution may also affect goods initially taxed at 5%, which were excluded from our main control group. Figure A.8 assesses the impact on related products including those taxed at 5%, and shows qualitatively similar results.

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—namely, the four-digit HSN classification—even though they do not preclude such behavior for a narrower subset of goods. Regarding misreporting, despite the frequent policy concerns, the evidence is limited, and such behavior occurs primarily in settings involving very closely related goods, often in the context of border tariff evasion, where classification margins are very detailed. The absence of detectable substitution at the HSN-4 level aligns with the tax-incidence literature, which typically finds substitution responses only among narrowly defined and highly substitutable products (e.g., specific hygiene products, soda varieties, or nicotine delivery devices), rather than across broader product categories. Another possibility is that many related products are complements rather than substitutes.

To test these hypothesis, we implement a machine-learning approach to classify close substitutes and complements for each treated product (see Appendix). The model more readily identifies untreated complements than untreated substitutes,

[17] The reporting rate declines slightly, but the estimate (–0.005) is economically small and statistically insignificant, again inconsistent with meaningful mislabeling or substitution.

suggesting that many "related" products may in fact be consumed jointly rather than as substitutes. Under the most restrictive matching criteria, we detect positive effects for complements, while estimates for substitutes are negative but statistically insignificant. These results indicate that cross-price responses exist only at a level of product similarity far narrower than that targeted by the reform, and are therefore unlikely to be first order in the context of India's multi-rate GST system.

4.5 Within Firm Spillovers

Our identification exploits within-firm variation in treatment status across products, absorbing firm–time shocks. A potential concern is that the rate cut may affect untreated products through firm-level spillovers, violating SUTVA. To further assess this issue, we exploit heterogeneity in firms' exposure to the reform. Prior to November 2017, firms differed in the share of their sales accounted for by treated products. If firm-level demand spillovers were quantitatively important, control products sold by firms with greater pre-reform exposure might exhibit larger post-reform increases in quantities and sales. The median control product is sold by a firm for which treated products account for 3.5% of pre-reform sales (mean: 11.7%). Conditional on the firm selling at least one treated product, the median share increases to 5.0% (mean: 13.1%). We estimate a regression restricted to control products, using the firm's baseline share of treated sales as the independent 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; 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.

Figure A.9 reports estimates for consumer prices, quantities, and sales of -0.022, -0.027, and -0.049 log points, respectively. Only the price coefficient is statistically significant at conventional levels. For interpretation, we rescale these coefficients by the average baseline share of treated sales among exposed firms (11.7%): this implies indirect effects of -0.0028, -0.0035, and -0.0064 log points. These are an order of magnitude smaller than the corresponding direct effects of the rate reduction (-0.027,

0.028, and 0.054). This implies that SUTVA violation due to within-firm spillovers could account at most for 10–13% of the estimated treatment effects.[18]

4.6 Robustness

Figure 8 summarizes key robustness checks (see Table A.13 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.[19]

The third 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 fourth 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 *product* level, which tax rates are defined (HSN-4) even when firms report at narrower levels; (iii) aggregating all treated and all control products within each firm; and (iv) aggregating sales across firms at the product level.

Across all four sets of robustness checks, the estimates remain close to those in

[18]The results on substitution and the robustness checks restricting the control group to unrelated products indicate that SUTVA is also not violated across products.

[19]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.

the main specification, reinforcing the robustness of these results. Moreover, the findings are not driven by a small number of outlier products: Figure A.10 displays *product*-level regression estimates for changes in consumer prices, quantities, and sales, which are centered around the average effects and close to normally distributed.

4.7 Second Reform: Rate reduction in July 2018

A few products that remained in the 28% tax bracket after the November 2017 reform experienced a rate cut to 18% in July 2018. This later reform mainly affected electrical appliances (e.g. Vacuum Cleaners) and house renovation items (e.g. paints).[20] While the reform covered fewer products and was more targeted, it complements the November 2017 analysis: it occurs well after the GST introduction and extends to higher-ticket products representing a non-trivial share of total sales (Table A.4).

We estimate the effects of the later reform using the same specification as for the earlier reform (equation 1).[21] The results are consistent with those from the earlier reform, despite differences in timing and product composition. Figure A.11 presents event-study estimates for consumer prices (Panel A), quantities (Panel B), and sales (Panel C). Consumer prices for treated products fall by 0.028 log points, implying a 55% pass-through (given the partial reduction in the effective tax rate of 0.052 log points). Quantities increase by 0.037 log points—slightly larger than the elasticity estimated earlier, and pre-tax sales rise by 0.062 log points. As with the earlier reform, we find no evidence of extensive-margin responses and can reject even small increases in reporting more treated products (Figure A.12).

These results reinforce the interpretation that the elasticities estimated around the November 2017 reform are not driven by the GST introduction five months earlier and extend to a broad set of products, although some series exhibit mild pre-trends.

[20] Table A.15 lists the main treated products in July 2018. Products can be divided in two main categories: electrical appliances, house renovation materials, and a third miscellaneous group. The first category includes *Heating And Thermic Appliances, Hair-Dressers*, *Refrigerators And Freezers*, *Laudry Washing Machines*, *Food Grinders And Mixers, Juice Extractors*, and *Vacuum Cleaners*. The second includes *Glaziers Putty, Calking, Surface Preparations* and *Paints And Vanishes*.

[21] The control group is constructed analogously, using the entire pre–July 2018 period to measure sales to final consumers. Table A.14 describes the treated and control groups. Treated products are larger in terms of sales and unit prices, consistent with their being more semi-durable items.

4.8 Firm Level Results and Heterogeneity by Size

Firm-product tax returns, used in the main analysis, are mandatory only for firms with annual sales above 15 million rupees (approximately USD 234,000). Firms filing product level returns—176,000 out of 981,000 registered firms (18%)—account for 74% of total sales (Table A.17).[22] We turn here to firm-level tax returns, to cover all registered firms, and to analyze value added, only reported at the firm level.

At the firm-level, we don't observe whether firms sell treated products in a given month, nor their reported quantities or prices. Instead, we proxy treatment status with the GST registration form which lists up to five goods firms intend to sell. The information is non-binding. We classify firms as treated if at least one listed good faced a rate cut. The firm-level regression is:

$$\log(y_{it}) = \sum_{t=\text{Jul17},\, t\neq\text{Oct17}}^{\text{Mar19}} \beta_t \cdot \text{treatment}_i \times \text{month}_t \; + \; \alpha_i \; + \; \gamma_t \; + \; \epsilon_{it}. \tag{3}$$

$\log(y_{it})$ is the logarithm of sales for firm i in month t, treatment_i is the firm's treatment status based on registration, and α_i and γ_t are firm and time fixed effects.

The specification excludes firms that list products treated in other periods, exempt, or services, to mirror the firm–product analysis. The resulting sample includes 281,000 firms, which report an average of 2.6 products at registration (Table A.18). 28% of firms are classified as treated, having listed at least one product that faced a rate cut in November 2017. Treated and control firms are similar across most baseline characteristics, but treated firms list more products (3.2 versus 2.4), which is expected given that treatment is defined as reporting at least one treated product.

Despite the coarse assignment, treated firms exhibit a clear post-reform decline in ETRs of 3 percentage points, as opposed to only a 0.8 percentage points reduction for firms assigned to the control (Figure A.13.A).[23] Thus the registration form is informative of tax variation. Treated firms raise reported sales by 0.064 log points relative to control firms (Figure A.13.B); when restricting to final consumers sales for

[22] Of the remaining 26% of sales, 20% correspond to firms filing detailed firm-level returns that distinguish sales to firms from sales to final consumers. Only 6% of total sales are reported exclusively in aggregated firm-level returns, although these firms file a third of the tax returns.

[23] In this context 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.

the sample that reports it, we estimate an increase of 0.054 log points (Figure A.13.C).[24]

How do the firm-level results compare to the firm–product estimates? Because treatment is defined differently and the unit of observation changes, the comparison is not direct. In the firm-level sample, treatment assignment is noisier and the implied reduction in ETR is smaller, since any firm listing at least one treated product is classified as treated. To clarify these differences, Table A.20 replicates the firm-level regressions separately for (i) returns with only firm-level information, (ii) returns with detailed firm-level information, and (iii) returns with firm–product information. These reporting types differ sharply in firm size (Table A.17). Firms filing the firm–product return are therefore most relevant for aggregate elasticities, while responses among smaller, marginally formal firms remain informative about heterogeneity.

Comparing, with the same methodology, total sales responses for the firms that also file firm–product returns allows a direct benchmarking across specifications. Two results emerge. First, firms filing only firm-level returns exhibit large behavioral responses, with elasticities of -3.6 (unweighted) and -5.1 (weighted). Second, firms filing firm–product returns display smaller elasticities (-2.2 unweighted; -1.6 weighted), close to the weighted firm–product estimate (-1.5). Firms filing the detailed firm-level return but not the firm–product return fall between these two groups, indicating that differences reflect firm composition rather than specification.

4.9 Heterogeneity by Proxy for Informal Sector Exposure

In economies with large informal sectors, indirect taxes may be more distortionary. Consumers can substitute from formal to informal varieties of goods, and registered firms facing informal competition may absorb a larger share of the tax or become more prone to evasion (Di Marzio et al. 2025). While all firms in our tax data are formal by construction, exposure to informal competition varies across products and firms. We proxy this exposure using measures of firm size and product complexity that have been shown to predict informality in other contexts (Kleven et al. 2016).[25]

At the firm level, we examine heterogeneity based on whether a firm is located in the capital, Bangalore, and whether its total sales are above the sample median

[24] Table A.19 presents robustness checks varying control group and balance, yielding similar results.

[25] The literature shows that larger firms and more complex products typically involve longer value chains, more specialized inputs, and reliance on external financing which generate third-party information trails and make fully informal production and trade more difficult.

at baseline. At the product level, we proxy exposure to informality using product complexity (Atlas of Economic Complexity, Hausmann et al. 2013). We also classify products by median sales in our data, splitting firm–HSN pairs according to median pre-reform sales; products sold by larger median firms are more likely to be supplied predominantly by formal firms. Finally, we map products to sectors to exploit representative data from the 2014–15 Annual Survey of Industries (ASI), which covers both registered and unregistered firms. We divide sectors by median employment size, under the same rationale that sectors characterized by larger firms are more formal.

Table A.16 reports the price and quantity responses across the five heterogeneity dimensions. The drop in effective tax rates and in consumer prices is very similar across splits. The results on quantity and sales responses are more mixed. Splits based on product complexity and sector-level employment—proxies for lower exposure to informal competition—yield slightly larger responses among products that are more likely to be supplied formally. By contrast, splits based on median product sales and firm location display the opposite pattern. Across all specifications, estimated quantity and sales elasticities remain moderate. Overall, even in environments with substantial informality, formal firms exhibit relatively low elasticities, consistent with a degree of segmentation between formal and informal markets.

This conclusion should be qualified by the previous firm-level evidence showing that firms filing only the basic firm-level return respond more strongly than those filing the product-level return (Table A.20). These smaller firms are more likely to have unreported sales, suggesting that informality operates mainly through the behavior of marginally formal firms rather than through competitive pressure on large, fully registered firms. Since sales are concentrated among large firms, economy-wide estimates are primarily shaped by their behavior, attenuating the role of informality.

4.10 Heterogeneity by Products' Engel Curve Slopes

An important dimension of the equity–efficiency trade-off in indirect taxation is product-level income elasticity. Governments may want to levy higher taxes on goods disproportionately consumed by richer households, unless these goods exhibit higher demand elasticities and thus larger efficiency costs. To assess this trade-off, we first estimate the slope of the Engel curve for each survey product. Next, we match

HSN codes to survey product categories.[26] Figure A.15 shows the estimated Engel elasticities by survey product, highlighting treated items. Treated products span most of the elasticity distribution. On average, treated products have an Engel elasticity of 1.8, although the 25% percentile is below one, while the 75% is above three. This dispersion indicates that the reform did not target only highly luxury goods.

Figure A.14 plots the relationship between our estimated elasticities with respect to the ETR and each product's Engel-curve slope. 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.

To quantitatively asses this point, we regress survey-product elasticities on the corresponding Engel curve coefficients, estimated as:

$$\hat{\beta}_j^{\text{elasticity}} = \alpha + \delta \hat{\beta}_j^{\text{Engel}} + \varepsilon_j \tag{4}$$

Table A.23 reports the results weighted by survey consumption.[27] 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 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%.

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 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

[26] This correspondence is imperfect: HSN codes tend to be more detailed than survey categories.

[27] We also report results under alternative weighting schemes: by equally weighting products, by consumption shares in the survey, by the number of matched HSN-level observations or by the share of pre-tax sales in the tax data.

directly overcomes these limitations: it provides a long pre-treatment period;[28] 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:

$$y_{bdt} = \sum_{t=\text{Aug16}, t\neq\text{Nov17}}^{\text{Mar19}} \beta_t \cdot (\text{treat}_b \times \text{month}_t) + \alpha_b + \gamma_{td} + \epsilon_{bdt} \tag{5}$$

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.21 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.[29]

Figure 9.A shows the event-study estimates from equation 5 for consumer prices.[30]

[28] 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.

[29] Table A.22 reports the number of barcodes at the hypermarket's *Family* level. The leading *Families* mirror the HSN-4 classification (Table A.7). For instance, the *Personal Care & Cosmetics* department contains the *Hair Care* and *Skin Care* families, corresponding to HSN-4 *Hair Preparations* and *Make-Up*, two of the key treated *products* in the tax data.

[30] 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

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 heterogeneous patterns observed there, where larger firms exhibit lower pass-through (Table A.16).

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 administrative 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 the size and formality of the hypermarket, quantity responses likely reflect genuine increases in sales rather than newly reported informal activity. This quantity response

tax rates. This limitation does not affect quantities or consumer prices, which are invariant to tax reporting and include all tax components.

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

Estimated Elasticities In the preceding analysis we estimated elasticities of quantities and sales with respect the tax rate. First, note that these elasticities are uncompensated and therefore incorporate income effects. In our setting, these are likely close to compensated elasticities as the size of the income shock is modest. Goods subject to the tax cut account for 14.7% of total sales in the final sample of the tax data and 8.8% of total consumption in the survey data (combining the November 2017 and July 2018 tax cuts). Even a large 10-percentage-point rate cut only implies a rise in real income between 0.9% and 1.5%. We will initially not make a distinction between compensated and uncompensated elasticities.

Second, the actual reduction in ETRs was only two-thirds of the statutory rate cut, the policy parameter directly set by the central government. As a result, the elasticity with respect to the STR is two-thirds of the elasticity with respect to the ETR. We use elasticities with respect to the ETR in what follows, both because they generalize better to other settings (with close STR-ETR alignment) and because they yield larger elasticity estimates, implying more conservative optimal rate differentiation.

Tax Revenues The tax revenue impact of a rate change directly depends on the sales elasticity. The change in revenue can be written as:

$$dR = (1 + t \cdot \varepsilon_s)\, dM.$$

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 with respect to the ETR, $\varepsilon_s = -1.12$, and $t = 0.18$, the revenue change is $-(1 - 1.12 \times 0.18) \approx -0.80$. Thus, one dollar of mechanical revenue loss translates into \$0.80 of actual loss, implying a behavioral offset of \$0.20. At the upper bound of the 95% confidence interval, $\varepsilon_s = -1.72$, behavioral responses offset \$0.31 per mechanical dollar. Using instead the STR elasticity, which is one-third smaller ($\varepsilon_s^{STR} \approx -0.73$), the revenue loss is $-(1 - 0.73 \times 0.18) \approx -0.87$. 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 (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:

$$dW = -\,t \cdot \varepsilon_q \cdot dM.$$

The Marginal Cost of Public Funds (MCPF), defined as the total private cost per dollar of revenue raised, is:

$$\text{MCPF} = 1 + \frac{-\,dW}{dR} = 1 + \frac{t \cdot \varepsilon_q}{1 + t \cdot \varepsilon_s}$$

With $\varepsilon_q = 0.61$ and $t = 0.18$, the welfare gain per dollar of revenue forgone with the rate reduction was about $\frac{dW}{dM} \approx 0.14$, implying an MCPF of 1.14. This means that raising one dollar of revenue reduces private welfare by about \$1.14. Even considering the elasticity with respect to ETR and the upper bound of the 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.16). 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.17).

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.18). The reduction amounts to 0.7 percentage 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.19 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.20). 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.21).

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.[31]

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

$$\frac{\tau_j}{1+\tau_j} = \frac{1 - \int_i \frac{g^i}{\mu} a_j^i \varepsilon_j^p \, di}{\varepsilon_j^s} \tag{6}$$

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 elasticity of pre-tax sales with respect to the tax (i.e., the tax base elasticity).[32]

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

[31] We extend the derivation to include the utility of firm owners in the Appendix B.
[32] The derivation is in the Appendix B.

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.[33]

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 (Section 4.10), we adopt a conservative approach and use the point estimate from the unweighted specification (-0.120), implying that goods with steeper Engel curves are more elastic in absolute value.[34] 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%).

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

[33] 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.

[34] 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 (4).

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 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,[35] 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).

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

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

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

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.

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

Di Marzio, I., Mocetti, S., & Rubolino, E. (2025). *Market externalities of tax evasion* (CESifo Working Paper No. 11896). CESifo. Munich.

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), None. https://doi.org/10.1016/j.jinteco.2023.103793

Gupta, M., & Rajaraman, I. (2020). Is the 14% revenue guarantee to states justified? *Economic and Political Weekly*, *55*, 18–21.

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.

Hausmann, R., Hidalgo, C. A., Bustos, S., Coscia, M., Simoes, A., & Yildirim, M. A. (2013). *The atlas of economic complexity: Mapping paths to prosperity* (2nd). MIT Press. https://doi.org/10.7551/mitpress/9647.001.0001

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

Kleven, H. J., Kreiner, C. T., & Saez, E. (2016). Why Can Modern Governments Tax so Much? An Agency Model of Firms as Fiscal Intermediaries. *Economica*, (15218), 51. https://doi.org/10.3386/w15218

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. https://doi.org/None

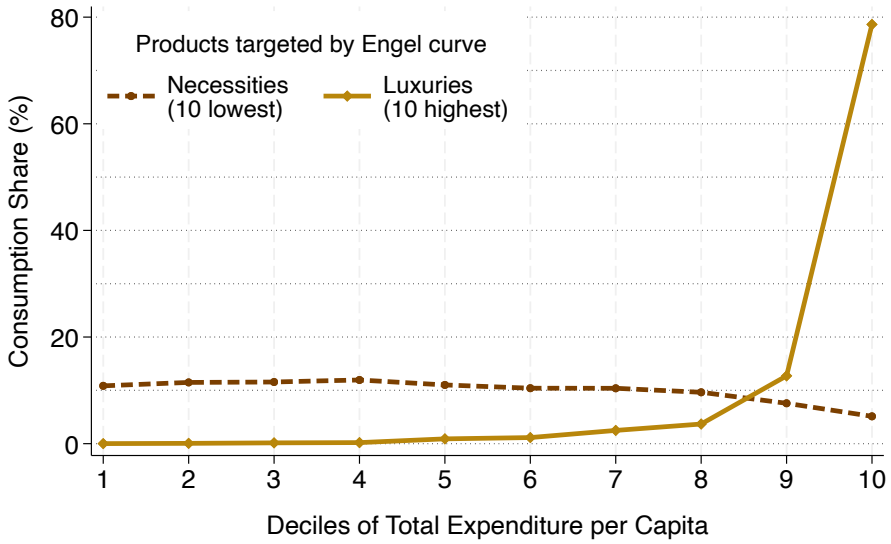
State of State Finances. (2024). *State of State Finances* (tech. rep.) [FY 2024–25]. PRS Legislative Research.

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

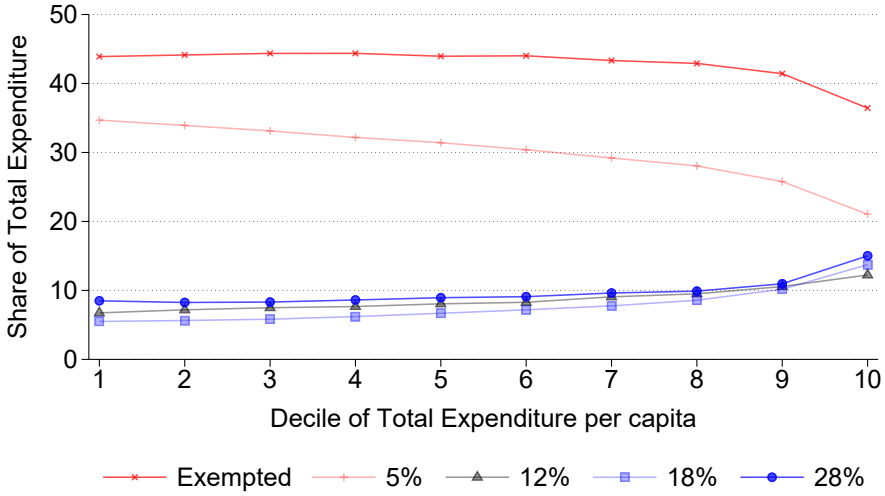
Figures

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.

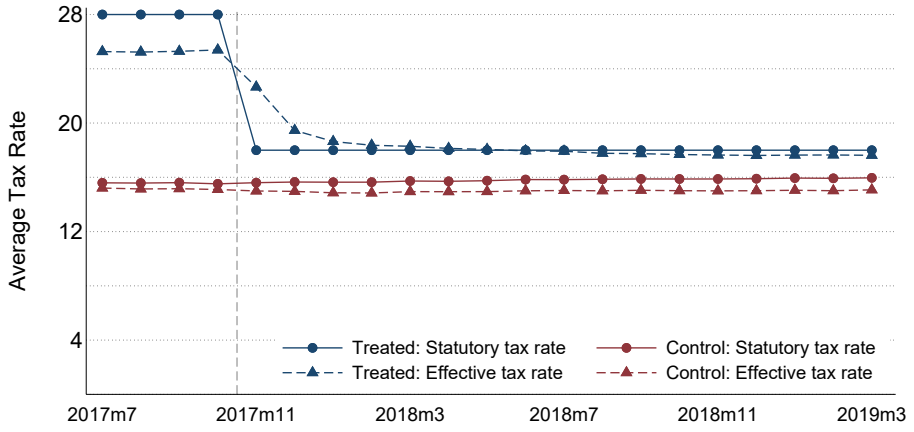
Figure 2: Distributional Expenditure Shares by Baseline GST rate



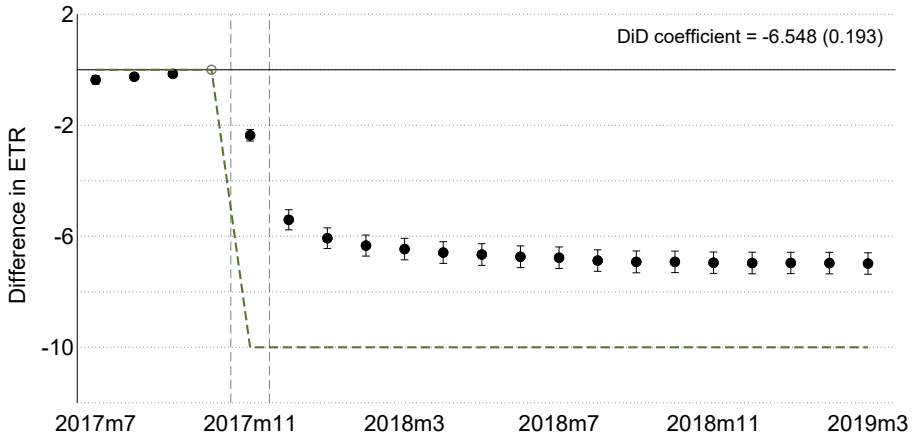
Note: This figure presents the share of total consumption by product GST rate in July 2017 across deciles of household total expenditure per capita, based on the 2011-12 household expenditure survey. Products received as gifts or social transfers and jewelry are excluded.

Figure 3: Average Statutory and Effective Tax Rates

(a) Average Level



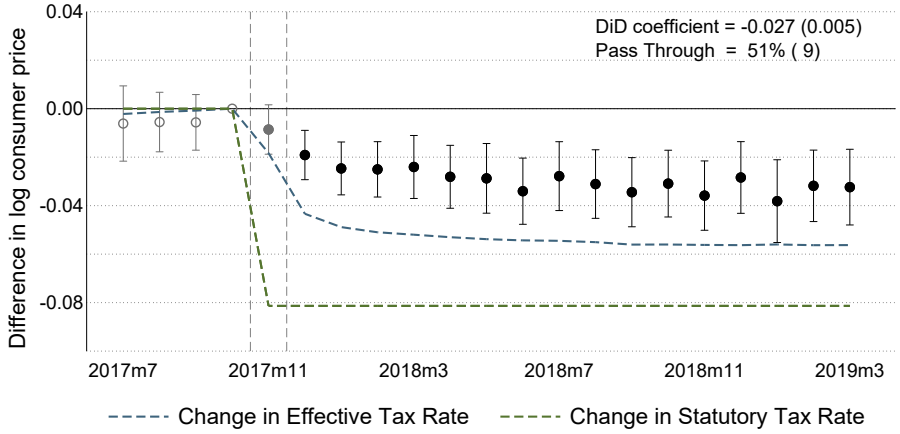
(b) Event-study



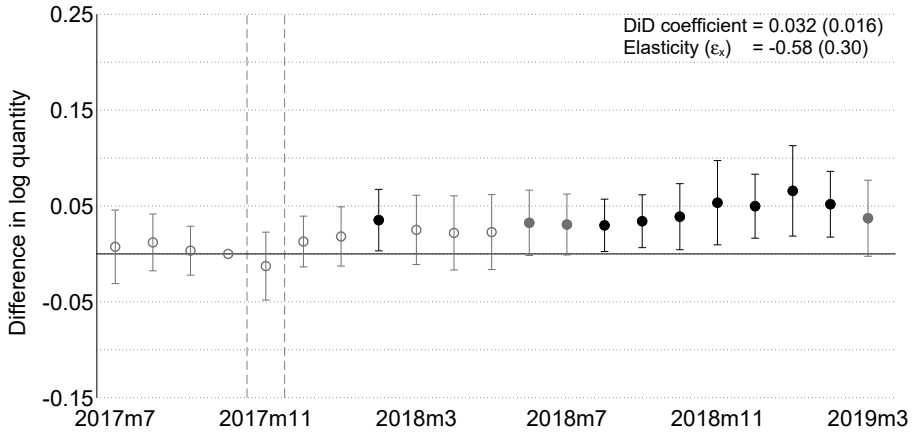
Note: Panel A presents the average Statutory Tax Rate (STR) and the average Effective Tax Rate (ETR) by treatment group. Panel B presents the event-study coefficients for changes in ETRs. The dashed green line represents the change in STR. Treated products are those that experienced a reduction in STR 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



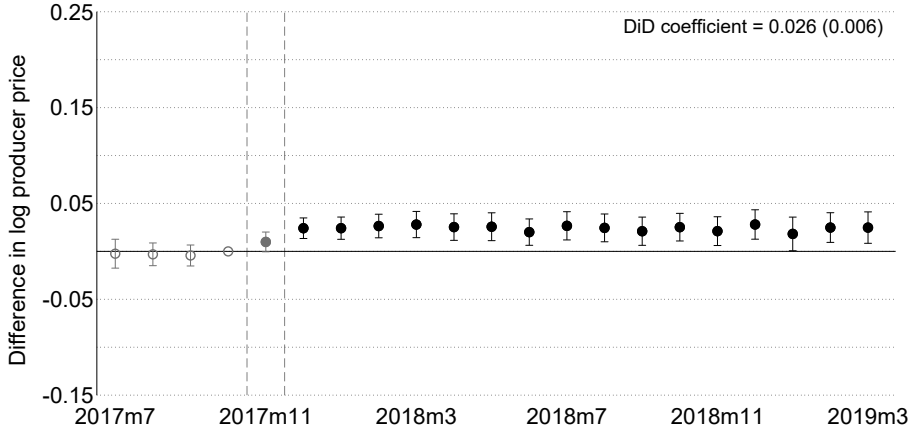
(b) Quantity



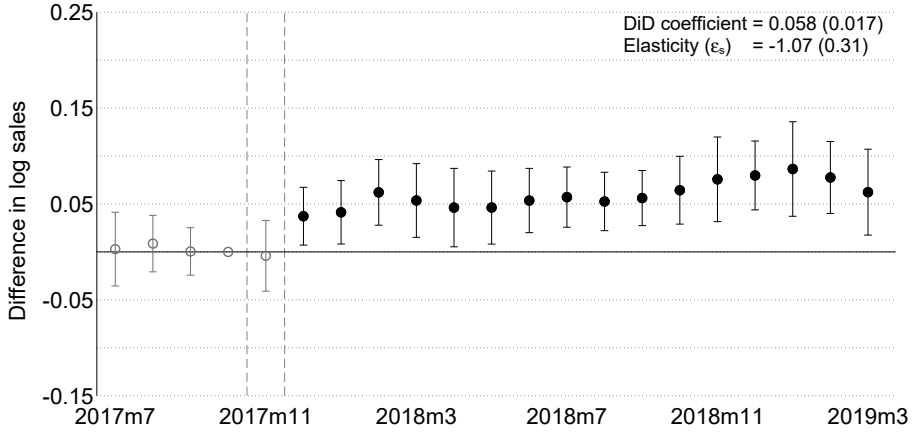
Note: Panel A presents the event study coefficients for log consumer price. The green and the blue dashed lines indicate respectively the difference in log STR and log Effective Tax Rates (ETR) between the two groups. Pass-Thought is defined as the ration between the reduction in consumer price and the 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)



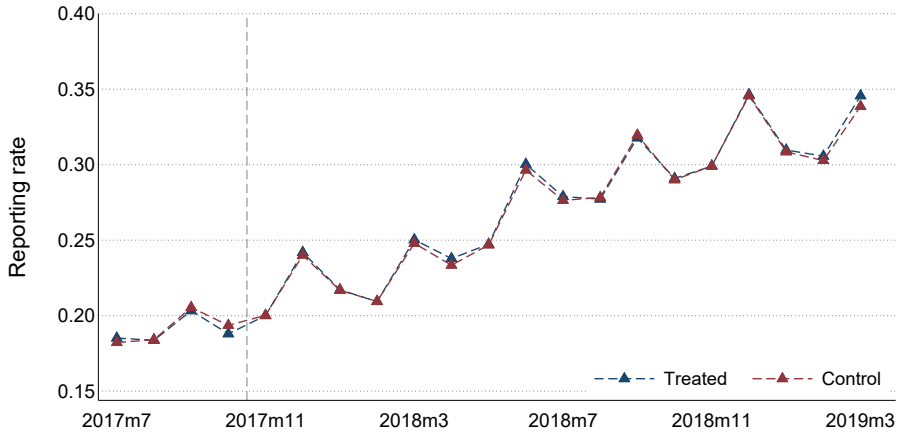
(b) Sales Value



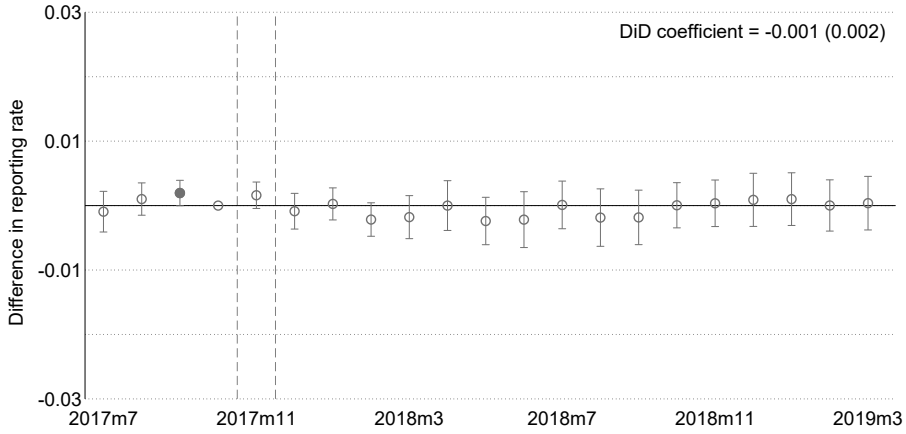
Note: This figure presents the event study coefficients for log pre-tax prices and log sales. Sales amounts are pre-tax. Treated products are those that experienced a reduction in Statutory Taxes Rates (STRs) from 28% to 18%. Control products are those taxed at 12% or 18% that did not experience a rate 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

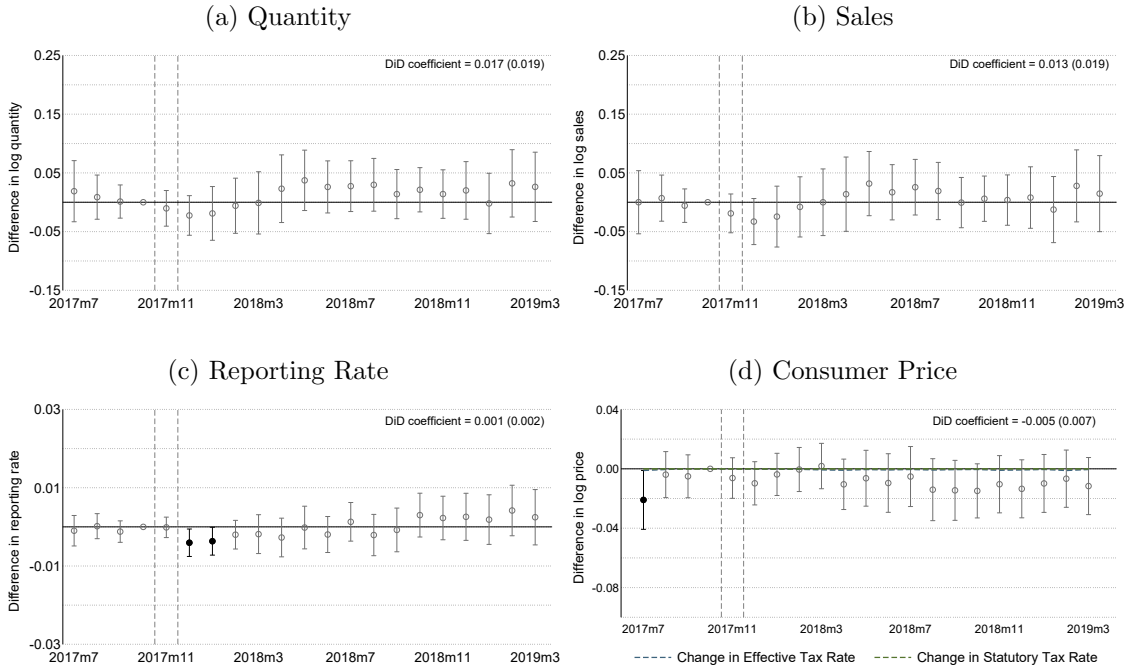


(b) Probability of reporting a product: Event-study coefficients



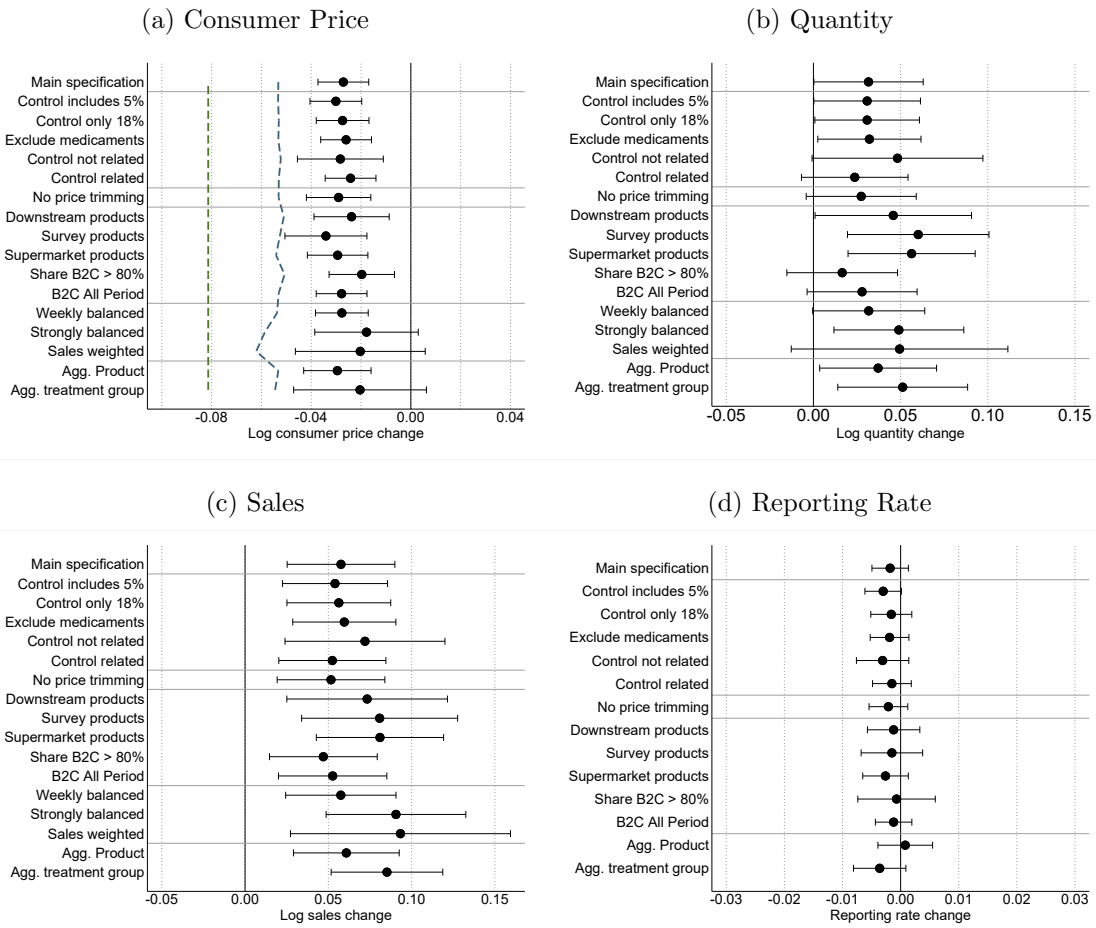
Note: This figure presents the results for the reporting rate. The reporting rate is defined as the probability of reporting selling a product in a month conditional on having sold it any month in the period analyzed. Treated products are those that experienced a reduction in Statutory Taxes Rates (STRs) 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.

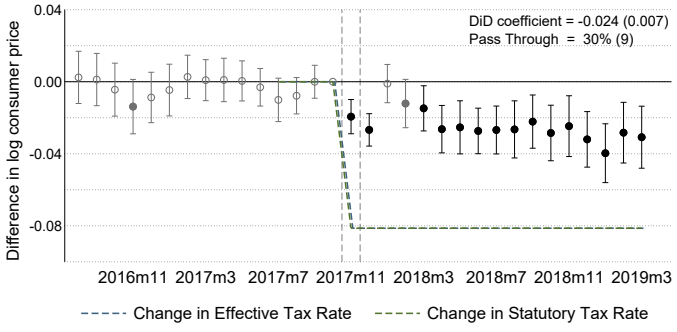
Figure 8: Alternative specifications



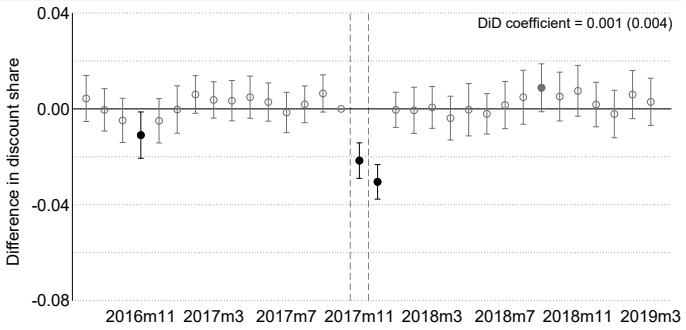
Note: The figure reports the differences-in-differences coefficients of a series of robustness tests. Each coefficient reports the point estimates from a separate robustness specification. Estimates come from difference-in-differences regressions comparing treated and control products across three periods: 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

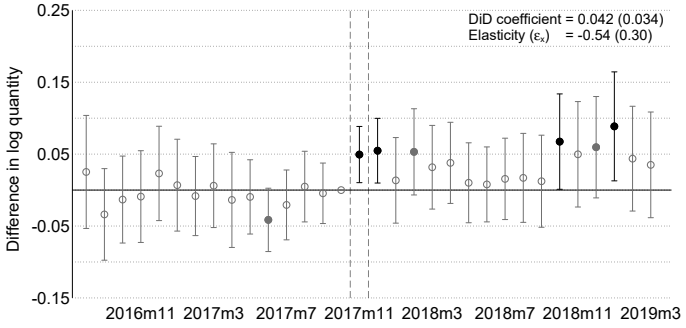
(a) Consumer Price (Inclusive of discounts)



(b) Discounts

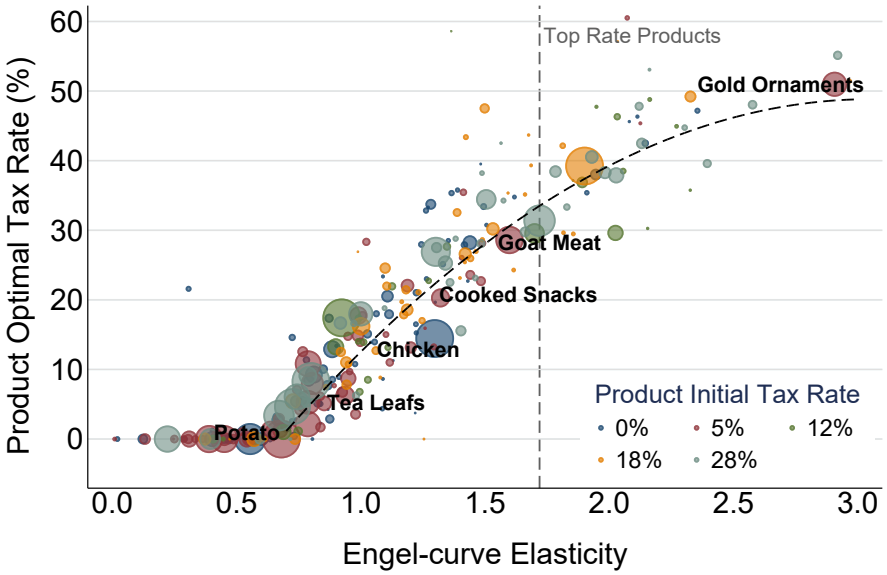


(c) Quantity



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

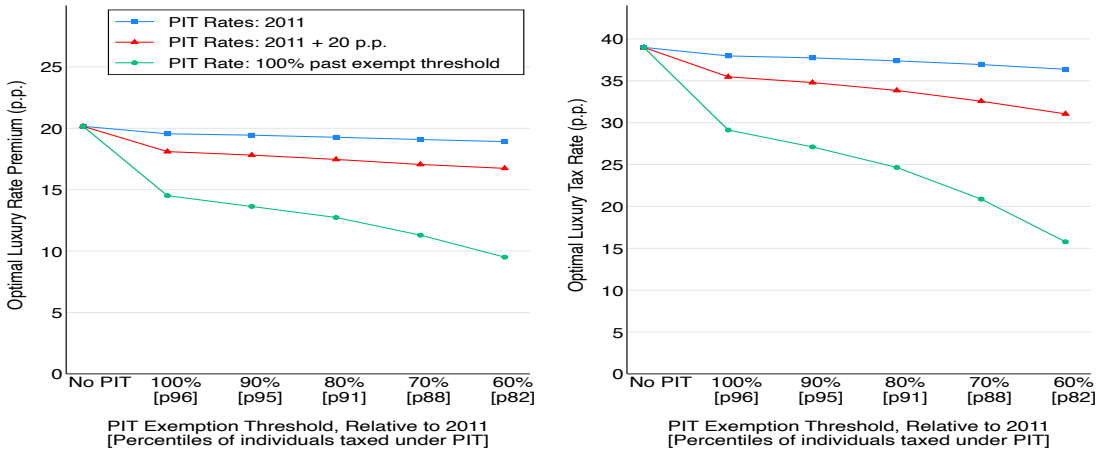
Figure 10: Optimal Tax Rate for Each Survey Product



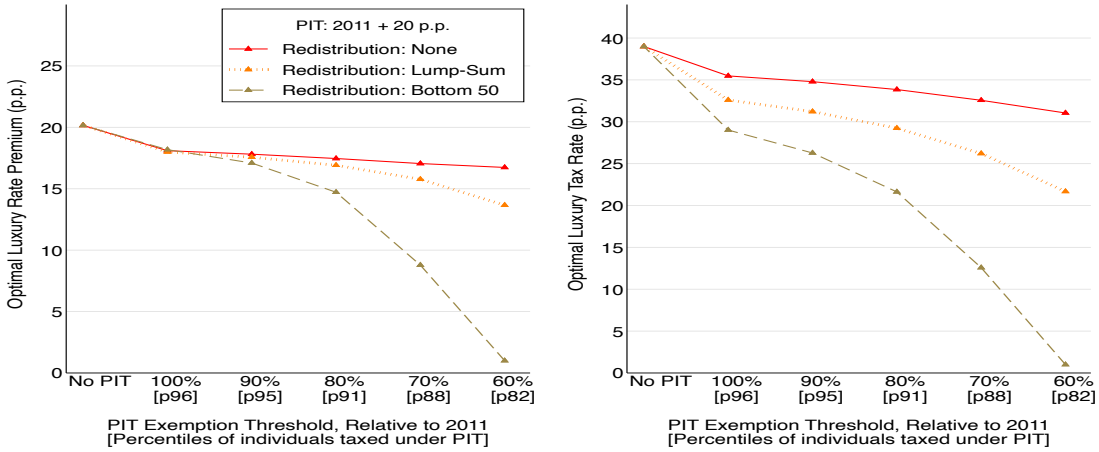
Note: This figure presents our estimates of optimal tax rates by survey product as a function of the Engel-curve elasticity parameter. Optimal tax rates below zero (i.e., subsidies) are winsorized at zero. Each circle represents one survey product, and circle sizes reflect total consumption. Product 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	165	681	846
Firm-HSNs	144,304	347,172	491,476
Of which before Nov 17	56,962	138,596	195,558
Before November 2017:			
Months reported	2.06 (1.17) [2.00]	2.05 (1.16) [2.00]	2.05 (1.16) [2.00]
HSN level reported	6.18 (1.98) [8.00]	6.18 (1.98) [8.00]	6.18 (1.98) [8.00]
Final consumers sales (Firm share)	0.81 (0.16) [0.84]	0.80 (0.16) [0.83]	0.81 (0.16) [0.83]
Sales ('000 rupees)	218 (2338) [7.22]	485 (17718) [6.60]	407 (14955) [6.78]
Sales (Log value)	8.98 (2.44) [8.88]	8.93 (2.60) [8.79]	8.95 (2.56) [8.82]
Quantity (Log units)	3.87 (2.43) [3.53]	4.19 (2.57) [3.87]	4.10 (2.54) [3.76]
Price pre-tax (Log value per unit)	5.27 (1.96) [5.01]	4.91 (2.03) [4.66]	5.01 (2.02) [4.78]
Observations	801,090	1,927,118	2,728,208
Of which before Nov 17	117,579	284,136	401,715

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.

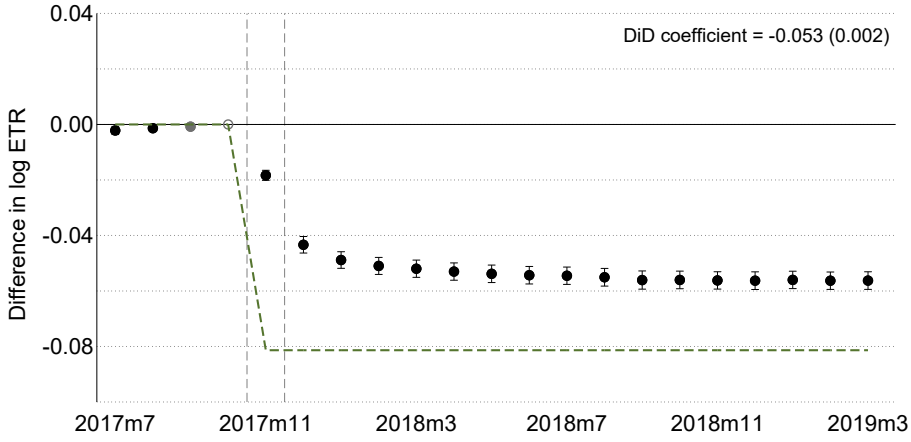
Table 2: Description of firms in the tax data

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.

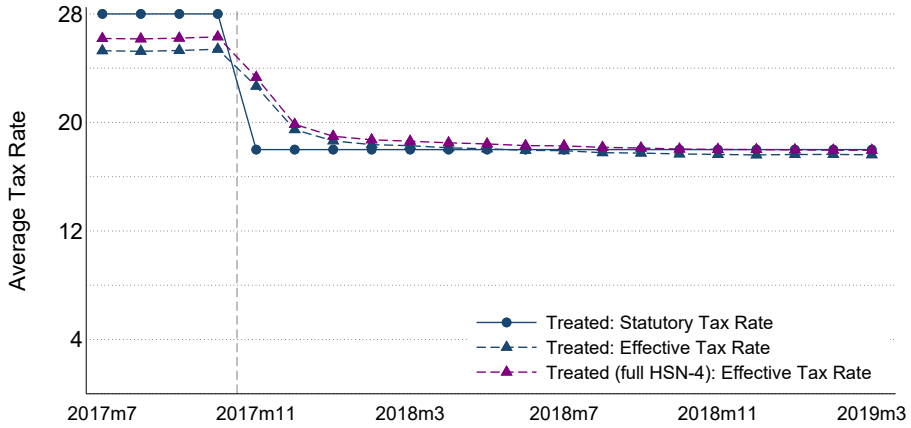
Appendix A: Additional Figure and Tables

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



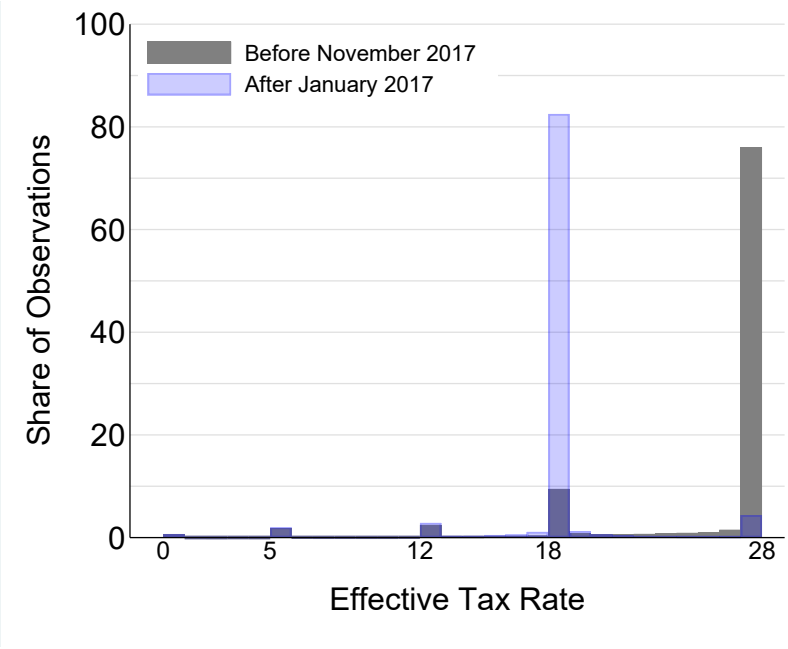
Note: The figure presents the event-study coefficients for changes in log ETRs. The dashed green line represents the change in log STR. Treated products are those that experienced a reduction in STR from 28% to 18%. Control products are those taxed at 12% or 18% that did not experience a rate change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

Figure A.2: Average Effective Tax Rates of fully treated HSN-4 level product



Note: This figure presents the average Statutory Tax Rate (STR) and the average Effective Tax Rate (ETR) for treated products. Treated products are those that experienced a reduction in STR from 28% to 18% in November 2017. Full HSN-4 treated products restrict to codes in which the entire HSN-4 level underwent the STR reduction from 28% to 18% in November 2017. ETRs are winsorized at the 1^{st} and 99^{th} percentiles.

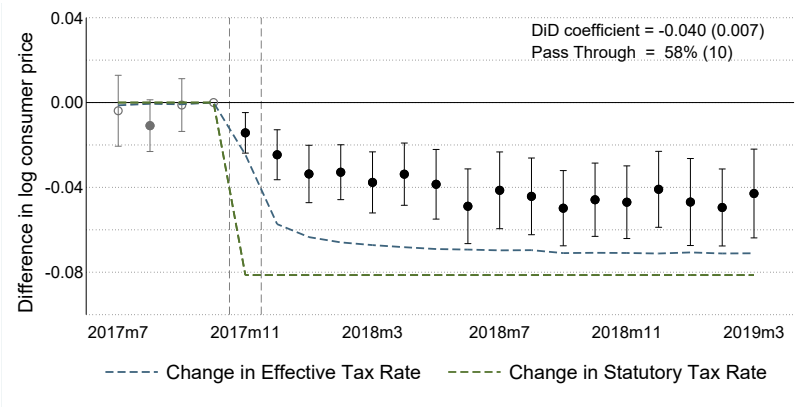
Figure A.3: Histogram of treated products Effective Tax Rates before and after rate reduction



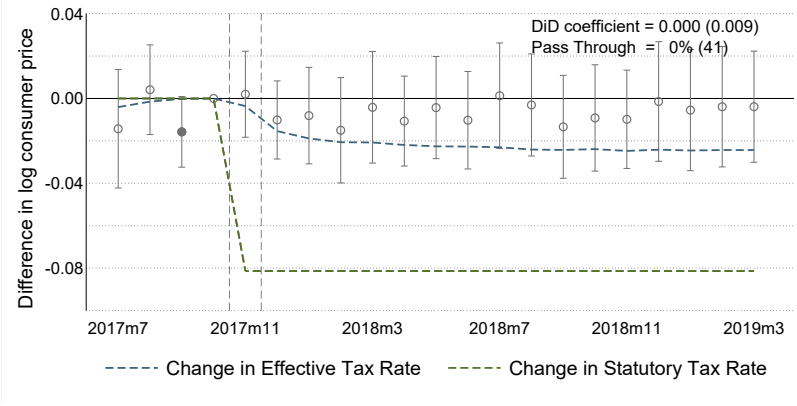
Note: This figure presents the distribution of Effective Tax Rates (ETR) before and after November 2017 rate reduction. Treated products are those that experienced a reduction in STR from 28% to 18% in November 2017. ETRs are winsorized at the 1st and 99th percentiles and rounded, such that the 18% bin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

Figure A.4: Consumer Price and Pass-Thought by STR-ETR match status

(a) ETR matches STR



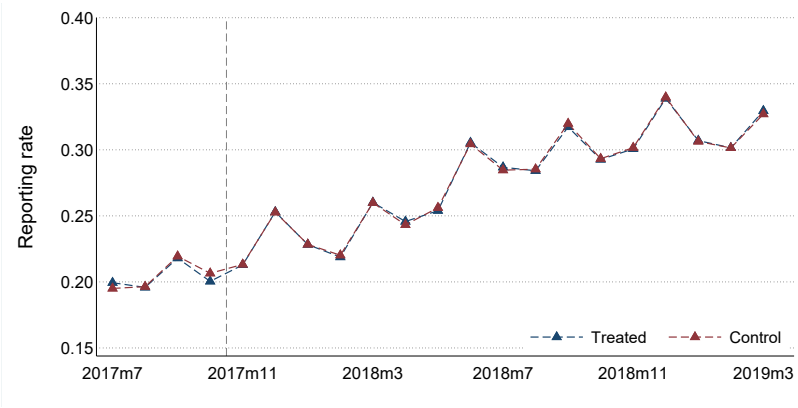
(b) ETR does not match STR



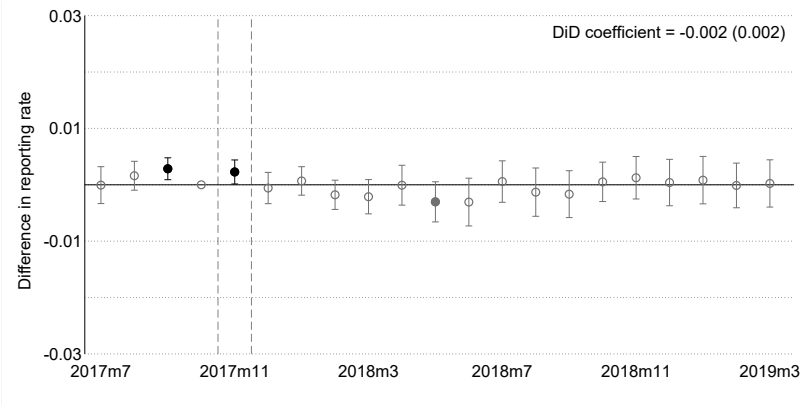
Note: This figure presents the event study coefficients for the difference in log consumer price. Treated products are those that experienced a reduction in Statutory Taxes Rates (STRs) from 28% to 18%. Control products are those taxed at 12% or 18% that did not experience a rate change during the period. 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.5: Reporting Rate

(a) Level

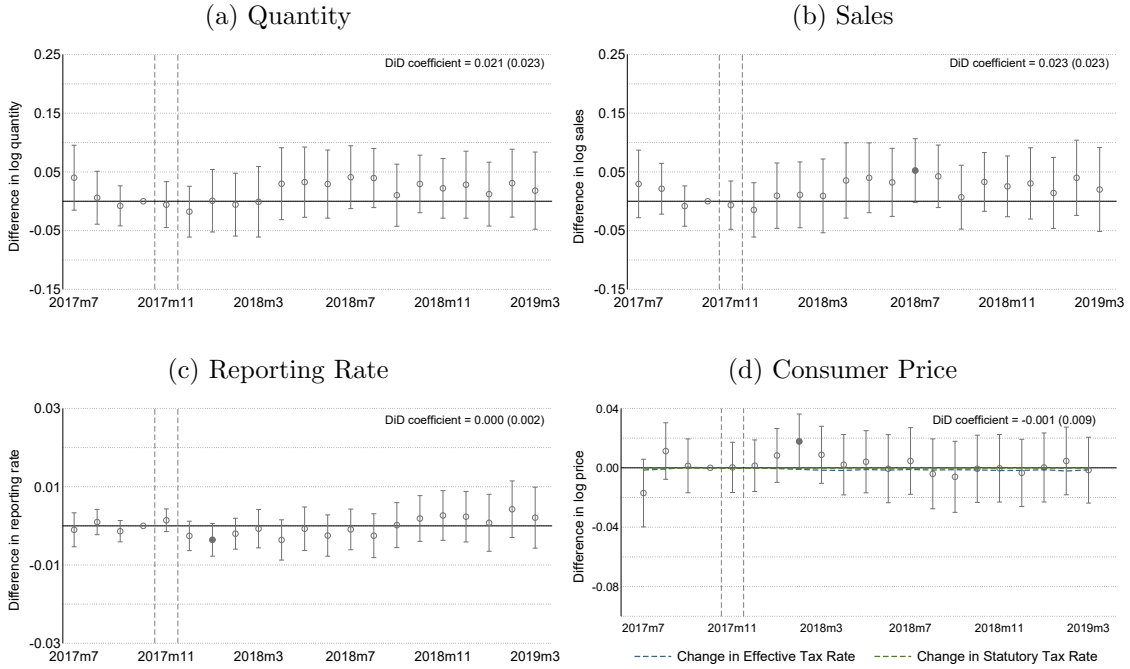


(b) Event-study



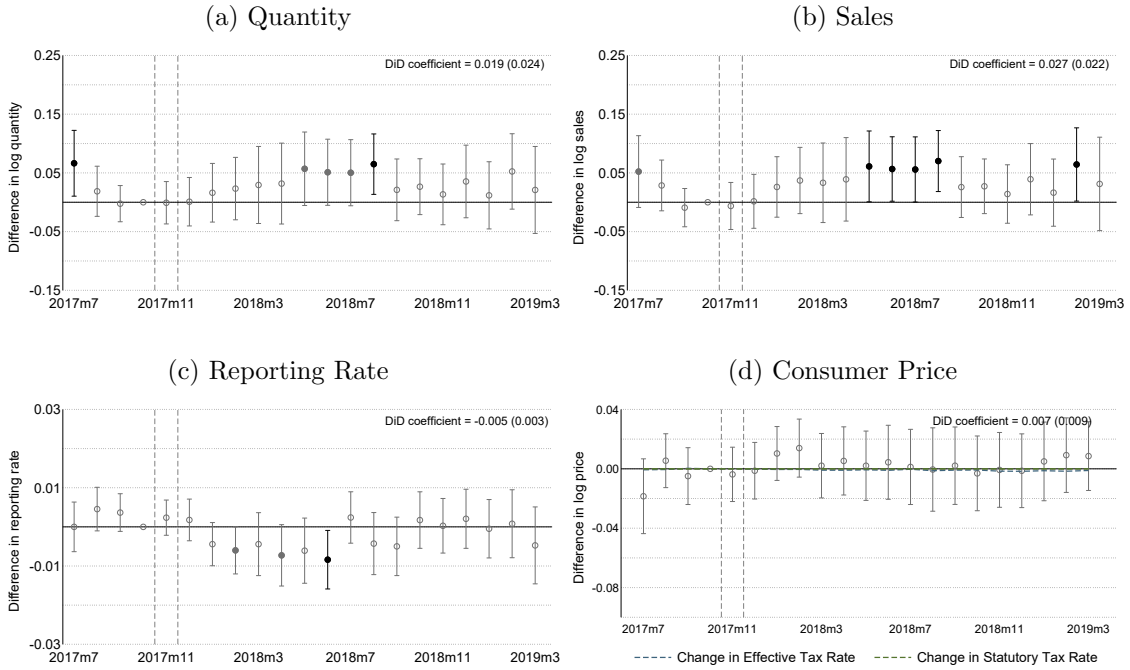
Note: This figure presents the results for the reporting rate. The reporting rate is defined as the probability of reporting selling a product in a month conditional on having sold it any month in the period analyzed. The sample definition uses the average share of final consumer sales before November 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.6: Related products affected by higher relevant treated products



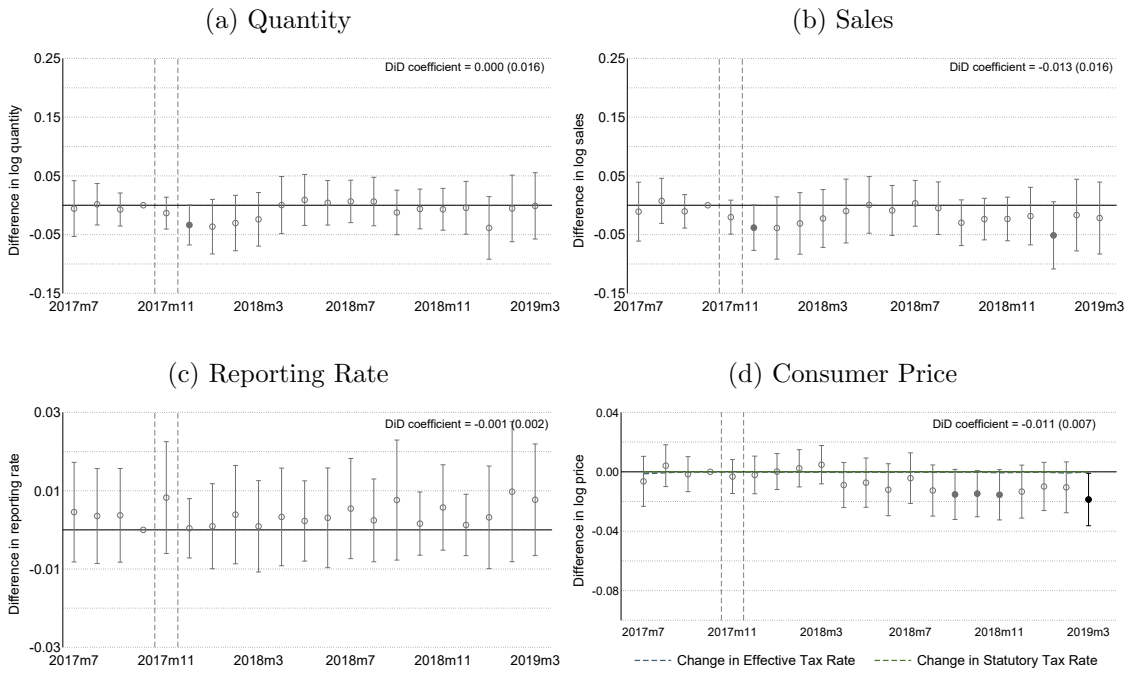
Note: This figure presents the results for log quantity, log sales, log price, and the reporting rate, comparing related and non-related control products. Control products are those taxed at 12% or 18% that did not experience a rate change during the period. Related products are control products in the same HSN-2 category as a treated product that accounted for more than 13.3% of the HSN-2 pre–November 2017 sales (median). 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 A.7: Related Products affected within firm



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 sold by the same firm before November 2017. 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.

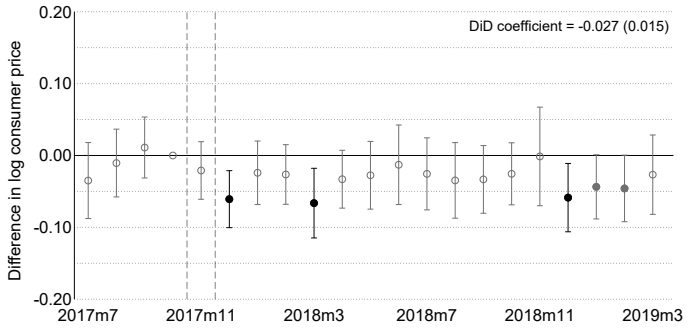
Figure A.8: Results for Related but Untreated Products, Including Taxed at 5%



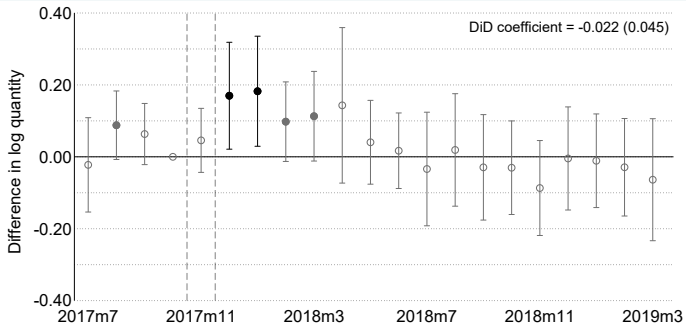
Note: This figure presents the results for log quantity, log sales, log price, and the reporting rate, comparing related and non-related control products. Control products are those taxed at 5%, 12% or 18% that did not experience a rate change during the period. Related products are control products 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 A.9: Within-Firm Spillover

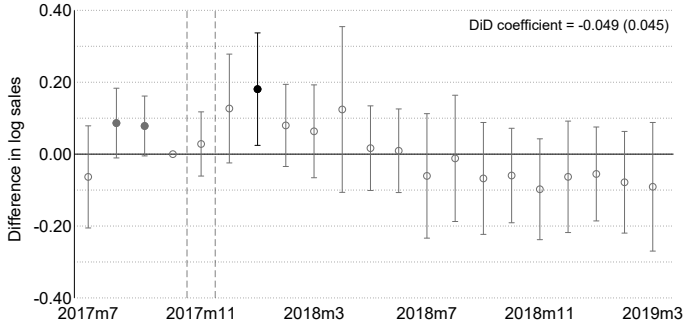
(a) Consumer Price



(b) Quantity

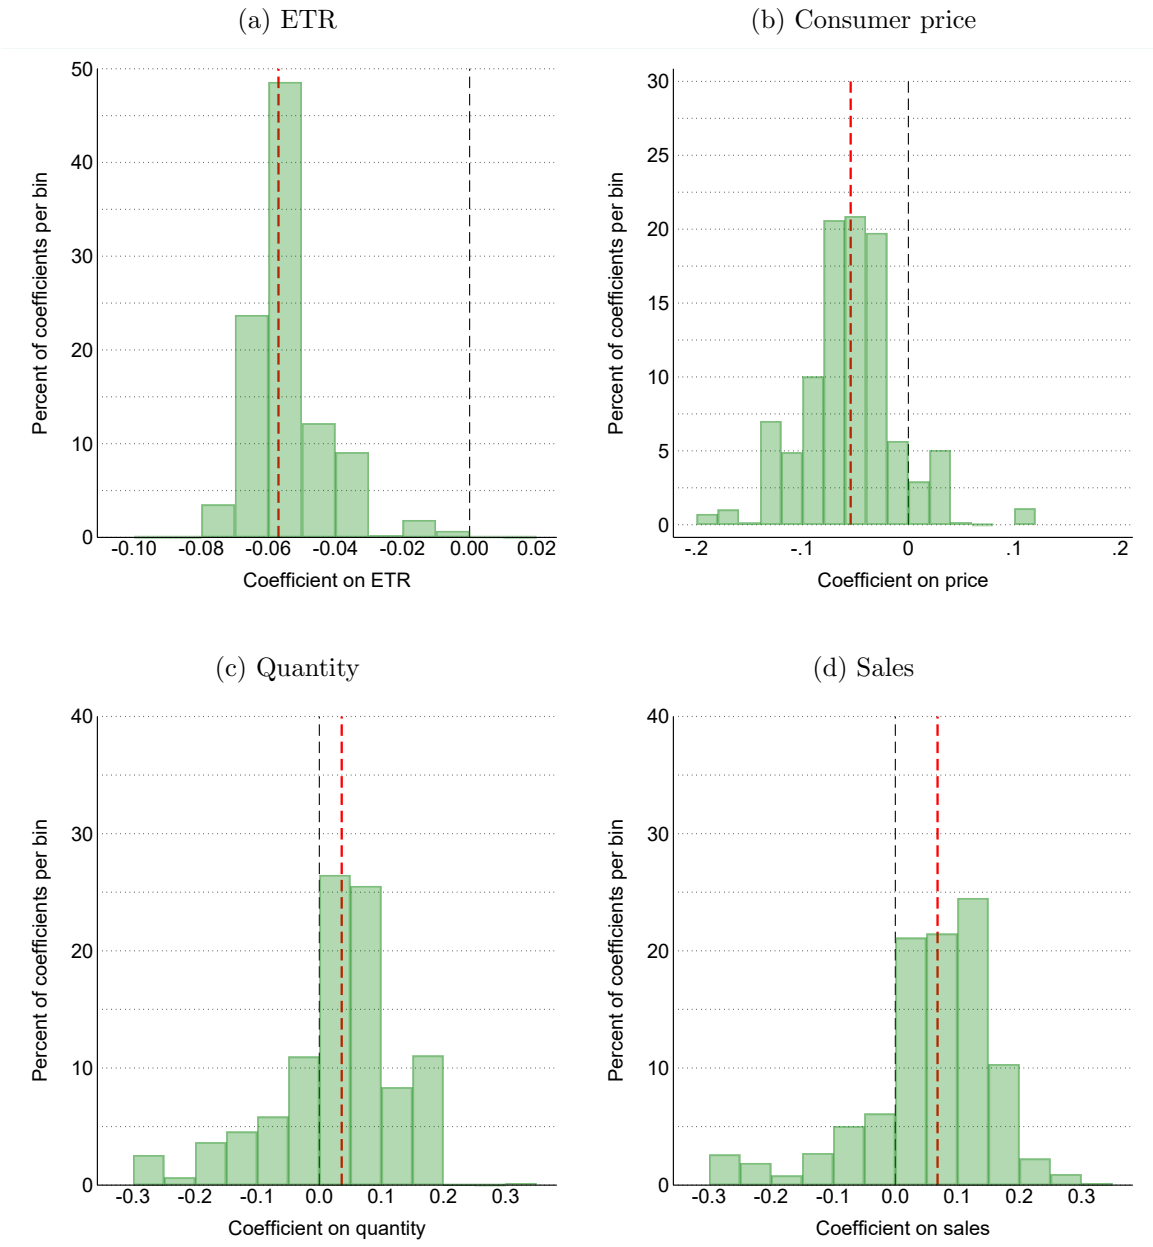


(c) Sales



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.

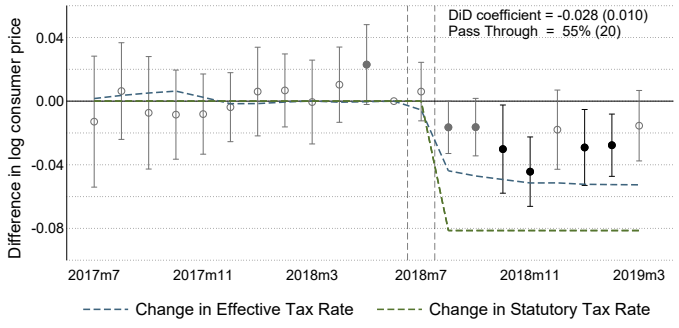
Figure A.10: Distribution of product coefficients



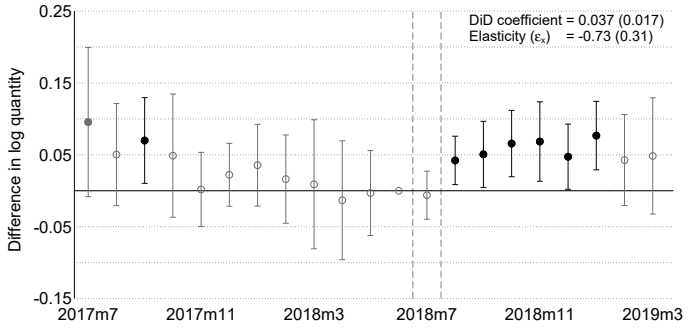
Note: This figure displays coefficients of the difference-in-differences for each treated product for log sales, quantity, effective tax rates and consumer prices. Coefficients are weighted by number of observations before November 2017 and winterized on each histogram maximum and minimum values. The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

Figure A.11: Treated in July 2018

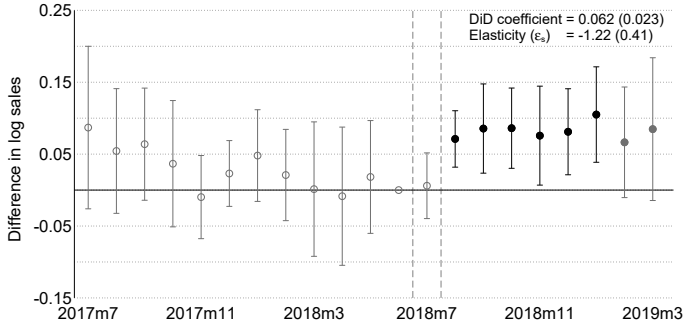
(a) Consumer Price



(b) Quantity



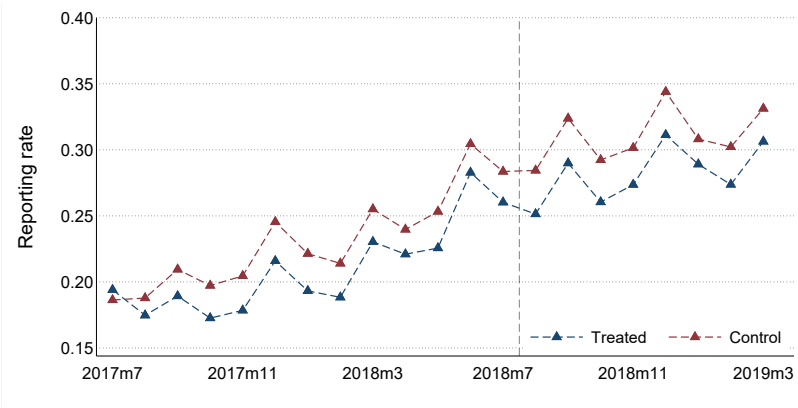
(c) Sales



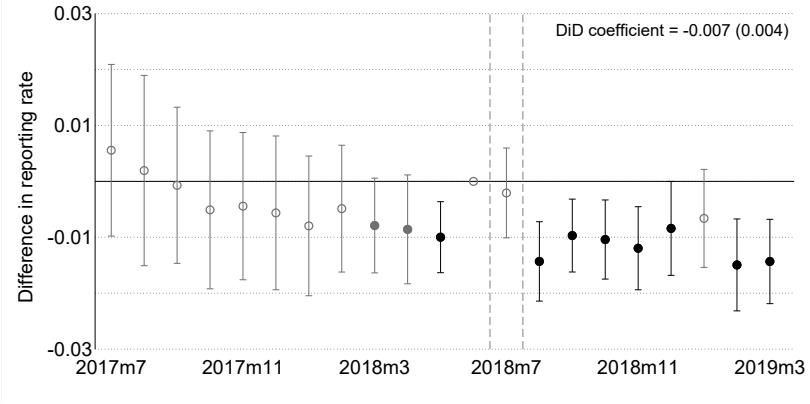
Note: This figure presents the results for log quantity, log sales and log price. reated products are those that experienced a reduction in Statutory Taxes Rates (STRs) from 28% to 18% in July 2018. Control products are those taxed at 12% or 18% that did not experience a rate change between July 2017 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.12: Reporting Rate Treated in July 2018

(a) Level



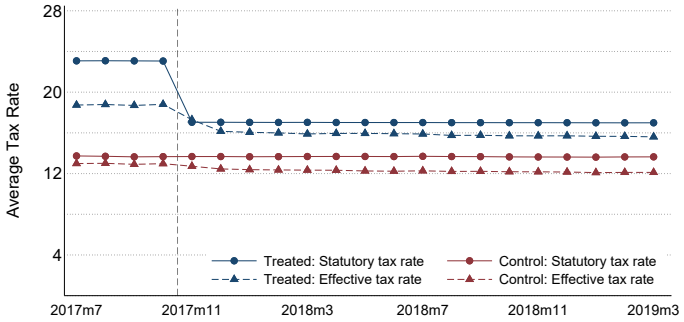
(b) Event-study



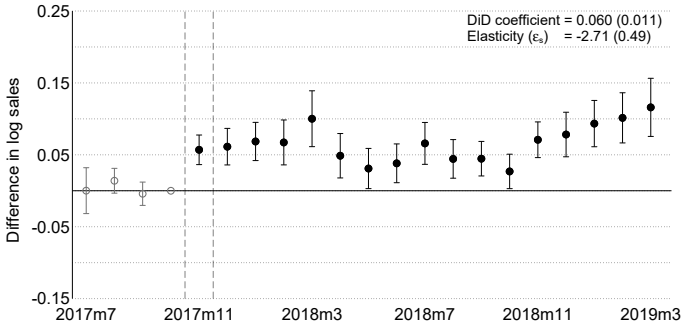
Note: This figure presents the results for the reporting rate. The reporting rate is defined as the probability of reporting selling a product in a month conditional on having sold it any month in the period analyzed. Treated products are those that experienced a reduction in Statutory Taxes Rates (STRs)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

Figure A.13: Firm tax returns responses

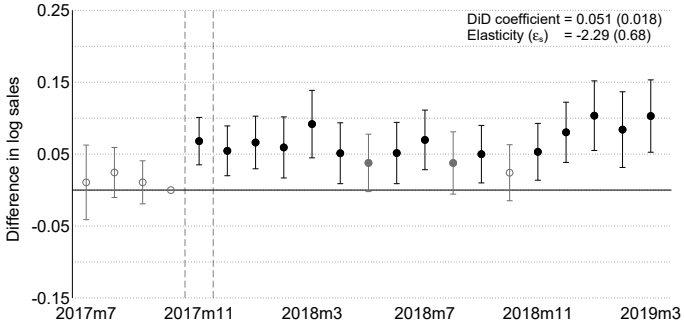
(a) Tax Rates



(b) Total sales



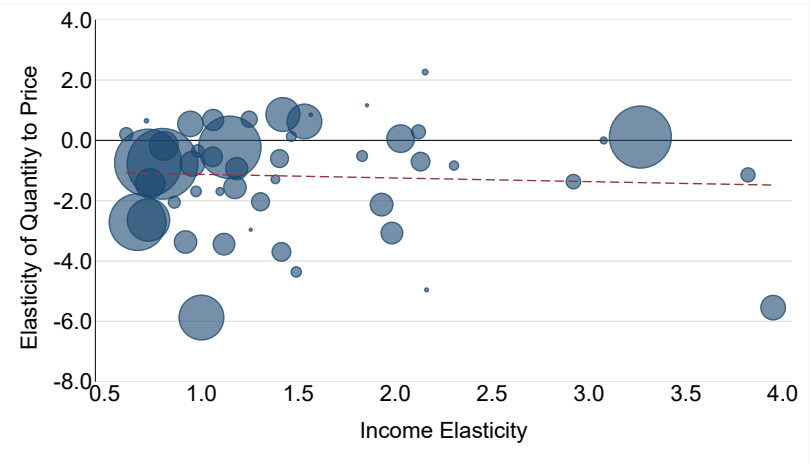
(c) Sales to final consumers



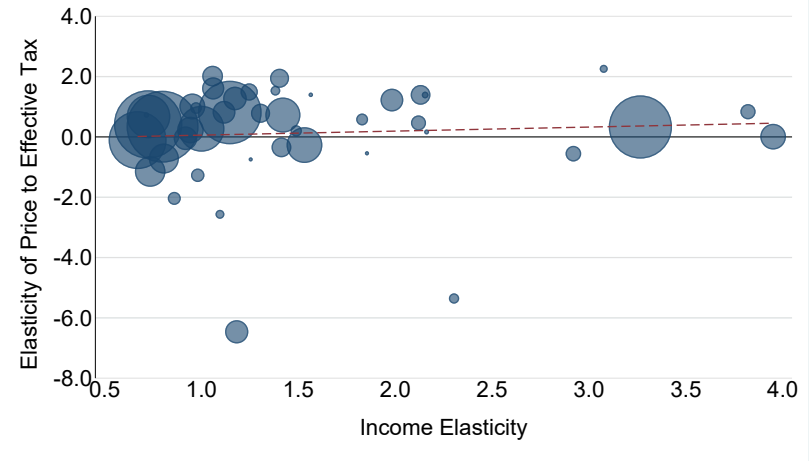
Note: Panel A presents the average statutory tax rate and the average effective tax rates by treatment status using firm-level tax returns (GSTR3). Effective tax rates are winsorised at the 1st and 99th percentile. Panel B and Panel C present the event-study coefficients compating treated and control firms for the log total sales and log sales to final consumers, respectively, firm-level tax returns (GSTR3 and GSTR1). Treated firms are those that, in the registration form, report selling at least one treated product. Sales to final consumers include sales to consumers, to composition scheme firms, firms selling exempted products and firms never claiming ITC.

Figure A.14: Correlation between product elasticity and Engel-curve slope

(a) Quantity Elasticity

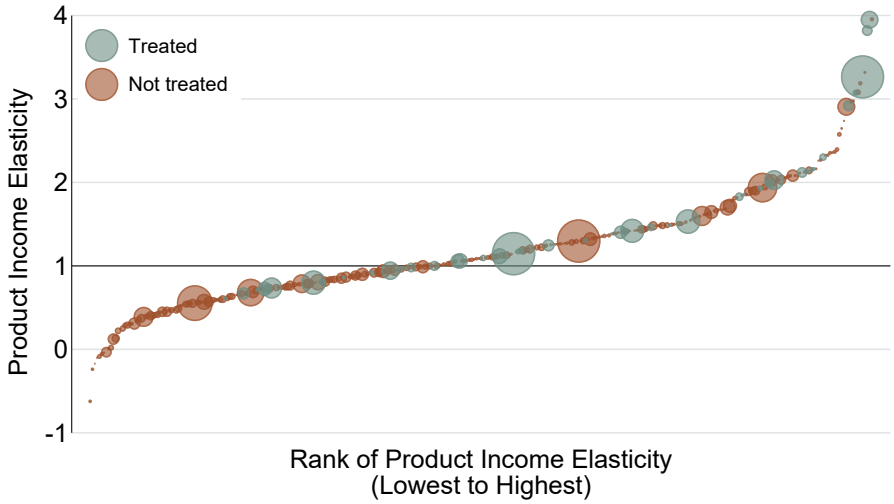


(b) Consumer Price Elasticity



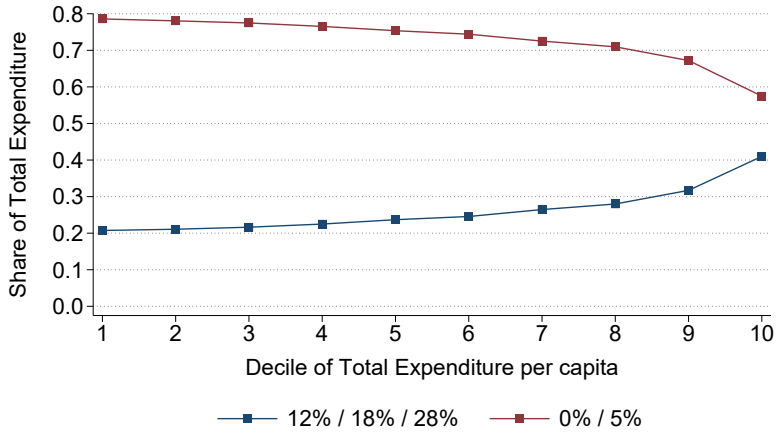
Note: This figure plots the correlation between product quantity and consumer price elasticities with respect to effective tax rates and the Engel-curve slope. Each circle represents a survey product, and circle size increases with the product's total consumption in the survey. The red dashed line shows 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.

Figure A.15: Product Income Elasticity by Treatment Status



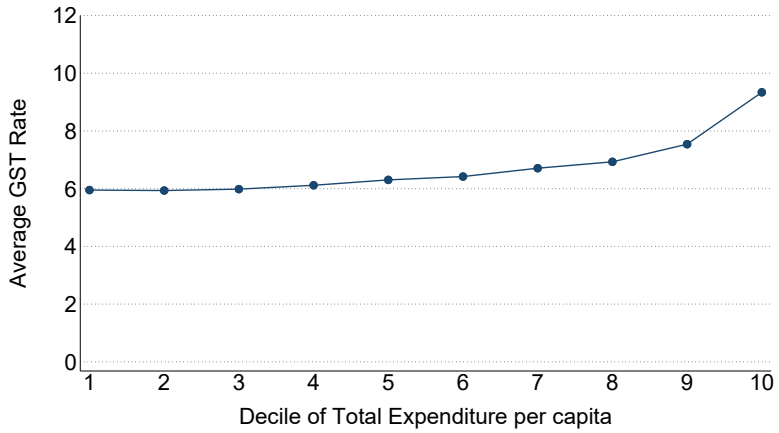
Note: This figure presents Engel beta coefficient of survey products. Treated products are those with a rate reduction from 28% to 18% in November 2017. The circle size increases with product consumption in the survey.

Figure A.16: Expenditure share by total consumption decile and groups of baseline tax rate



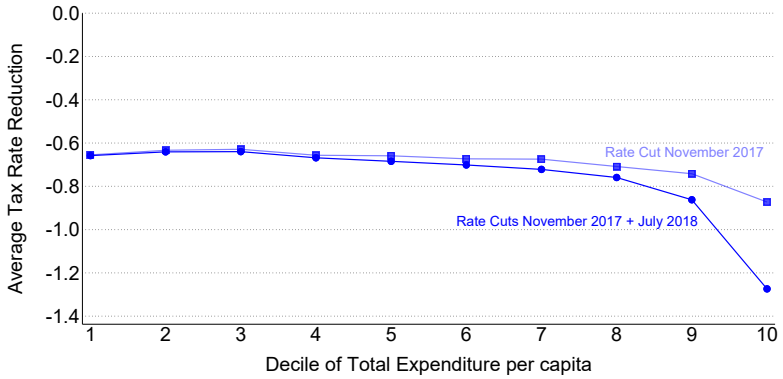
Note: This figure presents the share of total consumption by groups of GST rate in July 2017 across deciles of household total expenditure per capita. Services, goods received as gifts or social transfers and jewellery are excluded from the share of total expenditure.

Figure A.17: Average baseline GST tax rate by total consumption decile



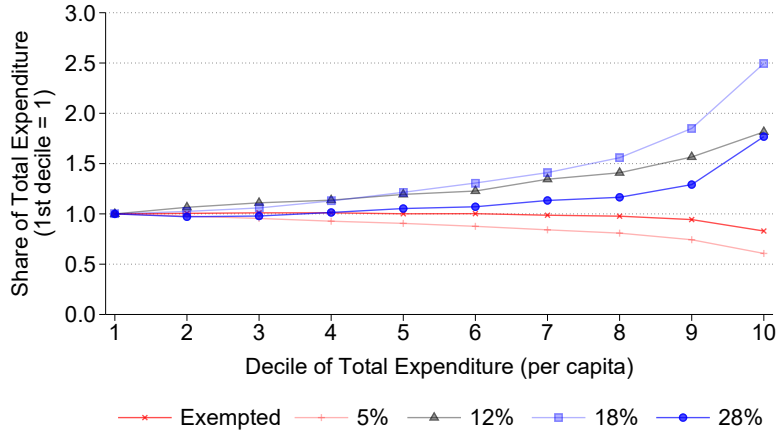
Note: This figure presents the average GST rate in July 2017 across deciles of household total expenditure per capita. The average rate is defined as total taxes divided by total consumption in each decile. Services, goods received as gifts or social transfers and jewellery are excluded from the share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

Figure A.18: Reduction in average tax rate after rate reductions



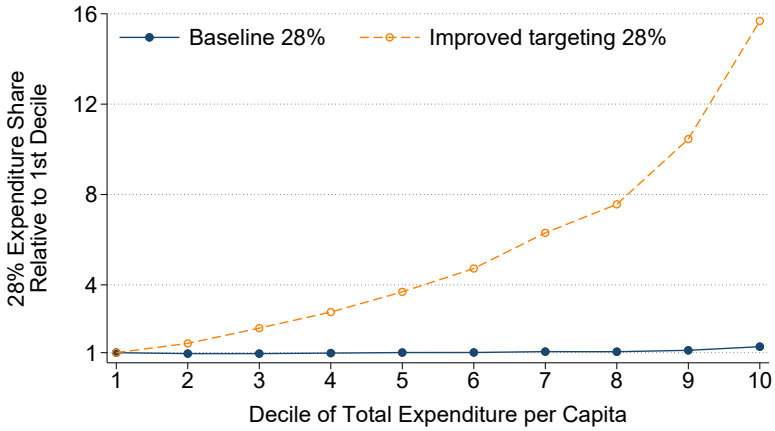
Note: This figure presents the reduction GST rate across deciles of household total expenditure per capita following the November 2017 and July 2018 rate reductions. The average rate reduction is defined as total taxes before minus total taxes after the reduction divided by total consumption in each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

Figure A.19: Normalized expenditure share by total consumption decile and baseline tax rate



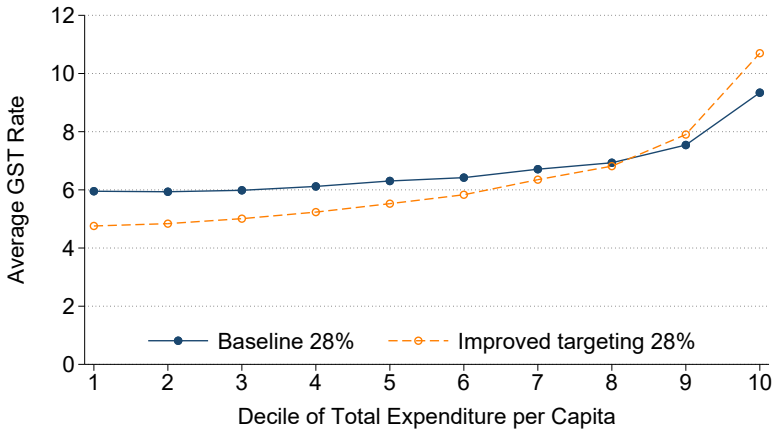
Note: This figure presents the share of total consumption by GST rate in July 2017 across deciles of household total expenditure per capita normalized by the consumption share of the first decile. Services, goods received as gifts or social transfers and jewellery are excluded from the share of total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

Figure A.20: 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.21: 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.

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	1,137	91.40
2	92	7.40
3	13	1.05
4	2	0.16
Total	1,244	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
Initial sample	12,815,721
Services	832,835
Incomplete Tax Returns	
Missing product code	361,296
Non-existant code	384,889
HSN 2-digit code	105,158
Multiple rates	154,946
No quantity	236,195
Final consumers focus	
Final consumer sales share < 0.5	5,837,308
Final consumer sales share missing	468,690
Extreme price variation	227,969
Final sample	4,206,435
Before November 2017	619,739

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 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	1368	619,739	100	100
Treated products	165	117,579	19.0	6.1
Control products	681	284,136	45.9	32.7
Not related	328	119,539	19.3	14.4
Related	353	164,597	26.6	18.3
Other treated	105	52,932	8.5	11.1
Aug 2018	33	21,210	3.4	7.9
Other periods	18	10,110	1.6	1.6
Other rates	54	21,612	3.5	1.6
Other control	417	165,092	26.7	50.1
Rate 0% (exempted)	140	40,195	6.5	21.9
Rate 5%	248	98,446	15.9	17.0
Rate 28%	29	26,451	4.3	11.2

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	0	0	0	0	9	0
5%	11	266	0	0	0	20	4
12%	4	16	167	0	0	14	11
18%	1	4	33	514	0	40	7
28%	0	0	6	195	29	17	87
Total after (%)	10	21	15	52	2		
Treated after (%)	12	7	19	28	0		

Note: This table presents the share of products 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. We restrict the sample to products with more than 100 observations in the four months preceding the reform. The table shows that 18% of products were taxed at 28% at GST onset. Among them, 5.1% experienced a rate reduction from 28% to 18%, while 2% remained taxed at 28%.

Table A.6: Tax rates before and after rate changes

Before \ After	0%	5%	12%	18%	28%	Total before (%)	Treated before (%)
0%	9.3	0.0	0.0	0.0	0.0	9	0
5%	0.3	29.7	0.0	0.0	0.0	30	1
12%	0.2	0.0	15.2	0.0	0.0	15	2
18%	0.0	0.1	1.5	17.5	0.0	19	8
28%	0.0	0.0	0.1	14.9	11.2	26	57
Total after (%)	10	30	17	32	11		
Treated after (%)	5	0	10	46	0		

Note: This table presents the share of sales before November 2017 in each combination of initial (July 2017) and final (March 2020) GST rates. The table shows that 22% of sales were taxed at 28% at GST onset. Among them, 19.1% experienced a rate reduction from 28% to 18%, while 3.2% remained taxed at 28%.

Table A.7: Top products with a rate reduction in November 2017

Product name	Product code (4-digit HSN)	Treated products (%)	Cumulative share (%)
Make-Up	3304	8.1	8.1
Hair Preparations	3305	6.5	14.6
Washing, Cleaning Or Degreasing Preparations	3402	5.8	20.3
Shaving Preparations	3307	5.6	26.0
Plastic Articles (Office And School Supplies)	3926	3.7	29.7
Chocolate	1806	3.5	33.2
Rubber Articles (Gaskets, Seals, Bands)	4016	3.1	36.3
Insulated Electric Conductors (Wires)	8544	3.0	39.3
Electrical Circuits (Switchers, Lamp-Holders, Fuses)	8536	2.9	42.2
Polishes And Creams	3405	2.7	44.9
Eletrical Transformers	8504	2.3	47.2
Malt Extract	1901	2.3	49.5
Mountings	8302	2.1	51.6
Perfumes And Toilet Waters	3303	2.0	53.6
Air Vacuum Pumps	8414	2.0	55.5
Liquid Or Cream Soap	3401	1.9	57.5
Suitcases, Vanity Cases, Briefcases	4202	1.8	59.3
Lamps	8539	1.7	61.0
Lighting Or Signalling Equipment	8512	1.7	62.7
Ceramic Flags And Paving	6907	1.6	64.3
Furniture	9403	1.5	65.8
Printing Machinery	8443	1.3	67.1
Extraxts Of Coffee	2101	1.3	68.3
Mattress Supports	9404	1.2	69.5
Baths And Sinks	3922	1.1	70.7
Sanitary Fixtures	6910	1.1	71.8
Batteries	8506	1.1	72.8
Inner Tubes	4013	1.0	73.8
Bread, Pastry, Cakes, Biscuits	1905	0.9	74.8
Painters' Colours	3213	0.9	75.7
Other Products	.	24.3	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.8: Top control products

Product name	Product code (4-digit HSN)	Treated products (%)	Cumulative share (%)
Medicaments In Measured Doses	3004	22.1	22.1
Insecticides, Rodenticides, Herbicides	3808	6.5	28.6
Tubes, Pipes And Hoses Of Plastic	3917	3.2	31.8
Brooms And Brushes	9603	3.1	34.9
Bar Soap	3401	3.1	38.1
Screws, Bolts, Nuts Of Iron/Steel	7318	3.0	41.1
Medical Instruments	9018	3.0	44.0
Petroleum And Bituminous Oil	2710	2.6	46.6
Prepared Glues And Adhesives	3506	2.5	49.1
Telephone Sets	8517	2.3	51.4
Packing And Closure Goods Of Plastic	3923	2.1	53.5
Medicaments (Not In Measured Doses)	3003	2.1	55.7
Computers	8471	2.1	57.8
Centrifuges And Filtering Apparatus	8421	2.0	59.8
Wadding, Gauze, Bandages	3005	1.8	61.6
Household And Hygienic Articles Of Plastic	3924	1.8	63.3
Prepared Or Preserved Fruit And Nuts	2008	1.7	65.1
Household And Hygienic Articles Of Iron Or Steel	7323	1.7	66.8
Butter And Other Milk Fats And Oils	0405	1.7	68.4
Pens	9608	1.6	70.0
Taps, Cocks, Valves For Pipes	8481	1.5	71.5
Concretes, Cements And Additives	3824	1.5	73.0
Padlocks And Locks	8301	1.5	74.5
Self-Adhesive Items Of Plastics	3919	1.5	76.0
Ball Or Roller Bearings	8482	1.4	77.4
Bread, Pastry, Cakes, Biscuits	1905	1.4	78.7
Plates, Sheets, Film Or Strip Of Plastic	3920	1.3	80.1
Jams And Fruit Or Nut Pastes	2007	1.3	81.4
Tubes, Pipes And Hoses Of Rubber	4009	1.3	82.7
Paper And Paperboard Cut	4823	1.2	83.9
Other Products	.	16.1	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.9: Comparison of ETR compliers and non-compliers

HSN-4 Treated Full	Statistic	ETR match STR	ETR do not match STR
Firm-HSN-Month	N	48,322	13,897
Firm-HSN sales ('000 rupees)	Mean	190	226
	Median	8	5
Firm-HSN code level	Mean	6.18	5.84
	Median	8	4
Firm-HSN months reported	Mean	2.69	2.75
	Median	3	3
Firm-Month	N	13,123	7,267
Firm is urban	Mean	0.48	0.50
	Median	0	1
Firm share no ITC	Mean	0.81	0.84
	Median	0.83	0.89
Firm sales ('000 rupees)	Mean	19,577	28,288
	Median	1,631	1,825
HSN-4 Treated Partial	**Statistic**	**ETR match STR**	**ETR do not match STR**
Firm-HSN-Month	N	34,395	20,107
Firm-HSN sales ('000 rupees)	Mean	209	299
	Median	8	7
Firm-HSN code level	Mean	6.66	5.60
	Median	8	4
Firm-HSN months reported	Mean	2.73	2.79
	Median	3	3
Firm-Month	N	11,858	9,931
Firm is urban	Mean	0.46	0.50
	Median	0.00	1.00
Firm share no ITC	Mean	0.81	0.83
	Median	0.83	0.87
Firm sales ('000 rupees)	Mean	20,642	20,718
	Median	1,827	1,717

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 ETR no lower than 27% and no higher 29%, that is, those with an ETR error lower than 1%. The column *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%. Full HSN-4 treated products restrict to codes in which the entire HSN-4 level underwent the STR reduction from 28% to 18% in November 2017. Partial HSN-4 treated products restrict to codes in part of the HSN-4 level did not underwent the STR reduction from 28% to 18% in November 2017.

Table A.10: Determinants of the STR-ETR initial gap

	STR-ETR Initial Gap
Share of non-ITC sales	1.846
	(0.480)
Sales (log)	-0.103
	(0.044)
Bangalore urban	0.111
	(0.155)
No. periods reported	-0.122
	(0.044)
Firm sales (log)	-0.043
	(0.047)
HSN x Month FE	Yes
Observations	114,611
R^2	0.292

Note: The table reports coefficients from regressions of the STR-ETR gap in percentage points on firm-HSN and firm characteristics. The regression controls for HSN code combined with month fixed effects. The sample is treated products in the months before November 2017 and an observation is a firm-HSN-month. *Share of non-ITC sales* is the average share of firm sales to no-ITC claiming consumers (final consumers) before November 2017. *Sales (log)* is total log sales of this firm-HSN before November 2017. *Bangalore urban* is a dummy if the firm is localted in Urban Bangalore. *No. periods reported* is the number of months a firm-HSN was reported by the firm before November 2017. *Firm sales (log)* is total log sales of the firm before November 2017.

Table A.11: Description of Related and Not Related groups

	Related	Not Related	Control
Products	353	328	681
Firm-HSNs	222,883	124,289	347,172
Of which before Nov 17	82,290	56,306	138,596
Before November 2017:			
Months reported	2.00 (1.14) [2.00]	2.12 (1.18) [2.00]	2.05 (1.16) [2.00]
HSN level reported	6.04 (1.99) [8.00]	6.38 (1.94) [8.00]	6.18 (1.98) [8.00]
Final consumers sales (Firm share)	0.81 (0.16) [0.83]	0.80 (0.17) [0.82]	0.80 (0.16) [0.83]
Sales ('000 rupees)	468 (17853) [5.94]	508 (17529) [7.71]	485 (17718) [6.60]
Sales (Log value)	8.82 (2.55) [8.69]	9.09 (2.66) [8.95]	8.93 (2.60) [8.79]
Quantity (Log units)	3.87 (2.47) [3.47]	4.63 (2.64) [4.41]	4.19 (2.57) [3.87]
Price pre-tax (Log value per unit)	5.13 (2.12) [4.92]	4.60 (1.85) [4.34]	4.91 (2.03) [4.66]
Observations	1,167,114	760,004	1,927,118
Of which before Nov 17	164,597	119,539	284,136

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.12: 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	5.54	6.31
90	Instruments For Demonstrational Purposes	9023	Medical Instruments	9018	8.01	3.12
34	Washing, Cleaning Or Degreasing Preparations	3402	Soap	3401	46.00	3.00
96	Lighters	9613	Brooms And Brushes	9603	1.76	2.81
39	Plastic Articles (Office And School Supplies)	3926	Tubes, Pipes And Hoses Of Plastic	3917	24.49	2.62
73	Sanitary Ware Of Iron/Steel	7324	Screws, Bolts, Nuts Of Iron/Steel	7318	0.69	2.54
84	Air Vacuum Pumps	8414	Computers	8471	15.53	2.05
85	Insulated Electric Conductors (Wires)	8544	Telephone Sets	8517	17.48	1.73
19	Malt Extract	1901	Bread, Pastry, Cakes, Biscuits	1905	18.48	1.30
83	Mountings	8302	Padlocks And Locks	8301	59.34	1.22
40	Rubber Articles (Gaskets, Seals, Bands)	4016	Tubes, Pipes And Hoses Of Rubber	4009	13.37	1.21
33	Make-Up	3304	Dental Hygiene Preparations	3306	81.91	1.19
48	Wallpaper	4814	Paper And Paperboard Cut	4823	0.99	1.10
17	Sugar Confectionery	1704	Sugars, Sugar Syrups, Artificial Honey, Caramel	1702	0.40	0.94
68	Worked Monumental Or Building Stone	6802	Abrasive Powder Or Grain	6805	94.59	0.93
22	Waters Containing Added Sugar Or Flavoured	2202	Natural Or Artificial Mineral Waters And Aerated Waters	2201	0.28	0.81
76	Roofs, Dors And Windows	7610	Table, Kitchen Or Household Articles Of Aluminium	7615	0.38	0.74
21	Extraxts Of Coffee	2101	Sauces And Mixed Condiments	2103	5.80	0.68
82	Razors	8212	Hand Tools (Blow Lamps, Clamps, Anvils)	8205	21.78	0.67
95	Festive Articles	9505	Wheeled Toys And Dolls	9503	0.65	0.64
32	Painters' Colours	3213	Synthetic Organic Colouring Matter	3204	3.99	0.62
70	Mirrors	7009	Glassware For Kitchen, Toilet, Office, Decoration	7013	32.39	0.49
74	Household Items Of Copper	7418	Nails, Tacks, Drawing Pins, Staples	7415	62.08	0.29
91	Watches	9102	Alarm Clocks, Wall Clocks	9105	95.28	0.23
44	Plywood	4412	Wood Sawn Or Chipped Lengthwise	4407	26.51	0.23
94	Furniture	9403	Medical Furniture	9402	79.37	0.16
36	Ferro-Cerium And Spark Alloys	3606	Safety Fuses, Igniters, Detonators	3603	22.60	0.05
69	Ceramic Flags And Paving	6907	Bricks, Blocks, Tiles	6901	81.31	0.04
42	Suitcases, Vanity Cases, Briefcases	4202	Clothing Acessories Of Leather	4203	92.36	0.03
67	Artificial Flowers	6702	Bird Skins And Feather	6701	21.02	0.01
43	Articles Of Furskin	4303	Raw Furskins	4301	0.02	0.01
25	Granite, Basalt, Sandstone (Building Stone)	2516	Marble, Travertine, Ecaussine (Building Stone)	2515	1.67	0.00

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 "Washing, Cleaning Or Degreasing Preparations" and the main control product in HSN-2 34 was "Soap". 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 "Soap" accounted for 2.3% of observations in the related group before November 2017.

Table A.13: 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	-0.053	-0.053	-0.053	-0.053	-0.052	-0.053	-0.053	-0.053	-0.051	-0.053	-0.054	-0.051	-0.053	-0.054	-0.059	-0.062	-0.053	-0.054	-0.059
	(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	-0.027	-0.030	-0.027	-0.026	-0.028	-0.024	-0.029	-0.027	-0.024	-0.034	-0.029	-0.020	-0.028	-0.028	-0.018	-0.020	-0.030	-0.020	-0.018
	(0.005)	(0.005)	(0.005)	(0.005)	(0.009)	(0.005)	(0.007)	(0.005)	(0.008)	(0.008)	(0.006)	(0.007)	(0.005)	(0.005)	(0.011)	(0.013)	(0.007)	(0.014)	(0.064)
Quantity	0.032	0.031	0.031	0.032	0.048	0.024	0.027	0.033	0.046	0.060	0.056	0.016	0.028	0.032	0.049	0.049	0.037	0.051	0.024
	(0.016)	(0.016)	(0.015)	(0.015)	(0.025)	(0.016)	(0.016)	(0.016)	(0.023)	(0.021)	(0.019)	(0.016)	(0.016)	(0.016)	(0.019)	(0.032)	(0.017)	(0.019)	(0.075)
Sales	0.058	0.054	0.056	0.060	0.072	0.052	0.052	0.058	0.073	0.081	0.081	0.047	0.053	0.058	0.091	0.093	0.061	0.085	0.065
	(0.016)	(0.016)	(0.016)	(0.016)	(0.024)	(0.016)	(0.016)	(0.016)	(0.025)	(0.024)	(0.020)	(0.016)	(0.017)	(0.017)	(0.021)	(0.034)	(0.016)	(0.017)	(0.055)
Reporting Rate	-0.002	-0.003	-0.002	-0.002	-0.003	-0.002	-0.002	-0.002	-0.001	-0.002	-0.003	-0.001	-0.001	-0.007	_	-0.007	0.001	-0.004	_
	(-0.002)	(-0.003)	(-0.002)	(-0.002)	(-0.003)	(-0.002)	(-0.002)	(-0.002)	(-0.001)	(-0.002)	(-0.003)	(-0.001)	(-0.001)	(-0.007)		(-0.007)	(0.001)	(-0.004)	
Consumer Price Elasticity (ϵ_p)	0.503	0.553	0.513	0.484	0.534	0.455	0.539	0.539	0.456	0.635	0.536	0.389	0.515	0.503	0.302	0.321	0.552	0.371	0.430
	(0.100)	(0.101)	(0.104)	(0.100)	(0.173)	(0.100)	(0.129)	(0.129)	(0.152)	(0.165)	(0.120)	(0.133)	(0.100)	(0.101)	(0.180)	(0.209)	(0.134)	(0.246)	(0.967)
Quantity Elasticity (ϵ_q)	-0.577	-0.558	-0.571	-0.586	-0.899	-0.445	-0.499	-0.499	-0.881	-1.107	-1.019	-0.318	-0.508	-0.559	-0.828	-0.744	-0.682	-0.930	-0.438
	(0.305)	(0.299)	(0.294)	(0.288)	(0.481)	(0.299)	(0.308)	(0.308)	(0.452)	(0.404)	(0.354)	(0.322)	(0.309)	(0.313)	(0.326)	(0.484)	(0.325)	(0.344)	(1.168)
Sales Elasticity (ϵ_s)	-1.071	-1.003	-1.053	-1.105	-1.361	-0.989	-0.961	-0.961	-1.432	-1.513	-1.481	-0.922	-0.980	-1.051	-1.542	-1.454	-1.130	-1.559	-1.007
	(0.313)	(0.306)	(0.305)	(0.299)	(0.466)	(0.312)	(0.312)	(0.312)	(0.475)	(0.457)	(0.364)	(0.328)	(0.315)	(0.320)	(0.364)	(0.515)	(0.305)	(0.309)	(0.841)
Quantity-Price Elasticity (ϵ_q^p)	-1.245	-1.106	-1.137	-1.296	-2.065	-0.981	-1.045	-1.045	-1.862	-1.851	-1.990	-0.856	-1.110	-1.186	-2.755	-2.395	-1.311	-2.523	-0.554
	(0.542)	(0.442)	(0.556)	(0.570)	(0.734)	(0.634)	(0.498)	(0.498)	(0.848)	(0.756)	(0.632)	(0.788)	(0.506)	(0.555)	(1.864)	(2.073)	(0.565)	(1.478)	(0.868)
N	2,513,097	3,114,993	1,841,160	2,223,839	1,427,447	1,801,952	2,664,518	2,520,735	1,216,140	838,997	2,230,799	1,387,973	2,414,841	1,535,022	258,384	1,678,232	1,623,849	199,472	10,227

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.14: Description of treatment and control group

	Treated	Control	All
Products	25	681	706
Firm-HSNs	20,392	344,666	365,058
Of which before Nov 17	14,003	243,764	257,767
Before November 2017:			
Months reported	3.59 (3.39) [2.00]	3.84 (3.53) [2.00]	3.83 (3.52) [2.00]
HSN level reported	5.76 (1.98) [4.00]	6.12 (1.98) [8.00]	6.10 (1.99) [8.00]
Final consumers sales (Firm share)	0.83 (0.16) [0.87]	0.81 (0.17) [0.84]	0.81 (0.17) [0.85]
Sales ('000 rupees)	1796 (21985) [14.38]	388 (20419) [5.79]	459 (20504) [6.04]
Sales (Log value)	9.70 (2.71) [9.57]	8.81 (2.57) [8.66]	8.86 (2.59) [8.71]
Quantity (Log units)	3.10 (2.00) [2.71]	4.10 (2.52) [3.76]	4.05 (2.51) [3.71]
Price pre-tax (Log value per unit)	6.80 (2.20) [6.74]	4.87 (1.98) [4.65]	4.97 (2.04) [4.73]
Observations	101,457	1,889,752	1,991,209
Of which before Nov 17	50,211	935,974	986,185

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.15: 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	18.0	18.0
Heating And Thermic Appliances, Hair-Dressers	8516	16.2	34.2
Electric Accumulators	8507	11.3	45.5
Paints And Vanishes (Non-Aqueous Disolved)	3208	10.7	56.2
Paints And Vanishes (Aqueous Disolved)	3209	7.1	63.3
Refrigerators And Freezers	8418	6.2	69.5
Suitcases, Vanity Cases, Briefcases	4202	5.3	74.8
Laudry Washing Machines	8450	4.8	79.6
Food Grinders And Mixers, Juice Extractors	8509	4.5	84.1
Shavers	8510	2.9	87.0
Motor Cars For The Transport Of Persons	8703	2.7	89.7
Scent And Toilet Sprays	9616	2.5	92.3
Vacuum Cleaners	8508	1.8	94.0
Paints And Vanishes (Enamels, Lacquers, Distempers)	3210	1.4	95.4
Wooden Frames For Art Or Mirrors	4414	1.3	96.7
Trailers And Semi-Trailers	8716	0.7	97.4
Chains Of Copper	7419	0.5	97.9
Framed Mirrors	7009	0.4	98.3
Carving And Moulding Material And Articles	9602	0.4	98.7
Works Trucks	8709	0.4	99.1
Globes For Lamps And Lanterns	7020	0.4	99.5
Festive Articles	9505	0.2	99.8
Other Articles Of Wood	4421	0.1	99.9
Special Purpose Vehicles	8705	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.16: Heterogeneity on proxies for informality exposure

	Product Complexity		Product Sales		Employees (Sector)		Firm Location		Firm Size	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Below median	Above median	Below median	Above median	Below median	Above median	Not Capital	Capital	Below median	Above median
ETR	-0.056	-0.048	-0.052	-0.054	-0.051	-0.055	-0.054	-0.052	-0.053	-0.053
	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Consumer Price	-0.027	-0.024	-0.020	-0.029	-0.025	-0.026	-0.030	-0.022	-0.032	-0.021
	(0.007)	(0.007)	(0.007)	(0.008)	(0.007)	(0.007)	(0.007)	(0.007)	(0.006)	(0.007)
Quantity	0.028	0.048	0.038	0.002	0.020	0.050	0.045	0.016	0.028	0.034
	(0.019)	(0.021)	(0.019)	(0.022)	(0.017)	(0.027)	(0.020)	(0.017)	(0.018)	(0.020)
Sales	0.056	0.075	0.070	0.027	0.047	0.079	0.069	0.047	0.050	0.066
	(0.020)	(0.022)	(0.019)	(0.024)	(0.018)	(0.026)	(0.020)	(0.018)	(0.019)	(0.020)
N	1,237,895	1,212,339	1,266,260	1,198,803	1,248,011	1,220,826	1,381,497	1,118,245	1,029,602	1,084,993

Note: Each cell reports estimates from separate difference-in-differences regressions on subsamples defined by proxies for exposure to informal sector competition. Product Complexity classify products based on the Atlas of Economic Complexity. Product Sales classify products based on the average firm-HSN-period sales of the HSN-4 before November 2017. Employees (Sector) classify products based on median number of firm employees in the sector using ASI data harmonized to 16 aggregated HSN-based sectors. Firm Location separates firms in Bangalore Urban (state capital) from those elsewhere. Firm size divides firm-HSNs based on firm sales before November 2017. All outcomes are in logs.

Table A.17: Description of firm-level tax returns

	Firm Return	+ Detailed Firm Return	+ Firm-Product Return	All
Firms	187,907	285,265	107,727	399,494
Tax returns	370,199	701,485	259,596	1,331,280
	27.8%	52.7%	19.5%	100%
Sales	1,099	2,171	21,048	5,554
	(27,651)	(67,329)	(339,952)	(158,729)
	[217]	[339]	[1,275]	[362]
	5.5%	20.6%	73.9%	100%
Taxes	60	225	2,261	576
	(1,838)	(2,969)	(48,816)	(21,702)
	[15]	[38]	[137]	[36]
	2.9%	20.6%	76.5%	100%
Products registered	2.63	2.70	3.16	2.77
	(1.63)	(1.67)	(1.91)	(1.72)
	[2.00]	[2.00]	[3.00]	[2.00]
	26.4%	51.4%	22.3%	100%

Note: This table divides firm-level tax returns into three mutually exclusive groups. The All column pools all groups. Sales and tax amounts are reported in thousand rupees. 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.

Table A.18: Description of firm-level tax returns

	Treated	Control	All
Firms	78,151	203,244	281,395
Tax returns	1,066,453	2,904,585	3,971,038
with detailed	751,704	2,018,033	2,769,737
with firm-product	205,752	520,983	726,735
Sales	3,308 (53,034) [352]	4,843 (223,044) [414]	4,431 (192,728) [397]
Sales Final Consumer	991 (21,323) [113]	1,013 (18,622) [112]	1,007 (19,394) [112]
Taxes	433 (6,858) [51]	457 (18447) [37]	451 (16,172) [40]
Products registered	3.18 (1.52) [3.00]	2.35 (1.50) [2.00]	2.57 (1.55) [2.00]

Note: This table divides firm-level tax returns into three mutually exclusive groups. *Only Firm Returns* includes firm-months with only the aggregate firm-level return (GSTR-3). *Detailed Firm Returns* includes firm-months for which both the aggregate and detailed firm-level returns are available (GSTR-3 and GSTR-1). *Firm and Firm-Product Returns* includes firm-months the aggregate and detailed firm-level returns (GSTR-3 and GSTR-1) and firm-HSN returns (GSTR-1 Table 12). Sales and tax amounts are reported in thousand rupees. For each variable, the table reports the mean, standard deviation in parentheses, and median in brackets.

Table A.19: Robustness of firm-level results

		Control		Balance		
	(1)	(2)	(3)	(4)	(5)	(6)
	Main Sample	Include Exempted & Services	Remove 5%	Weekly Balanced	Strongly Balanced	Sales Weighted
ETR	-0.019	-0.018	-0.019	-0.019	-0.020	-0.018
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Sales	0.060	0.063	0.046	0.057	0.048	0.049
	(0.011)	(0.010)	(0.014)	(0.011)	(0.011)	(0.019)
Sales Elasticity (ϵ_s)	-2.706	-3.001	-2.047	-2.621	-2.209	-2.341
	(0.486)	(0.486)	(0.576)	(0.488)	(0.436)	(1.022)
ETR Final consumers	-0.019	-0.018	-0.019	-0.019	-0.020	-0.020
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Sales Final consumers	0.051	0.034	0.055	0.047	0.040	0.020
	(0.018)	(0.015)	(0.023)	(0.018)	(0.017)	(0.024)
Sales Elasticity (ϵ_s) Final consumers	-2.788	-2.034	-2.840	-2.626	-2.256	-1.619
	(0.850)	(0.793)	(1.108)	(0.853)	(0.833)	(1.411)
N detailed	2,383,885	3,911,072	1,630,624	2,140,541	1,581,135	2,182,815
N	3,954,623	6,839,119	2,623,739	3,400,234	2,243,556	3,499,897

Note: This table reconciles the firm-level estimates across alternative samples and balancing choices. Column (1) reports the baseline specification. Column (2) expands the sample by including exempted products and services. Column (3) drops firms selling at least one product taxed at 5%. Column (4) restricts the estimation to a weekly balanced panel, defined as reporting at least once before and once after November 2017. Column (5) restricts the sample to a strongly balanced panel, defined as reporting in every period. Column (6) weights observations by total sales before November 2017, with weights winsorised at the 97.5th percentile. All regressions include firm and time fixed effects, and cluster standard errors using a random HSN reported at registration and the firm level.

Table A.20: Reconcialiation of firm-level results

	Not sales weighted			Sales weighted		
	(1) Firm Return	(2) + Detailed Firm Return	(3) + Firm-Product Return	(4) Firm Return	(5) + Detailed Firm Return	(6) + Firm-Product Return
ETR	-0.017	-0.018	-0.024	-0.011	-0.014	-0.023
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Sales	0.074	0.058	0.059	0.065	0.053	0.036
	(0.012)	(0.012)	(0.014)	(0.038)	(0.029)	(0.020)
Sales Elasticity (ϵ_s)	-3.575	-2.927	-2.177	-5.087	-3.698	-1.590
	(0.585)	(0.548)	(0.466)	(3.233)	(1.742)	(0.714)
N	1,170,398	2,014,630	709,747	987,948	1,806,860	662,204

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. Treatment is defined using the registration form, and firms are considered treated if they list at least one treated product. Columns (4), (5) and (6) report sales-weighted regressions using pre–November 2017 sales, with weights winsorised at the 97.5th percentile.

Table A.21: 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.22: 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.23: 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	-0.120	0.285	0.448	0.303
	(0.331)	(0.283)	(0.431)	(0.322)
Dependent variable: Elasticity of consumer price to effective taxes				
Income elasticity	-0.125	-0.113	0.095	0.193
	(1.005)	(0.603)	(0.169)	(0.167)
Observations (Survey Products)	48	48	48	48

Notes: This table reports estimates from regressions of survey-product elasticities on the corresponding Engel-curve coefficients, as specified in equation (4). 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."[36] In contrast, he referred to his version of GST as the "Genuine Simple Tax."[37] 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?"[38]

[36] 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

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

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

B Optimal Commodity Taxation

B.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 is given by $v(p, y^i)$. 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

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

where aggregate demand for good j is

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

and consumer prices satisfy

$$p_j = (1+\tau_j)q_j.$$

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

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

Assuming perfect pass-through ($\partial p_j / \partial \tau_j = q_j$), 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

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

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

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

Define the (own-price) demand elasticity

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

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

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

Interpretation This expression makes explicit the trade-off faced by the government. The denominator, ε_j^x, captures the efficiency cost of taxation: 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

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

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.

B.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

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

the derivative of welfare with respect to τ_j is

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

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

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

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

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

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

$$\int_i g^i a_j^i\, \varepsilon_j^p\, di = \mu\left(1 - \tau_j \varepsilon_j^s\right),$$

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

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

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.

B.3 Firm Ownership

We now extend the analysis by allowing individuals to be firm owners. Individual income includes profit income, so that

$$y^i = y_0^i + \pi^i, \qquad \pi^i = q_j x_j^i,$$

where we assume production is costless and profits proportional to sales. 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 τ_j yields

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

where the additional term reflects the effect of the tax on profit income.
Define social marginal welfare weights

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

with aggregate profits $\pi_j = q_j x_j$.
Dividing by $q_j x_j$ and rearranging, the first-order condition becomes

$$\int_i g^i \left(a_j^i \varepsilon_j^p + s_j^i \varepsilon_j^\pi\right) di = \mu\left(1 + \tau_j \varepsilon_j^s\right),$$

where

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

Given the profit structure implied by costless production and profits proportional to sales, $\pi_j = q_j x_j$, profit responses to taxation satisfy

$$\frac{d\pi_j}{d\tau_j} = \frac{d(q_j x_j)}{d\tau_j}, \qquad \Rightarrow \qquad \varepsilon_j^\pi = \varepsilon_j^s.$$

Substituting this relationship into the first-order condition and rearranging yields the optimal tax rule with firm ownership:

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

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.

B.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,

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

Differentiating welfare with respect to τ_j yields

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

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

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

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

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

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

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

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:

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

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.