Why VAT Pass-Through Varies Across Countries: The Role of Market Power*

Tiago Bernardino
IIES, Stockholm University

Pedro Brinca
Nova SBE

Saman Darougheh
Danmarks Nationalbank

Márcia Silva-Pereira
Banco de Portugal and Nova SBE

October, 2025

Abstract

We show that VAT pass-through rates depend systematically on market concentration using data from 16 European countries covering 1999-2019. Low-concentration industries exhibit 50% contemporaneous pass-through to consumer prices, while high-concentration industries show near-zero transmission. Market concentration varies substantially across countries, and this variation explains 16% of cross-country differences in pass-through rates. Our results suggest that optimal VAT policy should account for local market structure differences.

Keywords: VAT pass-through, Market concentration, Cross-country variation
JEL Classification: H22, H25, H87, L11, L13

*T. Bernardino acknowledges financial support from the Jan Wallanders & Tom Hedelius Foundation (H21-0020). P. Brinca acknowledges funding from Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2013, UID/ECO/00124/2019, PTDC/EGE-ECO/7620/2020, and Social Sciences DataLab, LISBOA-01-0145-FEDER-022209), POR Lisboa (LISBOA-01-0145-FEDER-007722, LISBOA-01-0145-FEDER-022209), POR Norte (LISBOA-01-0145-FEDER-022209) and CEECIND/02747/2018. M. Silva-Pereira acknowledges financial support from Fundação para a Ciência e a Tecnologia (2021.04898.BD). The views expressed herein are those of the authors and should not be attributed to any of the affiliated institutions.

1 Introduction

Value-added taxes (VAT) represent the largest source of government revenue in most developed economies and have increasingly been used as macroeconomic stabilization tools. Central to the effectiveness of VAT policy is the pass-through rate, which determines how much of a tax change is actually passed on to consumers versus being absorbed by firms. Yet empirical evidence reveals striking heterogeneity in pass-through estimates across sectors and countries. As an example for the former, Carbonnier (2007) studies French VAT changes and finds a pass-through of 57% for cars and 77% for housing repairs. Meanwhile, VAT reductions in the food sector during the 2021-2022 high-inflation episode were fully passed through in Portugal (Bernardino et al., 2025), but only approximately 70% passed through in Germany (Fuest, Neumeier, and Stöhlker, 2024).[1]

This heterogeneity creates fundamental uncertainty about when VAT policies will effectively reach consumers versus primarily benefit firm profits (for a discussion of different pass-through rates, see Benzarti, 2025). While economic theory suggests market structure should be a key determinant of tax pass-through (Dierickx, Matutes, and Neven, 1988), systematic empirical evidence linking concentration to VAT transmission has been limited.

We utilize comprehensive data on VAT rates (Benzarti and Tazhitdinova, 2021) and market concentration, covering 16 European countries and 50 product categories from 1999 to 2019, to document two main findings. First, market power is a significant determinant of the pass-through of value-added taxes. VAT changes in low-concentration markets exhibit a contemporaneous pass-through of 50%, whereas high-concentration markets show near-zero transmission to consumer prices. Second, the (large) cross-country differences in average market concentration explain 16% of the variation in country-level pass-through rates.

Our findings show that seemingly similar VAT reforms can produce dramatically different consumer outcomes across European countries. While the VAT Directive (2006/112/EC) has harmonized VAT frameworks across EU member states, our results suggest that optimal VAT policy should account for local market structure differences.

[1] The heterogeneity in pass-through estimates is further exemplified by 50% for Finnish haircare services (Kosonen, 2015) and 100% in Norwegian food sectors (Gaarder, 2019).

We complement recent work by Dimitrakopoulou et al. (2024), who show in the case of Greece that gasoline pass-through increases from 50% in monopolistic markets to 80% in competitive ones. The vertical supply chain matters, too: Rozema (2018) finds that downstream firms only bear one third of the firm share of the tax burden, and, in concurrent work, Bellon, Copestake, and Zhang (2024) demonstrate that upstream market competition affects overall VAT pass-through. Beyond market structure, firm-level characteristics also shape pass-through heterogeneity. Harju, Kosonen, and Skans (2018) analyze VAT reductions for restaurants in Finland and Sweden and document differences in price adjustment behavior between independent restaurants and chain restaurants, reflecting differences in management and price-setting practices. Kosonen (2015) finds heterogeneity by firm size; with a larger pass-through for larger firms. We contribute to this literature by providing a systematic cross-country evidence linking downstream market concentration to VAT pass-through across 16 European countries and 50 product categories.

In the remainder of this paper, we first discuss the theoretical relationship between market power and pass-through in Section 2. Here we also show that this relationship is very similar whether estimated using the Herfindahl Index—commonly used in empirical work—or the Lerner Index, the theory-consistent measure of market power. Section 3 contains our main empirical exercise. Section 4 concludes.

2 Theory

While our primary contribution lies in the empirical analysis, we provide a simple theoretical framework to characterize the relationship between market power and VAT pass-through. We rely on a Logit demand model, a well-studied environment for which we present only key equations. For a complete description of this model, we refer to Anderson, Palma, and Thisse (1992).

We use the model to make two points. First, we reiterate previous findings that pass-through of value-added taxes is decreasing in market power in the context of Logit demand: firms with less elastic demand find it less profitable to raise prices proportionally when already extracting significant consumer rents. Second, we demonstrate that the Herfindahl Index, despite not being structurally derived from this model, captures the market power-pass-through relationship sufficiently well for empirical analysis.

Model Setup. We model a single sector which comprises $j = 1, \ldots, N$ firms facing outside competition from a good indexed $j = 0$. Each household i has logit preferences over these $N + 1$ goods, with utility from good $j \in [0, \ldots, N]$ given by

$$u_{ij} = \alpha - \beta p_j^c + \epsilon_{ij}, \tag{1}$$

where p_j^c is the after-tax price of good j, α captures average preference for market goods relative to the outside option (e.g., not purchasing, self-production, or alternative markets), β is price sensitivity, and ϵ_{ij} are idiosyncratic taste shocks distributed Type I extreme value. We normalize the outside good to have zero price and zero mean utility. The resulting choice probabilities are

$$s_j = \frac{\exp(\alpha - \beta p_j^c)}{1 + \sum_{k=1}^{N} \exp(\alpha - \beta p_k^c)} \tag{2}$$

$$s_0 = \frac{1}{1 + \sum_{k=1}^{N} \exp(\alpha - \beta p_k^c)}, \tag{3}$$

where s_0 denotes the market share of the outside good.

Firm j produces with a constant marginal cost c_j. Firm j chooses before-tax price p_j^f to maximize profits π_j given competitors' prices $\mathbf{p}_{-j}^c$:

$$\max_{p_j^f} \pi_j(\mathbf{p}_{-j}) = (p_j^f - c_j) \cdot s_j(p_j^c, \mathbf{p}_{-j}^c), \tag{4}$$

where the optimal before-tax price is given by

$$p_j^f = c_j + \frac{1}{\beta(1+\tau)(1-s_j)}. \tag{5}$$

The after-tax price is given by $p_j^c = p_j^f(1+\tau)$. An equilibrium is a set of prices $\mathbf{p}^c$ where each price is optimal given all others. This model has a unique equilibrium which can be solved for using its contraction mapping property (Caplin and Nalebuff, 1991).

Pass-through and market power. Using a first-order approximation (holding market shares fixed), the pass-through elasticity of firm j can be written as

$$\gamma_j = \frac{d\ln p_c}{d\ln(1+\tau)} = \frac{1}{1 + \frac{1}{\beta c_j (1+\tau)(1-s_j)}}. \tag{6}$$

In this context, to compute the market power of firm j we begin with its own-price elasticity of demand, $\eta_j \equiv -\frac{p_j^c}{s_j}\frac{\partial s_j}{\partial p_j^c} = \beta p_j^c(1-s_j)$. The market power is then commonly defined as the Lerner Index, the inverse of the own-price elasticity: $\mu_j \equiv 1/\eta_j$.

The pass-through γ_j decreases with market power: firms with lower price elasticity η_j have a higher Lerner Index μ_j and a lower pass-through γ_j (Weyl and Fabinger, 2013). The intuition is as follows: firms with substantial market power already earn high markups over marginal cost. When VAT changes increase costs, these firms face a tradeoff: maintain their large markup by fully passing through the tax (and losing customers), or compress their markup slightly while raising prices less. Absorbing part of the tax change through margin compression lowers profits by less than passing on the full price change and losing customers. In contrast, firms with low market power operate on thin margins and cannot afford significant margin compression—it is optimal for them to pass through a much higher share of the tax change.

In a symmetric equilibrium (where all firms have the same marginal cost $c_j = c$), it is also straightforward to see that the average pass-through rate is higher in industries where firms, on average, have higher market concentration. In asymmetric equilibria, this relationship is more nuanced, and we will turn towards a numerical simulation.

Empirical Measurement Challenge. Our goal is to measure this relationship empirically across European countries, where firm-level own-price elasticities are notoriously difficult to estimate. Researchers thus commonly use the Herfindahl-Hirschman Index (HHI) as a proxy for market power. This approach has important limitations. First, concentration is an equilibrium outcome rather than a structural parameter. Second, industry concentration may imperfectly capture market power in the presence of import competition or other outside forces. Despite these shortcomings, the HHI's broad availability and comparability across countries and industries makes it a natural choice, consistent with recent work on VAT incidence (Bellon, Copestake, and Zhang, 2024). We now examine whether HHI correlates sufficiently with pass-through to serve as a useful empirical proxy.

Numerical Validation. Variation in market power can stem from two factors. First, differences in the number of firms. Changing N holds the *relative* distribution of market shares fixed, but proportionately scales down market shares of each firm. Second, for a given number of firms, we can increase the heterogeneity of market shares. This will increase the market power of some firms and decrease that of others.

Turning to the model, we vary these two dimensions as follows. We fix $\beta = 2$ and $\tau = 0.15$, and simulate 100 markets. For each market, we draw the number of firms N uniformly from $\{5, \ldots, 50\}$ and cost dispersion $\sigma \sim U(0, 1)$. We then draw heterogeneous costs c_j for all N firms with standard deviation σ.[2] For each market, we compute the revenue-weighted average Lerner Index μ and HHI using revenue shares r_j of firms $j = 1, \ldots, N$, excluding the outside option: $\text{HHI} = \sum_{j=1}^{N} r_j^2$. Pass-through is computed numerically by increasing τ from 0.15 to 0.16. These parameter values are not meant to quantitatively match any specific country; rather, we test whether the HHI captures the relationship between pass-through and market power across for reasonable parameter values.

Figure 1 shows results for two values of α, where both the HHI and the Lerner Index have been standardized for comparability. We first turn to panel (a), where we have chosen a high preference for the market good, $\alpha = 3$. We show in blue the relationship between the Lerner index μ and the pass-through rate: this relationship is relatively linear with only minor dispersion around the regression line. We can see that a one-standard-deviation increase in market power corresponds to approximately a fall in pass-through by 30 percentage points. We plot in red the relationship where we instead estimate market power using the HHI. We find that the HHI approximates the relationship very well, with both very similar point estimates and a low dispersion around the estimated regression line.

[2] To ensure that the c_j are well-behaved, we assume that $c_j = 1 + \zeta_j$, where $\zeta_j \sim N(0, \sigma^2)$. We furthermore bound c_j from below by 0.1.

Figure 1: Model-simulated Market Power vs Pass-Through Elasticity

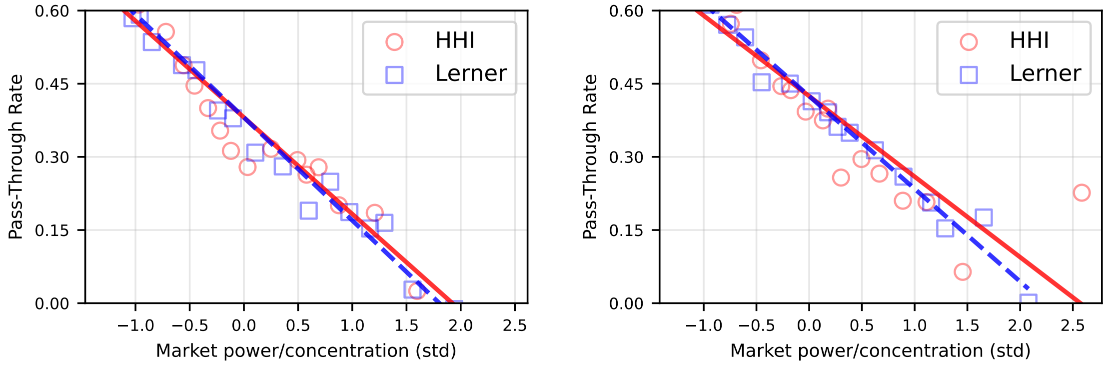

(a) Low outside competition ($\alpha = 3$)

(b) High outside competition ($\alpha = 1$)

Notes: The figure plots pass-through elasticity against two measures of market power using simulated data, for two values of α that determine the attractiveness of the outside option. For each value of α, we simulate 100 markets, drawing the number of firms uniformly from $\{5, \ldots, 50\}$ and cost dispersion σ uniformly from $(0, 1)$ for each market. We compute the revenue-weighted HHI and Lerner index for each market and simulate a tax increase from 15% to 16%. We standardize both indices and aggregate the observations into 20 bins to estimate pass-through against average market power within each bin. Blue squares show the Lerner index; red circles show the HHI. The solid red line and dashed blue line indicate the corresponding regression lines.

Robustness to Outside Competition. We compare these findings to $\alpha = 1$, where the outside option is more attractive. This raises the outside option's average market share from 3% to 20% (varying from 9% to 30% across markets). This tests a key concern: what if firms face substantial outside competition that HHI fails to capture? We find that increasing outside competition does not significantly alter our conclusions.

The simulation uses a simplified demand model with specific parameter values; the relationship between HHI, Lerner index, and pass-through could differ under alternative parameterizations. Are our choices reasonable? The average (non-standardized) HHI is approximately 0.1, within the range observed in our empirical analysis in the next section. The simulated VAT increase (15% to 16%) matches typical policy changes discussed in the next section. Finally, the simulated elasticity—a one-standard-deviation increase in HHI decreases pass-through by 30pp—closely matches our baseline empirical estimate of the contemporaneous pass-through, 26pp.

We conclude that market power and VAT pass-through exhibit a negative relationship, consistent with standard oligopoly theory. The numerical analysis demonstrates that the HHI, despite being an aggregate measure rather than a structural parameter, serves as an effective empirical proxy for market power in predicting pass-through, even when outside competition is substantial.

3 Empirical analysis

In this section, we first introduce the data used. We then study the relationship between market concentration and the pass-through of value-added taxation. Finally, we ask to what extent cross-country differences in market concentration explain cross-country differences in VAT pass-through.

3.1 Data

Prices. We use monthly price data from the Harmonized Index of Consumer Prices (HICP) published by Eurostat at the COICOP 5-digit level (Classification of Individual Consumption by Purpose). Data span 1999–2019, matching the period for which VAT rates are available.

VAT rates. We merge price data with the historical VAT rates database compiled by Benzarti, Carloni, et al., 2020 and extended by Benzarti and Tazhitdinova, 2021, containing all value-added tax changes by category for each European country from 2000–2019.

Concentration. We use industry-level revenue-based Herfindahl-Hirschman Indices (HHI) for two-digit NACE industries from the Competitiveness Research Network (CompNet) to proxy market power.[3] We focus on industries operating at the bottom of the supply chain that interact directly with consumers. For COICOP5 categories associated with multiple NACE2 industries, we take the average HHI across these industries. The previous section demonstrates that HHI can serve as a reasonable proxy for market power despite its limitations. This, combined with its broad availability and comparability across countries and industries, makes it a natural choice for our setting, consistent with other recent work on VAT incidence (e.g., Bellon, Copestake, and Zhang, 2024).

Our final sample contains 16 countries with VAT changes in up to 50 product groups (COICOP5 categories) during our period, summarized in Table 1.

[3] Two-digit is the highest disaggregation at which CompNet provides the HHI.

Table 1: Summary statistics by country

Country	#Products	Avg. HHI	#VAT Events	Estimated $\hat{\gamma}$	Estimated $\hat{\gamma}$ (s.e.)
CZ	49	0.037	75	0.598	0.204
DE	50	0.019	32	0.185	0.076
ES	47	0.077	64	0.122	0.087
FI	50	0.022	79	0.717	0.250
FR	50	0.013	46	−0.617	0.477
HR	49	0.052	68	0.129	0.108
HU	48	0.048	95	0.272	0.113
IT	49	0.013	50	1.074	0.496
LU	50	0.073	15	−0.021	0.213
LV	45	0.061	102	0.882	0.199
NL	48	0.038	23	0.612	0.122
PL	50	0.030	36	0.507	0.262
PT	48	0.036	79	0.443	0.114
RO	48	0.029	33	0.502	0.058
SI	49	0.102	49	0.130	0.141
SK	48	0.051	27	−0.104	0.051

Notes: Sample: 1999–2019. We compute country-level HHI by averaging sector-level HHI, weighted by employment shares. Estimated pass-through ($\hat{\gamma}$) according to equation (9).

3.2 Results

We first examine VAT changes through an event study, assigning each category-country pair (i, c) to either "high concentration" or "low concentration" groups based on whether its median concentration exceeds the sample median.

Extending the setup by Benzarti, Carloni, et al., 2020, we estimate:

$$\begin{aligned}\Delta \log(p_{ict}) = \alpha + &\sum_{j=-6}^{6} \beta_{1j}\, \Delta \log(1 + \tau_{ic,t+j}) \cdot \mathbb{1}(\text{HHI}_{ic} < \text{HHI}_{\text{Median}}) \\ &+ \sum_{j=-6}^{6} \beta_{2j}\, \Delta \log(1 + \tau_{ic,t+j}) \cdot \mathbb{1}(\text{HHI}_{ic} > \text{HHI}_{\text{Median}}) + \kappa_{ct} + \eta_{it} + \gamma_{ci} + \epsilon_{ict},\end{aligned} \tag{7}$$

where $\Delta \log(p_{ict})$ is the change in log (after-tax) price of category i in country c in month t, $\Delta \log(1 + \tau_{ic,t+j})$ is the change in the log tax rate, κ_{ct}, η_{it}, and γ_{ci} are country-month, category-month, and country-category fixed effects. The sets of coefficients β_{1j} and β_{2j} estimate the dynamic pass-through separately for low-concentration and high-concentration markets. Standard errors are clustered at the country-category level.

Identification As common within the literature, we have no quasi-random variation in tax rates. Instead, the inclusion of the saturated set of two-way fixed effects ζ_{ct}, η_{it}, and γ_{ci} implies that our estimates stem from simultaneous heterogeneity across all three dimensions. For instance, the country-month fixed effects ζ_{ct} absorb all aggregate shocks to country c in month t (e.g., macroeconomic conditions, exchange rate movements, or concurrent policy changes), while the category-month fixed effects η_{it} control for global trends affecting category i in month t (e.g., worldwide demand shifts or supply chain disruptions). The country-category fixed effects γ_{ci} absorb time-invariant differences across markets, such as persistent differences in market structure, consumer preferences, or baseline regulatory environments. Our identifying variation therefore comes from the differential timing of tax changes across country-category pairs, conditional on these time-varying and cross-sectional controls. In other words, we compare the price response in a given country-category pair when it experiences a tax change to the counterfactual price evolution implied by (i) aggregate conditions in that country-month, (ii) global trends in that category-month, and (iii) the fixed characteristics of that country-category market.

Figure 2: Dynamic pass-through by market concentration

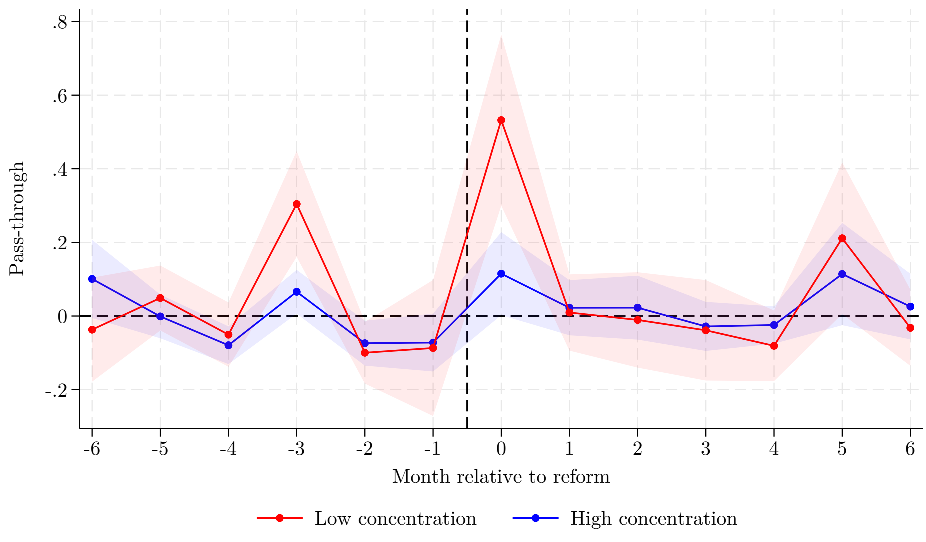

Notes: The estimated contemporaneous pass-through of a value-added tax introduced in period 0 over time – the $\beta_{i,j}$ coefficients in equation (7). The shaded area highlights the 95% confidence interval.

Figure 2 shows clear differences between concentration groups: high-concentration sectors exhibit near-zero pass-through (the contemporaneous pass-through of 10% is barely significant), while low-concentration sectors show 50% contemporaneous pass-through and 30% anticipatory pass-through three months prior to implementation.

Anticipation. Anticipatory effects are consistent with prior findings (Benedek et al., 2020). Buettner and Madzharova (2021) document that the median announcement-to-implementation period is 3 months, aligning with our observed pre-implementation effects. In an alternative specification, we split all value-added changes into those that occur during January, and those that occur during the rest of the year. We find that anticipation effects are purely driven by January changes (see Figure A1). January changes are typically legislated and announced several months in advance, giving firms time to adjust prices before the policy takes effect.

To isolate the role of market concentration in pass-through heterogeneity, we focus hereinafter on contemporaneous effects and do not attempt to quantify the full dynamic pass-through. To examine the precise relationship across the full distribution of market structures, we estimate:

$$\begin{aligned}\Delta \log(p_{ict}) =& \alpha + \beta \Delta \log(1+\tau_{ict}) + \gamma \Delta \log(1+\tau_{ict}) \cdot \mathrm{HHI}_{ict} + \delta\, \mathrm{HHI}_{ict} \\ &+ \zeta_{ct} + \eta_{it} + \gamma_{ci} + \epsilon_{ict},\end{aligned} \tag{8}$$

where γ captures differential pass-through by concentration. We standardize the HHI using the entire sample to render estimates interpretable.[4]

[4]The inclusion of lags and leads yields similar results.

Table 2: Pass-through of Value Added Taxes by Concentration

	(1)	(2)	(3)	(4)
Tax change	0.233***	0.100	0.242***	0.248***
	(0.054)	(0.087)	(0.075)	(0.075)
Market concentration (std) × Tax change	-0.102*	-0.330***	-0.255***	-0.260***
	(0.053)	(0.092)	(0.085)	(0.085)
Market concentration (std)	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Time×Country FE	No	Yes	Yes	Yes
Time×Product FE	No	No	Yes	Yes
Country×Product FE	No	No	No	Yes
Observations	149983	149983	149312	149312
R-squared	.001	.050	.355	.359

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses and clustered at the category-country level. The table does not report the constant term. We report under "Tax change" the contemporaneous effect of changes in taxes on changes in prices. Results are virtually identical when including lagging and leading tax changes. Market concentration is measured using HHI and standardized according to the average within-country standard deviation of HHI.

Table 2 builds the results by adding controls step-by-step. Column (4) presents our main specification: pass-through at mean concentration is 24.8 percentage points. A one-standard-deviation increase in concentration reduces pass-through by 26 percentage points—more than eliminating baseline transmission entirely. This reconciles with our event study: high and low concentration groups differ by 1.4 standard deviations, yielding a predicted 36pp difference, close to the observed 40pp gap.

Market concentration varies much more across countries than across industries: cross-country variation explains 71.6% of the variance in HHI.[5] This motivates the final part of our analysis: how much do cross-country differences in market concentration explain cross-country pass-through variation?

[5] When repeating the previous exercise after standardizing the HHI using only within-country dispersion, we find that the above- and below-median concentration groups differ by 4.3 within-country standard deviations. One within-country standard deviation increase in market concentration lowers the pass-through by 8.8 percentage points.

3.3 Cross-country analysis

We now demonstrate that country-level concentration differences explain a significant share of cross-country pass-through variation. We proceed in two steps. First, we estimate average contemporaneous effects for each country, controlling for industry-time fixed effects:

$$\Delta \log(p_{ict}) = \alpha + \sum_{c'} \gamma_{c'} \cdot \mathbb{1}(c = c') \cdot \Delta \log(1 + \tau_{ict}) + \eta_{it} + \epsilon_{ict}, \tag{9}$$

where c' is the country index and $\mathbb{1}(c = c')$ indicates whether an observation belongs to country c'.

Table 1 reports the first-stage coefficients and standard errors for each country. Despite some imprecise estimates, we find economically and statistically significant differences in pass-through rates across most countries. The first-stage estimates exhibit heterogeneous precision: while most countries have tightly estimated coefficients, a few (most notably France and Italy) have considerably larger standard errors.

In the second stage, we regress each country's average concentration level—computed as the employment-weighted average of industry-level HHIs—on the estimated country-specific coefficient $\hat{\gamma}_c$. We address this heterogeneous precision of our estimated $\hat{\gamma}_c$ in two ways. First, we weight each observation inversely by the variance of its estimate in an otherwise standard OLS regression.[6] This accounts for differences in statistical precision but remains sensitive to outliers. Second, by employing a Huber-weighted regression, which limits the influence of outliers without excluding them from the regression.[7]

Figure 3 shows the results. Country-level pass-through declines significantly with market concentration. For instance, moving from Portugal's HHI of 0.03 to Spain's HHI of 0.07 corresponds to a decline in estimated pass-through from 0.50 to 0.12. The precision-weighted OLS yields a p-value of 0.09, while the Huber-robust regression yields a p-value of 0.03, suggesting that the relationship is not driven by outliers. We find an R^2 of 0.16, indicating that concentration heterogeneity explains a substantial

[6] This approach is standard in two-stage estimation procedures where first-stage estimates have heterogeneous precision. This is called "inverse variance weighting" because it optimally reflects the amount of information each estimate contains (Burke, Ensor, and Riley, 2017; Higgins et al., 2019).

[7] The Huber regression uses an M-estimator that applies least squares weights to observations with small residuals and downweights those with large residuals, providing robustness without the complete exclusion of influential points.

Figure 3: Country-level contemporaneous pass-through and market power

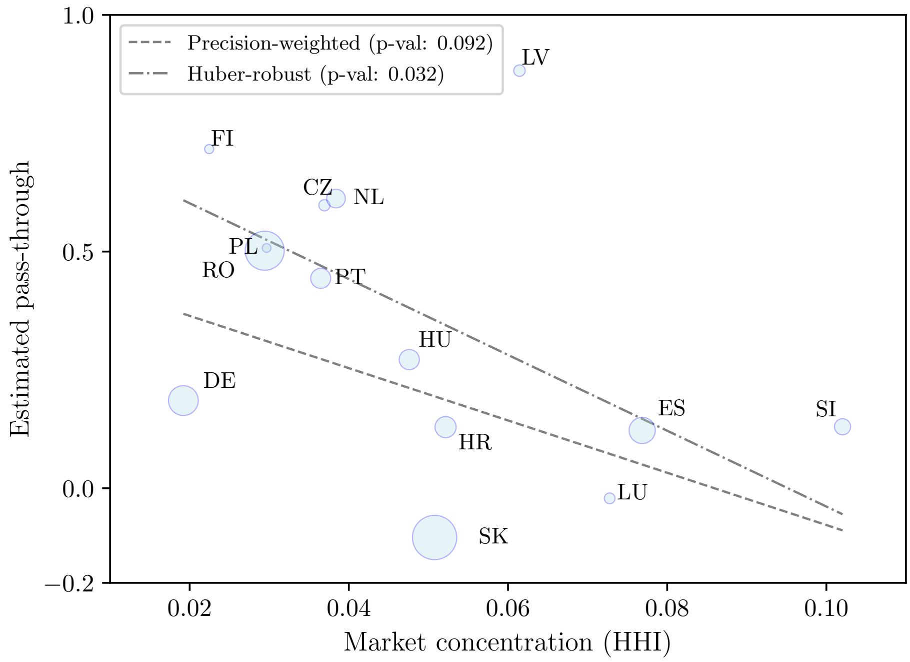

Notes: The figure plots country-level average HHI (employment-weighted) against contemporaneous pass-through estimates. Circle size is inversely proportional to the variance of the first-stage estimate and thus reflects its precision. The dashed line shows a precision-weighted regression (p=0.092, R^2=0.16), while the dash-dotted line shows a Huber-weighted robust regression (p=0.032). Italy and France, which have imprecise first-stage estimates, are included in the regression analysis (Table 1) but excluded from this figure for clarity.

share of cross-country pass-through variation. We consider this explanatory power large, when taking into account that VAT changes differ across countries in many other aspects (sign, permanence, anticipation, etc.).

4 Conclusion

This paper demonstrates that market concentration is a key determinant of VAT pass-through to consumer prices. Using data from 16 European countries over 1999-2019, we show that low-concentration industries pass through 50% of VAT changes to consumers contemporaneously, while high-concentration industries exhibit near-zero transmission. Since market concentration varies primarily across countries rather than industries—with cross-country differences explaining 72% of concentration variation—these differences in market structure account for 16% of the observed cross-country heterogeneity in pass-through rates.

Our findings help resolve the empirical puzzle of heterogeneous VAT pass-through by identifying market structure as one key mediating factor. The policy implications are clear: VAT cuts designed for consumer relief work best in competitive industries, while VAT increases in concentrated sectors raise revenue with minimal consumer burden. This suggests countries should consider their market structure when designing VAT policy rather than applying uniform rates.

Our results also question EU attempts to harmonize VAT rates. Since cross-country concentration differences explain a significant share of pass-through variation, optimal policy would allow countries with competitive markets to set lower rates for consumer benefit while permitting higher rates where firms absorb the burden. The welfare implications are substantial: VAT reductions deliver meaningful consumer savings in competitive sectors but provide minimal consumer benefit in concentrated ones.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the authors used Artificial Intelligence (Claude) in order to proofread the text. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

Anderson, Simon P., André de Palma, and Jacques-François Thisse (1992). *Discrete Choice Theory of Product Differentiation*. Cambridge, MA: MIT Press.

Bellon, M., A. Copestake, and W. Zhang (2024). "Supply and Demand Determinants of Heterogeneous VAT Pass-Through". In: *SSRN*.

Benedek, D. et al. (2020). "Varieties of VAT Pass-Through". In: *International Tax and Public Finance* 27, pp. 890–930.

Benzarti, Y., D. Carloni, et al. (2020). "What Goes Up May Not Come Down: Asymmetric Incidence of Value-Added Taxes". In: *Journal of Political Economy* 128.12, pp. 4438–74.

Benzarti, Y. and A. Tazhitdinova (2021). "Do Value-Added Taxes Affect International Trade Flows? Evidence from 30 Years of Tax Reforms". In: *American Economic Journal: Economic Policy* 13.4, pp. 469–489.

Benzarti, Youssef (2025). "Tax incidence anomalies". In: *Annual Review of Economics* 17.

Bernardino, Tiago et al. (2025). "The full, persistent, and symmetric pass-through of a temporary VAT cut". In: *Journal of Public Economics* 248, p. 105416.

Buettner, Thiess and Boryana Madzharova (2021). "Unit Sales and Price Effects of Preannounced Consumption Tax Reforms: Micro-level Evidence from European VAT". In: *American Economic Journal: Economic Policy* 13.3, pp. 103–134.

Burke, Danielle L, Joie Ensor, and Richard D Riley (2017). "Meta-analysis using individual participant data: one-stage and two-stage approaches, and why they may differ". In: *Statistics in Medicine* 36.5, pp. 855–875.

Caplin, Andrew and Barry Nalebuff (1991). "Aggregation and Imperfect Competition: On the Existence of Equilibrium". In: *Econometrica* 59.1, pp. 25–59.

Carbonnier, C. (2007). "Who Pays Sales Taxes? Evidence from French VAT Reforms, 1987–1999". In: *Journal of Public Economics* 91.5-6, pp. 1219–1229.

Dierickx, I., C. Matutes, and D. Neven (1988). "Indirect taxation and cournot equilibrium". In: *International Journal of Industrial Organization* 6.3, pp. 385–399. ISSN: 0167-7187.

Dimitrakopoulou, Lydia et al. (2024). "VAT pass-through and competition: evidence from the Greek Islands". In: *International Journal of Industrial Organization* 97, p. 103110.

Fuest, C., F. Neumeier, and D. Stöhlker (2024). "The Pass-Through of Temporary VAT Rate Cuts: Evidence from German Supermarket Retail". In: *International Tax and Public Finance*, pp. 1–47.

Gaarder, Ingvil (2019). "Incidence and Distributional Effects of Value-Added Taxes". In: *The Economic Journal* 129.618, pp. 853–876.

Harju, Jarkko, Tuomas Kosonen, and Oskar Nordström Skans (2018). "Firm Types, Price-Setting Strategies, and Consumption-Tax Incidence". In: *Journal of Public Economics* 165, pp. 48–72.

Higgins, Julian PT et al. (2019). "Analysing data and undertaking meta-analyses". In: *Cochrane handbook for systematic reviews of interventions*. John Wiley & Sons. Chap. 10, pp. 241–284.

Kosonen, T. (2015). "More and Cheaper Haircuts after VAT cut? On the Efficiency and Incidence of Service Sector Consumption Taxes". In: *Journal of Public Economics* 131, pp. 87–100.

Rozema, Kyle (2018). "Tax Incidence in a Vertical Supply Chain: Evidence from Cigarette Wholesale Prices". In: *National Tax Journal* 71.3, pp. 427–450.

Weyl, E. Glen and Michal Fabinger (2013). "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition". In: *Journal of Political Economy* 121.3, pp. 528–583.

A Appendix Figures

Figure A1: Dynamic pass-through: VAT changes in January versus other months

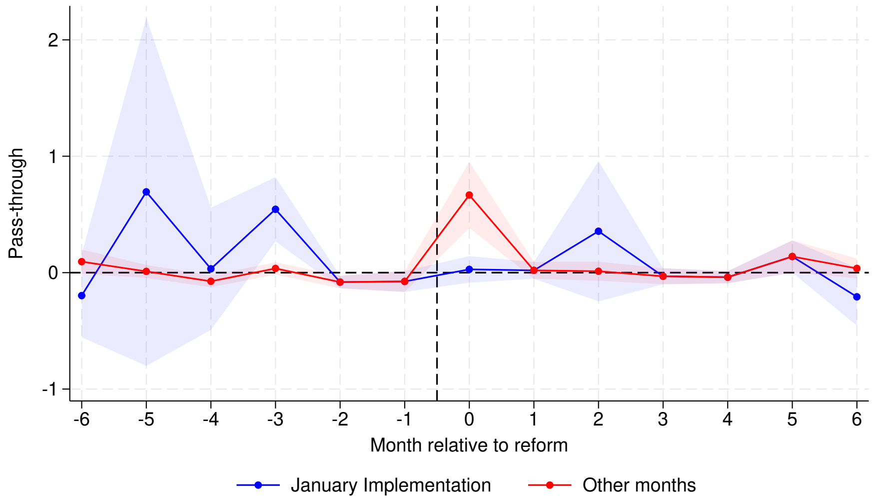

Notes: We classify as January adjustments tax changes that occur during January, and estimate (7) separately for January adjustments and adjustments in other calendar months. The figure plots the $\beta_{i,j}$ coefficients, the shaded area highlights the 95% confidence interval. We note that anticipation effects only occur for January adjustments, and non-January adjustments only trigger contemporaneous in prices.