

Practical 4: Markups, taxes, and forty years of the cola war

Practical session, Day 4 — data, code and regressions

Dr. Ernil Sabaj

Warwick Summer School
International Business and Finance

24 July 2026

The last hour. Three measurements the textbook does not contain.

- **1. Markups:** the production approach, the raw ratio, and why the level is contested but the trend is not.
- **2. Transfer pricing, measured:** same-HQ firms, haven vs domestic incorporation, a 15-point gap in the **ETR** (effective tax rate = tax paid / pre-tax profit).
- **3. The cola war:** the repeated game in the accounts, plus an honestly underpowered reaction function.

State what would falsify your claim before you look at the estimate.

The plan: from this morning's slides to tests

Claim / diagram from the lecture	Data	What we estimate	The claim is in trouble if...
Monopoly pricing: $\frac{P-MC}{P} = \frac{1}{\varepsilon}$ (pricing rules)	Sales/COGS, 45 years	Markup level, trend, dispersion (production approach)	levels were measurable from accounts alone. (They are not; trends are.)
Transfer pricing shifts profit to low-tax homes (TP slides)	ETR of US-HQ firms by country of incorporation	Median ETR: US- vs haven-incorporated	the gap is small or wrong-signed. (It is 15pp, $t = 4.7$.)
Advertising as repeated prisoner's dilemma (pay-off matrices)	Coke & Pepsi accounts, 1985–2024	40 years of SG&A and margins; a reaction function	either firm ever defected to 'don't advertise'. (Neither did.)
Tit-for-tat is detectable (repeated games)	Same, first differences	$\Delta\text{Coke on } \Delta\text{Pepsi: } \beta = 0.36, t = 0.96, N = 39$	— a lesson in power, not in cola

Everything is pre-computed on the slides; every number re-runs from replication/ in class.

The data this hour

File	stata/ib_panel.dta (same panel as all week).
Markups	sale/cogs, all firm-years 1980–2024; the θ -sensitivity table varies the output elasticity.
ETR	US-HQ firms; txt and pretax cumulated over FY2015–24, <i>profit years only</i> , cumulative pre-tax > \$50m; groups by country of incorporation (fic): $n = 3,576$ US / 50 haven / 44 other.
Cola	Coca-Cola and PepsiCo annual SG&A, first-differenced: $N = 39$ year-pairs.

Three measurements, one file — and each one's sample restrictions do half the work.

One file, same as all week

Everything this hour comes from `ib_panel.dta` — one row per firm per year, 441,198 rows, dead firms included, Monday's guards throughout.

- **Markups** use two columns you met on Monday: `sale` (revenue) and `cogs` (cost of goods sold — what it cost to make what was sold).
- **The tax comparison** uses two new columns: `txt` (the year's tax expense) and `pretax` (profit before tax) — plus `fic`, the country a firm is *legally registered* in, vs `hqloc`, where it is actually *run* from. Monday's nationality lesson becomes today's research design.
- **The cola war** uses `sga` — selling, general & administrative expense, the accounting bucket that contains advertising.

Software: Stata on the slides; identical Python in `replication/student_pack_day4/`. (No World Bank data today.)

What we do to those rows today

Three questions, three different shapes of the same file.

firm	fyear	sale (\$m)	cogs (\$m)	markup
APPLE INC	2024	391,035	198,907	1.97

One row like that is all a markup is: $391,035/198,907 = 1.97$ — in 2024 Apple charged about twice what the goods cost to make.

- **Markups (part A):** keep *every* row — one markup per firm per year.
- **Tax rates (part B):** *add up* each firm's rows over 2015–24 — one number per firm, so a single odd year cannot distort it.
- **The cola war (part C):** keep just *two* firms' rows, over forty years.

From this morning's pricing rules to a measurable object

Monopoly pricing (this morning): $\frac{P - MC}{P} = \frac{1}{\varepsilon}$ — the **Lerner index** of market power, unobservable without MC .

The production approach (De Loecker–Warzynski): cost minimisation over a flexible input V implies


$$\mu \equiv \frac{P}{MC} = \theta_V \cdot \frac{P Q}{P_V V} = \theta_V \cdot \frac{\text{Sales}}{\text{COGS}},$$

where θ_V is the **output elasticity** of the input — in plain words, *how much extra output one extra unit of that input buys*. It cannot be read off any balance sheet; that is the whole difficulty.

- Set $\theta_V = 1$: the raw **Sales/COGS** ratio you can compute in one line. De Loecker, Eeckhout & Unger — “**DEU**” in the tables that follow — instead *estimate* θ for each industry and year from production data.
- Same accounting object, different scaling. The **trend** is robust to the choice; the **level** is not — that is the next twenty minutes.

The modern measure: markups

Markups of US-listed non-financial firms, 1980-2025



This morning's Lerner index defined market power as $(P - MC)/P$ — price minus marginal cost, over price. You cannot see P or MC in company accounts. But under a mild assumption (firms don't deliberately waste inputs), **sales ÷ cost of goods sold** tracks the same object — that ratio is what the chart shows. Source: own calculations, Compustat quarterly, US-listed non-financial firms, ratios winsorised to [0.5, 10].

Read that chart carefully

- The **typical firm** barely moved: median $1.37 \rightarrow 1.64$ over 45 years.
- The **90th percentile firm** doubled its markup: $2.00 \rightarrow 3.98$.
- The **sales-weighted mean** rose $1.37 \rightarrow 1.78$ — it moves with the tail because the tail firms are large.

“Markups rose” is true.
“Firms raised their markups” is false.
A few firms did, and they got big.

This reallocation — market share shifting toward high-markup firms — is the central finding of De Loecker, Eeckhout & Unger (2020), who estimate the aggregate markup rose from **1.21 in 1980 to 1.61 in 2016**.

Transfer pricing: from anecdote to a difference in means


This morning's transfer-pricing slides asserted that profit moves to low-tax jurisdictions. Testable implication: among firms **headquartered in the same country**, those **incorporated in havens** should pay lower effective tax rates.

$$ETR_i = \frac{\sum_{t=2015}^{2024} \text{tax expense}_{it}}{\sum_t \text{pre-tax income}_{it}}$$

(Cumulate first: year-by-year ETRs explode when income passes through zero.) **Design:** US-headquartered firms, positive cumulative pre-tax income > \$50m, grouped by country of incorporation (`fic`): United States / tax havens (CYM, BMU, VGB, IRL, JEY, LUX, MHL, ...) / other foreign.

Prediction: $ETR_{\text{haven}} < ETR_{\text{US}}$. Let's see by how much.


The answer: fifteen points



- Median effective tax rate: US-incorporated **21.3%**, haven-incorporated **6.3%** — a gap of roughly **15 points**.
- Is that just luck? A **t-test** asks whether a difference between two group averages is bigger than chance would produce. Here it gives $t = 4.7$: far too big to be chance. (We use *Welch's* version, which allows the two groups to be differently spread — essential when one has 3,576 firms and the other 50.)

Own calculations, Compustat, FY2015–2024, US-HQ firms with cumulative pre-tax income > \$50m. $n = 3,576 / 44 / 50$.

The same comparison, as a picture: every firm is a dot



How to read it: one dot = one firm, placed by its effective tax rate over 2015–24 ($\text{tax} \div \text{pre-tax profit}$); the black marks are the group medians. **Why some dots sit below zero:** a firm can finish the decade having booked a net tax *credit* rather than a tax bill — losses set against earlier profits, research credits, a settled dispute. Those are real firms, not errors (the sample is trimmed to -20% to $+80\%$). **And note the crowd:** 50 red dots against 3,576 navy — five unusual haven firms could move that red median.

Compustat via `ib_panel.dta`; sample exactly as in `04_practical_regressions.do`, Part D. Regenerate: `python/06_new_lecture_figures.py`.

Before you tweet that number

- ❶ **The firms chose their address; we did not assign it.** Companies decide whether to register in a haven, and the ones that bother are those whose profits are easiest to move — firms built on patents and brands rather than factories. So the 15-point gap mixes together *what the address does* and *what kind of firm picks it*. We cannot separate the two here.
- ❷ **Only 50 firms are in the haven group.** The difference is real, but a handful of very large firms carry it. Look at how many observations you have before you trust a comparison.
- ❸ **The accounts show tax charged, not tax paid.** The two differ because some tax is deferred to later years. A cash answer would need the cash-flow statement.
- ❹ **The rules changed during our window.** The 2017 US reform cut the headline rate from 35% to 21%, and the OECD's 15% global minimum starts in 2024 — both shrink this gap in future. Our number is a ten-year average over a moving rulebook.

Still: a 3× difference in tax rates for the same-country firms
is the transfer-pricing lecture, written in the accounts.

The cola war is a repeated game we can actually watch



The advertising game we just drew has a familiar structure:

		PepsiCo	
		Advertise	Don't
Coca-Cola	Advertise	(2, 2)	(4, 1)
	Don't	(1, 4)	(3, 3)

- Both advertise. Shares barely move. Both would be better off if neither did.
- **(Advertise, Advertise)** is the unique Nash equilibrium, and it is Pareto-dominated.
- Coke and Pepsi have played this game, quarterly, since the 1890s.
Forty years of it are in the accounts. Let's look.

Forty years, and neither one blinked

The cola war: forty years of the same prisoner's dilemma



Source: own calculations, Compustat quarterly, fiscal years 1985–2024. SG&A (xsgaq) and operating income after depreciation (oiadpq) as % of sales, 4-quarter fiscal years.

What the data says

	1985		2024	
	Coca-Cola	PepsiCo	Coca-Cola	PepsiCo
Sales (\$bn)	7.9	7.7	47.1	91.9
SG&A (% of sales)	33.7	49.1	31.1	39.1
Operating margin	13.2	9.7	30.0	15.8
Markup (sales/COGS)	1.9	2.4	2.7	2.4

- In 1985 they were **the same size**. Today Pepsi's revenue is **twice** Coke's.
- Yet Coke's operating margin is **double** Pepsi's — the gap widened from 3.5 points to 14.2.
- Both spent roughly **37–39% of sales** on SG&A for three decades straight. **Neither ever defected to “don't advertise”**.
- Coke's markup rose 1.9 → 2.7. **Pepsi's did not move at all.**

Back to this morning's game: advertising is the classic prisoner's dilemma — both would earn more if *neither* advertised heavily, yet each does better by advertising, so both do. The table is that game *repeated*: forty years holding at 37–39% of sales, neither cutting, neither escalating — tit-for-tat, visible in the accounts.

Before you believe me — three objections

- ❶ **They are not the same business.** Coke mostly sells *concentrate* to independent bottlers; Pepsi owns bottling *and* Frito-Lay, which is roughly half its revenue. Snacks carry lower margins than syrup.
Some of that 14-point gap is business model, not strategy.
- ❷ **SG&A is not advertising.** It is advertising *plus* sales force, distribution, admin, legal. Compustat's advertising item (xad) is not in our extract, and most firms stopped disclosing it after 1994 anyway.
We are using a proxy. Say so.
- ❸ **Coke's 2017 drop is an accounting event.** It refranchised its bottling operations — lower revenue, higher margin, same underlying business.
Look at the level shift around 2017 in the left panel. That is a reorganisation, not a strategic choice about advertising.

None of this kills the point. All of it changes how hard you can push it.

A reaction function for the cola war — and a lesson in small samples

A **reaction function** asks: when one firm moves, does the other follow? Regress the *change* in Coke's advertising ratio on the *change* in Pepsi's:

$$\Delta\left(\frac{\text{SG\&A}}{\text{Sales}}\right)_{\text{Coke},t} = \alpha + \beta \Delta\left(\frac{\text{SG\&A}}{\text{Sales}}\right)_{\text{Pepsi},t} + \varepsilon_t$$

$\hat{\beta}$	+0.355 (0.368)
t	0.96
N (years, 1986–2024)	39
R^2	0.14

- $\hat{\beta} = +0.36$: when Pepsi raises its advertising ratio by 1 percentage point, Coke raises its own by about 0.36 — the direction expected. But the $\pm$ (0.37) is as big as the estimate, so it is **nowhere near significant**.
- A “no result”, handled honestly. The theory is *not* rejected: with 39 yearly observations this test could only have detected a response above about **0.7**, so a sensible tit-for-tat of 0.3–0.5 was invisible from the start. The honest sentence is “*this data cannot tell moderate interaction apart from none*” — always ask what size of effect your test could have found before reading anything into a null.

Discussion: whose payoff matrix was right?

Three questions for the room — there is no single right answer, which is the point.

- **Did the game turn out symmetric?** The prisoner's dilemma we drew this morning gives both players the *same* mediocre payoff. But the accounts show Coke with double the margin and Pepsi with double the sales. So did they really play the same game — or was one of them playing for size and the other for margin?
- **Why has neither ever stopped advertising?** If heavy advertising is genuinely a trap that hurts both, a hundred years of repeated play should have let them find their way out of it. They never did. What does that tell us — that cooperation is harder than the model suggests, or that the model has the payoffs wrong?
- **Or is the advertising not wasteful at all?** If it is what builds the brand that lets Coke charge 2.7 times its costs for sugared water, then advertising is not a trap — it is an *investment*, and calling it a prisoner's dilemma is simply the wrong frame.

Links back to Wednesday (is heavy advertising a barrier to entry?) and to Monday (is the Coke brand a VRIO resource — valuable, rare, hard to imitate, organised for?).

- De Loecker, J. & Warzynski, F. (2012), 'Markups and Firm-Level Export Status', *AER* 102(6), 2437–2471.
- De Loecker, J., Eeckhout, J. & Unger, G. (2020), 'The Rise of Market Power and the Macroeconomic Implications', *QJE* 135(2), 561–644.
- De Ridder, M., Grassi, B. & Morzenti, G. (2026), 'The Hitchhiker's Guide to Markup Estimation', *Econometrica* 94(1), 137–168.
- Bresnahan, T. & Lau, L. — markup identification; see Bresnahan (1989), *Handbook of IO*, Vol. 2.
- Zucman, G. (2014), 'Taxing across Borders: Tracking Personal Wealth and Corporate Profits', *JEP* 28(4), 121–148.
- Tørsløv, T., Wier, L. & Zucman, G. (2023), 'The Missing Profits of Nations', *Review of Economic Studies* 90(3), 1499–1534.
- OECD Pillar Two (global minimum tax), in force from 2024.
- Axelrod (1984); Schelling (1960) — repeated-game cooperation and commitment.
- Welch, B.L. (1947), 'The generalization of "Student's" problem when several different population variances are involved', *Biometrika* 34(1–2), 28–35.