

Practical 3: Concentration — the theorem, the measurement, and the regression that fails

Practical session, Day 3 — data, code and regressions

Dr. Ernil Sabaj

Warwick Summer School
International Business and Finance

23 July 2026

This morning gave you the **HHI** (Herfindahl–Hirschman index: the sum of squared market shares, 0–10,000) and its thresholds. This hour: the Cournot theorem behind it, the honest computation, and the regression the theorem implies — which fails.

- 1. Why the HHI was ever taken to measure market power. The theory says: *average markup* = $\text{HHI} / \text{how price-sensitive customers are}$. (The markup measure is the **Lerner index** = $(\text{price} - \text{marginal cost})/\text{price}$; price-sensitivity is the **elasticity of demand**, ε ; and **Cournot** is the standard model of firms competing by choosing quantities.)
- 2. Compute it: half of US industries breach the “highly concentrated” line, every year since 1980.
- 3. Test it: **SCP** (structure–conduct–performance — the old claim that concentrated *structure* causes high-margin *performance*) in four columns, then with firm fixed effects. $R^2 = 0.000$, then negative.

Porter tells you which room you are in. The data says the room explains 15%.

The plan: from this morning's slides to tests

Claim / diagram from the lecture	Data	What we estimate	The claim is in trouble if...
HHI thresholds: <1500 competitive, >2500 concentrated (HHI slide)	All listed non-financials, by SIC-3 industry	The full HHI distribution, 1980–2024	a sane share of industries breach 2500. (Half do, every year.)
Cournot: average markup = HHI/ε (today's derivation)	Margins + HHI panel	The SCP regression the theorem implies, 4 specifications	$\beta_{\text{HHI}} \leq 0$ — which is what we find, incl. within-firm
Five forces $\Rightarrow$ industry determines profit (semi-conductor diagram)	ROA variation, 1995–2024	Industry vs firm variance decomposition	firm effects dwarf industry effects (they do: 61% vs 15%)
Generic strategies: differentiate or cost-lead (Porter triangle)	Gross margin $\times$ asset turnover, 2024	The margin–turnover frontier	no trade-off appears (corr = -0.37 : it does)

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

The data this hour

File	stata/ib_panel.dta — 441,198 firm-years, 38,954 US-listed firms, 1980–2025.
Unit today	the industry-year : firm sales summed to SIC-3 $\times$ year cells (280 industries) to build HHI. SIC-3 = Standard Industrial Classification, 3-digit: the official code that sorts firms into industries; then back to firm-years for the SCP and within-firm regressions.
Variables	sale (to build the HHI shares), op_margin (the outcome, y), hhi (the explanator, x), controls lnsale and rd_int; fixed effects (a separate dummy for each) on firm gvkey, year fyear and sector gsector.
Caveat	listed firms only — HHI here is concentration <i>among listed firms</i> , not the census HHI. That truncation is itself on the syllabus.

Every regression this hour states its cluster level — watch for it.

One file, same as Monday

Everything this hour comes from `ib_panel.dta` — the panel you already know:

- One row = one firm in one year; 441,198 rows; every firm ever listed on a US exchange, dead ones included.
- Monday's guards still apply throughout: `full_year==1`, margins kept in $[-1, 1]$, and we always say which filters a result uses.
- What is *new* today is the **unit of analysis**: we will repeatedly *add firms up into industries* (sum sales within an industry-year, compute shares) and then drop back down to firm-level regressions. Watch for that switch — it decides what a number means.

Software: Stata on the slides; identical Python in `replication/student_pack_day3/`.
(No World Bank data today.)

Recap: the dataset (same as Monday)

One row = one firm-year. 441,198 rows, dead firms included.

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

- Today we **aggregate rows into industry-years**: sum sales within SIC-3 $\times$ year, then square the shares.
- Monday's guards (winsorisation, full_year, assets > \$10m) still apply throughout.

Why anyone ever believed HHI measures market power

Not superstition — a theorem. The two terms again: **Cournot competition** = firms choose *how much to produce*, and the price follows from total output. The **Lerner index** = $(P - MC)/P$: the margin on the last unit sold as a share of the price — the cleanest measure of “market power”. In that model firm i ’s profit-maximising choice implies:

$$\frac{P - MC_i}{P} = \frac{s_i}{\varepsilon}.$$

Multiply by s_i and sum over firms — the **share-weighted industry Lerner index** is

$$\sum_i s_i \frac{P - MC_i}{P} = \frac{\sum_i s_i^2}{\varepsilon} = \frac{\text{HHI}}{\varepsilon}.$$

- In words: **under Cournot the HHI is not folklore — it equals the industry’s average profit margin, scaled by how price-sensitive customers are (ε)**. That equation launched forty years of Structure–Conduct–Performance regressions.
- Every assumption is load-bearing: quantity competition, one identical good, a correctly-defined market, a common ε . Break one and the mapping dissolves.

Today: compute the HHI honestly, run the SCP regression it implies, and watch what happens.

Let's actually compute that index

The previous slide gave you the formula and the thresholds:

$$\text{HHI} = \sum_{i=1}^N s_i^2 \quad \left\{ \begin{array}{ll} < 1500 & \text{competitive} \\ 1500\text{--}2500 & \text{moderately concentrated} \\ > 2500 & \text{highly concentrated} \end{array} \right.$$

So let's do it. Every US-listed non-financial firm, grouped into 3-digit SIC industries, sales shares, 2024. Industries with at least five listed firms.

What share of US industries do you expect to be
“highly concentrated”?


Computing HHI from the panel (Stata)

```
use ib_panel, clear                // open the panel
keep if full_year==1 & sale>0 & !missing(sic3) // the usual guards
drop if inlist(real(gsector),40,60) // drop financials & real estate
egen indid = group(sic3)           // number the industries

bysort indid fyear: egen ind_sales = total(sale)
                        // within each industry-year: add up all sales
gen share = sale/ind_sales         // each firm's market share
gen share2 = share^2              // square it...
bysort indid fyear: egen hhi = total(share2)
                        // ...and sum the squares: that IS the HHI
replace hhi = hhi*10000           // conventional 0-10,000 scale
bysort indid fyear: gen nfirms = _N // count firms per cell
keep if nfirms >= 5              // keep cells with enough firms
```

Five lines of arithmetic — and three research decisions hiding in them: SIC-3 as the “market”, listed firms as the universe, a 5-firm minimum. Keep such choices visible in your own work.

Applying the textbook thresholds to Compustat: everything is a monopoly




And it has been about half in **every year since 1980**. If the rule were right, the US economy has been a collection of monopolies for forty-five years. Source: own calculations, Compustat quarterly, fiscal 2024. 158 SIC-3 industries with ≥ 5 listed firms.

Three reasons the number is wrong

- 1 **Compustat only sees listed firms.** SIC 573 (computer stores) looks concentrated because Best Buy is listed and ten thousand private dealers are not. The denominator is missing.
- 2 **An SIC code is not a market.** SIC 367 lumps Nvidia's data-centre GPUs with Micron's memory and Jabil's contract assembly. Nobody substitutes between them.
The HHI of a non-market is a non-number.
- 3 **The thresholds are for antitrust markets, not economies.** The US competition authorities (the DOJ, Department of Justice, and FTC, Federal Trade Commission) apply the HHI to a market defined by the **hypothetical monopolist test** — usually local, usually narrow. Cement in one metro area. Not “chemicals, nationally”.

The formula is not the problem.
The market definition is always the problem.

So is the measure useless? No — use it in differences



The **level** is contaminated by everything on the last slide. The **change** is much more informative, because the measurement error is broadly stable over time. Source: own calculations, Compustat quarterly. SIC-3 industries with ≥ 5 listed firms.

US industries did concentrate — and the debate about why is live

Industries observed in both 1997 and 2024	156
Share whose HHI rose	72.4%
Mean HHI, 1997 → 2024	2,116 → 2,881
Median change	+649 points

- This replicates Grullon, Larkin & Michaely (2019), who report that over 75% of US industries became more concentrated.
- **Two readings, and they have opposite policy implications:**
 - **Market power** — weak antitrust let incumbents entrench (Philippon, *The Great Reversal*).
 - **Superstar firms** — technology lets the most productive firms serve the whole market. Concentration is efficiency showing up (Autor, Dorn, Katz, Patterson & Van Reenen, 2020).
- Same chart. **Same coefficient. Opposite conclusion.** You cannot settle this with a concentration ratio — which is the whole reason the field moved on to markups.

Making the five forces measurable


Porter's model is a checklist until you attach a number to each force. Here is one way, using data you can pull yourself:

Force	A measurable proxy	Semiconductors, 2024
Rivalry	HHI of sales; dispersion of margins within the industry	HHI = 591 (looks competitive); margins range -9% to +62%
Threat of entry	New listings per year; capital intensity; R&D as % of sales	R&D = 12.1% of sales ; a leading-edge fab costs \$20bn+
Substitutes	Cross-industry margin correlation; product life-cycle length	Weak — no substitute for compute
Buyer power	Customer concentration (10-K disclosures); gross margin	Nvidia's top customers are a handful of hyperscalers
Supplier power	Supplier's margin vs yours; single-source inputs	ASML: sole supplier of EUV lithography

Source: own calculations, Compustat quarterly, SIC 367x, fiscal 2024, 154 US-listed firms, combined sales \$802bn.

The same industry, two very different lives

Semiconductors (SIC 367x), 2024: HHI = 591, i.e. "unconcentrated"



Source: own calculations, Compustat quarterly, fiscal 2024. Share of combined sales of US-listed SIC 367x firms. Operating margin = operating income after depreciation / sales.

Two questions for the room

- 1 The HHI here is **591**. The rule you will meet in the next lecture says anything under 1,500 is a **competitive market**.
 - Nvidia earns a **62% operating margin** and a markup of **4.2×** cost.
 - **Is this a competitive market?**
- 2 Intel and Nvidia sit in the **same SIC code**. They face the **same five forces**.

Operating margin	2014	2019	2022	2023	2024
Intel	28.0%	31.2%	4.4%	0.1%	-8.9%
Nvidia	16.2%	26.2%	20.9%	54.1%	62.4%

- In 2014 Intel was the more profitable firm. **What does industry structure explain here?**

Industry or firm? Decomposing profitability

Take every US-listed non-financial firm, 1995–2024. Profitability (here **ROA** = Return On Assets = profit ÷ assets) varies enormously across firms; ask how much of that spread each label accounts for (“share of variance explained” = the R^2 from Monday, used as a measuring cup): **which industry you are in** versus **which firm you are**:

Share of the variation in ROA explained by	Annual	5-year averages
Year (the macro cycle)	1.4%	—
Industry (SIC-3)	15.2%	20.1%
Firm identity	60.7%	75.6%

- Firm effects **include** the firm’s industry — each firm sits in one. So what is *specific to the firm* adds roughly **45 points beyond** its industry.
- Rumelt (1991) and McGahan & Porter (1997) reached the same conclusion with different data and different methods.

Porter tells you which room you are standing in.
It does not tell you who you are.

Source: own calculations, Compustat quarterly. 138,593 firm-years, 10,190 firms, 224 SIC-3 industries; firms with ≥ 5 years, industries with ≥ 5 firms.

The SCP regression, as the theorem implies

The Cournot slide says margins should rise with HHI. So:

```
gen hhi_k = hhi/1000          // HHI in thousands: easier slopes
gen lnsale = ln(sale)         // firm size (log)
gen rd_i   = cond(missing(rd),0,rd)/sale // R&D intensity
keep if inrange(fyear,2015,2024) & sale>10 // recent decade, real firms
encode gsector, gen(sec)     // sector is text ("40.0"): make it numeric

* four versions of the same regression, adding controls each time:
eststo m1: reg op_margin hhi_k, cluster(indid)
eststo m2: reg op_margin hhi_k lnsale rd_i, cluster(indid)
eststo m3: reg op_margin hhi_k lnsale rd_i ///
           i.sec i.fyear, cluster(indid)
eststo m4: reg op_margin hhi_k share lnsale rd_i ///
           i.sec i.fyear, cluster(indid)
esttab m1 m2 m3 m4, se r2 keep(hhi_k share lnsale rd_i)
* "eststo" stores each result; "esttab" prints them side by side.
* "i.sec i.fyear" = a dummy for each sector and each year.
* cluster(indid): HHI is the SAME for every firm in an industry-year,
*   so we tell Stata industry blocks move together (next slide: why).
```

Why cluster by industry here (and how to choose in general)

The SCP regressor — HHI — is **identical for every firm in an industry-year**. Any industry-level shock therefore hits whole blocks of observations together:

- Treat the 36,431 firm-years as independent and the effective sample for the HHI coefficient is not 36,431 — it is closer to the **number of industries** (~ 200). Ignoring this makes the $\pm$ several times too small — you'd be claiming precision you don't have. (The textbook name, if you meet it: Moulton's problem.)
- Rule that generalises: **cluster at the level of the regressor's variation** (or the shock's), not the observation's. Industry-level $x \Rightarrow$ cluster by industry; policy by state $\Rightarrow$ by state.
- Corollary you can check in the do-file: clustering barely moves the *coefficients*, only the SEs — if your “result” lives or dies on the clustering choice, you don't have a result, you have a standard-error choice.

One-line version to keep: *cluster where your x varies*.

The oldest regression in industrial organisation

Structure–Conduct–Performance says concentration → market power → higher margins.
So run it. US-listed non-financials, 2015–2024, standard errors clustered by industry:

Dependent variable: operating margin	(1)	(2)	(3)	(4)
Industry HHI (per 1,000 points)	−0.0005 (t = −0.1)	−0.0133 (t = −3.1)	−0.0087 (t = −2.6)	−0.0061 (t = −1.6)
Firm's own market share				−0.123 (t = −2.1)
Log sales, R&D intensity		yes	yes	yes
Sector and year fixed effects			yes	yes
N	36,431	36,431	36,431	36,431
R ²	0.000	0.251	0.265	0.267

- **How to read the number.** In column (2), −0.0133 means: an industry that is 1,000 HHI points *more* concentrated has an operating margin about **1.3 percentage points lower**. The theory predicted *higher*.
- **What the result shows.** Column (1): concentration explains **none** of the differences in margins — the R^2 (share of the variation accounted for) is 0.000. Once we control for size and R&D the sign turns **negative**: more concentrated industries earn *lower* margins, not higher.

Statistical vs economic significance: put units on the coefficient

Column (2) says $b_{\text{HHI}} = -0.013$ per 1,000 HHI points, $t = -3.1$. Significant! And...tiny:

- A rise of **one standard deviation** in industry HHI (about 1,500 points here — a standard deviation being the typical distance from the average) moves the predicted margin by roughly **2 percentage points**, while margins themselves vary by more than 20 points. Concentration “matters” at $t = 3$ and explains almost nothing.
- Compare the same table's size effect: log sales carries R^2 from 0.000 to 0.25. **Rank regressors by variance explained and by scaled effect, never by stars.**
- Portable template for reading *any* coefficient: *a one-standard-deviation change in x moves y by $b \times sd(x)$ — now express that as a percentage of a standard deviation of y .* If a paper won't do this arithmetic, do it to the paper.

With 36,431 observations, t -statistics are cheap. Magnitudes are the argument.

The stronger test: follow the same firm as its industry concentrates

Comparing *different* industries mixes in everything else that differs between them. So compare each firm *with itself* over time — a separate dummy for every firm and every year (“fixed effects”), which removes anything about a firm that does not change:

$$\text{margin}_{it} = \beta \text{HHI}_{j(i),t} + \mu_i + \gamma_t + \varepsilon_{it}$$

	Firm FE + Year FE
Industry HHI (per 1,000 pts)	−0.0055 (t = −2.5)
N	139,908

Standard errors clustered by SIC-3 industry. Sample 1995–2024, non-financials.

- **What it shows:** −0.0055 means that when **the same firm’s** industry becomes 1,000 HHI points more concentrated, its margin falls by about **half a percentage point**. Small, but the wrong sign again — the SCP prediction fails within firms too.
- Consistent with concentration reflecting **competitive shakeout** (efficient rivals absorbing share), not softening competition.

Why SCP failed — and why that is the lesson

- 1 **Demsetz (1973) got there first.** If firms differ in efficiency, the efficient ones take share *and* earn margins. Concentration and profit are both **outcomes**. Regressing one on the other identifies nothing.
- 2 **The measurement is broken** — see the earlier “three reasons the number is wrong” slide. If the HHI is not a real number, its coefficient is not a real result.
- 3 **Firm size does the work.** Adding log sales and R&D takes R^2 from 0.000 to 0.251. **Who you are beats where you are** — the same answer as the variance table earlier this hour (firm 61%, industry 15%).

This is why researchers now measure how price-sensitive demand is and work out markups directly, instead of correlating structure with performance.


What would actually identify the effect of concentration?

Everything so far is *association* — things moving together. Concentration is itself an **outcome** of technology, entry costs and who wins. To show it actually *causes* higher prices or margins, you need concentration to move for some reason unconnected to how well firms are doing:

- **Merger event studies:** compare prices in markets where a national merger changed local concentration a lot vs a little (the workhorse in banking, airlines, hospitals, beer).
- **Regulatory shocks:** entry deregulation, licensing changes — concentration moves by law, not by success.
- **The modern route:** estimate how price-sensitive demand actually is, and back out each firm's markup directly — dispensing with HHI entirely. (That is where the field went, and where tomorrow's markup measurement starts.)

“We regressed performance on structure” proved nothing about cause in 1970 and proves nothing now. Today's exercise is *why* the field moved on.

Entry and exit: how easily can outsiders get in — and out?



- **Bankruptcy is 3.3% of exits. Acquisition is 48.6%.** The typical public company does not fail — it gets bought.
- That reframes Porter's "threat of entrants". An entrant's realistic exit is a **sale to an incumbent**, which is also the incumbent's cheapest way to remove a threat.
- A market is called **contestable** when entry and exit are easy, because the mere *threat* of entry disciplines prices. **So is heavy merger-and-acquisition activity evidence of a contestable market — or of one where competition simply gets bought?**

Source: own calculations, Compustat quarterly, delisting reason (d1rsn) for all US-listed firms 1980-2025 (n = 26,051 delistings).


One more diagram made measurable: Porter's generic strategies

This morning's triangle said firms choose **differentiation** or **cost leadership**. DuPont accounting gives each a fingerprint:

$$\text{ROA} = \underbrace{\frac{\text{profit}}{\text{sales}}}_{\text{margin (differentiate)}} \times \underbrace{\frac{\text{sales}}{\text{assets}}}_{\text{turnover (cost-lead)}}$$

- If the strategy choice is real, margin and turnover should be **substitutes** across firms: you can have one, not both.
- Prediction for the scatter: a downward-sloping cloud with industries occupying different ends — pharma top-left, retail and utilities bottom-right.

$\text{corr}(\text{gross margin, asset turnover}) \text{ in } 2024 = -0.37, n = 3,222.$ Next slide.



Sector medians in red. “Stuck in the middle” is visible too: the mass of firms near the origin of both axes — low margin *and* low turnover — which is where the negative ROA lives.

Own calculations, Compustat 2024, non-financials with sales > \$50m. Turnover = sales/assets; margin winsorised to [−20, 100].

Your turn (15 min)

Do-file 01_concentration_markup.do:

- ① Pick one industry you care about (SIC-3). Plot its HHI 1980–2024 and annotate the two biggest jumps — then look up what merger or shock caused each.
- ② Re-run the SCP table using **4-digit SIC**. Finer markets $\Rightarrow$ what happens to the HHI coefficient? What does that teach about market definition?
- ③ Compute the **4-firm concentration ratio** (CR4) alongside HHI. Rank industries by both. Where do they disagree, and why does HHI's squaring matter?

The point of (2) is that the “effect of concentration” is not robust to the definition of the market
—
which is the polite way of saying it was never identified.

- Demsetz, H. (1973), 'Industry Structure, Market Rivalry, and Public Policy', *Journal of Law and Economics* 16(1), 1–9.
- Schmalensee, R. (1985), 'Do Markets Differ Much?', *American Economic Review* 75(3), 341–351.
- Rumelt, R. (1991), 'How Much Does Industry Matter?', *Strategic Management Journal* 12(3), 167–185.
- McGahan, A. & Porter, M. (1997), 'How Much Does Industry Matter, Really?', *Strategic Management Journal* 18(S1), 15–30.
- Grullon, G., Larkin, Y. & Michaely, R. (2019), 'Are US Industries Becoming More Concentrated?', *Review of Finance* 23(4), 697–743.
- Autor, D., Dorn, D., Katz, L., Patterson, C. & Van Reenen, J. (2020), 'The Fall of the Labor Share and the Rise of Superstar Firms', *QJE* 135(2), 645–709.
- Philippon, T. (2019), *The Great Reversal*, Harvard University Press.
- Bresnahan, T. (1989), 'Empirical Studies of Industries with Market Power', in *Handbook of Industrial Organization*, Vol. 2.
- U.S. DOJ & FTC (2023), *Merger Guidelines* (HHI thresholds).
- Selling, T. & Stickney, C. (1989), 'The Effects of Business Environment and Strategy on a Firm's Rate of Return on Assets', *Financial Analysts Journal* 45(1), 43–52.
- Cameron & Miller (2015), *JHR* 50(2) — clustering field guide.
- Moulton, B. (1990), 'An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units', *REStat* 72(2), 334–338.

Take this with you

- Every framework this morning was a **hypothesis**. Today you saw which parts survive estimation, which fail, and which the data cannot even address.
- A regression is an association until you have argued identification. We wrote that argument out loud each time — copy the habit.
- The dataset determines the question. Knowing what Compustat **cannot** answer (private firms, input prices, geographic segments) is half of empirical literacy.

Code and data: [replication/](#). Do-files 01–04 reproduce every number shown today.