

Practical 2: Institutions — measurement, manipulation, and identification

Practical session, Day 2 — data, code and regressions

Dr. Ernil Sabaj

Warwick Summer School
International Business and Finance

21 July 2026

This morning ranked countries by institutional quality. This hour: where those numbers come from, how they were once faked, and what a regression on them can and cannot say.

- **Part 1.** The Doing Business scandal and its replacement (B-READY).
- **Part 2.** What firms report — and a live regression: bribery on income, cross-section vs fixed effects.
- **Part 3.** Political risk in a listing count: 25 years of Chinese firms in New York.

If the measure itself is corrupted, the regression measures the corruption.

Monday's six statistics ideas are all you need today; each is re-explained where it appears.

The plan: from this morning's slides to tests

Claim / diagram from the lecture	Data	What we estimate	The claim is in trouble if...
Countries ranked by institutional quality (Doing Business slide)	WilmerHale report; B-READY	— (a case, not an estimate)	the index was produced independently of politics. (It was not.)
Corruption is an environmental constant firms adapt to (PESTEL framing)	WBES (World Bank Enterprise Surveys): % of firms asked for a bribe, 400 country-waves	Bribery on $\ln GDP_{pc}$: cross-section, then country FE	the slope is zero — or so large only reporting drift can explain it
Development cures corruption (implicit in the readings)	Same, repeated waves	Within-country slope —19.7 vs cross-section —4.9	we can tell real improvement from quieter respondents. (We cannot.)
Political risk is a scenario exercise (Russia slides)	China-HQ listings, 1995–2025	A count over time; incorporation structure (VIE)	the series tracks fundamentals rather than politics

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

The data this hour

Today's main dataset is **tiny** — it fits on one screen:

- `wbes_bribery_gdp.csv`: one row = **one country in one survey year** (a “country-wave”). 400 rows; 165 countries; 2006–24.
- It comes from the **World Bank Enterprise Surveys**: interviewers visit thousands of firms per country and ask, among other things, whether officials requested an informal payment. Our variable is *the % of a country's firms who said yes* — what firms *experienced*, not what experts perceive.
- The second column, `gdppc`, is GDP per person (constant US\$) in the survey year.
- Most countries were surveyed **2–4 times** — that repetition is what Part 2's “compare a country with itself” trick runs on.

Monday's big firm panel (`ib_panel.dta`) returns in Part 3, for the legal-origins test. Same software note as Monday: slides show Stata; the identical Python is in `replication/student_pack_day2/`.

One row = one country-survey. 400 of them, 165 countries.

country	year	% of firms asked for a bribe
Georgia	2008	9.8
Georgia	2023	1.0
Denmark	2020	3.4
China	2024	0.1

- Georgia's -8.8pp across 15 years is exactly the **within-country variation** the FE regression consumes.
- China reporting *less bribery than Denmark* is the measurement lecture in one cell.

About that Doing Business ranking...

- On **16 September 2021** the World Bank **discontinued Doing Business entirely**, after the law firm WilmerHale — reviewing some 80,000 internal documents — found the rankings had been manipulated.
- **What was changed:**
 - **China** (DB 2018): ranked 78th. Without the irregularities: **85th**. The changes were made while China was expected to play a central role in a World Bank capital increase.
 - **Saudi Arabia** (DB 2020): named “top improver”. It would not have been.
 - **Azerbaijan** (DB 2020): listed among the top-10 improvers. It would not have been.
 - **UAE** (DB 2020): Paying Taxes indicator altered; rank unchanged at 16.
- The report named the Bank’s then-CEO and the office of its then-President. Both disputed the findings.

World Bank Group statement, 16 September 2021; WilmerHale, *Investigation of Data Irregularities in Doing Business 2018 and 2020*, 15 September 2021.

Why we are spending a slide on this

This is not gossip. It is the **cleanest teaching case in the whole course**.

- We use indices like this to tell firms **where to invest**. Governments rewrote actual laws to climb this ranking.
- The index that measured **institutional quality** was itself **captured by the politics it was supposed to measure**.
- **Every number in every lecture — including mine — was produced by someone with an interest.** Ask who, and what they wanted.

What replaced it: Business Ready (B-READY)

Three pillars, scored 0–100:

- ➊ **Regulatory framework** — the rules as written (*de jure*)
- ➋ **Public services** — what the state actually provides
- ➌ **Operational efficiency** — how the two combine in practice (*de facto*)

Ten topics following a firm's life cycle: entry, location, utilities, labour, finance, international trade, taxation, dispute resolution, market competition, insolvency.

Key design changes vs Doing Business:


- Firm-level survey data, not just expert opinion — built on the Enterprise Surveys
- Balances ease of business against workers, consumers, environment
- **No single headline ranking** to game

Rollout:

- 2024: 50 economies
- 2025: 101 economies (interim)
- 2026: full edition, $\approx$ 180 by 2027
- 43 indicators now in the WDI

World Bank, *Business Ready 2025*; worldbank.org/en/businessready. B-READY improves upon and replaces Doing Business.

What firms actually report: bribe solicitation



- Not a perception index. Firms are asked whether they were asked for a gift or informal payment across six specific transactions: electricity, water, construction permits, tax meetings, import licences, operating licences.
- A refusal to answer counts as a yes.

World Bank Enterprise Surveys, latest available survey per economy since 2014; regional figures are simple country averages. >219,000 firms, 159 economies.

Now be suspicious of the data I just showed you

Reported share of firms asked for a bribe:

	Earlier survey	Latest survey	
Azerbaijan	52.0% (2009)	0.7% (2024)	↓ 51 pts
China	11.6% (2012)	0.1% (2024)	↓ 12 pts
Vietnam	38.3% (2009)	31.1% (2023)	↓ 7 pts
Ukraine	38.5% (2008)	37.4% (2019)	flat
Georgia	9.8% (2008)	1.0% (2023)	↓ 9 pts

- By this measure **China (0.1%)** and **Azerbaijan (0.7%)** are cleaner than **Denmark (3.4%)**, **Belgium (6.9%)** and **Canada (5.8%)**.
- **Do you believe that?** What does a manager say to a stranger with a clipboard when the state can hear?
- **Azerbaijan was one of the four countries in the Doing Business scandal.** Georgia's fall, by contrast, tracks a documented reform programme.

World Bank Enterprise Surveys. Survey years in parentheses.

Self-reported data is weakest exactly where the thing you are measuring is most dangerous to report.

Bribery, censorship, forced labour, expropriation risk, state interference — the countries with the most of it will report the least of it.

That is not a reason to throw the data away. It is a reason to read the **direction of the bias** before you read the number.

Part 2. From bar charts to a regression

The regional chart is description. The question underneath it is causal:

Does economic development reduce corruption?

Two designs, with very different assumptions:

- ① **Cross-section:** compare rich vs poor countries at a point in time.
Identifies a causal effect only if nothing else drives both income and bribery. (Everything else drives both.)
- ② **Country fixed effects:** compare each country *with itself* as it grows, using repeated survey waves.
Differences out everything permanent — geography, colonial history, culture. Still vulnerable to anything that *trends*.

We will run both, live, on 400 country-surveys.

Building the dataset (Stata)

```
import delimited wbes_bribery_gdp.csv, clear
                                // open a CSV (it's in replication/data/)
gen lgdp = ln(gdppc)           // log income: slopes then read
                                // "per doubling" (Monday, slide 9)
encode iso3, gen(cid)          // turn country codes ("GEO") into
                                // numbers -- the panel tools need that
* (1) cross-section: one row per country
bysort cid (year): gen last = _n == _N
                                // within each country, sorted by year,
                                // flag its LATEST survey
reg bribe lgdp if last, robust
                                // best-fit line across countries;
                                // "robust" = a +/- that stays valid
                                // when spread differs across countries
* (2) within-country: repeated waves
xtset cid year                 // declare: this country, this year
xtreg bribe lgdp, fe cluster(cid)
                                // "fe" = compare each country WITH
                                // ITSELF; cluster = a country's waves
                                // move together (Monday's correction)
```


Fixed effects with two waves is just looking at changes

“Fixed effects” sounds technical; with two surveys per country it is nothing scarier than this: **for each country, subtract the earlier survey from the later one, and relate the change in bribery to the change in income.** In symbols:

$$\Delta \text{bribe}_c = \beta \Delta \ln \text{GDPpc}_c + \Delta \varepsilon_c.$$

- Everything permanent about the country — geography, colonial history, culture, a *constant reporting bias* — subtracts out. That is the whole sales pitch.
- **What does not subtract out:** anything that *changes* over the window. Georgia's −8.8pp (9.8 → 1.0) enters as a genuine within-country move — but so would a change in respondents' willingness to talk. FE removes level bias, **not trending bias**.
- One practical detail: only the **123 countries surveyed at least twice** actually contribute — a country with a single survey has no change to compare. Always know *which observations* your estimate really comes from.

The cross-section



World Bank Enterprise Surveys (latest wave per country, $N = 165$) $\times$ WDI GDP per capita. Labelled points discussed next.

Read the scatter before you read the line

Habits of the trade, applied to the last slide:

- **What is a point?** One country's *latest survey*. So $N = 165$; countries are not equally reliable (some surveys are small); and poor countries cluster in particular regions — part of the slope is a between-region comparison.
- **The spread is not the same everywhere:** huge scatter at low incomes, tight at high. That doesn't change the slope — it changes how sure we can be about it, which is exactly what the robust option fixes. (The textbook name, if you meet it: heteroskedasticity.)
- **Zeros pile up on the floor:** a straight line happily predicts *negative* bribery for rich countries — harmless for summarising the slope, silly for prediction. Know which you are doing.
- **Spot the extreme points before arguing:** the labelled countries (China and Azerbaijan implausibly low; Israel and Iraq high) are where the measurement argument will be won or lost — the next three slides.

Both regressions on one slide


Dep. var.: % of firms asked for a bribe	(1) Cross-section	(2) Country FE
ln GDP per capita	−4.88 (0.58)	−19.7 (3.4)
Country fixed effects	no	yes
N	165 countries	358 surveys / 123 countries
R^2	0.28	— (within)

(1) robust SEs — the $\pm$ stays valid when the spread differs across countries; (2) SEs clustered by country — a country's waves move together. Read the slopes per doubling of income: −3.4pp in (1), −13.6pp in (2).

Data: wbes_bribery_gdp.csv — (1) latest wave per country, $N = 165$ countries; (2) all 400 country-waves, 165 countries, 2006–24. y : % of firms asked for a bribe; x : ln GDP per capita (constant US\$).

- Both negative, both large, both significant. And the within-country slope is **four times steeper** than the cross-section.
- Before you believe (2), name what else changed *inside* these countries, in these same 15 years, that also moves reported bribery...

See the within-country variation the FE estimate uses



Every grey line is one country's bribery series across survey waves (countries with ≥ 3 waves). The FE regression lives *entirely* inside these lines. Bangladesh 60→23; Georgia 10→1; **China 11.6→0.1**. Real improvement — or growing caution about what firms tell a World Bank enumerator?

Data: wbes_bribery_gdp.csv, WBES 2006–24. Regenerate: python/06_new_lecture_figures.py.

Two stories fit the FE estimate. The data cannot separate them.

Story A: real improvement

Growth builds tax capacity, formalises firms, raises audit quality, raises the opportunity cost of extortion. Georgia is the poster case: documented reform programme, bribery 9.8% → 1.0%.

Story B: reporting drift

As states get richer and stronger, firms get *more careful talking to surveyors*. China: 11.6% → **0.1%** while GDP tripled. Azerbaijan: 52% → 0.7% — one of the four countries in the Doing Business manipulation.

- Both stories predict a negative FE coefficient. **The regression is consistent with either.**
- The measurement error here is not classical noise — it is *correlated with the regressor* (state capacity grows with income). That biases the slope in a direction you cannot sign without outside information.

An identified regression on a corrupted measure identifies the corruption of the measure.

When mismeasurement is harmless — and when it lies with direction

Data is never measured perfectly. What matters is *what kind* of error it is:

- **Pure noise** (random, unrelated to anything): it just *drags slopes toward zero*. Annoying, but it makes you *understate* effects — a conservative error.
- **Error that moves with the thing you're studying** — our case: as states grow stronger (which happens as countries grow richer), firms plausibly *report less* regardless of the truth. Now the error is not noise: it **pushes the slope in a particular direction — the same direction as the effect we are claiming**. Real improvement and growing silence are added together in one number, inseparably.
- That is why the “two stories” slide is not hand-wringing: without an *outside* check (audit data, conviction records, a field experiment), the steep within-country slope simply cannot be split into its two parts.
- Portable lesson: before trusting any self-reported number, ask **who has a reason to misreport — and whether that reason grows or shrinks along your x-axis**.

The reverse-causality problem, and the famous attempt to solve it

Cross-country income $\leftrightarrow$ institutions runs both ways: rich countries can afford good courts. The famous attempted fix uses an **instrument** — a variable that moves institutions but plausibly affects income *only through* them:


- **Acemoglu, Johnson & Robinson (2001, AER)**: instrument institutions with **settler mortality**. Where colonisers could survive, they built inclusive institutions; where they died, they built extractive ones. Mortality 200+ years ago plausibly affects income today *only* through institutions.
- IV estimate: institutions explain most of the income gap between the top and bottom of the distribution — “institutions rule” over geography and trade (Rodrik, Subramanian & Trebbi 2004).
- Not uncontested: Albouy (2012) shows the mortality series is fragile; the exclusion restriction (mortality $\rightarrow$ income only via institutions) is an *assumption*, not a finding.

Why it matters for you: this literature is the intellectual backbone of every “institutional quality” index your future employer will buy.

Political risk, priced

The Russia slides showed firms leaving under sanctions.
Here is the same story, running for twenty-five years.

Chinese firms on US exchanges



Nothing about the *economics* of these firms explains this line. **Politics does:** reverse-merger fraud scandals, the 2020 Holding Foreign Companies Accountable Act (audit inspection or delisting), the Didi IPO and Beijing's response, then a partial 2022 audit agreement. Source: own calculations, Compustat quarterly. Firms with headquarters (1oc) in China with a fiscal-year record. 2025 is partial.

The structure underneath: you do not own what you think you own

Of the 685 China-headquartered firms ever listed in the US:

Incorporated in	Firms
Cayman Islands	392
United States	202
British Virgin Islands	62
China	23
Hong Kong	5

- Chinese law restricts foreign ownership in tech, education and media. The workaround is the **Variable Interest Entity**: you buy shares in a Cayman company that holds *contracts* with the Chinese operating company — not its equity.
- When Beijing banned for-profit tutoring in July 2021, the VIE contracts of the listed education firms became worth roughly nothing. **Nobody expropriated anything.** The contracts simply pointed at a business that was no longer legal.
- This is the “home country” problem from Lecture 1 with real money attached.

One practical warning before you merge World Bank data to firms

The natural next step is firm outcomes (left) on country institutions (right):

```
use ib_panel, clear
merge m:1 iso3 fyear using wdi_country_year
```

- Merge on `hqloc` (headquarters), never `fic` (incorporation). 6.2% of US-listed firms are incorporated somewhere other than where they operate, and the gap is systematic: Cayman, BVI, Bermuda.
- Merge on `fic` and you will estimate the effect of **Cayman Islands institutions** on Houston oil-services firms. This exact mistake appears in published work.
- And remember what this sample *is*: firms that raised capital in New York. It is not “firms in country X”. For that, use the Enterprise Survey microdata or Orbis.

Part 3. LLSV inside this panel: legal origin at the firm level

This morning's legal-origins chart was **country-level**. The same claim has a firm-level version: firms from weaker-investor-protection (civil-law) countries should trade at lower valuations. Test it on the cross-listed foreign firms in `ib_panel.dta`:

```
use stata/ib_panel, clear
keep if hqloc != "USA" & inrange(fyear, 2010, 2024)
        // foreign-HQ firms only, 2010-24
gen origin = 1*inlist(hqloc,"GBR","CAN","AUS","IND","SGP","HKG","ISR","IRL") ///
+ 2*inlist(hqloc,"FRA","ITA","ESP","BEL","NLD","BRA","MEX","TUR") ///
+ 3*inlist(hqloc,"DEU","AUT","CHE","JPN","KOR","TWN") ///
+ 4*inlist(hqloc,"DNK","FIN","NOR","SWE")
        // code each HOME country into its legal
        // family: 1 English 2 French 3 German 4 Scand.
gen q = (mktcap + assets - equity)/assets // Tobin's q (Monday)
gen lnq = ln(q) if inrange(q, 0.2, 20) // log + the guard
areg lnq ib1.origin i.gsector if origin>0, absorb(fyear) cluster(gvkey)
// regression of lnq on the origin groups; "ib1" makes English
// the comparison group (all results are RELATIVE TO IT);
// i.gsector + absorb(fyear) = sector & year effects;
// cluster(gvkey) = Monday's same-firm correction
```

Before you run it: which origin group do you predict lowest? By how much?

Civil-law-origin firms: a 12–13% Tobin's-q discount.

- French origin -12.1% ($t = -3.8$); German origin -13.1% ($t = -3.2$); English base. Scandinavian $+6\%$ but $t = 0.7$ on just **43 firms**. $N = 13,345$ firm-years, year + sector FE, SEs clustered by firm.
- **Selection cuts against us finding this:** every firm here chose to cross-list in a common-law market — yet home legal origin still prices.
- Link to Practical 1: the -13% foreignness discount is not uniform — **legal origin splits it**.
- What it is *not*: causal proof. Origin is bundled with everything else about a country. LLSV's critics (Rajan–Zingales' 1913 reversals) live in exactly this gap.

Your turn (15 min)

Using `wbes_bribery_gdp.csv` (do-file 04, Part B):

- 1 Add **region fixed effects** to the cross-section. How much of the income slope survives within-region comparison?
- 2 Drop China, Azerbaijan and Uzbekistan — the three most suspicious trajectories — and re-run the FE regression. Does the slope move? What does your answer imply about how much of the result those three countries carry?
- 3 Regress the **change** in bribery on the **initial level** (convergence regression). Do high-bribery countries improve faster, or is there hysteresis?

Discussion target: by the end you should be able to say, in one sentence each, what (1), (2) and (3) can and cannot identify.

- World Bank Enterprise Surveys, enterprisesurveys.org; microdata at microdata.worldbank.org.
- WilmerHale (2021), *Investigation of Data Irregularities in Doing Business 2018 and Doing Business 2020*, 15 September.
- World Bank (2024, 2025), *Business Ready* reports.
- Acemoglu, Johnson & Robinson (2001), *AER* 91(5), 1369–1401.
- Albouy, D. (2012), 'The Colonial Origins of Comparative Development: An Investigation of the Settler Mortality Data', *AER* 102(6), 3059–3076.
- Rodrik, D., Subramanian, A. & Trebbi, F. (2004), 'Institutions Rule', *Journal of Economic Growth* 9(2), 131–165.
- Fisman, R. & Svensson, J. (2007), 'Are corruption and taxation really harmful to growth? Firm level evidence', *Journal of Development Economics* 83(1), 63–75.
- WDI GDP per capita series via Gapminder/Open Numbers mirror (constant US\$).
- Angrist, J. & Pischke, J.-S. (2009), *Mostly Harmless Econometrics*, Princeton University Press (chs. on FE and measurement error).

- 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.