Pablo Fiorentini

PhD candidate in Economics, Rome Economics Doctorate (RED: EIEF, Luiss, Sapienza, Tor Vergata)
pfiorentini@luiss.it · pfiorentini@redphd.it · pablofiorentini.com

Application to the NBER Digital Economics and AI Tutorial 2027

Graduate study

PhD in Economics, Rome Economics Doctorate (RED) 2024–present
Third year in 2026–27. Advisor: Andrea Pozzi. Four-year RED scholarship (2024–2028).

Master in Economics and Finance, CEMFI, Madrid 2022–2024
Thesis supervised by Pau Roldán-Blanco.

Bachelor in Economics, Universidad Carlos III de Madrid 2018–2022

Research interests

Industrial Organization and Political Economy. I study rules written or shaped by interested parties, such as safety codes drafted by industry committees and lobbying of the EU institutions, and industrial policy for digital and green investment. I use language models to measure economic content in large text records. The self-regulation of AI is the application that motivates my main project.

Main project

Adversarial Peer Review: Safety Rules Written by Competing Firms

AI companies write their own safety frameworks and help draft the technical standards that regulators rely on, such as the European harmonised standards for the AI Act. How such rules depend on the interests of the firms that write them is hard to study in AI, where these arrangements are only a few years old. I study a long running case of the same arrangement, where the votes of the people who write the rules are on record.

The setting is a product safety code that governments adopt as law. Technical committees revise it in regular cycles, and their members work for the organisations the code binds: makers of competing products, the builders and utilities that buy them, installers, testing laboratories and inspectors. A change needs a supermajority of the committee; otherwise the old text stays. Rivals review each other's proposals, and the rules that pass bind them all; I call this adversarial peer review. I ask whether members vote against changes that would restrict their own organisation's products or raise its costs more often than other members voting on the same change, and whether the changes that fail are sunk by the seats they burden or by the supermajority rule.

I built the data from the committees' printed ballots for ten editions over more than twenty-five years: about 50,000 balloted proposals, comments and revisions; rosters with each member's organisation and interest class (5,900 entries linked into 1,681 persons); and about 730,000 member votes. Where a ballot does not name the yes voters, yes votes are imputed from the roster (three quarters of all votes); named and imputed votes reproduce every printed count on 89.5 percent of the items with imputed votes.

The stakes come from text. Language-model agents (Claude) code which products a yes would require or burden, who bears its cost, and what each organisation makes or pays for. Two independent codings agree on what a yes does for 89 percent of the items coded twice, but on what an organisation makes for only 61 percent of the organisations coded twice; no human coder has checked the codes yet. An open-weight model run on a computing cluster classified the stated reasons in about 30,000 written explanations of votes.

First patterns, all of them associations: 82 percent of the items with printed counts drew not one no vote; most changes that failed drew more yes than no votes, and the members a failed change burdens never cast enough no votes to sink it alone, though they seldom held enough seats to do so. On the same proposal, burdened members vote no 9 to 40 percentage points more often than members without a coded stake, depending on the sample; on the recent editions, where every vote is named, the gaps are smaller and imprecise. Next come a human coder on a random sample, to measure the accuracy of the language-model codes, and a model of the committee, written after the data and fully solved. Currently at the analysis stage.

Other work in progress

Green Conditionality in Industrial Policy: Certified Energy Savings, Threshold Incentives, and Rationing in Italy's Transizione 5.0

Italy's Transizione 5.0 tax credit (2024–2025) paid for digital and energy investment only if a certifier attested a minimum reduction in energy consumption, with higher rates for higher certified savings. I study whether this conditionality led to additional investment and real energy savings, or mainly to strategic certification around the thresholds. I plan to compare ex ante and ex post certified savings near each threshold, look at differences across certifiers, and use the moment the budget ran out to measure the cost of rationing for firms. Currently at the design and data access stage.

Gender Bias in AI Systems

Joint work with Martina Logi Cervera and Carolina Polito, within the ENGAGE.EU Research Label project *Digital Technologies and Democratic Practices for Gender Equality*. We drafted a critical framework for understanding gender bias in AI across data and model levels, arguing these biases are structural externalities for governance and regulation. Presented at the ENGAGE.EU Research Label roundtable, Luiss, April 2026.

Lobbying and Industry Concentration in the European Union

I study how concentrated lobbying of the EU institutions is across firms, and how this compares with the concentration of sales within industries, with three public sources: the EU Transparency Register, the Have Your Say consultation portal, and the European Commission's records of its meetings with interest representatives.

Other projects

A Monetary Model of Money Laundering, joint work with Fabrizio Mattesini: a theoretical model to be linked to moments drawn from the Financial Intelligence Unit datasets of the Bank of Italy. *Urban Mobility, with a Focus on the City of Rome*: a scraper I started in my first year collects the positions of shared bikes and e-scooters in Rome, together with the city's public transport feeds. *Automation, Population Growth, and the Future of Human Capital*, my Master thesis at CEMFI, currently under revision: a capital accumulation model with human capital in which automation is possible.

Why this tutorial

My main project studies rules written by interested firms, and the governance of AI is its motivation. At the tutorial I would like to learn how economists are approaching AI standards and safety frameworks and which data exist on them, and to get feedback on measurement with language models, on which my project relies.