Seminar 11 (Simulation) — The Lindmark Decree

IR227 Politics of Change — alternative seminar format

July 2026

The setting

[Bridge to seminar] *Lindmark is a mid-sized country in the heart of the European Union, home to roughly 9 million people. It is spring 2026. Lindmark's AI and data-protection regulator is four months from finalising the national decree implementing the EU AI Act's high-risk obligations. Eight weeks ago, Vesla — a European AI-and-robotics firm — announced a €2.4bn manufacturing campus in Lindmark, and separately licensed its Vesla OS platform to retrofit two of Valexo Industrial's three plants with robotic, AI-managed assembly lines.*

Lindmark, Vesla, and Valexo are fictional. The EU AI Act and GDPR are not.

What you'll do in 90 minutes

- Three teams: the Lindmark Autoworkers' Federation (LAF) writes a demand memo; the national regulator (LDGA) writes an implementing white paper; Vesla writes a market access and compliance plan.
- Each team works from the same background material — including the EU AI Act's real risk tiers — but applies it independently, to build its own case. Nothing about how Vesla OS is classified is settled in advance.
- Plenary is a live three-way hearing.

Three methodological points up front

1. The AI Act's risk tiers are real, citable law — but how they apply to Vesla OS's specific functions is genuinely contestable. Let teams argue different readings; don't hand them a settled classification.
2. AI/automation is not just a labour-market shock — Farrell's reframing applies directly: Vesla OS allocates work, scores performance, and shapes termination decisions. That is a governance function exercised over workers, not only a technology that displaces them.
3. The same workforce faces two distinct harms that don't collapse into one demand: job loss for the displaced, and algorithmic supervision for the retained. Push every team to keep them separate.

Three teams, one hearing

	LAF (union)	LDGA (regulator)	Vesla
Deliverable	Demand memo	Implementing white paper	Market access plan
Core ask	Statutory worker voice	Build the missing floor	Predictable, light enforcement
Primary tool	Acemoglu-Restrepo	Lecture 8 mechanisms	Streeck-Thelen mechanisms

$\sim$8 students per team (class of $\sim$24). Each team owns one role brief from `ta/worksheet-11-sim.md` plus a short intelligence note on the *other two* sides.

Recap: the EU AI Act's risk tiers (real law, background material only)

Framework: Risk-based regulation under Regulation 2024/1689

1. Unacceptable risk — prohibited outright.
2. High risk — Annex III, including AI used to recruit, allocate tasks, monitor/evaluate performance, or materially influence termination decisions.
3. Limited risk — transparency obligations only.
4. Minimal risk — no specific obligations.

High-risk obligations bind across the EU from 2 August 2026 — four months after the date this simulation is set. Hand this out as reference; do not resolve it as a class. Each team argues its own reading of where Vesla OS's four functions land.

Recap: Acemoglu-Restrepo's displacement-reinstatement framework

- Robots and AI systems substitute for workers in specific tasks (displacement); new tasks and industries sometimes create new demand for labour (reinstatement) — but not automatically for the *same* workers.
- Effects concentrate in routine-manual and routine-cognitive occupations, not evenly across a workforce.
- Apply directly: which Valexo tasks does the robotic line substitute for? Does the campus's retraining pipeline reach *those* workers, or a different, differently-qualified pool?

Recap: the commitment-problem lens

- Employers and governments cannot always credibly promise future retraining or compensation; workers who doubt it will materialise have every reason to resist change now, while they still have leverage.
- A retraining fund or a "commitment to consult workers" is only as good as the institutional guarantee behind it: statutory entitlement vs. discretionary guidance; independent monitoring; enforceable timelines.
- Carried over from Lecture 5 and Lecture 9 — LAF's central weapon, LDGA's central design constraint.

Recap: Streeck-Thelen's institutional-change mechanisms

Framework: Five modes of gradual institutional change

Drift (rule unchanged, environment changes); conversion (rule redirected to new purposes); layering (new rule added without repealing the old); displacement (old rule actively replaced); exhaustion (institution withers).

Ask, for each side, which mechanism it is trying to engineer for *itself*:

- Vesla: is its arrival at Valexo *displacement* of existing management practice, or *layering* — a new system beside the old, formally leaving it intact?
- LAF: is a statutory worker-voice right a *displacement* of management's current free hand, or *layering* a new institution alongside it?
- LDGA: is the national decree *layering* onto the AI Act's floor — and onto what else?

This is Vesla's primary tool this session — the market access plan should name one

Recap: Farrell's reframing of AI as governance

- AI is not just an object of governance (something to regulate) — it is a governance mechanism in its own right, allocating work, attention, and legitimacy.
- Vesla OS allocates tasks, scores performance, and shapes termination decisions — a governance function exercised over workers, whatever Vesla or the government calls it.
- Push every team: is your side treating Vesla OS only as a labour-market technology, or also as a governance system whose exercise of power is itself the thing at stake? The two framings lead to different demands.

Recap: Meckling's coalition logic (secondary, for LDGA and Vesla)

- Large-scale industrial investment builds political coalitions that make harder regulatory measures sustainable later — but only if someone deliberately connects the investment to the coalition that needs building.
- Vesla's €2.4bn campus is this kind of asset whether or not either side frames it that way.
- LDGA's question: spend it as coalition capital toward *what* — union acquiescence, or binding worker-voice rights? Vesla's question: how much does being visibly "the carrot" cost if the Valexo displacement numbers go public?

Tasks

- Setup and role assignment (10 min)
- Team drafting (35 min) — all three fill in their deliverable, plus a short cross-transition reflection.
  - LAF: demand memo.
  - LDGA: national implementing white paper.
  - Vesla: market access and compliance plan.
- Hearing prep (8 min) — each team selects two spokespeople.
- The Lindmark hearing (27 min) — 5 min presentation each, 12 min TA-moderated cross-examination.
- Debrief and real-world parallels (10 min).

Debrief — real-world parallels

- Lindmark is fictional; the legal backdrop is not. Three EU states have already answered pieces of tonight's fight.
- Germany's *Betriebsverfassungsgesetz* §87(1)(6) gives works councils a statutory co-determination right over introducing any technical device designed to monitor employee performance or behaviour — a binding answer to LAF's "statutory, not discretionary" demand.
- Spain's 2021 Rider Law requires employers to disclose algorithmic-management parameters to worker representatives — a binding transparency precedent.
- France's data-protection authority (CNIL) fined Amazon France Logistique in 2024 over excessive warehouse-worker monitoring — enforcement precedent against exactly this kind of system, under GDPR alone, before the AI Act's obligations even bind.
- None of these existed because the EU AI Act required them — member states built them separately. That is exactly the missing-condition gap tonight's LDGA

Three things to take away

1. A real, binding EU-level floor still leaves the classification of a specific system genuinely contestable — and the national fight over enforcement and worker voice happens regardless of how that argument resolves.
2. AI/automation governance is labour-market and governance-over-workers simultaneously (Farrell) — treating it as only one collapses half the analysis.
3. Displacement and algorithmic supervision are not the same harm, even when they hit the same workforce — a settlement that only addresses one leaves the other unresolved.

Tomorrow: polycrisis synthesis — if one wealthy, institutionally strong EU state still can't fully resolve one company's worth of AI-driven displacement and algorithmic governance, what happens when climate, demographic, migration, and AI transitions land on the same workforce at once?