

Civilian AI Is Already at the Front. Europe's Rules Can't See It.

Europe controls the export of dangerous things. The technology now shaping surveillance and targeting isn't a thing — it is a set of weights that can be retrained. Brussels takes evidence until 15 October.

1. The current issue

An object-recognition model trained to count warehouse boxes can, with modest retraining, classify people and vehicles in drone footage. Two European instruments govern this and neither reaches it. [Regulation \(EU\) 2021/821](#), the Dual-Use Regulation, controls export of items listed in Annex I, amended yearly; models, weights and training data are not listed. The [AI Act](#) does not apply to systems placed on the market exclusively for military, defence or national security purposes. A list that catches specified items, an AI law that steps back from military use: the retrained model falls between them.

The costs land unevenly: developers cannot tell whether downstream repurposing obliges them; twenty-seven [national authorities](#) license cases the list does not describe; people classified abroad have no forum; parliaments approve a cloud contract rather than a weapons programme. Licensing is the one point at which the Union's human-rights commitments bind *before* deployment.

2. Historical lineage

- **1949–1996 — the list outlives its war.** COCOM controlled listed items bound for the Eastern bloc. It [dissolved in March 1994](#) and its lists passed into the Wassenaar Arrangement, whose method — technical specifications revised by consensus — Annex I still inherits.
- **2000 — GPS becomes a utility.** On 1 May 2000 the United States [switched off Selective Availability](#), the deliberate degradation of the civilian signal. A military system became consumer infrastructure, and dual-use stopped meaning civilian goods with military potential and came to mean both at once.
- **2011–2013 — correction after normalisation.** European spyware exports to Libya, Egypt and Bahrain became public; Wassenaar answered in 2013 with controls on “intrusion software”. Researchers soon showed the definition also [captured penetration testing and vulnerability research](#), and it was renegotiated for years.
- **2018 — Project Maven.** Google applied commercial cloud machine learning to Pentagon drone-footage analysis. Four thousand employees petitioned; the contract was [not renewed](#). Briefly it was visible that targeting already ran on commercial infrastructure.
- **2021–2025 — where the new controls landed.** The recast added [Article 5](#), requiring authorisation for *unlisted* cyber-surveillance items — written after the scandals it answers, barely used since. [Delegated Regulation \(EU\) 2025/2003](#) then added quantum, semiconductor equipment and advanced chips: hardware chokepoints transposing Wassenaar's 2024 decisions, silent on [models and weights](#).

3. Mechanism of path dependence

[Paul David](#)'s account of QWERTY names technical interrelatedness, economies of scale and quasi-irreversible investment. All three operate here, plus two mechanisms specific to law.

- **Installed base.** Once cloud platforms, open weights and off-the-shelf drone components are the default on both sides of the line, a separate military stack costs more and performs worse.
- **Entrenchment through procedure.** Annex I changes by annual delegated act transposing Vienna. The lock-in is the amendment *procedure* — reactive, consensus-bound, item-shaped — not the list's contents.
- **Professional expertise.** Licensing officials match objects against specifications; 2013 showed the vocabulary does not transfer to software.
- **Procurement and funding.** Agencies buy AI “as a service”, binding capability to a vendor's update cycle, while ReArm Europe money rewards “dual-use by design”.
- **Political imaginary.** “Civilian” and “military” are administrative categories now treated as facts about objects.

4. Politics and distribution

[Langdon Winner](#) argued in 1980 that artefacts have politics in two senses: a design can settle a political question quietly, outlasting the debate, and some technologies need a particular structure of authority to work at all. Reconnaissance-as-a-service does both.

Authority accrues to cloud and AI firms, co-authors of military capability scrutinised as ordinary vendors; to procurement officials, who buy capability on a commercial contract and skip the scrutiny a weapons programme attracts; and to states with privileged access to frontier infrastructure. This is what [James Beniger](#) called the control revolution: owning the information infrastructure becomes owning the capacity to act.

[Barbrook and Cameron](#) supply the ideological half. The Californian Ideology's self-image — a civilian industry of markets and garage genius — always obscured the public and military money beneath it, and it now does legal work: “civilian technology” is a category claim, and the claim decides which regulator applies. Fewer choices go to people classified inside a commercial black box, firms whose models are repurposed without their knowledge, small states renting capability they cannot govern, and parliaments reading a budget line.

5. Contingency and alternatives

None of this was inevitable, and reading lock-in as destiny is the error this history warns against. Each mechanism above raises the price of a different path; none sets one. 2013 proves it: states *did* extend a control list to software when they chose to. Models are unlisted because of drafting difficulty and political appetite, not physics.

The clearest fork was 2018. Google's AI principles restricted weapons work, and the industry could have turned a firm-specific concession into a binding sectoral code on military end-use, as arms manufacturers have. It didn't: there was no forward market in restraint, competitors absorbed the contracts, the moment passed. In February 2025 Google [deleted the commitment](#). A voluntary pledge proved reversible because it was voluntary; a 2018 contract condition would have had nothing to reverse.

6. Governance implication

One proposition: put the trigger in procurement rather than the list. Make disclosure of civilian AI lineage a condition of public contracting for reconnaissance and targeting support — the authority discloses model provenance and cloud dependencies, obtains authorisation on a *foreseeable end-use* test modelled on Article 5, and publishes the human-rights review before deployment.

History shapes that design. Annex I is amended reactively, so classification always arrives after repurposing and the trigger cannot be a listing. 2013 showed technical definitions of software over-reach, so the test must attach to use. 2018 showed voluntary restraint has no forward market, so the obligation must sit where a state is already the counterparty. The [open public consultation](#) runs to 15 October 2026; the Commission must report under Article 26(4) from 10 September 2026.

Further reading and sources

Theory. Winner, "Do Artifacts Have Politics?" (1980) · David, "Clio and the Economics of QWERTY" (1985) · Beniger, *The Control Revolution* (1986) · Barbrook & Cameron, "The Californian Ideology" (1995/96) · [Thirty years on](#) (Tech Policy Press, 2025)

Law and institutions. [Dual-Use Regulation \(EU\) 2021/821](#) · [AI Act \(EU\) 2024/1689](#) · [Delegated Regulation \(EU\) 2025/2003](#) · [White Paper on Export Controls \(2024\)](#) · [Recommendation on national control lists \(2025\)](#) · [Genesis of the Wassenaar Arrangement](#)

Cases and analysis. [GPS Selective Availability \(gps.gov\)](#) · [Maven non-renewal](#) (The Intercept, 2018) · [Google drops the weapons pledge \(2025\)](#) · [EFF on the intrusion-software controls](#) · [Ruohonen & Kimppa on the Wassenaar debate](#) · [SIPRI on controlling algorithms, data and models](#) · [SIPRI on the Article 5 catch-all](#)

Have your say. [Open public consultation, to 15 October 2026](#) · [Call for evidence](#)