

Escuela de Negocios

Tipo de documento: Preprint



The Inference Bottleneck: Antitrust and Neutrality Duties in the Age of Cognitive Infrastructure

Autorías: Besanson, Gastón; Celani, Marcelo

Fecha: 26/02/2026

Publicado originalmente en: Arxiv

¿Cómo citar este documento?

Besanson, G., Celani, M. (2026). *The Inference Bottleneck: Antitrust and Neutrality Duties in the Age of Cognitive Infrastructure*. Arxiv. <https://doi.org/10.48550/arXiv.2602.22750>

El presente preprint se encuentra alojado en el Repositorio Digital Universidad Torcuato Di Tella, bajo una licencia Creative Commons Atribución 4.0, según lo indicado en la fuente original del documento.

Dirección: <https://repositorio.utdt.edu/handle/20.500.13098/14326>

The Inference Bottleneck: Antitrust and Neutrality Duties in the Age of Cognitive Infrastructure

Gaston Besanson
Universidad Torcuato Di Tella
Buenos Aires, Argentina

Marcelo Celani
Universidad Torcuato Di Tella
Buenos Aires, Argentina

Abstract

As generative AI commercializes, competitive advantage is shifting from one-time model training toward continuous inference, distribution, and routing. At the frontier, large-scale inference can function as cognitive infrastructure: a bottleneck input that downstream applications rely on to compete, controlled by firms that often compete downstream through integrated assistants, productivity suites, and developer tooling. Foreclosure risk is not limited to price. It can be executed through non-price discrimination (latency, throughput, error rates, context limits, feature gating) and, where models select tools and services, through steering and default routing that is difficult to observe and harder to litigate.

This essay makes three moves. First, it defines cognitive infrastructure as a falsifiable concept built around measurable reliance, vertical incentives, and discrimination capacity, without assuming a clean market definition. Second, it frames theories of harm using raising-rivals'-costs logic for vertically related and platform markets, where foreclosure can be profitable without anticompetitive pricing. Third, it proposes Neutral Inference: a targeted, auditable conduct approach built around (i) quality-of-service parity, (ii) routing transparency, and (iii) FRAND-style non-discrimination for similarly situated buyers, applied only when observable evidence indicates functional gatekeeper status.

Keywords

Antitrust; vertical foreclosure; self-preferencing; raising rivals' costs; non-price discrimination; inference; APIs; gatekeepers; DMA; AI Act; FRAND; auditing; routing transparency

1 Introduction: From Models to Infrastructure Power

Antitrust meets new technologies with an old question: when does control over a core input become the power to pick winners in vertically related markets? In the generative AI era, that question increasingly travels through inference and the interfaces that deliver it.

The commercial control point is not only who can train large models, but who can reliably serve frontier-grade inference at scale under technical and contractual terms downstream developers can depend on, and who can route users through AI interfaces that increasingly substitute for traditional distribution channels. Enforcement remarks and policy discussions have highlighted the risk pattern often described as "open early / closed late"—generous early access that may induce dependence, followed by tighter access or degraded external quality once switching costs and distribution consolidation increase. [1]

This essay's claim is narrow: if inference becomes a bottleneck input for a set of downstream applications, foreclosure, when it

occurs, will often be implemented through non-price discrimination and steering, not only through posted prices.

2 Why Inference Is the Bottleneck Surface

Inference is the commercial surface through which providers differentiate and monetize: token-based pricing, throughput tiers, context limits, realtime endpoints, and capability wrappers (tool use, batch modes, enterprise governance). Public pricing and documentation show how access is materialized primarily as inference consumption and service features, not as model weights. [2]

This packaging matters for competition because it creates multiple discrimination levers that are: (i) high impact (latency and reliability can determine viability for copilots, assistants, and realtime systems), (ii) low visibility (outsiders cannot easily observe serving priority, queue placement, or eligibility policies), and (iii) plausibly justified (differences can be attributed to capacity, safety, or abuse prevention).

OECD work in late 2025 emphasizes that competition risks are not confined to training and that AI adoption can reshape downstream competitive dynamics. [3, 4] Market-definition ambiguity reinforces the case for focusing on measurable conduct: foundation models can be inputs across many sectors while also competing with downstream applications at the same time. [1] The functional question is whether a provider can act as a discriminatory bottleneck for downstream rivals.

3 Defining "Cognitive Infrastructure" as a Falsifiable Concept

Definition (Cognitive Infrastructure). An inference service functions as cognitive infrastructure when all three conditions hold:

- (A) **Material reliance with costly substitution.** Downstream products rely on the service's performance and governance envelope (reliability, latency, context, tool-use features, safety/compliance assurances) such that switching to available alternatives would require non-trivial reengineering and re-validation and/or would cause a measurable degradation in cost, quality, or compliance outcomes.
- (B) **Control with vertical incentives.** Access terms and technical quality are controlled by a firm (or small set of firms) that is vertically integrated into downstream complements competing with its API customers. The concern is "referee-and-player" incentives where upstream governance choices can predictably shift downstream rivalry. [1]
- (C) **Discrimination capacity (non-price).** The controller can raise rivals' effective costs or reduce rivals' quality through non-price discrimination—QoS differentials, feature gating, model/version access, routing and eligibility policies—while maintaining facially neutral pricing.

Testable predictions and minimal audit protocol. (1) *QoS wedge*: holding region, load, and contracted tier constant, first-party complements exhibit systematically better QoS (p50/p95/p99 latency, timeout/error rates, quota denials) than similarly situated third-party API users. (2) *Routing bias*: where the model routes tasks to tools or services, routing outcomes favor the provider’s own services or preferred partners at rates inconsistent with objective quality, coverage, or relevance benchmarks. A feasible audit uses repeated measurements over a sampling window; controls for region, tier, and request classes; and flags sustained deltas above a practical threshold (e.g., >15% in p95/p99 or denial rates over a rolling 7-day window) absent documented objective criteria. This does not require private user data or model-weight disclosure. Foreclosure concerns do not require strict indispensability in an essential-facilities sense; the question is whether conduct materially raises rivals’ effective costs or degrades their ability to compete. [6–8]

4 Theory of Harm and Mechanisms

Foreclosure in inference markets is unlikely to appear as predatory pricing or explicit refusal to supply. It may appear as small, persistent degradations and steering that shift demand over time. Raising-rivals’-costs (RRC) theory provides a framework: a vertically integrated firm can profit by increasing rivals’ effective costs (or degrading their quality) and recapturing demand, even without below-cost pricing. [7, 8] Platform economics strengthens this logic in two-sided and ecosystem markets, where access rules and visibility/routing conditions can raise rivals’ costs without changing a headline price. [9, 10] Related work on competing ecosystems predicts that total exclusion may be suboptimal; keeping rivals “living but not thriving” can preserve ecosystem breadth while relaxing competition against the gatekeeper’s own downstream complements. [12] Self-preferencing can likewise be implemented through product design that degrades interoperability or selectively withholds advanced features. [11]

Under cognitive-infrastructure conditions, three mechanisms matter: (i) *QoS discrimination* (latency, throughput, error rates), where persistent QoS deltas reduce rival quality and force costlier mitigations; (ii) *feature gating* (context, tool use, enterprise wrappers), where capabilities or timing advantages are withheld from similarly situated third parties; and (iii) *steering and default routing* in agentic flows, where opaque selection reduces rivals’ effective access to demand absent auditable routing artifacts. First-party latency advantages can also reflect legitimate architectural efficiency (colocation, shared caches, fewer hops). The audit protocol must therefore distinguish architectural efficiency from artificial degradation by controlling for region/topology and focusing on residual, persistent wedges absent objective criteria.

5 Neutral Inference: Targeted, Auditable Neutrality Duties

Neutral Inference is a conduct approach designed to preserve contestability in downstream markets by constraining discrimination at a bottleneck input only when observable evidence indicates functional gatekeeper status.

Institutional fit. U.S. antitrust enforcers are “enforcers, not regulators,” and courts are reluctant to impose broad ex ante access




Figure 1: Inference bottleneck architecture. A provider controlling model serving and an API gateway/router can implement non-price discrimination and steering through (i) QoS knobs (queue priority, rate limits), (ii) feature gating (wrappers/version access), and (iii) routing/steering (tool eligibility/defaults), without changing headline API prices.

mandates. [1] Trinko reinforces limits around compelled dealing as a general remedy. [18] The EU has adopted contestability and fairness duties for gatekeepers under the DMA, including constraints relevant to self-preferencing and steering. [14, 15] The AI Act and the GPAI Code of Practice create transparency and risk-management artifacts that can support technical auditing without requiring disclosure of proprietary weights. [16, 17]

Positioning. Neutral Inference is best read as an operational design pattern: it translates non-discrimination and anti-steering concerns into auditable technical artifacts (QoS parity tests and routing logs). Routing-log transparency concerns selection outcomes and eligibility constraints, not chain-of-thought disclosure.

Designation triggers. Dependency and switching costs across a meaningful set of downstream products; sustained performance or governance gaps versus close substitutes on metrics customers buy; and vertical expansion into distribution through major downstream complements that compete with API customers and shape routing. [1, 3, 4]

Three obligations. (1) *QoS parity*: no materially better inference QoS for first-party complements than for similarly situated third parties, after controlling for objective cost and risk drivers, with standardized aggregated disclosures. (2) *Routing transparency*: auditable routing records capturing eligible tool sets, selected tools (and rank where applicable), and whether commercial relationships constrained eligibility or influenced ranking. (3) *FRAND-style non-discrimination*: non-discriminatory access terms for similarly situated buyers, with fast dispute resolution. A workable enforcement model resembles compliance auditing more than decade-long litigation: periodic audits; standardized test suites for QoS parity and routing bias; interim measures when discrimination is detected; and penalties tied to inference revenue.

6 Conclusion

The central competitive risk in the next phase of AI is not only concentration in training. It is the inference bottleneck: scalable cognition integrated into downstream products and mediated through non-price discrimination and opaque routing. Neutral Inference offers a narrow response—QoS parity, routing transparency, and FRAND-style non-discrimination—applied only when observable

evidence indicates functional gatekeeper status. The objective is not to punish scale. It is to preserve contestability before the infrastructure layer quietly acquires the power to determine which applications can exist.

References

- [1] Melissa Holyoak, FTC Commissioner, *Antitrust and AI: Foundational Principles Meet Foundation Models* (remarks; discusses “open early / closed late”).
- [2] OpenAI, *API Pricing* (webpage), accessed 13 Jan. 2026.
- [3] OECD, *Artificial Intelligence and Competitive Dynamics in Downstream Markets*, DAF/COMP/GF(2025)4 (7 Nov. 2025).
- [4] OECD, *Competition in Artificial Intelligence Infrastructure* (background note for OECD Competition Committee, Dec. 2025).
- [5] Holyoak, *supra* note 1 (market-definition ambiguity; AI as input and competitor across sectors).
- [6] “Aplicación de la teoría de raising rivals’ costs” (uploaded memo), *raising rivals’ costs.docx*.
- [7] Steven C. Salop & David T. Scheffman, Raising Rivals’ Costs, 73 *American Economic Review* 267 (1983).
- [8] Steven C. Salop & David T. Scheffman, Cost-Raising Strategies, 36 *Journal of Industrial Economics* 19 (1987).
- [9] Mark Armstrong, Competition in Two-Sided Markets, 37 *RAND Journal of Economics* 668 (2006).
- [10] Jean-Charles Rochet & Jean Tirole, Platform Competition in Two-Sided Markets, 1 *Journal of the European Economic Association* 990 (2003).
- [11] Massimo Motta, Self-Preferencing and Foreclosure in Digital Markets (2023), *International Journal of Industrial Organization*.
- [12] Michele Bisceglia & Andrea Piccolo, Vertical Integration, Innovation and Foreclosure with Competing Ecosystems (2022), *International Journal of Industrial Organization*.
- [13] Holyoak, *supra* note 1 (FTC/DOJ as enforcers, not regulators; reluctance to ex ante rules).
- [14] Regulation (EU) 2022/1925 (Digital Markets Act).
- [15] European Commission, Commission opens non-compliance investigations against Alphabet, Apple and Meta under the Digital Markets Act (Press release, 25 Mar. 2024).
- [16] Regulation (EU) 2024/1689 (AI Act) (Official Journal, 13 Jun. 2024).
- [17] European Commission, The General-Purpose AI Code of Practice (published 10 Jul. 2025).
- [18] *Verizon Communications Inc. v. Law Offices of Curtis V. Trinko, LLP*, 540 U.S. 398 (2004).