

THE NEXT FIVE YEARS OF COMPUTATIONAL ANTITRUST



BY THIBAUT SCHREPEL¹

¹ Thibault Schrepel is an Associate Professor of Law at Vrije Universiteit Amsterdam and the Director of Stanford Computational Antitrust at Stanford University's CodeX Center. The author acknowledges funding from the European Research Council under the ATLANTIS project (ERC Consolidator Grant no. 101228709). Project outputs are collected at teamatlantis.eu.

CPI ANTITRUST CHRONICLE

September 2026

THE NEXT FIVE YEARS OF COMPUTATIONAL ANTITRUST

By Thibault Schrepel



WHEN TWO AI AGENTS TALK: A GAP IN DETECTION CAPABILITIES

By Alba Ribera Martínez



WHEN INNOVATION COMPETITION HAS NO PRODUCT YET: MAKING GENERAL INNOVATION COMPETITION OPERATIONAL

By Mariateresa Maggiolino



COMPUTATIONAL ANTITRUST FOR COMPLEX ADAPTIVE MARKETS

By Filip Lubinski



INSIDE THE BLACK BOX: A PRACTITIONER'S GUIDE TO MACHINE LEARNING USED IN PRICING ALGORITHMS

By Ted Tatos & Hal Singer



ALGORITHMIC COLLUSION OR ALGORITHMIC CONFUSION? RETHINKING THE RISKS OF AI PRICING FOR COMPANIES AND POLICYMAKERS

By Ai Deng



THE NEXT FIVE YEARS OF COMPUTATIONAL ANTITRUST

By Thibault Schrepel

In five years, computational antitrust has gone from a proposition to a field. More than eighty agencies now use computational tools, a literature has formed around them, a scattered set of researchers and enforcers has come to recognize itself as one. The next five years will decide what kind of field it becomes.

Visit www.competitionpolicyinternational.com for access to these articles and more!

CPI Antitrust Chronicle September 2026

www.competitionpolicyinternational.com

Scan to Stay Connected!

Scan or click here to sign up for CPI's **FREE** daily newsletter.



Computational antitrust has stopped being a proposition to defend. It is now a field to reckon with. More than eighty antitrust agencies across six continents take part in the Stanford Computational Antitrust project, and many of them run these tools in live cases.² The OECD has devoted successive roundtables to the questions the subject raises, among them Data Screening Tools for Competition Investigations in 2022, Algorithmic Competition in 2023, and Artificial Intelligence, Data and Competition in 2024.³ The International Competition Network has set up a dedicated group of agency technologists.⁴ The G7 competition authorities now meet in a standing summit, and their 2024 Rome Communiqué commits them to strengthening their digital capacity.⁵

Agencies have convened their own public events. The United Kingdom's Competition and Markets Authority and South Africa's Competition Commission have organized panels on the subject, the Saudi Competition Agency has hosted a dedicated workshop, and Poland's UOKiK has run one on market monitoring, with the Greek, Australian, Swiss, Portuguese, and Hong Kong agencies among the others holding public events.⁶ This special issue, which Competition Policy International has devoted to the subject, is the latest sign.

The question is no longer whether antitrust will become computational. It already has. The question is what kind of computational antitrust we build.

This piece does two things. Part I looks back at what the last five years produced. Part II looks forward to the questions that now deserve attention. The early years of any research area are spent exploring widely. The next stage calls for sustained work on a few hard questions.

I. WHAT THE LAST FIVE YEARS PRODUCED

Five years is enough time to ask whether something has been built.

The starting wager was narrow. It held that the central difficulty in antitrust had become computational before it had become doctrinal. Enforcers were not short of legal theories. They were short of the capacity to process the scale and speed of the markets those theories were meant to govern. The moment suited the claim. Data had grown abundant while machine learning matured, and agencies were beginning to feel the limits of their existing methods.

The Stanford Computational Antitrust project did more than supply a name. Computational antitrust is defined as the use of computational methods in the analysis and enforcement of competition law. It runs from the screens agencies build to flag collusion in procurement data to the simulations researchers run on markets. What the project added was a research agenda and a venue. A team of editors, contributors, academics on the advisory board and agency partners let common problems surface across contributions, so that scholarship began to answer scholarship rather than accumulate in isolation. Even so, we were pushing on a door the moment had already begun to open.

The wager is now old enough to judge. A name and a venue are not yet a field, and neither is a vocabulary. The claim that computational antitrust has become one must be tested rather than asserted. So the sections that follow weigh the evidence. Section A asks whether there is a literature. Section B asks whether there is an international conversation. Section C asks whether there is a practice inside the agencies. Section D puts the three together. It asks whether they amount to a field rather than a topic. I argue that they do.

2 Stanford Computational Antitrust, CodeX, Stanford Law School, <https://law.stanford.edu/computationalantitrust>. The list of participating agencies is available at <https://law.stanford.edu/codex-the-stanford-center-for-legal-informatics/projects/computational-antitrust/computational-antitrust-agencies/>.

3 OECD, "Data Screening Tools for Competition Investigations", OECD Roundtables on Competition Policy Papers No. 284, OECD Publishing, Paris (2022), <https://doi.org/10.1787/4c5bbb9d-en>; OECD, "Algorithmic Competition", OECD Competition Policy Roundtable Background Note (2023); OECD, "Artificial Intelligence, Data and Competition", OECD Competition Committee Roundtable, 12 June 2024; OECD, "Competition in Artificial Intelligence Infrastructure", OECD Roundtable, 4 December 2025.

4 International Competition Network, Technologists Network, <https://www.internationalcompetitionnetwork.org/working-groups/icn-operations/technologists/>.

5 G7 Digital and Tech Ministers' Meeting, Ministerial Declaration, Takasaki, 30 April 2023, paras 49-54, <https://g7g20-documents.org/database/document/2023-g7-japan-ministerial-meetings-ict-ministers-ministers-language-ministerial-declaration-the-g7-digital-and-tech-ministers-meeting>, resolving to convene a summit of competition authorities and policymakers and to establish a point of contact group among them. The first such summit was held in Tokyo in November 2023 and the sixth G7 Digital Competition Summit in Rome on 3 and 4 October 2024, which concluded with the G7 Digital Competition Communiqué, among whose commitments is the strengthening of the authorities' digital capacity.

6 Among others: "Computational Antitrust", Competition and Markets Authority, United Kingdom (April 2023), and "The Use of Agent-Based Modeling in Antitrust", CMA (June 2025); Computational Antitrust workshop, Saudi Competition Agency (December 2025); UOKiK Workshop on Market Studies, Monitoring and Digital Markets, Warsaw (December 2023); "Computational Antitrust", Germany Competition Day, Düsseldorf (September 2023); "The Promise of Computational Antitrust", Hellenic Competition Commission (2021); Australian Competition and Consumer Commission (November 2023); Swiss Competition Commission (November 2024); Portuguese Competition Authority, Open Seminar Series (November 2025); Hong Kong Competition Commission (May 2025).

A. A Literature

The computational antitrust literature divides into several kinds.

One group of papers works on the doctrine. It asks what computational methods do to the categories antitrust already uses. Anthony Casey and Anthony Niblett argued that predictive technologies could recast merger notification and review, so that bright-line thresholds give way to a micro-directive tailored to the transaction.⁷ Cary Coglianese and Alicia Lai set out how machine learning could help antitrust agencies detect evasive anticompetitive conduct. They insist that officials weigh case by case whether an algorithm should displace human judgment.⁸ Robert Mahari, Sandro Lera, and Alex Pentland argue that data ownership has become a better signal of monopoly power than market share or price, and they propose a network-science measure of a firm's capacity to turn size into growth as an early warning for anticompetitive mergers.⁹ Herwig Hofmann and Isabella Lorenzoni examine what automated decision-making does to basic principles of law, from the duty to give reasons to the right to be heard.¹⁰

A further group builds the instruments themselves. Joseph Harrington and David Imhof showed how machine learning can screen public tenders for cartels.¹¹ Klaus Gugler, Florian Szücs, and Ulrich Wohak used natural language processing, which is the automated reading of ordinary text, to define relevant markets from firms' own business descriptions. They found their results close to the European Commission's expert assessments across a sample of start-up acquisitions.¹² Yann Guthmann, Adrien Frumence, and Camille Hoogterp built a tool at the French Competition Authority that maps citations across its own decisions as a network.¹³ Alba Ribera Martínez proposes computational presumptions, grounded in measurable privacy thresholds, as an alternative to blunt data-combination bans in AI markets.¹⁴ Philip Hanspach complements these contributions by identifying the conditions under which such instruments should inform legal decision-making.¹⁵ He argues that machine-learning evidence requires transparent methodologies.

A third group takes the institutions themselves as its subject. It asks what adoption does to an agency rather than to a doctrine. This is where computational antitrust has drawn institutional economics into its core, since it treats enforcement bodies as organizations with their own constraints and their own capacity to absorb new methods, rather than as neutral appliers of rules. Renato Nazzini and James Henderson argue that computational antitrust requires competition authorities to develop new investigative capabilities, data infrastructures, and international cooperation to reduce the empirical uncertainty surrounding algorithmic pricing.¹⁶ Jorge Padilla turns to the psychology of the enforcers themselves. He argues that computational tools can help agencies correct for motivated skepticism when the tools acknowledge uncertainty and are revised as new information arrives.¹⁷

The term has traveled well beyond any single journal. A search on Google Scholar returns more papers using it than can readily be counted, and they take the concept as a settled starting point rather than a proposal.¹⁸ Will Carpenter and his co-authors have examined machine learning in cartel damages estimation.¹⁹ James Andrews and Peter Ormosi have reviewed a quarter-century of merger simulations in Commis-

7 Anthony J. Casey & Anthony Niblett, "Micro-Directives and Computational Merger Review", I Stanford Computational Antitrust 132 (2021).

8 Cary Coglianese & Alicia Lai, "Antitrust by Algorithm", II Stanford Computational Antitrust 1 (2022).

9 Robert Zev Mahari, Sandro Claudio Lera & Alex Pentland, "Time for a New Antitrust Era: Refocusing Antitrust Law to Invigorate Competition in the 21st Century", I Stanford Computational Antitrust 52 (2021).

10 Herwig C.H. Hofmann & Isabella Lorenzoni, "Future Challenges for Automation in Competition Law Enforcement", III Stanford Computational Antitrust 37 (2023).

11 Joseph E. Harrington Jr. & David Imhof, "Cartel Screening and Machine Learning", II Stanford Computational Antitrust 133 (2022).

12 Klaus Gugler, Florian Szücs & Ulrich Wohak, "Using Natural Language Processing to Delineate Digital Markets", IV Stanford Computational Antitrust 33 (2024).

13 Yann Guthmann, Adrien Frumence & Camille Hoogterp, "Deploying Network Analysis in Antitrust Law", III Stanford Computational Antitrust 1 (2023).

14 Alba Ribera Martínez, "Computational Presumptions Applied to AI Markets", VI Stanford Computational Antitrust 32 (2026).

15 Philip Hanspach, "Economics in the Era of Machine Learning: What Do Competition Lawyers Need to Know?", IV Stanford Computational Antitrust 175 (2024).

16 Renato Nazzini & James Henderson, "Overcoming the Current Knowledge Gap of Algorithmic 'Collusion' and the Role of Computational Antitrust", IV Stanford Computational Antitrust 1 (2024).

17 Jorge Padilla, "The Psychology of Competition Agencies: Using Computational Tools to Address Motivated Skepticism", V Stanford Computational Antitrust 152 (2025).

18 The count fluctuates and is sensitive to how the query is framed. The point is one of magnitude rather than precision.

19 Will Carpenter, Anna Lane, Joshua Hia, Steffen Reinhold, Iain Boa & Martin Spindler, "Machine Learning in Cartel Damages Estimation: Challenges and Opportunities", 17 Journal of European Competition Law & Practice 46 (2026).

sion decisions.²⁰ Victoriia Noskova and Oliver Budzinski have mapped computational tools across each step of merger review.²¹ These are just some examples among many.

B. An International Conversation

The institutions that convene the debate have done a great deal to move it forward. The OECD has returned to the subject year after year. Its roundtable on Data Screening Tools for Competition Investigations examined the empirical methods agencies use to detect cartels in procurement data.²² Its roundtable on Algorithmic Competition explored how computational techniques can strengthen enforcement against anti-competitive algorithmic practices.²³ The subsequent roundtable on Artificial Intelligence, Data and Competition extended that agenda by analyzing AI markets themselves and by identifying concrete applications of AI for competition authorities.²⁴ The OECD also hosted dedicated computational antitrust events in 2023, 2024, and 2025, the last of which saw the launch of the fourth cross-agency report.²⁵ In the meantime, the International Competition Network has set up a dedicated group of agency technologists that allows agencies to share computational approaches across jurisdictions rather than reinventing them independently.²⁶

The agencies are also organizing among themselves. Austria's Federal Competition Authority hosted an AI-focused workshop within the European Competition Authorities meeting in Vienna in February 2025. The workshop brought together more than eighty experts from participating agencies and DG Competition.²⁷ The Italian authority participates in the European Competition Network working group on digital investigation and artificial intelligence and in the International Competition Network Technologist Forum.²⁸ Hungary's GVH submitted a joint project with other EU authorities under the European Commission's Technical Support Instrument to support digital transformation in competition law enforcement and provide staff with up-to-date knowledge on the use of AI.²⁹ In the Americas, COFECE (now the Comisión Nacional Antimonopolio) reports working with fellow agencies to launch the Group of Competition Agencies of America, where COFECE and CADE co-chair an agency-only working group on digital markets and digitalization.³⁰ Canada's Competition Bureau co-founded the OECD Working Group on Behavioral Science and Competition.³¹ None of these bodies existed when the first report was published.

20 James Andrews & Peter L. Ormosi, "Merger Simulations in EU Merger Control: What Have We Learned?", Centre for Competition Policy Working Paper 25-03 (2025).

21 Victoriia Noskova & Oliver Budzinski, "Computational Methods in the Evaluation of Mergers and Acquisitions", in Artificial Intelligence & Competition Policy 231 (Alden Abbott & Thibault Schrepel eds., Concurrences 2024).

22 OECD, *Data Screening Tools for Competition Investigations* (OECD Roundtables on Competition Policy Papers No 284, OECD Publishing 2022) <https://doi.org/10.1787/4c5bbb9d-en>.

23 OECD, *Algorithmic Competition* (OECD Competition Policy Roundtable Background Note, OECD 2023) https://www.oecd.org/content/dam/oecd/en/publications/reports/2023/05/algorithmic-competition_2be02d00/cb3b2075-en.pdf.

24 OECD, *Artificial Intelligence, Data and Competition* (OECD Artificial Intelligence Papers No 18, OECD Publishing 2024) https://www.oecd.org/en/publications/2024/05/artificial-intelligence-data-and-competition_9d0ac766.html.

25 "Computational Antitrust", OECD Competition Committee, Paris (June 2023); "Computational Antitrust", OECD, Paris (June 2024); "The Computational Antitrust Event", OECD, Paris (18 June 2025). See CodeX, "Computational Antitrust Project at the OECD", Stanford Law School (25 September 2025), <https://law.stanford.edu/2025/09/25/codex-computational-antitrust-project-at-the-oecd/>.

26 International Competition Network, Technologists Network, <https://www.internationalcompetitionnetwork.org/working-groups/icn-operations/technologists/>.

27 Austrian Federal Competition Authority, contribution to the Fourth Cross-Agency Report. It describes an AI-focused workshop held within the framework of the European Competition Authorities (ECA) Meeting, Vienna, February 2025, which brought together over eighty experts from participating competition authorities, including DG Competition.

28 Autorità Garante della Concorrenza e del Mercato (Italy), contribution to the Fourth Cross-Agency Report. It reports that its Data Science Unit participates in the European Competition Network working group on digital investigation and artificial intelligence and in the International Competition Network Technologist Forum.

29 Gazdasági Versenyhivatal (Hungary), contribution to the Fourth Cross-Agency Report. It features a joint project submitted with other EU competition authorities under the European Commission's 2025 Technical Support Instrument programme, entitled "Supporting Digital Transformation in Competition Law Enforcement". The Italian authority reports its engagement in a Technical Support Instrument project under the same banner. In its contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, the GVH describes the project as providing staff with up-to-date knowledge on the use of artificial intelligence, with the first sessions scheduled to start in the summer of 2026.

30 Comisión Federal de Competencia Económica (Mexico), contribution to the Second Cross-Agency Report. It describes the launch of the Group of Competition Agencies of America (GrACA) and an agency-only working group on digital markets and digitalization co-chaired with CADE, itself divided into a sub-group on digital markets and a sub-group on digitalization within the agencies. COFECE also reports working meetings on technology and digitalization held in November 2022 with the Federal Trade Commission, the Antitrust Division of the Department of Justice, and the Canadian Competition Bureau. See also the contribution of the Instituto Nacional de Defensa de la Competencia y de la Protección de la Propiedad Intelectual (Peru) to the Fourth Cross-Agency Report, on the OECD Regional Centre for Latin America, a joint initiative with the OECD running since 2019.

31 Competition Bureau Canada, contribution to the Fourth Cross-Agency Report. It reports that its Behavioural Insights Unit co-founded the OECD Working Group on Behavioural Science and Competition, and that the Bureau participates in the Canadian Digital Regulators Forum alongside the Office of the Privacy Commissioner, the Canadian Radio-television and Telecommunications Commission, and the Copyright Board of Canada.

The cooperation now reaches the level of code. France's Autorité de la concurrence has launched Project Tardis within the ICN Technologist Group, an open-source effort to aggregate and standardize agencies' documents into machine-readable formats.³² Latvia's Competition Council runs a retrieval tool built on source code shared by the Danish authority and has offered its own cartel screen to the Danish and Austrian authorities in return, while Canada's Competition Bureau has shared the source code of Eagle Eye, its geospatial merger-analysis tool, with interested agencies.³³ The European authorities, for their part, now build capacity jointly, from the DATACROS consortium to the Technical Support Instrument projects on bid rigging and digital transformation to the fifteen-agency DICE training program.³⁴

The agency-hosted events named at the outset, from panels in London and Pretoria to a workshop in Riyadh and Paris at the French Competition Authority, were organized on the agencies' own initiative. So were the working groups. In the first report, Romania's Competition Council listed among its lessons learned the need for cooperation between authorities and the regular exchange of information on tested screening methodologies.³⁵ It was a request, addressed to no one in particular. Four reports later, the forums exist. Nobody granted the request. The agencies built it themselves, which is precisely the point.

C. A Practice Inside the Agencies

The most instructive record is the one the agencies wrote themselves. An annual conference at Stanford has gathered this work since 2021, the first edition of which, *Exploring Antitrust 3.0*, drew agencies from many countries.³⁶ The project's podcast invites the people who run the agencies to explain their practice in their own words, with recent guests including the heads of the Spanish, Brazilian, Dutch, Portuguese, Austrian, Singaporean, Egyptian, and Taiwanese agencies.³⁷

Each year since 2021, the participating agencies have been asked to describe what they have built.³⁸ Read in sequence, the four reports show how fast the ground has moved. In the first, a good deal of what agencies reported was the groundwork of digitization. Argentina's competition commission described the shift from paper files to electronic records as a foundational step.³⁹ By the Fourth Cross-Agency Report, which gathers twenty-five agencies, the reported work had changed character entirely.⁴⁰ The fifth, published in September 2026, gathers some thirty.⁴¹

The range of what agencies now report is wide. Machine-learning screens flag potentially collusive tenders, and Spain's BRAVA is a leading example, trained on previously sanctioned cases and built so that every flag carries an account of the reasoning that produced it, for

32 Autorité de la concurrence (France), contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 208.

33 Competition Council (Latvia), contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 220. It describes a Retrieval Augmented Generation tool built on source code shared by the Danish Competition and Consumer Authority. The Council reports that its cartel screening tool, developed with the University of Latvia, Riga Technical University, and the Corruption Prevention and Combating Bureau, has been offered to the Danish and Austrian authorities for testing on their own datasets; Competition Bureau Canada, contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 186. It showcases "Eagle Eye", a geospatial application that automates local market analysis using drive times based on the road network. The Bureau has shared its source code with interested agencies.

34 Úřad pro ochranu hospodářské soutěže (Czech Republic), Croatian Competition Agency, Lithuanian Competition Council, and Antimonopoly Office of the Slovak Republic, contributions to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, on the DATACROS III consortium, the Technical Support Instrument projects "Strengthening Detection and Reporting of Bid Rigging in Czechia, France, Ireland, Latvia, Poland, and Portugal" and "Supporting Digital Transformation in Competition Law Enforcement", the "DICE" training project gathering fifteen national competition authorities, and the transfer of web-scraping know-how from the Czech to the Slovak authority.

35 Consiliul Concurenței (Romania), contribution to the First Cross-Agency Report. It lists among its lessons learned the "need for a specific cooperation between competition authorities" and the "regular exchange of information on valid and tested screening methodologies".

36 "Computational Antitrust: Exploring Antitrust 3.0", Stanford University (2021), <https://conferences.law.stanford.edu/computational-antitrust-exploring-antitrust/>. The conference has been held annually since.

37 Stanford Computational Antitrust podcast, https://www.youtube.com/channel/UCSr7ZOKjCHzf6_yYV69CKlg. Recent guests include Cani Fernández (President, CNMC, Spain), Gustavo Augusto Freitas de Lima (Interim President, CADE, Brazil), Martijn Snoep (Chairman, ACM, Netherlands), Nuno Cunha Rodrigues (President, Autoridade da Concorrência, Portugal), Natalie Harsdorf (Director General, Federal Competition Authority, Austria), Alvin Koh (Chief Executive, CCCS, Singapore), Mahmoud A. Montaz (Chairman, Egyptian Competition Authority), and Andy Chen (Vice Chair and Acting Chair, Taiwan Fair Trade Commission).

38 The annual reports reflect the practices of participating agencies. At the time of writing, the European Commission has chosen not to participate in the Stanford Computational Antitrust project and is therefore not represented in these reports.

39 "Computational Antitrust: First Annual Report", II Stanford Computational Antitrust (2022), reporting on agency activity in 2021, including the contribution of the Comisión Nacional de Defensa de la Competencia (Argentina) on the transition to electronic case files.

40 Thibault Schrepel & Teodora Groza (eds.), "Computational Antitrust Worldwide: Fourth Cross-Agency Report", V Stanford Computational Antitrust 1 (2025). See also the second and third cross-agency reports (2023, 2024), available at <https://law.stanford.edu/computationalantitrust>.

41 Thibault Schrepel & Alba Ribera Martínez (eds.), "Computational Antitrust Worldwide: Fifth Cross-Agency Report", V Stanford Computational Antitrust 167 (2026).

analysts and eventually for courts.⁴² Catalonia's competition authority runs ERICCA, which clusters companies by their behavior in public tenders to surface likely collusion. Brazil's CADE has run Cérebro for years to mine public procurement data for bid-rigging patterns. Its screens have now carried the Novo Rumo bid-rigging investigation from detection to court-upheld search warrants and formal proceedings. CADE has since turned them on algorithmic fare-setting in the airline industry.⁴³ Colombia's Superintendence of Industry and Commerce runs a family of tools with names like Sabueso, Sherlock, and Búho, which scrape and normalize prices across supermarkets and airlines. Chile's FNE monitors procurement at the regional level through data-sharing agreements. Singapore's competition commission has built a toolkit that tests AI systems for explainability and fairness as part of a compliance program.⁴⁴ Its bid-rigging screen has since sifted 114 tenders down to six suspect firms in under ten minutes, and experiments run with the city's media regulator have shown LLM pricing agents converging on high prices with no instruction to collude.⁴⁵

Generative AI models are now entering daily operations, and the agencies are building rather than buying. Canada's Competition Bureau is developing COMPASS, a secure generative AI platform built in-house so that staff can work on sensitive case material without exposing it to third-party models.⁴⁶ The platform is now being extended with automated data-extraction pipelines, cloud computing, and AI agents.⁴⁷ France's Autorité de la concurrence has built a retrieval-augmented generation tool that lets case handlers query the authority's own decisional corpus in natural language, on top of the open dataset it also publishes.⁴⁸ The tool has since left prototype status. It draws on the Commission's open data and now answers questions about the European Commission's cases from the same interface.⁴⁹ Poland's UOKiK went further still. To build a dark patterns detector, it swept 316 websites to assemble a training set, ran neuromarketing experiments to study how the designs actually affect users, and fine-tuned GPT-4 against its own guidelines.⁵⁰ The detector, DRAKE, went live in March 2026, joined by ARBUZ, a system that screens standard form contracts for abusive clauses by their semantic similarity to a register of court decisions.⁵¹ A competition agency running neuromarketing experiments is a fact worth pausing on, and I return to it below. The Hellenic Competition Commission, for its part, is developing communication-forensics tools that use machine learning to classify email content and graph-based network analysis to identify suspicious patterns of contact.⁵² Graph analytics of this kind map the webs of relationships between firms and bids and individuals, and the more ambitious architectures reach to graph neural networks. The inventiveness runs wider still.

42 Comisión Nacional de los Mercados y la Competencia (Spain), contribution to the Fourth Cross-Agency Report. BRAVA is the Bid-Rigging Algorithm for Vigilance in Antitrust. It is a supervised machine-learning model trained on previously sanctioned cases. It draws on the Spanish public procurement platform and the commercial registries, and covers a database of more than seven million tenders. Its explainability relies on LIME and SHAP techniques, which attribute each prediction to the features that produced it.

43 Conselho Administrativo de Defesa Econômica (Brazil), contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 179. In the Novo Rumo operation, Cérebro combined network analysis, document-metadata analysis, and unsupervised machine learning to support a judicial search-and-seizure warrant against twelve firms, later upheld against defense challenges. In December 2025, an Administrative Proceeding was opened against sixteen companies and fifteen individuals over an estimated US\$2.2 billion in affected tenders. In the airline investigation, time-series and complex-systems methods identified persistent interdependence between GOL's and LATAM's pricing trajectories. Contractual analysis then identified shared providers of tariff intelligence and dynamic-pricing services as a centralized informational hub. The Commission opened a formal Administrative Proceeding.

44 Fourth Cross-Agency Report, contributions of the Autoritat Catalana de la Competència (Catalonia), the Superintendencia de Industria y Comercio (Colombia), the Fiscalía Nacional Económica (Chile), and the Competition and Consumer Commission of Singapore.

45 Competition and Consumer Commission of Singapore, contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 267. It reports that its Bid-Rigging Detection Tool sifted 114 tenders and identified six suspicious businesses in under ten minutes, work that would otherwise have taken weeks. The Commission also describes experiments run with the Infocomm Media Development Authority. Large language models prompted only to maximize profits converged on coordinated high prices, and no collusive intent appeared in their reasoning. For a converging finding from CADE's Cérebro team, see the contribution of the Conselho Administrativo de Defesa Econômica (Brazil). It cites research in which LLM pricing agents in a repeated Bertrand duopoly converge to supracompetitive prices, absent any inter-agent communication channel.

46 Competition Bureau Canada, contribution to the Fourth Cross-Agency Report. It describes COMPASS (Competition AI Secure System), a secure generative AI platform developed in-house to allow staff to work with sensitive material without exposure to third-party large language models.

47 Competition Bureau Canada, contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 186.

48 Autorité de la concurrence (France), contribution to the Fourth Cross-Agency Report. It showcases a Retrieval Augmented Generation tool that allows case handlers to query the Autorité's own decisional corpus in natural language, alongside a full-text search engine built on the same dataset. The underlying data is released as open data.

49 Autorité de la concurrence (France), contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 208.

50 Urząd Ochrony Konkurencji i Konsumentów (Poland), contribution to the Fourth Cross-Agency Report. It describes an internet sweep of 316 websites to build a training dataset, neuromarketing experiments to study how dark patterns affect users, and a fine-tuned GPT-4 model to detect dark patterns against defined guidelines.

51 Urząd Ochrony Konkurencji i Konsumentów (Poland), contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 257. The Office describes two systems. DRAKE, a dark-pattern detection system deployed in March 2026, combines structural website analysis, behavioral research, and code-level examination. The behavioral research draws on eye-tracking, facial expression analysis, and EEG-based measurement. ARBUZ is an AI-supported system that detects abusive clauses in standard-form contracts through semantic similarity to a register of court decisions.

52 Hellenic Competition Commission (Greece), contribution to the Fourth Cross-Agency Report, on machine-learning classification of email content and graph-based network analysis to identify suspicious communication patterns, together with the entity-resolution and domain-attribution problems these raise.

Hungary's GVH operates Cartel Chat, an anonymous reporting channel, and Italy's AGCM has reverse engineered the ranking algorithms of hotel and e-commerce platforms in live cases.⁵³

The Stanford Computational Antitrust project's fifth report extends the record. Pakistan's competition commission has built a bid-rigging system that joins statistical screens and machine learning to a graph of directors and shareholders. The graph exposes corporate families that split bids across nominally competing firms.⁵⁴ Kazakhstan's agency credits its Ormek system with more than doubling cartel investigations, and it has cut the identification of collusion indicators from months to days.⁵⁵ Portugal's authority runs a family of in-house tools.⁵⁶ They scrape online markets without code, they screen procurement for staff who do not program, and they read the daily news for mergers nobody notified. Detection of non-notified mergers is itself becoming a shared capability, built independently by Portugal, the Netherlands, Latvia, and Australia in the same year.⁵⁷ Catalonia has extended ERICCA's data to cover the whole of Spain. It is developing AI tools that anonymize documents, that review seized WhatsApp exchanges, and that screen regulation itself for restrictions on competition.⁵⁸ Taiwan's commission scraped reservoir water levels to test a suspected cartel among cruise-boat operators, and it mapped asphalt delivery ranges to draw geographic markets.⁵⁹ The breadth now runs from deep-learning screens across 555 sectors, to publicly accessible screening tools, to agencies that report no computational tools at all.⁶⁰

The distance between the two ends of that sequence is the real story of the past five years, and the credit belongs to the agencies. The scale has become institutional. The Dutch authority runs a forty-five-person data taskforce under a chief data officer, and Australia's commission counts some fifty staff working regularly in data and code.⁶¹ In several places, they are now building ahead of the academic literature rather than behind it. The reports are also candid about what remains hard. The fourth names three obstacles the agencies themselves identify. Secure cloud computing is one. The explainability of deployed tools is another. The quality of the interaction between people and machines is the third.⁶²

53 Gazdasági Versenyhivatal (Hungary), contribution to the Fourth Cross-Agency Report, on the Cartel Chat anonymous reporting platform and the Virtual Data Room. Autorità Garante della Concorrenza e del Mercato (Italy), contribution to the Fourth Cross-Agency Report, on the analysis of booking and pricing algorithms of the Italian railways and the reverse engineering of ranking algorithms used by hotel and e-commerce platforms.

54 Competition Commission of Pakistan, contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 249. It describes the Bid-Rigging Detection System. The system combines statistical screens with a random forest classifier and a graph that links procurement records to the corporate registry. It detects collusion across director networks and bid-splitting within corporate families that investigators could not previously see.

55 Agency for Protection and Development of Competition (Kazakhstan), contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 217. It reports that the "Ormek" search system raised the number of cartel investigations from 67 in 2018–2020 to 139 by 2025, and that it cut the identification of indirect collusion indicators from one and a half to two months down to one or two days.

56 Autoridade da Concorrência (Portugal), contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 262. It describes four in-house tools. Scrap-IT is a no-code web scraper with AI-assisted data structuring. Screen-IT is a public procurement screening platform for staff without programming knowledge. Detect-IT is an AI-based tool that cross-checks a daily news feed against financial data to flag potentially non-notified mergers. Foreense++ cut forensic file certification from roughly two hours to twenty minutes.

57 On the detection of non-notified mergers, see the contributions to the Stanford Computational Antitrust project's Fifth Cross-Agency Report of the Autoridade da Concorrência (Portugal), on Detect-IT; the Autoriteit Consument & Markt (Netherlands), on the LLM-based screening of news sources for concentration activity; the Competition Council (Latvia), on a merger monitoring tool that tracks all company ownership changes lodged in state registers, and that flagged thirty cases for further inspection in 2025; and the Australian Competition and Consumer Commission, on merger surveillance drawn from administrative data and public stock market announcements.

58 Autoritat Catalana de la Competència (Catalonia), contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 195, on the extension of ERICCA to data covering the whole of Spain. The Authority is also developing AI tools that anonymize documents, that review WhatsApp and other messaging exchanges obtained during inspections, and that screen regulatory provisions for unjustified restrictions on competition.

59 Taiwan Fair Trade Commission, contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 275. It showcases several applications. These include the scraping of historical reservoir water-level data in a suspected concerted action among reservoir cruise-boat operators, a cointegration analysis of corn and feed prices in the Feed Industry Association case, and the GIS mapping of asphalt plants' delivery ranges to delineate geographic markets.

60 Turkish Competition Authority, contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report. The Authority describes a deep-learning model that monitors producer price indices across 555 sectors, and an AI-supported public procurement project that produces a cartel risk score for each completed electronic tender. Comisión Nacional de Defensa de la Competencia (Dominican Republic), contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report. The Commission describes the Semáforo Colusorio, a publicly accessible screening tool that processes forty-nine variables to classify procurement procedures into five risk-alert levels. Georgian Competition and Consumer Protection Agency, Fair Competition Commission (Tanzania), and Trinidad and Tobago Fair Trading Commission, contributions to the Stanford Computational Antitrust project's Fifth Cross-Agency Report. They report respectively the absence of computational tools due to budgetary constraints, foundational digital infrastructure work, and early exploratory use of generative AI tools.

61 Autoriteit Consument & Markt (Netherlands), contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 245. Australian Competition and Consumer Commission, contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 169. The Commission reports around fifty staff who regularly use data and code, alongside approximately thirty economists and ten data engineers.

62 Fourth Cross-Agency Report, above. The three obstacles most commonly reported are secure cloud computing, the explainability of deployed tools, and the quality of the interaction between people and machines.

That practice has already changed what antitrust is, though the change is not always visible. It rarely announces itself in a new legal test or a headline decision. It shows up in which cases get opened and in which patterns become perceptible at all. Above all, it shows up in what an agency will treat as evidence.

D. All in All: A Field

The three preceding sections can be taken together. There is a literature. There is an international conversation. There is a practice inside the agencies. What remains is to ask whether the sum of them amounts to a field rather than a topic.

The sociology and philosophy of science have asked that question for decades, and their answers converge. Looking at it from different angles, computational antitrust is now a field indeed. Pierre Bourdieu located a field in its autonomy, the point at which a body of producers begins to work for one another rather than for an outside audience.⁶³ Computational antitrust has crossed that line. Its scholarship now answers scholarship, and its practitioners write first for each other.⁶⁴ Diana Crane described the same passage as the formation of an invisible college, the communication network through which a scattered set of researchers and enforcers comes to recognize itself as one.⁶⁵ The cross-agency reports and the standing forums are that network made visible. Tony Becher and Paul Trowler add that a mature discipline is also a culture, with its own tribe and its own territory; the agency technologists now meeting across borders have acquired both.⁶⁶ Further, Stephen Toulmin offered the account that fits this field best. He treats a discipline as an evolving population of concepts and methods, selected and refined as practice tests them.⁶⁷ That last description matters most here, since it fits a vocabulary and a set of tools that change as the work feeds back into them, and I return to it at the close. The remaining marks are the ones Kuhn and Whitley set out.

Thomas Kuhn located a field in its shared exemplars, the worked problems that show new entrants what good work looks like.⁶⁸ Computational antitrust has them, and they are increasingly agency-built rather than academic. BRAVA is one. An enforcer anywhere who wants to build a bid-rigging screen now starts from what the CNMC did, and from the explainability constraints it accepted. Cérebro is another, and ERICCA, and now COMPASS. These are reference designs, adapted rather than reinvented, and their circulation between agencies is the clearest evidence that the work has a shared standard of what counts as done well.⁶⁹

Finally, Richard Whitley pointed to organizational conditions. A field needs dedicated venues. It needs practitioners who depend on each other's judgment, and it needs a stock of problems the group treats as its own.⁷⁰ All three now hold. The venues exist, from a dedicated journal to an annual conference to the OECD roundtables and the ICN Technologist Forum.⁷¹ The mutual dependence is real, since an agency deciding whether to trust an explainability method looks to what other agencies have accepted, and a researcher building a screen looks to what failed in a live case. And the stock of problems is now settled enough that the agencies can name it without prompting. Stanford's fourth report identifies

63 Pierre Bourdieu, "The Specificity of the Scientific Field and the Social Conditions of the Progress of Reason", 14 *Social Science Information* 19 (1975). Bourdieu defines a scientific field through the autonomy of a set of producers who take one another as their primary audience.

64 E.g., Niall MacMenamin, Vendela Fehrm and Zhaoning (Nancy) Wang, 'The Role of AI in Litigation and Competition Expert Analysis: A Conversation with Anindya Ghose' (Compass Lexecon, 25 April 2025) <https://www.compasslexecon.com/insights/publications/the-role-of-ai-in-litigation-and-competition-expert-analysis-a-conversation-with-professor-anindya-ghose>.

65 Diana Crane, *Invisible Colleges: Diffusion of Knowledge in Scientific Communities* (University of Chicago Press, 1972), on the communication networks through which a research effort recognizes itself as one. The term "invisible college" is Derek de Solla Price's, from *Little Science, Big Science* (Columbia University Press, 1963).

66 Tony Becher & Paul Trowler, *Academic Tribes and Territories: Intellectual Enquiry and the Cultures of Disciplines* (2nd ed., Open University Press, 2001), treating disciplines as cultures with their own norms and boundaries.

67 Stephen Toulmin, *Human Understanding, Volume I: The Collective Use and Evolution of Concepts* (Princeton University Press, 1972), setting out a populational account of a discipline as an evolving body of concepts and methods.

68 Thomas S. Kuhn, "Second Thoughts on Paradigms", in Frederick Suppe (ed.), *The Structure of Scientific Theories* 459 (University of Illinois Press, 1974), reprinted in *The Essential Tension* 293 (University of Chicago Press, 1977). See also the Postscript to the second edition of *The Structure of Scientific Revolutions* (University of Chicago Press, 1970), where Kuhn separates the "disciplinary matrix" from the exemplars that instantiate it.

69 The exemplars are increasingly agency-built rather than academic. BRAVA, Cérebro, ERICCA, and COMPASS are cited across contributions to the cross-agency reports and function as reference designs that other agencies adapt. The fifth report documents the circulation directly. Latvia's Competition Council runs a tool built on Danish source code and has offered its own screen to the Danish and Austrian authorities. Canada's Competition Bureau has shared Eagle Eye's source code with interested agencies. Slovakia's Antimonopoly Office acquired its web-scraping methods from the Czech authority.

70 Richard Whitley, *The Intellectual and Social Organization of the Sciences* (Clarendon Press, 1984; 2nd ed., Oxford University Press, 2000).

71 The dedicated venues now include the Stanford Computational Antitrust journal, the annual Stanford conference, the OECD roundtables and events, the ICN Technologist Forum, the ECN working group on digital investigation and artificial intelligence, and the cross-agency reports themselves.

the three named earlier, namely secure cloud computing, the explainability of deployed tools, and the quality of human-machine interaction. That is a shared agenda, not a list of individual preoccupations.

By these measures, computational antitrust qualifies as a field. It has its own vocabulary and its own tools, and those tools have entered practice. Practice, in turn, is reshaping the tools, as agencies report back what fails in a real case and researchers rebuild accordingly. Vocabulary and tools and practice and agencies now form a system with feedback running in every direction, which is to say a complex system. It adapts without any single hand steering it. The clearest sign that the work has taken root is that the strongest recent contributions are increasingly not from those who were there at the start.

II. WHAT THE NEXT FIVE YEARS WILL DECIDE

So there is a field. That settles the backward-looking question and opens the forward-looking one. A field that exists is not a field whose direction is fixed, and the shape computational antitrust takes for the years ahead is still open. The five items that follow are the ones that will decide it. Each is a question the field now has to answer for itself, and the answer given to each will define what computational antitrust becomes. Section A asks how to build the legal framework computational tools require. Section B asks how the field earns legitimacy in its instruments. Section C asks what antitrust is, and should be. Section D charts the cognitive frontier. Section E asks how much of antitrust should be automated.

What follows makes no claim to cover every challenge. It picks out the five I take to be most defining. Both the research and the practice are moving quickly, and the next gain lies in connecting them more tightly.

A. Building the Legal Framework

Antitrust agencies have adopted computational tools faster than the law governing their use has been written. This is not a failure of anyone's diligence. Law moves more slowly than technology. The rules of procedure took decades to settle. So did the standards of proof, and so did the categories through which a practice is characterized and documented. None of them can, nor should, be rewritten in a season, and technology does not wait for them.

The gap between the two is widening, and it cuts in both directions. It holds agencies back. A tool whose output the law offers no settled way to handle is a tool an agency cannot safely deploy, however useful it would be. Careful enforcers are therefore leaving capability unused. Where tools are already in use, the same gap opens loopholes in how decisions get constructed, since the safeguards that would ordinarily govern such evidence have not been articulated. The fifth annual report addresses the practical side of this question by asking agencies what computational methods have enabled them to achieve what would otherwise have been impossible.⁷² The gap can also invert. Where the law is absent altogether, the methods stand in for it. The CARICOM Competition Commission had no jurisdiction over a dominant acquisition in a member state with no national competition law. It ran the econometric analysis anyway, and built the evidence base as a substitute for the missing instrument.⁷³

The legal difficulties are not unknown. Scholarship has flagged them from the start. The first is the admissibility of evidence whose creation process cannot be verified.⁷⁴ What no one has yet done is to build a coherent legal framework for the whole, one that addresses these vulnerabilities together rather than one at a time. Until that framework exists, decisions resting on computational analysis remain more exposed to challenge than the agencies producing them may appreciate, and the exposure grows with every new tool deployed.

This is what ATLANTIS is for. Under a European Research Council Consolidator Grant at the Vrije Universiteit Amsterdam, the project will build that framework, a legal regime solid enough that agencies can rely on computational analysis and the parties before them can trust the result. The grant provides two million euros over five years, enough for a team of three doctoral researchers and one postdoctoral researcher.

72 Stanford Computational Antitrust project's Fifth Cross-Agency Report. For the agencies' answers, see in particular the contributions of the Competition Commission of Pakistan, the Competition and Consumer Commission of Singapore, the Conselho Administrativo de Defesa Econômica (Brazil), the Agency for Protection and Development of Competition (Kazakhstan), the Competition Council (Latvia), and the Taiwan Fair Trade Commission.

73 CARICOM Competition Commission, contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 190. It describes an interrupted time series and ARDL analysis of a dominant firm's acquisition in a Member State that has no national competition legislation. The Commission frames computational methods as a substitute for formal regulatory instruments where those instruments are absent. The analysis found the observed price dynamics consistent with regional macroeconomic trends rather than with the exercise of market power.

74 Schrepel, "Computational Antitrust: An Introduction and Research Agenda", I Stanford Computational Antitrust 1, at 1-9. It raises the question of the evidentiary status of computationally generated findings, and the literature that has since developed around it.

They focus on that single objective rather than on many questions at once. The work will appear at teamatlantis.eu as it is released.⁷⁵ It will be one effort among the several the field needs.

B. Earning Legitimacy in Its Instruments

A science is defined by its shared instruments. Physics agrees on units of measure. Clinical medicine agrees on trial protocols. Computational antitrust has yet to settle on either. A merger simulation that no one else can reproduce, run on data no one else can inspect, tells the reader less than it should, however careful the work behind it. Common corpora and shared benchmarks are the precondition for the field being able to disagree productively. Until two researchers can run competing methods against the same benchmark and compare the results, the field accumulates demonstrations rather than knowledge. Data access sits underneath this, and it is one of the pillars of computational antitrust, since the methods can only compute on what agencies and researchers can lawfully obtain.

The same demand for verifiability returns, with higher stakes, in the courtroom. There it stops being a methodological question and becomes a question of legitimacy. An agency's authority to decide rests on the decision being contestable. Companies and individuals hold procedural and fundamental rights, and those rights must not weaken because an agency has adopted new tools. A defendant is entitled to understand and contest the basis of a decision, and a court of appeal that will ultimately have to agree is entitled to distrust a process it cannot inspect. The first adversarial tests are arriving. In Brazil, search warrants that rested on network indicators and econometric screens were granted and then upheld against defense challenges. The computational evidence survived scrutiny at the warrant stage.⁷⁶

Handled well, computational methods strengthen the rights of the defense rather than eroding them, because the same techniques that help build a case can test whether it was built fairly. The hard question is not whether analysis can reach an answer. It is whether it can reach one that a respondent and a court will accept, which is to say whether the field can earn the authority it is already exercising.

C. Agreeing on What Antitrust Will Be

The deeper cluster of questions goes to the essence of the discipline, to what antitrust measures and how it is practiced. Computational methods force both into the open. The field will have to agree on answers.

Start with what we measure. Antitrust practice still leans heavily on industrial organization and on neoclassical theory, the framework in which markets settle into a stable resting point that analysis can then describe. That framework mismeasures an economy defined by feedback loops and increasing returns, the condition in which early advantages compound rather than erode.⁷⁷ Brian Arthur and the researchers associated with the Santa Fe Institute spent decades showing that many real economies behave as adaptive systems, path-dependent and often far from any equilibrium.⁷⁸

Computational tools are unusually good at putting these assumptions to the test. An agent-based model can show whether a market actually converges to the equilibrium the standard analysis presumes, and simulation can reveal dynamics no closed-form model predicts. The tools can supply the challenge. What they cannot supply is the willingness to take it seriously. Enforcement institutions sit on decades of case law built on the standard framework, and precedent creates its own path dependency. A theory embedded in a thousand decisions is not dislodged by a better model, because every departure from it must be argued against the accumulated weight of what courts have already accepted. The obstacle is institutional rather than analytical, and it will take a desire for the challenge, inside agencies and inside courts, for the better analysis to matter. The economists who have built this alternative have been saying so for years, and their case is worth hearing in their own words.⁷⁹

⁷⁵ ATLANTIS, ERC Consolidator Grant no. 101228709, Vrije Universiteit Amsterdam (2026-2031). Project outputs are collected at <https://teamatlantis.eu>.

⁷⁶ Conselho Administrativo de Defesa Econômica (Brazil), contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 179. The Commission reports that judicial authorization for the Novo Rumo dawn raids required presenting network-based indicators, econometric screens, and digital infrastructure signals in a form both legally sufficient and intelligible to a nonspecialist court, and that subsequent challenges by defense counsel were rejected.

⁷⁷ Nicolas Petit & Thibault Schrepel, "Complexity-Minded Antitrust", 33 *Journal of Evolutionary Economics* 541 (2023), <https://doi.org/10.1007/s00191-023-00808-8>.

⁷⁸ W. Brian Arthur, "Competing Technologies, Increasing Returns, and Lock-In by Historical Events", 99 *Economic Journal* 116 (1989); W. Brian Arthur, "Foundations of Complexity Economics", 3 *Nature Reviews Physics* 136 (2021); Thibault Schrepel, "The Evolution of Economies, Technologies, and Other Institutions: Exploring W. Brian Arthur's Insights", 20 *Journal of Institutional Economics* e21 (2024).

⁷⁹ See the *Scaling Theory* podcast, episodes with J. Doyné Farmer (Institute for New Economic Thinking, University of Oxford; Santa Fe Institute), <https://www.youtube.com/watch?v=xmcntt8w1b8>; Eric Beinhocker (Blavatnik School of Government, University of Oxford), <https://www.youtube.com/watch?v=Agia1LUSYpc>; Scott E. Page (University of Michigan; Santa Fe Institute), <https://www.youtube.com/watch?v=dZOiD-2wlo4>; W. Brian Arthur (Santa Fe Institute), <https://www.youtube.com/watch?v=F9K6vm67nY>; David Krakauer (Santa Fe Institute), <https://www.youtube.com/watch?v=-dTclfTSD0w>; and Melanie Mitchell (Santa Fe Institute), https://www.youtube.com/watch?v=nX_i9fyC5oM.

Then there is how we practice. The parties agencies investigate will adopt these methods too. Compliance will be automated, and the analysis a defendant brings to a hearing will increasingly be produced by the same kinds of systems the agency used to build its case. At some point, enforcement becomes machine against machine. Nothing in current procedure was designed for that encounter. What does cross-examination mean when the analysis on both sides was produced by systems neither advocate fully reconstructs? What does equality of arms require when it is measured in computation?⁸⁰ The answer will also depend on the expertise parties bring into litigation. Law firms play a central role in selecting the experts and arguments that reach the court. In adversarial systems, and especially under the European principle of party disposition, courts can address only the arguments and evidence the parties place before them. Computational antitrust therefore depends not only on agencies and courts, but also on the work of the lawyers and experts who shape the cases before them.

D. Charting the Cognitive Frontier

The three preceding challenges are ones I take to be common ground, in the sense that the field recognizes them as its own whether or not it accepts my account of them. What follows is different. It is a personal reading of where the work is heading, offered with correspondingly less confidence and correspondingly more conviction. I set it apart for that reason.

One thing markets will not do is hold still, and this is precisely where computational antitrust earns its keep. Foundation-model markets change structure faster than an investigation can document, and the conduct under examination often adapts while the examination is under way. A static snapshot cannot capture a system that never sits still, which is why such markets are not a problem for computational antitrust but its clearest justification.

The frontier beyond that is cognitive. The problem is not that firms are getting better at observing consumers. It is that they are changing what consumers want. Recommendation systems and generative interfaces do not merely satisfy preferences. They form them, and a preference formed by the practice under investigation is no longer independent of it.

That breaks something antitrust relies on more than it admits. Almost every assessment of harm rests on a counterfactual, on what the market would have looked like absent the practice. The counterfactual presupposes that preferences exist and sit still while conduct varies around them. Once the conduct is itself shaping what people want, there is no fixed point left to compare against. The question shifts from whether consumers got what they wanted at a competitive price to what the practice did to the wanting itself, and antitrust has no vocabulary for that.

Answering it means measuring effects on cognition. This is less speculative than it sounds. Poland's UOKiK has already run neuromarketing experiments to establish how dark patterns act on users, and it found the effects held regardless of education or income. Singapore's commission is following it into the laboratory. It runs a controlled experiment that varies how hard a subscription trap is to escape, in order to measure what the design does to choice.⁸¹ Canada's Bureau, for its part, is building a toolkit that maps harmful online choice architecture onto the provisions of its Competition Act.⁸² The tools exist in principle, since cognitive science can observe how attention and choice form in the brain, and that work is itself increasingly computational.⁸³ The difficulty is generalization. A finding about how a small sample of people respond does not translate into a claim about a market of millions, and moving from an individual effect to a collective one has no analytical shortcut. It will require simulation and computational inference. Cognitive science and computational antitrust will have to meet, and I expect this to become the field's defining terrain.

There is a second sense of cognition here, and it points inward. The systems agencies adopt have an interior of their own, and until recently that interior was closed. This may be beginning to change. Anthropic's July 2026 interpretability work identifies what its researchers call a J-space inside a language model, which is a small internal region holding the concepts the model is working with before it produces any

80 The principle of equality of arms is a component of the right to a fair trial under Article 6 of the European Convention on Human Rights and Article 47 of the Charter of Fundamental Rights of the European Union.

81 Competition and Consumer Commission of Singapore, contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 267. The Commission describes a research collaboration with the Civil Service College of Singapore that deploys a bait-and-switch experiment on a representative sample of Singaporean adults. The experiment systematically varies the difficulty of opting out of a subscription, in order to measure the behavioral effect of dark patterns.

82 Competition Bureau Canada, contribution to the Stanford Computational Antitrust project's Fifth Cross-Agency Report, at 186. The Bureau describes a toolkit under development on harmful online choice architecture (HOCA). The toolkit maps manipulative digital design onto provisions of Canada's Competition Act.

83 On the computational turn in the study of choice, see Antonio Rangel, Colin Camerer & P. Read Montague, "A Framework for Studying the Neurobiology of Value-Based Decision Making", 9 *Nature Reviews Neuroscience* 545 (2008), and Ian Krajbich, Carrie Armel & Antonio Rangel, "Visual Fixations and the Computation and Comparison of Value in Simple Choice", 13 *Nature Neuroscience* 1292 (2010), both of which model the formation of preference and attention as a computational process.

output.⁸⁴ If such a region can be read reliably, the black box may not be permanently sealed. The implication for antitrust would be considerable. An agency able to show what a system was attending to when it flagged a firm could offer a degree of explainability that antitrust has never had, exceeding what any human decision-maker can provide about their own reasoning. That is a large claim and it rests on very new work. It is far too early to say how solid the finding is, or whether it generalizes beyond the models in which it was found. It is precisely the kind of frontier research computational antitrust is positioned to carry into law, and precisely the kind that should be treated with caution until it is confirmed.

E. Deciding How Much to Automate

One question ahead is not on the list above, because it is not a problem to be solved. It is a choice to be made. How much of antitrust are we willing to automate?

The comfortable answer is that a human stays in the loop and the machine only assists. Said this way, it settles very little, because everything depends on what that human is doing. The formula is now universal. Agency after agency in the fifth report volunteers that computational outputs inform but never determine its decisions.⁸⁵ Some outputs can genuinely be verified. When an analysis points to one case, a person can read the case and check the reasoning end to end. But when an analysis detects a pattern across tens of thousands of documents, no person can redo that work. The reviewer is then trusting the system and validating its finding, which is a different act, however carefully performed. The phrase covers both situations and tells us nothing about which one we are in. Aviation faced this question decades ago and answered it honestly. A modern pilot does not fly the aircraft for most of a flight, but monitors a system that flies it, and the profession rebuilt its training and its liability rules around that fact rather than pretending otherwise. Antitrust has not yet had that conversation.

The real question is whether we should insist that everything be verifiable by a person. Insisting on it has an obvious appeal and a real cost, because some of what computational tools do best is precisely what we cannot check by hand, the pattern across a corpus too large for any reader. Demanding human verification of everything means giving up much of that. Not demanding it means accepting findings on trust. Neither position is comfortable, both have serious arguments behind them, and the field has not resolved the question. My own expectation is that we will automate a great deal more than we do today, and considerably more than we currently think we should. The pressure will come from swelling caseloads and from markets that move faster than any investigator, and it will be difficult to resist.

Whether there is a limit, and where it should fall, is not a question the tools can answer. It is ours to decide, and I do not think the answer is obvious. That, more than any single method or model, is the work of the next five years.

III. CONCLUDING THOUGHT

Toulmin's account is the one to return to here. A discipline, on his reading, is a population of concepts and methods that survives or does not according to whether practice finds a use for it. Nothing in that account has an author. The vocabulary of computational antitrust was not fixed by whoever proposed a term but by the agencies that kept some words and abandoned others, and the methods that count as standard are the ones that held up in live cases. A name can be offered. Whether it holds is settled by everyone who does the work afterward, and over five years they settled it.

Whatever computational antitrust becomes, it will be the work of the people now doing it. Among them are the agencies that had the foresight to invest in new methods before anyone else validated them, the economists and computer scientists who crossed into a field that was not theirs, the case handlers who tested tools against real files and said plainly when they failed, and the editors and reviewers who did the unpaid work of holding the standard. The field exists because they decided it was worth their time.

Computational antitrust has become a field that has changed what counts as knowledge in competition law. For decades, antitrust advanced through legal doctrine and economic theory. It now advances through computation as well. That shift is unlikely to be reversed. If that is right, the question facing the field is larger than whether to adopt new tools. It is whether competition law is prepared to rethink how it produces knowledge about markets.

⁸⁴ Anthropic, "Verbalizable Representations Form a Global Workspace in Language Models", Transformer Circuits Thread (6 July 2026), <https://transformer-circuits.pub/2026/workspace/>. The finding is recent and awaits independent replication.

⁸⁵ See, among others, the contributions to the Stanford Computational Antitrust project's Fifth Cross-Agency Report of the Taiwan Fair Trade Commission, the Turkish Competition Authority, the Comisión Nacional Antimonopolio (Mexico), the Urząd Ochrony Konkurencji i Konsumentów (Poland), and the Competition Council (Latvia). Each reports that computational outputs inform but do not determine enforcement decisions.

CPI Subscriptions

CPI reaches more than 35,000 readers in over 150 countries every day. Our online library houses over 23,000 papers, articles and interviews.

Visit competitionpolicyinternational.com today to see our available plans and join CPI's global community of antitrust experts.

