The Four-Path Evidence-Gap Diagnostic

For any firm whose prices are set or recommended by software — and for its advisers — a structured way to ask what the evidence can actually show.

Saurabh Bhalla · independent researcher, competition law & technology · v0.3 · August 2026 · private draft for discussion — not legal advice

Four paths can influence a software-set or software-recommended price: *runtime, training, control-plane, and human/vendor*. For one defined price decision: which paths existed? Which can the firm's present evidence reconstruct? And where must the answer remain *indeterminate*?

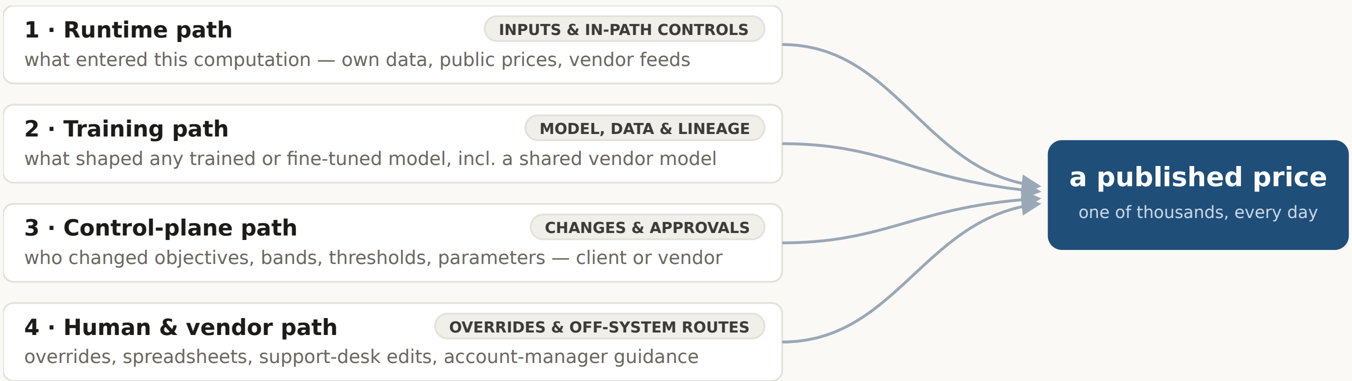

THE DIAGNOSTIC — QUESTIONS AN ADVISER CAN PUT TO A CLIENT VERBATIM

1 Runtime

- Choose one specific published price from a defined period. Can you identify every recorded input to that decision — its origin, age, transformations, and the issuer of any public/non-public classification?
- Which controls actually gated the computation, and which existed only as policy or documentation?

2 Training

- Did the deployed system use a trained or fine-tuned model? If so, what data, transformations and model version shaped that price?
- If it used a shared vendor model, did data from other clients or competitors enter training — and what evidence can the firm inspect, obtain under audit rights, or independently corroborate? If no trained model existed, what record supports that conclusion?

3 Control-plane

- Which objectives, price bands, thresholds, parameters, models or policies changed during the defined period?
- Who or what made each change, who approved it, and what record links the change to the price decisions it affected?

4 Human & vendor

- When a price was overridden or set outside the normal computation — during an outage, through a spreadsheet, or by manual repricing — what evidence survived, and how was it reconciled to the published-price population?
- Could support tickets, vendor configuration, account-manager guidance, notes or calls influence pricing? If so, where was that influence recorded?

If present evidence cannot resolve a path within the stated boundary and period, the diagnostic reports that path as **indeterminate**. If records support that a path did not exist, state only that bounded fact — absence is never inferred from silence. The output is a path-by-path account of **what present records support, which claims depend on named issuers or assurances, and what cannot be shown**: input to counsel and economists, never a verdict. A record that reports its own gaps is worth more than an assurance that has none.

WHY THIS QUESTION IS LIVE

Entered and proposed remedies, statutes and guidance now interrogate these paths directly.

Selected official sources make parts of this question live. They do not adopt the four-path taxonomy, and they do not establish any firm's evidentiary position. Status checked 1 August 2026.

United States v. RealPage — proposed remedy · RUNTIME + TRAINING + OVERSIGHT

Proposed final judgment (filed 24 November 2025, No. 1:24-cv-00710, M.D.N.C.): subject to specified exceptions, §IV.A.1 requires RealPage to cease using current or historical unaffiliated-property data at runtime. §IV.A.3 permits unaffiliated-property data in model training only when it is historical or backward-looking, at least twelve months old, and not from active leases. §IV.A.5 limits the geographic granularity of models trained on such data; §VII provides for a monitor with access to code, pseudocode, model-training documentation and runtime pricing logic. The judgment remains proposed; a stipulation and order entered 26 March 2026 requires compliance with its provisions pending entry.

Proposed Final Judgment (24 Nov 2025) · Response to Public Comments, 91 FR 25373 (8 May 2026) · DOJ case page

Same case, as to Greystar — entered judgment · CUSTOMER SIDE, INCL. TRAINING

Final judgment entered 2 March 2026: no licensing or use of revenue-management products that use others' non-public data at runtime, pool non-public data across owners, *or whose models were trained on others*' *non-public data* (§IV.A.1). The judgment operationalises these requirements through vendor certification (§IV.E), general-counsel certification, employee attestations (§VI.C) and specified monitorship triggers (§VIII.B).

Final Judgment as to Greystar (entered 2 Mar 2026)

California AB 325 — statute · COMMON ALGORITHMS

Signed 6 October 2025. Unlawful to use or distribute a "common pricing algorithm" — a methodology using competitor data to recommend, align, stabilise, set or otherwise influence a price or commercial term — as part of a contract, combination or conspiracy; a separate coercion prohibition; and a complaint need not allege facts tending to exclude independent action (§§16729, 16756.1).

AB 325, Cartwright Act: violations — chaptered text (2025–2026 Reg. Sess.)

DOJ guidance and enforcement remarks · EVIDENCE EXPECTATIONS

The Antitrust Division's compliance-program evaluation (November 2024) asks how a company assesses the antitrust risk of AI and algorithmic revenue-management software, whether compliance personnel join deployment, and how quickly it can detect and correct decisions by AI or other new technologies that are not consistent with the company's values. A senior Division enforcement official, May 2026: "Every automated bid is a structured record. Every tool that ingests competitor-data inputs is an artifact."

Evaluation of Corporate Compliance Programs in Criminal Antitrust Investigations (Nov 2024) · D. W. Glad, remarks, Antitrust West Coast Conference (14 May 2026)

United Kingdom — CMA · USERS AND SUPPLIERS

Business guidance on pricing algorithms (November 2024): a common algorithmic system can exchange confidential information and help coordinate behaviour; "Not knowing what an algorithm is doing is no excuse: both end users and suppliers must understand how the tools and technologies they are using or supplying work." On AI and collusion (March 2026): algorithmic screening is "part of a wider evidential toolkit rather than as a substitute for legal and economic analysis."

Pricing algorithms and competition law: what you need to know (8 Nov 2024) · AI and collusion: frontiers, opportunities and challenges (4 Mar 2026)

The wider field · MOVING

The FTC and DOJ opened a joint request for public comment toward new business-collaboration guidance (February 2026, extended to 21 May 2026), expressly naming algorithmic pricing and information sharing. The OECD's review of algorithmic pricing across G7 jurisdictions (October 2025) finds shared concerns, divergent responses and limited enforcement to date.

FTC-DOJ request for public comment (Feb 2026) · Deadline extension (Apr 2026) · OECD, Algorithmic pricing and competition in G7 jurisdictions (Oct 2025)

WHAT THE DIAGNOSTIC YIELDS

A bounded, provisional map of the firm's evidentiary position — path by path: what present records support, which claims depend on named issuers or assurances, and what cannot be shown. That map is an input to counsel and economists, never a verdict. This diagnostic cannot by itself establish that a price was lawful, independently determined or non-collusive. I am researching evidence architectures for exactly this boundary, and circulating this note to be corrected — if parts of it are wrong, or useless, I would like to hear that first.