



Deterministic Auto-Segmentation

A Framework for Explainable and Governable Personalized Pricing

REGULATORY CONTEXT, OPERATING PRINCIPLES AND IMPLEMENTATION FRAMEWORK

October 2026

Important notice: This whitepaper is provided for informational purposes only and does not constitute legal advice. Laws and regulatory requirements concerning personalized pricing vary by jurisdiction and continue to develop. Organizations should obtain advice from qualified legal counsel regarding their specific practices.

Contents

Executive Summary.....	3
The Regulatory Context for Personalized Pricing	4
A Fragmented Regulatory Landscape	5
What Is Deterministic Auto-Segmentation?	7
What Makes Deterministic Segmentation Different	8
Three Operating Principles.....	9
The Five-Stage Framework.....	11
Stage 1: Map and Classify	13
Stage 2: Audit.....	19
Stage 3: Build Deterministic Segments	21
Stage 4: Disclosure	26
Stage 5: Governance	28
Transitioning from Legacy Black-Box Models	30
Technology Requirements for Deterministic Segmentation	31
Scope and Limitations	33
Frequently Asked Questions	34
Appendix A: Alignment with the FTC Proposed Enforcement Policy Statement	37
Appendix B: Implementation Blueprint	39
Sources and References.....	40

Executive Summary

Regulators are increasingly asking businesses to explain how a personalized price was reached: what data was used, why the customer received that treatment and whether the decision can be reproduced later.

Deterministic Auto-Segmentation, or DAS, is one way to structure pricing so those answers are available from the decision record itself.

The FTC's August 2026 proposal and the state laws now coming into force point in the same general direction:

- **Disclose the basis.** Where the rules apply, tell customers that a price is personalized, why, and which types of data were used.
- **Know the inputs.** Identify every data point behind a price, including inferences, and review any that reveal sensitive circumstances such as financial distress, urgency or vulnerability.
- **Stand behind the decision.** Be ready to show a regulator, a court or a customer how a past price was set, which in practice means reconstructing the rule, data and explanation in force at the time.

Deterministic segmentation is intended to make that easier. Customers are assigned to defined segments using explicit rules, and each segment is linked to a price or offer. Because the rule is part of the production decision, the reason for the treatment is available when the price is set rather than generated later as a separate explanation. Versioning preserves the rule, treatment and supporting data over time.

The commercial value of personalization remains important. In a field experiment, Dubé and Misra found that personalized prices increased expected profits by 19% compared with the best uniform price, while more than 60% of consumers benefited. Deterministic segmentation provides a clear, auditable, easily governed record for legal, compliance and business review.

A simple example shows how that works. A customer might see: "Repeat-buyer price, 8% off, based on your purchase history with us." Behind that offer is a record containing the segment rule, the data used, the treatment, the basis for the treatment and the version in force.

The framework reflects practical experience building and operating segment-based pricing systems at Klary, including the challenge of turning analytical customer groups into pricing rules that can be deployed, monitored and explained. The principles are presented here independently of any particular product so organizations can evaluate how deterministic segmentation could apply within their own pricing, legal and governance environment.

The main body focuses on the regulatory context, the deterministic segmentation architecture and the five implementation stages. More technical material on segment discovery, FTC alignment and implementation follows for readers who need the detail.

The Regulatory Context for Personalized Pricing

Personalized pricing and offers increasingly sit at the intersection of pricing strategy, consumer data and regulatory scrutiny.

In August 2026, the Federal Trade Commission proposed an enforcement policy statement addressing circumstances in which personal data is used to determine an individual consumer's price. Where consumers reasonably expect a common price, the proposed policy states that a business should clearly disclose the fact that a price is personalized, the basis for the personalization and the types of data on which the personalization is based.¹

At the same time, states are developing their own approaches. New York already requires a prescribed disclosure for certain algorithmically determined prices. Maryland has enacted restrictions affecting personalized food pricing. Connecticut and New Jersey have adopted measures that take effect in 2027. Colorado already gives consumers the right to opt out of certain profiling, and its replacement AI law takes effect in January 2027. Additional proposals remain under consideration in other states, while California's AB 2564 passed both chambers but did not receive a final Assembly vote before the 2026 session ended. These laws differ materially in scope, definitions, exclusions and treatment of discounts.²

The resulting challenge is not simply whether personalization is used. It is whether an organization can identify how a particular price or offer was determined, what data was used, how consumers were grouped or differentiated, and whether the same decision can be reconstructed later.

Deterministic segmentation is one architectural approach to addressing that problem. It structures personalization around explicit customer segments and readable production rules rather than relying exclusively on individual-level or opaque model outputs.

Regulatory expectation

Identify why a price or offer differs and what data was used to determine it.



Traceable decision logic

Connect the customer data, applicable rule and resulting treatment in a single decision record.



Reproducible decision

Reconstruct the price or offer, the rule that produced it and the explanation in force at the time.

The regulatory issue is not simply whether personalization occurs, but whether the organization can trace and reconstruct how a particular decision was made.

¹ Federal Trade Commission, Proposed Enforcement Policy Statement Regarding Personalized Pricing (Aug. 19, 2026), Docket No. FTC-2026-1057.

² Greenberg Traurig, "Algorithmic Pricing Under Fire: State Restrictions on Personalized and Surveillance Pricing" (Sept. 3, 2026).

A Fragmented Regulatory Landscape

There is no single U.S. rule governing personalized pricing. Federal policy remains proposed, while state laws take different approaches to disclosure, prohibited practices, discounts and group-based pricing. As a result, the legal treatment of a pricing practice may depend on the jurisdiction, the data used, the nature of the price difference and the structure of the offer.

JURISDICTION	STATUS	CORE APPROACH
FTC proposed policy	Proposed, Aug. 2026	Disclosure where personal data determines an individual price in circumstances where consumers reasonably expect a common price
New York, GBL § 349-a	Enforced since Nov. 10, 2025	Prescribed disclosure for certain personalized algorithmic pricing
Maryland, HB 895	Effective Oct. 1, 2026	Restrictions on higher personalized prices in covered food retail/delivery contexts, with specified exclusions
Connecticut, P.A. 26-130	Effective July 1, 2027	Restrictions on surveillance pricing, including certain group-based pricing
New Jersey, Fair Price Protection Act	Pricing provisions effective Aug. 1, 2027	Restrictions relating to grocery pricing, with provisions for qualifying discounts based on disclosed, uniformly applied criteria
New York, One Fair Price Act	Passed legislature June 4, 2026; not signed as of early October 2026 (Governor's deadline Dec. 31, 2026)	Would impose broader surveillance-pricing restrictions while preserving qualifying discounts
Colorado, SB 26-189 (replaces the Colorado AI Act)	Effective Jan. 1, 2027 (signed May 14, 2026)	Disclosure, access, correction and human-review rights for automated decision tools that materially influence consequential decisions; the listed decision categories may not reach retail pricing
Colorado Privacy Act, profiling opt-out	In force	Right to opt out of profiling behind decisions with legal or similarly significant effects, including access to essential goods or services, with a limited exception for human-involved automated processing
California, AB 2564	Introduced Feb. 20, 2026; passed both chambers but died when the Assembly did not vote on Senate amendments before the session ended Aug. 31, 2026	Would have banned surveillance pricing by retailers

Colorado deserves particular attention for segmentation. SB 26-189 lists education, employment, housing, financial services, insurance, healthcare and essential government services as consequential decisions, so it may not reach retail pricing. The Colorado Privacy Act is the more direct hook: consumers can opt out of profiling in furtherance of decisions that produce legal or similarly significant effects, which include access to

essential goods or services. The implementing rules let a controller decline an opt-out request, subject to conditions, where the profiling is based on human-involved automated processing. Whether a segment rule that people have reviewed and approved qualifies is a question for counsel, but a deterministic architecture makes that human review visible and documented.³

These regimes are not equivalent, and segmentation by itself does not create a universal safe harbor.

IMPORTANT DISTINCTION: Deterministic segmentation is a pricing architecture, not a legal exemption. Whether a particular practice is permitted depends on the applicable law, the data used, the treatment produced and the surrounding circumstances.

³ Colo. S.B. 26-189 (signed May 14, 2026), as summarized by Holland & Knight (May 18, 2026); Colo. Rev. Stat. § 6-1-1306(1)(a)(I)(C); 4 Colo. Code Regs. 904-3, Rule 9.04(C).

What Is Deterministic Auto-Segmentation?

Deterministic Auto-Segmentation is a pricing architecture in which customers are assigned to defined, mutually exclusive segments using explicit rules, with a predetermined price or offer treatment associated with each segment.

The defining feature is not segmentation alone. The production decision itself is rule-based and reproducible.

For each pricing decision, the organization can maintain a record of:

- the segment rule
- the data used
- the source of that data
- the price or offer treatment
- the explanation associated with the decision
- the effective version and date

A customer who meets the same rule under the same version receives the same treatment.

FIGURE 1: ONE TRAVEL OFFER, ONE REPRODUCIBLE RECORD

1 Customer and pricing data



Customer attributes

Loyalty tier, market, account status



Booking behavior

Prior stays, booking frequency, average booking value



Prices and offers

Public rate, member rate, targeted offers



Outcomes

Bookings, repeat stays

2 determines the applicable segment

Illustrative example: Frequent direct bookers

Rule 4+ direct bookings in 12 months
AND loyalty member

Data Booking history and membership status

Source First-party

Treatment Member offer below public rate

Version v3 · Sept. 2026

3 creates a reproducible decision record



Why this offer?



What data was used?



How did the treatment differ?



Could sensitive data be involved?



What can we explain to the traveler?



Can we reproduce it later?

Illustrative example. Whether a particular pricing practice is legally permissible depends on the applicable jurisdiction and circumstances.

A segment record ties data, rule, treatment, explanation and version together.

What Makes Deterministic Segmentation Different

Traditional customer segmentation is commonly used for analytics, campaign planning and targeting. Those segments may influence a business decision without directly determining the price or offer shown to an individual customer.

Deterministic segmentation is different because the segment rule becomes part of the production decision itself. The customer's treatment is determined by an explicit rule that can be applied consistently, versioned and reconstructed later.

Four characteristics distinguish deterministic segmentation from conventional segmentation.

1

Explicit Rules

Segment membership is determined by defined conditions and thresholds. The rule itself is readable and can be inspected directly rather than inferred from an opaque score.

2

Mutually Exclusive Execution

For a given pricing decision, one operative segment determines the applicable treatment. This creates a single decision path rather than several overlapping explanations.

3

Versioned Treatment

Each segment rule is linked to a defined price or offer treatment and an effective version. When the rule changes, the prior version can be retained for historical review.

4

Reproducibility

The organization can reconstruct how a historical price or offer was determined using the rule, data and version that were active at the time.

Together, these characteristics make the segment part of the production pricing architecture, rather than simply an analytical grouping used upstream.

Three Operating Principles

Deterministic segmentation rests on three principles that apply to every segment and every pricing decision.

1

Deterministic Execution

The same inputs evaluated against the same active rule produce the same segment assignment and treatment.

This creates a direct relationship between the customer's data, the applicable rule and the resulting price or offer.

2

Full Explainability

The reason for the pricing decision is embedded in the production logic.

The explanation does not need to be inferred after the decision from hundreds of weighted variables or reconstructed from a separate explanatory model.

3

Modular Governance

Rules can be reviewed, approved, versioned, tested and replaced individually.

Changes to one segment or one pricing treatment do not require the entire personalization architecture to become opaque.

Deterministic Execution establishes what happened. Full Explainability establishes why it happened. Modular Governance preserves how that decision can be reviewed over time.

Why These Principles Matter

Each principle addresses a different challenge in personalized pricing, and together they create a decision structure that can be inspected, explained and governed over time.

1 Deterministic Execution creates reproducibility.

Because identical inputs always produce the same result, the organization has a defined basis for each price or offer and can reconstruct a historical decision later. It also makes clear which change in the underlying data or rule would have produced a different outcome.

2 Full Explainability connects the decision to its reason.

The explanation is derived from the same logic that determines the treatment. This is important in the context of the FTC's proposed approach, which focuses on explaining the basis for personalization and the types of data

used. Instead of generating a plausible explanation after the decision, the organization can point to the rule that actually produced the result.

3 Modular Governance preserves control as the system changes.


Pricing rules, customer behavior and regulatory requirements do not remain static. Treating each segment as a versioned governance object allows individual rules to be reviewed, tested, approved, updated or retired while preserving the record of what was previously in production. This supports both ongoing monitoring and historical auditability.

Together, the three principles create a traceable chain from data → rule → treatment → explanation → record.

The Five-Stage Framework

Deterministic segmentation connects the data used for personalization, the rule that determines a customer's treatment and the explanation associated with that decision. The implementation process can be organized into five stages.

FIGURE 2: FIVE STAGES FOR IMPLEMENTING PERSONALIZED PRICING



Classify: Identify every data point and feature used in the pricing or offer decision, including its source, whether it is raw or inferred, and whether it may require heightened review. This supports the FTC's proposed expectation that businesses disclose the types of data on which personalization is based. Recording each source also supports the separate task of verifying data provenance and consent.

Output: a documented feature inventory showing provenance, derivation and potential treatment effect.

Audit: Examine how the existing pricing logic behaves in practice, including who receives different treatment, who pays above or below the relevant baseline, and whether customer-facing descriptions match the actual logic. This helps identify practices that may require further legal or fairness review before they are translated into deterministic rules.

Output: an audit record of current decision logic, treatment differences and identified areas for review.

Segment: Translate personalization into explicit, mutually exclusive production rules so that each customer-facing decision can be traced to a defined set of criteria. This is important because group-based pricing itself is not automatically outside regulatory scope; some state approaches can apply to differentiated treatment at

both the individual and group level. Each segment's price or offer is then set from its measured response to price.

Output: an approved segment rule linked to a defined price or offer treatment.

Disclose: Derive the customer-facing explanation from the same production rule that determined the treatment. This is the stage that addresses the FTC and New York disclosure expectations described in The Regulatory Context for Personalized Pricing.

Output: disclosure language that can be traced directly to the active rule and underlying data record.

Govern: Preserve each segment rule, data source, treatment, disclosure text, effective date and version, and monitor how those elements change over time. This creates the record needed to test whether a historical pricing decision can be reconstructed and whether the explanation shown to the customer matched the rule actually in force.

Output: a versioned segment registry supporting monitoring, review, change control and historical reproducibility.

The versioned segment record links all five stages. The objective is not simply to make pricing easier to explain, but to create a common operational record that can support execution, disclosure, review and historical reconstruction.

Stage 1: Map and Classify

01 MAP / CLASSIFY

02 AUDIT

03 SEGMENT

04 DISCLOSE

05 GOVERN

Before Classification, Map Where Personalization Occurs

Personalization is rarely confined to the displayed list price.

Personal data can influence several forms of customer economic treatment:

- base or list price;
- targeted discounts and coupons;
- promotions and bundles;
- loyalty and membership pricing;
- retention and win-back offers;
- fees and surcharges;
- offer ranking and eligibility;
- introductory or negotiated prices.

The first step is therefore to map every workflow in which personal data can alter what a customer pays or which economic offer the customer receives.

Mapping works backward from what the customer sees, in five steps:

- **Start at the price surface.** List every screen, email and offer where a consumer sees a price or discount, and trace each back to the system that produced it.
- **Interview the owners.** Pricing, promotions, CRM, loyalty and e-commerce teams each hold a piece. Personalization often sits in a team that does not call itself “pricing.”
- **List every feature.** A feature is any piece of information the system uses, including raw facts such as past orders and computed scores such as “likely to cancel.” Inferred scores count: FTC staff describe firms inferring financial sensitivity from what shoppers leave in their carts.⁴
- **Include vendors.** Ask each intermediary which data it uses. Under the FTC’s proposed approach, the business itself is expected to verify consent for vendor data rather than rely on contract assurances.⁵
- **Record one line per feature.** Capture the workflow, source, derivation, owner and whether the feature can raise a price or only lower it.

⁴ FTC staff, Surveillance Pricing 6(b) Study: Research Summaries – A Staff Perspective (Jan. 2025).

⁵ Federal Trade Commission, Proposed Enforcement Policy Statement (Aug. 19, 2026); Baker McKenzie, “United States: FTC Proposes Enforcement Framework for Personalized Pricing” (Aug. 26, 2026).

The map determines the disclosure. The FTC expects a disclosure to name the types of data behind a price, so a disclosure can be no more accurate than this inventory. The NIST AI Risk Management Framework makes “Map” a formal function for the same reason: inputs and context must be known before risk can be measured.⁶

FIGURE 3: PERSONALIZATION EXTENDS BEYOND THE PRICE

Map every workflow that can change a customer’s price or offer



TAKEAWAY

Map personalization wherever personal data can change a customer’s economic treatment.

Workflows in which personal data can change a customer’s economic treatment.

Classify the Inputs Behind Each Workflow

Each input should then be classified across several dimensions:

SOURCE	FORM	SENSITIVITY	EFFECT
• first-party; • third-party.	• raw data; • inferred variable or score.	• ordinary commercial information; • information or inference requiring heightened review because it may reveal health, financial condition, household circumstances, urgency or vulnerability.	• can raise a price; • can lower a price; • affects eligibility or ranking only.

The purpose of classification is visibility. It does not determine by itself whether a particular input is legally permissible.

Why the source matters. First-party data comes from the consumer’s own relationship with the seller: purchases, account activity and stated preferences. Third-party data comes from another entity, such as data

⁶ NIST, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (Jan. 2023).

brokers, other websites, device signals or public records. The source *drives* both consumer trust and regulatory risk:

Consumer Trust: Consumers find information obtained outside a site less acceptable, and in experiments, revealing such data flows reduced ad effectiveness.⁷ In two experiments, consumers judged prices based on purchase history fairer than prices based on location.⁸ Secondary use of information lowers trust in a firm, and privacy expectations are strongest when data brokers are involved.⁹

Regulatory Risk: The FTC’s own deception example is a price presented as reflecting purchase history that actually reflects shopping habits with other firms.¹⁰ A practical default is to prefer first-party features and to use a third-party feature only with documented consent for pricing use and a reason it cannot be replaced.

Why sensitivity matters. Sensitive features reveal financial condition, health, household circumstances, urgency or vulnerability, directly or by inference. Even an innocent-looking signal can qualify. In 2016, Uber’s then head of research said that users with low phone batteries accept surge prices more readily. Uber says it does not use the signal, but the remark still drew congressional criticism.¹¹ Consumers object to this kind of targeting: in a January 2025 Consumer Reports survey, 76% of Americans opposed loyalty discounts based partly on demographics such as age and income.¹²

Regulators are focused here as well. The FTC’s illustrations repeatedly turn on personal data or conclusions about a consumer’s circumstances, such as a medical emergency, a funeral, children in the household or the ability to comparison shop.¹³ Commenters have urged the FTC to base unfairness on assessed vulnerability or protected class rather than on price differences alone.¹⁴ Sensitivity should therefore be judged by use as well as by data type, because the same field can be harmless in one context and sensitive in another.¹⁵

⁷ T. Kim, K. Barasz & L. K. John, “Why Am I Seeing This Ad? The Effect of Ad Transparency on Ad Effectiveness,” 45 J. Consumer Research 906 (2019).

⁸ A. Priester, T. Robbert & S. Roth, “A Special Price Just for You: Effects of Personalized Dynamic Pricing on Consumer Fairness Perceptions,” 19 J. Revenue & Pricing Mgmt. 99 (2020).

⁹ K. Martin, “The Penalty for Privacy Violations: How Privacy Violations Impact Trust Online,” 82 J. Business Research 103 (2018); K. Martin & H. Nissenbaum, “What Is It About Location?,” 35 Berkeley Tech. L.J. 251 (2020).

¹⁰ Federal Trade Commission, Proposed Enforcement Policy Statement Regarding Personalized Pricing (Aug. 19, 2026), Docket No. FTC-2026-1057.

¹¹ CBC News, “Uber users with low phone batteries more likely to accept surge pricing” (May 21, 2016); Snopes, “Uber Increases Prices for Customers with Low Phone Batteries?” (Nov. 18, 2024); Letter from Sen. Sherrod Brown, Chairman, Senate Banking Committee, to Dara Khosrowshahi, CEO, Uber (July 3, 2024).

¹² Consumer Reports nationally representative surveys (May 2024; Jan. 2025).

¹³ Federal Trade Commission, Proposed Enforcement Policy Statement Regarding Personalized Pricing (Aug. 19, 2026), Docket No. FTC-2026-1057.

¹⁴ S. Bulla, Foundation for American Innovation, “Comment on the FTC’s Proposed Personalized Pricing Policy” (Sept. 2026).

¹⁵ K. Martin & H. Nissenbaum, “Measuring Privacy: An Empirical Test Using Context to Expose Confounding Variables,” 18 Colum. Sci. & Tech. L. Rev. 176 (2016).

FIGURE 4: MAP EVERY INPUT BEHIND A PERSONALIZED PRICE OR OFFER



The Hidden Risk of Inferred Data

A model can infer what it never collected. Classification therefore has to address two different risks:

- **Sensitive data collection.** The organization holds a field that states a sensitive fact, such as a health condition or pregnancy status. The risk sits in what is requested and stored, and it can be checked by reading the data schema.
- **Sensitive inference.** The organization holds no sensitive field, but ordinary features, alone or in combination, let a model or rule estimate a sensitive fact. The risk sits in what the features reveal together and in the price or offer they then drive. It can be checked only by naming every feature and testing combinations before a segment ships.

CASE STUDY: THE TARGET INFERENCE MODEL

Target reportedly built a pregnancy-prediction score from about 25 ordinary products, such as lotion and mineral supplements, that look harmless one by one. Taken together, the purchases produced a score and an estimated due date, and Target sent coupons timed to stages of pregnancy. In one widely reported case, a father saw the coupons before he knew his daughter was pregnant.¹⁶ No field said “pregnant.” Ordinary purchases revealed a sensitive fact, and that fact drove the offer.

¹⁶ C. Duhigg, “How Companies Learn Your Secrets,” New York Times Magazine (Feb. 16, 2012); K. Hill, “How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did,” Forbes (Feb. 16, 2012).

Judge the inference, not the field. As the FTC’s illustrations above show, what matters is what the data reveals about a consumer’s circumstances, not the field itself.¹⁷

Why opaque models make this harder. In a black-box model, an inference sits inside the weights, spread across hundreds of features and thousands of paths, and no one can state a short, accurate reason for one customer’s price. Explanations produced after the fact can misstate what the model actually does.¹⁸ In a deterministic segment, the same inference would have to appear in a short rule on named features, each with a cutoff, where a lawyer can locate it and a reviewer can challenge it before the segment sets a single price. This follows directly from the way the groups are formed: membership is defined by named features rather than hidden embeddings (see Technology Requirements for Deterministic Segmentation).

Keep Each Organization’s Data Separate

Where a third party provides pricing or segmentation software, data provenance also has an antitrust dimension. Shared pricing software has been challenged as a “hub” for coordination, with plaintiffs alleging that competing companies aligned prices through a common vendor. RealPage, Cendyn, Zelis and MultiPlan have all faced such claims.¹⁹ Companies using shared pricing or segmentation software can be sued alongside the vendor, and the law remains unsettled: in July 2026 the Third Circuit reversed the dismissal of a claim involving Cendyn, while in 2025 the Ninth Circuit had affirmed the dismissal of a similar one.²⁰

A segment architecture should therefore be able to demonstrate isolation at every layer:

- **Data.** Each organization’s segments use only that organization’s records, with no pooling across other organizations or competitors and no external enrichment. In New Jersey’s RealPage litigation, a contract term of this kind led the court to dismiss the antitrust claims against one landlord in March 2026, although in September 2026 the court allowed the state to renew those claims on new allegations that the landlord shared data directly with competitors.²¹ Any feature weighting or preprocessing should be documented and available for review.
- **Code and models.** Segments for each organization are trained and stored separately, so the code shows no path from one organization’s records to another’s prices.

¹⁷ Federal Trade Commission, Proposed Enforcement Policy Statement Regarding Personalized Pricing (Aug. 19, 2026), Docket No. FTC-2026-1057.

¹⁸ C. Rudin, “Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead,” 1 Nature Machine Intelligence 206–215 (2019).

¹⁹ Cornish-Adebiyi v. Caesars Entertainment, Inc., No. 24-3006 (3d Cir. July 29, 2026), as summarized by Mayer Brown (Aug. 3, 2026); New Jersey v. RealPage, Inc., No. 25-cv-3057 (D.N.J. Mar. 31, 2026), as summarized by Milbank (June 2026); United States v. RealPage, Inc., No. 1:24-cv-00710 (M.D.N.C. May 19, 2026), as summarized by Arnold & Porter (June 2026); In re Zelis Repricing Antitrust Litigation (D. Mass. Mar. 30, 2026) (order denying motion to dismiss); In re MultiPlan Health Insurance Provider Litigation (N.D. Ill. June 2025) (order denying motions to dismiss).

²⁰ Cornish-Adebiyi v. Caesars Entertainment, Inc., No. 24-3006 (3d Cir. July 29, 2026), as summarized by Mayer Brown (Aug. 3, 2026); Gibson v. Cendyn Group, LLC, 148 F.4th 1069 (9th Cir. 2025).

²¹ New Jersey v. RealPage, Inc., No. 25-cv-3057 (D.N.J. Mar. 31, 2026), as summarized by Milbank (June 2026); order granting leave to amend (D.N.J. Sept. 29, 2026), as reported by Law360 (Sept. 30, 2026).

- **Documents.** Discovery can compare marketing materials, business plans, policies and drafts. Each should say the same thing, and none should describe using one organization's data for another, whether as inputs, benchmarks or training data.
- **Pricing authority.** The organization, not the vendor, sets each price. Courts have weighed auto-accept features, override limits and deviation scoring in the RealPage and Cendyn cases.²²

These controls reduce particular coordination risks, but they do not by themselves establish antitrust compliance. Market structure, information exchange outside the software, contract terms, how recommendations are used and actual conduct all remain relevant and are matters for antitrust counsel.

²² Cornish-Adebiyi v. Caesars Entertainment, Inc. (3d Cir. 2026), as summarized by Mayer Brown; United States v. RealPage, Inc., No. 1:24-cv-00710 (M.D.N.C. May 19, 2026), as summarized by Arnold & Porter (June 2026).

Stage 2: Audit

01 MAP / CLASSIFY

02 AUDIT

03 SEGMENT

04 DISCLOSE

05 GOVERN

Goal

Establish what the current pricing system actually does before it is rebuilt.

Output

An audit record of decision logic, treatment differences and areas requiring legal or fairness review.

For Counsel

Determine whether any treatment difference, including at group level, raises concerns under the applicable regime.

Audit Both the Logic and the Treatment

The audit has two parts: the logic behind each pricing decision and the treatment it produces.

What the audit examines. The audit traces each personalized price back to the logic that produced it and measures the results. It asks who receives different treatment, who pays above or below the relevant baseline, which inputs drive those differences, and whether any explanation given to customers matches what the system actually does.

Why group-level results matter. Some state regimes, including Connecticut's, can reach pricing differences applied to groups as well as to individuals. An audit that examines only individual prices can miss a pattern that appears only when outcomes are compared across groups.

What the audit produces. The resulting audit record sets the baseline for the rules built in Stage 3 and documents what changed when they replaced the previous system.

Why the audit still applies. The audit still matters when segments are formed without a predefined target. Unsupervised discovery describes the client's data faithfully but does not correct it: if historical data carries bias or sensitive proxies, the groups will reflect it. The audit is where that is caught, before a group is turned into a production rule.

FIGURE 5: SEVEN QUESTIONS FOR AUDITING PERSONALIZED PRICING

A defensible pricing program must explain both the decision logic and the treatment it produces.

THE LOGIC**How the decision is made**

- 1** Does the model estimate willingness to pay?
- 2** Does the customer explanation match the actual logic?
- 3** Where did the data come from, and was it authorized for pricing use?

THE TREATMENT**What the decision produces**

- 4** Could the treatment exploit urgency or vulnerability?
- 5** Do outcomes differ across protected groups or likely proxies?
- 6** Who pays more and who pays less than the baseline?

7**THE GOVERNING TEST**

Can the company reproduce and explain any past price or offer?

TAKEAWAY

Audit both the logic of personalization and the treatment it produces.

Stage 3: Build Deterministic Segments

01 MAP / CLASSIFY

02 AUDIT

03 SEGMENT

04 DISCLOSE

05 GOVERN

From Individual Prediction to Readable Production Rules

A deterministic segment is a group of customers defined by a short production rule containing a limited number of conditions and thresholds.

For example:

EXAMPLE SEGMENT RECORD

SEGMENT	Repeat buyers
RULE	Three or more purchases during the previous 12 months AND average basket value of at least \$40
DATA	Purchase history on this account (first-party)
SENSITIVE	None identified
TREATMENT	8% below list price
BASIS FOR TREATMENT	Lifted repeat purchases 20% versus a control group
MEMBERS	14,200 customers
VERSION	v3, effective Sept. 1, 2026

Each field in the record answers a question a regulator is likely to ask:

- **Why this price?** The basis has two parts: the rule explains why the customer belongs to the segment, and the basis for treatment explains why that segment receives that price or offer. Together they supply the basis the FTC asks for.
- **What data drove it?** The data and source name the data types.
- **Who pays more than the baseline?** The treatment and membership show who pays more or less.
- **Does it exploit someone's circumstances?** The sensitivity flag and the rule show what is targeted.
- **What do we tell the customer?** The rule and the basis for treatment, in plain words, form the disclosure.
- **Can a past price be reproduced?** The rule version, the pricing-model version and a reference to the input data as it stood on that date together rebuild the decision.

Each customer should belong to one operative segment for a given pricing decision.

That matters because the rule becomes the reason the customer receives the segment's treatment, while the basis for treatment records why the segment receives that treatment. Instead of attempting to explain the contribution of many variables after the fact, the organization can identify the specific rule that applied when the price was set.

The underlying framework describes this as using mutually exclusive segments so each customer has one applicable reason rather than a list of potentially competing explanations.

ROLE OF MACHINE LEARNING: Machine learning can still be used to discover potentially useful segments, thresholds or relationships. In a target-agnostic approach, discovery is unsupervised: groups are formed without a predefined target, and the pricing goal is applied only afterwards. Deterministic execution means that once deployed, the customer-facing pricing decision is governed by the approved production rule.

From Segments to Prices

A deterministic segment establishes who receives a treatment. A separate pricing step determines what that treatment should be. Segment-based pricing can be organized in three steps:

- **Find segments.** Deterministic segmentation forms groups from customer features, without a purchase target. Each group is then checked for a distinct propensity to purchase and described by a short rule: a few features, each with a cutoff, joined by AND or OR. Each group becomes a named segment, such as “returning value seekers.” The goal shapes the use, not the groups: the pricing goal comes in only at the next step, when each group's response to price is measured.
- **Measure each segment's price response.** A weighted formula estimates buying from segment membership, rather than from hundreds of raw features, with price added. All segments are estimated together in one model, so every estimate draws on the full customer base rather than only one segment's customers. Fitting a separate model for each segment would discard that shared information. The joint model can take either of two forms, and in both each segment receives a recorded price response: how much that segment buys less as the price rises.²³

²³ J. H. Friedman & B. E. Popescu, “Predictive Learning via Rule Ensembles,” 2 Annals of Applied Statistics 916 (2008); P. E. Rossi, R. E. McCulloch & G. M. Allenby, “The Value of Purchase History Data in Target Marketing,” 15 Marketing Science 321 (1996). These sources support estimating response from segment membership; the advantage of joint over per-segment estimation reflects the modelling approach described here.

FORM A**Shared price sensitivity**

$$\text{buying} = \text{base} + a_1 \cdot S_1 + a_2 \cdot S_2 + \dots + a_p \cdot \text{Price}$$

All segments share one price coefficient (a_p) and differ only in their baseline level of buying.

FORM B**Segment-specific price sensitivity**

$$\text{buying} = \text{base} + a_1 \cdot S_1 + a_2 \cdot S_2 + \dots + b_1 \cdot (S_1 \cdot \text{Price}) + b_2 \cdot (S_2 \cdot \text{Price}) + \dots$$


A segment-by-price term (b_i) gives each segment its own price sensitivity.

In both forms, $S_i = 1$ if the customer is in segment i and 0 otherwise. Either way, the coefficients are estimated together in one model rather than in separate per-segment models.

- **Set each segment's price.** Each segment receives the price or discount its measured sensitivity supports, subject to commercial and legal constraints, and everyone in the segment is treated the same way.

FIGURE 6: FROM SEGMENTS TO PRICES

DAS forms segments from customer features, checked by propensity to purchase. One joint model then prices them.

**ILLUSTRATIVE EXAMPLE****SEGMENTS FOUND BY DAS**

S_1 Returning value seekers

S_2 Loyal full-price buyers

S_3 New customers

JOINT PRICE MODEL

$$\text{buying} = \text{base} + a_1 \cdot S_1 + a_2 \cdot S_2 + a_3 \cdot S_3 + a_p \cdot \text{Price}$$

All segments estimated together

PRICE OFFER PER SEGMENT

S_1 8% below list

S_2 List price

S_3 Intro offer

Not one model per segment: estimating jointly keeps the shared information.

Segment names and treatments are illustrative.

Why this helps explanation. Because segments do not overlap, each customer has one reason rather than a list, and the rule is the explanation rather than a proxy for it. Short reasons are also easier to understand: in experiments with 3,800 people, participants shown a simple model built on few facts could better predict what it would do. The same study cautions that transparency did not, on its own, help participants catch the model's mistakes.²⁴ Each segment's response to price is a single recorded number, so the basis for its price can be

²⁴ F. Poursabzi-Sangdeh, D. G. Goldstein, J. M. Hofman, J. Wortman Vaughan & H. Wallach, "Manipulating and Measuring Model Interpretability," Proc. ACM CHI (2021).

checked and reproduced. Each price also traces to a precise cutoff, which tells the customer what would change it (“one more order would qualify you”), the kind of specific reason credit law already requires.²⁵

Why this helps fairness. In two experiments, consumers rated prices that differ by customer group as fairer than prices set for each individual.²⁶ A meta-analysis of price-change studies finds that the perceived reason for a price change shapes how fair it seems,²⁷ and knowing how a price was set measurably changes perceived fairness.²⁸ Within each segment, everyone who meets the same rule pays the same price, so like customers are treated alike.

Choosing the Level of Personalization

Deterministic segmentation does not require organizations to choose between one uniform price and individual-level personalization. The business can determine the appropriate level of granularity.

Broader segments are easier to review and govern but may reduce pricing precision. Finer segments preserve more differentiation but increase the number of rules, treatments and records that must be maintained.

Finer segmentation can preserve more of the precision associated with individualized pricing while retaining explicit rule-based execution.²⁹

The design objective is therefore not maximum segmentation. It is an appropriate balance between:

- pricing precision;
- explanation complexity;
- governance burden;
- fairness and consumer-perception considerations.

²⁵ CFPB, Circular 2022-03, “Adverse Action Notification Requirements in Connection with Credit Decisions Based on Complex Algorithms” (May 26, 2022); Equal Credit Opportunity Act, 15 U.S.C. § 1691(d); Regulation B, 12 C.F.R. § 1002.9(b)(2); S. Wachter, B. Mittelstadt & C. Russell, “Counterfactual Explanations Without Opening the Black Box,” 31 Harv. J.L. & Tech. 841 (2018).


²⁶ A. Priester, T. Robbert & S. Roth, “A Special Price Just for You: Effects of Personalized Dynamic Pricing on Consumer Fairness Perceptions,” 19 J. Revenue & Pricing Mgmt. 99 (2020), cited in the FTC’s proposed statement.

²⁷ F. Tarrahi, M. Eisend & F. Dost, “A Meta-Analysis of Price Change Fairness Perceptions,” 33 Int’l J. Research in Marketing 199 (2016).

²⁸ S. Maxwell, “Rule-Based Price Fairness and Its Effect on Willingness to Purchase,” 23 J. Economic Psychology 191 (2002).

²⁹ J.-P. Dubé & S. Misra, “Personalized Pricing and Consumer Welfare,” 131 J. Political Economy 131 (2023).

FIGURE 7: THE PERSONALIZATION GRANULARITY TRADE-OFF



DAS increases pricing precision beyond uniform pricing while retaining explicit, explainable segment rules rather than moving to individual-level pricing.

Conceptual spectrum of personalization granularity. The ratings are illustrative and do not represent empirically measured rankings.

Stage 4: Disclosure

01 MAP / CLASSIFY

02 AUDIT

03 SEGMENT

04 DISCLOSE

05 GOVERN

The Explanation Should Come From the Production Rule

A central design principle of deterministic segmentation is explanation fidelity: the explanation provided about a personalized price should describe the logic that actually produced the decision.

A deterministic segment allows the same underlying record to support different levels of explanation.

LEVEL 1

At the Point of Pricing

“Member price, based on your purchase history with us.”



LEVEL 2

Additional Detail

“This offer applies to customers with at least six purchases in the previous 12 months who recently used a coupon. Only purchase history associated with this account was used.”



LEVEL 3

Internal or On-Request Record

The complete rule, data sources, treatment, effective dates and version, and a mechanism to correct inaccurate data, request human review, or opt out of automated profiling.

Note: The Level 3 mechanisms are designed to support consumer rights under the Colorado Privacy Act (CPA) and Colorado’s AI legislation (SB 26-189).

The important point is that each level originates from the same production rule, rather than from a separate explanation generated after the fact. The same structured record can produce a short disclosure or the full decision basis.

What the customer is told. The disclosure combines three elements: the segment the customer is in, any other information that moved that customer’s price individually, and the price itself.

Specific reasons do the work. In a lab test of personalized-pricing disclosures, generic labels alone barely improved consumer understanding.³⁰ A segment rule supplies the FTC’s “basis” element in one short, specific and verifiable sentence.

³⁰ OECD, “The Effects of Online Disclosure About Personalised Pricing on Consumers: Results from a Lab Experiment in Ireland and Chile,” OECD Digital Economy Papers No. 303 (2021).

AI-drafted wording needs a fixed reference. Language models write fluent explanations that can drift from the real logic.³¹ Checking an AI-drafted disclosure against a short, structured segment record is a simple and repeatable control.

LEGAL CAUTION: Whether any particular wording satisfies a legal disclosure requirement is a question for qualified counsel. Deterministic segmentation provides the underlying decision record from which accurate wording can be developed.

³¹ M. Turpin, J. Michael, E. Perez & S. R. Bowman, “Language Models Don’t Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting,” NeurIPS (2023); G. R. Pratama & K.-K. Tseng, “Accurate Ensembles, Fragile Narratives: Multi-Scale Stacking and a Fidelity Audit of LLM-Generated Explanations for Credit Risk,” arXiv preprint arXiv:2608.08126 (Aug. 8, 2026).

Stage 5: Governance

01 MAP / CLASSIFY

02 AUDIT

03 SEGMENT


04 DISCLOSE

05 GOVERN

One Record, Multiple Governance Functions

Each versioned segment record holds the segment identifier, production rule, data sources, treatment, disclosure text, effective date, version, and owner and approval record. To reconstruct a specific historical decision, not just the rule, the record should also preserve, where appropriate: the relevant input values or a reference to the data snapshot, the feature-derivation version, the pricing-model or treatment version, the jurisdictional rule set applied, the disclosure actually displayed and the decision timestamp.

FIGURE 8: ONE RECORD, FIVE GOVERNANCE FUNCTIONS



THE GOVERNING TEST

Can the company reconstruct any past price or offer from the rule, data and disclosure in force at the time?

Plug In, Plug Out: Change One Segment at a Time

Because each segment is a separate, versioned record, segment rules can change one at a time. A segment can be added, edited or retired without rewriting the other segments' rules or retraining a large model end to end. Two layers should be kept distinct: rule modularity, under which each membership rule is versioned independently, and pricing-model dependency, under which a change to the set of segments may require the joint price model to be re-estimated and every segment's treatment to be revalidated.

- **Plug in a segment.** Write a new rule, test it against historical data and approve it. The existing segments' rules stay as they are.

- **Plug out a segment.** Remove a segment that no longer performs, no longer passes review or relies on a feature that has been withdrawn. Its record is kept for historical reconstruction, and its customers fall back to a defined default treatment until they are reassigned.
- **Change one rule.** Adjust one threshold or one treatment, issue a new version with its own effective date, and leave every other segment's rule untouched.
- **Roll back one version.** If a change causes a problem, restore the previous version of that one segment's rule, together with the pricing-model version that was in force alongside it.

What still needs updating. After a segment is added or removed, or a rule change shifts which customers belong to a segment, the joint price model is re-estimated. It is a short model built on segment membership rather than hundreds of raw features, so the update is quick and its effect on every segment's price can be reviewed before release.

Why this matters for governance. In an opaque model, by contrast, removing one input or correcting one behavior usually means retraining the whole model. Every customer's price can shift, and the change has to be revalidated across the entire system. With modular segments, the scope of review matches the scope of the change.

Transitioning from Legacy Black-Box Models

Deterministic Segmentation Does Not Necessarily Require Immediate Replacement


Many organizations already operate pricing systems that use machine-learning models, weighted formulas or other complex decision mechanisms.

Post-hoc explainability tools can be valuable for model analysis and diagnostics. Their role, however, is different from a deterministic production rule. A post-hoc explanation estimates or describes why a model produced an output; it does not necessarily become the mechanism that produced the customer's price.³²

For organizations that want to move toward deterministic execution, one possible transition pattern is mimic, then migrate.

FIGURE 9: FROM BLACK BOX TO DETERMINISTIC PRICING

Mimic, then migrate



³² C. Rudin, "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead," 1 Nature Machine Intelligence 206–215 (2019).

Technology Requirements for Deterministic Segmentation


The five-stage framework depends on how segments are found in the first place. The underlying method can be understood in three steps: what the grouping is not steered by, how groups are formed, and what that means for anyone reviewing the result.

Grouping Without a Predefined Target

Standard models learn from target events. Most predictive models are built with supervised learning. They are trained to predict a target variable chosen in advance, such as whether a customer accepts a cross-sell offer, and each feature matters to the extent that it predicts that outcome.

A target-agnostic approach learns the processes behind the data. It analyzes the data without relying on a predefined target variable while training. The general technique is unsupervised learning, applied in a novel way to tabular customer data: instead of fitting the data to an outcome, it looks for the structure already present in it, at an unusually granular level of detail. The segments it produces remain fully explainable.

FIGURE 10: LEARNING FROM A TARGET VS. LEARNING FROM THE DATA



Top: a supervised model learns which features predict the target. Bottom: a target-agnostic approach uses the same features without a target and finds the groups the data itself supports, such as “returning value seekers.”

How Groups Are Formed

A population-centric approach prioritizes stable, persistent customer groups over geometric similarity or density.

- **Statistical consistency, not distance.** Groups are defined by sub-populations that are statistically consistent and persist in the data, rather than by optimizing a geometric or density-based objective, as k-means or DBSCAN clustering do.
- **Evidence-driven hierarchy.** Group formation is hierarchical. A sub-group is introduced only when it shows a statistically meaningful divergence from its parent population.
- **No preset structure.** The approach does not rely on a preset number of groups, distance thresholds or linkage criteria, and the order of the data does not change the result. A significance standard still determines when a divergence is meaningful, so the method is not free of every setting; what it avoids are the settings that would let an analyst shape the groups directly.
- **Readable membership.** Membership is governed by stable feature-level regularities rather than latent embeddings, so each group can be described by named features.

Target-Agnostic Discovery and Its Limits

These properties have a direct consequence for personalized pricing: customers are grouped by what their data shows, not by a target chosen in advance.

- **Groups form without a purchase target.** Customers are grouped from their recorded features. Propensity to purchase is used afterwards to evaluate and describe each group, and price comes in only when each group's response is measured.
- **Little to steer.** Because the group structure is not preset, analysts have little discretion over how groups form. Feature selection, preprocessing and the significance standard can still affect the resulting groups, so they remain subject to governance.
- **Transparent input handling.** The features, preprocessing and any weighting used to form groups should be documented and available for review. Where third-party data is used, its source should be recorded and assessed through the Classify and Audit stages.
- **Membership stays readable.** Because each group is defined by named features, a reviewer can find a sensitive inference before it sets a price.

WHERE IT STOPS: as Stage 2 notes, the approach describes the data faithfully but does not correct it, so the Classify and Audit stages still apply.

From here, the framework takes over. Each group is written as a short production rule (Stage 3), its response to price is measured and its treatment set (From Segments to Prices), and the resulting record supports disclosure and governance (Stages 4 and 5).

Scope and Limitations

Ensuring a pricing model complies with a patchwork of state and federal laws remains a massive operational hurdle. While heavily regulated industries like insurance already automate these checks, retail pricing still leans almost entirely on manual legal review.

Current Limitations: Architectural vs. Legal Scope

At present, deterministic segmentation provides the structural foundation for compliance, rather than the final legal determination. It makes pricing logic visible, reproducible and governable, but it does not independently determine whether:

- the underlying practice is permitted in a particular jurisdiction
- personal data was lawfully collected
- consent or another legal basis is sufficient
- a feature or proxy produces unlawful discrimination
- a higher or lower personalized price is permitted
- a particular disclosure satisfies applicable law
- a specific retention period is legally required

Today, these determinations remain matters for your organization's legal, privacy and compliance functions.

Future Scope: Automated Compliance Agents

A deterministic architecture is designed to close this exact gap. By converting opaque, black-box models into structured, readable rules, deterministic segmentation creates the necessary prerequisite for automated legal review.

In the foreseeable future, this architecture will enable the deployment of AI-driven compliance agents designed to evaluate deterministic segments against state and federal laws, potentially automating the bulk of the regulatory and fairness review that currently burdens legal teams.

Frequently Asked Questions

Regulatory Scope

Does the FTC's proposal ban personalized pricing?

No. The FTC proposal addresses circumstances in which personalized pricing may become unfair or deceptive and proposes disclosure expectations in specified situations. Separate state laws may impose additional restrictions or requirements.

Is the FTC proposal already final?

No. It is a proposed enforcement policy statement, not a final rule of general applicability.

Does using deterministic segmentation make a pricing practice compliant?

No. Deterministic segmentation is a technical and governance architecture. Legal permissibility depends on the applicable jurisdiction, data, treatment and circumstances.

Is segment-based pricing automatically permitted?

No. Some state regimes can apply to group-based pricing. The applicable criteria, exclusions and disclosures must be assessed under the relevant law.

Data and Fairness

Can deterministic segmentation use third-party data?

Yes. Deterministic segmentation can incorporate third-party data where it is provided as an input to the segmentation process. Whether that data should be used requires separate analysis of its provenance, authorization, applicable law, consumer expectations and governance risk. Any third-party inputs should also be documented and reviewed as part of the Classify and Audit stages.

Do inferred variables count as data used in the decision?

They should be included in the system inventory where they influence the pricing decision. Inferred variables and scores are features that must be mapped.

How should sensitive characteristics be handled?

Inputs that reveal or infer health, financial condition, household circumstances, urgency, vulnerability or protected characteristics should receive heightened review. The precise legal definition of sensitive data varies by jurisdiction.

Technical Architecture

Does deterministic segmentation require abandoning machine learning?

No. Machine learning can be used for analysis, segment discovery, threshold identification and validation. Deterministic segmentation concerns how the customer-facing pricing decision is executed. In Klary's approach, segment discovery itself is unsupervised and does not fit groups to a chosen outcome; see Technology Requirements for Deterministic Segmentation.

Why not simply explain the existing black-box model?

Post-hoc tools can help interpret a complex model, but the resulting explanation is distinct from the mechanism that generated the output. Deterministic segmentation makes the readable rule part of the production mechanism itself.

What happens if two segment rules match?

The production design should ensure mutually exclusive assignment or define deterministic precedence. The objective is one operative treatment and one reproducible decision path for each pricing event.

Can deterministic segmentation be used only for discounts?

No. The framework can be applied to any defined pricing or offer treatment. Whether higher prices are permitted is a separate legal and policy question.

Governance and Implementation

How often should segments be updated?

There is no universal interval. Changes should follow the organization's validation, monitoring and change-control process.

What happens when a customer's underlying data is corrected?

Because the decision is rule-based, correcting an input allows the organization to determine how the corrected data would affect segment assignment under the applicable rule.

What records should be retained?

At minimum, the architecture is designed to preserve the rule, relevant inputs and sources, treatment, disclosure wording, effective dates and version needed to reconstruct the decision. Legal retention requirements should be determined separately.

Does an existing pricing model have to be replaced immediately?

No. A deterministic model can be developed and evaluated in parallel before production migration.

What is the central governance test?

Can the organization reconstruct why a specific customer received a specific price or offer under the rule and version in force at that time?

Appendix A: Alignment with the FTC Proposed Enforcement Policy Statement

The table maps each position in the FTC’s proposed enforcement policy statement to the deterministic segmentation principles that address it, separating what the proposal says (paraphrased, not quoted) from our view of what it implies for deterministic segmentation. The statement is a proposal: it does not bind the FTC or the public, and any enforcement action would still have to establish a violation of existing law. Coverage is addressed by architecture where the rule itself supplies the answer, supported by architecture where deterministic segmentation makes it possible to act on the position without supplying the answer itself, and outside DAS scope where the position sits outside the pricing rule.

#	WHAT THE PROPOSAL SAYS	DAS PRINCIPLE	IMPLICATION FOR DAS	COVERAGE
1	The disclosure expectation applies where consumers reasonably expect that prices will not vary based on their personal data.	—	Whether the expectation applies is a market and legal question, not something the pricing rule can answer.	Outside DAS scope
2	Personalized pricing covers prices informed by personal data or inferences about the consumer, such as willingness to pay or likelihood of comparison shopping.	Full Explainability	Inferred features must appear in the explicit membership rule or pricing logic, so they are mapped and disclosed like any other input.	Addressed by architecture
3	Where the statement applies, a business should clearly and conspicuously disclose that the price is personalized, the basis for the personalization and the types of data on which it is based.	Full Explainability	The membership rule supplies why the customer is in the segment, the basis for treatment supplies why the segment receives its price, and the data record names the data types.	Addressed by architecture
4	Concealed personalization can cause injury that consumers cannot reasonably avoid, because they cannot see or respond to how the price arose.	Deterministic Execution	Fixed thresholds mean the factor that would change a customer’s price is known, so the record can help a consumer identify and correct relevant data.	Supported by architecture

#	WHAT THE PROPOSAL SAYS	DAS PRINCIPLE	IMPLICATION FOR DAS	COVERAGE
5	An effective disclosure conveys the relevant information about how the personalized price arose, not a generic label.	Deterministic Execution; Full Explainability	Each price traces to a specific rule and treatment basis, so the disclosure can state the actual reason rather than a category.	Addressed by architecture
6	Presenting a price as based on one kind of data, such as purchase history, when it reflects another, such as shopping habits with other firms, can be deceptive.	Deterministic Execution	Because the disclosure is generated from the same record that set the price, the stated basis can be checked against the data actually used.	Addressed by architecture
7	Collecting, using or disclosing personal data for pricing without obtaining and verifying consent can raise separate unfairness concerns, including for data obtained from vendors.	—	Deterministic segmentation records each feature's source, but verifying consent and provenance remains a separate legal and privacy task.	Outside DAS scope
8	The illustrations include personalization tied to circumstances such as emergencies, funerals, children in the household or an inability to comparison shop, described as situations raising Section 5 concerns where adequate disclosure is absent.	Full Explainability; Modular Governance	A rule keyed to such a signal would be visible in plain logic and can be reviewed and removed before it sets a price.	Addressed by architecture
9	The more sophisticated personalized pricing practices become, the less likely consumers are to benefit.	Deterministic Execution; Full Explainability	Keeping the production logic short and explicit makes the treatment easier for reviewers and consumers to evaluate, whatever model performance is claimed.	Supported by architecture
10	The FTC expressly declines, for now, to take a position on whether some personalized pricing practices are unfair even when fully disclosed.	Deterministic Execution	Deterministic segmentation makes treatment differences visible for that review, but whether a disclosed practice is fair remains a legal judgment.	Supported by architecture

Appendix B: Implementation Blueprint

Ten building blocks, each paired with a control objective, for implementing the five stages.

#	BUILDING BLOCK	CONTROL OBJECTIVE
1	Establish Scope	List every surface where a price, fee, discount, bundle or offer can vary. Identify the system and owner behind each surface.
2	Build the Feature Inventory	Record all raw, inferred and external inputs, their sources, derivation and directional impact on treatment.
3	Classify Risk	Mark first-party vs. third-party inputs, sensitive data, protected-class data, likely proxies and high-risk combinations.
4	Audit Outcomes	Measure who pays above and below baseline; compare groups; test whether customer-facing descriptions match actual logic.
5	Design Segments	Define mutually exclusive groups using short, readable rules and select an intentional level of granularity.
6	Assign Treatment	Estimate segment price response and set the price or offer subject to commercial and legal constraints.
7	Generate Disclosure	Derive customer-facing explanations from the same production rule and record the data categories used.
8	Version and Govern	Store rule, data sources, treatment, wording, effective dates and change history; monitor over time.
9	Apply Jurisdiction Policy	Separate analytical segmentation from market-level permissioning so state-specific restrictions and disclosures can be enforced.
10	Reproduce Before Launch	Demonstrate that a historical decision can be reconstructed from the stored record before the system goes live.

Sources and References

- [Federal Trade Commission, Proposed Enforcement Policy Statement Regarding Personalized Pricing, Aug. 19, 2026](#)
- [FTC public comment docket, FTC-2026-1057](#)
- [New York General Business Law § 349-a](#)
- [Maryland HB 895, Protection From Predatory Pricing Act, 2026 Md. Laws ch. 154](#)
- [New Jersey Fair Price Protection Act, P.L. 2026, c. 55 \(A4085 / S3612\)](#)
- [Connecticut Public Act 26-130](#)
- [New York One Fair Price Act, S8623B / A9349B](#)
- [NIST, Artificial Intelligence Risk Management Framework 1.0](#)
- FTC staff, Surveillance Pricing 6(b) Study: Research Summaries – A Staff Perspective (Jan. 2025)
- Greenberg Traurig, “Algorithmic Pricing Under Fire: State Restrictions on Personalized and Surveillance Pricing” (Sept. 3, 2026)
- Baker McKenzie, “United States: FTC Proposes Enforcement Framework for Personalized Pricing” (Aug. 26, 2026)
- Colorado S.B. 26-189 (signed May 14, 2026); Holland & Knight, “Colorado Governor Signs SB 189, Significantly Amending the State’s AI Law” (May 18, 2026)
- Colorado Privacy Act, Colo. Rev. Stat. § 6-1-1306(1)(a)(I)(C); 4 Colo. Code Regs. 904-3, Rule 9.04(C)
- California A.B. 2564 (introduced Feb. 20, 2026; not enacted, legislative session ended Aug. 31, 2026)
- CFPB, Circular 2022-03, “Adverse Action Notification Requirements in Connection with Credit Decisions Based on Complex Algorithms” (May 26, 2022); Regulation B, 12 C.F.R. § 1002.9(b)(2)
- Cornish-Adebisi v. Caesars Entertainment, Inc., No. 24-3006 (3d Cir. July 29, 2026)
- New Jersey v. RealPage, Inc., No. 25-cv-3057 (D.N.J. Mar. 31, 2026)
- United States v. RealPage, Inc., No. 1:24-cv-00710 (M.D.N.C. May 19, 2026)
- S. Bulla, Foundation for American Innovation, “Comment on the FTC’s Proposed Personalized Pricing Policy” (Sept. 2026)
- Consumer Reports nationally representative surveys (May 2024; Jan. 2025)
- J.-P. Dubé & S. Misra, “Personalized Pricing and Consumer Welfare,” 131 J. Political Economy 131 (2023)
- A. Priester, T. Robbert & S. Roth, “A Special Price Just for You: Effects of Personalized Dynamic Pricing on Consumer Fairness Perceptions,” 19 J. Revenue & Pricing Mgmt. 99 (2020)
- S. Maxwell, “Rule-Based Price Fairness and Its Effect on Willingness to Purchase,” 23 J. Economic Psychology 191 (2002)
- F. Tarrahi, M. Eisend & F. Dost, “A Meta-Analysis of Price Change Fairness Perceptions,” 33 Int’l J. Research in Marketing 199 (2016)
- T. Kim, K. Barasz & L. K. John, “Why Am I Seeing This Ad? The Effect of Ad Transparency on Ad Effectiveness,” 45 J. Consumer Research 906 (2019)
- K. Martin, “The Penalty for Privacy Violations: How Privacy Violations Impact Trust Online,” 82 J. Business Research 103 (2018)

- K. Martin & H. Nissenbaum, “Measuring Privacy: An Empirical Test Using Context to Expose Confounding Variables,” 18 Colum. Sci. & Tech. L. Rev. 176 (2016)
- K. Martin & H. Nissenbaum, “What Is It About Location?,” 35 Berkeley Tech. L.J. 251 (2020)
- C. Duhigg, “How Companies Learn Your Secrets,” New York Times Magazine (Feb. 16, 2012); K. Hill, “How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did,” Forbes (Feb. 16, 2012)
- CBC News, “Uber users with low phone batteries more likely to accept surge pricing” (May 21, 2016)
- J. H. Friedman & B. E. Popescu, “Predictive Learning via Rule Ensembles,” 2 Annals of Applied Statistics 916 (2008)
- P. E. Rossi, R. E. McCulloch & G. M. Allenby, “The Value of Purchase History Data in Target Marketing,” 15 Marketing Science 321 (1996)
- F. Poursabzi-Sangdeh et al., “Manipulating and Measuring Model Interpretability,” Proc. ACM CHI (2021)
- S. Wachter, B. Mittelstadt & C. Russell, “Counterfactual Explanations Without Opening the Black Box,” 31 Harv. J.L. & Tech. 841 (2018)
- C. Rudin, “Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead,” 1 Nature Machine Intelligence 206–215 (2019)
- OECD, “The Effects of Online Disclosure About Personalised Pricing on Consumers: Results from a Lab Experiment in Ireland and Chile,” OECD Digital Economy Papers No. 303 (2021)
- M. Turpin, J. Michael, E. Perez & S. R. Bowman, “Language Models Don’t Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting,” NeurIPS (2023)
- G. R. Pratama & K.-K. Tseng, “Accurate Ensembles, Fragile Narratives: Multi-Scale Stacking and a Fidelity Audit of LLM-Generated Explanations for Credit Risk,” arXiv preprint arXiv:2608.08126 (Aug. 8, 2026)
- Snopes, “Uber Increases Prices for Customers with Low Phone Batteries?” (Nov. 18, 2024)
- Letter from Sen. Sherrod Brown, Chairman, Senate Banking Committee, to Dara Khosrowshahi, CEO, Uber (July 3, 2024)
- New Jersey v. RealPage, Inc., No. 25-cv-3057 (D.N.J. Sept. 29, 2026) (order granting leave to amend)
- In re Zelis Repricing Antitrust Litigation (D. Mass. Mar. 30, 2026) (order denying motion to dismiss)
- In re MultiPlan Health Insurance Provider Litigation (N.D. Ill. June 2025) (order denying motions to dismiss)
- Gibson v. Cendyn Group, LLC, 148 F.4th 1069 (9th Cir. 2025)
- Mayer Brown, “Algorithmic Price Fixing After Cornish-Adebiyi: The Third Circuit Revives Hub-and-Spoke Claims Against Shared Pricing Software” (Aug. 3, 2026)
- Milbank, “The Latest Intelligence: Antitrust Developments Impacting Providers and Users of Algorithmic Tools” (June 2026)
- Arnold & Porter, “Algorithmic Pricing: Navigating Antitrust and Consumer Protection Risks” (June 2026)

This whitepaper is provided solely for general informational and educational purposes. It does not constitute legal, regulatory, compliance or other professional advice and should not be relied upon as a determination that any pricing practice, data use, disclosure or technical architecture complies with applicable law.

Laws and regulatory guidance concerning personalized, algorithmic and surveillance pricing vary by jurisdiction and are continuing to evolve. The applicability and interpretation of any law may depend on the specific facts, data, pricing practices and consumer interactions involved.

Organizations should consult qualified legal counsel and other appropriate professional advisers before implementing or modifying personalized pricing practices.