Your Data, Their Profit:
The Consumer Cost of AI Surveillance Pricing

Written Testimony before the Senate Judiciary Committee
Sub-Committee on Crime and Counterterrorism

August 4, 2026

Robert B. Hedges, Jr.
Digital Fellow at the Initiative for the Digital Economy,
Massachusetts Institute of Technology
Former Chief Data Officer, Visa (Retired)

In this paper, we will define surveillance pricing, (I) review the history of corporate efforts to develop the data collection and analysis capabilities that support the practice, (II) mechanically explain how it works, (III) assess the consumer harm, (IV) review potential opportunities to reduce the harm, and (V) propose specific policy solutions. The goal of this testimony is to provide a matter-of-fact explanation of how the complex analytical techniques and technology platforms that support surveillance pricing work as well as offer guidance on what policy solutions should be pursued. Investing in developing a stronger public awareness and understanding of how surveillance pricing works will hopefully develop the momentum for the required actions by consumers, data scientists, corporates, tech platforms, and policymakers to end this consumer-harming practice.

The exuberant promise around artificial intelligence (AI) is that it will make everyday life easier, more efficient, and less expensive for everyone. The promise is real and we should embrace the transformational societal benefits that responsible AI adoption can deliver. For two decades, however, accompanying the growth of the digital economy has been the expanding practice of digital platforms to routinely collect, analyze, and ultimately sell consumers' identity and behavior data. The commercial process has earned the term "data monetization". Select firms have learned to "monetize" the data that they collect about consumers – they analyze the data to influence the messaging that you see and what you buy – and translate those behavioral insights into algorithms to be implemented in their marketing and pricing systems. "Algorithmic pricing" is a term used to define the broad set of data and analytical model-based approaches to pricing. The recent rapid developments in real-time computing, messaging, and AI have extended the reach and speed of retailers' and digital platforms' ability to real-time make pricing adjustments. Within the broad category of algorithmic pricing, this real-time practice is called "dynamic pricing". In that specific domain, the use of extensive consumer behavior data, collected without the consumers' awareness or informed consent, is called "surveillance pricing". This testimony focuses on surveillance pricing – how it works, the harms it causes, and what solutions can stop it.

Importantly, algorithmic or dynamic pricing is not inherently bad. Making real-time price adjustments in response to measured demand surges (or declines) and market data make inherent good sense to a well-functioning market. The ability to rapidly and precisely make price adjustments can help a market function more efficiently for the benefit of all. What deserves scrutiny, however, is the practice of surveillance pricing when the pricing changes are based on consumer-specific personal data, collected without the consumers' informed consent, and often used counter to the best interests of the consumer. The power of the data science and technology is compelling. The lack of transparency regarding both the collection of consumer data and how that data is then used is problematic and delivers real consumer harm.

I. Consumer Behavior Data Has Become a Valuable Signal to Inform Retail Marketing and Pricing

To begin, let's first establish a clear understanding about what the large digital marketing platforms are in the business of doing. The large digital platforms collect data about consumers from many sources, including shopping history, web browsing activities, locational data, the content of email messages, inferred emotional state, inferred purchase intent, device type, battery life, mouse clicks, loyalty programs, and payment card activity, and they assemble and analyze it into detailed consumer profiles.[1] The digital platforms assemble at scale what they call "identity graphs" that map all the data that they collect about you to your consumer profile. These detailed profiles know, or infer, your behaviors, preferences, vacation ambitions, restaurant choices, taste in fashion, politics, health conditions, relationship status, presence of children, and countless other behavior and preference attributes.

Your consumer profile can then be managed to produce "signals" about what a consumer might want. Signals project your future shopping needs, your sense of urgency to make a purchase, your brand loyalty, your typical shopping behavior, and your "willingness to pay" or price sensitivity. This is not normal "maximization of profits" in a market economy. Competition authorities all around the world have identified and published reports on the anticompetitive effects of algorithmic and surveillance pricing, including the risk of collusion to fix prices. For a market economy to provide optimal results, transparency is a key factor, consumers need to be able to compare prices objectively and switch if something better comes along. Surveillance pricing is intrinsically opaque. The signals are sold to marketers, advertisers, digital marketplaces, and retailers to inform their strategy,

[1] Federal Trade Commission, A look at What ISPs Know About You: Examining the Privacy Practices of Six Major Internet Service Providers, FTC Sta ff Report (Washington, DC: Federal Trade Commission, October 2021).

marketing, and pricing – not just in general, but when they are dealing with specific consumers. It is simply the business model running exactly as designed.

As you would imagine, the signals are valuable. The signals inform marketers, retailers, and digital platforms, on a wide range of potential shopping or consumer purchase issues – for example, you might need suntan lotion because you are about to go on a beach vacation, or you need to purchase a port-a-crib because you just had a child, or you are enthusiastic about purchasing new furniture because you just bought a new home. The use of signals can increase the effectiveness of targeted advertising, strengthen the conversion rate on marketing leads, inform messaging and behavioral manipulation techniques to encourage purchases, and fuel product bundling and pricing models. Signals, derived from consumer data, are the vehicles through which insights about consumer behavior and preferences are traded. The process of collecting and analyzing consumer data, converting it to sellable signals, and then broadly making the signals available either through data marketplaces or one-off strategic data relationships is what has come to be called data monetization.

Online shopping behavior has been the focus of the collection and application of consumer preference and behavior data for decades. What websites you browse, how much time you spend reviewing products, what items you click on, what items make it into your shopping cart – these are all consumer behavior data that can be converted into signals, seeking to understand consumers' intentions and urgency, and then marketed to third parties to allow them to influence your future purchase behavior.

The emergence of agentic commerce takes these decades-old digital marketing practices and puts them into an even stronger closed loop network. The digital platforms make money from the consumer profiles that they have developed by leveraging the profiles' insights to decide what products to show you, what to recommend, and how much to charge. The new wave of AI assistants and shopping agents sits inside this same arrangement. When a platform offers you an agent that will book your travel or fill your cart, it is worth remembering who that agent works for. It is not necessarily searching for your best deal. It is working to produce the best outcome based on the algorithms that the agent has been programmed to optimize, usually for the company that built the agent. And every prompt or question a consumer asks the agent adds to the digital platform's profile of that consumer, further refining the platform's knowledge of the consumer – making the platform even more effective in messaging and marketing to the consumer.

1. Reflecting the value of consumer behavior-based signals, the digital platforms and leading retailers have made big commitments to pursuing R&D and patents

While the actual scale of the practice of surveillance pricing is hard to assess from outside the world of data science and technology practitioners, the ambition of the digital platforms and tech players in this space is very visible. Today, the essentially unconstrained access to consumer behavior data, extraordinary gains in compute power, and real-time connectivity have produced a broad array of creative data collection techniques and algorithmic pricing strategies designed to maximize revenue based on previously untrackable or uncollectable consumer behavior data. Given the potential value of the ability to price based on consumer behavior signals, we have seen significant corporate effort and resources dedicated to developing and building-out the capabilities.

While the actual consumer data collection process may be intentionally designed to be invisible and the pricing algorithms to be opaque, the level of energy to pursue innovations in surveillance pricing is quite visible. All formal innovation patent requests are submitted to the US Patent Office. The resulting historical record of corporate effort to develop innovative vehicles to collect and assess consumer behavior data and then apply consumer behavior signals to product pricing is quite well documented. The volume of patent applications filed are a good indicator of corporate intent to build strong capabilities in this space – even as the same companies work to keep specific practices opaque.

We can see the data monetizers' ambitions in the patent record, where new methods of collecting and modeling consumer behavior to inform product pricing are claimed, filing by filing, as proprietary inventions. The instinct to claim ownership of proprietary pricing tactics is not new. As early as 2005 Amazon applied for – and in 2010 was granted – a "method and system for dynamic pricing" of its web-services offerings, with Jeff Bezos himself listed as a co-inventor (U.S. Patent 7,743,001). Uber followed in 2013 with its surge-pricing application, "System and method for dynamically adjusting prices for services" (U.S. Patent Application 2013/0246207 A1), co-invented by then-CEO Travis Kalanick, which claimed the now-familiar mechanics of multiplying a default fare when demand outruns driver supply. (Uber later abandoned the application; the practice it described became an industry standard anyway.)

What was innovative and novel in the early 2010s has today become widespread and industrialized. Over the period 2015 to 2025, twenty-three of the leading digital platform, retail, payments and technology companies filed more than 3,100 U.S. patent applications, 3,143 based on our recent scan of the US Patent and Trademark Office database, each staking a proprietary claim to a distinct method of collecting consumer data, modeling and segmenting consumer behavior, forecasting demand, determining prices, and targeting personalized discounts and incentives. The chart below summarizes where that activity concentrates. Notably, the deepest vein of patent filings is not price-setting itself but the surveillance techniques and supporting technology that supports and informs it.[2]

[2] Author's analysis of Google Patents and USPTO Patent Public Search data, queried July 31, 2026: U.S. pre-grant publications, filing dates January 1, 2015–December 31, 2025, CPC subclasses G06Q 30/0201–30/0207, twenty-three named assignees with name variants queried separately and deduplicated by publication number. Full query set and methodology caveats are provided in the accompanying appendix.

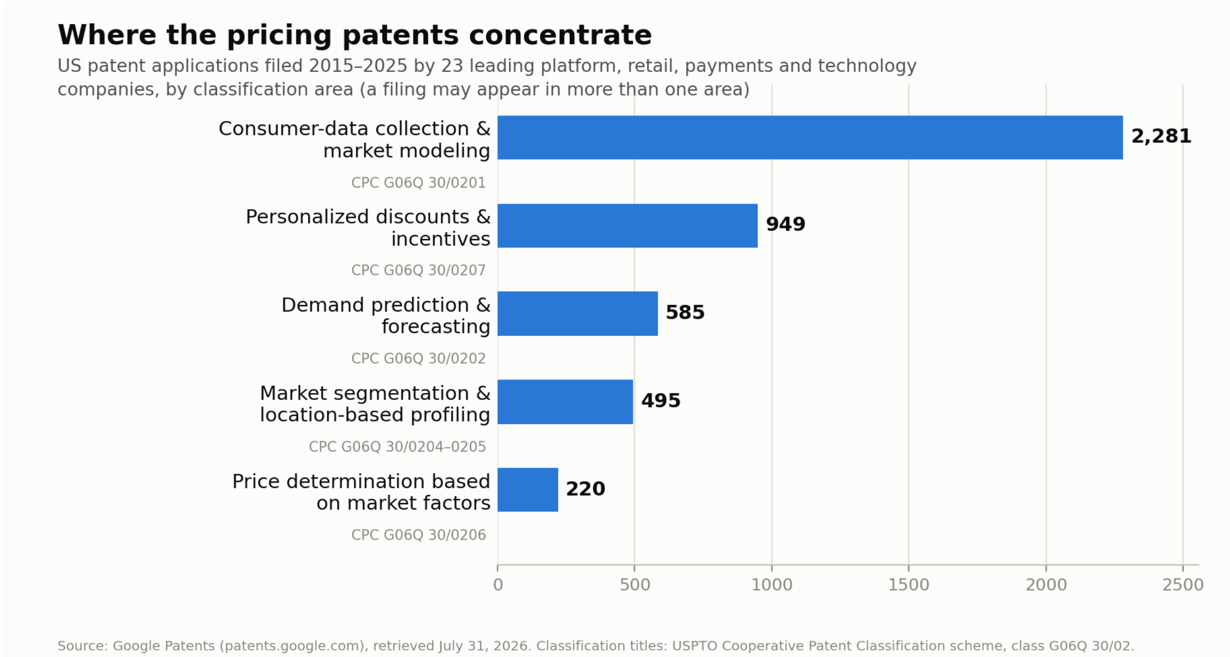

Reading the individual filings themselves is very instructive. Each patent filing documents quite deliberate, sometimes startling, efforts to collect and harvest consumer behavior data, usually from consumers who have no idea that their behavior is being tracked and measured, to infer purchase intent, price awareness and willingness to pay, and to convert those inferences into tactics to have the consumer pay more for an item and generate increased revenue for the retailer. Out of the 3,143 applications filed over the last eleven years, consider these four:

- **Instacart** (Maplebear Inc.), U.S. Patent 12,417,465 (filed 2024, granted September 2025): uses data from in-store sensors, including cameras and weight sensors mounted on smart shopping carts, together with app interactions and item-replacement behavior to compute an individual shopper's *price-sensitivity score*, then targets prompts and offers to that shopper accordingly.

- **IBM**, U.S. Patent 11,182,835, “Individual online price adjustments in real time” (filed 2019, granted 2021, and since assigned to DoorDash): monitors a user's online communications to identify interest in a product *and how much the user already knows about its price*, then adjusts the price shown to that individual in real time with pricing calibrated to each consumer's ignorance.

- **Walmart**, U.S. Patent 11,429,992, “Systems and methods for dynamic pricing” (filed 2018, granted 2022): estimates demand and price elasticity for the items in a shopper's electronic basket using a multi-armed bandit model in plain terms, continuous live pricing experiments run on customers and sets the displayed price from the results.

- **Amazon**, U.S. Patent 10,102,547, “Deal scheduling based on user location predictions” (filed 2015, granted 2018): aggregates users' location check-ins, travel patterns, mobile-device location data, preferences and purchase histories to predict where consumers will congregate, then schedules, and sizes deals for those places and moments.

The corporate intent, visible in even a basic review of patent records, is clear. The ambition of the digital platforms’ consumer behavior data collection efforts is striking. More extensive scrutiny of the patents and research and innovation agendas of the large digital platforms is warranted, and the right policymaking body should explore in detail the intent, capabilities, and current implementation of the programs outlined in the patent filings

2. Major digital platforms and retailers have also made major acquisitions to support surveillance pricing

The ability to execute surveillance pricing is highly specialized. Success requires the ability to assemble the required consumer shopping behavior data sets, the expertise to synthesize that data into high-fidelity signals regarding consumers’ willingness to pay, the modeling and data engineering skills to real-time deliver the willingness to pay signals to digital marketing platforms and online retailers during consumers’ shopping journeys, and the data science ability to test and learn from continuous consumer testing regarding how consumers react to alternative pricing. Internal R&D and innovation and ultimately the filing of patents, is one path towards the development of the required specialized skills and platforms to support surveillance pricing. Given the value and uniqueness of the skills, however, it is not surprising that major digital platforms and retailers have also been acquiring established companies in the “pricing technology” space rather than pursuing the slower course of organically building internal capabilities. Since 2015, at least seven significant acquisitions of pricing, personalization, and customer-data-science companies have been made by major digital platforms and retailers, with disclosed deal values exceeding $4 billion (several deal terms were not disclosed). The most strategic deals include:

- Kroger closed its dunnhumby USA joint venture with dunnhumby Ltd. (2015) acquiring the business’s US assets and analytics team to form a Kroger wholly-owned customer-data-science business subsidiary, 84.51*; terms undisclosed
- Mastercard acquired Applied Predictive Technologies/APT (2015) for $600 million
- Walmart acquired Jet.com (2016) for approximately $3.3 billion, then the largest U.S. e-commerce acquisition
- Walmart (Flipkart) acquired Upstream Commerce (2018); terms undisclosed
- Instacart acquired Eversight (2022) for ~$59 million, the AI pricing-and-promotion testing platform later identified as the engine behind Instacart’s 2025 price experiments

- Mastercard acquired Dynamic Yield, an AI-powered personalization platform (2022); terms undisclosed
- Interpublic acquired Intelligence Node, a real-time competitive pricing and e-commerce intelligence platform (2024); terms undisclosed

Merger and acquisition activity in the real-time consumer data collection and analytics space has recently declined. Corporate innovation and M&A activities efforts have more focused on opportunities in Gen AI applications and cybersecurity. The high-profile nature of surveillance pricing-related acquisitions may also have become viewed as risky in terms of generating public attention. The major players now prefer to develop and power internally and minimize public awareness of activities in this space.

II. How the Analytical Mechanics of Surveillance Pricing Work

The inputs used by dynamic pricing models are quite specific. Leveraging data from corporate marketing profiles, or third-party signals available from Google, Meta, Amazon, or a set of lesser-known digital platforms and consultancies, detailed assessments of consumers' willingness to pay can be developed. Willingness to pay is the data science and marketing analytics terms used to characterize a consumer's price sensitivity. A "high" willingness to pay means the consumer is less price sensitive, really wants to purchase the product, has a strong and urgent need, or simply declines to negotiate. A high willingness to pay suggests that the consumers' desire for a product or service is not going to be dampened by a higher price. In contrast, a "low" willingness to pay suggests price sensitivity. A consumer who has low willingness to pay is going to carefully comparison price shop, look for a better deal, ask for a discount, or simply leave the item unpurchased in their shopping cart. For the revenue-maximizing firm, with the technical ability to act on the awareness of consumers' price sensitivity, the ability to accurately assess and infer willingness to pay is a powerful capability.

Relative willingness to pay can be inferred using a variety of sources of consumer behavior data. Data scientists have studied this issue extensively. A consumers' willingness to pay can be inferred from a variety of data sources[3] – a person's precise location, browser search history, purchase history, and even small behavioral cues such as mouse movements and the frequency and pattern of items left sitting in an online shopping cart.

[3] Aniko Hannak, Gary Soeller, David Lazer, Alan Mislove, and Christo Wilson, "Measuring Price Discrimination and Steering on E-commerce Web Sites," in Proceedings of the 2014 Internet Measurement Conference (New York: ACM, 2014), 305–18, https://personalization.ccs.neu.edu/static/pdf/imc151-hannak.pdf; Jakub Mikians, László Gyarmati, Vijay Erramilli, and Nikolaos Laoutaris, "Detecting Price and Search Discrimination on the Internet," in Proceedings of the 11th ACM Workshop on Hot Topics in Networks (New York: ACM, 2012), 79–84.

The consumer data can be sorted and understood across a few categories:

- Location: The ZIP code or an internet address serving as a proxy for your income and buying power
- Time of day: The time of day that the research, search, or purchase is made
- Purchase history: Your account status, such as your purchase history with a retailer, your loyalty tier, and your subscription status to understand your commitment to the brand or product category
- Digital access: The type of phone or computer, its operating system, and the fingerprint it leaves behind, which can identify you even when you are not signed in
- Shopping intent or referral source: The shopping search path that brought you to the item (e.g., a direct Google search for a specific model item) signaling an urgent need vs. consumers arriving from coupon aggregator sites who are probably deal shopping
- Behavioral cues: The trackable digital behaviors, such as how long you linger on a web page, how often you return, the speed with which you maneuver through an online product catalog, and what you abandon in your digital shopping cart, all provide a measure of urgency
- Cross-retailer shopping behavior: The number of product detail pages (PDPs) that you open on a shopping journey signals active comparison shopping. Opening multiple PDPs during the same session indicates the consumer may be comparing products' features and pricing

These are all behavioral cues from which your willingness to pay can be inferred. And there is also external data bought from brokers, which fills in demographics and financial estimates.

Some of the technology to capture consumer behavior is relatively intuitive and straight-forward, other solutions are less so. To track real-time consumer shopping behavior and derive context-aware willingness to pay, retailers and digital platforms use an integrated stack of front-end data capture and real-time consumer profile integration technologies. For example, to assess how actively a consumer is comparison-shopping for a specific product, these data capture tools track the count of product detail page (PDP) views in real time and immediately synthesize them into a willingness-to-pay signal. Tools like Google Analytics 4 and Adobe Analytics use JavaScript trackers to capture consumers' page views. Software solutions like Amplitude or Mixpanel track user-specific shopping behaviors. First-Party Cookies enable website-specific tracking and store it over time. Server-Side Tracking, or Edge Infrastructure, allows platforms to manage data directly within their own secure cloud environment, ensuring privacy compliance and data security at the first-party level. AWS CloudFront and Cloudflare Workers are important tools for this effort.

All this data is captured in real-time in customer data platforms (CDPs) that operate with sub-second latency, allowing it to be streamed essentially instantaneously into retailers' pricing and personalization engines, including those running multi-armed bandit programs.

High profile vendors of this technology include Segment (Twilio), Tealium, mParticle, and ActionIQ.

Ultimately the behavioral data and willingness to pay signals need to be acted upon through real-time personalization and pricing engines. Dynamic Yield (owned by MasterCard) and Insider One are examples of successful companies in this space.

Beyond the consumer data collection and analysis, retailers can then actively test and measure how alternative pricing translates into consumer purchase volume. Retailers build models based on willingness to pay signals, test in real-time how the derived pricing options affect consumer demand and revenue, and then extends the resulting pricing strategy across the applicable shoppers.

Data scientists conduct real-time quantitative tests of alternative pricing based on the consumer signals to maximize revenue. The analytical techniques include A/B Split Testing where different price points, based on consumer-specific signals, are offered on identical products, to measure price elasticity; Matched-Pair testing where consumers with common demographic and behavioral attributes are offered different pricing on a controlled basis to measure the potential for revenue growth. The most sophisticated approach is multi-armed bandit (MAB) testing.[4] The name comes from the colloquial term for a slot machine ("one-armed bandit"). Imagine a gambler standing in front of a row of slot machines, trying to figure out which one has the highest payout percentage, or yield, while minimizing the amount of money wasted on losing machines.

Traditional A/B testing splits consumer traffic evenly between Price A and Price B, for a fixed period, regardless of which price is performing better.

Multi-armed bandit testing continuously shifts consumer traffic toward the winning price in real-time. A/B testing applies the randomized controlled experiment which was formalized by R.A. Fisher in the 1920s to commercial decisions; the first online A/B test was run by Google in 2000[5]. Multi-armed bandit testing is the modern, real-time equivalent, in which digital platforms and retailers invest significant technology and compute resources to analyze consumer purchase behavior and act on it immediately.

The allocation and routing of traffic across alternative price points is managed using a set of data science algorithms – Thompson (or Bayesian) sampling, Upper Confidence Bound

---

[4] Richard S. Sutton and Andrew G. Barto, Reinforcement Learning: An Introduction, 2nd ed. (Cambridge, MA: MIT Press, 2018), chap. 2.

[5] Ronald A. Fisher, The Design of Experiments (Edinburgh: Oliver and Boyd, 1935); Ron Kohavi, Diane Tang, and Ya Xu, Trustworthy Online Controlled Experiments: A Practical Guide to A/B Testing (Cambridge: Cambridge University Press, 2020), chap. 1.

(UCB), or epsilon-greedy. These allocation strategies are used to discover, in real-time, the revenue-maximizing price for each consumer signal and segment.[6]

A whole ecosystem of companies are active in this Pricing Optimization and Modeling space – the experimentation platforms Kameleoon, Optimizely, and the retail pricing specialists DynamicPricing.ai, Competera, and Eversight are all examples. Major consultancies – McKinsey, through its Periscope pricing solutions; Bain, through its Advanced Analytics practice and proprietary Dynamic Pricing Engine; and the pricing specialist Simon-Kucher all maintain pricing optimization practices that build the required data infrastructure, deploy adaptive experimentation, and run digital commerce operations.

Advanced multi-armed bandit software supports willingness to pay consumer signals. The systems use what are called "contextual bandits" to ingest upstream consumer behavioral signals that will indicate high or low willingness to pay as the consumer visits the online store.

These consumer behavior data elements are collected, analyzed, and employed to develop a wide range of signals to inform a full spectrum of strategic and tactical pricing moves. willingness to pay becomes a critical and actionable consideration:

- Developing a price for a specific consumer or segment of consumers for a specific product
- Making pricing adjustments triggered by a consumer's shopping behavior, in response to specific activity, such as a discount that appears only after you walk away from a full cart
- Revising pricing upward or downward for segments or groups of consumers, however, defined. A whole group can be priced by proxy, such as everyone on a certain device or in a certain neighborhood. Or the price can be retuned continuously as you move across sessions and days

Willingness to pay as a variable in developing pricing strategies and tactics has become commonly deployed across a wide range of retail and hospitality categories including groceries and household goods, in apparel, in airfare and hotels, and increasingly in residential rent set by software.[7] As mentioned above, a small industry of specialist

---

[6] William R. Thompson, "On the Likelihood That One Unknown Probability Exceeds Another in View of the Evidence of Two Samples," Biometrika 25, nos. 3–4 (1933): 285–94; Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer, "Finite-Time Analysis of the Multiarmed Bandit Problem," Machine Learning 47, nos. 2–3 (2002): 235–56.

[7] United States v. RealPage, Inc., No. 1:24-cv-00710 (M.D.N.C. filed August 23, 2024); U.S. Department of Justice, "Justice Department Sues RealPage for Algorithmic Pricing Scheme That Harms Millions of American Renters," press release, August 23, 2024.

vendors and consultants has developed who sell analytical platforms and willingness-to-pay estimates to the retailers and digital platforms.

Finally, the technology for detecting and defending against surveillance pricing is starting to catch-up with the software and tools that enable it. To research and test for evidence of surveillance pricing, one no longer needs to rely on volunteer consumer "mystery shoppers" to collect pricing data from their own digital shopping experiences. As an alternative, synthetic digital consumer profiles can be created and used to conduct shopping research. By varying the synthetic consumer profiles' key data signals that surveillance pricing systems rely on (e.g., location, device type, browsing history), one can simulate the purchase activities of different consumers and see how the product pricing varies by consumer. One company, Fair Patterns, has shared that they have very specifically been developing and testing this capability. Using AI capabilities trained to recognize "dark patterns" web design features, Fair Pattern has developed the capability to shop and test how prices offered for products can vary based on the different consumer behaviors profiles. Earlier academic audits that first developed this methodology, Hannak et al. Mikians et al., used exactly this technique. On the commercial side, Fair Patterns, founded by Marie Potel-Saville, uses AI trained to recognize dark-pattern design features to detect and remediate manipulative interfaces.[8] The same technology capabilities can now be applied to the detection of surveillance pricing.

With companies like Fair Patterns showing the way in development of new tools and protocols to detect and alert consumers to price offers that appear disadvantaged, consumers have reason to hope that technology will catch-up to better protect their interests.

1. The effectiveness of consumer behavior data collection and analysis has led to corporate overreach

Much can happen when corporates think that their practices are outside of the public view or can avoid public scrutiny. The analytics are sophisticated. The tech is cutting edge. The business managers may be a bit enraptured with their own thinking. The resulting public outcry about what is commonly experienced as corporate over-reach comes almost as a surprise.

Since 1999, there have been multiple high-profile episodes of documented instances in which companies varied prices or offers based on algorithmic pricing routines or inferred consumer characteristics:

---

[8] Fair Patterns, accessed August 2, 2026, https://www.fairpatterns.ai/; Hannak et al., "Measuring Price Discrimination and Steering on E-commerce Web Sites."

- Varying the price of a cold drink in relationship to weather temperature: In 1999, Coca-Cola confirmed it had been developing vending machines that would automatically raise the price of a drink in hot weather, and its chief executive publicly defended the idea as fair. The company said the technology had not been placed in any consumer market, and it retreated within days of the public backlash.[9]
- Levering the digital devices that you use as a proxy for affluence: In 2012, the travel site Orbitz was found to be steering online shoppers who browsed on Mac computers toward pricier hotels. Orbitz inferred that Mac users would seek more upscale hotel brands than PC users and said Mac users were about 40 percent more likely to book four- or five-star hotels.[10]
- Using geolocation as a proxy for affluence: In 2012, a Wall Street Journal investigation found that multiple retailers including Staples varied online prices based on a shopper's inferred location (from ZIP code or IP address), and specifically for Staples, whether the shopper lived near competing office supply stores. Customers farther from competitors often saw higher prices. The Wall Street Journal also found that the ZIP codes shown the discounted price had a weighted average income of roughly $59,900, while those shown the higher price averaged $48,700, meaning lower-income areas systematically paid more.[11] In 2015 ProPublica reported that The Princeton Review had charged different prices for identical online SAT tutoring packages based on ZIP codes with larger Asian-American populations were significantly more likely to receive higher prices, although the company said the differences reflected the business costs rather than discriminatory intent.[12]
- Using sense of urgency: Online sellers have long employed scarcity and urgency cues such as countdown timers, low-stock warnings, "only 2 left", and cart-abandonment discounts. These tactics are designed to influence purchasing behavior by increasing perceived urgency or fear of missing out, rather than by reflecting changes in the seller's underlying costs. When such scarcity claims are fabricated, in 2022, the Federal Trade Commission's staff report identified them as deceptive dark patterns.[13][14]
- Using inferred consumer price sensitivity: Consumer Reports' two-year investigation on car insurance in 2015 analyzed over two billion price quotes and documented

---

[9] Constance L. Hays, "Variable-Price Coke Machine Being Tested," New York Times, October 28, 1999, sec. C.

[10] Dana Mattioli, "On Orbitz, Mac Users Steered to Pricier Hotels," Wall Street Journal, June 26, 2012.

[11]Jennifer Valentino-DeVries, Jeremy Singer-Vine, and Ashkan Soltani, "Websites Vary Prices, Deals Based on Users' Information," Wall Street Journal, December 24, 2012.

[12]Julia Angwin, Surya Mattu, and Jeff Larson, "The Tiger Mom Tax: Asians Are Nearly Twice as Likely to Get a Higher Price from Princeton Review," ProPublica, September 1, 2015

[13]Federal Trade Commission, Bringing Dark Patterns to Light: An FTC Staff Report (Washington, DC: Federal Trade Commission, September 2022).

[14]Robert B. Cialdini, Influence: The Psychology of Persuasion, rev. ed. (New York: Harper Business, 2007).

insurers setting premiums partly on a consumer's predicted tolerance for a price increase rather than on their risk of filing a claim. Consumer Reports named the practice the "schmo tax"; by late 2015 eleven states, plus D.C. had banned it. The vendor, Earnex, reported that 45% of large insurers used it.[15]

- Public deployment of dynamic pricing capabilities: In February 2024, Wendy's announced its investment in approximately $20 million in new digital menu boards that could change prices through the day, with dynamic pricing tests to begin as early as 2025. The public reaction to what looked like surge pricing on burgers was swift and negative, and within days the company clarified it would not raise prices at peak times. That same year, the Federal Trade Commission opened a formal study of surveillance pricing, ordering a group of intermediary firms to explain how their products work, releasing initial findings the following winter. Retailers were using consumers' precise location, demographics, browsing histories, even mouse movements and abandoned shopping carts, to set "targeted, tailored prices."[16][17][18]
- Using consumer behavior data and willingness to pay signals: In December 2025, the grocery-delivery service Instacart was spotlighted in a joint investigation by Consumer Reports, Groundwork Collaborative, and More Perfect Union for charging different shoppers' different amounts for the same grocery items (e.g., eggs) at the same time, Instacart ended the program later that month.[19][20][21][22]
- Using consumer behavior data and willingness to pay signals: In June 2026, Consumer Reports turned to ride-hailing, testing identical Uber and Lyft routes across seventeen states. All thirty routes tested produced at least two distinct price groupings, separated by at least five percent, for the same ride at the same time; one Kansas City Route returned twenty-nine different prices for the fifty-five would-

---

[15] Consumer Reports, "The Truth About Car Insurance," Consumer Reports, July 30, 2015; Jeff Blyskal, "Price Optimization Helps Car Insurers Figure Out Whether You're a 'Schmo,'" Consumer Reports, September 3, 2015; "The Price of Price Optimization," Insurance Journal, October 19, 2015.

[16] "Wendy's Clarifies No Surge Pricing after CEO 'Dynamic Pricing' Comment." *Washington Post*, February 27, 2024.

[17] Federal Trade Commission. "FTC Issues Orders to Eight Companies Seeking Information on Surveillance Pricing." Press release, July 23, 2024.

[18] Federal Trade Commission. "FTC Surveillance Pricing Study Indicates Wide Range of Personal Data Used to Set Individualized Consumer Prices." Press release, January 17, 2025.

[19] Wells, Katie, Lindsay Owens, Angel Han, and Alan Smith. *Same Cart, Different Price: Instacart's Price Experiments Cost Families at Checkout*. Washington, DC: Groundwork Collaborative, with Consumer Reports and More Perfect Union, December 9, 2025.

[20] Consumer Reports and Groundwork Collaborative. "Instacart Pricing Investigation: Data and Methodology." GitHub, December 2025.

[21] Associated Press. "Instacart Ends a Program That Tested How Much Shoppers Would Pay by Showing Different Prices for the Same Items." *Fortune*, December 22, 2025.

[22] Consumer Reports Advocacy. "Consumer Reports Calls on FTC to Investigate Uber and Lyft over Pricing Practices." June 2026.

be-riders, the median gap between the lowest and highest price grouping on a route was roughly fifty percent, and about eleven percent of advertised discounts were fake, calculated against inflated base prices. Consumer Reports had asked the FTC and state attorneys general to investigate.[23][24]

As evidenced by the December 2025 exposure of Instacart's practices, there is significant public trust and brand risk to the companies that engage in surveillance pricing programs. As a result, not surprisingly, retailers' and digital platforms' surveillance pricing practices are very carefully managed to minimize the risk of public awareness. The complexity and opaqueness of the underlying practices, however, make casual public scrutiny very difficult. Consumers are regularly exposed to the practice and left feeling frustrated. The scale of the practice, however, is hard to detect from the consumer's perspective. What consumers are left with is the lingering notion that they are being taken advantage of and that they are paying more for goods than what would seem fair.

2. How surveillance pricing fails the standards for "competitive markets" and "fair pricing" that at we learned in high school economics

We were all taught in high school that the price of a good moves with market conditions such as supply, demand, time of day, and how much inventory is left. That is Economics 101. Airline seats get more expensive as a flight fills for the holidays. Hotel rooms cost more in peak holiday season. When a retailer has excess inventory, prices drop to move the inventory. A fundamental economics assumption of how markets work is that prices respond to the market, and at any given moment the price of a product for sale is the same for everyone looking at the product. You may not love the price, but it is visible, and it applies to all shoppers equally. On a rainy day in Manhattan, the price of umbrellas sold at busy street corners goes up.

Of course, this fundamental Economics 101 pricing dynamic depends on a set of assumptions. We assume that pricing is transparent; that everyone knows and understands the pricing. We assume that the market is functioning fairly, with all participants sharing equal information. We assume that the market is competitive; that it is not an oligopoly.

The model of competitive pricing taught in introductory economics descends from Adam Smith's account in his legendary 1776 book, *The Wealth of Nations,* but was formalized only in the late nineteenth and early twentieth centuries as "perfect competition." That

---

[23] Consumer Reports. "Consumer Reports Investigation Reveals Uber and Lyft AI-Driven Pricing Tactics Lead to Significantly Different Prices." Press release, June 16, 2026.

[24] "How Uber and Lyft Use Artificial Intelligence to Price Rides." *Consumer Reports*, June 2026. Accessed July 31, 2026.

model assumes several key requirements for competitive and fair pricing – consumers are treated universally the same, no single seller dominates the market, there is full transparency around pricing and costs, there is no collusion between sellers, and competition stops sellers from overcharging consumers. For the relatively simple markets of that era, more than 250 years ago, the principles were well-served. Smith's assumptions became core tenants of the Economics 101 principles that we all learned.[25]

How digital markets work, where surveillance pricing tactics are active, do not pass the tests for "competitive market" and "fair pricing" that we learned in high school economics. Consumers are not treated universally the same in how they are marketed and priced. Significant concentration of sellers exists across major retail product categories, and the digital platforms play a dominant role in matching consumers to products. Pricing is increasingly not transparent. Understanding whether you are being overcharged has become a tedious and complex process. The Economics 101 world that we were all taught in high school increasingly fails to exist.

Of all the changes since Adam Smith wrote about in the Wealth of Nations, the most profound difference in commerce today may be that in digital commerce, the consumer is not an anonymous consumer. Quite the opposite, in the digital marketplace, the consumer is usually very actively analyzed, studied, extremely well understood by the retailer.

What has replaced that anonymity in the digital marketplace is a complex system of consumer data collection, modeling, and messaging. In today's world, anonymity has been replaced by "hyper-personalization". Today, the price is built around you. The seller uses signals and data about your behavior and circumstances to estimate how much you personally are likely to pay, then shows you a price shaped by that estimate. You get tagged as someone who will pay more for reasons that have nothing to do with the seller's costs. Maybe you are loyal to the brand, or you are in a hurry, or you are committed to a specific travel date, or you have an observed track record of not comparing prices. The price responds to you rather than to the market, and you almost never know it happened. That last part is what makes it so hard to argue about. With ordinary dynamic pricing, everyone can see the fare went up. With surveillance pricing, the number is set for an audience of one.

---

[25] Adam Smith, An Inquiry into the Nature and Causes of the Wealth of Nations (London: W. Strahan and T. Cadell, 1776), bk. I, chap. VII, "Of the Natural and Market Price of Commodities"; on the later construction of the perfect-competition model, see George J. Stigler, "Perfect Competition, Historically Contemplated," Journal of Political Economy 65, no. 1 (1957): 1–17.

III. Surveillance Pricing's Consumer Harm is Real; We Can Estimate the Impact

The Consumer Reports and Groundwork Collaborative study of Instacart put a concrete figure on the consumer harm. Their work estimated that the price variation from surveillance pricing could cost a US household of four roughly $1,200 a year, with some items prices as much as 23 percent higher and nearly three out of four tested items shown at different prices to different shoppers.[26] Surveillance pricing is not abstract; it delivers measurable consumer harm. It is easy to treat all of this as abstract, so it helps to put a number on it.

According to the Bureau of Labor Statistics' Consumer Expenditure Survey, the typical United States consumer unit spent approximately $78,535 in 2024. Roughly $33,000 of that fell in the categories where digital and algorithmic pricing now operates: food at home and away from home ($10,169), apparel and services ($2,001), household furnishings, supplies, and operation ($5,212), transportation ($13,318), entertainment ($3,609), and personal care ($978).[27] E-commerce accounted for 16.9 percent of total U.S. retail sales in the first quarter of 2026, up from 16.0% a year earlier. E-commerce penetration varies widely by consumer spending category, with the percent of e-commerce exceeding 30% in categories like apparel and 50% in electronics, but even grocery, which has been the slowest to adopt category, had reached roughly 19% of weekly grocery spending by December 2025.[28]

When shopping occurs online, the price a consumer sees is seen only by that consumer. The price is delivered to the consumer via their personal laptop or mobile phone.

Across the various spending categories, where consumers are making their purchases online, the surveillance pricing algorithms operate. Past research studies indicate price increases ranging from 5% to 20%, varying by industry category. No comprehensive review of surveillance pricing price adjustments, by spending category, has been conducted. For

[26] Katie Wells, Lindsay Owens, Angel Han, and Alan Smith, Same Cart, Different Price: Instacart's Price Experiments Cost Families at Checkout (Washington, DC: Groundwork Collaborative and Consumer Reports, December 9, 2025).

[27] U.S. Bureau of Labor Statistics, "Consumer Expenditures — 2024," news release USDL-25-2085, December 19, 2025, table A.

[28] U.S. Census Bureau, "Quarterly Retail E-Commerce Sales: 1st Quarter 2026," news release CB26-77, May 18, 2026; Brick Meets Click and Mercatus, "U.S. eGrocery Sales Surge 32% YoY to a Record $12.7 Billion in December 2025," January 2026,

the purposes of estimating the overall cost of algorithmic pricing to consumers, it is reasonable to assume a typical range of 3% to 12%.[29]

To estimate the total cost to consumers, we can go spending category by spending category. We separate the consumer spending by channel – online purchase from the in-store purchase volume. We can apply estimated price adjustment percentages at the by consumer spending category, by channel. This basic, but detailed modeling translates into the average US household carrying an average annual higher cost of approximately $1,800.

For the average US household, $1,800 in increased household spending is material. It is not a rounding error in a household's budget. The harm is not hypothetical. It comes directly from US consumers.

1. Unfettered access to consumer behavior data is the critical success factor for surveillance pricing

No consumer would willingly supply personal data to third parties to be used against them. If a consumer was informed that data was being collected on their shopping behavior and purchase patterns, to be used in pricing models that sought to maximize retailers' profitability, the consumers would most certainly say "no". Which is probably why consumers are never asked. Tremendous corporate energy goes into ensuring the "opt-in" and regulatory requirements for using consumer data in pricing is minimal. Pricing is rarely even explicitly mentioned in corporate data use disclosures. A formal requirement to secure consumer approval to collect and use consumer shopping behavior data in pricing models would kill the practice.

In parallel, no retailer would want to the obligation to disclose how their pricing adjustments are managed. In private, retailers will acknowledge the headline and reputational risk issues associated with consumer-specific dynamic pricing. Freed from the obligation of reporting how the pricing works, retailers are at liberty to manage their pricing however they deem most advantageous and appropriate. The only meaningful guardrails are the conscience of the data scientists and executive teams that develop and harness the algorithms that are fueled by the data unwittingly supplied by consumers.

The result is a pricing system that operates on consumer data without consent, sets prices without disclosure, and is accountable to no one from outside the firms that run it. Consumers cannot see it, cannot measure it, and cannot opt out of it. The combination of material economic effect, zero transparency, and no enforceable duty of care is what

---

[29] Jean-Pierre Dubé and Sanjog Misra, "Personalized Pricing and Consumer Welfare," Journal of Political Economy 131, no. 1 (2023): 131–89; Hannak et al., "Measuring Price Discrimination and Steering on E-commerce Web Sites"; Le Chen, Alan Mislove, and Christo Wilson, "An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace," in Proceedings of the 25th International Conference on World Wide Web (2016), 1339–49.

distinguishes surveillance pricing from ordinary commercial pricing, and it is why this is properly a matter for Congress rather than for the market to resolve on its own.

2. Surveillance pricing is not creating value; it is taking value away from consumers. We need solutions.

The public awareness of bad corporate conduct takes time to grow to enable legislative or regulatory action. The challenge of building public awareness, however, is particularly challenging when the problematic practices are both complex and generally opaque. Complexity and opaqueness are certainly two key characteristics of surveillance pricing.

In the context of growing public awareness of AI applications, there is an increasing realization (and uneasiness) that surveillance pricing is best understood as being more about "AI shifting value" and less about "AI creating value". Getting consumers to pay more for a product than they should does not create any new value; it just shifts value – from the consumer to the digital platform or retailer.

IV. Multiple Opportunities Exist for Progress in Returning to Competitive and Fair Product Pricing

While the ability to engage in surveillance pricing tactics has expanded with technology, we also now see with greater clarity the potential opportunities to limit and reduce its impact. A combination of consumer behavior adjustments, new technologies and tools, more consumer-oriented corporate practices, and appropriate regulation could all help. Each potential solution is considered below.

1. Given the current lack of transparency, consumers have only limited and tedious options to defend themselves

The collection, analysis, and operationalizing of consumer behavior data is pervasive in the digital marketplace. Consumers have only limited options to reduce their vulnerability to surveillance pricing. By understanding how the current consumer data collection and monitoring system works, and how it can be thwarted, consumers can take some steps to protect themselves. The whole surveillance pricing ecosystem runs on data-derived signals about individual consumers. The more confidently the system can identify you and read your history, the more precisely it can develop a willingness to pay signal with which to price you. To avoid falling vulnerable to surveillance pricing, you want to minimize the digital platforms' ability to collect data and develop behavioral profiles. New and signed-out shoppers, with little attached behavioral history, tend to get coarser and less targeted prices, because the system has to rely on what essentially become rough guesses. With this awareness, consumers can apply four principles to defend against being taken advantage of by surveillance pricing systems. While the systematic adoption of the four principles is tedious, it is far better than simply surrendering to opaque practices that you cannot today escape.

The first principle is to reduce the shopping behavior data that is collected by denying the system a clean read on you. When shopping and the product price matters, check it in more than one way, on a different mobile phone, a different laptop, a different browser, or in a private window where less of your history is attached, and compare what you see.

The second principle is to avoid broadcasting urgency and intent. The consumer behavior that surveillance pricing systems read most eagerly is the behavior that says you have already decided to make the purchase and you are in a hurry. Slowing down, comparing options, opening multiple product detail pages (PDPs), and not circling back to the same item all signal you to be an active and scrutinizing consumer. Let items sit in your cart for a few days to see if the price drops, or price the item on a second device or in a guest session before buying.

The third principle is to treat loyalty with healthy skepticism. Loyalty programs are sold as savings, and sometimes they deliver, but we need to recognize that a loyal customer who never examines prices or comparison shops is exactly the customer that a "revenue optimizing" pricing system is most comfortable charging more. Demonstrating a willingness to switch brands or products is itself a form of protection.

The fourth principle is to reduce the trail of consumer shopping behavior that you leave behind, by limiting tracking where you reasonably can and by not assuming that staying signed in always works in your favor. Reject non-essential cookies whenever a site gives you the option and understand that consent banners are usually built so that accepting takes one click and declining takes several. Clear the cookies you have already accepted regularly, to reset some of the behavioral profile attached to your browser. Most browsers now let you block third-party cookies by default, which takes the decision out of your hands and costs you very little in daily use. Consider using a "burner" email and only use it for shopping to limit digital platforms' profile-building efforts. Disable the advertising ID on your phone to thwart tracking. Or even shop in person for categories where it's practical (without using a loyalty card) to deny the collection of digital behavior data.

Consumers should not have to defend themselves from harm through the disciplined adoption of tedious online shopping practices. The available consumer steps can only help at the margin, and no individual can fully opt out. The information asymmetry and practical imbalance is real. The shopping consumer has a browser, very limited visibility into how online pricing works, and some patience. The key actors in the digital commerce ecosystem have a complex and opaque tech platform, a well-developed history of your shopping behaviors, a detailed profile and set of signals, and up-to-date real time knowledge of your product preferences and needs. Disciplined personal habits can make you a little less vulnerable, but the overall playing field, however, still remains very much tilted against the consumer.

2. The practitioners of data Science should be the first line of defense. We need a “Hippocratic Oath” for data scientists

Unfortunately, but importantly, in today’s technology and data science, there is no shared common definition or ethic regarding what is an appropriate use of data and what is not. Regulatory guidance is missing on the topic, and unlike many other professions and fields, there are not institutions that establish the conduct standards for data scientists and the users of data.

Contrast how we govern data scientists and their use of consumer data with how the medical profession governs its conduct. Medicine has a professional ethical tradition, expressed in the Hippocratic Oath and its modern successors and in the widely invoked principle of *primum non nocere* – first, do no harm. Similarly, pharmacists must pass licensing exams and adhere to strict integrity standards to operate in the best interests of their patients. In the financial industry, investment advisors are governed by Fiduciary Responsibility – one must always act in the best interests of one’s clients. In the Legal profession, there is formal professional testing (the Bar exam) and ongoing licensing requirements. In the various trades – electricians, plumbers, and commercial truck driving, there are formal licensing requirements and oversight review boards.

Data science has produced ethical codes like the ACM Code of Ethics and Professional Conduct, the American Statistical Association’s Ethical Guidelines for Statistical Practice, the NIST AI Risk Management Framework, and the OECD AI Principles. But unlike medicine, law, or the licensed trades, none is tied to licensure, none is enforceable against an individual practitioner, and the NIST framework is expressly voluntary. There is no review or oversight board that can bar a data scientist from practice.[30]

The absence of clear “conduct” standards around the management of data and application of data science and AI tools contributes to the immaturity of industry practices. Conduct issues in the application of data science have been well documented for decades. The critical need for judgement was first spotlighted in the 2012 story of the Target Stores marketing team inadvertently sharing the “forecasting” of a woman being pregnant to the woman’s father.[31] More recently, the 2018 story of Facebook and Cambridge Analytica is all about the undisciplined application of data to an inappropriate use case.[32] Over the past

[30] Association for Computing Machinery, ACM Code of Ethics and Professional Conduct, adopted June 22, 2018, American Statistical Association, Ethical Guidelines for Statistical Practice, rev. 2022; National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (January 2023); Organisation for Economic Co-operation and Development, OECD AI Principles, adopted May 2019, updated May 2024.

[31] Charles Duhigg, "How Companies Learn Your Secrets," New York Times Magazine, February 16, 2012.

[32] Carole Cadwalladr and Emma Graham-Harrison, "Revealed: 50 Million Facebook Profiles Harvested for Cambridge Analytica in Major Data Breach," Observer, March 17, 2018.

several years, there has been a wave of popular books – *An Ugly Truth*, *Your Face Belongs to Us*, *Careless People, and Empire of AI*[33] – that tell the story of corporate entities allowing data scientists to collect and lever consumer data without a disciplined process to review and evaluate the appropriateness of the specific use cases being pursued. The common conclusion is that we need frameworks and guardrails on the commercial use of data to ensure that consumers' interests are appropriately weighed.

With all the sensitive and critical data that can be collected, analyzed, and modeled to infer future consumer behavior – with potentially exceptional value being created – we cannot escape the need, indeed the responsibility, to apply oversight judgement in how this work is pursued. Judgement needs to be applied by the data scientist, the marketing manager, and the business executive to ensure the right data is collected, under the right consent and disclosure conditions, and then only used in the appropriate, "consented", consumer use cases.

We should think through both what would be an appropriate Hippocratic Oath for data scientists as well as how to institutionalize that "first, do no harm" mindset. Proposals for the development of a fiduciary standard for the conduct of data science have been made by some leading academics. Professor Jack Balkin of Yale Law School has explicitly that the collectors and analyzers of consumer data should operate with the same fiduciary responsibilities of doctors, lawyers, and accountants.[34] Professors Neil Richards of Washington University School of Law and Woodrow Hartzog of Boston University School of Law have proposed a "Duty of Loyalty" standard for data management.[35] These basic concepts need to be promoted and matured for broad adoption. The colleges and universities that develop and train our data professionals should play a key role. Today, most graduate programs in data science and data engineering may not even require a class on "responsible data use" or the relevant data privacy or digital commerce regulations. The professional associations that support ongoing data scientist education and development should step up. The technology companies that provide the platforms and tools used by data scientists should embrace and communicate this necessary consumer commitment. And finally, policymakers and regulators should support the process by both promoting awareness of the institutional programs as well as mandating participation with appropriate regulation.

---

[33] Karen Hao, Empire of AI: Dreams and Nightmares in Sam Altman's OpenAI (New York: Penguin Press, 2025); Sheera Frenkel and Cecilia Kang, An Ugly Truth: Inside Facebook's Battle for Domination (New York: Harper, 2021); Kashmir Hill, Your Face Belongs to Us: A Secretive Startup's Quest to End Privacy as We Know It (New York: Random House, 2023); Sarah Wynn-Williams, Careless People: A Cautionary Tale of Power, Greed, and Lost Idealism (New York: Flatiron Books, 2025).

[34] Jack M. Balkin, "Information Fiduciaries and the First Amendment," *UC Davis Law Review* 49, no. 4 (2016): 1183-1234.

[35] Neil M. Richards and Woodrow Hartzog, "A Duty of Loyalty for Privacy Law," *Washington University Law Review* 99, no. 4 (2021): 961-1021.

There is plenty that can be done to support and inspire consumer-focused thinking on the part of our data science and data technology professionals.

3. Tech and AI innovators can pursue innovations to develop consumer solutions

While the strong examples from recent history mostly spotlight the ambition of data and technology innovation efforts to monetize data, the same tech expertise and creativity could be applied to the pursuit of consumer-empowering and data privacy protecting solutions. Through consumers' use of tools such as Global Privacy Control (GPC), an opt-out signal built into browsers including Brave, Firefox, and DuckDuckGo, supports a consumer to broadcast a single refusal to have their data sold or shared by covered businesses, rather than having to negotiate that refusal site by site. It is treated as a legally valid opt-out request in California, Colorado, and Connecticut. It is equivalent to a consumer manually clicking "Do Not Sell or Share My Personal Information," and a covered business that receives the signal and keeps selling that consumer's data is violating state laws.

Venture capital investors principally back data monetization efforts, but there is an increasing number of new tech start-ups seeking to support consumers' interests in this space and specifically support the maintenance of a consumer's anonymity and data privacy. Recent efforts that have gained some traction with consumers and businesses, including DuckDuckGo's Duck.ai and Tim Berners-Lee's Charlie, both which seek to thwart and limit the collection of the consumer behavior data that surveillance pricing models depend on. A recent example is Permission Slip, which acts as an authorized agent on the consumer's behalf to send deletion and do-not-sell requests in bulk to companies that hold their data. Launched by Consumer Reports in 2022, Permission Slip facilitated nearly five million privacy requests across more than 400 companies. DeleteMe acquired the company from Consumer Reports in July 2026 and continues to operates its services.

Beyond protecting against a consumer's behavior being tracked, new innovations, like the new capabilities from Fair Patterns, are creating the ability to directly test and check for evidence of algorithmic-driven pricing outcomes. More consumer awareness of the practice and the harm will inevitably lead to more innovation in the space.

More innovation in the interest of consumers would be immensely helpful. Past consumer innovations like Volvo's development of the 3-point seatbelt, the development of unleaded gasoline, and the child-resistant safety cap are all good examples of industry innovating on behalf of consumers. We need to inspire a similar wave of consumer-focused innovation from the big players in digital commerce and data science.

4. Retailers and digital platforms should proactively strengthen their practices

The practicers of surveillance pricing could proactively take steps to limit its harmful impacts on consumers. The simplest step that retailers and digital platforms could take is to proactively adopt greater transparency in their data collection and use practices. Accurately informing consumers what behavior data is being collected, and how that data is used, would at least start to inform the marketplace regarding how widespread the practice of surveillance pricing is.

Besides increasing the transparency of their practices, retailers and digital platforms could boldly offer consumers the opportunity to opt-out of the practice of their consumer behavior data being collected and used in their pricing programs. Going further, retailers and digital platforms could proactively offer consumers the opportunity to review what personal behavior has been collected and used for marketing, messaging, and pricing. By earnestly increasing the transparency around their practices, retailers and digital platforms have the opportunity to gain consumer trust and confidence, gain customer loyalty, and re-set marketplace expectations.

5. Ultimately, policymakers and regulators should require more transparency and consumer control

Since the Groundwork Collaborative's investigation of Instacart's dynamic pricing practices that was exposed in December 2025, there has been a wave of legislative proposals at both the State and Federal levels. At the state level, lawmakers have introduced more than 40 bills across at least 24 states in 2026 — already outpacing all of 2025 — including California's AB 2564 and New York's Algorithmic Pricing Disclosure Act. Maryland and Connecticut have enacted the first state bans, and New York's legislature has passed a One Fair Price Act, now awaiting the governor's signature. (Note: this legislative movement predates the Instacart report — the FTC's 2024 study and New York's November 2025 disclosure law came earlier; the Groundwork Collaborative's Instacart report accelerated an existing wave.) As of August 1st, three states now have formal regulations on surveillance pricing and 24 states have proposed legislation under consideration. We are now also encouragingly seeing activity on the federal level. At least three relevant bills have now been introduced: the Stop AI Price Gouging and Wage Fixing Act of 2025 (H.R. 4640), the One Fair Price Act of 2025 (S. 3387), and the Stop Price Gouging in Grocery Stores Act of 2025.

While prohibitions on surveillance pricing would be worthy public policy, how to best operationalize and monitor such a policy is less clear. If consumers are not able to easily discern when the tactic is being used, it will be hard to explicitly prohibit the practice. In parallel, there are certainly scenarios where a benign application of traditional dynamic pricing would be acceptable.

Establishing a policy framework for what consumer data can be employed in real-time models to inform pricing versus when it should be prohibited is a necessary step. In parallel, the policy framework should identify and codify in what household spending categories, or consumer spending use cases, should the practice of surveillance pricing be prohibited. We may actually find that regulating the practice of surveillance pricing is prohibitively complex versus having clear rules of how providers can only use consumer behavior data for pricing modeling when they have the explicit opt-in consent of consumers. Strong transparency disclosures and the requirement of informed consumer consent for the use of their data would be an impactful solution.

V. The Best Solution is Transparency that Enables Consumer Control and Choice over How their Data is Collected and Used

Because the digital commerce ecosystem in the US has evolved to enable the collection, analysis, and application of consumer behavior data without any meaningful disclosures or constraints, we need to take regulatory steps to provide that very basic transparency and consumer protection. As mentioned earlier, the basic conditions that Adam Smith identified as being necessary for competitive markets to work have been lost in the evolution of the digital marketplace. For the good of both individual consumers as well as for the more vigorous performance of markets, we need a stronger approach that ensures clarity of what data is being collected and consumer choice regarding for what use cases the data can be used.

To reduce the consumer harm from surveillance pricing, relying on consumers' personal caution in dealing with complex and opaque system will not be effective. Personal caution is necessary, but it would be unfair to tell consumers that the only available approach to minimize personal harm in dealing with an opaque, data-driven, system is simply to shop more carefully.

The stronger answer is to implement focused regulation to change how consumer behavior data is allowed to be collected, mandate simple and clear disclosures about data collection and its use and establish use case-specific consumer consent requirements for collecting the data in the first place. In parallel, we should strictly define what AI-driven marketing and pricing systems are allowed to know and do.

If retailers' and digital platforms' access to the consumer-specific behavior data that feeds surveillance pricing is controlled by the consumer to begin with, consumers can protect their interests by denying access to the sensitive willingness to pay data. Regulation need not tediously define what pricing practices are acceptable, but rather, and much more effectively, define the need for affirmative consumer consent to use the data in the first place. Lacking the access to data to “price optimize”, the product price must go back to reflecting the product rather than the person.

Taking one more step towards the transparency required for competitive and fair pricing, regulation also could simply require retailers and digital platforms to inform consumers that their product pricing will not be derived from consumer behavior data and signals unless a consumer agrees to allow it. Through such simple and affirming “product labeling”, we could bring clarity to a healthy marketplace practice as well as establish clear consumer rights around how their behavioral data can be only used with their explicit permission. A consumer would only be subjected to the collection of their behavior data for use in pricing models if they had affirmatively “opted-in’ for such a program.

From a consumer perspective, there is a real cost from the broad use of consumer behavior-based surveillance pricing against the typical US household’s budget items. We estimate the annual cost per household at approximately $1,800. Prohibiting consumers’ behavior data from being used to inform pricing would return to US households the estimated $1,800 annually. This potential consumer savings is not abstract; consumers would see it the next time they shop for groceries online or buy an airline ticket.

By requiring transparency around the collection and use of consumer behavior data in pricing, along with consumer decision rights on whether their personal behavior data can be used for pricing optimization, we are re-establishing that retailers and digital platforms should win consumers’ business on price and quality, the way markets are supposed to work. Such a regulatory move would simplify an increasingly complex system, relieve consumers from the burden that surveillance pricing currently imposes, and help to ensure the effective function of competitive and fair markets as they were originally envisioned.

Appendix – Patent Analysis Details

Methodology: Patent filings on consumer-data collection, modeling and algorithmic pricing

Prepared	July 31, 2026
Primary source	Google Patents (patents.google.com), queried July 31, 2026 (US Eastern time)
Secondary validation	USPTO Patent Public Search Basic (ppubs.uspto.gov/basic/), spot checks queried July 31, 2026 — see 'Cross-checks' tab
What was counted	US pre-grant patent application publications (Google Patents filter status=APPLICATION, country=US), filing date Jan 1, 2015 - Dec 31, 2025. Counting published applications approximates 'patents filed': nearly all US filings publish 18 months after filing, and counting applications avoids double-counting a filing that later also appears as a granted patent.
Classification scope (CPC)	Cooperative Patent Classification subgroups under G06Q 30/02 ('Marketing; Price estimation or determination; Fundraising'): 30/0201 'Market modelling; Market analysis; Collecting market data'; 30/0202 'Market predictions or forecasting for commercial activities'; 30/0204 'Market segmentation' incl. 30/0205 '{based on location or geographical consideration}'; 30/0206 '{Price or cost determination based on market factors}'; 30/0207 'Discounts or incentives, e.g. coupons or rebates' (incl. child codes). Titles verified against the USPTO CPC scheme, uspto.gov/web/patents/classification/cpc/html/cpc-G06Q.html. CPC matching in Google Patents is hierarchical (verified: adding a child code to a parent-code query does not change the result count).
Company set	23 companies. Candidate list of 26 large digital-platform, retail, payments and technology companies chosen editorially; the final set is every candidate with at least one matching filing. Assignee names queried per corporate entity, incl. legacy names (Wal-Mart Stores / Walmart Apollo; Facebook / Meta Platforms; Maplebear Inc. = Instacart). Booking Holdings, Kroger, Live Nation/Ticketmaster returned <5 and were dropped from the candidate scan; Netflix (3) retained. Shown merged under the parent brand.

Deduplication	Every result page of every query was retrieved and publication numbers collected; totals are unique publication numbers (3,143). 289 documents matched more than one assignee-name query (e.g., both Wal-Mart and Walmart Apollo) and are counted once.
Headline figure	3,143 unique US patent application publications, filed 2015-2025, by the 23 companies, in the five CPC areas.
Key caveats	(1) 2024-2025 filings are undercounted because applications publish ~18 months after filing (counts retrieved July 2026 capture filings only through roughly January 2025). (2) A document with several CPC codes counts once in the total but once in each area on the 'By area' tab, so area counts sum to more than the total. (3) G06Q 30/0201 is broad: it includes consumer-behavior modeling that is not exclusively pricing-related (spot-check of titles confirms some general market-analytics filings). The article text should describe the cluster as 'consumer-data collection, modeling, segmentation, pricing and personalized-offer' patents, not solely 'pricing patents'. (4) Assignee matching is by name substring; subsidiaries with different names are not captured. (5) Filings under AI classes (e.g., G06N) that avoid G06Q marketing codes are not captured - the figure is conservative in that respect.

Application publications by CPC area (chart data)

CPC code	Official CPC title (USPTO scheme)	Plain-English label used in chart	Application publications
G06Q 30/0201	Market modelling; Market analysis; Collecting market data	Consumer-data collection & market modeling	2,281
G06Q 30/0207 (incl. children)	Discounts or incentives, e.g. coupons or rebates	Personalized discounts & incentives	949
G06Q 30/0202	Market predictions or forecasting for commercial activities	Demand prediction & forecasting	585

G06Q 30/0204–0205	Market segmentation; based on location or geographical consideration	Market segmentation & location-based profiling	495
G06Q 30/0206	Price or cost determination based on market factors	Price determination / dynamic pricing	220

A document with multiple CPC codes appears in each matching area; areas sum to 4,530 vs 3,143 unique documents.