February 2026 Issue

AI-BASED PRICING

TAIWAN GLOBAL AI REPORT

AI Report

Author:
Dr. CHEN Pei-yu (Sharon)

Editors:
Dr. CHEN Yen-Kuang
Dr. TUNG Chen-Yuan

TAIWAN
駐新加坡台北代表處
Taipei Representative Office in Singapore

FOREWORD

Artificial intelligence (AI) has become an indispensable part of modern life, powering systems from conversational agents to autonomous vehicles and complex financial infrastructures. While these innovations offer tremendous opportunities, they also introduce complex risks that extend beyond technical vulnerabilities, affecting national security, economic stability, and public safety and trust.

The Global Alliance for Taiwan Technology Diplomacy (GATTD) seeks to facilitate cooperation, helping governments, industry, and academia confront these challenges together. Our mission is to strengthen global and regional security through research, partnerships, talent development, and commercialization, while fostering economic growth through collaboration between Taiwan and other technology-driven economies. In partnership with, and under the leadership of the Taipei Representative Office in Singapore, we aim to share insights widely and connect audiences across Taiwan, Singapore, and beyond.

This report features an in-depth analysis of the increasing implementation of AI-based pricing across industries, and explores outcomes and solutions for policymaker when AI-based pricing crosses from maximization to exploitation.

We welcome your feedback and comments, which will help us improve and expand future reports in this series.

Dr. CHEN Yen-Kuang
CEO, Global Alliance for Taiwan Technology and Diplomacy (GATTD)

Dr. TUNG Chen-Yuan
Representative, Taipei Representative Office in Singapore

For feedback purposes, please email Dr. Chen at: ykchen@gmail.com

TAIWAN
駐新加坡台北代表處
Taipei Representative Office in Singapore

01

Optimization vs. Exploitation: The Fine Line in AI-Based Pricing

When youshop online, book a flight, or hail a ride, there's a good chancethe price you see isn't fixed—it's the product of dynamic pricing, a strategy in which prices change based on market conditions—such as demand, season, or inventory levels. Traditionally, these adjustments have been governed by human-defined rules or formulas—for instance, "if inventory drops to 50%, increase price by 10%." Increasingly, these adjustments are no longer governed by simple rules but by artificial intelligence (AI) algorithms that learn, adapt and optimize prices in real time, a practice known as AI-based pricing.

This shift marks more than a technical upgrade. AI-based pricing fundamentally changes how prices are set, how value is extracted, and how fairness is perceived. While such systems promise efficiency and revenue gains, they also raise profound ethical, legal, and governance challenges.

At the heart of the debate lies a critical question: where does legitimate optimization end and exploitation begin?

02

What Is AI-Based Pricing?

AI-based pricing refers to the application of artificial intelligence techniques—including machine learning, deep learning, and reinforcement learning—to optimize pricing decisions through data-driven prediction, pattern recognition, and adaptive learning.

Unlike traditional rule-based or econometric pricing models, AI-based pricing models don't just follow instructions. Instead, they ingest large volumes of structured and unstructured data (e.g., demand patterns, consumer behavior, market conditions, competitor actions), continuously experiment and learn how prices influence outcomes, and based on the information learned, they dynamically determine or recommend prices that maximize specified objectives such as revenue, profit, or market share.

Figure 1 : AI Based Pricing

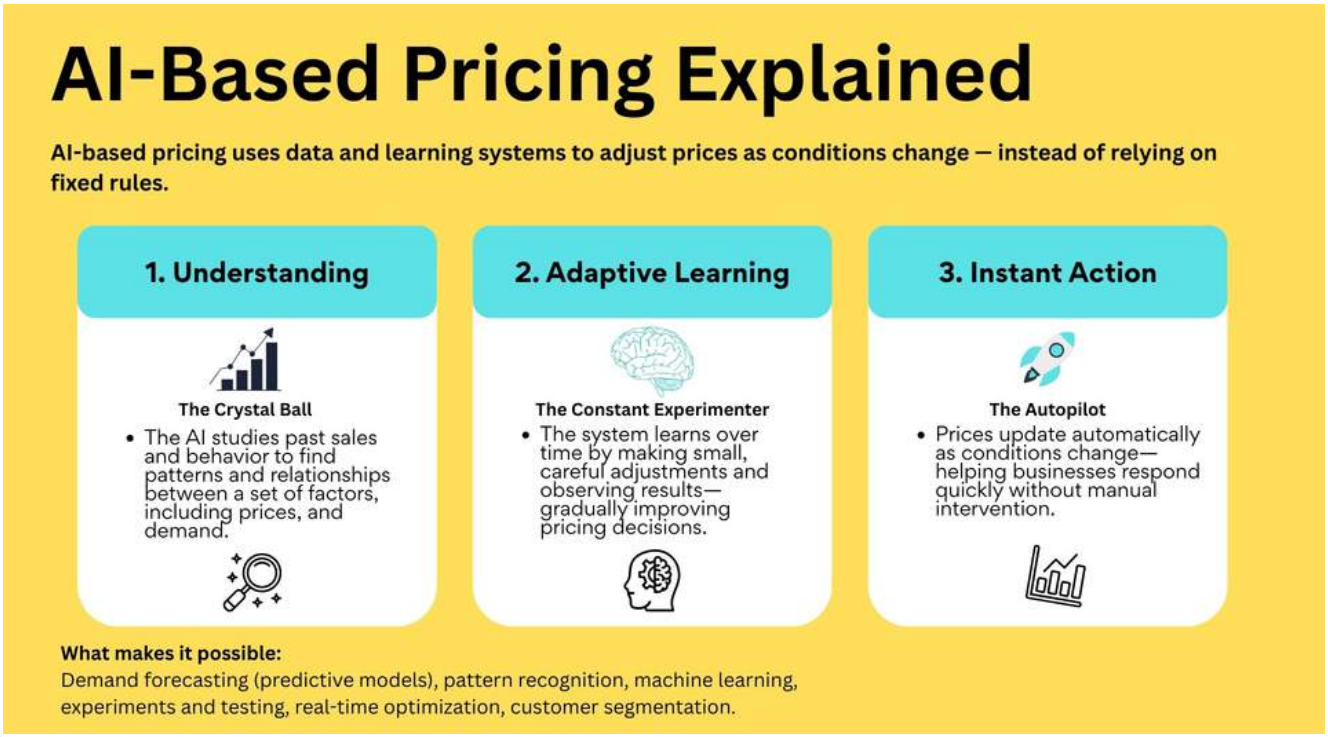

Figure 2: AI Based Pricing

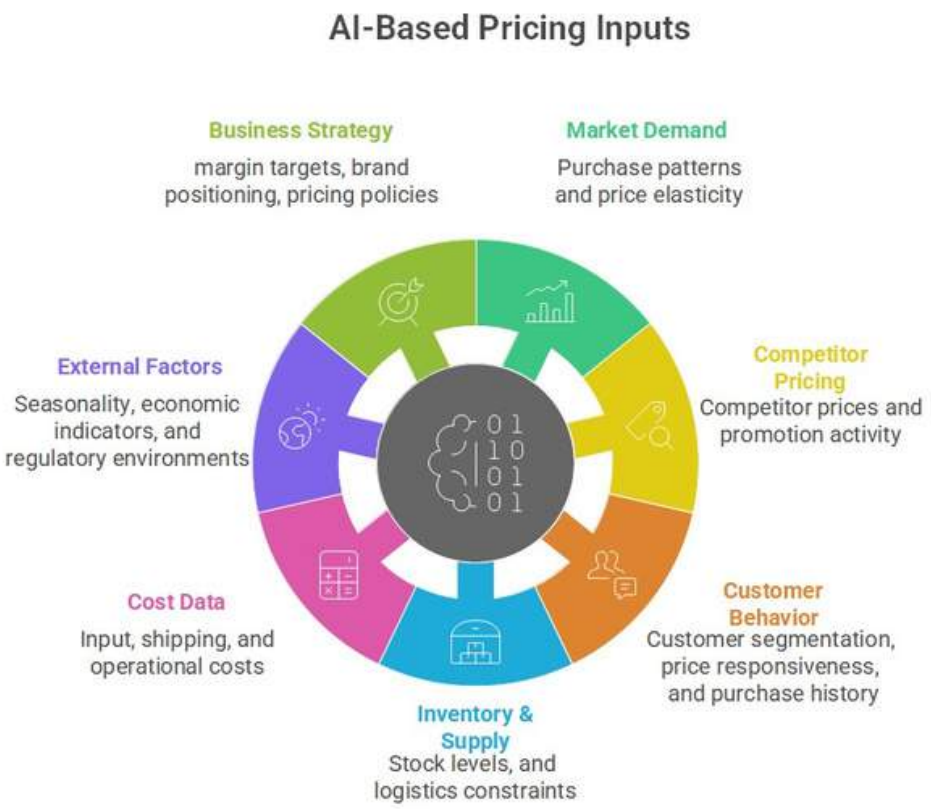

03

The Business Value

Dynamic pricing is not new. Industries such as airlines, hospitality, fuel retail, and ride-hailing have relied on it for decades. What has changed is the scale, speed, and granularity made possible by AI and real-time data. AI-enabled pricing practices are growing rapidly: research using Lightcast job posting data shows that the share of jobs related to AI pricing has increased more than tenfold over the past decade, suggesting widespread and expanding use of AI pricing algorithms across industries (Adams et. al., 2026).

In e-commerce and retail, AI-powered pricing engines continuously recalibrate prices across thousands of products using real-time signals such as competitor moves, promotions, inventory constraints, and customer behavior and past interactions. Recent implementations in the retail sector have demonstrated that AI-driven pricing solutions can increase revenue by up to 25% and improve profit margins by up to 16% (Sahoo, 2025). In energy and utilities, empirical research finds that dynamic tariff structures—particularly time-of-use and real-time pricing, can correct roughly 10% of mispricing relative to marginal costs in electricity markets and encourage demand shifting from peak to off-peak periods, reducing the need for costly capacity expansions and improving overall system efficiency. In logistics and manufacturing, algorithmic pricing allows firms to respond quickly to volatile input costs, supply disruptions, and shifting demand.

Taken together, these developments suggest that AI-based pricing is becoming a core infrastructure capability—integrating demand forecasting, marketing, supply chain management, and sales into a unified revenue optimization framework, enabling firms to capture significant competitive advantages through precision and responsiveness.

Yet greater sophistication also brings greater risk. Without appropriate oversight, personalization can cross a critical boundary—shifting from market-based differentiation to individualized exploitation.

Figure 3: Business value across sectors

04

The Economic Logic —and Its Limits

From an economic standpoint, dynamic and personalized pricing is often viewed as a mechanism for enhancing market efficiency.

Take a simplified example: imagine an airline has three seats to sell. One traveler is willing to pay $500, another $300, and a third only $100. If the airline sets one fixed price—say, $300—it sells two seats and earns $600. But if it tailors prices to each traveler's willingness to pay, it sells all three seats for a total of $900. This outcome increases both total revenue and overall market participation—an efficiency gain in economic terms.

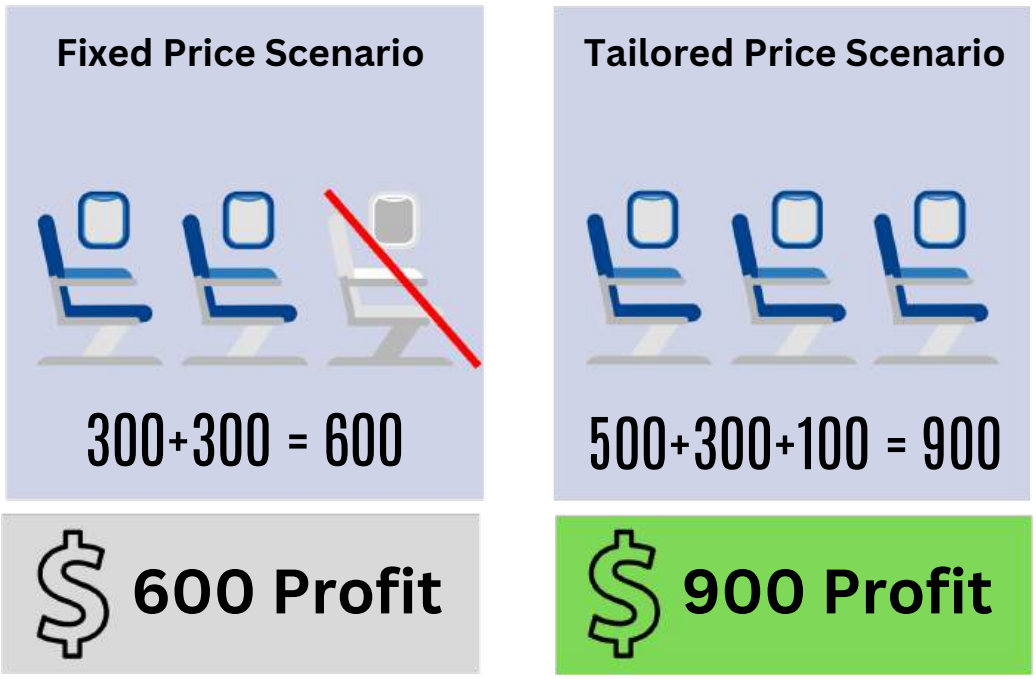

Yet, such personalized pricing in pursuit of efficiency also raises a deeper question: is this form of pricing fair? What does fairness mean in this context? Is it fair if all customers pay the same price, or is it fairer that those who place higher value on the product pay more?

Fairness in pricing can be interpreted through multiple lenses. Distributive fairness concerns whether outcomes are equitable—whether customers receive value commensurate with what they pay. Procedural fairness focuses on whether the process by which prices are determined is transparent, consistent, and free of bias. From a distributive perspective, differentiated prices may seem fair if each consumer voluntarily pays a price aligned with their perceived value. However, from a procedural standpoint, the same practice may appear unfair if consumers are unaware of the personalization mechanism or perceive it as manipulative. Thus, what appears economically efficient may still violate social expectations of fairness, challenging the legitimacy of algorithmic pricing in the eyes of consumers.

Moreover, the presumed efficiency gain depends on a big assumption: that the pricing mechanisms are implemented ethically. When personalization crosses into manipulation—charging higher prices solely because the algorithm detects a lack of alternatives or an urgent need—the practice shifts from optimization to exploitation. In such cases, the pursuit of efficiency can undermine fairness, eroding consumer trust and legitimacy in the marketplace.

05

When Data Knows Too Much

The power—and risk—of AI-based pricing lies in its dependence on data. Unlike traditional pricing models that rely primarily on aggregate variables such as supply, demand or inventory levels, AI-driven systems frequently draw on fine-grained behavioral and contextual digital signals, including a user's location, device type, browsing history, and time spent viewing a product, or patterns inferred from past interactions. Individually, such signals may appear innocuous; collectively, they enable algorithms to infer urgency, price sensitivity, or willingness to pay with remarkable precision, even when explicit personal identifiers are not used.

Recent findings from the U.S. Federal Trade Commission underscore that these practices are no longer theoretical. In its ongoing Section 6(b) study on what it terms "surveillance pricing," the FTC documents how firms and third-party intermediaries use extensive personal and behavioral data to tailor prices, discounts, or offers at the individual level. According to the FTC's research summaries, pricing tools may incorporate information such as browsing and purchase histories, location data, device identifiers, and real-time interaction signals—including cursor movements or dwell time—to determine how prices or promotions are presented to specific consumers. Crucially, these practices often occur without consumers' explicit awareness that such data is being used for pricing purposes. Importantly, not all AI-based pricing constitutes surveillance pricing; the latter refers specifically to individualized price setting based on extensive behavioral profiling rather than observable market conditions.

Not all data inputs, however, raise the same ethical or regulatory concerns. As illustrated in Figure 4 (Data Use Spectrum), AI- based pricing exists along a continuum. At one end are widely accepted, market-based inputs—such as inventory levels, historical demand patterns, time, or competitor prices—that are generally perceived as legitimate, transparent, and fair. Moving along the spectrum, firms increasingly rely on data that is legal but ethically sensitive, including coarse location data, device type, clickstream behavior, and inferred measures of price sensitivity. While permissible, these inputs may feel opaque or unsettling to consumers, particularly when they influence prices in ways that are difficult to observe or understand.

Further along the spectrum lie practices that the FTC characterizes as surveillance pricing, where personalization is driven by extensive behavioral profiling or cross-context data aggregation. Even when protected characteristics such as race, gender, or income are not explicitly used, proxies embedded in behavioral data can correlate with sensitive attributes, making discriminatory or exploitative outcomes difficult to detect and harder to contest. At the far end of the spectrum are data uses that are clearly prohibited, including the use of protected characteristics, precise geolocation without consent, or practices that violate privacy, consumer protection, or competition laws.

A defining feature of AI-based pricing systems is their opacity. Because these models operate as complex “black boxes,” consumers rarely know whether a higher price reflects genuine market conditions, unfavorable timing, or algorithmic inferences about their personal circumstances—the issue is further confounded when prices change frequently in small time segment, using extremely fine-grained data.

For example, Amazon has been observed changing prices for popular items as often as every 10 minutes to maintain its competitive edge (Sellbery, 2025). This scale is staggering: while a traditional brick-and-mortar retailer might make 50,000 price changes in a month, Amazon conducts more than 2.5 million price changes daily. The FTC has highlighted this opacity as a core consumer protection concern: when individuals cannot see or understand how prices are determined, they are unable to assess whether they are being treated fairly or to exercise meaningful choice. As data granularity increases, so too does the imbalance of information between firms and consumers—amplifying both predictive power and ethical risk.

Taken together, these developments illustrate how AI-based pricing can quietly shift from legitimate optimization toward individualized exploitation. What appears efficient from a revenue-maximization perspective may, in practice, erode procedural fairness and consumer trust. The FTC's findings make clear that the question is no longer whether data-driven personalization could cross this boundary, but whether existing governance frameworks are sufficient to ensure that it does not.

Figure 4: Data Use Spectrum

06

Will AI-Based Pricing Become Universal?

It is increasingly expected that organizations will adopt AI-based pricing as a standard practice to optimize costs, revenues and profits across operations. However, the individualized or personalized application of such pricing, where algorithms tailor prices to specific customers, is likely to remain sector-specific rather than universal. AI-based personalized pricing is most effective in markets where transactions are closely linked to individual identities or devices—such as airline tickets, ride-hailing services, and hotel bookings—where opportunities for direct price comparison or resale are limited. In these contexts, personalization can be implemented discreetly, and firms can capture value through individualized offers without immediate consumer resistance.

By contrast, in markets characterized by high price transparency, standardized products, or strong secondary markets—such as books, groceries, or consumer electronics—personalized pricing is less practical and more easily exposed. When consumers can easily compare prices or exchange information, differential pricing may be perceived as unfair, prompting consumer backlash and reputational harm. Moreover, in markets with active resale channels, price dispersion can generate arbitrage opportunities, where consumers with access to lower prices resell to others, thereby undermining profitability. As a result, the reach of AI-based pricing will likely expand selectively—thriving in data-rich, low-transparency environments, while remaining constrained in markets where openness and comparability protect consumers and discipline pricing behavior.

07

Implications for Managers and Policy Makers

Proponents contend that AI-based pricing can enhance resource allocation and market efficiency when implemented responsibly. By continuously adjusting prices in real time to reflect changes in costs, demand, and market conditions, AI pricing systems help firms allocate resources more effectively across products, regions, and customer segments. They also enable more precise alignment between supply and demand —reducing waste, optimizing capacity utilization, and improving operational responsiveness.

By tailoring prices to individual preferences and circumstances, such systems can serve more consumers, improve market matching, and offer more flexible or cost-effective bundles. Consumers with greater price sensitivity may benefit from lower prices, while firms can allocate capacity more efficiently. However, realizing these benefits depends critically on responsible design, governance, and oversight. AI-based pricing systems must not only optimize for profitability but also operate within ethical, legal, and societal boundaries that sustain long-term trust.

To ensure responsible use, firms should institutionalize a set of best practices grounded in transparency, accountability, compliance, and consumer protection (Figure 5).

“By tailoring prices to individual preferences and circumstances, such systems can serve more consumers, improve market matching, and offer more flexible or cost-effective bundles.”

- **Transparency:** To mitigate risks of exploitation, companies should clearly disclose how AI pricing systems use data-- what inputs are included and how they are used, and adopt explainable AI (XAI) frameworks that make pricing decisions interpretable, auditable, and aligned with corporate ethics. Guardrails are essential to prevent excessive volatility, discriminatory outcomes, or manipulative pricing patterns that could damage consumer trust. This also ensures procedural fairness.

Transparency

- **Accountability:** Regular, independent audits are essential to detect and mitigate algorithmic bias. Publicly releasing audit findings or summaries can strengthen credibility and demonstrate a commitment to ethical governance.

Accountability

- **Compliance:** Firms must ensure full adherence to consumer protection, data privacy, and anti-discrimination laws, supported by detailed documentation of pricing logic and decision-making processes.

Compliance

- **Recourse:** Providing consumers with mechanisms to question or appeal pricing outcomes enhances procedural fairness and reinforces perceptions of legitimacy.

Recourse

Without these safeguards, even well-intentioned AI-based pricing systems can easily become exploitative, leveraging data asymmetries to extract maximum consumer surplus rather than to create mutual value. Over time, such practices can erode trust, trigger consumer backlash, and invite regulatory intervention.

Figure 5: Management Action Plan

Pillar	Strategic Action	Long-term Benefit
Transparency	Implement Explainable AI Frameworks	Mitigates perceptions of manipulation and builds trust
Accountability	Independent Algorithmic Audits	Detects unintentional bias/discrimination
Compliance	Privacy-by-Design logic	Ensures adherence to GDPR/CCPA and emerging standards
Recourse	Human-in-the-loop overrides	Strengthens perceived fairness and marketplace legitimacy

For policy makers, the challenge lies in balancing innovation with protection. Overregulation may stifle technological progress and harm economic efficiency, but inaction risks allowing subtle forms of algorithmic discrimination to proliferate. Policymakers should therefore establish ethical and legal guardrails that promote fairness while preserving the efficiency gains of AI-based pricing. This may include mandatory algorithmic transparency standards, independent auditing requirements, restrictions on the use of protected traits or proxy data for personalization, and penalties for discriminatory or deceptive practices (Figure 6).

Equally important, policy makers must ensure that markets remain competitive. When consumers have viable alternatives —multiple sellers or platforms competing for their business— they can exert natural market discipline on firms, discouraging excessive price discrimination or exploitative behavior. Promoting competition through open market access, interoperability standards, and anti-collusion enforcement thus complements regulatory oversight by reinforcing market-based fairness.

Figure 6: Strategic Actions for Policy Makers

Policy Area	Objective & Description
Transparency Standards	Establish **Explainable AI (XAI)** requirements so pricing logic is interpretable and auditable.
Independent Auditing Requirements	Require **third-party bias audits** to detect and mitigate discriminatory outcomes in algorithmic pricing.
Data Guardrails & Compliance	Implement **strict restrictions** on using protected traits or sensitive proxies for individualized pricing.
Consumer Recourse Mechanisms	Ensure **clear, accessible mechanisms** for consumers to question, challenge, or appeal automated pricing outcomes.
Foster Market Competition	Maintain **open market access** and interoperability to ensure consumers have viable alternatives and prevent platform lock-in.

08

Broader Implications for AI Security and Governance

Concerns about data exploitation in AI-based pricing have rapidly moved to the forefront of global policy debates, reflecting a convergence of regulatory, consumer, and expert concerns.

From a regulatory standpoint, state authorities are increasingly focused on the competition and consumer protection implications of algorithmic pricing. In the United States, the Federal Trade Commission has begun scrutinizing whether shared data sources and third-party pricing intermediaries could produce “effectively collusive” outcomes or unfair competitive advantages, even in the absence of explicit coordination among firms. These investigations reflect growing concern that AI systems may amplify market power, obscure price formation, and weaken traditional antitrust safeguards.

Consumer advocates, meanwhile, emphasize the erosion of perceived price fairness. Reports documenting price variations of up to 23% for identical items within the same store have intensified public skepticism toward algorithmic pricing practices (Consumer report, 2025). This perceived arbitrariness undermines confidence in market integrity, particularly when consumers suspect that prices reflect personal profiling rather than transparent supply-and-demand dynamics. Such reactions align with broader survey evidence indicating widespread discomfort with dynamic and personalized pricing, reinforcing the idea that trust—not price alone—plays a central role in consumer welfare.

08

A central risk lies in the data pipelines that underpin these systems. If inadequately secured or governed, sensitive behavioral data becomes vulnerable to breaches, model inversion attacks, and adversarial manipulation. The leakage or misuse of such data can expose individuals' private habits, purchasing power, or inferred socioeconomic status—transforming personalization from a commercial tool into a significant privacy and security liability.

From a policy perspective, these dynamics underscore that AI security must extend beyond technical robustness to encompass data integrity, ethical governance, and social resilience. Secure AI systems are not only those that function accurately, but those that protect the confidentiality and integrity of their data while remaining transparent, accountable, and auditable throughout the algorithmic lifecycle. This includes safeguards to prevent AI- driven personalization from being weaponized for deceptive pricing practices or large-scale consumer fraud.

Ultimately, AI-based pricing serves as a microcosm of the broader AI governance challenge. Protecting against exploitation and misuse requires more than resilient algorithms; it demands secure, ethical, and transparent data practices embedded within strong institutional oversight. The long-term sustainability of AI adoption will depend on whether businesses and policymakers can ensure that the data powering these systems is collected, managed, and governed responsibly—in ways that protect individual rights, uphold public trust, and preserve the legitimacy of data-driven markets.

08

AI experts and governance observers focus on a different but related vulnerability: the widening oversight gap. While 82% of executives plan to deploy task-specific AI agents within the next three years, only 12% of organizations currently describe their AI governance structures as "mature." (Cisco, 2026). This imbalance creates material risk. The incremental gains from AI optimization can be rapidly offset by governance failures, including data breaches, regulatory penalties, or reputational damage. The financial stakes are significant: the average global cost of a data breach reached $4.44 million in 2025 (IBM, 2025), illustrating how fragile the business case for poorly governed AI systems can be.

At the same time, evidence suggests that effective governance can reverse this dynamic. Organizations that invest in robust Responsible AI (RAI) programs often benefit from a measurable "trust halo," including a 50% reduction in adverse AI incidents and higher enterprise valuations and revenues—up to 4% above peers that focus narrowly on baseline compliance (PWC, 2025). These outcomes suggest that trust and security are not merely defensive expenditures, but strategic assets.

Beyond pricing itself, AI-based pricing systems expose broader challenges for AI security and governance. The same data-driven mechanisms that enable granular personalization also create opportunities for deception, manipulation, and fraud. Highly detailed behavioral profiles— built from browsing activity, purchase histories, location data, and device information—can be repurposed by malicious actors to execute targeted scams, including fake discounts, phishing campaigns, or deceptive "limited-time" offers tailored to inferred urgency or financial vulnerability. In this sense, AI-based pricing does not merely pose fairness concerns; it expands the attack surface for AI-enabled exploitation.

09

Conclusion: Walking the Fine Line

AI-based pricing represents one of the most powerful yet contentious applications of artificial intelligence in modern commerce. At its best, it enhances efficiency, aligns prices with consumer preferences, and optimize resource allocation. At its worst, it risks transforming everyday transactions into mechanisms of digital surveillance, exploitation, and misuses in which algorithms silently test the boundaries of human behavior to maximize profit.

The future of AI-based pricing will depend on how effectively businesses, regulators, and society define and enforce the boundary between legitimate optimization and unethical manipulation. The technology itself is here to stay; its legitimacy, however, will hinge on whether it is governed in ways that protect data integrity, prevent exploitation and manipulation, and ensure transparent, accountable use.

Acknowledgements

Please feel free to reach out to the Economic Division of the Taipei Representative Office in Singapore should you have any enquiries or are seeking partnership opportunities of investment or collaboration in the field of semiconductors in Taiwan.

Cover Art generated by ChatGPT @Serena, Fang Ching Liu
Interior Design and Layout: Serena, Fang Ching Liu

Contact

Published:TaipeiRepresentative Office in Singapore
Address: 460 Alexandra Road, #23-00 mTower, Singapore 119963

Email: singapore@sa.moea.gov.tw
Telephone: +65 6500-0128

TAIWAN
駐新加坡台北代表處
Taipei Representative Office in Singapore