AI-Enabled Personalized Pricing: Efficiency, Rent Extraction, and the Conditions for a New Market Equilibrium

Zijing Xu

Suzhou North American High School, No. 268 Tian'edang Road, Wuzhong District, Suzhou, Jiangsu, China
Correspondence: ting.liu@worldshaper.cn

Abstract. Artificial intelligence and behavioral data make personalized pricing increasingly scalable, but its welfare effects depend on more than predictive accuracy. This conceptual paper develops an integrated framework linking economic surplus, information asymmetry, algorithmic decision-making, and consumer behavior. It also provides formal criteria that distinguish personalized pricing from ordinary dynamic pricing by data input, objective, decision unit, mechanism, and counterfactual test. Evidence from economic models, field and online experiments, marketplace data, and computational studies shows that personalization can expand participation and lower prices for some consumers, while also transferring surplus, reducing trust, and encouraging strategic data concealment. Reinforcement-learning agents may sustain supra-competitive prices through repeated feedback, although exploration rules and market structure can also produce competitive outcomes. Amazon is retained only as a historical illustration, not as proof of current individualized pricing. The paper concludes that efficiency is most plausible when personalization expands output, markets remain contestable, data use is transparent, and consumers retain meaningful choice.

Keywords: personalized pricing; dynamic pricing; information asymmetry; consumer welfare; reinforcement learning; algorithmic collusion

1 Introduction

Digital platforms can condition offers on purchase histories, browsing behavior, device or location data, and inferred customer profiles. These capabilities make behavior-based pricing more feasible, but they also create information asymmetry: the seller may understand how a profile affects an offer while the consumer cannot observe the classification rule or the relevant counterfactual price (Acquisti & Varian, 2005; Vasserman, 2023). The resulting welfare question is therefore not simply whether algorithms predict demand accurately, but who gains from that information and whether consumers can contest its use.

© The Author(s) 2026
B. P. Siba et al. (eds.), *Proceedings of the 2026 5th International Conference on Public Service, Economic Management and Sustainable Development (PESD 2026)*, Advances in Economics, Business and Management Research 411,
https://doi.org/10.2991/978-94-6239-825-2_23

Personalized pricing and dynamic pricing are often conflated. Dynamic pricing changes a product's price in response to market-wide states such as time, demand, inventory, capacity, or competitors. Personalized pricing uses individual- or profile-level information so that two consumers facing the same market state may receive different prices or discounts (Kannan & Kopalle, 2001; Chen et al., 2020). Hybrid systems can use both logics, making formal classification criteria necessary.

Existing evidence also shows that the effects are conditional. A randomized field study found that machine-learning personalization increased firm profit relative to uniform pricing, while slightly reducing aggregate consumer surplus; nevertheless, more than 60% of customers received a price below the optimized uniform price (Dubé & Misra, 2023). Online experiments report negative fairness and trust responses among both disadvantaged and price-favored consumers (Richards et al., 2016; Hufnagel et al., 2022), whereas marketplace data indicate that automated repricing does not uniformly increase prices (Hanspach et al., 2024).

This paper asks when AI-enabled personalized pricing creates efficiency gains and when it produces welfare loss. It contributes a four-part analytical framework, a formal boundary between pricing regimes, a synthesis of empirical evidence, and a mechanism-based account of how supervised learning, reinforcement learning, and multi-agent pricing can generate discrimination or supra-competitive outcomes.

2 Conceptual Framework and Methodology

2.1 Research Design

This study uses a theory-driven narrative review and conceptual synthesis of peer-reviewed journal articles. The evidence base includes economic models of price discrimination and competition, field and online experiments, consumer surveys, marketplace pricing data, and computational studies of learning agents. Theory identifies mechanisms; experiments estimate causal behavioral responses; field and marketplace studies assess observed outcomes; and simulations show what learning systems can generate under specified assumptions. Amazon is used only as a historical illustration, not as the principal empirical basis for the conclusions.

2.2 Formal Boundary Between Pricing Regimes

A pricing rule is classified by the marginal cause of the price difference. Holding product, time, inventory, capacity, and competitive conditions constant, a change caused by consumer identity or profile is personalized; a change applied to all comparable consumers because the market state changed is dynamic. Table 1 operationalizes this distinction.

Table 1. Formal criteria distinguishing dynamic from personalized pricing.

Criterion	Dynamic pricing	Personalized pricing
Primary inputs	Time, demand, inventory, capacity, competitors	Identity, history, behavior, device, location, inferred profile
Pricing objective	Respond to market state; clear capacity or inventory	Estimate individual/segment response and capture or expand demand
Decision unit	Product-time-market or broad segment	Consumer, account, session, or narrow profile
Mechanism	Same market state normally yields the same offer	Same market state can yield different offers because profiles differ
Counterfactual test	Change the market state while holding the consumer constant	Change the consumer/profile while holding the market state constant

2.3 Integrated Analytical Framework

The framework follows a causal chain from data and information asymmetry to prediction, pricing action, consumer response, and market feedback. Welfare is evaluated across total surplus, consumer surplus, autonomy, and competition. Efficiency is more likely when personalization expands output or participation and consumers can switch or withhold data; loss is more likely when market power, opaque profiling, or behavioral responses allow persistent surplus transfer, as shown in table 2.

Table 2. Integrated framework for evaluating efficiency and welfare loss.

Dimension	Core mechanism	Efficiency pathway	Welfare-loss pathway
Economics	Price discrimination, coverage, competition	Lower prices expand participation; rivalry passes gains to consumers	Market power converts information into rent and reduced consumer surplus
Information asymmetry	Firm observes profile and rule; consumer does not	Voluntary disclosure, clear notice, usable alternatives	Opaque or sensitive inference, bundled consent, no effective opt-out
Algorithmic decision	Prediction, experimentation, repeated updating	Better matching, inventory use, rapid competitive response	Proxy discrimination, feedback loops, tacit coordination
Consumer behavior	Fairness, trust, privacy, strategic response	Accepted discounts and informed participation	Lower trust, reduced purchase, greater search, data concealment

3 Evidence and Mechanisms

3.1 The Boundary in Hybrid Pricing Systems

Many systems are hybrid. A platform may first set a market-wide base price from demand and inventory, then personalize a discount using purchase history or predicted

willingness to pay. Granularity alone is insufficient: rapid prices can remain dynamic if all users in the same state see the same offer. The decisive question is whether personal information has a causal effect after market conditions are held constant. This criterion prevents ordinary repricing from being treated as evidence of individualized discrimination.

3.2 Empirical Evidence on Welfare and Consumer Response

Empirical results reject both an entirely efficient and an entirely exploitative account. Dubé and Misra's (2023) randomized pricing experiments show that personalization can raise profit and lower prices for many low-valuation customers, while slightly reducing aggregate consumer surplus relative to optimized uniform pricing. Theoretical work likewise shows that welfare depends on market coverage and competition: personalization may benefit consumers in highly covered competitive markets but harm them when coverage is low or only some firms possess data advantages (Rhodes & Zhou, 2024). Consumer control over disclosure can improve welfare, but the result depends on the available disclosure technology and market competitiveness (Vasserman, 2023).

Behavioral evidence identifies an additional channel. Incentive-compatible and online experiments find that interpersonal price differences reduce perceived fairness and purchase probability, although participation in price formation can mitigate the response (Richards et al., 2016). Three e-commerce experiments found adverse reactions among both price-disadvantaged and price-favored customers (Hufnagel et al., 2022), and two experiments showed that the type of personal data used shapes privacy and justice perceptions (Pizzi et al., 2022). Research on algorithmic dynamic pricing further reports lower retailer trust and longer price search, although adaptation can weaken the effect over time (Vomberg et al., 2025).

3.3 How Learning Algorithms Generate Coordination or Discrimination

Supervised pricing models map consumer and market features to predicted demand or willingness to pay, then choose the price that maximizes expected revenue. Discrimination emerges when identity-linked variables or proxies enter this state vector, so otherwise similar market situations produce different offers. Reinforcement learning adds experimentation: a price is an action, profit is the reward, and the model updates the expected value of actions after observing demand. When several agents learn simultaneously, each agent's environment includes rivals' prior prices and reactions.

In repeated interaction, Q-learning agents can discover a reward-punishment pattern: maintaining a high price yields sustained profit, whereas an undercut is followed by lower rival prices and a temporary price war. This can support supra-competitive outcomes without communication (Calvano et al., 2020). The mechanism is not inevitable. Better demand forecasts may increase the payoff from deviation and destabilize coordination (Miklós-Thal & Tucker, 2019), while insufficient exploration can create high prices that resemble collusion without a genuine punishment strategy; more thorough

learning may restore competition (Abada et al., 2024). Algorithm design, memory, exploration rate, update timing, and market concentration must therefore be examined before high prices are attributed to autonomous collusion.

3.4 Marketplace Evidence and the Amazon Illustration

Amazon's historical DVD experiment demonstrates the capability to test differentiated prices and the reputational sensitivity of such practices, but it cannot establish current individualized pricing (Huang et al., 2005). Broader marketplace evidence is more informative about automated repricing: data from approximately 2,800 products on a major European marketplace found heterogeneous effects and no general association between algorithm use and higher Buy Box prices (Hanspach et al., 2024). Even this evidence identifies automation, not necessarily personalization. The Amazon case is therefore illustrative, while causal claims rest on the broader empirical literature.

4 Discussion and Policy Implications

4.1 When Does Personalization Improve Welfare?

The integrated framework yields four conditions. Personalization is more likely to improve welfare when it expands output or access, competition remains strong enough to pass gains to consumers, data use is understandable and avoidable, and consumer responses do not destroy trust or induce costly defensive behavior. Welfare loss is more likely when market power permits persistent surplus capture, sensitive proxies classify vulnerable consumers, prices cannot be meaningfully compared, or repeated learning weakens competitive deviation. Predictive accuracy alone is therefore not an efficiency criterion.

4.2 Regulation Should Target Mechanisms

A blanket ban would eliminate legitimate discounts and capacity-management benefits. Governance should instead require notice when personal data materially affect an offer, a usable non-personalized option, limits on sensitive and proxy variables, audit trails linking inputs to decisions, and accessible challenge procedures. Competition authorities should examine the state variables, objectives, exploration rules, shared pricing vendors, and retaliation patterns of learning systems rather than treating every rapid price change as personalization or collusion.

4.3 Limitations

This synthesis does not estimate the prevalence or net welfare effect of personalized pricing across sectors. Field evidence remains limited, marketplace studies often identify automated rather than individualized pricing, and computational results depend on

model design. Future work should use matched-user audits, randomized disclosure designs, and multi-market field studies that hold market conditions constant while varying consumer profiles.

5 Conclusion

AI-enabled pricing should be evaluated as a linked economic, informational, algorithmic, and behavioral system. Dynamic pricing changes offers because market conditions change; personalized pricing changes them because consumer-level information changes the decision. Personalization can expand participation and lower prices for some consumers, but it can also transfer surplus, reduce trust, and weaken competition. Reinforcement learning explains how repeated feedback may generate discrimination or supra-competitive prices, while competing evidence shows that these outcomes depend on exploration and market structure. A defensible equilibrium therefore requires contestable markets, transparent data use, meaningful consumer choice, and auditable pricing mechanisms.

References

1. Abada, I., Lambin, X., & Tchakarov, N. (2024). Collusion by mistake: Does algorithmic sophistication drive supra-competitive profits? European Journal of Operational Research, 318(3), 927–953. https://doi.org/10.1016/j.ejor.2024.06.006.
2. Acquisti, A., & Varian, H. R. (2005). Conditioning prices on purchase history. Marketing Science, 24(3), 367–381. https://doi.org/10.1287/mksc.1040.0103.
3. Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. American Economic Review, 110(10), 3267–3297. https://doi.org/10.1257/aer.20190623.
4. Chen, Z., Choe, C., & Matsushima, N. (2020). Competitive personalized pricing. Management Science, 66(9), 4003–4023. https://doi.org/10.1287/mnsc.2019.3392.
5. Dubé, J.-P., & Misra, S. (2023). Personalized pricing and consumer welfare. Journal of Political Economy, 131(1), 131–189. https://doi.org/10.1086/720793.
6. Hanspach, P., Sapi, G., & Wieting, M. (2024). Algorithms in the marketplace: An empirical analysis of automated pricing in e-commerce. Information Economics and Policy, 69, 101111. https://doi.org/10.1016/j.infoecopol.2024.101111.
7. Huang, J.-H., Chang, C.-T., & Chen, C. Y.-H. (2005). Perceived fairness of pricing on the Internet. Journal of Economic Psychology, 26(3), 343–361. https://doi.org/10.1016/j.joep.2004.03.002.
8. Hufnagel, G., Schwaiger, M., & Weritz, L. (2022). Seeking the perfect price: Consumer responses to personalized price discrimination in e-commerce. Journal of Business Research, 143, 346–365. https://doi.org/10.1016/j.jbusres.2021.10.002.
9. Kannan, P. K., & Kopalle, P. K. (2001). Dynamic pricing on the Internet: Importance and implications for consumer behavior. International Journal of Electronic Commerce, 5(3), 63–83. https://doi.org/10.1080/10864415.2001.11044211.
10. Miklós-Thal, J., & Tucker, C. (2019). Collusion by algorithm: Does better demand prediction facilitate coordination between sellers? Management Science, 65(4), 1552–1561. https://doi.org/10.1287/mnsc.2019.3287.

11. Pizzi, G., Vannucci, V., Shukla, Y., & Aiello, G. (2022). Privacy concerns and justice perceptions with the disclosure of biometric versus behavioral data for personalized pricing: Tell me who you are, I'll tell you how much you pay. Journal of Business Research, 148, 420–432. https://doi.org/10.1016/j.jbusres.2022.04.072.
12. Rhodes, A., & Zhou, J. (2024). Personalized pricing and competition. American Economic Review, 114(7), 2141–2170. https://doi.org/10.1257/aer.20221524.
13. Richards, T. J., Liaukonyte, J., & Streletskaya, N. A. (2016). Personalized pricing and price fairness. International Journal of Industrial Organization, 44, 138–153. https://doi.org/10.1016/j.ijindorg.2015.11.004.
14. Vasserman, S. (2023). Voluntary disclosure and personalized pricing. The Review of Economic Studies, 90(2), 538–571. https://doi.org/10.1093/restud/rdac033.
15. Vomberg, A., Homburg, C., & Sarantopoulos, P. (2025). Algorithmic pricing: Effects on consumer trust and price search. International Journal of Research in Marketing, 42(4), 1166–1186. https://doi.org/10.1016/j.ijresmar.2024.10.006.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (http://creativecommons.org/licenses/by-nc/4.0/), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

