Economics & Business Management
ISSN Print: 3079-5214
ISSN Online: 3079-5222
https://doi.org/10.63313/EBM.9240

Short-Term Optimal Recommendations and Long-Run Consumer Welfare: A Dynamic Model of Preference Evolution

Jiakai Liu
Department of E-commerce, Jinan University Shenzhen Campus, Shenzhen, China

How to cite this paper: Liu, J. (2026). Short-term optimal recommendations and long-run consumer welfare: A dynamic model of preference evolution. Economics & Business Management, 7(3), 56–62. ISSN Print: 3079-5214, ISSN Online: 3079-5222. https://doi.org/10.63313/EBM.9240
Published: 2026-11-16

Copyright © 2026 by author(s) and Erytis Publishing Limited.
This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).
http://creativecommons.org/licenses/by/4.0/



Abstract

Recommendation algorithms optimize click-through and conversion rates. Both are short-run objectives. As precision keeps rising, what happens to consumer welfare in the long run? We answer with a dynamic model in which preferences evolve endogenously through consumption experience. Goods fall into a mature category inside current tastes and an exploration category outside them, and even a forward-looking consumer can only choose within the recommendation set. Higher precision mechanically crowds out exploratory exposure. When the option value of exploration exceeds the current matching gain, long-run welfare decreases with precision and converges to a welfare lock-in level under full precision. Platform profit and consumer surplus then diverge systematically, which justifies a cap on precision. Numerical results show a welfare loss of up to 11 percent where exploration is valuable; the result flips when preferences are nearly exogenous. Regulation should target not the algorithm itself, but its compression of the exploration space.

Keywords

Recommendation Algorithm; Endogenous Preferences; Long-Run Consumer Welfare; Filter Bubble; Algorithm Governance

I. Introduction

On any major e-commerce or content platform, the first screen a user sees is not what she searched for. It is a sequence the algorithm has arranged for her. The 57th Statistical Report on China's Internet Development puts the country's internet users at 1.125 billion by December 2025 [1]. Most of the information and product flows these users face each day are generated by recommendation algorithms trained on click-through and conversion targets. Precision keeps improving, and platforms have every incentive to improve it further: a better prediction means more clicks and more sales, this period.

The common judgment is that more precise recommendations must also benefit consumers. In a static framework the claim is almost self-evident, since better

prediction lowers search costs and raises match quality. Pariser's "filter bubble" [2] opened the challenge to this view, and the debate that followed linked information homogeneity to the fragmentation of public discourse, yet most of this literature remains descriptive: bubbles are observed and measured, but rarely written down as an equilibrium object that one can solve and compare in welfare terms. Does long-run consumer welfare rise or fall with precision, and where do platform profit and consumer surplus start to diverge? Answering these questions requires a model.

We construct an infinite-horizon dynamic model. The mature category yields high utility today and changes nothing tomorrow; the exploration category yields little today, but consuming it improves future preferences. The algorithm predicts the item most likely to be clicked this period, and higher precision raises the share of in-preference goods in the recommendation set. The consumer is forward-looking yet confined to that set, and venturing beyond it costs extra search effort. This tension, foresight against a constrained choice set, drives the model.

Three findings follow. First, precision crowds out exploratory exposure, a mechanical channel that, to our knowledge, has not been formally characterized; when the option value of exploration exceeds the current matching gain, long-run welfare falls monotonically with precision. Second, under full precision, welfare converges to a level sustained by the mature category alone, which we call welfare lock-in. The better the algorithm knows you, the more firmly it locks you in. Third, platform profit rises with precision while consumer surplus falls, so beyond an upper bound on precision the platform's gains come at the consumer's long-run expense.

The paper also fills three gaps. Filter-bubble research is largely descriptive [3]; we provide a solvable steady state and a welfare metric. The endogenous-preference literature [4], [5] rarely meets algorithmic environments; embedding preference evolution inside a recommendation-set constraint gives "preference shaping" a technical meaning. And our welfare decomposition narrows the object of regulation from "the algorithm" to an observable, auditable compression of the exploration space.

II. Related Literature

The paper speaks to three strands of literature.

The first is endogenous preferences. Stigler and Becker handled changing tastes inside a stable-preference framework through adjacent complementarity [4]; Bowles argued that institutions shape preferences themselves [5]. This strand supplies our load-bearing assumption, that preference evolution depends on actual consumption experience, which the experience-good literature has verified in cosmetics, apparel, and content categories. What is new here is the input side: the fuel of preference evolution, exploratory opportunity, is rationed by the

recommendation set rather than chosen freely.

The second strand covers consumer search. Intermediaries have incentives to divert search deliberately, as Hagiu and Jullien showed [6], but preferences in these models are exogenous and the distortion is static. We place the distortion in a dynamic preference environment, so a static distortion compounds into a long-run welfare loss.

The third strand concerns algorithm governance. Pariser supplied the concept [2]; Chaney, Stewart, and Engelhardt showed by simulation that algorithmic confounding raises consumption homogeneity and lowers utility [3]; Kleinberg, Mullainathan, and Raghavan proved that engagement-based optimization departs from users' true interests when preferences are inconsistent [7]. Unlike [7], our consumer is fully rational and the loss comes purely from the choice-set constraint; unlike [3], we derive analytical comparative statics with explicit sign conditions.

The literature has treated "preferences change" and "algorithms constrain" separately, and we put them in one dynamic model to ask what happens to long-run welfare when the force that limits preference evolution comes from an algorithm that knows you ever better.

III. The Model

A. Environment

Time is discrete and infinite, $t = 1, 2, \ldots$ There is one representative consumer and one platform. Consuming one unit of the mature good M yields current utility $v > 0$ and no long-run benefit. The exploration good E delivers current utility that depends on the preference state $\theta \in [0, 1]$:

$$u_E(\theta) = \theta w - k, \quad w > 0, k \geq 0. \tag{1}$$

Assumption 1 (preference order): $u_E(\theta_0) < v < u_E(1)$, with initial state $\theta_0 \in (0, 1)$. The exploration category is worse today but better once preferences have evolved. This order generates all the tension: M dominates in the short run, E may dominate in the long run.

B. Preference evolution

Assumption 2 (experience dependence): the state updates only when E is actually consumed,

$$\theta_{t+1} = \theta_t + \eta(1 - \theta_t), \quad \eta \in (0, 1). \tag{2}$$

Consuming M leaves θ unchanged. Preference improvement comes from real consumption experience, not the passage of time, and the increment shrinks as θ rises, keeping the state bounded. If $\eta = 0$ or preferences are exogenously fixed, exploration is worthless and our results flip, as shown below; this is where the model bears its weight. As for functional form, any kernel that is monotone, bounded, and has diminishing increments leaves the sign conditions of the propositions intact and changes only magnitudes.

C. The algorithm and the recommendation set

Each period the platform shows the consumer one recommendation slot. The algorithm predicts the item she is most likely to click now, given her history, and precision $\alpha \in [0, 1]$ is the probability that the slot displays exactly that item. As long as $u_E(\theta) < v$, current click propensity is higher for M, so

$$\Pr\{\text{set} = \{M\}\} = \alpha, \quad \Pr\{\text{set} = \{E\}\} = 1 - \alpha. \tag{3}$$

Raising precision mechanically lowers the share of exploration goods in the set. In the benchmark we treat α as exogenous, corresponding to an algorithmic upgrade across generations; the extension makes it endogenous. Consuming M accumulates click data and pushes precision up, while consuming E shifts tastes and depreciates past data:

$$\alpha_{t+1} = \alpha_t + \lambda(1 - \alpha_t) \text{ after consuming M}; \quad \alpha_{t+1} = (1 - \mu)\alpha_t \text{ after consuming E}, \tag{4}$$

with $\lambda, \mu \in (0, 1)$. Equation (4) captures the feedback loop of "the more you use it, the sharper it gets."

D. The consumer

The consumer has discount factor $\delta \in (0, 1)$, full foresight, and knowledge of (2). Facing the slot, she acts as follows. Shown M, she may consume it (utility v), or pay a search cost $s > 0$ to explore E on her own (utility $u_E(\theta) - s$, and θ updates). Shown E, she may consume it (free exploration, θ updates), or abstain (utility 0, θ frozen). Assumption 3 (choice-set constraint): E outside the recommendation set is not free; active exploration costs s. This captures a world in which attention is intermediated by the platform.

Write $g(\theta) = \theta + \eta(1 - \theta)$. The Bellman equation is

$$V(\theta, \alpha) = \alpha \cdot \max\{ v + \delta V(\theta, \alpha'), \quad u_E(\theta) - s + \delta V(g(\theta), \alpha'') \} + (1 - \alpha) \cdot \max\{ u_E(\theta) + \delta V(g(\theta), \alpha''), \quad \delta V(\theta, \alpha) \}, \tag{5}$$

where α' and α'' follow (4). In the benchmark, $\alpha' = \alpha'' = \alpha$.

E. Two-period benchmark

A two-period version delivers closed-form results. The second-period value is

$$V_2(\theta) = \alpha v + (1 - \alpha)\tilde{u}(\theta), \quad \tilde{u}(\theta) \equiv \max\{u_E(\theta), 0\}. \tag{6}$$

Shown E in period one, exploring beats abstaining if and only if

$$u_E(\theta_0) + \delta V_2(g(\theta_0)) \geq \delta V_2(\theta_0). \tag{7}$$

The left side of (7) is the option value of exploration: accept a low or even negative payoff today for a better preference state tomorrow. With $\sigma \in \{0, 1\}$ denoting the exploration decision, two-period welfare is

$$W_2(\alpha) = \alpha(1 + \delta)v + (1 - \alpha)\sigma[u_E(\theta_0) + \delta V_2(g(\theta_0))]. \tag{8}$$

IV. Analysis

Proposition 1 (crowding out): the probability of a free exploratory exposure each

period is $1 - \alpha$, strictly decreasing in precision. When shown M, active exploration costs s on top. Exploration intensity never increases with α. Raising precision necessarily compresses the exploration space.

The proof is immediate from (3): the chance of a free exposure equals the chance that the algorithm is "wrong." Success and exploratory opportunity are two faces of the same event; every later result grows out of this fact.

Proposition 2 (welfare reversal): in the two-period model, given $\sigma = 1$,

$$dW_2/d\alpha = (1 + \delta)v - [u_E(\theta_0) + \delta V_2(g(\theta_0))] + (1 - \alpha)\delta[v - \tilde{u}(g(\theta_0))]. \quad (9)$$

The first term on the right is the matching gain from precision; the rest captures the loss from replacing the exploration path. Whenever the option value of exploration is large enough, namely

$$u_E(\theta_0) + \delta V_2(g(\theta_0)) - (1 - \alpha)\delta[v - \tilde{u}(g(\theta_0))] > (1 + \delta)v, \quad (10)$$

we have $dW_2/d\alpha < 0$. Condition (10) requires a large w, a large η, or a small s, all parameter regions where "exploration has a future." Equation (5) admits no closed form, but the same sign condition holds numerically for every parameterization we tested (Table I).

Proposition 3 (welfare lock-in): as $\alpha \to 1$, free exploratory exposure vanishes. If the option value of active exploration then fails to cover s, the consumer stops exploring, the preference state freezes at θ_0, and long-run welfare converges to $v/(1 - \delta)$. Under (4), continued M consumption feeds the data pipeline and pushes α_t higher still. Lock-in reinforces itself: the better the algorithm knows you, the more firmly it locks you in.

Admittedly, a fully rational, infinitely lived consumer could in principle wait for rare free exposures, so literal absorption holds strictly only when s is large or foresight is limited. We use "lock-in" in the welfare sense: high precision postpones preference upgrading until its discounted value is negligible, and regulators, who care about discounted welfare, cannot look away.

Proposition 4 (divergence and the precision cap): suppose platform profit $\pi(\alpha)$ is strictly increasing, since precision raises conversion and the price of attention, while consumer surplus $W(\alpha)$ is strictly decreasing under condition (10), so the platform's optimum $\alpha^P = 1$ and the social optimum $\alpha^S = \text{argmax}\{W(\alpha) + \omega\pi(\alpha)\}$, with $\omega \geq 0$ the regulator's weight on profit, diverge systematically. There exists $\tilde{\alpha} < 1$ such that beyond it, π still rises while W already falls. Relative to the socially optimal recommendation policy, the welfare gap decomposes into

$$\text{Gap}(\alpha) = [W(\alpha^S) - W(\alpha)]_\text{choice-set} + [W(\alpha^S) - W(\alpha)]_\text{preference}. \quad (11)$$

The first term is the direct loss from current mismatch; the second, the cumulative loss from suppressed preference evolution, dominates numerically, so a regulator who evaluates only current recommendation quality will systematically understate the welfare cost.

Numerical example. Set $\delta = 0.95$, $v = 1$, $s = 0.4$, $\lambda = 0.05$, $\mu = 0.10$, and solve (5) by value-function iteration with a 250-period truncation on a 41×41 grid with bilinear

interpolation. Table I reports W(α) at θ_0 under two parameter sets.

Table I. Long-Run Consumer Welfare W(α) (Discounted Utility Sum)

α	High exploration value (w=2.2, k=0.5, η=0.15, θ_0=0.2)	Low exploration value (w=1.6, k=0.9, η=0.06, θ_0=0.1)
0.0	24.86	0.54
0.2	24.31	7.74
0.4	23.75	12.01
0.6	23.20	15.20
0.8	22.65	17.79
1.0	22.10	20.00

[a]Note: computed by the authors in Python; see the text for the numerical setup.

In the high-value column, W falls from 24.86 at α = 0 to 22.10 at α = 1, a drop of 11.1 percent, still above the lock-in floor v/(1 − δ) = 20, toward which welfare drifts further once α evolves endogenously past the critical region of Proposition 3. The low-value column approximates exogenous preferences: matching gains dominate and W rises monotonically from 0.54 to 20.00. The result flips completely, which confirms the load-bearing role of Assumption 2.

Figure 1 plots both columns. On the high-value curve, pushing α toward 1 makes the consumer worse off; on the low-value curve, the same upgrade is a Pareto improvement. Precision, we argue, carries no welfare meaning by itself; the meaning depends on which curve an economy sits on.

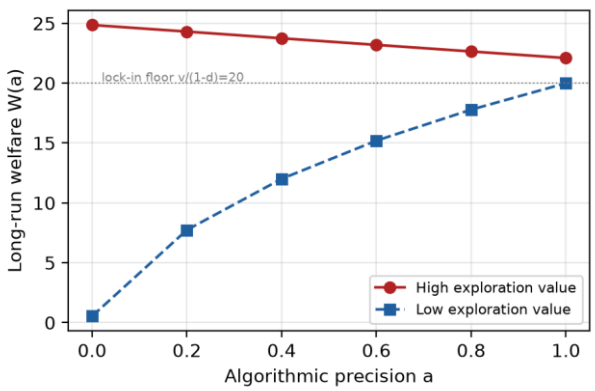


Fig 1. Long-run consumer welfare as a function of algorithmic precision. Data as in Table I.

V. Coclusion and Discussion

Placed in a dynamic environment with endogenous preference evolution, the static intuition that more precise recommendations are always better can fail systematically: higher precision crowds out exploratory exposure, less exposure stops preferences from updating, and frozen preferences lock the consumer inside the mature category the algorithm has drawn around her, with long-run welfare falling by up to 11 percent in our numerical example, a figure to be read with caution as an illustration rather than a measurement. Platform profit and consumer surplus diverge beyond a certain precision level, which gives regulation its footing:

the target is not the algorithm, but the algorithm's compression of the exploration space.

China's Algorithmic Recommendation Management Provisions, effective March 2022, already require platforms to offer users an opt-out and prohibit algorithmically induced addiction [8]. The direction is right, but an opt-out pushes the choice back to users, who may be poorly placed to weigh long-run costs. Our analysis points the other way, maintaining the exploration space directly: a mandatory exploration quota, which amounts to a lower bound on $1 - \alpha$; auditable diversity metrics in regulatory reporting; and data portability arrangements that weaken the lock-in built into (4). All three treat the exploration space as a public resource that needs maintenance.

One boundary deserves emphasis. Every counterintuitive result here hangs on Assumption 2; where preferences are close to exogenous, as in standardized durables and functional goods, the right column of Table I has already shown the reversal. We would argue that the policy implications are confined to preference-plastic categories and do not generalize into the claim that recommendation algorithms are harmful. Future work can proceed along three lines: platform competition, endogenous seller pricing, and more general preference kernels.

References

[1] China Internet Network Information Center, The 57th Statistical Report on China's Internet Development. Beijing: CNNIC, 2026.

[2] Pariser E. The filter bubble: What the Internet is hiding from you[M]. London: penguin UK, 2011.

[3] Chaney A J B, Stewart B M, Engelhardt B E. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility[C]//Proceedings of the 12th ACM conference on recommender systems. 2018: 224-232.

[4] Stigler G J, Becker G S. De gustibus non est disputandum[J]. The american economic review, 1977, 67(2): 76-90.

[5] Bowles S. Endogenous preferences: The cultural consequences of markets and other economic institutions[J]. Journal of economic literature, 1998, 36(1): 75-111.

[6] Hagiu A, Jullien B. Why do intermediaries divert search?[J]. The RAND Journal of Economics, 2011, 42(2): 337-362.

[7] Kleinberg J, Mullainathan S, Raghavan M. The challenge of understanding what users want: Inconsistent preferences and engagement optimization[J]. Management science, 2024, 70(9): 6336-6355.

[8] Cyberspace Administration of China, Ministry of Industry and Information Technology, Ministry of Public Security, and State Administration for Market Regulation, Algorithmic Recommendation Management Provisions for Internet Information Services, 2022.