Journal of Advances in Social Sciences

2026, Volume 2, Issue 3 | ISSN 3105-899X | eISSN 3105-9007

AI-Accelerated Opportunity Churn: Migration Frictions, Replenishment, and Governance in Digital Marketplaces

Mainak Mondal*

Courant Institute of Mathematical Sciences, New York University, New York, USA

Received
2026-5-21
Accepted
2026-6-23
Published
2026-6-28

Corresponding Author
Mainak Mondal*
Email: mm10845@nyu.edu

DOI: https://doi.org/10.65192/epk03p70

Abstract

Generative AI lowers listing and migration frictions in digital marketplaces. We ask whether faster seller churn improves market health or depletes opportunities faster. In an agent-based model, AI affects only launch cost C_L and migration time T_M; sellers chase transient product windows, consumers update trust, new opportunities arrive at rate λ_P , and the platform moderates at intensity φ. Higher AI assistance increases migrations but not bankruptcy survival when the catalog is finite and unreplenished. Collective churn can yield zombie markets: solvent sellers, near-zero active products, and collapsed assortment diversity. Replenishment and governance jointly bound outcomes: low λ_P drives exhaustion; only a narrow band of moderate φ keeps trust, participation, and diversity positive simultaneously. Platform health must be judged on activity and diversity together with trust, not on trust or seller survival alone.

Keywords
Agent-based model; Computational social science; Digital marketplaces; Generative AI; Opportunity churn; Platform governance

1. Introduction

Online marketplaces run on short-lived product windows. Sellers launch, advertise, learn, and exit; windows close through consumer learning, copycat entry, rising ad costs, and platform enforcement. Generative AI shortens listing and relaunch time. The system question is whether a platform can absorb many sellers migrating faster at once, not whether one seller gains an edge.

The ambiguity arises because AI changes both individual and collective time scales. For an isolated seller, faster listing and relaunching can extend survival by reducing time spent in failed niches. At the marketplace level, however, widespread adoption accelerates entry into the same visible opportunities, raising saturation and shortening the very windows sellers are trying to exploit. The effect of AI is therefore not signed ex ante: it depends on whether new product opportunities replenish faster than sellers collectively deplete them. This is not a trivial productivity story—faster sellers are not the same as a healthier market.

We term the process AI-accelerated opportunity churn. Lower CL and TM raise migration frequency, which in turn raises saturation and awareness and shortens the effective window $T_B =$

$min(T_L, T_C, T_P)$. A seller benefits only if $T_M < T_B$ and replenishment keeps new windows available. When λP is too low, the market can turn zombie: solvency remains while active products, seller participation, and diversity collapse—a state that trust-only metrics miss.

We contribute three results: AI as auditable friction reduction on (C_L, T_M), not a productivity shock; emergent zombie markets under weak replenishment; and legitimacy– vitality separation under governance (stable trust with collapsing participation).

The paper proceeds as follows. Section 2 situates the model. Section 3 states three propositions and the governance viability index. Section 4 describes the agent-based implementation. Section 5 lists experiments; Section 6 presents figures; Section 7 interprets mechanisms, governance, and limitations; Section 8 states implications.

2. Related Mechanisms and Modeling Approach

The model draws on platform economics (Rochet and Tirole, 2003; Armstrong, 2006; Hagiu and Wright, 2015; Zhu and Liu, 2019), trust under quality uncertainty (Akerlof, 1970; Shapiro, 1983; Valle et al., 2020), herding and saturation (Banerjee, 1992; Anderson, 2004), and agent-based regime analysis (Grimm et al., 2010; Railsback and Grimm, 2019). Empirical motivation is specific to marketplace churn rather than generic AI productivity: generative tools shorten seller research and listing tasks (Felten et al., 2023; Mollick et al., 2023; Grewal et al., 2021); social-commerce and trend cycles create burst demand and rapid niche closure (Liang and Turban, 2011; Anderson, 2004); copycat entry on visible winners is documented in platform retail (Zhu and Liu, 2019; Banerjee, 1992); fulfillment and quality uncertainty shape refunds and trust (Valle et al., 2020; Einav et al., 2014); and content moderation can remove harmful listings while imposing false-positive costs on legitimate sellers (Gillespie, 2018; Klonick, 2018; Gorwa et al., 2020). Generative AI is modeled only as lower C_L and T_M (OECD, 2023)—not as autonomous strategy, synthetic reviews, or ad optimization. Parameters are chosen for pattern-oriented reproduction (Grimm et al., 2010), not calibration to a named platform: the goal is comparative statics on migration, depletion, and governance, not GMV forecasting.

3. Theory

The theory is organized around competition among time scales. Let T_L denote the consumer-learning time scale, T_C the imitation or copycat-entry time scale, T_P the platform-penalty time scale, and $T_M(a_i)$ the migration time of seller i after exiting a product. The effective business-window duration is

$$T_B \approx min(T_L, T_C, T_P) \tag{1}$$

A seller survives a window transition when migration is completed before the next opportunity closes:

$$T_M(a_i) < T_B \tag{2}$$

3.1 Proposition 1: Individual Migration Advantage

AI assistance lowers launch cost and migration time:

$$C_L(a_i) = C_{L0}(1 - \gamma_L a_i) \tag{3}$$

$$T_M(a_i) = max\{1, bT_{M0}m_i(1 - \gamma_M a_i)c\} \quad (4)$$

Proposition 1. Holding other parameters fixed, higher ai weakly expands the set of states satisfying Equation (2). This is an individual friction advantage. It does not imply monotone bankruptcy survival when repeated relaunch costs and failed windows accumulate

3.2 Proposition 2: Collective Congestion and Window Shortening

When AI adoption becomes widespread, copycat entry can accelerate and the imitation time scale shortens:

$$T_B(\bar{a}) \approx \min\left(T_L, \frac{T_C}{1 + \gamma_C \bar{a}}, T_P\right) \quad (5)$$

where $\bar{a}$ is the adoption share.

Proposition 2. The same technology that helps one seller migrate can shorten windows for all sellers. Individual migration advantages can therefore be offset at the aggregate level by faster congestion and churn.

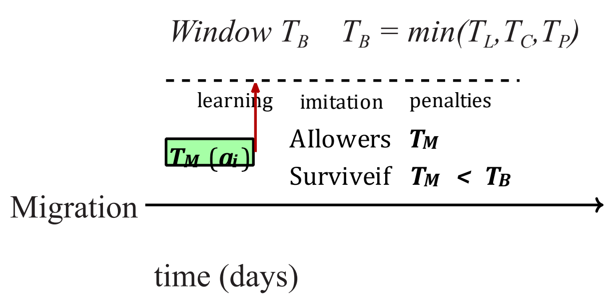

Figure 1. Time-scale competition. The dashed line is the product window T_B. The seller must complete migration (green bar, length T_M) before T_B closes; AI shortens T_M but does not extend T_B.

3.3 Proposition 3: Replenishment Threshold and Zombie Markets

Let new product opportunities arrive as

$$N_t^{new} \sim Poisson(\lambda_P) \quad (6)$$

If λ_P is too small, migration burns through the finite catalog faster than new windows appear. Define a zombie market as a late-window state with positive bankruptcy survival but near-zero seller activity, market-alive fraction, and Shannon diversity.

Proposition 3. The simulations imply an operational replenishment threshold λ^*_P under the reported calibration: for $\lambda_P < \lambda^*_P$, exhaustion and zombie regimes dominate long horizons; for $\lambda_P > \lambda^*_P$ and sufficiently low governance suppression, active diversity can persist. This is an emergent bifurcation in the ABM, not a closed-form existence proof.

Governance introduces a second constraint. Weak moderation allows bad listings and duplicate signals to depress effective trust; strong moderation can preserve apparent trust while suppressing participation through removals and false positives. For diagnostics only—not as a welfare objective—we summarize joint health with

$$V(\varphi) = (M_{alive} D_{Shannon} W_{eff} S_{active,norm})1/4 \quad (7)$$

where $S_{active,norm} = min(1,S_{active}/0.35)$. The index $V(\varphi)$ is a diagnostic for joint viability: it is high only when market-alive fraction, diversity, effective trust, and seller activity are all positive. It is not a social-welfare function and should not be read as proof that every margin peaks at moderate φ. Figure 1 visualizes the migration inequality against T_B; Figure 2 diagrams the migration–depletion loop (R_1) and the replenishment balance (B_1).

4. Model

The model is an agent-based marketplace ($|I| = 100$ sellers, $|J| = 10{,}000$ consumers, $T_{max} = 365$ days) with sellers, consumers, products, listings, and a platform. Sellers hold

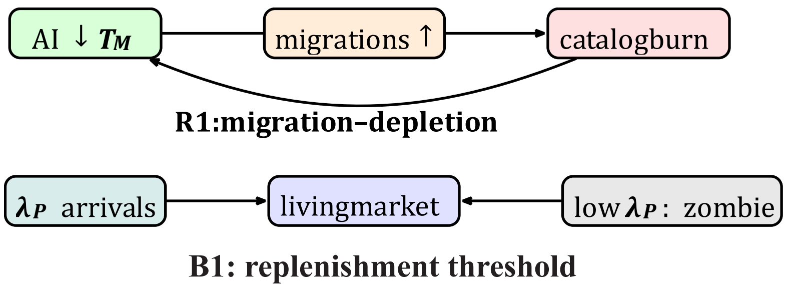

Figure 2. Core feedbacks. R1 (top): friction reduction accelerates migration and depletion. B1 (bottom): replenishment sustains vitality; insufficient λ_P yields zombie markets.

Table 1. Mechanism inventory and tests

Mechanism	Motivating fact	Experiment	Ablation
AI frictions (γ_L,γ_M)	Faster listing setup	2a	AI off
Awareness / saturation	Trend learning	2c	—
Trust dynamics	Lemons, refunds	2c, 4	Trust frozen
Imitation	Copycats	3b	$p_{imit} = 0$
Replenishment λ_P	New trends	4	$\lambda_P = 0$
Governance ecology	Moderation FP	4	Gov. off

cash K_i, ad budget B_i, quality q_i, risk r_i, exit threshold θ_i, AI level $a_i \in [0,1]$, and a migration timer. Products carry demand D_p, novelty N_p, awareness A_p, seller count n_p, saturation S_p, quality Q_p, refunds F_p, copyability χ_p, and unit cost c_p. Consumers carry trust w_j; the platform holds $W(t)$, moderation φ, transparency η, and ad-competition κ.

Each day: costs, seller launch/exit, visibility, purchases, awareness and trust updates, imitation, moderation, migration, Poisson arrivals at λ_P , and metrics—purchases precede trust updates. Sellers launch when expected ROI exceeds θi and cash covers $C_L(a_i)$; they exit on low ROI, refunds, high awareness, or ad spikes, then burn cash over $T_M(a_i)$ inactive days. Consumers choose listings via a logistic score over visibility, novelty, reviews, price, delay, refunds, and $W_{eff} = min(\bar{w}, W)$. Imitation adds copycats to visible winners; governance removes listings and generates false positives that feed Equation (11).

4.1 Seller Migration and Launch

Equations (3)–(4) define the AI friction channel. A seller launches a product when expected return exceeds θ_i and cash covers $C_L(a_i)$. A seller exits when return falls below threshold, refund signals exceed a tolerance, awareness exceeds an exit threshold, or advertising costs spike. Exiting sellers pay migration burn during $T_M(a_i)$ inactive days.

Bankruptcy occurs when $K_i \leq K_{floor}$.

Table 2. Seller archetype presets

Archetype	a_i	q_i	θ_i	m_i
Manual	0.0	0.75	−0.05	1.2
AI-assisted	0.5	0.65	−0.08	1.0
AI-churn	1.0	0.45	−0.15	0.7
Brand-builder	0.2	0.90	−0.02	1.5

4.2 Product-window Closure

Novelty decays each step. Product awareness evolves as

$$A_p(t+1) = clip[A_p(t) + \lambda_E E_p(t) + \lambda_N n_p(t) + \lambda_R F_p(t), 0, 1] \quad (8)$$

where $E_p(t)$ is exposure. Saturation is

$$S_p(t) = \frac{n_p(t)}{n_p(t) + k_s} \quad (9)$$

The product-level conversion rate is

$$CR_p(t) = CR_0 N_p(t)[1 - A_p(t)][1 - S_p(t)] W_{eff}(t)^{\rho} \quad (10)$$

with $W_{eff}(t) = min(\bar{w}(t), W(t))$. Purchases use a logistic choice model over visibility, novelty, reviews, price, delay, refunds, and trust.

4.3 Trust and Governance

Consumer trust averages to $\bar{w}(t) = |J|^{-1} P_j\ w_j(t)$. Platform trust updates from good purchases, bad purchases, refunds, duplicate signals, successful removals, and false positives:

$$W(t+1) = clip\ \Big(W(t) + \delta_G G_t - \delta_B B_t \sigma_B - \delta_R R_t^{\text{sig}} \sigma_R - \delta_D D_t^{\text{sig}} \sigma_D + \delta_\phi \phi n_t^{\text{rem}} - \delta_{\text{fp}} \phi^{\eta_{\text{fp}}} n_t^{\text{fp}} + b_{\text{gov}}(t), 0, 1\Big). \quad (11)$$

Trust collapse is recorded when $W(t) < W^*$ for at least τ consecutive steps. Moderation affects removal, false positives, imitation dampening, and cleanup. New products arrive according to Equation (6), with a contested-product cap to prevent inactive catalog size from being mistaken for economic vitality.

4.4 Regime Metrics

The late window is the final 185 days. The main responses are bankruptcy survival, selleractive fraction, seller-profitable fraction, active products, market-alive fraction, Shannon diversity, churn, late trust, trust-collapse probability, and the composite $V(\varphi)$. Regime labels are assigned hierarchically: exhaustion if market-alive is below threshold; trust degradation if trust falls or collapse persists; boom–bust if churn is high and diversity low; stable diversity if diversity, trust, and market-alive are jointly high; migration survival if seller activity is high with elevated churn; otherwise mixed/transitional.

5. Simulation Experiments

Experiments test Propositions 1–3 and Table 1. Unless noted, simulations use $|I| = 100$ sellers,

$|J|$ = 10,000 consumers, and T_{max} = 365 daily steps. Mixed markets use 25% of each archetype. Adoption fraction $\bar{a}$ sets the share of AI-churn sellers at initialization.

Table 3. Main experiment manifest

ID	Factors	λ_P	Role
2a	AI level a_i	0	Proposition 1: migration vs. survival
2c	Mixed archetypes	0	Trust–vitality decoupling
3b	T_{M0}, $\bar{a}$	0	Zombie markets without replenishment
4	λ_P , φ, $\bar{a}$	facets	Replenishment and governance
Abl.	Module toggles	0, 2	Mechanism necessity

Experiment 2a holds quality and exit rules fixed while varying $a_i \in \{0,0.5,1.0\}$. Experiment 2c runs four archetypes jointly with shared trust and awareness fields. Experiment 3b sweeps migration time and adoption with $\lambda_P = 0$. Experiment 4 crosses $\lambda_P \in \{0.5,2,5\}$, $\varphi \in \{0.05,...,0.7\}$, and $\bar{a} \in \{0,0.1,...,1.0\}$ with five seeds per cell. Ablation runs disable replenishment, trust dynamics, imitation, governance ecology, or AI frictions at representative cells.

Validation is pattern-oriented (Grimm et al., 2010). Simulator: market churn abm (Python 3.10+, YAML config; run manifest lists seeds and grids).

6. Results

Figures 3–7 and Table 4 test Propositions 1–3 respectively: individual migration vs. survival (Fig. 3), system depletion and trust–vitality decoupling (Figs. 4–5), replenishment and governance regimes (Figs. 6–7), and mechanism necessity (Table 4). All runs use five independent seeds where grids are reported; late-window metrics average days $t \in [181,365]$.

6.1 AI Raises Migration but not Survival

Figure 3 tests Proposition 1 (Experiment 2a, $\lambda_P = 0$, $a_i \in \{0,0.5,1.0\}$). The left panel shows bankruptcy survival: manual sellers ($a_i = 0$) remain highest through the late window; AI-churn sellers exit first. The center panel shows mean cash drawn down by repeated launches; the right panel shows cumulative migrations monotone in ai. Faster migration does not extend solvency when the catalog is finite.

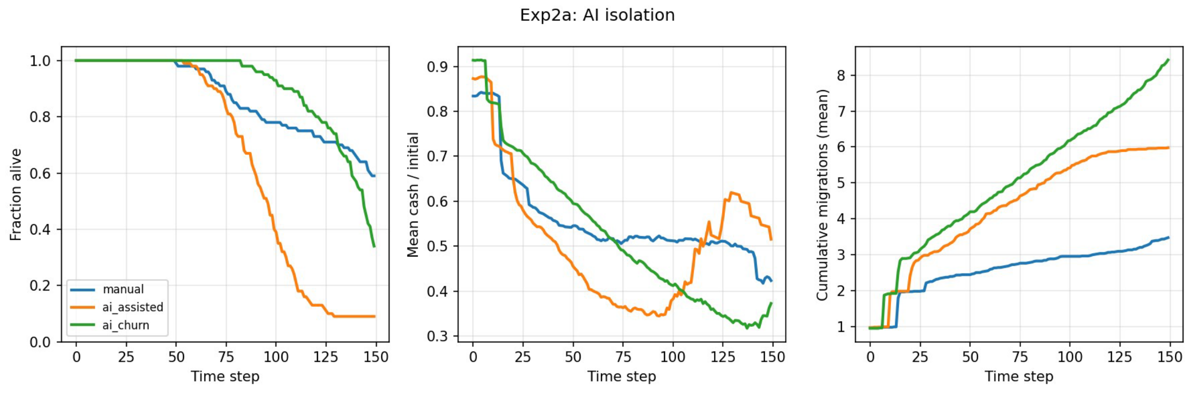

Figure 3. Experiment 2a ($\lambda P = 0$). Left: bankruptcy survival by AI level. Center: mean cash (normalized). Right: cumulative migrations. Manual sellers survive longest; AI-churn sellers migrate most.

6.2 Trust and Vitality Can Decouple

Experiment 2c mixes four archetypes (Table 2) with shared trust and awareness. Figure 4 plots aggregate trust, churn, active products, and Shannon diversity. Trust falls after early bad purchases and partially recovers toward 0.7; churn rises; active products and diversity go to zero by the late window. A platform monitoring only *W(t)* would miss catalog exhaustion.

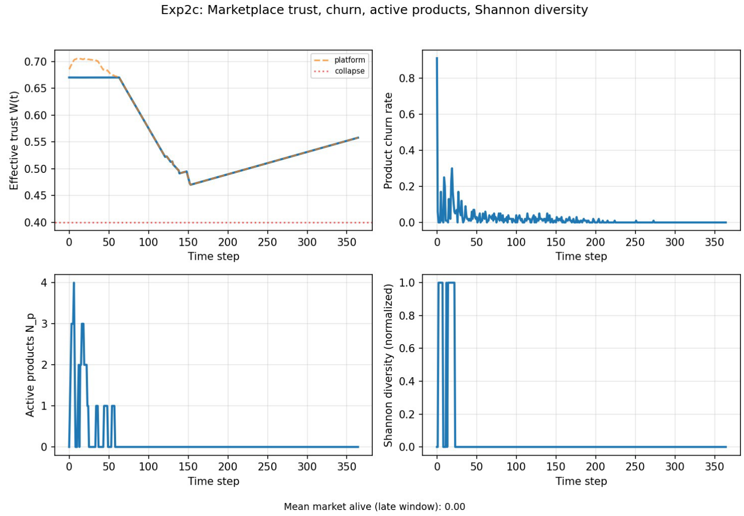

Figure 4. Experiment 2c. Platform trust *W(t)*, churn, active products, and Shannon diversity. Trust stabilizes; assortment metrics collapse.

6.3 Without Replenishment, Zombie Markets Emerge

Figure 5 (Experiment 3b, $\lambda_P = 0$) sweeps base migration time T_{M0} and adoption $\bar{a}$. The top row reports late-window active-product counts, M_{alive}, and Shannon diversity; the bottom row reports trust-collapse probability. The pattern is not that every cell is exactly zero—some cells show small nonzero active products or market-alive values—but diversity remains effectively absent across the grid while vitality metrics stay low. That is the zombie signature of Proposition 3: some sellers or listings can persist, but the market no longer supports a diverse, active assortment. Trust-collapse risk is highest where migration and adoption are both high.

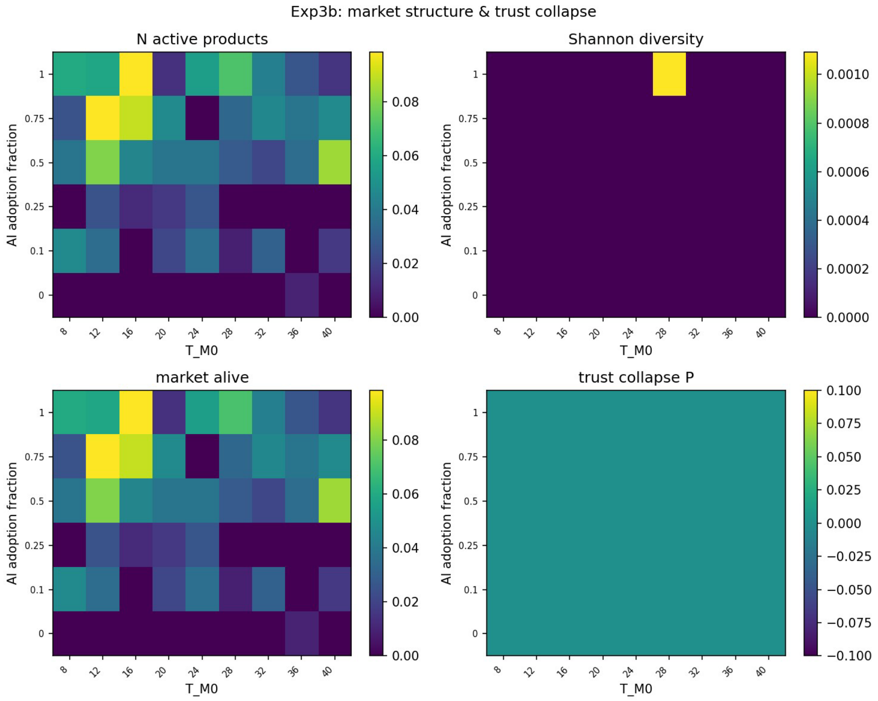

Figure 5. Experiment 3b ($\lambda_P = 0$). Top row: late active products, M_{alive}, and Shannon diversity over (T_{M0}, $\bar{a}$). Bottom: trust-collapse probability. Low diversity with depressed vitality across most cells.

6.4 Replenishment Separates Exhaustion from Stable Diversity

Figure 6 (Experiment 4) classifies late-window regimes over ($\bar{a}$,φ) for λ_P = 0.5,2,5. Low replenishment is exhaustion-dominated; at $\lambda_P \in \{2{,}5\}$, stable diversity appears at low-to-moderate φ.

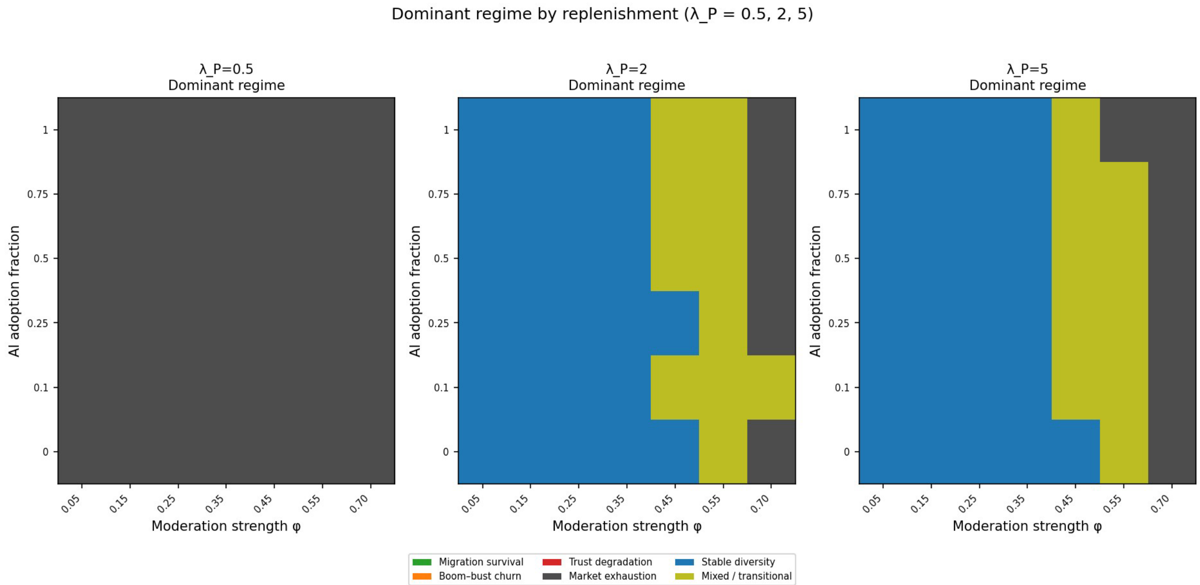

Figure 6. Experiment 4. Dominant regime over adoption and φ for λ_P = 0.5,2,5 (left to right).

6.5 Governance Separates Trust from Participation

Figure 7 slices φ at $\lambda_P = 2$, $\bar{a} = 0.5$. Effective trust collapses for φ . 0.15. The diagnostic $V(\varphi)$ is positive only near $\varphi \in [0.15, 0.25]$, where all margins in Equation (7) are jointly positive. For φ & 0.25, seller-active fraction falls sharply while $W_{eff} \approx 0.7$—trust can look healthy while participation thins.

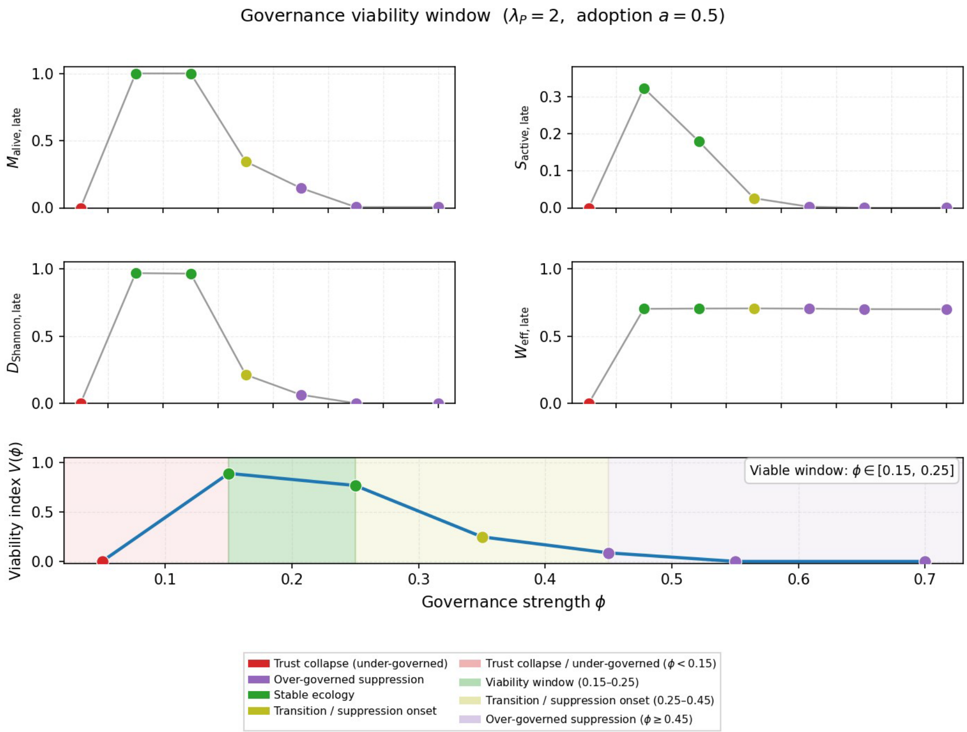

Figure 7. Governance diagnostics ($\lambda_P = 2$, $\bar{a} = 0.5$). Top: Malive, seller-active fraction, diversity, W_{eff} vs. φ. Bottom: diagnostic $V(\varphi)$, not welfare. Joint-viable band $\varphi \approx 0.15$–0.25; participation falls at higher φ while trust stays elevated.

6.6 Ablations

Table 4 toggles one module at a representative viable cell (*$\lambda_P = 2, \varphi = 0.15, \bar{a} = 0.5$*; five-seed means). Turning off replenishment ($\lambda_P = 0$) drives seller-active and marketalive fractions to zero—without new opportunities, the mechanism inventory confirms that B1 is necessary for any living market. Freezing trust or disabling imitation leaves the baseline nearly unchanged at this cell; that does not imply trust and imitation are globally irrelevant (Experiment 2c shows trust–vitality decoupling). It indicates that, in this representative viable cell, replenishment and governance dominate the local outcome. Disabling governance ecology lowers seller activity and market-alive fraction, showing that moderation rules affect participation even inside the viable band. Turning AI frictions off reduces seller activity without collapsing the market, consistent with Proposition 1:

AI raises migration intensity more than it guarantees economic vitality.

Table 4. Ablation snapshot at $\lambda_p = 2, \varphi = 0.15, \bar{a} = 0.5$ (five-seed means)

Ablation	S_{active}	M_{alive}
Baseline	0.31	1.00
No replenishment	0.00	0.00
Trust frozen	0.32	1.00
No imitation	0.31	1.00
Governance off	0.11	0.81
AI off	0.24	1.00

7. Discussion

7.1 Interpreting the Mechanism Chain

The substantive lesson is that seller-level productivity can be a misleading proxy for system health. Faster relaunching increases motion in the marketplace, but motion is not vitality. Vitality requires an active assortment, positive seller participation, and replenishment of profitable opportunities—the combination stressed in loop B1 (Figure 2). Generative AI compresses (C_L, T_M) and therefore amplifies migration frequency (loop R1). When λ_P is insufficient, collective migration depletes the shared catalog before new windows arrive. Trust and governance matter most once a living opportunity field exists; Experiment 2c already shows that trust can recover while assortment metrics die.

The results are intentionally conservative about AI. The model does not grant better product judgment, cheaper ads, or automated reputation management. Under that restriction, platforms that celebrate listing velocity alone risk missing zombie macro states in which M_{alive} and Shannon diversity trend toward zero while aggregate trust looks acceptable.

7.2 Governance: Trust–participation Separation

Figure 7 is the governance centerpiece. At $\lambda_P = 2$ and $\bar{a} = 0.5$, three φ-zones appear. Under-governed (φ . 0.15): effective trust and the diagnostic V (φ) collapse—weak cleanup lets bad signals dominate. Joint-viable ($\varphi \approx 0.15$–0.25): M_{alive}, seller-active fraction, diversity, and W_{eff} are simultaneously positive; $V(\varphi)$ is positive here but is not a welfare optimum—only a joint-feasibility check. Over-governed (φ & 0.45): $V(\varphi) \to 0$ from participation suppression while $W_{eff,late}$ can remain near 0.7.

That pattern is legitimacy–vitality separation. A platform can pass a trust survey while sellers stop listing. Policy and analytics should therefore treat aggregate trust as necessary but not sufficient.

7.3 Marketplace Health Metrics and Welfare Proxies

We do not integrate full social welfare, but the recorded metrics map to a stylized decomposition (using W for welfare to avoid confusion with platform trust $W(t)$):

$$W \approx CS + \Pi_S + \Pi_P - H - C_M \tag{12}$$

where CS is consumer surplus (variety and match quality), Π_S seller profit, Π_P platform revenue, H harm from refunds and delays, and C_M moderation cost including false positives. Market-alive fraction and Shannon diversity proxy variety in CS; seller-active and seller-profitable fractions proxy Π_S; $W(t)$ proxies perceived quality and platform legitimacy; removals and false positives proxy C_M. The diagnostic $V(\varphi)$ in Equation (7) is separate from W: it flags joint feasibility of margins, not dollar welfare. Zombie markets waste seller-side surplus (Π_S low) while consumers face thin choice sets; over-governance can raise C_M and collapse Π_S with stable $W(t)$.

A practical monitoring bundle is $(S_{active}, M_{alive}, D_{Shannon}, churn, W(t))$ reported on the same dashboard and time horizon. None of these alone detects trust–participation separation; together they flag exhaustion, zombie states, and participation cliffs.

7.4 Relation to Platform Practice

Platforms rolling out generative listing tools face a systems trade-off, not a seller-level productivity question. If λ_P is endogenously low because trend discovery is slow or enforcement freezes categories, faster migration only speeds depletion. If λ_P is adequate, moderation must avoid both trust collapse at low φ and seller exit at high φ. The regime map suggests that stable diversity is a region in $(\lambda_P, \varphi, \bar{a})$ space, not a default outcome of AI adoption.

For research, the model offers testable comparative statics: raising $\bar{a}$ should shift R1 upward; raising λ_P should expand B1; raising φ should move platforms along the slice in Figure 7. Empirical work could calibrate λ_P from category inflow rates and φ from enforcement intensity, then test whether participation and trust diverge as predicted.

7.5 Limitations and Extensions

The study is pattern-oriented (Grimm et al., 2010), not fitted to a named marketplace. AI is limited to friction reduction; autonomous agents, synthetic reviews, and targeting algorithms are excluded. Suppliers are not strategic—unit costs and delays are product attributes. Regime labels and the diagnostic $V(\varphi)$ depend on threshold choices in the classifier; sensitivity to those thresholds is reported in the replication package. Experiment 4 grids should be re-run after parameter changes because viable φ bands can shift with adoption and replenishment.

Extensions include endogenous λ_P tied to platform category policy, multi-platform competition, richer AI channels, and welfare integrals over simulated purchase paths. CoMSES Net archival of market churn abm, seeds, and figure scripts is planned for peer review.

8. Conclusion

Generative AI compresses the frictions of listing and relaunching in digital marketplaces. This paper asks what that compression does to *marketplace* health when many sellers react at once. An agent-based model with explicit time scales, Poisson replenishment, and platform governance yields three durable findings.

First, AI is a migration-frequency amplifier, not a survival guarantee. In a finite, unreplenished catalog, higher assistance increases relaunches while manual sellers can survive longer by migrating less and spending less on failed windows. Second, collective churn produces zombie markets: bankruptcy survival can remain positive while marketalive fraction, seller activity, and assortment diversity vanish. Replenishment at rate λ_P is the primary bifurcation parameter separating exhaustion from stable diversity. Third, governance separates trust from participation: weak moderation collapses trust; strong moderation can preserve aggregate trust near 0.7 while seller participation collapses—a failure mode invisible to trust-only KPIs.

The policy implication is that AI seller tools should be evaluated on joint vitality metrics, not on listing counts, migration speed, or platform trust alone. The scientific implication is that marketplace ecology under AI is a systems problem: migration speed, opportunity replenishment, competition, and governance must align for a living market to persist.

Funding

There is no funding to report.

Conflicts of Interest

The author(s) declare no conflicts of interest regarding the publication of this paper.

Acknowledgements

The author thanks colleagues at the Courant Institute of Mathematical Sciences and New York University for feedback on agent-based modeling and platform economics. Computational experiments used the market churn abm replication package bundled with this manuscript.

References

Akerlof, G. A. (1970). The market for lemons: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, *84*(3), 488–500. https://doi.org/10.2307/1879431

Anderson, C. (2004). *The long tail*. Wired.

Armstrong, M. (2006). Competition in two-sided markets. *RAND Journal of Economics*, *37*(3), 668–691. https://doi.org/10.1111/j.1756-2171.2006.tb00037.x

Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics*, *107*(3), 797–817. https://doi.org/10.2307/2118364

Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, *99*(suppl. *3)*, 7280–7287. https://doi.org/10.1073/pnas.082080899

Edelman, B. and Luca, M. (2014). Digital discrimination: The case of Airbnb.com. Harvard Business School NOM Unit Working Paper. https://doi.org/10.2139/ssrn.2377353

Edelman, B. and Geradin, D. (2016). *Platforms and regulation*. In Oxford Handbook of Digital Economics.

Einav, L., Knoepfle, D., Levin, J., and Sundaresan, N. (2014). Sales taxes and internet commerce. *American Economic Journal: Economic Policy*, *6*(4), 1–26.

Eisenmann, T., Parker, G., and Van Alstyne, M. W. (2006). Strategies for two-sided markets. *Harvard Business Review*, *84*(10), 92–101.

Epstein, J. M. (2006). Generative Social Science: Studies in Agent-Based Computational Modeling. Princeton University Press. https://doi.org/10.1515/9781400842872

Farmer, J. D. and Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460, 685–686. https://doi.org/10.1038/460685a

Felten, E., Raji, I. D., et al. (2023). *Generative AI and the future of work*. Technical report, Princeton University.

Gillespie, T. (2018). Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media. Yale University Press. https://doi.org/10.12987/9780300235029

Gorwa, R., Binns, R., and Katzenbach, C. (2020). Algorithmic content moderation: Technical and political challenges in the automation of platform governance. *Big Data & Society*, *7*(1), 1–14. https://doi.org/10.1177/2053951719897945

Grewal, D., Bart, Y., Rosell, F., and Kopalle, P. K. (2021). The future of technology and marketing in retailing. *Journal of Retailing*, *97*(4), 539–549.

Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 83, 1420–1443. https://doi.org/10.1086/226707

Grimm, V. et al. (2010). The ODD protocol: A review and first update. *Ecological Modelling*, 221, 2760–2768. https://doi.org/10.1016/j.ecolmodel.2010.08.019

Klonick, K. (2018). The new governors: The people, rules, and processes governing online speech. *Harvard Law Review*, 131, 1598–1670.

Liang, C., and Turban, E. (2011). Introduction to the special section social commerce: A research framework for social commerce. *International Journal of Electronic Commerce*, *16*(2), 5–14. https://doi.org/10.2753/jec1086-4415160201

Hagiu, A. and Wright, J. (2015). Multi-sided platforms. *International Journal of Industrial Organization*, 43, 162–174. https://doi.org/10.1016/j.ijindorg.2015.03.003

Macal, C. M. and North, M. J. (2010). Tutorial on agent-based modelling and simulation. *Journal of Simulation*, 4, 151–162. https://doi.org/10.1057/jos.2010.3

Mollick, E. R. et al. (2023). *Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality.* Harvard Business School Working Paper.

OECD (2023). *Artificial intelligence in e-commerce*. Policy discussion on AI tools in online retail.

Railsback, S. F. and Grimm, V. (2019). Agent-Based and Individual-Based Modeling: A Practical Introduction. Princeton University Press. https://doi.org/10.2307/jj.28274141

Rochet, J.-C. and Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, *1*(4), 990–1029. https://doi.org/10.1162/154247603322493212

Shapiro, C. (1983). Premiums for high quality products as returns to reputations. *Quarterly Journal of Economics*, *98*(4), 659–679. https://doi.org/10.2307/1881782

Stiglitz, J. E. (1984). *Information and economic analysis: A perspective*. *Economic Journal*, 94, 23–41.

Valle, F. et al. (2020). *E-commerce and the market for lemons. Information Economics and Policy*, 52, 100883.

Zhu, F. and Liu, Q. (2019). Competing with complementors: An empirical look at Amazon.com. *Strategic Management Journal*, 40, 1827–1851.