SECTION 3.

MARKETING AND LOGISTICS ACTIVITIES

Ivitskiy Igor
PhD, Associate Professor, Founder Doctor Ads LTD
Doctor Ads LTD, London

FEEDBACK LOOP AMPLIFICATION IN ALGORITHMIC BUDGET ALLOCATION: A DISCRETE LOGISTIC MODEL OF THE ADVERTISER DEATH SPIRAL

Modern automated bidding systems on Google Performance Max and Meta Advantage+ behave as contextual bandits whose reward channel is the advertiser's own conversion stream [1]. This architecture assumes label fidelity. Yet the Angluin-Laird PAC framework shows that symmetric label noise at rate η inflates sample complexity by a factor of $1/(1-2\eta)^2$, diverging as $\eta \to 0.5$ [3]. The question this abstract addresses is narrower and, to the author's reading, underexplored in the applied literature: given a nonzero initial contamination, how quickly does algorithmic reinforcement drive the training distribution across the learnability boundary? Industry telemetry places invalid and fraudulent activity in the double digits across programmatic lead inventory, with the ANA's supply chain transparency study finding that only 36 cents of every dollar entering a DSP reaches a quality impression [2]. The question this abstract addresses is narrower and, to the author's reading, underexplored in the applied literature: given a nonzero initial contamination, how quickly does algorithmic reinforcement drive the training distribution across the learnability boundary?

Let $f_n \in [0,1]$ denote the fraction of fraudulent conversions entering the optimizer's training buffer at cycle n, where a cycle corresponds to one budget reallocation step. Because each observed conversion is treated as evidence of audience quality, the platform redirects marginal spend toward the segment that produced it. Fraud sources therefore receive reinforcement proportional to their current share, while their saturation is bounded above by unity. The resulting recurrence is the discrete logistic map

$$f_{n+1} = f_n + \alpha f_n(1 - f_n)$$

where $\alpha > 0$ *is the feedback amplification rate, an aggregate of the platform's* *learning rate, its exploration-exploitation balance, and the elasticity of fraud supply. The map has two fixed points. The origin $f^* = 0$ is unstable for any $\alpha > 0$, since the Jacobian reduces to $1 + \alpha > 1$. The upper fixed point $f^* = 1$ is stable: once fraud saturates the training buffer, no self-correction pathway exists inside the closed loop. The intermediate value $f_c = 0.5$, corresponding to the Angluin-Laird impossibility boundary, is not a fixed point but a transit threshold beyond which PAC-learnability formally collapses.*

Approximating the discrete recurrence by its continuous analogue $df/dn = \alpha f(1-f)$ and integrating yields the familiar closed form $f(n) = f_0/[f_0 + (1-f_0)e^{-\alpha n}]$. Solving for the cycle at which $f(n) = f_c$ gives

$$n_c = \frac{1}{\alpha} ln\left(\frac{f_c(1-f_0)}{f_0(1-f_c)}\right)$$

which at $f_c = 0.5$ simplifies to $n_c = \alpha^{-1}\, ln\big((1-f_0)/f_0\big)$. Two numerical instantiations illustrate the practical horizon. At a moderate feedback rate $\alpha = 0.15$ and an initial contamination $f_0 = 0.25$ consistent with the Anura benchmark [2], collapse occurs at $n_c \approx 7.3$ cycles. Under the same α but a cleaner starting point of $f_0 = 0.05$, the advertiser retains roughly twenty cycles of usable signal before the boundary is crossed. The sensitivity to initial conditions is logarithmic, while the sensitivity to α is strictly inverse. A practitioner who doubles the platform's aggressiveness halves the runway.

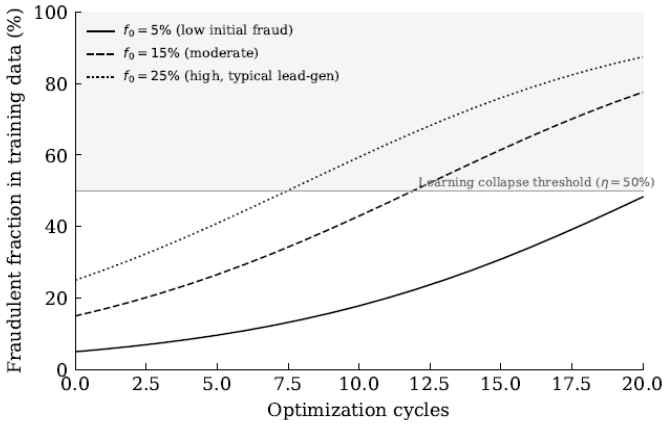

Fig. 1. **Feedback loop amplification: growth of the fraudulent fraction in training data over optimization cycles, for three initial contamination levels. At $f_0 = 5\%$ (low initial fraud, solid line), the learning collapse threshold is not reached within 20 cycles. At $f_0 = 15\%$ (moderate, dashed), the threshold is crossed around cycle 15. At $f_0 = 25\%$ (typical lead-gen, dotted), the threshold is crossed around cycle 7. The horizontal line at 50% marks the theoretical boundary beyond which learning becomes impossible (Eq. (10)).**

Validation of the reverse dynamic is consistent with the theoretical result of Natarajan and colleagues, who show that asymmetric label noise, precisely the class produced by conversion fraud, where negatives are flipped to positives but not the reverse, degrades learning roughly 1.7 times more severely than symmetric noise at the same aggregate rate, so that its excision produces disproportionate, nonlinear gains rather than proportional denoising [4]. This is the qualitative signature the logistic model predicts: abrupt detrapping from a late-stage attractor rather than incremental recovery.

Two implications follow. First, the operative metric for campaign health is not the point estimate of f_n but its trajectory: a 10 percent fraud rate with $\alpha = 0.20$ is strictly worse than a static 20 percent rate. Second, because the data processing inequality forbids recovery of information lost upstream [5], interventions placed after ingestion cannot reset f_n; only pre-ingestion verification, aligned with the signal hygiene paradigm of the M.A.T.H. framework [6] and its API-native operationalization [7], acts on the recurrence itself. The death spiral is not a failure of the optimizer. It is the optimizer functioning exactly as designed on a corrupted objective.

References:

1. Calvano E., Calzolari G., Denicolò V., Pastorello S. Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*. 2020. Vol. 110, No. 10. P. 3267–3297. DOI: https://doi.org/10.1257/aer.20190623.
2. Association of National Advertisers. Programmatic media supply chain transparency study. ANA, 2023. URL: https://www.ana.net/miccontent/show/id/rr-2023-12-ana-programmatic-media-supply-chain-transparency-study (accessed: 06.04.2026).
3. Angluin D., Laird P. Learning from noisy examples. *Machine Learning*. 1988. Vol. 2, No. 4. P. 343–370. DOI: https://doi.org/10.1007/BF00116829.
4. Natarajan N., Dhillon I. S., Ravikumar P., Tewari A. Learning with noisy labels. *Advances in Neural Information Processing Systems*. 2013. Vol. 26. P. 1196–1204. URL: https://papers.nips.cc/paper/2013/hash/3871bd64012152bfb53fdf04b401193f-Abstract.html (accessed: 06.04.2026).
5. Cover T. M., Thomas J. A. Elements of information theory. 2nd ed. Hoboken: Wiley, 2006. 776 p. DOI: https://doi.org/10.1002/047174882X.
6. Ivitskiy I., Savchenko D., Sydorenko D. Data integrity as the terminal constraint in AI-driven advertising: an information-theoretic analysis of conversion fraud and agentic threat evolution. *Zenodo*. 2026. DOI: https://doi.org/10.5281/zenodo.18675362.
7. Ivitskiy I. API-native marketing operations: a reference architecture for automated decision loops. *Global Trends in Economics, Finance and Management*: proceedings of the III International scientific and practical conference (Kyiv, September 30, 2025). Kyiv: Research Europe, 2025. P. 101–104. URL: https://researcheurope.org/wp-content/uploads/2025/10/re-30.09.2025-101-104.pdf (accessed: 06.04.2026).