The Matchmaker's Fixed Point: Engagement-Optimized Matchmaking as Performative Prediction

Adnan Selimović, Ph.D. Dividual — Working Paper WP-2026-01 (v1.1) · July 2026 *Working paper. Comments welcome: hello@dividual.studio. Formal statements in §2 are quoted with their exact hypotheses from the cited sources; §3 is a mapping, offered as one, with its load-bearing assumption stated in §5.*

Abstract

Engagement-optimized matchmaking (EOMM) estimates each player's churn risk, assembles matches that minimize predicted churn, and retrains its churn model on the behavior that follows. The behavior that follows is behavior the matchmaker itself produced. We show that this pipeline is a *performative predictor* in the sense of Perdomo, Zrnic, Mendler-Dünner & Hardt (2020): the deployed model shifts the distribution it is trained on, and naïve retraining converges to a *performatively stable point* — a model that is optimal on the data it induces. Stable points and performative optima are distinct solution concepts, and the gap between them can be large; Miller, Perdomo & Zrnic (2021) exhibit cases where the retraining fixed point *maximizes* the risk it was built to minimize, with a suboptimality gap that can be made arbitrarily large. Mapped onto matchmaking, the consequence is concrete: a studio can run a technically sound retraining loop, watch its metrics agree with its model, and be parked at an equilibrium that is suboptimal for engagement itself — and further still from player welfare. To our knowledge this mapping is absent from the published literature: the canonical EOMM formulation treats churn risk as an exogenous supervised-learning target, and the nearest neighbor (Chen, Elmachtoub & Lei, 2026) models the induced distribution as a control lever in the hands of an informed planner; our subject is the trap it sets for an uninformed one. We state the mapping, mark its load-bearing assumption — that the matchmaker retrains on its own induced outcomes — and derive falsifiable consequences a telemetry team could check against its own retraining cadence.

Keywords: matchmaking, churn, performative prediction, feedback loops, live-service games, EOMM, SBMM

---

1. A prediction that moves its own target

A churn model in a live-service game occupies an unusual epistemic position. It predicts whether a player will leave; its prediction feeds a matchmaker; the matchmaker assembles the player's next match; and the match is among the strongest short-run influences on whether the player leaves. The model's output flows, through the policy it serves, into the ground truth it will be scored against. Prediction here has a property that classical supervised learning assumes away: it moves its own target.

The canonical formulation makes the classical assumption explicit. Chen, Xue, Kolen, Aghdaie, Zaman, Sun & Seif El-Nasr (2017) introduce EOMM — Engagement Optimized MatchMaking — as a two-stage pipeline: learn a churn-risk function $c(s_i, s_j)$ from historical match data by

supervised learning, then solve a minimum-weight perfect matching to minimize summed predicted churn across the pool. The churn function is estimated once, from data, as a static object. Nowhere does the formulation register that deploying the matchmaker changes the distribution of matches, hence the distribution of churn, hence the data the next churn model will be fitted to. The pipeline predicts passively and acts consequentially, and the two stances are never reconciled.

The empirical premise — that match assembly moves churn — is documented, with a stated caveat. Kang, Suh & Kim (2024), in a two-way fixed-effects panel over roughly six million matches and 262,698 players of a commercial mobile game, find churn responds to opponent assignment: being matched against stronger opponents raises exit probability, weaker opponents lower it, and the effect survives player and time fixed effects. The design is observational — nothing is randomized, so "causal" would overstate what the panel identifies — but the existence and direction of the response is all the argument below requires, and no step leans on its magnitude. In these estimates, match assignment behaves like a treatment. A matchmaker that assigns matches on the basis of predicted churn is administering the treatment its own training data records the response to.

Once both pieces are on the table — the model retrains on outcomes, the outcomes respond to the model — one question is settled: the loop exists. The live question is where it settles, and whether the settling point is anywhere a studio would choose on purpose. That question has a formal home, and the formal home has results in it. Our contribution is the mapping: we name the loop with the machinery built for it, transfer the machinery's central distinction to matchmaking, and state what the distinction predicts about telemetry. The intuition that "the matchmaker shapes what it measures" will strike any practitioner as familiar once stated. The contribution is the precise identification of *which* equilibrium the loop selects, why that equilibrium can be bad on the optimizer's own terms, and what would count as evidence either way.

2. The formal spine: two solution concepts, three results

Perdomo, Zrnic, Mendler-Dünner & Hardt (2020) call a prediction *performative* when it influences the outcome it aims to predict, and formalize the setting with one object added to the classical picture: a **distribution map** $\mathcal{D}(\cdot)$, assigning to each model $\theta \in \Theta \subseteq \mathbb{R}^d$ the data distribution $\mathcal{D}(\theta)$ over instances $z \in \mathcal{Z}$ that deploying θ induces. The map's strength is measured by *sensitivity*: $\mathcal{D}(\cdot)$ is ε-sensitive if for all $\theta, \theta' \in \Theta$,

$$W_1(\mathcal{D}(\theta), \mathcal{D}(\theta')) \;\leq\; \varepsilon \, \|\theta - \theta'\|_2,$$

where W_1 is the Wasserstein-1 (earth-mover's) distance. With a loss $\ell(z; \theta)$, define the *decoupled performative risk* $\mathrm{DPR}(\theta, \theta') = \mathbb{E}_{z \sim \mathcal{D}(\theta)} \, \ell(z; \theta')$ — the loss of model θ' evaluated on the world that model θ makes — and the *performative risk* $\mathrm{PR}(\theta) = \mathrm{DPR}(\theta, \theta)$. Two solution concepts follow, and their separation carries everything below.

A model θ_{PS} is **performatively stable** if it minimizes expected loss on the distribution it itself induces:

$$\theta_{\mathrm{PS}} \;\in\; \arg\min_{\theta}\; \mathbb{E}_{z\sim\mathcal{D}(\theta_{\mathrm{PS}})}\,\ell(z;\theta) \;=\; \arg\min_{\theta}\; \mathrm{DPR}(\theta_{\mathrm{PS}},\theta).$$

Stable points are exactly the fixed points of *repeated risk minimization* (RRM), the retraining iteration $\theta_{t+1} = \arg\min_\theta \mathrm{DPR}(\theta_t,\theta)$: deploy, collect the data the deployment generates, refit, repeat. A system at θ_{PS} looks converged, calibrated, and healthy from the inside. Retraining confirms it; evaluation on freshly induced data confirms it; the loop is closed and quiet.

A model θ_{PO} is **performatively optimal** if it minimizes the performative risk itself:

$$\theta_{\mathrm{PO}} \;\in\; \arg\min_{\theta}\; \mathrm{PR}(\theta) \;=\; \arg\min_{\theta}\; \mathbb{E}_{z\sim\mathcal{D}(\theta)}\,\ell(z;\theta).$$

The optimum accounts for the world each candidate model *would* make; the stable point answers only to the world the incumbent model *has* made. Three results, stated with their exact hypotheses, fix how far apart the two can sit.

Result 1 — convergence (Perdomo et al. 2020, Thm. 3.5). If ℓ is β-jointly smooth and γ-strongly convex in θ, and $\mathcal{D}(\cdot)$ is ε-sensitive with $\varepsilon < \gamma/\beta$, then RRM converges at a linear rate to a unique performatively stable point. The threshold is tight: at $\varepsilon \geq \gamma/\beta$ there are instances — with arbitrarily smooth losses and arbitrarily small sensitivity relative to that threshold — where RRM oscillates forever and no convergence occurs. Retraining pipelines land on stable points when they land anywhere.

Result 2 — proximity, in the tame regime (Perdomo et al. 2020, Thm. 4.3). If ℓ is L_z-Lipschitz in z and γ-strongly convex in θ, and $\mathcal{D}(\cdot)$ is ε-sensitive, then every stable point and every optimum satisfy

$$\|\theta_{\mathrm{PO}} - \theta_{\mathrm{PS}}\|_2 \;\leq\; \frac{2L_z\,\varepsilon}{\gamma}.$$

When performative effects are weak relative to the loss's curvature ($\varepsilon \ll \gamma$), retraining is approximately optimal and the classical picture survives.

Result 3 — the trap (Miller, Perdomo & Zrnic 2021, Prop. 2.1). For any $\gamma > 0$ and any $\Delta > 0$ there exists a performative prediction problem whose loss is γ-strongly convex in θ, in which the unique stable point *maximizes* the performative risk, with suboptimality gap $\mathrm{PR}(\theta_{\mathrm{PS}}) - \mathrm{PR}(\theta_{\mathrm{PO}}) \geq \Delta$. Strong convexity of the loss buys nothing by itself: outside the small-ε regime, retraining to convergence can be a procedure for locating the *worst* model that agrees with itself, and the gap has no ceiling.

The asymmetry between Results 2 and 3 is the paper's hinge, so it deserves one plain sentence: whether a studio's retraining loop is harmlessly near-optimal or parked at an arbitrarily bad fixed point is governed by the ratio of the distribution map's sensitivity to the loss's curvature — a ratio nobody estimates, because the standard pipeline contains no object that represents $\mathcal{D}(\cdot)$ at all. Escaping the trap requires exactly that missing object: optimizing PR directly needs added structure even for convexity (a mixture-dominance condition, with its own tight threshold $\varepsilon < \gamma/(2\beta)$; Miller et al. 2021, Thm. 3.1), and the algorithms that provably reach θ_{PO} do it

by estimating how the deployed model moves the distribution (performative gradient descent; Izzo, Ying & Zou, 2021) — they model the loop the naïve pipeline omits.

Two extensions matter for the transfer to games. Brown, Hod & Kalemaj (2022) treat the *stateful* case, where the population's response depends on both the deployed model and the population's current state — the version that fits a player base whose disposition carries over between matches, sessions, and seasons. Narang, Faulkner, Drusvyatskiy, Fazel & Ratliff (2023) treat *multiplayer* performative prediction, multiple decision-makers with coupled distribution maps — the version that fits a matchmaker whose pairing decisions couple players by construction. A formally minded reader should anchor on these two variants; the original single-decision-maker model is the pedagogical spine, and its assumptions are the ones a matchmaking transfer must generalize.

The recommender-systems literature reached the same structure empirically before it had this name. Chaney, Stewart & Engelhardt (2018) show that recommenders trained on data from users already exposed to recommendation enter a confounded loop that "homogenizes user behavior without increasing utility." Jiang, Chiappa, Lattimore, György & Kohli (2019) formalize the degeneracy conditions. Carroll, Dragan, Russell & Hadfield-Menell (2022) make the mechanism causal and measurable: reinforcement-learning recommenders learn to *shift user preferences* in directions that make their own future predictions easier. And the welfare gap at the equilibrium has an empirical flagship: in a pre-registered audit, Milli, Carroll, Wang, Pandey, Zhao & Dragan (2025) find that engagement-based ranking on a major platform amplifies content users themselves report disliking — the engagement optimum and the users' stated preferences come apart in the field, on production systems, measurably.

3. The mapping

The transfer is nearly mechanical, which is why its absence from the literature is notable. Take θ to be the parameters of the matchmaker's churn model c_θ, and let π_θ be the assembly policy it feeds — in the EOMM formulation, the minimum-weight perfect matching on the pool under predicted pairwise churn. Let $\mathcal{D}(\theta)$ be the distribution over telemetry instances — match compositions, outcomes, and subsequent player behavior: retention, session cadence, exit — that deploying π_θ induces. One structural note before the checklist: the pipeline is two-stage (a prediction loss trains c_θ; a decision objective consumes it), so the mapping runs at the level of the composite $\theta \mapsto \pi_\theta \mapsto \mathcal{D}(\theta)$ — which is the object the sensitivity definition measures, and the only object retraining ever touches. Then:

1. **The distribution map is real and strong.** Match assignment moves churn (Kang et al., 2024); assembly policy determines assignment; therefore θ moves $\mathcal{D}(\theta)$. In the formal vocabulary, the map's W_1-sensitivity is bounded away from zero — and everything a studio knows about the salience of match quality says it is large. Since the tame regime of Result 2 requires ε *small* relative to the loss's curvature, the case for complacency has to be argued, and nobody has an estimate of either constant.

2. **The retraining loop is the stability dynamic.** A pipeline that periodically refits its churn model on accumulated telemetry, under a live assembly policy, is running (empirical, approximate) repeated risk minimization — the iteration $\theta_{t+1} = \arg\min_\theta \mathrm{DPR}(\theta_t, \theta)$ whose fixed points are precisely the performatively stable models. Where the pipeline converges, it converges to a θ_{PS}.

3. **The two concepts come apart here as everywhere.** There is no special dispensation for games. The stable matchmaker is optimal against the player behavior it manufactures; the performative optimum is defined against the behavior each candidate policy *would* manufacture. The gap between them is invisible to any evaluation run on induced data — which is to say, invisible to the standard evaluation stack.

4. **The equilibrium can be bad on the optimizer's own terms.** This is the Miller et al. point, and it deserves emphasis because it closes off a comfortable reading. The claim is stronger than "engagement-optimal play may be bad for players" — that gap is already established one level up (Kleinberg, Mullainathan & Raghavan, 2024; Milli et al., 2025). The claim is that the retraining loop can converge to a point that is poor *for engagement itself*: a self-confirming matchmaker that no longer explores, scores well against the population it has shaped, and forgoes basins of genuinely higher retention that a distribution-aware optimizer could reach. Góis, Mofakhami, Santos, Gidel & Lacoste-Julien (2025) demonstrate the same structure in a mechanism-design setting: the search for stable, accurate predictions can minimize social welfare with high probability. Stability is an agreement between the model and its own consequences. Nothing in the agreement mentions quality.

A dashboard consequence follows immediately, and it is the one practitioners feel before they name it. At a stable point, the metrics computed on induced behavior — retention curves, session length, quit rates — are functionals of the distribution the optimizer selected. They are the system grading its own homework. Their agreement with the model is an equilibrium condition, and reading that agreement as external validation mistakes a fixed point for a fact about players.

4. Prior art, and the increment stated precisely

The nearest published neighbor deserves careful separation, because it reaches the endogeneity and stops. Chen, Elmachtoub & Lei (2026) model a matchmaker whose pairing policy endogenously steers the population's skill and churn composition, and solve for the policy that exploits this — matching myopically for short-run reward while adjusting the skill distribution for long-run engagement, with results on 5.4 million Lichess matches. The induced distribution appears in their model as a *lever*: a fully informed planner steers it deliberately, by dynamic programming, toward higher engagement. The framing is performative *optimality* — what the steering is worth to a planner who models it. Our subject is the complement: performative *stability* — where a pipeline lands when it does the steering *without* modeling it, retraining naïvely on the distribution it has already moved. Same endogeneity, opposite normative pole. Their planner never retrains on induced data, so the trap we describe cannot appear in their model; conversely, nothing in their optimal-control results protects a production pipeline that retrains from converging to a stable, suboptimal point. The two papers describe the ceiling and the floor of the same loop. (The paper uses none of the performative-prediction vocabulary and cites none of that literature — consistent with a mapping that sits unwritten, and unopposed.)

Beyond that: the EOMM source itself, as noted, treats churn as exogenous — the passive-prediction stance is in the canon, unremarked. Optimization extensions (neural-combinatorial and RL matchmakers) inherit it. Searches across the games and machine-learning literatures return, to our knowledge, no co-occurrence of the performative-prediction formalism with matchmaking: the claim "matchmaking has not previously been framed as a performative

predictor" is a bounded negative — an absence in the published record as of mid-2026, stated as such, and falsified the day someone surfaces the paper we could not.

So the increment, claimed narrowly: **(i)** the explicit mapping of the EOMM-style churn-retraining pipeline onto performative prediction, with the stateful and multiplayer variants identified as the correct formal homes; **(ii)** the transfer of the stable-$\neq$-optimal separation — including its sharpest form, stability at a risk-maximizing point — to match assembly; **(iii)** the normative inversion relative to Chen–Elmachtoub–Lei: they read the induced distribution as the informed planner's instrument; we read it as the naïve loop's pathology. The underlying intuition is claimed by nobody and belongs to everybody; a reader who finds the top-line thought obvious once stated is agreeing with us about its truth, and the fixed-point identification is where the content lives.

5. What the mapping predicts, and what would soften it

A reframe earns its keep by being falsifiable. This one carries one load-bearing assumption and several checkable consequences.

The assumption. The pathology requires the loop to actually close: the matchmaker (or the churn model feeding it) retrains on outcomes its own policy induced. Where matchmaking is static skill-and-connection SBMM — tuned rarely, never refit on induced engagement — the performative-stability claim softens to the milder observation that evaluation data is policy-induced (a confounding problem, still real, with the Chaney et al. structure) without the fixed-point dynamics. The strength of the claim, system by system, is an empirical fact about the retraining architecture, and the first diagnostic question is simply: what, in this pipeline, is refit on what, and on what cadence?

Consequences, where the loop closes. Each is stated so that telemetry could contradict it. First, *retraining-cadence dependence*: if the system is converging to a stability point, successive refits should show geometrically shrinking update norms $\|\theta_{t+1} - \theta_t\|$ — the linear-rate contraction of Result 1, directly observable in a pipeline's own model-registry history — while the model's counterfactual performance, evaluated on behavior induced by a *different* policy (a holdout region, an A/B arm running an alternative matchmaker), need not improve and can degrade. Agreement-with-self and quality part ways, measurably. Second, *exploration asymmetry*: policies that deliberately perturb assembly (randomized match windows, forced exploration) should reveal engagement gradients the converged policy reports as flat — the signature of a local basin the stable point cannot see over. Third, *homogenization*: under the Chaney–Jiang mechanism, the distribution of realized match compositions should narrow over successive retraining generations faster than the player population itself narrows — the optimizer converging on the population it has manufactured. Each of these is an existence-level claim about system dynamics, estimable in principle from experiment logs and retraining histories a live-service studio already keeps; each needs no new instrumentation, just a design that treats the optimizer, for once, as the object of measurement.

What the mapping does *some* work against even when soft. Even without closed-loop retraining, the metrics point stands: engagement metrics computed on policy-induced behavior are functionals of the policy's objective and cannot arbitrate between that objective and player welfare (Milli et al., 2025, is the field demonstration). The fixed-point machinery sharpens this from a confounding complaint into an equilibrium prediction, but the complaint is valid alone.

6. Why a studio should care

The commercial reading requires no appeal beyond the studio's own interest. A live-service game sells continued participation in a market where the player can always leave; its long-run revenue is a functional of genuine retention, and genuine retention is exactly what a performatively stable matchmaker is no longer optimizing for. The stable point defends itself against the evidence — its metrics agree with it — while the game next door, under a different policy, discovers the basin it could not. Churn is where the gap cashes out: players who exit from an equilibrium that was locally quiet and globally wrong, exits the converged evaluation stack had no way to attribute. The diagnosis and the business case are the same sentence. A studio that wants to know whether it is parked at a fixed point has, in its retraining history and its experiment logs, the data to ask; the asking is a measurement-and-identification program, and the first result it returns is which kind of loop the studio is actually running.

No conspiracy is required at any step. A sincere team, a sound pipeline, an honest metric — and a fixed point that files the loop's own signature under system health.

References

Brown, G., Hod, S. & Kalemaj, I. (2022). Performative prediction in a stateful world. *AISTATS*, PMLR 151:6045–6061.

Carroll, M., Dragan, A., Russell, S. & Hadfield-Menell, D. (2022). Estimating and penalizing induced preference shifts in recommender systems. *ICML*, PMLR 162:2686–2708.

Chaney, A.J.B., Stewart, B.M. & Engelhardt, B.E. (2018). How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. *RecSys '18*, 224–232.

Chen, M., Elmachtoub, A.N. & Lei, X. (2026). Matchmaking strategies for maximizing player engagement in video games. *Management Science*. DOI 10.1287/mnsc.2023.02957.

Chen, Z., Xue, S., Kolen, J., Aghdaie, N., Zaman, K.A., Sun, Y. & Seif El-Nasr, M. (2017). EOMM: An engagement optimized matchmaking framework. *WWW '17*, 1143–1150.

Góis, A., Mofakhami, M., Santos, F.P., Gidel, G. & Lacoste-Julien, S. (2025). Performative prediction on games and mechanism design. *AISTATS*, PMLR 258.

Izzo, Z., Ying, L. & Zou, J. (2021). How to learn when data reacts to your model: Performative gradient descent. *ICML*, PMLR 139:4641–4650.

Jiang, R., Chiappa, S., Lattimore, T., György, A. & Kohli, P. (2019). Degenerate feedback loops in recommender systems. *AIES '19*, 383–390.

Kang, J., Suh, E. & Kim, J. (2024). Match experiences affect interest: Impacts of matchmaking and performance on churn in a competitive game. *Heliyon* 10(3):e24891.

Kleinberg, J., Mullainathan, S. & Raghavan, M. (2024). The challenge of understanding what users want: Inconsistent preferences and engagement optimization. *Management Science* 70(9):6336–6355.

Miller, J., Perdomo, J.C. & Zrnic, T. (2021). Outside the echo chamber: Optimizing the performative risk. *ICML*, PMLR 139.

Milli, S., Carroll, M., Wang, Y., Pandey, S., Zhao, S. & Dragan, A.D. (2025). Engagement, user satisfaction, and the amplification of divisive content on social media. *PNAS Nexus* 4(3):pgaf062.

Narang, A., Faulkner, E., Drusvyatskiy, D., Fazel, M. & Ratliff, L.J. (2023). Multiplayer performative prediction: Learning in decision-dependent games. *JMLR* 24(202):1–56.

Perdomo, J.C., Zrnic, T., Mendler-Dünner, C. & Hardt, M. (2020). Performative prediction. *ICML*, PMLR 119:7599–7609.