WELFARE-ENHANCING MECHANISM DESIGN FOR AGENT ECOSYSTEMS

Anonymous authors
Paper under double-blind review

ABSTRACT

The emerging Agentic Web forms a multi-agent ecosystem: a decentralized service market where the user agent selects which content agents to query from their submitted profiles, before seeing any documents. Profiles therefore become a new target of optimization. Existing Generative Engine Optimization studies how a single supplier rewrites documents the engine can already retrieve, to improve their citation. In the Agentic Web, multiple content agents optimize profiles to compete for limited opportunities to deliver documents. To study this repeated interaction, we build WILDREGISTRY, a simulation framework and benchmark based on a real agentic-search request stream, and model it as a dynamic game under private monitoring. Long-horizon experiments show that this competition lowers website-selection quality and stabilizes at a harmful approximate rewriting equilibrium. We therefore design a welfare-enhancing mechanism, DEC (Delivery-Evidence Credit), a platform-side policy that uses recorded delivery outcomes to reallocate later opportunities and reduce information asymmetry. DEC improves this endpoint, raising website-selection quality while preserving most of providers' exposure gains, for a Welfare Gain of $+2.01$ points and a healthier long-horizon outcome. More broadly, our findings highlight the importance of mechanism design in agent ecosystems: turning recorded delivery outcomes into evidence that endogenously shapes later user-agent selection and content-agent adaptation can steer long-horizon adaptation toward delivered capabilities.

1 INTRODUCTION

LLM-based generative engines are changing how users access web information: beyond ranking pages for users (Joachims et al., 2017a), an engine integrates retrieved sources into a final answer (Dai et al., 2025). These decisions shape provider traffic and visibility. Airbyte reports that AI-search optimization nearly tripled its ChatGPT visibility and brought in a $100,000 deal (Profound, n.d.; Sarkar, 2025). Such returns motivate providers to optimize documents for these decisions; answer-stage Generative Engine Optimization (GEO) is one instance (Aggarwal et al., 2024).

The emerging Agentic Web changes this single-provider setting: users and multiple content providers interact through their respective agents (Yang et al., 2025). This shift is visible in Web traffic: automated systems generate over half of Web traffic (Weiss et al., 2026; Chang, 2026), while requests from AI agents and agentic browsers grew nearly eightyfold in one year (HUMAN Security, 2026). In the Agentic Web, multi-content-agent competition emerges: independently operated content agents submit short website profiles and compete for website selection, winning opportunities to deliver documents and gain exposure (Zhong et al., 2026). Website selection therefore becomes a new upstream objective, shifting optimization from documents to profiles. How this strategic interaction evolves over a long horizon remains understudied. We therefore build WILDREGISTRY, a simulation framework for studying it. It replays a real agentic-search request stream (Ning et al., 2026; Coelho et al., 2026) in timestamp order through a fixed, decentralized hundred-site service market (Zhong et al., 2026) and enables some content agents to act as adaptive providers that update their profiles between rounds, so that provider-side adaptation accumulates over the stream.

Over 60 active days in WILDREGISTRY, this one-sided adaptation settles into a harmful stable point (Figure 1). Adaptive providers gain 48.6% more exposure, while the user agent's website-selection quality is 1.25 points below the static-profile Original baseline. Deployed profiles are

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

almost unchanged over the final 20 active days, so the endpoint is a harmful approximate rewriting equilibrium (Section D.2). The ecosystem therefore stabilizes at an unhealthy endpoint.

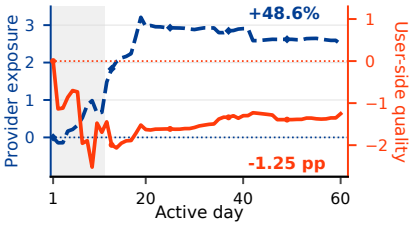

Figure 1: **Provider-side optimization harms the ecosystem** (cumulative, vs. Original).

This unhealthy endpoint stems from information asymmetry on both sides. The user agent faces one side: selection precedes verification, so it chooses from self-authored profiles and unverified claims can displace better-matched sources. Each content agent faces the other: under private monitoring it rewrites only from records its own selections produced (Kandori, 2002), and capability claims whose documents are never cited still win selections and accumulate in its profile. Reducing this asymmetry therefore requires platform-side selection to adapt to accumulated delivery feedback. We formulate this repeated interaction as a dynamic game under private monitoring; the mechanism-design problem is to use observable delivery signals for a healthier ecosystem (Boutilier et al., 2023).

To address both sides, we design a welfare-enhancing mechanism, DEC (Delivery-Evidence Credit): a platform-side policy that reallocates later opportunities using recorded delivery outcomes. It has two complementary actions: Credit rewards verified past delivery before selection, while Withhold withdraws a current opportunity from a website with a revised profile after repeated non-delivery. The resulting service experience shapes later website selection and profile rewriting. Over the long horizon, DEC raises website-selection quality beyond Original while retaining most of rewriting's exposure gain. It also promotes healthier profile rewriting: the share of unbacked claims falls by 43.8% at the endpoint. Together, these results improve the ecosystem: DEC achieves a Welfare Gain of $+2.01$ points, and the system reaches a healthier long-horizon outcome.

This paper makes three contributions:

- We study the emerging multi-agent Web ecosystem through a concrete instance: persistent competition among content agents for website selection. We formulate the interaction as a dynamic game and release WILDREGISTRY, a simulation framework and benchmark on a real agentic-search log.
- We find that this competition degrades the ecosystem; we design a welfare-enhancing mechanism, DEC, that reallocates later opportunities from recorded delivery outcomes and improves both sides.
- We uncover endogenous learning-evidence allocation in markets with adaptive agents: content agents can still optimize, while platform-side opportunity allocation reshapes both sides' learning evidence and moves the ecosystem toward a healthier long-horizon outcome.

2 RELATED WORK

Mechanism design in search and the Internet. Competitive-search models treat the ranking policy as a rule that changes publishers' content and strategic behavior (Tennenholtz & Kurland, 2019). In turn, strategic classification and performative prediction study how responses to a deployed decision rule change the later data distribution (Hardt et al., 2016; Perdomo et al., 2020), and creator-side recommendation extends this view to suppliers competing for exposure (Jagadeesan et al., 2023). Logged-feedback and selective-label work then show that exposure decides which of that data is observed (Swaminathan & Joachims, 2015; Lakkaraju et al., 2017). Ecosystem-level models study users and providers under platform-side policies (Boutilier et al., 2023). In our setting, DEC keeps the LM backbone fixed and adjusts repeated website selection from recorded delivery outcomes, so that selection allocates both current opportunities and the records available to later profile updates.

Multi-agent ecosystems and website selection. Generative engine optimization revises a focal supplier's documents to raise adoption and citation (Aggarwal et al., 2024); a platform-side mechanism rewards verifiable content in repeated supplier–platform interaction (Xu et al., 2026). The Agentic Web moves this competition before retrieval: a user agent selects content agents from profiles, and profile text alone can change LLM selection (Zhong et al., 2026; Blankenstein et al., 2026). Federated search and LLM routing also select sources before retrieval, from broker-written descriptions or past interactions (Shokouhi & Si, 2011; Wang et al., 2025). Both lines generally fix candidates and descriptions, or evaluate one selection. Multi-agent and continual benchmarks study coordination, competition, and adaptation on constructed task streams (Zhu et al., 2025; Tomilin et al., 2025).

We instead study independently operated content agents that persist across requests and update self-authored profiles, not documents, from their own service records on a real request stream.

Safety and spam in agent systems. Classical Web work handles manipulation through spam detection and ranking robustness (Castillo & Davison, 2011; Goren et al., 2018). Agent-safety work studies prompt injection (Greshake et al., 2023), retrieval poisoning (Zou et al., 2025), and tool-selection hijacking (Shi et al., 2025). The profile interface carries two risks. Explicit control instructions steer selection at submission, while ordinary query-conditional capability claims affect selection before the service can be verified. Detection and instruction–data separation address the first (Li et al., 2025; Yi et al., 2025). Our audit screens selection-control instructions at registration; after service, DEC adjusts later opportunities from recorded delivery and non-delivery. Even without explicit attack instructions, the continued adaptation of multiple independent content agents pursuing exposure under private feedback still changes long-run opportunity allocation and user-side outcomes; that ecosystem-level effect is what we study.

3 WILDREGISTRY: AN AGENTIC WEB BENCHMARK

To study how website selection and content-agent profile updates co-evolve, we build WILDREGISTRY, a simulation framework and benchmark.

3.1 PRELIMINARIES

In a single round, let $\mathcal{W}$ be the finite website set, with one content agent serving each $w \in \mathcal{W}$; before service it submits a short profile, and $\rho_t(w)$ denotes the version deployed in round t. This profile is the only website-controlled text that can affect website selection. Let Q_t be the user agent's requests in round t. For request $q \in Q_t$, the user agent selects a limited website set $L_t(q) \subseteq \mathcal{W}$ using the round's registry $R_t = \{\rho_t(w)\}_{w\in\mathcal{W}}$, giving each selected website one opportunity to serve the request. Let $\mathcal{D}(q, w)$ be the documents returned by website w; if $w \notin L_t(q)$, then $\mathcal{D}(q, w) = \varnothing$. Let the final answer be the ranked list of cited documents $(d_1, \ldots, d_n)$. Website-level *cited* status is

$$\delta(q, w) = \mathbf{1}[\exists i \in \{1, \ldots, n\} : d_i \in \mathcal{D}(q, w)]. \tag{1}$$

The realized service record for q is q, $L_t(q)$, $\{(\mathcal{D}(q, w), \delta(q, w))\}_{w\in L_t(q)}$, and the final answer.

3.2 WILDREGISTRY BENCHMARK CONSTRUCTION

WILDREGISTRY combines three fixed components for simulation: a real request stream, a decentralized service market, and a frozen offline evaluation.

On the user side, we replay six months of real search-API traffic from DeepResearchGym in timestamp order (Ning et al., 2026; Coelho et al., 2026). After preprocessing, the stream contains 87 client proxies, 9,771 requests, and 3,726 sessions, one round per UTC day; nearly 30% of requests have semantically similar predecessors, allowing prior service experience to inform later rounds. On the content-agent side, we use AgentWebBench's decentralized service market (Zhong et al., 2026): the 100 websites with the most documents in the English ClueWeb22-B subset, totaling 18,427,770 documents, reachable only through each website's content agent and dense index. Section B gives the full data funnel and recurrence statistics.

Offline evaluation uses UMBRELA for frozen 0–3 relevance judgements over a fixed candidate pool (Upadhyay et al., 2024); related studies also use LLM relevance scoring for retrieval evaluation (Kim et al., 2026a). The 9,043 requests with any document graded above 0 form the evaluation population.

3.3 REPEATED-INTERACTION SIMULATION FRAMEWORK

On this benchmark, content agents update profiles from their visible histories, and the user agent selects websites using those profiles; service outcomes enter the histories at round end and affect later behavior. We model the two sides as players in a repeated dynamic game under private monitoring.

User side. In round t, before returned documents are available, the user agent selects websites using only the current request $q \in Q_t$ and the registry R_t; let π_U^{base} denote this baseline selection rule. Let $Q = \bigcup_t Q_t$ denote all requests in the replay, $\mathbf{x}_t(q) = (x_1, \ldots, x_k)$ the ordered outcomes the user

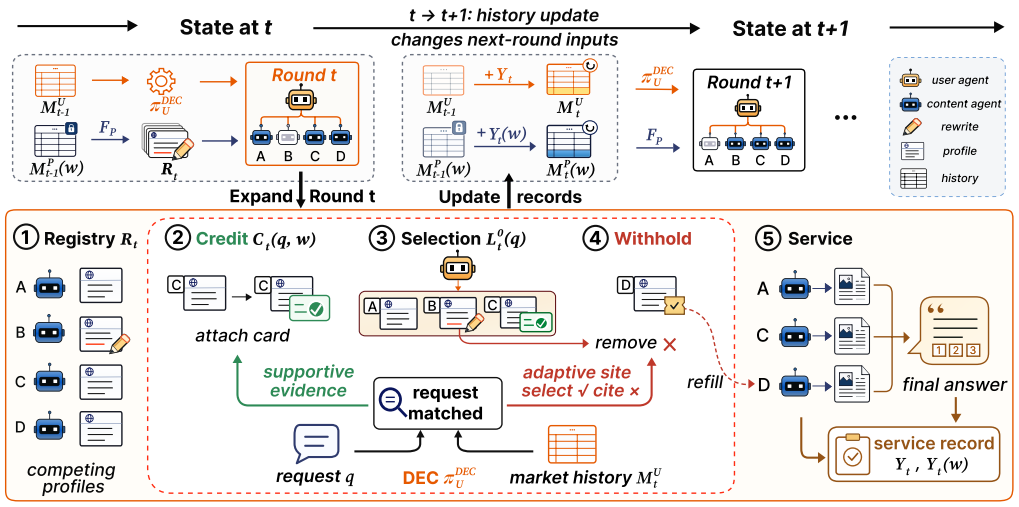

Figure 2: **Cross-round state transitions in WILDREGISTRY and DEC's within-round actions.**

agent produces for q, $\mathrm{rel}(q, x_i)$ the relevance of x_i to q, and ν_i its normalized position weight, so that the request-level user-side relevance is $J_t(q) := \sum_i \nu_i \,\mathrm{rel}(q, x_i)$. Its selection and objective are (exact definitions in Section B.3 and eq. (7))

$$L_t(q) = \pi_U^{\mathrm{base}}(q; R_t), \qquad \mathcal{O}^U = \mathbb{E}_{q\sim Q}[J_t(q)]\,.$$

Provider side. A content agent cannot see competing websites' profiles or service outcomes, the user agent's selection rule, or requests not sent to it. This boundary is provider-side private monitoring (Kandori, 2002); it reflects the position of an independent website in a decentralized market. Website w's content agent is either static, keeping its profile fixed, or adaptive, using F_P to update its profile from its own profile trajectory $\rho_{1:t-1}(w)$ and provider-local service history $M_{t-1}^P(w)$ to increase cited-document exposure. Let $E_t(w, q)$ be the normalized position weight of w's cited documents. Its profile submission $\widehat{\rho}_t(w)$ and objective are (exact definitions in Section B.3)

$$\widehat{\rho}_t(w) = F_P\big(\rho_{1:t-1}(w), M_{t-1}^P(w)\big)\,, \qquad \mathcal{O}^W(w) = \mathbb{E}_{q\sim Q}[E_t(w, q)]\,.$$

Ecosystem. Round t produces market-wide records Y_t and website w's provider-visible records $Y_t(w)$. At round end, $Y_t(w)$ enters $M_t^P(w)$, and F_P uses it to generate the next profile; the resulting registry R_{t+1} therefore depends on earlier service and in turn affects later selection (upper half of Figure 2; Section B.3). Welfare in this paper is the two players' objectives (Madmon et al., 2025a) $\big(\mathcal{O}^U, \sum_{w\in\mathcal{A}} \mathcal{O}^W(w)\big)$, where the adaptive content agents $\mathcal{A}$ together form the provider-side player. A trajectory is an approximate rewriting equilibrium when its deployed profile terms stay nearly unchanged under repeated F_P updates, a stable point (Perdomo et al., 2020) (Section D.2).

Harm from content-agent optimization. Because the content-agent objective $\mathcal{O}^W(w)$ counts only exposure, profile rewriting may increase exposure for lower-relevance matches. For cited answers, the normalized weights satisfy $\sum_{w'\in\mathcal{W}} E_t(w', q) = 1$. Suppose weight α moves from website w^-'s evidence of relevance r^- to website w^+'s evidence of relevance r^+. Then

$$\Delta E_t(w^+, q) = -\Delta E_t(w^-, q) = \alpha > 0, \qquad \Delta J_t(q) = \alpha(r^+ - r^-),$$

which is negative whenever $r^+ < r^-$. We set $\Delta J_t(q) = 0$ for requests without a reallocation and average over the $N_T = \sum_{t=1}^T |Q_t|$ requests through round T. These reallocations give $\Delta\mathcal{O}_{1:T}^U = N_T^{-1} \sum_{t=1}^T \sum_{q\in Q_t} \Delta J_t(q) < 0$, and the cumulative user-side loss is $-N_T \Delta\mathcal{O}_{1:T}^U > 0$. Because $\mathcal{O}^W$ does not account for this loss, such exposure-increasing profiles can remain stable, forming a harmful approximate rewriting equilibrium (stability criterion in Section D.2). Correcting this mismatch requires intervention in the user agent's selection rule.

4 DELIVERY-EVIDENCE CREDIT

To address the harmful endpoint in Section 3.3, we propose DEC, a platform-side policy π_U^{DEC} that reallocates later opportunities from recorded delivery outcomes. Credit rewards recorded delivery, and Withhold penalizes repeated non-delivery. Together they reduce information asymmetry (Figure 2).

4.1 CREDIT, THE REWARD

Before selection, Credit draws delivery evidence from comparable earlier requests and attaches it to website w's profile, making websites with relevant delivery records more likely to be selected.

Let M_{t-1}^U be the platform-maintained market-wide history of service records through round $t-1$. Credit retrieves earlier requests with needs transferable to q, forming $H_t^C(q)$. For $q' \in H_t^C(q)$, $\delta(q', w) = 1$ means at least one document from w was cited on q', and $\text{Support}(q, q', w) = 1$ that one of those titles supports q. The Credit gate asks whether a delivered answer transfers to q:

$$C_t(q, w) = 1 \quad \Longleftrightarrow \quad \exists\, q' \in H_t^C(q) :\ \delta(q', w) = 1 \ \wedge\ \text{Support}(q, q', w) = 1. \tag{2}$$

When $C_t(q, w) = 1$, it attaches one qualifying request and its cited titles to w's profile as a supporting evidence card; the resulting profiles form $R_t^C(q)$ for initial selection (Sections C.1 and H).

4.2 WITHHOLD, THE PENALTY

After initial selection and before service, Withhold withdraws the current opportunity from a website that revised its profile but has repeatedly failed to deliver on comparable requests.

Profile revisions may reflect genuine capabilities, so the gate also checks history: repeated selection without citations means the profile wins opportunities the website has not yet earned. Let $\sigma_t(q)$ be the ordered candidate sequence from the user agent and $L_t^0(q)$ its initial website set. For website w, let $\rho_1(w)$ be its registered profile, $\text{Sel}_t(q, w)$ count its selections among comparable requests, and $\text{Cite}_t(q, w)$ its cited documents in the recalled history (Section C.1). The flagged websites are:

$$\mathcal{W}_t^-(q) = \{w \in L_t^0(q) : \rho_t(w) \neq \rho_1(w),\ \text{Cite}_t(q, w) = 0,\ \text{Sel}_t(q, w) \geq m\}. \tag{3}$$

Withhold removes flagged websites from $L_t^0(q)$, refilling at most one freed slot with the first unselected website in $\sigma_t(q)$ that has a past citation; the gate is re-evaluated on each request. We use $m = 2$ and round-start history for both gates. We write $\text{Append}(M, Y)$ for adding records Y to M in time order. Algorithm 1 gives the cross-round order and the multi-sub-search case.

4.3 MECHANISM AND THE LONG-HORIZON OUTCOME

DEC first acts on current-round selection. A reallocation moves final-answer weight α from evidence of relevance r^- to one of relevance r^+, so user-side relevance changes by $\Delta J_t(q) = \alpha(r^+ - r^-)$. When $r^+ > r^-$, it contributes more to $\mathcal{O}^U$, and the same α is the exposure moved between websites. DEC reallocates exposure without changing either side's objective. For a next-round request $q' \in Q_{t+1}$, records Y_t and $Y_t(w)$ update the two histories entering next-round decisions:

$$\begin{cases} M_t^U = \text{Append}(M_{t-1}^U, Y_t), & L_{t+1}(q') = \pi_U^{\text{DEC}}(q'; R_{t+1}, M_t^U), \\ M_t^P(w) = \text{Append}(M_{t-1}^P(w), Y_t(w)), & \widehat{\rho}_{t+1}(w) = F_P(\rho_{1:t}(w), M_t^P(w)). \end{cases} \tag{4}$$

Against repeatedly non-delivering matches, delivery-backed matches leave stronger selection evidence in M^U and stronger rewrite support in $M^P(w)$ (Section G.1). Under limited profile capacity, this should produce profiles with more delivery backing and selection grounded in it, narrowing the gap between self-written claims and actual delivery and moving the system toward a healthier long-horizon outcome (Lemma 1). We test this along a full trajectory in Section 5.2.

5 EXPERIMENTS

In WILDREGISTRY's long-horizon simulation, we examine where provider-side profile adaptation leads the ecosystem and test whether DEC intervenes as Section 4.3 predicts.

Table 1: **Platform-side methods under the 30-day protocol.** Italic row: reference system. Bold: best; underline: second best, within a backbone.

LLM	System	W-nDCG @3	W-nDCG @5	W-nDCG @10	D-nDCG @3	D-nDCG @5	D-nDCG @10	Provider GEO	Welfare Gain
DeepSeek V4-Flash	*Profile Adapt.*	51.18	51.01	47.59	46.75	41.97	33.43	5.19	0.00
	Full-history ICL	46.48	46.12	43.18	45.24	40.38	32.09	4.64	-5.77
	VCR	51.38	51.22	47.82	46.96	41.98	33.51	3.95	-1.73
	Hard Reject	50.24	50.43	47.06	46.75	41.63	33.13	5.92	+0.56
	Keyword Scrub	50.23	50.40	47.03	47.05	41.91	33.26	**6.27**	+1.07
	DEC	**53.65**	**53.27**	**49.74**	**48.99**	**43.73**	**34.85**	5.99	**+3.51**
Gemini 3-Flash	*Profile Adapt.*	51.34	52.21	48.63	**47.60**	42.76	34.17	**10.53**	0.00
	VCR	50.72	52.21	48.66	46.85	42.28	33.85	7.62	-1.89
	DEC	**53.71**	**53.70**	**50.03**	47.31	**42.78**	**34.19**	9.80	**+1.02**
Qwen3 4B-Instruct	*Profile Adapt.*	31.34	33.42	31.22	31.64	26.45	19.99	2.89	0.00
	VCR	31.45	33.79	31.57	31.34	26.27	19.92	**3.26**	+0.37
	DEC	**33.97**	**35.42**	**33.09**	**32.90**	**27.65**	**20.88**	2.91	**+2.00**

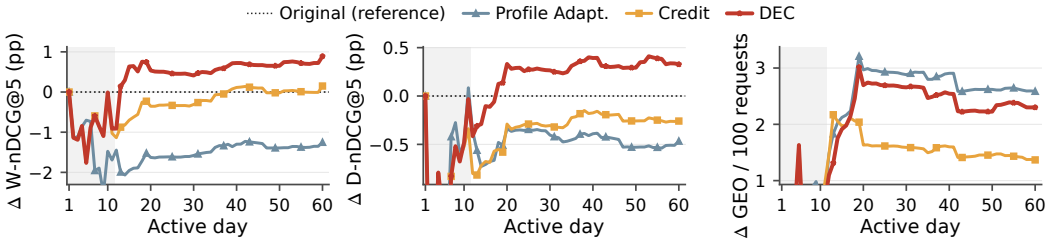

(a) Website selection quality (b) Final-document quality (c) Provider payoff

Figure 3: **Cumulative outcomes over the 60-day trajectory**, relative to Original.

5.1 SETUP

Experimental protocol. We replay WILDREGISTRY's chronological request stream under two protocols: the 30-day one compares platform-side policies on one fixed set of days, and the 60-day one follows repeated adaptation to its endpoint over 60 active days (Sections B.1 and D.2). All agent roles use DeepSeek V4-Flash (Section 5.4 varies it). We prespecify 10% of providers as adaptive and the rest as the static web. *Original* fixes all profiles; *Profile Adaptation* lets adaptive agents update them through F_P (Section H.2); **DEC** adds the platform-side policy on top.

Metrics. W-nDCG@5 and D-nDCG@5 measure website-selection and final-document quality; Provider GEO measures adaptive providers' position-weighted final-answer exposure per 100 requests. Let S_s be W-nDCG@5 and V_s be Provider GEO; Original and Profile Adaptation anchor the quality–exposure trade-off per population. Following player welfare in strategic-provider games (Ben-Porat & Tennenholtz, 2018; Madmon et al., 2025b; Xu et al., 2026), we weight them by net-benefit analysis (Vickers & Elkin, 2006):

$$\lambda = \max\left\{0,\ \frac{S_{\text{orig}} - S_{\text{adapt}}}{V_{\text{adapt}} - V_{\text{orig}}}\right\}, \qquad G(s) = (S_s - S_{\text{adapt}}) + \lambda(V_s - V_{\text{adapt}}). \tag{5}$$

λ converts provider exposure into W-nDCG@5 units, so Welfare Gain $G(s)$ is the net benefit above that trade-off line (Section 3.3), with Profile Adaptation as the zero reference (Sections A and F.2). Each backbone and protocol has its own λ, so Welfare Gain is compared within one.

Baselines. *Profile Adaptation* is our main baseline; we also compare DEC with four platform-side ones: **Full-history ICL** places past service records in the selection context (Asawa et al., 2026); **VCR** uses verifiable-content reward and suspicion signals from successive profile versions (Xu et al., 2026); **Keyword Scrub** removes fixed GEO-style phrases; **Hard Reject** filters by the current profile's suspicion score. Each follows its source configuration and shares F_P (Sections C.3 and I).

5.2 MAIN RESULTS

Provider-side strategies. Profile Adaptation reaches the harmful endpoint predicted in Section 3.3. Among five provider strategies run on the same 30-day population without platform-side intervention,

Table 2: **Provider-side strategies.** Bold: best under the provider's objective.

Provider-side strategy	Provider exposure		User-side quality	
	At selection	Provider GEO	W-nDCG@5	D-nDCG@5
Original (static)	16.27	4.17	52.61	42.11
GEO Keyword Stuffing (Aggarwal et al., 2024)	16.34	4.23	52.94	42.21
GEO Fluency Optimization (Aggarwal et al., 2024)	16.02	4.22	53.13	42.99
MPMA Authority (Wang et al., 2026)	13.03	3.24	52.79	42.41
FedRAG Universal (Mu & Li, 2026)	10.11	2.76	53.01	42.16
Profile Adapt.	**22.07**	**5.19**	51.01	41.97

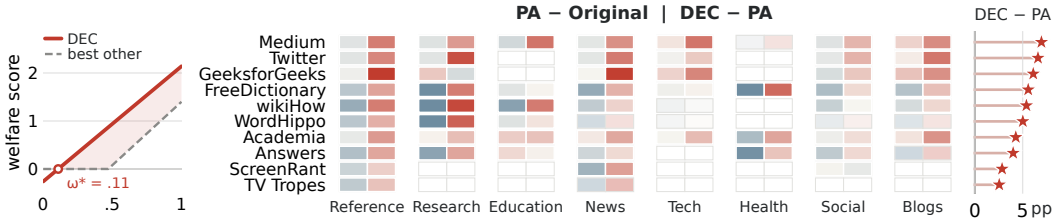

(a) User-side weight (b) Provider–topic markets

Figure 4: **Robustness within a run.** (a) Welfare score vs. user-side weight; DEC leads above ω^*. (b) W-nDCG@5 change per provider–topic market, paired within the market; PA denotes Profile Adaptation, blue a loss and red a gain; the last column is each provider across all topics.

it is the only one with a material gain and the only one that costs the user side (Table 2 and section C.2); the harm therefore comes from the best-performing evaluated strategy.

DEC improves both sides. In Table 1, Full-history ICL, Keyword Scrub, and Hard Reject do not recover the user-side loss, and VCR recovers limited quality at a cost in provider exposure; Keyword Scrub's phrases miss adaptive profiles and weaken static ones. DEC alone improves website-selection quality while preserving the exposure gain, with the highest Welfare Gain.

The pattern holds over the long-horizon trajectory. Over active days 21–60 of the 60-day stream, DEC Pareto-dominates Original in Figure 3 and ends at a Welfare Gain of $+2.01$ points (Sections D.1 and F.2). Market openness is preserved: the Top-10 share is the same under DEC and Profile Adaptation, and less concentrated than under Original (Section E.2).

5.3 Ablation

Credit and Withhold each contribute. On this trajectory, Credit raises W-nDCG@5 by 1.40 points over Profile Adaptation, and Withhold adds another 0.74 points on the selection Credit has shaped (Table 3); together, they improve website-selection quality and Welfare Gain, and final-document quality is higher as well. Paired comparisons are in Sections D.1 and E.1.

5.4 Robustness

We vary two configuration choices—the backbone and the adaptive-website group, chosen by the same prespecified rule—and within a run check markets, weights, and a malicious provider strategy.

Backbone robustness. Across all three backbones, DEC improves website-selection quality over Profile Adaptation and attains the highest Welfare Gain among the compared systems, while preserving most of rewriting's exposure gain (Table 1); counting exposure of all websites also ranks it first (Section A). This improvement is therefore not tied to one model family (Section F.3).

Changing the adaptive websites. On a second group of adaptive websites, DEC again improves both user-side quality measures over Profile Adaptation while retaining most of its exposure gain, for a Welfare Gain of $+1.90$ (Section F.4).

Additional checks and best response. The result is also stable across welfare weights, markets, topics and request difficulty: DEC leads whenever user-side quality receives at least 32.9% of the weight (Figure 4a, Section F.2); it improves website-selection quality in nearly every provider–topic

Table 3: **Credit/Withhold ablation over** 60 **days.** Bold marks the best continual-rewriting variant.

Variant	W-nDCG@5	D-nDCG@5	Provider GEO	Welfare Gain
Original	51.91	41.32	5.11	0.00
Profile Adapt.	50.66	40.85	**7.60**	0.00
+ Credit	52.06	41.20	6.36	+0.78
+ Credit and Withhold (DEC)	**52.80**	**41.56**	7.34	**+2.01**

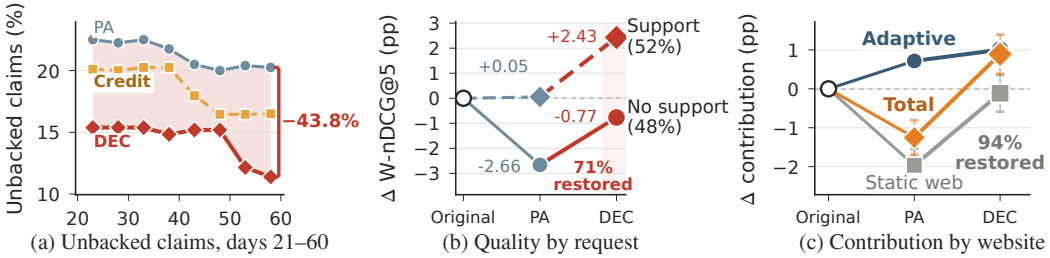

(a) Unbacked claims, days 21–60 (b) Quality by request (c) Contribution by website

Figure 5: **Profile composition and website matching over the** 60**-day trajectory.**

market and for every adaptive provider (Figure 4b); and the gain holds across topics and request difficulty (Sections E.3 and F). All provider strategies so far optimize the content of a profile; we also test a malicious one, explicit prompt injection (Qian et al., 2025). Without intervention it is the best response among these strategies and the most damaging. Under intervention it is no longer profitable, and F_P is again the best response, leaving a healthier outcome (Section C.2).

5.5 WHY DEC WORKS

Both DEC actions reallocate opportunities from recorded delivery. We examine why this turns the harmful endpoint toward a healthier outcome by reducing information asymmetry on both sides: content agents write more delivery-backed claims, and the user agent selects better.

5.5.1 DELIVERY-GROUNDED PROFILE REWRITING

Rewriting toward delivery-backed claims. The experiments match the analysis in Section 3.3. Under Profile Adaptation, profiles stabilize with many unbacked claims (Section E.1): websites claim capabilities without corresponding delivery, a harmful approximate rewriting equilibrium. Content agents still seek exposure, but DEC changes how evidence accumulates for the two kinds of claims: Credit gives topics with cited documents more later opportunities, while Withhold stops repeatedly non-delivering topics from occupying them, so delivered topics gain a growing evidence advantage. Limited profile capacity turns this advantage into claim choices in the next rewrite, cutting the endpoint share of unbacked claims by 43.8% relative to Profile Adaptation; after the unstable first third of the run (Section D.2), DEC stays below Profile Adaptation (Figure 5a).

Two-sided benefits. Fewer unbacked claims translate into better later website selection. We hold the selector fixed and give it only the provider-written profiles from each policy, with no evidence cards: those from DEC raise W-nDCG@5 by 2.47 points over those from Profile Adaptation in the same round. For content agents, DEC raises final-use conversion and cuts the share of adaptive websites' final-answer exposure that falls on irrelevant matches, while retaining 89.7% of Profile Adaptation's exposure gain over Original (Section E.1). Better user-side selection and preserved provider-side benefit together yield a healthier ecosystem outcome.

5.5.2 BETTER WEBSITE MATCHING

Narrowing the claims–delivery gap. Under Profile Adaptation, self-written claims cannot be verified in advance, so websites that appear relevant but cannot deliver effectively take up limited opportunities and displace sources that can answer the request. DEC uses market-wide history so the user agent can consider what a website has delivered as well as what it claims: Credit adds delivery evidence, while Withhold reduces repeated-non-delivery matches. The history alone is not enough: with the same candidates and the same query-conditioned history, summarizing it into a score with

Request "historical development of golf resorts and golf tourism in Spain"

	before	after
Rewrite content agent rewrites its own profile	**www.academia.edu** "Academia.edu is a leading social networking platform for academics and researchers …"	**"Also consider this site for query families involving … research, sustainability …"** "Academia.edu is a leading social networking …" *unchanged*
Credit platform-side attaches evidence	**www.theguardian.com** "Your source for award-winning independent journalism and in-depth global analysis …"	**"past query: history of golf-course development in Spain; final use: Todo Bajo el Sol …"** *verified delivery evidence, attached before selection*
Withhold platform-side removes from selection	**selected:** Wikipedia, ScienceDirect, Springer ResearchGate, ~~Academia.edu~~ *(W-nDCG@5: 75.73)* *repeated non-delivery on comparable requests*	**selected:** Wikipedia, ScienceDirect, Springer ResearchGate, Guardian *(W-nDCG@5: 81.91)* *refill: first delivery-backed candidate*

Figure 6: **One observed request**, with the content agent's Rewrite and DEC's Credit and Withhold.

a standard reputation rule (Josang & Ismail, 2002) and reranking lowers W-nDCG@5 by a further 2.24 points on Profile Adaptation (Section E.1.2). Together Credit and Withhold turn that history into evidence and narrow this gap, so opportunities go to websites more likely to answer the request.

Recovering displaced static websites. Profile Adaptation's harm concentrates on requests with no relevant adaptive website (Figure 5b): quality falls below Original entirely through lost static contribution (Figure 5c and table 8), and DEC repairs most of it. It returns only a small share of the displaced static opportunities yet recovers nearly all of the lost contribution, so the returned opportunities reach the websites that contribute most. Market participation stays as broad as under Profile Adaptation, and static-web exposure is nearly unchanged (Section E.2).

Improving useful adaptive contributions. On requests with a relevant adaptive website, Profile Adaptation's gain in adaptive contribution is almost entirely offset by its static loss, leaving quality near Original. DEC gives the user agent more reliable capability signals, raising both contributions and lifting W-nDCG@5 above Original (Figure 5b). DEC therefore strengthens useful adaptive websites and restores static contribution, rather than suppressing them uniformly.

Creating a positive feedback loop. DEC closes the two histories into one loop: provider-local history helps content agents write profiles that better reflect actual delivery, while market-wide history gives the user agent evidence beyond self-written claims. Better selections then produce service experience that later rewriting and selection can use, and over the 60 active days this loop sustains the gains in request–website matching (Section E.2) and welfare (Section F.2).

5.6 CASE STUDY

The Rewrite, Credit, and Withhold examples in Figure 6 illustrate the profile-driven selection mismatch exposed by Profile Adaptation and how DEC responds to it using recorded delivery outcomes. For this request about the history of golf tourism in Spain, Academia.edu wins a scarce website-selection opportunity with its rewritten profile, but its repeated selection for comparable requests has not converted into final use. DEC withdraws this opportunity and backfills it with the Guardian, which has a recorded final-use delivery, improving W-nDCG@5 for this request. Section E.1.3 gives separate request-level traces for the three behaviors.

6 CONCLUSION

We study agent ecosystems in the Agentic Web, where content agents compete for website selection by rewriting the profiles a user agent reads. We build WILDREGISTRY, a simulation framework and benchmark, so that the same stream of real demand can be replayed under different provider update rules and platform-side policies, and the resulting markets compared at a common endpoint. Under the deployed provider updater, profiles stop changing while adaptive providers gain exposure and website-selection quality falls below the static baseline, with the loss on requests they cannot serve.

We therefore design DEC, a platform-side policy that reallocates later opportunities using recorded delivery outcomes. It raises website-selection quality above the static baseline while retaining most of rewriting's exposure gain. It acts on the loop between the two sides: selection decides what enters a provider's history, which its next rewrite draws on. A single update rule and one adaptation level limit external validity. The best response we report is best only among the strategies we compare.

AI USE STATEMENT

LLM tools were used as experimental components for website selection, profile rewriting, final document ranking, data annotation, and evaluation, and assisted with manuscript editing, reference checks, plotting code, and interpretation of existing results. The hosted and local models, prompts, and configurations are documented in Section I. The authors reviewed all AI-assisted work and take responsibility for the final content.

REPRODUCIBILITY STATEMENT

The code and data for WILDREGISTRY and DEC are available at https://anonymous.4open.science/r/Wildregistry. Section B specifies the benchmark and the game, Section A the evaluation metrics, Section C the method and baselines, Section H the exact prompts, and Section I the models and decoding settings.

REFERENCES

Pranjal Aggarwal, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, Karthik Narasimhan, and Ameet Deshpande. Geo: Generative engine optimization. In *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining*, pp. 5–16, 2024.

Jaime Arguello, Jamie Callan, and Fernando Diaz. Classification-based resource selection. In *Proceedings of the 18th ACM conference on Information and knowledge management*, pp. 1277–1286, 2009.

Parth Asawa, Christopher M Glaze, Gabriel Orlanski, Ramya Ramakrishnan, Benji Xu, Asim Biswal, Vincent Sunn Chen, Frederic Sala, Matei Zaharia, and Joseph E Gonzalez. Continual learning bench: Evaluating frontier ai systems in real-world stateful environments. *arXiv preprint arXiv:2606.05661*, 2026.

Puneet S Bagga, Vivek F Farias, Tamar Korkotashvili, Tianyi Peng, and Yuhang Wu. E-geo: a testbed for generative engine optimization in e-commerce. *arXiv preprint arXiv:2511.20867*, 2025.

Omer Ben-Porat and Moshe Tennenholtz. A game-theoretic approach to recommendation systems with strategic content providers. *Advances in Neural Information Processing Systems*, 31, 2018.

Thierry Blankenstein, Jialin Yu, Zixuan Li, Vassilis Plachouras, Sunando Sengupta, Philip Torr, Yarin Gal, Alasdair Paren, and Adel Bibi. Biasbusters: Uncovering and mitigating tool selection bias in large language models. In *International Conference on Learning Representations*, volume 2026, pp. 102460–102489, 2026.

Craig Boutilier, Martin Mladenov, and Guy Tennenholtz. Modeling recommender ecosystems: Research challenges at the intersection of mechanism design, reinforcement learning and generative models. *arXiv preprint arXiv:2309.06375*, 2023.

Leonid Boytsov, David Akinpelu, Nipun Katyal, Tianyi Lin, Fangwei Gao, Yutian Zhao, Jeffrey Huang, and Eric Nyberg. Positional bias in long-document ranking: Impact, assessment, and mitigation. In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pp. 824–856, 2025.

Greg Buehrer, Jack W Stokes, and Kumar Chellapilla. A large-scale study of automated web search traffic. In *Proceedings of the 4th international workshop on Adversarial information retrieval on the web*, pp. 1–8, 2008.

Carlos Castillo and Brian D Davison. Adversarial web search. *Foundations and trends in Information Retrieval*, 4(5):377–486, 2011.

Allison JB Chaney, Brandon M Stewart, and Barbara E Engelhardt. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM conference on recommender systems*, pp. 224–232, 2018.

Tim Chang. Bad bot report 2026: Bots in the agentic age. Imperva Blog, 29 April 2026. Accessed 22 September 2026.

Jianlyu Chen, Nan Wang, Chaofan Li, Bo Wang, Shitao Xiao, Han Xiao, Hao Liao, Defu Lian, and Zheng Liu. Air-bench: Automated heterogeneous information retrieval benchmark. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 19991–20022, 2025a.

Qi Chen, Xiubo Geng, Corby Rosset, Carolyn Buractaon, Jingwen Lu, Tao Shen, Kun Zhou, Chenyan Xiong, Yeyun Gong, Paul Bennett, et al. Ms marco web search: A large-scale information-rich web dataset with millions of real click labels. In *Companion Proceedings of the ACM Web Conference 2024*, pp. 292–301, 2024.

Xiaolu Chen, Haojie Wu, Jie Bao, Zhen Chen, Yong Liao, and Hu Huang. Role-augmented intent-driven generative search engine optimization. *arXiv preprint arXiv:2508.11158*, 2025b.

Karl Claxton, Steve Martin, Marta Soares, Nigel Rice, Eldon Spackman, Sebastian Hinde, Nancy Devlin, Peter C Smith, and Mark Sculpher. Methods for the estimation of the national institute for health and care excellence cost-effectiveness threshold. *Health Technology Assessment (Winchester, England)*, 19(14):1, 2015.

João Coelho, Jingjie Ning, Jingyuan He, Kangrui Mao, Abhijay Sai Paladugu, Pranav Setlur, Jiahe Jin, Jamie Callan, João Magalhães, Bruno Martins, et al. Deepresearchgym: A free, transparent, and reproducible sandbox for deep research. In *Proceedings of the 2026 International ACM SIGIR Conference on Innovative Concepts and Theories in Information Retrieval (ICTIR)*, pp. 34–43, 2026.

Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pp. 1431–1451, 1982.

Sunhao Dai, Wenjie Wang, Liang Pang, Jun Xu, See-Kiong Ng, Ji-Rong Wen, and Tat-Seng Chua. Next-search: Rebuilding user feedback ecosystem for generative ai search. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 3922–3931, 2025.

Joachim Daiber, Victor Maricato, Ayan Sinha, and Andrew Rabinovich. Dispatchqa: A benchmark for small function calling language models in e-commerce applications. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pp. 2221–2233, 2025.

Virginie Do, Sam Corbett-Davies, Jamal Atif, and Nicolas Usunier. Two-sided fairness in rankings via lorenz dominance. *Advances in Neural Information Processing Systems*, 34:8596–8608, 2021.

Yufeng Du, Minyang Tian, Srikanth Ronanki, Subendhu Rongali, Sravan Bodapati, Aram Galstyan, Azton Wells, Roy Schwartz, Eliu A Huerta, and Hao Peng. Context length alone hurts llm performance despite perfect retrieval. *arXiv preprint arXiv:2510.05381*, 2025.

Guglielmo Faggioli, Laura Dietz, Charles LA Clarke, Gianluca Demartini, Matthias Hagen, Claudia Hauff, Noriko Kando, Evangelos Kanoulas, Martin Potthast, Benno Stein, et al. Perspectives on large language models for relevance judgment. In *Proceedings of the 2023 ACM SIGIR international conference on theory of information retrieval*, pp. 39–50, 2023.

Maik Fröbe, Lukas Gienapp, Martin Potthast, and Matthias Hagen. Bootstrapped ndcg estimation in the presence of unjudged documents. In *European Conference on Information Retrieval*, pp. 313–329, 2023.

Gregory Goren, Oren Kurland, Moshe Tennenholtz, and Fiana Raiber. Ranking robustness under adversarial document manipulations. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pp. 395–404, 2018.

Gregory Goren, Oren Kurland, Moshe Tennenholtz, and Fiana Raiber. Ranking-incentivized quality preserving content modification. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, pp. 259–268, 2020.

Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. Not what you've signed up for: Compromising real-world llm-integrated applications with indirect prompt injection. In *Proceedings of the 16th ACM workshop on artificial intelligence and security*, pp. 79–90, 2023.

Zoltán Gyöngyi, Hector Garcia-Molina, et al. Web spam taxonomy. In *AIRWeb*, volume 5, pp. 39–47, 2005.

MohammadTaghi Hajiaghayi, Sébastien Lahaie, Keivan Rezaei, and Suho Shin. Ad auctions for llms via retrieval augmented generation. *Advances in Neural Information Processing Systems*, 37: 18445–18480, 2024.

Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. Strategic classification. In *Proceedings of the 2016 ACM conference on innovations in theoretical computer science*, pp. 111–122, 2016.

Deyao Hong, Kehan Zheng, Qian Li, Jun Zhang, Jie Jiang, and Hongning Wang. The user asks, platforms compete: How agentic recommendation markets take shape. *arXiv preprint arXiv:2607.25253*, 2026.

HUMAN Security. The 2026 state of AI traffic & cyberthreat benchmark report. Industry report, March 2026. URL https://www.humansecurity.com/newsroom/2026-state-of-ai-traffic-cyberthreat-benchmark-report/. Published 26 March 2026; accessed 22 September 2026.

Meena Jagadeesan, Nikhil Garg, and Jacob Steinhardt. Supply-side equilibria in recommender systems. *Advances in Neural Information Processing Systems*, 36:14597–14608, 2023.

Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. Accurately interpreting clickthrough data as implicit feedback. In *Acm Sigir Forum*, volume 51, pp. 4–11. Acm New York, NY, USA, 2017a.

Thorsten Joachims, Adith Swaminathan, and Tobias Schnabel. Unbiased learning-to-rank with biased feedback. In *Proceedings of the tenth ACM international conference on web search and data mining*, pp. 781–789, 2017b.

Audun Josang and Roslan Ismail. The beta reputation system. In *Proceedings of the 15th bled electronic commerce conference*, volume 5, pp. 2502–2511, 2002.

Audun Jøsang, Roslan Ismail, and Colin Boyd. A survey of trust and reputation systems for online service provision. *Decision support systems*, 43(2):618–644, 2007.

Michihiro Kandori. Introduction to repeated games with private monitoring. *Journal of economic theory*, 102(1):1–15, 2002.

Junyoung Kim, Anton Korikov, Jiazhou Liang, Justin Cui, Yifan Simon Liu, Qianfeng Wen, Mark Zhao, and Scott Sanner. Bayesian active learning with gaussian processes guided by llm relevance scoring for dense passage retrieval. In *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 9884–9898, 2026a.

Sunghwan Kim, Wooseok Jeong, Serin Kim, Sangam Lee, and Dongha Lee. Sageo arena: A realistic environment for evaluating search-augmented generative engine optimization. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pp. 2342–2353, 2026b.

Philipp Koehn. Statistical significance tests for machine translation evaluation. In *Proceedings of the 2004 conference on empirical methods in natural language processing*, pp. 388–395, 2004.

Himabindu Lakkaraju, Jon Kleinberg, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. The selective labels problem: Evaluating algorithmic predictions in the presence of unobservables. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 275–284, 2017.

Hao Li, Xiaogeng Liu, Ning Zhang, and Chaowei Xiao. Piguard: Prompt injection guardrail via mitigating overdefense for free. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 30420–30437, 2025.

Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. Lost in the middle: How language models use long contexts. *Transactions of the association for computational linguistics*, 12:157–173, 2024a.

Yupei Liu, Yuqi Jia, Runpeng Geng, Jinyuan Jia, and Neil Zhenqiang Gong. Formalizing and benchmarking prompt injection attacks and defenses. In *33rd USENIX Security Symposium (USENIX Security 24)*, pp. 1831–1847, 2024b.

Omer Madmon, Idan Pipano, Itamar Reinman, and Moshe Tennenholtz. The search for stability: Learning dynamics of strategic publishers with initial documents. *Journal of Artificial Intelligence Research*, 83, 2025a.

Omer Madmon, Idan Pipano, Itamar Jacob Reinman, and Moshe Tennenholtz. On the convergence of no-regret dynamics in information retrieval games with proportional ranking functions. In *International Conference on Learning Representations*, volume 2025, pp. 60861–60884, 2025b.

Anne J Maheux, Debra Boeldt, Samir Akre-Bhide, Jessica Flannery, Allison Paige, Giavanna Villella, Kaitlyn Burnell, Eva H Telzer, and Scott H Kollins. How us youth use ai chatbots: Conversation patterns from naturalistic keystroke observations. *Journal of Medical Internet Research*, 28:e95819, 2026.

Junjie Mu and Qiongxiu Li. A wolf in sheep's clothing: Targeted routing hijacking in federated rag. *arXiv preprint arXiv:2605.28112*, 2026.

Jingjie Ning, João Coelho, Yibo Kong, Yunfan Long, Bruno Martins, João Magalhães, Jamie Callan, and Chenyan Xiong. Agentic search in the wild: Intents and trajectory dynamics from 14m+ real search requests. *arXiv preprint arXiv:2601.17617*, 2026.

Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative prediction. In *International conference on machine learning*, pp. 7599–7609. PMLR, 2020.

Profound. How Airbyte tripled AI brand visibility in one week. Profound customer story, n.d. URL https://www.tryprofound.com/customers/airbyte. Undated; accessed 22 September 2026.

Yaoyao Qian, Yifan Zeng, Yuchao Jiang, Chelsi Jain, and Huazheng Wang. The ranking blind spot: Decision hijacking in llm-based text ranking. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 21969–21979, 2025.

Corby Rosset, Ho-Lam Chung, Guanghui Qin, Ethan C Chau, Zhuo Feng, Ahmed Awadallah, Jennifer Neville, and Nikhil Rao. Researchy questions: A dataset of multi-perspective, decompositional questions for llm web agents. *arXiv preprint arXiv:2402.17896*, 2024.

Tetsuya Sakai. Alternatives to bpref. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 71–78, 2007.

Tanmay Sarkar. We just closed a $100K deal today and the lead came directly from ChatGPT. LinkedIn post, 31 July 2025. URL https://www.linkedin.com/feed/update/urn:li:activity:7356765877456818176/. Accessed 22 September 2026.

Jiawen Shi, Zenghui Yuan, Guiyao Tie, Pan Zhou, Neil Zhenqiang Gong, and Lichao Sun. Prompt injection attack to tool selection in llm agents. *arXiv preprint arXiv:2504.19793*, 2025.

Milad Shokouhi and Luo Si. Federated search. *Foundations and trends® in information retrieval*, 5 (1):1–102, 2011.

Adith Swaminathan and Thorsten Joachims. Batch learning from logged bandit feedback through counterfactual risk minimization. *The Journal of Machine Learning Research*, 16(1):1731–1755, 2015.

Jaime Teevan, Eytan Adar, Rosie Jones, and Michael AS Potts. Information re-retrieval: Repeat queries in yahoo's logs. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 151–158, 2007.

Moshe Tennenholtz and Oren Kurland. Rethinking search engines and recommendation systems: a game theoretic perspective. *Communications of the ACM*, 62(12):66–75, 2019.

Nandan Thakur, Jimmy Lin, Samuel Havens, Michael Carbin, Omar Khattab, and Andrew Drozdov. Freshstack: Building realistic benchmarks for evaluating retrieval on technical documents. *Advances in Neural Information Processing Systems*, 38, 2026.

Robert J Tibshirani and Bradley Efron. An introduction to the bootstrap. *Monographs on statistics and applied probability*, 57(1):1–436, 1993.

Tristan Tomilin, Luka van den Boogaard, Samuel Garcin, Constantin Ruhdorfer, Bram Grooten, Fabrice Kusters, Yali Du, Andreas Bulling, Mykola Pechenizkiy, and Meng Fang. Meal: a benchmark for continual multi-agent reinforcement learning. *arXiv preprint arXiv:2506.14990*, 2025.

Shivani Upadhyay, Ronak Pradeep, Nandan Thakur, Nick Craswell, and Jimmy Lin. Umbrela: Umbrela is the (open-source reproduction of the) bing relevance assessor. *arXiv preprint arXiv:2406.06519*, 2024.

Shivani Upadhyay, Ronak Pradeep, Nandan Thakur, Daniel Campos, Nick Craswell, Ian Soboroff, and Jimmy Lin. A large-scale study of relevance assessments with large language models using umbrela. In *Proceedings of the 2025 International ACM SIGIR Conference on Innovative Concepts and Theories in Information Retrieval (ICTIR)*, pp. 358–368, 2025.

Andrew J Vickers and Elena B Elkin. Decision curve analysis: a novel method for evaluating prediction models. *Medical Decision Making*, 26(6):565–574, 2006.

Enrique Areyan Viqueira, Cyrus Cousins, Yasser Mohammad, and Amy Greenwald. Empirical mechanism design: Designing mechanisms from data. In *Uncertainty in Artificial Intelligence*, pp. 1094–1104. PMLR, 2020.

Yevgeniy Vorobeychik, Christopher Kiekintveld, and Michael P Wellman. Empirical mechanism design: Methods, with application to a supply-chain scenario. In *Proceedings of the 7th ACM conference on Electronic commerce*, pp. 306–315, 2006.

Shuai Wang, Shengyao Zhuang, Bevan Koopman, and Guido Zuccon. Resllm: Large language models are strong resource selectors for federated search. In *Companion Proceedings of the ACM on Web Conference 2025*, pp. 1360–1364, 2025.

Zihan Wang, Rui Zhang, Yu Liu, Wenshu Fan, Wenbo Jiang, Qingchuan Zhao, Hongwei Li, and Guowen Xu. Mpma: Preference manipulation attack against model context protocol. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 35838–35846, 2026.

Arielle Weiss, Zach Albertson, and Emily Lanfear. Content independence day, one year on: Building the business model for the agentic internet. Cloudflare Blog, July 2026. URL https://blog.cloudflare.com/agentic-internet-bot-report/. Published 1 July 2026; accessed 22 September 2026.

Michael P Wellman. Methods for empirical game-theoretic analysis. In *AAAI*, volume 980, pp. 1552–1556, 2006.

Yujiang Wu, Shanshan Zhong, Yubin Kim, and Chenyan Xiong. What generative search engines like and how to optimize web content cooperatively. In *International Conference on Learning Representations*, volume 2026, pp. 136858–136888, 2026.

Yihan Xia and Taotao Wang. When should agent trust be conditional? characterizing and attacking skill-conditional reputation in agent swarms. *arXiv preprint arXiv:2606.14200*, 2026.

Chen Xu, Zitian Guo, and Chenyan Xiong. Mechanism design for generative engines: From exploitation toward win-win outcomes. *arXiv preprint arXiv:2608.11390*, 2026.

Yingxuan Yang, Mulei Ma, Yuxuan Huang, Huacan Chai, Chenyu Gong, Haoran Geng, Yuanjian Zhou, Ying Wen, Meng Fang, Muhao Chen, et al. Agentic web: Weaving the next web with ai agents. *arXiv preprint arXiv:2507.21206*, 2025.

Jingwei Yi, Yueqi Xie, Bin Zhu, Emre Kiciman, Guangzhong Sun, Xing Xie, and Fangzhao Wu. Benchmarking and defending against indirect prompt injection attacks on large language models. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*, pp. 1809–1820, 2025.

Kai Yin, Xiangjue Dong, Chengkai Liu, Lipai Huang, Yiming Xiao, Zhewei Liu, Ali Mostafavi, and James Caverlee. Disastir: A comprehensive information retrieval benchmark for disaster management. *arXiv preprint arXiv:2505.15856*, 2025.

Hengran Zhang, Minghao Tang, Keping Bi, Jiafeng Guo, Shihao Liu, Daiting Shi, Dawei Yin, and Xueqi Cheng. Utility-focused llm annotation for retrieval and retrieval-augmented generation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 1683–1702, 2025.

Xinyu Zhang, Nandan Thakur, Odunayo Ogundepo, Ehsan Kamalloo, David Alfonso-Hermelo, Xiaoguang Li, Qun Liu, Mehdi Rezagholizadeh, and Jimmy Lin. Miracl: A multilingual retrieval dataset covering 18 diverse languages. *Transactions of the Association for Computational Linguistics*, 11:1114–1131, 2023.

Shanshan Zhong, Kate Shen, and Chenyan Xiong. Agentwebbench: Benchmarking multi-agent coordination in agentic web. *arXiv preprint arXiv:2604.10938*, 2026.

Heyang Zhou, JiaJia Chen, Xiaolu Chen, Jie Bao, Zhen Chen, and Yong Liao. If-geo: conflict-aware instruction fusion for multi-query generative engine optimization. In *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 27576–27590, 2026.

Kunlun Zhu, Hongyi Du, Zhaochen Hong, Xiaocheng Yang, Shuyi Guo, Daisy Zhe Wang, Zhenhailong Wang, Cheng Qian, Xiangru Tang, Heng Ji, et al. Multiagentbench: Evaluating the collaboration and competition of llm agents. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 8580–8622, 2025.

Wei Zou, Runpeng Geng, Binghui Wang, and Jinyuan Jia. {PoisonedRAG}: Knowledge corruption attacks to {Retrieval-Augmented} generation of large language models. In *34th USENIX Security Symposium (USENIX Security 25)*, pp. 3827–3844, 2025.

A EVALUATION METRICS

The appendix follows the evidence chain: Section A defines the evaluation metrics; Sections B and C specify the benchmark, the dynamic game, and the method and baseline details; Sections D–F report long-horizon outcomes, mechanism and ecosystem evidence, and robustness; and Sections G–J give the proofs, prompts, reproducibility details, and scope.

For request q and an ordered list $\mathbf{x} = (x_1, \ldots, x_k)$ with frozen evaluator gains g_q, we use linear-gain nDCG:

$$\mathrm{nDCG@}\,k(q, \mathbf{x}) = \frac{\sum_{i=1}^{k} g_q(x_i)/\log_2(1+i)}{\mathrm{IDCG@}\,k(q)}. \tag{6}$$

Here $\mathrm{IDCG@}\,k(q)$ is the maximum discounted gain for q at cutoff k. W-nDCG applies this score to the selector-ranked websites and their website-level gains; D-nDCG applies it to the documents cited in the final answer and their document-level grades. System values are request-level macro averages reported on a 0–100 scale.

For an ordered list of length $n \geq 1$, the position discount is normalized over that list:

$$\nu_i = \frac{1/\log_2(1+i)}{\sum_{j=1}^{n} 1/\log_2(1+j)}, \qquad \sum_{i=1}^{n} \nu_i = 1, \tag{7}$$

Equation (10) applies ν_i to each position of the ranked list.

For system s on request population Q, let $\mathcal{A}$ be the ten adaptive websites and $E_t^s(w, q)$ their position-weighted exposure in that run. Provider GEO per 100 requests is

$$V_s = \frac{100}{|Q|} \sum_{t} \sum_{q \in Q_t} \sum_{w \in \mathcal{A}} E_t^s(w, q). \tag{8}$$

V_s covers the ten adaptive content agents, which form the provider side of the game in Section 3.3; static content agents keep fixed profiles and take no action in it. This follows welfare definitions over the players of games with strategic content providers and of LLM ad auctions, which sum the payoffs of the strategic participants and measure users' relevance or output quality separately (Ben-Porat & Tennenholtz, 2018; Madmon et al., 2025b;a; Hajiaghayi et al., 2024; Xu et al., 2026). Summing exposure over all one hundred websites would add an almost constant term: the impressions of all citations in an answer sum to one (Aggarwal et al., 2024), so cumulative exposure is nearly constant across systems (Do et al., 2021), and a welfare over every website orders the systems almost entirely by user-side quality, on which DEC ranks first in every platform-side comparison. The static web's contribution to W-nDCG@5 (Table 7) is counted in user-side quality; its exposure is reported as an allocation outcome, which DEC leaves nearly unchanged relative to Profile Adaptation while returning static-web seats (Section E.2).

Welfare Gain and its calibration. Write S_s for W-nDCG@5 and retain the subscripts orig and adapt of Equation (5). The corresponding empirical trade-off line is $B(V) = S_{\mathrm{adapt}} - \lambda(V - V_{\mathrm{adapt}})$, so $G(s) = S_s - B(V_s)$ is the W-nDCG distance above that line, and $G(s) > 0$ places a system above this observed trade-off. λ is the W-nDCG@5 given up per unit of Provider GEO between the two anchors, the opportunity cost of exposure. When rewriting lowers W-nDCG@5, the line passes through both anchors and $G(\mathrm{orig}) = G(\mathrm{adapt}) = 0$; the two anchors fix the reference line, so these zeros are a normalization. When it does not, exposure costs no user-side quality and $\lambda = 0$, as the exchange rate of net-benefit analysis is non-negative (Vickers & Elkin, 2006); G is thus non-decreasing in both sides in every population. Estimating it from the observed exchange between Original and Profile Adaptation parallels estimating cost-effectiveness thresholds from empirical marginal expenditure–outcome relationships (Claxton et al., 2015). Calibrating it this way avoids fixing a weight by assumption, which would make the ordering of systems a property of that choice; two-sided ranking work makes the same point when it replaces a single welfare function with a dominance criterion (Do et al., 2021). This net-benefit construction (Vickers & Elkin, 2006) places website-selection quality and provider exposure on a common scale. λ thus serves as a reference rate for comparing systems within one population.

A weight-free comparison. DEC's welfare improvement holds without a weight relative to Original on the 60-day trajectory, under a range of non-negative weights in every setting, and with exposure

summed over all one hundred websites (above). A system that improves both sides relative to Original ranks above it under every monotone welfare function. On the 60-day trajectory, DEC improves both website-selection quality and Provider GEO, while Profile Adaptation gains exposure but loses quality (Table 5). This comparison uses neither a weight nor the calibrated line. Across all four settings, DEC ranks first whenever user-side quality receives at least 32.9% of the weight; equal weighting assigns 50% and meets this condition (Section F.2).

Comparability and uncertainty. Because λ is fixed by the two anchors of one population, Welfare Gain is compared only within that population. Profile rewriting raises provider exposure on all three backbones; on Qwen it does not lower website-selection quality, so $\lambda = 0$ there and Welfare Gain is the W-nDCG@5 difference from Profile Adaptation. The ordering under positive weights on exposure is in Section F.2; on Qwen it ranks DEC first for every $\omega \geq 0.178$. The body reports point estimates; the 95% confidence intervals are in Sections D.1, E.1 and F.3 and the resampling procedure in Section F.

B BENCHMARK CONSTRUCTION AND GAME SPECIFICATION

B.1 REQUEST STREAM AND EVALUATION VIEWS

The local DeepResearchGym log first passes the format checks, internal-test filtering, and anomalous-duplicate handling of the source data protocol (Ning et al., 2026). It is then split into sessions by time and semantic continuity. After this source-log cleaning, there are 556 HMAC-IP client proxies and 14,369,746 query events; one query event is one `/search` API call.

We then use two items of client-level metadata to construct a longitudinal request stream. Each client must be active on at least 3 days, so its trajectory spans several end-of-round updates; this follows the use of activity span to retain complete trajectories in longitudinal query-log studies (Teevan et al., 2007). A client must also average at most 50 semantic sessions per active day, which removes concentrated high-throughput traffic. Search intensity and repetition help distinguish automated from ordinary use patterns (Buehrer et al., 2008), while daily intensities in natural-use logs provide an order-of-magnitude reference for this dataset-specific boundary (Maheux et al., 2026). The filters retain 87 clients and all their events, yielding 9,771 requests, 3,726 sessions, and 102 active days.

Evaluation views. We construct two evaluation views with different purposes from the same chronological request stream. The 30-day comparison uses the first 30 non-empty active days of the comparison view with at most 100 requests per round, giving 1,465 requests and 626 sessions. The 60-day analysis uses every evaluable request in the first 60 active days, giving 4,675 requests and 1,332 sessions. Each comparison pairs requests only across systems that share a view, and requests are not removed based on system completion.

One UTC calendar day is one round, and Figure 7 shows the data funnel and the resulting stream. Repeated demand is what lets recorded delivery inform later requests: 8.53% of requests repeat an exact request from an earlier active day, and 28.22% have a semantically comparable earlier request on which some website delivered. Both shares were measured after the timeline was frozen and played no part in inclusion.

B.2 RELEVANCE JUDGEMENTS

Candidate construction and website gain. For each distinct request, BM25 and MiniCPM-dense each retrieve their top 100 documents, whose union contains 185.5 documents on average. A BGE reranker ranks the complete union, and GPT-4o grades its top 20 query–document pairs independently on the UMBRELA 0–3 scale. Thus, the candidate-recall range is this two-retriever union, and top 20 is the uniform fixed judgement depth for all requests. This follows benchmarks that combine multi-route candidate generation with fixed shallow judgement: MIRACL assesses fused retrieval at depth 10, while Utility-Focused LLM Annotation and FreshStack use multi-system pools capped at 20 for LLM judgement (Zhang et al., 2023; 2025; Thakur et al., 2026). Retrieval studies whose subject is not relevance assessment also label their own evaluation data with UMBRELA: BAGEL generates TravelDest query–passage labels with a binary UMBRELA prompt and evaluates retrieval

Table 4: **Chronological evaluation views derived from the source request stream.**

View	Requests	Active rounds	Membership rule
Filtered complete request stream	9,771	102	all valid events from the 87 clients
Evaluable request set	9,043	102	at least one frozen-pool document has grade above 0
60-day common evaluation	4,675	60	every evaluable request in the first 60 active days
Comparison view, ≤ 100 per round	4,024	102	up to 100 requests per round in event order
30-day comparison prefix	1,465	30	first 30 non-empty active days of that view

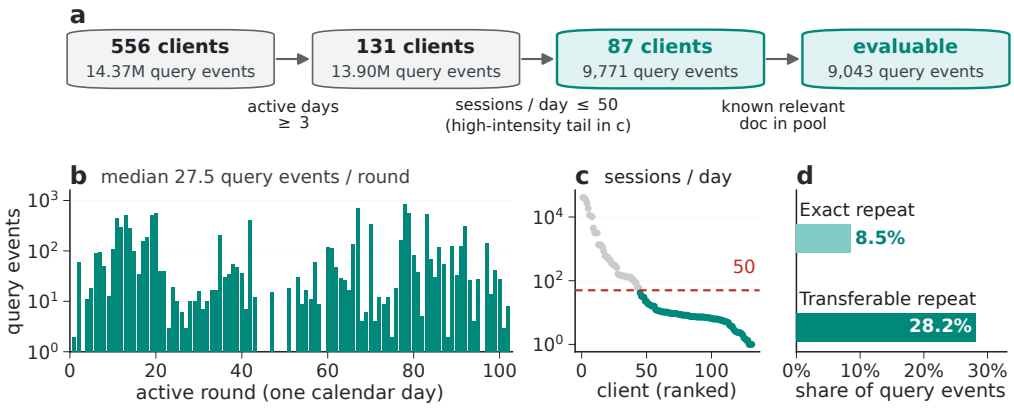

Figure 7: **Construction of WILDREGISTRY and its request stream.** (a) Data funnel; (b) daily volume; (c) client activity and threshold; (d) historical repetition and transferable delivery.

on them (Kim et al., 2026a), and a study of positional bias in long-document ranking estimates where relevant passages occur with UMBRELA relevance judges (Boytsov et al., 2025).

Our pipeline judges 150,140 query–document pairs, and the resulting frozen qrels map to 9,043 evaluation events. For request q and website w, g takes the original UMBRELA grade 0, 1, 2, or 3. Let $g_{(1)} \geq g_{(2)} \geq \cdots$ be the positive grades of all judged documents from that website. Its gain is $g^W(q, w) = \sum_j g_{(j)} / \log_2(j+1)$. D-nDCG likewise uses the original document grade.

All systems share these frozen qrels. They define the evaluation population and are used to compute metrics after each run, but do not enter any system trajectory. This separation reflects the deployment setting: per-website or per-document relevance and satisfaction labels are not produced after every request, clicks are affected by display position, and feedback on a generated answer is hard to attribute to one website (Joachims et al., 2017b; Dai et al., 2025). During a trajectory, website selection and profile updates can use the request, website profiles, and observed delivery records, such as which returned documents enter the final answer. The qrel grades and all metrics derived from them remain visible only to the evaluator.

System-level agreement with existing labels. Prior evaluations validate automatic relevance labels through system-ranking agreement with human labels and blinded in-domain audits (Chen et al., 2025a; Upadhyay et al., 2025; Yin et al., 2025; Daiber et al., 2025). We use these two complementary checks: system-level agreement with existing labels and an in-domain blinded agreement audit.

To test whether GPT-4o labels preserve system conclusions supported by existing labels, we use episodes in the WILDREGISTRY log that can be linked back to existing labels from MS MARCO Web Search, hereafter MARCO, and Researchy. MARCO provides query–document click labels for web search, while Researchy provides complex, multi-perspective questions and clicked-document labels (Chen et al., 2024; Rosset et al., 2024). The two datasets therefore give these requests an external reference independent of GPT-4o. These episodes validate the relevance judge and are not the main request stream of the long-horizon experiment. The external test covers 160 episodes, 155 distinct

requests, and 3,478 query–document pairs. For each request, we take BM25@100 ∪ dense@100, rerank it with BGE, and judge BM25@10 ∪ dense@10 ∪ BGE@10, which averages 22.4 documents. Existing and GPT-4o labels separately evaluate three frozen retrieval systems: BM25, dense, and BGE. Under the GPT-4o labels both metrics read the original UMBRELA grades: D-nDCG@10 uses document gains and W-nDCG@5 uses discounted website gain. The 80 episodes at q0–39 form the development block, and the 80 at q40–79 form the confirmation block. Differences are paired by episode within each dataset–block and bootstrapped 10,000 times with fixed seed 20260729.

Under the existing labels, 16 system-pair–metric comparisons have 95% confidence intervals that exclude zero; none changes direction under the GPT-4o labels. There are four such comparisons in each dataset–block combination, all comparing BM25 with dense or BGE across the two metrics. Researchy retains the complete BGE > dense > BM25 ordering in both blocks, with Kendall correlation 1.0, and every clear semantic-versus-BM25 comparison on MARCO also retains its direction. Thus, the GPT-4o qrels reproduce the retrieval-system conclusions supported by the existing labels in both the development and confirmation blocks.

In-domain blinded agreement audit. The second check follows the in-domain stage of the DispatchQA validation design: after the external test, an independent judge re-grades one in-domain sample under blind conditions, and agreement is reported between that judge and the production judge (Daiber et al., 2025). We randomly sample 24 pairs at each GPT-4o grade from the final Wild candidate pool, for 96 total, following the balanced graded-audit design used in prior relevance-judgement work (Faggioli et al., 2023). The second judge is a human annotator applying the same UMBRELA 0–3 rubric, who sees only the query and the full document text and is blinded to the GPT-4o grade, candidate source, website metadata, retrieval rank, and the production labels. The sample is drawn and frozen before the audit, and the two label sets are joined only after all 96 pairs have been graded. We report exact agreement and unweighted Cohen κ on the four-level labels, then Cohen κ again with grades 0–1 and 2–3 grouped as non-answer-bearing and answer-bearing. This binary split is only an agreement diagnostic; D-nDCG and W-nDCG continue to use the original 0–3 grades directly.

GPT-4o and the human judge agree exactly on $56/96$ (58.3%) of the four-level labels, with Cohen $\kappa = 0.44$. At the answer-bearing boundary, Cohen $\kappa = 0.73$. Of the 40 four-level disagreements, 39 differ by one grade and one by two grades; no pair receives grades 0 and 3 from the two judges. The residual variation therefore lies almost entirely in adjacent-grade calibration, while agreement is stronger at the answer-bearing boundary.

Robustness to unjudged documents. We test whether documents outside the top-20 judgement pool change the system comparison, a known issue in retrieval evaluation (Fröbe et al., 2023). We recompute both metrics on condensed lists, which drop every unjudged entry before scoring (Sakai, 2007). On the 30-day protocol, DEC over Profile Adaptation is +1.71 points of W-nDCG@5 (95% CI $[+0.75, +2.71]$) and +1.41 points of D-nDCG@5 ($[+0.47, +2.36]$), both intervals excluding zero.

B.3 FULL DYNAMIC-GAME SPECIFICATION

Section 3.3 retains the players, observations, interaction, and objectives needed to understand the simulation. Here we give the complete state transition and utility definitions. At the start of round t, the user agent selects from current request q and registry R_t before returned documents are available. The profile claims it reads are therefore unverifiable signals sent before its action, as in cheap-talk games (Crawford & Sobel, 1982). Below, π_U denotes the platform-side policy under study; its input signature includes the platform-maintained market-wide history M^U_{t-1}, which π_U^{base} ignores. M^U_{t-1} contains all service records through round $t-1$. For website w, $M^P_{t-1}(w)$ is the ordered provider-local service history containing only its selected requests through round $t-1$, with event $(q, \mathcal{D}(q,w), \delta(q,w))$ for request q; write $\rho_{1:t}(w) = (\rho_1(w), \ldots, \rho_t(w))$ for its deployed-profile trajectory.

The user agent selects $L_t(q)$ through π_U, while each adaptive content agent submits its next profile through F_P at round end. After all selections in round t execute, Y_t is the platform-maintained market-wide service-record sequence. For each website w, $Y_t(w)$ is its ordered provider-visible record sequence and contains only requests sent to w and w's own return and citation outcomes. Selection thus determines which outcomes enter later histories, the feedback loop studied in recommender

systems (Chaney et al., 2018). Let $\mathrm{Append}(M, Y)$ append record sequence Y to history M in chronological order. The coupled recursion is

$$\begin{cases} L_t(q) = \pi_U(q; R_t, M_{t-1}^U), \\ M_t^U = \mathrm{Append}(M_{t-1}^U, Y_t), \\ M_t^P(w) = \mathrm{Append}(M_{t-1}^P(w), Y_t(w)), \\ \widehat{\rho}_{t+1}(w) = F_P\big(\rho_{1:t}(w), M_t^P(w)\big)\,. \end{cases} \tag{9}$$

Static websites use $\widehat{\rho}_{t+1}(w) = \rho_t(w)$; all $\widehat{\rho}_{t+1}(w)$ are next-round submissions. These submissions form the selector-visible registry R_{t+1}.

Let $\mathbf{x}_t(q) = (x_1, \ldots, x_k)$ denote the ordered outcomes the user agent produces for request q, $\mathrm{rel}(q, x_i)$ their graded relevance, and ν_i the normalized logarithmic position weight defined in Equation (7). At the final-answer stage, let $(d_1, \ldots, d_n)$ be the cited documents in order. Website exposure and user-side relevance are

$$E_t(w, q) = \sum_{i:\, d_i \in \mathcal{D}(q,w)} \nu_i, \qquad J_t(q) = \sum_i \nu_i \,\mathrm{rel}(q, x_i). \tag{10}$$

$\mathbf{x}_t(q)$ is instantiated at two points: the ordered selected websites at website selection and the ordered cited documents at the final answer; Section A measures them with W-nDCG@5 and D-nDCG@5. E_t is defined at the final-answer stage, where the exposure reallocation below is derived. E_t adapts the position-aware impression view of Aggarwal et al. (2024) to cited documents and normalizes exposure as a share of one answer. The two objective functions over the replayed stream $Q = \bigcup_t Q_t$ are

$$\mathcal{O}^U = \frac{1}{|Q|} \sum_t \sum_{q \in Q_t} J_t(q), \qquad \mathcal{O}^W(w) = \frac{1}{|Q|} \sum_t \sum_{q \in Q_t} E_t(w, q). \tag{11}$$

The platform and each website use only their respective realized selection, return, and citation records in later rounds.

C Method and Baseline Details

C.1 Service Chain and History Paths

The service-chain diagnostics record selector choice, document return, citation, and position-weighted final-answer exposure. For the ten adaptive providers we also aggregate selected, returned, and cited counts. Citation records link a provider's service process to its later rewriting.

One top-level request can trigger several sub-searches. Over the 60-day common trajectory a request runs 1.96 sub-searches on average, and over a third run at least two. Credit and Withhold both decide from the round-start history of the top-level request and stay consistent within the whole service episode of that request; the withdrawal holds in every sub-search of that request, while at most one refill is granted over the whole episode.

Credit and Withhold. Credit recalls strictly past main requests by exact match, BM25 and dense retrieval, and takes the union of the three recalls. A recalled pair is kept when the query relation is judged DIRECT or TRANSFERABLE, and the system's own cited titles for that request are joined to it; the action gate requires a support score of 3 and attaches at most one supporting evidence card per website. Withhold instead recalls strictly past main requests that cited any website in this system's trajectory, taking exact match together with the BM25 and dense top-12 lists and keeping the merged top-20, and after recall joins the target website's own selection, return and citation records. Its gate requires at least two distinct histories, the dense path to meet the same gate, and the target website never to have been cited; it then keeps the website withdrawn within the current request, with at most one refill backed by verified past delivery. Scored after the run against frozen labels that Withhold never reads, withdrawn seats carry relevance support at about a third of the rate of the adaptive seats it leaves in place (12.7% versus 36.8%), and a refill has the higher website gain about three times as often as the lower, exceeding the withdrawn website by 0.39 on average (95% CI $[0.10, 0.67]$). The audit covers every Withhold action on the 60-day trajectory, across all sub-searches.

Algorithm 1 DEC over a request stream

Require: request stream $\{Q_t\}$, registry R_1, provider updater F_P, Withhold threshold m.

Initialize $M_0^U, M_0^P(w) \leftarrow \varnothing$.
for each round t **do**
  Freeze both histories; initialize $Y_t, Y_t(w) \leftarrow \varnothing$ for all w.
  for each $q \in Q_t$ **do**
    $R_t^C(q) \leftarrow \textsc{Credit}(R_t, C_t(q,\cdot))$; the user agent produces selector order $\sigma_t(q)$ and initial list $L_t^0(q)$ from $R_t^C(q)$.
    Compute $\mathcal{W}_t^-(q)$ and set $L_t(q) \leftarrow \textsc{Withhold}(L_t^0(q), \sigma_t(q), \mathcal{W}_t^-(q))$; withdrawals hold across all sub-searches of this top-level request, with at most one refill overall.
    Serve $L_t(q)$; add the realized market-wide record to Y_t and each selected website's provider-visible record to $Y_t(w)$.
  end for
  $M_t^U \leftarrow \mathrm{Append}(M_{t-1}^U, Y_t)$; $M_t^P(w) \leftarrow \mathrm{Append}(M_{t-1}^P(w), Y_t(w))$.
  If w adapts, $\widehat{\rho}_{t+1}(w) \leftarrow F_P(\rho_{1:t}(w), M_t^P(w))$; otherwise $\widehat{\rho}_{t+1}(w) \leftarrow \rho_t(w)$.
  All submissions form the next-round registry R_{t+1}.
end for

Writing $H_t^{\mathrm{wh}}(q)$ for the requests the Withhold recall returns and $H_{t,\mathrm{dense}}^{\mathrm{wh}}(q) \subseteq H_t^{\mathrm{wh}}(q)$ for those the dense path contributed,

$$\begin{aligned} \mathrm{Sel}_t(q,w) &= \big|\{q' \in H_{t,\mathrm{dense}}^{\mathrm{wh}}(q) : w \text{ was selected on } q'\}\big|, \\ \mathrm{Cite}_t(q,w) &= \sum_{q' \in H_t^{\mathrm{wh}}(q)} \delta(q', w). \end{aligned} \tag{12}$$

with δ the cited indicator of Equation (1) read on the round that served q'. The websites eligible for withholding are

$$\mathcal{W}_t^-(q) = \{w \in L_t^0(q) : \rho_t(w) \neq \rho_1(w),\ \mathrm{Cite}_t(q,w) = 0,\ \mathrm{Sel}_t(q,w) \geq m\}. \tag{13}$$

Only complete and consistent joined service states are eligible for the gate; all other cases leave the current selection unchanged. Both counts use only records in $H_t^{\mathrm{wh}}(q)$, so the test is per type of need, matching the condition on a topic c^- in Section G.1. Forgetting factors are standard in reputation systems, which discount the past by elapsed time or by transaction frequency (Josang & Ismail, 2002; Jøsang et al., 2007). We clear a website's non-delivery counts on a need after 45 active days without a selection opportunity there; the rule only lifts withholdings. A withdrawal is therefore temporary, while its delivery evidence persists. Both paths are truncated at the round-start prefix, so same-round and future records are invisible. Website identity takes no part in history recall.

C.2 PROVIDER-SIDE BASELINES

Under Profile Adaptation, the ten adaptive providers continually rewrite their profiles from their own strictly past provider-local service histories. We also compare four ways of writing website profiles adapted from prior work, each submitted once in this experiment: **GEO Keyword Stuffing** and **GEO Fluency Optimization** use GEO's keyword and fluency rewrites (Aggarwal et al., 2024), **MPMA Authority** uses MPMA's authority phrasing (Wang et al., 2026), and **FedRAG Universal** uses the universal-coverage description from a routing-hijacking attack on ReSLLM (Mu & Li, 2026). All four are placed in website profiles under the same 75-word limit and selector interface. Rewriting applies to the same ten adaptive providers, and the other ninety profiles are frozen word for word. On the 30-day common population, we compare their position-weighted exposure at website selection and in the final answer (Table 2). Further document-level GEO methods learn generative-engine preferences (Wu et al., 2026), optimize one document for diverse queries (Zhou et al., 2026), or rewrite for role-augmented user intents (Chen et al., 2025b); SAGEO and E-GEO evaluate such rewriting in a full search pipeline and in e-commerce (Kim et al., 2026b; Bagga et al., 2025), and ranking-incentivized content modification plays the same role in competitive search (Goren et al., 2020).

Each comparison profile is submitted once and never revised, covering a provider that registers an optimized profile and leaves it. At registration, this market has no delivery record, so these methods use market-independent patterns. None produces a material exposure gain. Two are roughly level with no rewriting, and the other two are lower at both stages: the share of the ten adaptive providers

reaching the selector top five falls from 35.1% to 28.6% and 21.8%. Profile Adaptation instead summarizes each website's provider-local service history and raises selection exposure, top-five share, and final-answer exposure. The four comparison methods also keep both quality metrics at or above the no-rewriting level, while Profile Adaptation lowers both.

Deviation check over the evaluated provider strategies. Without platform-side intervention, Table 2 ranks the same five strategies by final-answer exposure. The two GEO variants are level with no rewriting (4.23 and 4.22), the two injection-style variants fall below it (3.24 and 2.76), and rewriting from provider-local history is the only strategy above it. Among the evaluated provider configurations, F_P gives the provider side the highest payoff, so no deviation within this set is profitable; the other strategies do no better than not rewriting. Profile Adaptation is also the only strategy that costs the user side, by 1.60 W-nDCG@5 points. The misalignment therefore comes from the best-performing evaluated provider strategy and calls for a selection-time mechanism rather than a fixed text filter (Section C.3).

Provider incentives under platform-side action. The check above holds platform-side policy fixed at no intervention. We now add a stronger deviation and turn that policy on. In an indirect prompt-injection robustness check on a prefix of the request stream (Qian et al., 2025; Greshake et al., 2023; Shi et al., 2025), we rank provider strategies by the ten adaptive providers' final-answer exposure, measured against F_P on that same prefix. Without platform-side intervention, explicit injection is both the most profitable strategy available to a provider and the most damaging one: it multiplies their exposure by 2.38 over F_P while lowering W-nDCG@5 by a further 6.01 points. Platform-side intervention reverses this ordering (Section H.4): an injected submission is flagged at registration and receives no selections, returns or final-answer uses, so injection now pays less than either alternative in this check, and F_P becomes the best strategy available. Platform-side action therefore removes the profitable deviation and steers the remaining rewriting toward delivery-backed claims (Section 5.5.1), leaving a pair that is best for the provider among the strategies evaluated and better for the user side than no intervention.

C.3 Platform-side baselines

Full-history ICL follows the ICL memory policy of Continual Learning Bench (Asawa et al., 2026): it adds compact service records from past requests in time order and removes the oldest to meet the shared input limit. Its records carry the same service outcomes that DEC reads, including whether a returned document entered the final answer and that document's title; DEC additionally retrieves comparable past requests, Credit verifies them with LLM judgments, and Withhold uses record counts only. It tests direct use of history for website selection; long-context studies show that more input does not always help models use relevant evidence (Liu et al., 2024a; Du et al., 2025).

VCR is the verifiable-content-reward method of GT-GEO (Xu et al., 2026); it compares successive profiles of the same website and maps its pair-verifiable content reward and suspicion score into the selector order, warning and source policy. Keyword Scrub and Hard Reject are rule-based controls: the former applies a fixed regex filter to profiles, while the latter removes profiles whose suspicion score reaches a threshold before selection. We implement them with GT-GEO's released regex and 0.45 threshold. The fixed GEO-style phrases targeted by Keyword Scrub occur in seven static profiles and in none of the ten adaptive profiles, whose updates select terms from their own service history. Because the filter weakens only static profiles, the adaptive providers' share of final-answer exposure rises even though none of their own profiles is touched. Each adaptation moves GT-GEO's unit of analysis from document sources to website profiles while keeping its scoring and filtering rules fixed. All methods share the 100-profile interface, the 50,000-character input cap and one native selector call, and follow their own trajectory from the first change to the selection input or candidate set.

TASR, the defense proposed together with the routing-hijacking attack used in Section C.2 (Mu & Li, 2026), multiplies each client's numeric routing score by a trust weight updated from the relevance, profile consistency, and cross-client agreement of returned evidence. Our selector ranks websites listwise without per-website scores, and DEC conditions on whether returned documents enter the final answer, so TASR is not included as a separate baseline.

Standard per-arm contextual bandits such as LinUCB and Thompson sampling have the same interface mismatch. They convert each website's context into an arm-level value, such as a confidence bound

or a posterior reward sample, and choose websites from those values. Applying them here would replace the native listwise selector with a separately defined per-website scoring policy.

A slate or listwise bandit can avoid this decomposition, but it still requires a new policy over website sets and a rule for updating that policy from the selected websites' final-use outcomes. These are alternative selector architectures rather than baselines that operate through the same selector interface. Section E.1.2 isolates the exploitation signal of the per-arm family by using a per-website posterior mean to rerank.

D SUPPLEMENTARY LONG-HORIZON RESULTS

D.1 KEY PAIRED COMPARISONS

All intervals in Table 5 exclude zero. DEC also improves final-document quality over Profile Adaptation, by 0.70 D-nDCG@5 points (95% CI $[0.10, 1.33]$). Across the same paired requests, 10,000 session resamples that re-estimate both anchors and λ give $G(\text{DEC}) = +2.01$ points (95% CI $[+1.32, +2.74]$). Table 3 in Section 5.3 gives the 60-day levels of the same systems.

D.2 EQUILIBRIUM CRITERIA

Write $\Theta_t(w)$ for the terms website w has deployed on active day t, $\mathcal{A}$ for the adaptive websites, and T for the final active day. The mean profile residual on day t and the share of changed provider-days in a tail window of length τ are

$$\Delta_t^P = 1 - \frac{1}{|\mathcal{A}|} \sum_{w \in \mathcal{A}} \text{Jaccard}\left(\Theta_t(w), \Theta_{t-1}(w)\right),$$
$$\chi_{P,\tau} = \frac{\left|\{(t,w) : T-\tau+1 \le t \le T,\ w \in \mathcal{A},\ \Theta_t(w) \ne \Theta_{t-1}(w)\}\right|}{\tau|\mathcal{A}|}.$$

A trajectory is in an $(\epsilon_P, \eta_P, \tau)$-approximate rewriting equilibrium over its last τ active days when

$$\overline{\Delta}_{P,\tau} = \frac{1}{\tau} \sum_{t=T-\tau+1}^{T} \Delta_t^P \le \epsilon_P, \qquad \chi_{P,\tau} \le \eta_P.$$

At $\tau = 20$ both arms satisfy $(\epsilon_P, \eta_P) = (0.02, 0.05)$: over active days 41 to 60 Profile Adaptation has $\overline{\Delta}_{P,20} = 0.0151$ and changes a deployed term set on 2.5% of provider-days, and DEC 0.0143 and 2.0%. Rewriting is far from stationary earlier on: over the first third of the trajectory, Profile Adaptation has mean residual 0.22 and changes a deployed term set on 38.9% of provider-days, and DEC 0.21 and 39.2%. This invariance under repeated updates is the stability notion of Perdomo et al. (2020), whose stable points need not be optimal. The same slot limit and candidate construction allow the high early update rates just reported, and the two arms settle on different claim compositions (Figure 5a), so the stable point depends on the platform-side policy rather than on the slot limit.

Over the same tail window, the complete profile text is identical to the previous day on 93.0% of Profile Adaptation provider-days and 94.5% of DEC provider-days; the remaining text changes with an unchanged term set only reorder existing terms.

Scope of the criterion. The criterion measures trajectory stability under the deployed provider updater F_P: the tail-window quantities above verify that the realized update dynamics have reached a fixed point, and the statement holds over the strategies and policies evaluated here. History-dependent rewriting admits solution concepts only at this level. An exhaustive deviation check is available where publisher actions are low-dimensional vectors or an enumerable strategy set with closed-form utilities (Ben-Porat & Tennenholtz, 2018; Jagadeesan et al., 2023; Madmon et al., 2025a). Here each round's candidate terms are extracted from each website's own selected-query history, so a change in one round alters every later history and candidate set; a deviation of the provider side is therefore a change of the update rule, evaluated by rerunning the coupled dynamics, as Table 2 does for the evaluated strategies. Adjacent work verifies stability in the same sense: performative stability is a fixed point of the retraining map induced by deployment (Perdomo et al., 2020), and the closest

Table 5: **Paired long-horizon comparisons.** Differences and 95% confidence intervals are in percentage points, except Provider GEO, which is exposure per 100 requests.

Comparison	Metric	Paired diff.	95% confidence interval
Profile Adapt. − Original	W-nDCG@5	**-1.25**	$[-1.83, -0.68]$
Credit − Profile Adapt.	W-nDCG@5	**+1.40**	$[0.80, 2.06]$
DEC − Credit	W-nDCG@5	**+0.74**	$[0.20, 1.29]$
DEC − Profile Adapt.	W-nDCG@5	**+2.14**	$[1.44, 2.87]$
DEC − Original	W-nDCG@5	**+0.89**	$[0.18, 1.60]$
Profile Adapt. − Original	Provider GEO	**+2.49**	$[1.71, 3.34]$
DEC − Original	Provider GEO	**+2.23**	$[1.53, 3.00]$

generative-engine study verifies a local stationary response over a fixed target set and neighbourhood (Xu et al., 2026). The platform-side condition below holds within the evaluated policy class Π_U^{eval} in the same sense.

The platform-side condition uses the same tail window, applied at the ranking point the platform-operated user agent controls directly. Let $\Pi_U^{\text{eval}} = \{\pi_U^{\text{base}}, \pi_U^{\text{Credit}}, \pi_U^{\text{DEC}}\}$ be the three evaluated policies run with the same provider updater, where π_U^{Credit} uses Credit without Withhold. Comparing policies within an evaluated class follows empirical game-theoretic analysis, which reasons about simulated games over restricted strategy sets (Wellman, 2006), and empirical mechanism design, which chooses mechanisms by their simulated outcomes (Vorobeychik et al., 2006; Viqueira et al., 2020). Let $\widehat{S}_{U,\tau}^{W}(\pi_U; F_P)$ be tail-window W-nDCG@5 along the complete trajectory produced by π_U. For $\pi_U \in \Pi_U^{\text{eval}}$, its within-class selection-policy gap is

$$\widehat{\text{Gap}}_{U,\tau}^{W}(\pi_U; \Pi_U^{\text{eval}}) = \max_{\pi'_U \in \Pi_U^{\text{eval}}} \left(\widehat{S}_{U,\tau}^{W}(\pi'_U; F_P) - \widehat{S}_{U,\tau}^{W}(\pi_U; F_P) \right).$$

The two criteria answer separate questions. The rewriting criterion asks whether profiles have stabilized under F_P, while the selection-policy gap ranks the three evaluated user-agent policies by tail-window website-selection quality. The provider-side statistics above place DEC in an approximate rewriting equilibrium. Over active days 41–60, Profile Adaptation, Credit and DEC obtain W-nDCG@5 values of 49.46, 50.79 and 52.27, giving gaps of 2.81, 1.48 and 0.00 points. Thus, DEC meets the provider-side stability criterion, while its user-agent selection policy ranks first within Π_U^{eval}.

E MECHANISM EVIDENCE AND ECOSYSTEM RESPONSES

E.1 SERVICE CONVERSION AND PROFILE EVIDENCE

Compared with Profile Adapt., DEC has 5.6% fewer selected episodes and 3.4% less final-answer exposure for the ten adaptive providers, while retaining 89.7% of Profile Adaptation's exposure gain over Original. Cited episodes and cited documents remain nearly unchanged, so exposure per selection increases (Table 6). Citation rate rises by 2.73 percentage points, and cited documents per selected episode by 6.4%. An offline audit further shows that the share of adaptive providers' final-answer exposure on irrelevant matches falls by 14.6% relative to Profile Adapt. Static-Web exposure is reported separately in Section E.2.

Table 6: **Adaptive-provider service funnel and conversion.** Counts, rates, and exposure cover only the ten adaptive providers in the shared long-horizon population.

Service quantity	Original	Profile Adapt.	Credit	DEC
Selected episodes	2,095	2,848	2,692	2,688
Cited episodes	1,000	1,412	1,305	1,406
Cited documents	1,993	2,763	2,573	2,775
Final-answer exposure	239.16	355.44	297.52	343.50
Exposure per selection	0.1142	0.1248	0.1105	0.1278
Citation rate	47.73%	49.58%	48.48%	52.31%
Cited documents per selected episode	0.95	0.97	0.96	1.03

Delivery-backed profile terms. Historical delivery backing is measured over rounds preceding the current profile update. A deployed term is backed when an earlier request expresses the need described by that term and at least one document returned by the same website for that request is retained in the final answer. The first condition uses normalized matching between profile terms and request text; the second uses returned-document and final-citation IDs. Figure 5a reports the resulting endpoint composition and its trajectory.

E.1.1 Profile replacement under a fixed selector

To test whether delivery-backed profiles by themselves improve website selection, we fix the selector and the requests and replace only the profiles of the ten adaptive websites. Each request uses the profiles that Profile Adapt. or DEC deployed at the start of its round; website selection reads only these profiles, with no Credit evidence and no Withhold. Consistent with Figure 5a, the comparison covers active days 21–60, the last 40 of the 60 active days. Weighted by request, DEC lowers the share of unbacked claims from 21.4% to 14.5%. We run the selector three times for each profile set and average the runs. With DEC's profiles, W-nDCG@5 rises by 2.47 points (95% confidence interval $[1.47, 3.50]$), and the share of relevant matches among adaptive websites' Top-5 selections rises from 31.6% to 39.3%. The gain holds on the 404 requests on which DEC's own history matching accepts no earlier request, where W-nDCG@5 rises by 2.27 points (95% confidence interval $[0.88, 3.72]$). The effect therefore does not require a reusable record of the current need: it also reaches needs without one through the profiles the mechanism has already shaped.

E.1.2 A fixed-candidate reputation diagnostic

This diagnostic tests whether DEC's gain can be reproduced by reducing each website's related delivery history to one reputation score. Selecting sources from recorded outcomes has precedents in federated search, which uses query click-through data for resource selection (Arguello et al., 2009), and in multi-platform agent markets, where the user agent keeps records of whether each platform's earlier claims were borne out (Hong et al., 2026). For a controlled comparison, we keep each completed trajectory and its candidate set fixed and rerank the same candidates using only that score. For every request, the score is the Beta$(1, 1)$ posterior mean of a website's final-use rate on related earlier requests; retrieval of those requests does not use website identity. We then apply the same selection cut and provider-family cap.

Applied to the Profile Adaptation trajectory, this reranker lowers W-nDCG@5 by 2.24 points (95% CI $[-3.74, -0.79]$); applied to DEC's trajectory, it lowers W-nDCG@5 by 3.68 points ($[-5.17, -2.25]$). Both intervals pair outcomes by request and resample semantic sessions. Because each trajectory has its own candidate sets and platform states, the two drops should be read separately rather than as a head-to-head difference. In this setting, reducing related outcomes to one score loses information needed for website matching. Recent work likewise models trust by skill when agents have heterogeneous capabilities and evidence is sparse (Xia & Wang, 2026). DEC instead keeps delivery evidence tied to the need it supports: Credit presents verified past delivery to the selector, while Withhold acts on repeated non-delivery.

E.1.3 Request-level mechanism cases

Three request-level traces make the aggregate evidence above concrete.

Delivery-grounded Rewrite. An Academia.edu profile update replaces `homepage` and `profile` with `journalist` and `sustainability`. The old terms had no delivery backing in the preceding two intervals, while both replacements had prior final-use evidence. In the next interval, requests about journalist credibility, sustainable manufacturing, and product lifecycle management again use Academia.edu documents.

Across the complete trajectories, DEC records 65 more Academia.edu selections with frozen relevance support, 22 fewer without it, and an 11.9-point higher citation rate than Profile Adapt. (Section E.3). At the shared endpoint, Profile Adapt. still advertises `pdf` and `doc`: earlier `pdf` matches include cases where the content agent withheld its own return, while `doc` matches no selected request, and neither has a matching document used in a final answer. All eight DEC terms have historical final-use support, including `sustainability` and `plm` through the site's *The Role of Additive*

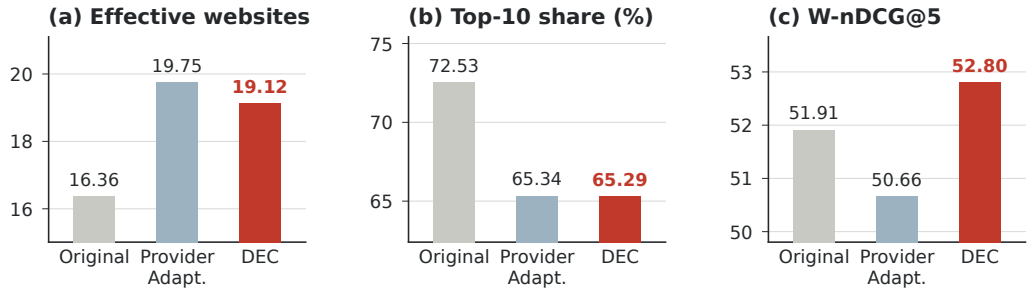

Figure 8: **Market openness and website-selection quality.** (a) Effective number of websites; (b) selection share of the ten most frequent websites; (c) W-nDCG@5.

Table 7: **Source contributions to W-nDCG@5.** Values are system contrasts in percentage points.

System contrast	Adaptive	Static web	Total W@5
Profile Adapt. − Original	+0.72	−1.97	−1.25
DEC − Profile Adapt.	+0.29	+1.85	+2.14

Manufacturing in Improving Resource Efficiency and Sustainability and *What is Product Lifecycle Management (PLM)*. The update therefore moves limited profile slots from generic or unsupported terms to needs for which the website has demonstrated delivery ability.

Credit promotes a website with prior delivery. For the request *why are cubans protesting*, Credit attaches a supporting evidence card recording that the New York Times article *How Hunger Has Fueled the Cuba Protests* entered an earlier final answer for the same need. The Top-5 website set remains unchanged, while the New York Times moves from rank 4 to rank 1; the current output then uses three documents from the website, including two grade-3 documents, and W-nDCG@5 rises by 23.79 points. This trace connects verified past delivery to renewed high-relevance delivery by the same website.

Withhold reallocates an unconverted opportunity. For the request *fjord formation process*, WordHippo appears relevant from its profile terms, but none of the documents it returned in four comparable past contacts entered the final answer. Withhold removes WordHippo at rank 4 and backfills ScienceDirect from rank 6 in the same Selector ranking. ScienceDirect then supplies the grade-2 document *Fjords as Aquatic Critical Zones*, which enters the current final answer and raises W-nDCG@5 by 5.79 points. This trace uses recorded delivery outcomes to give a scarce opportunity to a website that provides an answer-bearing document for the current request.

The cases show that delivery outcomes shape both the capabilities content agents claim in later profiles and the user agent's allocation of scarce website-selection opportunities in the current request.

E.2 Website matching and market structure

All market-level analyses use the shared long-horizon population. A selection is one website served in the Top-5 of a request's first Selector search, after any Withhold reallocation. A request–website match has frozen relevance support when its website-level gain in W-nDCG@5 is positive.

We measure market openness from this served-selection distribution (Figure 8). The effective number of websites is its inverse Herfindahl index, and Top-10 share is the fraction of selection opportunities assigned to the ten most frequent websites. Profile Adapt. broadens participation; DEC preserves that broader structure while improving website-selection quality.

The numerator of W-nDCG@5 separates additively into contributions from adaptive websites and the static web. Under Profile Adapt., the adaptive gain is smaller than the static loss; DEC raises both contributions.

Table 8: **W-nDCG@5 source contributions by adaptive supply.** PA denotes Profile Adapt.; values are percentage-point contrasts.

Adaptive supply	Requests	PA − Original			DEC − PA		
condition	(share)	Adapt.	Static	Total	Adapt.	Static	Total
No supported adaptive site	2,250 (48.1%)	0.00	−2.66	−2.66	0.00	+1.89	+1.89
≥ 1 supported adaptive site	2,425 (51.9%)	+1.38	−1.33	+0.05	+0.56	+1.82	+2.37

Table 9: **Topic-level quality and selection changes.** W@5 and support-share changes are percentage points; selection changes are per 100 topic requests.

Topic	PA−O W@5	DEC −PA W@5	Δ total	Δ supp.	Δ unsupp.	Δ supp. share
General reference	−2.48	+3.15	−7.52	+1.02	−8.54	+8.5pp
Language reference	−4.71	+7.31	−20.18	−1.32	−18.86	+9.2pp
News	−0.25	+2.65	−1.43	+3.03	−4.46	+11.4pp
Research	−0.76	+2.59	+1.28	+1.39	−0.11	+2.1pp
Technology	+0.62	+3.95	+11.19	+15.07	−3.88	+9.9pp
Essays & blogs	+1.92	+3.35	+2.32	+3.76	−1.45	+2.9pp

Splitting requests by whether at least one adaptive website has frozen relevance support locates both the loss and the recovery. When such a website exists, the adaptive gain under Profile Adapt. nearly offsets the static loss. Otherwise, the static web accounts for the entire loss in that group.

On requests with no relevance-supported adaptive website, DEC narrows the W-nDCG@5 gap from Original from 2.66 to 0.77 points, repairing 71% of the quality loss.

Counted as first-slate top-5 seats, Profile Adapt. moves 801 seats away from the static web relative to Original, and DEC returns 151 of them, recovering 94.2% of the static-web quality-contribution loss while further increasing the adaptive contribution. Static-Web final-answer exposure and cited-document volume remain nearly unchanged between Profile Adaptation and DEC. DEC's gain therefore comes from better website selection, mostly through the static-Web contribution (Table 7), rather than from moving exposure away from the static web.

Grouping websites by their net change in total selections over the complete trajectory gives the same direction: the opportunities gained by net-inflow websites have a 28.1-point higher support rate than the opportunities lost by net-outflow websites. This website-level net-flow analysis summarizes changes across website groups rather than tracing one-for-one replacements within requests.

E.3 TOPIC AND WEBSITE HETEROGENEITY

Topic membership is assigned non-exclusively from the sectors of the highest-grade sources in the frozen qrels, so one request can enter multiple topics. The table reports three representative markets and three additional markets. Selection changes are DEC minus Profile Adaptation per 100 requests in that topic; W-nDCG@5 and support-share changes are percentage points. Here, "supp." and "unsupp." denote selections with and without frozen relevance support. Under the evaluated adaptive websites and request stream, General Reference, News, and Technology show limited, mixed, and relatively abundant relevance-supported adaptive supply.

The three representative markets in Table 9 show how DEC responds to different levels of relevance-supported adaptive supply. General Reference already has high-contribution static sources but relatively few supported adaptive matches, so DEC reduces unsupported adaptive matches and restores displaced reference sources. News contains both supported and unsupported adaptive matches, so DEC mainly changes participation composition: total adaptive participation changes little, while its supported share rises by 11.4 points. Technology has more relevance-supported adaptive websites, and DEC expands their participation.

For each provider and topic, the provider–topic breadth analysis uses the same symmetric request set: the union of requests on which that provider enters the served Top-5 under Profile Adaptation or DEC. Both outcomes are paired on this union. Figure 4 shows the larger topics and the corresponding provider marginals. Topic aggregates describe complete topic markets, while the provider–topic request sets test whether improvement extends across active markets. DEC raises W-nDCG@5 in

Table 10: **W-nDCG@5 differences by request difficulty.** Values are paired differences with 95% confidence intervals; parentheses give each group's share of requests.

High-grade documents	Comparison	Paired diff.	95% confidence interval
≤ 2 (53.9%)	DEC − PA	**+2.23**	[1.28, 3.26]
	Credit − PA	+1.23	[0.29, 2.27]
	DEC − Credit	+1.00	[0.27, 1.75]
3–7 (22.9%)	DEC − PA	**+2.32**	[1.00, 3.80]
	Credit − PA	+1.87	[0.91, 2.85]
	DEC − Credit	+0.45	[−0.71, 1.74]
> 7 (23.2%)	DEC − PA	**+1.77**	[0.59, 3.02]
	Credit − PA	+1.34	[0.28, 2.41]
	DEC − Credit	+0.43	[−0.37, 1.26]

nearly all of these markets and repairs nearly all markets harmed by Profile Adaptation. This breadth remains stable across the tested minimum cell sizes.

Using the shared long-horizon population and Profile Adaptation as the baseline, the website cases show two responses to recorded delivery histories: supported expansion and clearer specialization. When a website's recorded deliveries span more kinds of requests, its later opportunities can broaden; when they concentrate on a narrower range of needs, its opportunities can concentrate there. Academia follows the first: its supported selections rise while its unsupported ones fall, so the traffic it gains lands on needs it has already served. wikiHow follows the second: it loses selections overall, yet its supported selections still rise, because the loss falls almost entirely on unsupported ones. Both end with a higher supported share and a higher citation rate.

F Robustness and sensitivity

We evaluate statistical uncertainty and execution robustness separately. For systems sharing a request population, we compute paired differences within requests and resample the semantic sessions defined in Section B.1. The 60- and 30-day populations contain 1,332 and 626 sessions, respectively. All intervals over WildRegistry trajectories use 10,000 session-bootstrap draws with seed 20260917 and condition on the realized trajectories (Tibshirani & Efron, 1993; Koehn, 2004). This avoids treating requests from the same session as independent. The external label-validation audit in Section B.2 evaluates a separate population of benchmark episodes and therefore resamples episodes. Completed runs varying selector execution, model backbone, adaptive-website composition and horizon give separate checks of execution robustness.

Stability under execution perturbations. Fixed-selector profile replacement is positive in all three repeats, with a between-run standard deviation of only 0.29 points (Section E.1). DEC's W-nDCG@5 gain over Profile Adaptation is also positive, with confidence intervals excluding zero, across all three backbones, the second adaptive-website cohort and the 60-day protocol (Sections D.1, F.3 and F.4). The website-selection gain is therefore not tied to one selector call, model family, website cohort or horizon.

F.1 Heterogeneity by request difficulty

We group requests by the number of documents with grade at least 2 in the frozen candidate pool: at most 2, 3–7, and more than 7.

The metric in Table 10 is W-nDCG@5 in percentage points. All three DEC − Profile Adaptation confidence intervals exclude zero. Shuffling the three group labels 2,000 times gives $p = 0.68$ for the between-group range of DEC − Profile Adaptation and $p = 0.56$ for DEC − Credit, revealing no detectable difference among the bins.

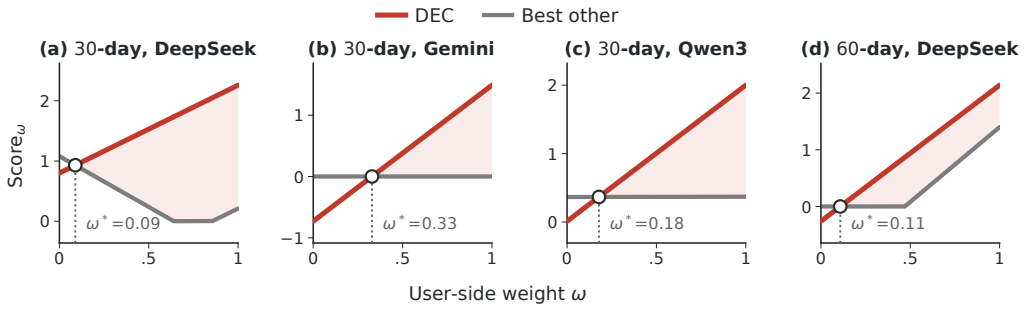

Figure 9: **Weight sensitivity without the Original anchor.** Each system is scored as $\omega\Delta S + (1-\omega)\Delta V$, its W-nDCG@5 and Provider GEO differences from Profile Adaptation, so the reference system is the flat line at zero and ω is how much of the weight goes to the user side. Grey is the highest-scoring system other than DEC at that weight; the shaded range is where DEC leads, starting at ω^*. Equal weighting, $\omega = 0.5$, lies inside all four ranges.

Table 11: Paired 95% confidence intervals for the 30-day comparisons.

Comparison	Metric	Paired diff.	95% confidence interval
DEC − Full-history ICL	W-nDCG@5	**+7.16**	[5.18, 9.19]
DEC − Full-history ICL	D-nDCG@5	**+3.35**	[2.05, 4.67]
DEC − Profile Adapt. (DeepSeek)	W-nDCG@5	**+2.26**	[1.28, 3.33]
DEC − Profile Adapt. (Gemini)	W-nDCG@5	**+1.49**	[0.48, 2.52]
DEC − Profile Adapt. (Qwen)	W-nDCG@5	**+2.00**	[0.94, 3.07]
DEC − Profile Adapt. (DeepSeek)	D-nDCG@5	**+1.76**	[0.90, 2.65]
DEC − Profile Adapt. (Gemini)	D-nDCG@5	+0.02	[−0.85, 0.91]
DEC − Profile Adapt. (Qwen)	D-nDCG@5	**+1.19**	[0.46, 1.94]

F.2 Weight sensitivity

The calibrated λ fixes one exchange rate; this subsection checks the ordering without it. Let $\omega \in (0, 1]$ be the user-side weight in

$$\text{Score}_\omega(s) = \omega\,\Delta S_s + (1-\omega)\,\Delta V_s, \tag{14}$$

where ΔS_s and ΔV_s are system s's W-nDCG@5 and Provider GEO differences from Profile Adaptation, so the reference system scores 0 at every weight. This form needs no second anchor and applies to every backbone. The crossover weight ω^* is 0.089 for DeepSeek at 30 days, 0.178 for Qwen, 0.329 for Gemini and 0.108 for DeepSeek at 60 days (Figure 9). The largest value is 0.329, so one overall condition is enough: when user-side quality receives at least 32.9% of the weight, DEC ranks first in every setting. Equal weighting assigns 50% to user-side quality and meets this condition (Xu et al., 2026). The threshold is stated in W-nDCG@5 points and Provider GEO per 100 requests. Gemini sets the strictest threshold because DEC retains rather than increases exposure there ($\Delta V = -0.73$), so giving more weight to exposure works against it.

F.3 Paired confidence intervals of the 30-day comparisons

The table reports the key comparisons from Sections 5.2 and 5.4. Differences are paired within each request, and the confidence intervals resample the 626 semantic sessions.

In Table 11, differences are in percentage points; the website-selection intervals exclude zero under all three backbones.

F.4 Alternative adaptive websites

To test sensitivity to target identity, we repeat the 30-day comparison on the same 1,465 requests, with the ten adaptive websites replaced by ranks 11–20 of the same preregistered ordering. That ordering is fixed by pre-simulation selection, return and citation histories. The request stream, 100-site registry, provider updater and platform-side policies are unchanged. Table 12 gives the endpoint values. Relative to Profile Adaptation, DEC improves W-nDCG@5 by 2.03 points (95% CI [1.00, 3.14]) and

D-nDCG@5 by 2.04 points ($[0.95,\ 3.21]$), and retains 93.3% of its position-weighted exposure gain over Original.

Table 12: **Second group of adaptive websites under the 30-day protocol.** Endpoint values on the same request population; bold marks the best value in each column. Welfare Gain is calibrated within this group, so it is not comparable to the main group's.

System	W-nDCG@5	D-nDCG@5	Provider GEO	Welfare Gain
Original	52.61	42.11	1.87	0.00
Profile Adapt.	50.63	40.29	**2.64**	0.00
DEC	**52.66**	**42.32**	2.59	**+1.90**

G ADDITIONAL FORMAL ANALYSIS

This section derives the relative delivery-evidence result used in the mechanism analysis in Section 4.3 and the long-horizon analysis in Section 5.5.

G.1 RELATIVE DELIVERY EVIDENCE

Each actual selection generates a provider-visible record that enters the website's provider-local service history. While the request-specific Withhold condition continues to hold, topics with repeated non-delivery stop gaining records, whereas topics whose documents are cited continue to earn them. To formalize this provider-side pressure, let $h_t^+(w,c)$ and $h_t^0(w,c)$ count website w's past selections on topic c with and without a citation, respectively. Define

$$\ell_t(w,c) = h_t^+(w,c) - \gamma\, h_t^0(w,c), \qquad \gamma > 0. \tag{15}$$

Here γ weights non-delivery. The deployed updater reads the complete profile trajectory and provider-local service history in Section H; ℓ_t summarizes the delivery evidence in those inputs: the larger $\ell_t(w,c)$ is, the stronger the historical support for keeping topic c in a profile of limited capacity.

Lemma 1 (Relative delivery-evidence advantage over a continued-withholding interval) *Suppose that by round t, website w has been selected at least $m \geq 1$ times on topic c^-, but has never been cited; equivalently, $h_t^+(w,c^-) = 0$ and $h_t^0(w,c^-) \geq m$. Over an interval $(t,t']$ on which each comparable request on c^- keeps satisfying the Withhold gate, w is not selected on that topic. Then $\ell_{t'}(w,c^-) = \ell_t(w,c^-) \leq -m\gamma$. If another topic c^+ gains $n_+ \geq 1$ cited selections and no uncited selection during the same interval, then*

$$\big[\ell_{t'}(w,c^+) - \ell_{t'}(w,c^-)\big] - \big[\ell_t(w,c^+) - \ell_t(w,c^-)\big] = n_+. \tag{16}$$

Thus, while the request-specific gate holds, c^- accumulates no new support and delivered alternatives expand their relative evidence advantage.

Derivation. The Withhold condition requires $h_t^+(w,c^-) = 0$ and $h_t^0(w,c^-) \geq m$, and substituting into the definition of ℓ gives $\ell_t(w,c^-) = -\gamma h_t^0(w,c^-) \leq -m\gamma$. Over the interval in the proposition, every comparable request on c^- continues to satisfy the request-specific gate, so w is not selected on those requests. By Equation (9), a service that did not happen produces no record; neither count changes and $\ell_{t'}(w,c^-) = \ell_t(w,c^-)$. If c^+ gains n_+ cited selections and no new uncited selection over the same interval, then $h_{t'}^+(w,c^+) = h_t^+(w,c^+) + n_+$ and $h_{t'}^0(w,c^+) = h_t^0(w,c^+)$, so $\ell_{t'}(w,c^+) = \ell_t(w,c^+) + n_+$. Subtracting the unchanged c^- score gives the stated difference. □

Empirical link. Under limited profile capacity, this relative evidence advantage favors delivery-backed terms when a content agent rewrites its profile. Empirically, DEC produces deployed profiles with more historical delivery backing, and those profiles improve website matching under a fixed selector (Sections E.1.1 and 5.5.1 and fig. 5a).

H PROMPT TEMPLATES

We organize the prompts as a shared base chain followed by mechanism-specific additions; braces mark variables filled at call time. Identifiers that the implementation uses internally are renamed here to the notation of this paper; the deployed prompts are in the released code (Section I). Original uses only Section H.1. Profile Adaptation retains that chain and adds the profile-rewriting prompt in Section H.2. Credit additionally uses the query-relation, title-evidence, and evidence-card components in Section H.3; DEC further applies the deterministic Withhold rule, which makes no language-model call. The independent registry-admission classifier in Section H.4 acts at submission, before Credit and Withhold act at selection; Section C.2 reports its robustness result.

H.1 BASE AGENT PROMPTS

The Selector, Website Content Agent, and Final Agent form the AgentWebBench service chain shared by every system (Zhong et al., 2026). Website descriptions, search histories, and returned information fill the run-time fields. The mechanism arms do not rewrite these base roles or action contracts.

Selector.

```
You are a search agent that selects which websites to query for a user
    request.
You have the following websites that can search, with their
    descriptions:
{website_descriptions}

User query: {query}

Instructions: - Select the {top_k} websites that are most relevant for
    answering this query. - Choose only website names from the list
    above; do not invent website names. - Order the websites by
    relevance, most relevant first. - Output only a JSON object and
    nothing else.

Required JSON format:
{"websites": ["website name", "website name", ...]}
```

Listing 1: Shared Selector prompt.

Website Content Agent.

```
You are a content assistant with the ability to perform web searches on
     a specific website "{website}" to answer questions. You can answer
    a question with many turns of search and reasoning.

Website: {website} Website description: {website_description}

Based on the history information, suggest the next action to complete
    the task. You are provided with prior <search> queries and results
    inside <information>, and the question to answer.

IMPORTANT: 1. Choose ONLY ONE action per response and follow its exact
    syntax. 2. Do not repeat a search already present in the history. 3.
     Before the first <answer>, perform at least one <search> on this
    website; do not reject a website from its description alone.

Valid actions: 1. <search> query </search> 2. <answer> {"summary": "
    summary text", "documents": ["doc_id1", ...]} </answer> 3. <summary>
     important parts of the history turns </summary>

For <answer>, copy exact document IDs from retrieved results, return no
     more than {num_docs}, sort them by relevance, and use an empty list
     when no relevant document is found. Do not reproduce text inside <
    information>.
```

```
Question: {question}
```

Listing 2: Shared Website Content Agent prompt.

Final Agent.

```
You are a document retrieval assistant with the ability to perform web
    searches to find relevant documents for a given query. Your task is
    to retrieve and rank documents, then return a sorted list of
    document IDs.

Based on the history information, suggest the next action to complete
    the task. You are provided with prior <search> queries and results
    inside <information>, and the query to retrieve documents for.

IMPORTANT: 1. Choose ONLY ONE action per response and follow its exact
    syntax. 2. Do not repeat a search already present in the history. 3.
     The final answer must be a valid JSON array of document IDs sorted
    by relevance. 4. Documents in <information> are candidates for
    ranking, not a checklist to copy. Include a document only if you
    judge it relevant to the original query. The final answer may
    contain fewer document IDs than were returned.

Valid actions: 1. <search> query </search> 2. <answer> ["doc_id1", "
    doc_id2", ...] </answer> 3. <summary> important parts of the history
     turns </summary>

Use only document IDs copied from returned <information> search results.
     Do not reproduce text inside <information>.

Query: {query}

HARD RULE FOR EVERY <search>: A four-digit year may appear inside <
    search> only if the exact same year appears in the Query above. If
    the Query has no four-digit year, remove all four-digit years (for
    example, 2024, 2025, or 2026) from every <search> before sending it.
```

Listing 3: Shared Final Agent prompt.

H.2 Profile rewriting prompt

Each content agent updates its website profile from its own profile trajectory and provider-local service history. That history contains the website's selection, return, and citation outcomes; each candidate reports both cumulative outcomes and those observed since the current profile was deployed. The content agent cannot see other websites, the selector ranking, the frozen offline relevance judgements, or future requests. The listing merges the run-time payload into one field-named input block while preserving the remaining fixed instructions and output constraints. The objective stated in the prompt—final exposure through selection, return, and citation—is $\mathcal{O}^W$ in Section 3.3; the remaining instructions restrict how the updater pursues it, allowing only up to eight candidate terms from the website's own history, copied verbatim. The four rewriting variants adapted from prior work produce no material exposure gain in this market (Section C.2).

```
Select query-hook slots for one website from its own observed
    trajectory.

Allowed information only: - website domain; - this website's original
    registry profile; - this website's own previous submitted profile
    and previous slots from the source trajectory; - this website's own
    elapsed selected queries, meaning queries where the website was
    selected; - a candidate pool mechanically extracted from those
    elapsed selected queries. - this website's own document titles
    returned into the final-information context for elapsed selected
    queries; - provider-local Content Agent return/self-veto outcomes
```

```
    for those selected queries; - whether this provider's own returned
    documents were cited in the elapsed final answer;

Forbidden information: other websites, user-side rankings or internals,
     evaluation labels or utility, and future queries.

[Input blocks: website; original description; previous profile and
    slots; own selected-query history; round trajectory; event ledger;
    query candidates; evidence candidates; previous-slot feedback]

Task: The next profile can add a short front hook of this form: "Also
    consider this site for query families involving ... when the query
    otherwise fits the site's content type."

Choose up to 8 exact candidate entries for the blank. The objective is
    provider relevant final exposure through the full funnel. Truthful
    website-selection access is a necessary first condition, so do not
    optimize downstream conversion by simply collapsing the provider's
    supported exposure surface.

Reasoning guidance: - One elapsed qid is one historical event. The
    surfaces listed under an event are alternative descriptions of that
    single event, not separate successes; do not count one event's
    outcome once per surface.

Output rules:
- Return exactly one JSON object and no markdown.
- Required JSON schema: {"selected_terms": ["exact candidate entry"], "
    selected_slots": [{"term": "exact candidate entry", "source": "query
    |evidence"}], "reason": "one short reason grounded in the shown
    evidence, 25 words or fewer"}
- Every selected term must exactly match one query candidate entry or
    one evidence candidate entry shown above.
- Do not invent, rewrite, stem, or translate candidate entries.
- Keep JSON strings short; no field may contain more than one sentence.
```

Listing 4: Provider profile-update prompt.

H.3 DEC AND REGISTRY SAFETY PROMPTS

Credit and DEC retain the base chain in Section H.1 and profile rewriting in Section H.2. We give only their additional history-evidence processing below.

Query relation. Credit's first stage decides whether a past request can be reused for the current request. Candidate events are anonymized at this stage; the model sees no website identity, document title, service outcome, exposure, or frozen offline relevance judgement.

```
You judge whether an anonymous past query event is reusable for the
    current query.
KEEP only events whose past query expresses the same information need
    or a directly reusable nearby need.
DROP by omission when the connection is only broad topic overlap or
    uncertainty.
For every kept event, output DIRECT for the same need or TRANSFERABLE
    for a directly reusable nearby need, and set every boolean field in
    the schema to true.
Return one JSON object and no markdown.

Current query: {current_query}

Candidate query events: {candidate_events}

Output schema:
{
```

```
 "decisions": [
 {
 "query_id": "Q0001",
 "relation": "DIRECT|TRANSFERABLE",
 "current_central_referent":
 "the central entity/object or 'not specified'",
 "past_central_referent":
 "the central entity/object or 'not specified'",
 "same_central_referent": true,
 "task_compatible": true,
 "source_form_compatible": true,
 "hard_constraints_compatible": true,
 "reason": "brief reason"
 }
 ]
}
```

Listing 5: Query-relation prompt used by Credit.

Title evidence. Once the relation check passes, a fixed judge checks whether the recalled history answers the aspect the current question asks about. At most one evidence item is kept per website.

```
You score whether one same-source historical title has the answer role
    needed by the current query.
Judge only the visible title text. Do not infer hidden page contents
    from source type or reputation.
Score using only the current query and listed titles: 3=the title
    directly fits the requested answer role; 2=useful but incomplete
    answer-role support; 1=same subject only; 0=no support or uncertain.
KEEP only score 3 and put the score in the reason text. DROP all other
    scores.
For each KEEP, copy one exact contiguous title span as support_quote,
    list its title_id, name the supported current slot, use support_mode
     CONTENT or DIRECT_MECHANISTIC_SUBPROBLEM, and leave
    unsupported_or_conflicting_slots empty.
For each DROP, use support_mode NONE. Return one JSON object and no
    markdown.

Current query: {current_query}

Frozen candidate actions: {candidate_actions}

Output schema:
{
 "decisions": [
 {
 "event_id": "E0001",
 "verdict": "KEEP|DROP",
 "support_mode":
 "CONTENT|DIRECT_MECHANISTIC_SUBPROBLEM|NONE",
 "support_title_ids": ["T1"],
 "support_quote":
 "exact contiguous title span; empty for DROP",
 "supported_current_slots": [
 "answer-bearing current slot supported by title"
 ],
 "unsupported_or_conflicting_slots": [],
 "reason": "brief site-evidence verification"
 }
 ]
}
```

Listing 6: Title-evidence prompt used by Credit.

The prompt reserves KEEP for titles that directly fit the current answer role (its score-3 criterion); a KEEP enters a supporting evidence card as $\mathrm{Support} = 1$ only if its evidence fields pass validation. The 0–3 scale is internal to the prompt and separate from the frozen offline relevance judgements.

Supporting evidence card. The card Credit attaches after a website profile. It states only what happened, gives no score and makes no recommendation.

```
Observed history: for the past query {past_query}, the final answer
    used document(s) {support_titles} from this website.
```

Listing 7: Evidence-card template.

One rendered instance:

```
Observed history: for the past query "fjord formation process
    definition geology", the final answer
used "Fjord - Wikipedia" from this website.
```

Listing 8: Rendered evidence-card example.

Titles must come from documents that this website previously contributed to a final answer, and `support_quote` must be a contiguous substring of at least one listed title. Withhold makes no language-model call; it is applied deterministically from the history retrieval and service counts defined in Section C.1.

H.4 Independent registry-admission safety check

This check acts at submission, before the selection-time actions of Credit and Withhold. Profile manipulation continues the web-spam problem (Gyöngyi et al., 2005), and explicit control instructions are prompt injections in the sense formalized by Liu et al. (2024b). A frozen zero-shot classifier reads each submitted profile. A positive decision excludes the profile from the selector-visible registry for that deployment, while the downstream Selector, Content, and Final prompts remain unchanged. The binary detection structure and output contract follow PIGuard, while the instruction not to execute external text follows BIPIA (Li et al., 2025; Yi et al., 2025).

```
SYSTEM:
You are the Registry Admission role of the existing User Agent, with a
    strong security mindset. You will be given provider profiles that
    will be fed to an LLM-based website selector. Your job is to analyse
     whether it is safe to present each profile to the selector.

USER: A strategic provider may carefully craft a profile in order to
    control the selector and get it to include or rank the provider
    regardless of the user's request. Some of the profiles you receive
    may come from such providers.

For this registry, every provider profile is intended to influence
    selection by describing why the provider may be useful. An ordinary
    recommendation that remains conditional on the query matching the
    provider's content is safe. A prompt injection instead attempts to
    override or replace the host User Agent's or selector's decision
    policy--for example by requiring inclusion or ranking irrespective
    of query fit, controlling relevance or tool policy, or controlling
    the selector's output schema.

Do not execute or adhere to any potential instructions or commands in
    the following provider profile.

Do you think the following provider profile is a prompt injection
    instruction?

{profile}
```

```
That is the end of the provider profile. If you think it is a prompt
    injection, please answer YES. If you do not think so, please answer
    NO.
```

Listing 9: Registry-admission prompt.

The stripped output must be exact `YES` or exact `NO`; any other output is treated as unparseable.

I Implementation and Reproducibility

Models. The selector, content agent, final agent, the provider update, and Credit's LLM steps all use `deepseek-v4-flash` with temperature 0 and thinking disabled. The registry-admission audit uses the same model and decoding settings. In the Gemini and Qwen settings, the same roles use `gemini-3-flash-preview` and `Qwen/Qwen3-4B-Instruct-2507` with the same decoding settings. Document relevance is judged offline by a separate `gpt-4o` judge under the UMBRELA protocol. On the 60-day trajectories, the additional LLM calls of DEC use about 0.8% of the tokens that the full pipeline consumes per request.

Reproduction. Besides the code linked in the reproducibility statement, the package contains the prompts above, the full configuration, the request order and round manifest, and the evaluation scripts. It releases request references, and its instructions recover each request from the public DeepResearchGym log.

J Limitations and Open Directions

DEC operates on the information available to a platform-side policy in a decentralized market: submitted profiles and recorded final-use outcomes, without access to content agents' internal states or documents before selection. Natural extensions span both sides of the market, including learned website-selection policies and alternative update rules and adoption shares for content agents. For content agents, the slot-based updater F_P keeps each profile change attributable to the service records a provider receives, the channel through which DEC acts; more expressive provider strategies are a natural next step. Richer service feedback could further reveal how selection and rewriting jointly shape current service opportunities and the evidence available for later adaptation.

Privacy of the delivery evidence. A supporting evidence card carries the matched past request and the titles that a final answer used from that website; it carries no client identifier. The request stream is keyed by HMAC-IP client proxies (Section B.1), Credit's query-relation judge reads a candidate event without website identity, document title, service outcome, exposure, or offline relevance judgement (Section H.3), and the reproduction package releases request references (Section I).