An Economic Framework for Generative Engines: Advertising or Subscription?

Luyang Zhang [1 2] **Cathy Jiao** [1] **Beibei Li** [1] **Chenyan Xiong** [1]

[1]Carnegie Mellon University
[2]Anaxi Labs

Abstract

Generative Engines (GEs) such as ChatGPT and Google's AI Overviews are rapidly reshaping search economics by delivering synthesized responses that allow users to bypass third-party websites, cutting those sites' advertising revenue. Yet this shift also leaves GEs facing their own monetization problem: whether to insert ads into synthesized responses or keep them ad-free to drive subscription conversions. In this paper, we introduce a dynamic framework to study this problem, which captures how query-level design choices shape user engagement, retention, and subscription conversion over time. Using this framework, we show that the optimal policy follows a cutoff rule: ads should only be shown to users only when the immediate ad payoff exceeds the long-term value of providing ad-free responses. This cutoff shifts toward *with-ad responses* when i) ad revenue is high or ii) users are less sensitive to ads, and toward *ad-free responses* when iii) subscription conversion becomes relatively more valuable. In addition, the presence of rival GEs shifts the optimal policy further toward ad-free responses, as ad-heavy monetization becomes less sustainable when users can freely switch to alternatives. Our findings reveal incentives for real-life generative engine providers to adopt designs that enhance user experience and long-term sustainability.

1. Introduction

Generative engines (GEs), such as ChatGPT and Google's AI Overviews, use large language models to present information directly in response to user queries (Zhou & Li, 2024; Scarcella, 2026), and are being rapidly adopted in web search (Venkatachary, 2024; OpenAI, 2024; Kaiser et al., 2025a). Consequently, GEs have changed user behavior by satisfying information needs directly, reducing how often users visit external websites, and thereby challenging the ad-based monetization models that websites and search engines have traditionally relied on (Gleason et al., 2023; Chapekis & Lieb, 2025; Mcdonald, 2025)

This paradigm shift presents a new problem for ad monetization in GEs. Current generative engine providers have already adopted diverse monetization approaches; for instance, OpenAI and Google have integrated advertising into their generative AI products (Google, 2025; OpenAI, 2026), while Anthropic remains strictly ad-free (Chmielewski et al., 2026). Yet it remains unclear which approach is most viable in the long term. In particular, because GEs incur substantial model inference costs (Luccioni et al., 2023), providers face a trade-off between *with-ad* and *ad-free* subscription models, a choice with consequences for user retention, subscription growth, and monetization sustainability (Babu & Kachwala, 2026).

To address these issues, we develop a framework to analyze the economics of monetization in GEs. Specifically, we model the interaction between GEs and users as a dynamic Stackelberg game. At each query, the GE leads by choosing to provide either an with-ad or ad-free response, after which the user decides whether to engage – and eventually become a subscribed customer – or take an outside option – such as a competing GE – following a discrete-choice model (McFadden, 1972). The GE anticipates these responses and adopts a monetization policy accordingly. Crucially, these interactions accumulate over time: ad-free responses can strengthen future engagement and subscription demand (Cox, 1972; Rust, 1987), while repeated ad exposure can erode retention. This dynamic structure allows us to analyze how optimal monetization design varies across user and query types, as well as across market conditions.

Next, within our framework, we characterize the generative engine's *optimal design policy* and study its economic implications. We show that the optimal policy takes a context-dependent threshold form. Specifically, the GE provides with-ad response only when i) its immediate ad revenue exceeds the long-term value of retaining the user ad-free (i.e., it exceeds the expected future revenue from stronger user retention and eventual subscription conversion), and ii) when users are less ad-sensitive (Gupta et al., 2006; Ascarza

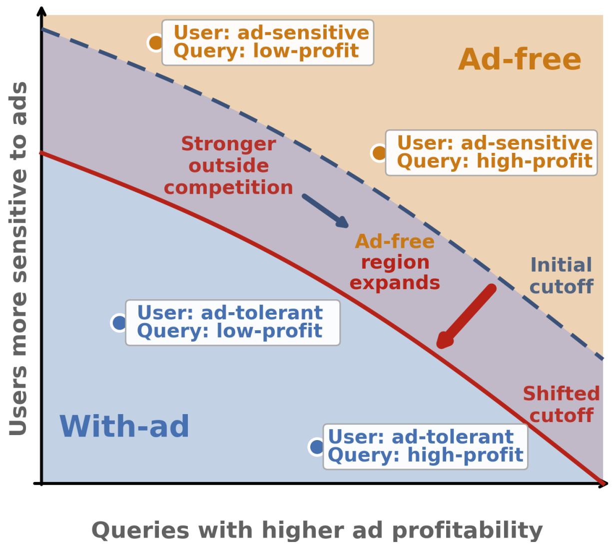

Figure 1: Optimal monetization policy thresholds for GEs. The dashed line shows the baseline cutoff between with-ad and ad-free responses; the solid line shows how outside competition shifts this cutoff, expanding the ad-free region.

& Hardie, 2013). Conversely, GEs shift toward ad-free responses when ad-free interactions generate more long-term value. Moreover, stronger external competition also shifts the threshold toward ad-free responses, since users can more easily switch to competing GEs (see Figure 1).

Finally, we conduct market simulations to examine how GE monetization policies, ranging from ad-heavy to ad-free, perform over time. We find that while ad-heavy policies generate higher short-term revenue, they erode the user base and slow subscription growth in the long run. Moreover, the costs of ad-heavy policies are not easily undone, as switching to ad-free responses later does not replicate the gains of having served them consistently from the start. These patterns hold across user types, query types, and market conditions, and intensify under stronger outside competition.

To summarize, our contributions are three-fold:

1. We provide a game-theoretic framework modeling the interaction between GEs and users around with-ad vs. ad-free responses
2. We characterize the optimal design policy in threshold form, highlighting trade-offs between short-term ad revenue and long-term user engagement.
3. We conduct market simulations comparing ad-heavy to ad-free policies over time, finding that ad-heavy strategies hurt long-term retention and subscription growth even if they boost short-term revenue

2. Related Work

Generative Engines and Monetization Design. Generative Engines synthesize cohesive responses from retrieved sources rather than ranking individual documents (Lewis et al., 2020; Aggarwal et al., 2024). Unlike traditional web search—primarily monetized through ads and user clicks at comparatively low serving costs (Varian, 2007; Evans, 2008; Athey & Ellison, 2011)—Generative Engines incur substantial inference costs (Kwon et al., 2023) and are also associated with a sharp decline in downstream clicks (Coyne, 2025; Ning et al., 2026). This shift motivates LLM-native advertising embedded directly in generated responses. Recent work proposes mechanism designs for ad auctions within retrieval-augmented generation streams (Hajiaghayi et al., 2024; Dubey et al., 2024; Duetting et al., 2024; Balseiro et al., 2025), with extensions to conversational and other generative auction formats (Mordo et al., 2024; Bhawalkar et al., 2025; Ma et al., 2025; Zhao et al., 2025).

User Behavior and Economic Model of Discrete Choice. Users' search engagement has traditionally been modeled through click and browsing models for ranked lists (Craswell et al., 2008; Dupret & Piwowarski, 2008; Chapelle & Zhang, 2009; Joachims et al., 2017), with related work using signals such as dwell time and query reformulation to evaluate user satisfaction (Agichtein et al., 2006; Li et al., 2009). With AI-generated content, users increasingly obtain answers without clicking (Chapekis & Lieb, 2025; Kaiser et al., 2025b). Engagement shifts away from link-based navigation (Gleason et al., 2023; Kirsten et al., 2025), motivating a focus on substitution across response formats rather than click allocation alone. Accordingly, economic models of random-utility discrete choice provides a natural foundation for modeling user choice among heterogeneous response formats (McFadden, 1972; Ben-Akiva & Lerman, 1985), consistent with its use in search-result substitution settings (Ghose et al., 2011; Yao & Mela, 2011; Ghose et al., 2012; Jerath et al., 2014; Jeziorski & Segal, 2015).

Dynamics, Retention, and Subscription Economics.

User churn models connect short-term engagement to future retention and lifetime value (Gupta et al., 2006; Fader & Hardie, 2009; Fader et al., 2010; Ascarza & Hardie, 2013). Building on these dynamics, subscription economics emphasize renewal costs and experience accumulation (Dubé et al., 2010; Della Vigna & Malmendier, 2006; Farrell & Klemperer, 2007; Miller et al., 2023), including freemium and trial-based paths to paid upgrades (Einav et al., 2025; Datta et al., 2015; Shi et al., 2019). Improving the free-tier experience highlights a core tension: richer free offerings can erode short-term ad revenue (Taylor, 2013).

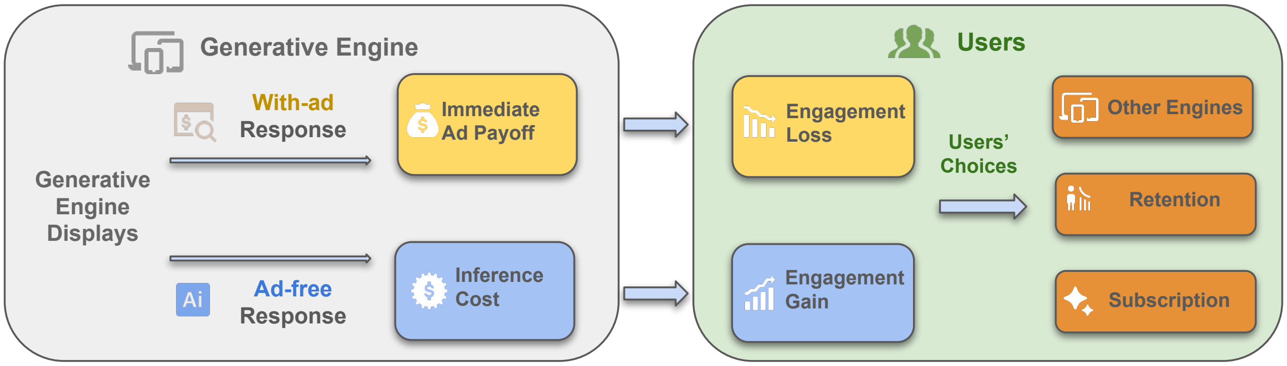

Figure 2: Overview of the generative engine's monetization framework. The generative engine chooses between a with-ad response (yielding immediate ad payoff) and an ad-free response (incurring inference cost). Each response affects user engagement: with-ad responses reduces engagement, while ad-free responses builds it. Users then choose among the generative engine, subscription, or outside alternatives, determining long-term retention and conversion.

3. Economic Framework

In this section, we introduce and formalize our dynamic generative engine monetization framework, which takes into account the ongoing interaction between a generative engine's design choices and its users' responses over time.

We build the framework in four parts. Section 3.1 formally defines the general market setting with GEs and users. Section 3.2 specifies how users respond to the generative engine's query-level design choice. Section 3.3 describes how users' per-query response accumulate their experiences over time, shaping their future engagement and subscription decisions. Section 3.4 formulates the generative engine's optimization problem over this repeated interaction.

3.1. Framework Setup and Interaction.

We begin by defining the basic elements of the framework, which includes generative engines, users, and their interaction sequence over time.

Generative Engines. Since current major generative engines offer differ tiers to its users, such as free, with-add, or paid subscriptions (Google, 2025; OpenAI, 2026), a central design decision in our framework is how the engine serves its non-subscribed users. Specifically, for each user query, the engine selects between two response formats: a with-ad response ($\mathrm{RE}_{\mathrm{ad}}$) or an ad-free response ($\mathrm{RE}_{\mathrm{free}}$), both at no charge. Over time, a non-subscribed user may convert to a paid subscription at price p, after which they receive ad-free responses ($\mathrm{RE}_{\mathrm{paid}}$) by default and the engine no longer faces a design choice for that user (Miller et al., 2023; Einav et al., 2025).

Users. On the user side, each user has (i) a *context*, which captures persistent user characteristics (e.g., demographics, usage patterns, and preferences), influencing how they engage with the generative engine and (ii) *state variables*, which captures their state at a given time. In particular, a user's context $x \in \mathcal{X}$, is drawn from a population distribution μ and remains stable over time, as is standard in user modeling (McFadden, 1972; Gupta et al., 2006). The user state variables, by contrast, evolve over time driven by their interactions with the generative engine. Details on user context and states are further discussed in Section 3.3.

Interaction sequence. Next, given a generative engine and an user, we now describe their interaction. Fix a user context x. In each period $t = 0, 1, 2, \ldots$:

1. **Query and display.** The user generates and submits a query $q_t \in \mathcal{Q}$ (we assume queries are conditionally i.i.d. given user context, i.e., $q_t \sim \mathcal{P}(\cdot \mid x)$ [1]) to the generative engine.
2. **Generative Engine Response.** If the user is active and non-subscribed, the generative engine chooses a response format $A_t \in \{\mathrm{RE}_{\mathrm{ad}}, \mathrm{RE}_{\mathrm{free}}\}$.
3. **User engagement.** The user either engages with the displayed response or takes an outside alternative such as other competing GEs. This choice determines the generative engine's immediate payoff (see Section 3.2).
4. **Cross-period dynamics: retention and subscription.** The engagement updates the user's state variables. Depending on the updated state, the user may convert to a paid subscription ($\mathrm{RE}_{\mathrm{paid}}$) according to the conversion rule in Section 3.3.

Finally, the interaction repeats: the updated state variables and subscription decision jointly determine the user's continued engagement with the generative engine in the next period (see Section 3.3).

[1]The i.i.d assumption can be relaxed without changing the framework's structure. More general stationary query dynamics can be handled by augmenting the state (e.g., including q_t). We assume i.i.d. queries to keep the state minimal; modeling query transitions is orthogonal to our main results. See Appendix A for details.

3.2. User Per-query Choice and generative engine's Payoffs

Building on the interaction model above, we now characterize the user engagement towards the generative engine, and the resulting payoff to the engine.

User per-query choice. Given a user context x, a query q, and a response $A_t = a$ displayed by the generative engine, the user decides whether to engage with the response a or not. We denote the decision variable as $Y_t \in 1, 0$, where $Y_t = 1$ denotes reading and using the generated content, and $Y_t = 0$ denotes taking an outside alternative, such as switching to a competing generative engine.

The user's engagement is driven by two utilities: that of the generative engine's displayed response, $v_a(x, q)$, and that of the outside alternative, $v_0(x, q)$. The engagement probability follows a binary logit, widely adopted in discrete-choice modeling (Luce et al., 1959; McFadden, 1972; Ben-Akiva & Lerman, 1985):

$$
\begin{aligned}
\Pr(Y_t = 1 \mid x, q, a) &= \frac{\exp\big(v_a(x,q)\big)}{\exp\big(v_a(x,q)\big) + \exp\big(v_0(x,q)\big)}, \\
\Pr(Y_t = 0 \mid x, q, a) &= 1 - \Pr(Y_t = 1 \mid x, q, a). \qquad (1)
\end{aligned}
$$

Generative engine payoffs. Given a generative engine's response a to user x on query q, its expected per-period payoff is denoted by $r_a(x, q)$, which is an expectation taken over the engagement decision Y_t. That is,

$$
r_a(x,q) = \begin{cases} -\kappa + R(x,q) & \\ \qquad \Pr(Y_t = 1 \mid x, q, a) & a = \mathrm{RE}_{\text{ad}}, \\ -\kappa & a = \mathrm{RE}_{\text{free}}, \end{cases} \qquad (2)
$$

where $R(x, q)$ is the expected ad revenue *conditional on* engagement, and the choice probability is given by (1). In addition, both response formats incur an inference cost κ (Luccioni et al., 2023) [2]. Only a response with ads, RE_{ad}, generates immediate ad revenue upon engagement. In contrast, displaying $\mathrm{RE}_{\text{free}}$ generates no immediate revenue; its value comes from retaining users and building toward subscription, as formalized in Section 3.3.

3.3. User Cross-Query Dynamics and Subscription

Next, we define the state variables that track a user's past interactions with the generative engine. These variables shape the user's future retention and subscription decisions.

State definition and update process. Each user's interaction history is summarized by three state variables: (1) an AI-experience state S_t, which captures cumulative experience from interactions with *ad-free* responses; (2) an ad-exposure state C_t, tracking cumulative exposure to *with-ad* responses; and (3) a subscription indicator $Z_t \in 0, 1$ ($Z_t = 1$ means subscribed to $\mathrm{RE}_{\text{paid}}$). Formally, the user enters period t with state:

$$
(S_t, C_t, Z_t) \in [0, \infty) \times [0, \infty) \times 0, 1.
$$

When $Z_t = 0$ (i.e., the user is unsubscribed), the generative engine chooses $A_t \in \{\mathrm{RE}_{\text{ad}}, \mathrm{RE}_{\text{free}}\}$ for query q_t, and the two user experience states S_t and C_t update only upon user engagement (i.e., $Y_t = 1$). Engaging with an ad-free response $\mathrm{RE}_{\text{free}}$ increments S_t, while engaging with a with-ad response RE_{ad} increments C_t. Formally,

$$
\begin{aligned}
S_{t+1} &= S_t + \mathbf{1}\{A_t = \mathrm{RE}_{\text{free}}, Y_t = 1\}\, \delta_S(x, q_t, S_t), \\
C_{t+1} &= C_t + \mathbf{1}\{A_t = \mathrm{RE}_{\text{ad}}, Y_t = 1\}\, \delta_C(x, q_t, C_t),
\end{aligned}
$$

where $\delta_S, \delta_C : \mathcal{X} \times \mathcal{Q} \times \mathbb{R}_+ \to [0, \infty)$. If the user chooses the outside option ($Y_t = 0$), neither state changes.

Conversion to the subscription. Over time, a user may convert to the paid subscription $\mathrm{RE}_{\text{paid}}$. Conversion occurs once the user's accumulated AI experience S_t exceeds a threshold that depends on user context x, subscription price p, and the ad exposure C_t, consistent with experience-driven adoption in freemium markets (Einav et al., 2025; Miller et al., 2023). Formally,

$$
Z_t = \mathbf{1}\{Z_{t-1} = 1 \text{ or } S_t \geq \tau(x, p, C_t)\}, \qquad (3)
$$

where $\tau(x, p, c)$ is the subscription threshold. The threshold is decreasing in c, reflecting that cumulative ad exposure can accelerate subscription demand (Cho & as, 2004; Anderson & Coate, 2005). Appendix A.1.1 replaces this threshold format with a smooth adoption hazard without changing the analysis.

We treat subscription as permanent: if $Z_t = 1$, then $Z_{t'} = 1$ for all $t' \geq t$. After subscription, the generative engine earns a per-period net payoff $r_{\text{sub}}(p) := p - \kappa_{\text{sub}}$, where $\kappa_{\text{sub}} \geq \kappa$ is the inference cost for subscription-tier. This setup is standard in subscription markets and keeps the dynamic program tractable (Einav et al., 2025; Della Vigna & Malmendier, 2006; Farrell & Klemperer, 2007). Appendix A.1.2 relaxes this by allowing post-subscription churn without changing the main theoretical arguments.

Cross-period retention and termination. Whether a user remains active depends on their accumulated interactions with the generative engine. After updating $(S_{t+1}, C_{t+1}, Z_{t+1})$, the user remains active in period $t + 1$ with probability $\rho_x(S_{t+1}, C_{t+1}, Z_{t+1}) \in [0, 1]$, where $\rho_x(s, c, 1) = 1$ for subscribed users. Thus, subscribed users remain active. Retention increases with AI experience S_{t+1}

[2]In practice, the inference cost κ can also depend on user context and query. Our framework and theoretical results extend to the cost $\kappa(x, q)$ (see Section A.2 for details).

and decreases with ad exposure C_{t+1}, reflecting that positive interactions encourage continued use while excessive ad exposure drives users away (Gupta et al., 2006; Ascarza & Hardie, 2013). We model churn as permanent, which is standard in customer-base analysis (Fader & Hardie, 2009; Fader et al., 2010).

3.4. Generative Engine's Objective and Design Policy

Given the user's per-query response and the cross-period dynamics defined above, the generative engine now faces a dynamic decision problem: how to choose response formats over time to maximize long-run value.

Policy class and value function. For pre-subscription users, the generative engine's policy maps the user's current state and query to a response format. Formally, a stationary Markov policy is a measurable map

$$g_x : [0,\infty) \times [0,\infty) \times \mathcal{Q} \to \{\mathrm{RE}_{\text{ad}}, \mathrm{RE}_{\text{free}}\}, g_x \in \mathcal{G}_x,$$

The value function $V_x(s,c,z)$ is the generative engine's maximum expected discounted payoff starting from state $(S_0, C_0, Z_0) = (s,c,z)$:

$$V_x(s,c,z) := \sup_{g_x \in \mathcal{G}_x} \mathbb{E}^{g_x}\left[\sum_{t\geq 0} \beta^t \Big(\mathbf{1}\{Z_t = 0\}\, r_{A_t}(x,q_t) + \mathbf{1}\{Z_t = 1\}\, r_{\text{sub}}(p)\Big)\right], \tag{4}$$

where future payoffs are discounted by $\beta \in (0,1)$, $A_t \in \mathrm{RE}_{\text{ad}}, \mathrm{RE}_{\text{free}}$ is the response format chosen by g_x when $Z_t = 0$, and $r_{A_t}(x,q_t)$ is the expected per-period payoff from (2). Since per-period payoffs are bounded and $\beta < 1$, the value function $V_x(s,c,z)$ is well-defined under standard contraction arguments, ensuring the Bellman equations below are well-posed.

Bellman equations. For a subscribed user ($z = 1$), the generative engine makes no display decision and receives $r_{\text{sub}}(p)$ each period, implying

$$V_x(s,c,1) = r_{\text{sub}}(p) + \beta V_x(s,c,1) = \frac{r_{\text{sub}}(p)}{1-\beta}. \tag{5}$$

For a non-subscribed user ($z = 0$), conditional on (x,s,c) and a query q, the generative engine chooses $a \in \{\mathrm{RE}_{\text{ad}}, \mathrm{RE}_{\text{free}}\}$. The user's within-query choice $Y \in \{1,0\}$ follows (1), and the one-period payoff is r_a as defined in (2). For each query, the generative engine compares the immediate ad revenue from RE_{ad} against the discounted continuation value from $\mathrm{RE}_{\text{free}}$, which generates no immediate revenue but boosts retention and subscription conversion. Using the state update rules from Section 3.3, we define the post-update states:

$$\begin{aligned}
S^+_{a,y}(x,q,s) &:= s + \mathbf{1}\{a = \mathrm{RE}_{\text{free}}, y = 1\}\delta_S(x,q,s),\\
C^+_{a,y}(x,q,c) &:= c + \mathbf{1}\{a = \mathrm{RE}_{\text{ad}}, y = 1\}\delta_C(x,q,c),\\
Z^+_{a,y}(x,q,s,c) &:= \mathbf{1}\{\, S^+_{a,y}(x,q,s) \geq \tau(x,p,C^+_{a,y}(x,q,c))\,\}.
\end{aligned} \tag{6}$$

Let $Y \sim \Pr(\cdot \mid x,q,a)$ denote the within-query choice, and write $(S^+, C^+, Z^+) := (S^+_{a,Y}(x,q,s),\, C^+_{a,Y}(x,q,c),\, Z^+_{a,Y}(x,q,s,c))$. The pre-subscription Bellman equation is

$$V_x(s,c,0) = \mathbb{E}_{q\sim\mathcal{P}(\cdot|x)}\Bigg[\max_{a\in\{\mathrm{RE}_{\text{ad}},\mathrm{RE}_{\text{free}}\}} \Big\{ r_a(x,q) + \beta\, \mathbb{E}_{Y\sim\Pr(\cdot|x,q,a)}\Big[\rho_x\big(S^+, C^+, Z^+\big) \cdot V_x\big(S^+, C^+, Z^+\big)\Big]\Big\}\Bigg], \tag{7}$$

Connection to a static benchmark. When cross-period interactions are absent, the design problem simplifies: each query is an independent decision with no future consequences. Specifically, the dynamic program reduces to a per-query static comparison between RE_{ad} and $\mathrm{RE}_{\text{free}}$ when (i) retention is exogenous (independent of accumulated experience) and (ii) subscription has no option value (e.g., $\tau = \infty$ or $r_{\text{sub}} = 0$). In this case, the generative engine simply chooses whichever response format yields higher immediate expected payoff on each query.

4. Theoretical Results

This section characterizes which users and queries should receive with-ad versus ad-free responses, and how this allocation shifts across market conditions. Understanding this variation helps identify the economic conditions under which different response formats are optimal. We first derive the generative engine's optimal display policy, then study how it varies with user and query types (Section 4.2) and market conditions, including AI economics for inference cost and subscription price (Section 4.3) and outside competition (Section 4.4).

4.1. Optimal Policy Characterization

We show that the optimal display policy takes a comparison rule: for a given user and query, the generative engine displays RE_{ad} if and only if its total value — immediate ad revenue plus long-run value from retention and subscription — exceeds that of $\mathrm{RE}_{\text{free}}$.

A policy is optimal if it maximizes the value function $V_x(s,c,0)$ in (4). For analysis, we impose standard regularity conditions (e.g., boundedness, measurability) throughout, which are common in dynamic programming and discrete-choice models (Puterman, 2014; Rust, 1987); see Appendix B.1 for details. The following proposition characterizes this optimal policy.

Proposition 4.1 *(Optimal policy as a comparison rule).* Under Assumption 2, for any query q and pre-subscription state $(s, c, 0)$, define $Q_x^a(s, c, q)$ as the generative engine's expected payoff from displaying response $a \in \mathrm{RE_{ad}}, \mathrm{RE_{free}}$, including the discounted continuation value:

$$\begin{aligned} &Q_x^a(s,c,q) \\ &\quad := r_a(x,q) + \beta \mathbb{E}_{Y\sim \Pr(\cdot|x,q,a)}\big[\\ &\qquad \rho_x\big(S^+_{a,Y}, C^+_{a,Y}, Z^+_{a,Y}\big)\, V_x\big(S^+_{a,Y}, C^+_{a,Y}, Z^+_{a,Y}\big)\big], \quad (8) \end{aligned}$$

which is the expected payoff from choosing action a on query q, including discounted continuation value.

Define $\mathrm{RE_{ad}}$'s value edge over $\mathrm{RE_{free}}$ as $\Delta_x(s, c, q) := Q_x^{\mathrm{RE_{ad}}}(s, c, q) - Q_x^{\mathrm{RE_{free}}}(s, c, q)$. Then the *optimal policy* g_x chooses $\mathrm{RE_{ad}}$ (i.e., $g_x^*(s, c, q) = \mathrm{RE_{ad}}$) if and only if $\Delta_x(s, c, q) \geq 0$. See Appendix B.3. □

Next, we decompose $\mathrm{RE_{ad}}$'s value edge over $\mathrm{RE_{free}}$ $\Delta_x(s, c, q)$ into short-term monetization and long-term value from retention and conversion.

Remark 4.2 *(Short-term vs. long-term trade-off).* For any pre-subscription state and query (x, s, c, q), we can identify the generative engine's short-term and long-term value trade-off by expanding the value edge $\Delta_x(s, c, q)$. We denote the long-term value function as

The value edge $\Delta_x(s, c, q)$ decomposes into two components: a short-term ad revenue edge and a long-term value edge. Define the expected future value from retaining the user as

$$\begin{aligned} &W_x(S^+_{a,Y}, C^+_{a,Y}, Z^+_{a,Y}) \qquad (9) \\ &\quad := \rho_x(S^+_{a,Y}, C^+_{a,Y}, Z^+_{a,Y})\, V_x(S^+_{a,Y}, C^+_{a,Y}, Z^+_{a,Y}). \end{aligned}$$

Then:

$$\begin{aligned} &\Delta_x(s,c,q) = \\ &\underbrace{R(x,q)\,\Pr(Y=1 \mid x,q,\mathrm{RE_{ad}})}_{\text{short-term monetization edge of } \mathrm{RE_{ad}} \text{ over } \mathrm{RE_{free}}} \\ &\underbrace{\begin{aligned}&+\beta\Big(\mathbb{E}_{Y\sim\Pr(\cdot|x,q,\mathrm{RE_{ad}})}\big[W_x(S^+_{\mathrm{RE_{ad}},Y}, C^+_{\mathrm{RE_{ad}},Y}, Z^+_{\mathrm{RE_{ad}},Y})\big] \\ &-\mathbb{E}_{Y\sim\Pr(\cdot|x,q,\mathrm{RE_{free}})}\big[W_x(S^+_{\mathrm{RE_{free}},Y}, C^+_{\mathrm{RE_{free}},Y}, Z^+_{\mathrm{RE_{free}},Y})\big]\Big)\end{aligned}}_{\text{long-term value edge (retention and subscription)}}. \\ &\qquad (10) \end{aligned}$$

The first component is $\mathrm{RE_{ad}}$'s expected ad revenue. The second component captures the long-term trade-off: $\mathrm{RE_{free}}$ builds AI experience and drives subscription conversion, whereas $\mathrm{RE_{ad}}$ raises ad exposure that may reduce future engagement. □

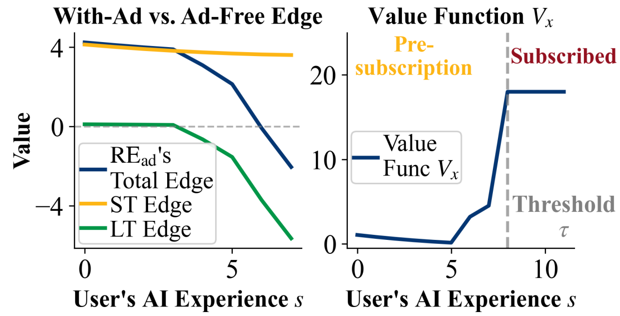

Figure 3: Simulation of the optimal policy in Proposition 4.1 for a fixed user and ad exposure state, as AI experience increases. *Left:* $\mathrm{RE_{ad}}$'s total edge over $\mathrm{RE_{free}}$ and its short-term (ST) and long-term (LT) components from Remark 4.2. *Right:* Pre-subscription value function V_x; the dashed line marks the subscription threshold, beyond which the value equals the subscribed level.

Figure 3 illustrates this decomposition. As AI experience s increases, the long-term component increasingly favors $\mathrm{RE_{free}}$ through higher retention and earlier subscription conversion (Figure 3, left). Once s exceeds the subscription threshold τ, the user subscribes and the value reaches the subscribed level (Figure 3, right).

4.2. Implications of User and Query Types

We now characterize how the optimal design policy varies across user and query types, identifying when the policy switches between $\mathrm{RE_{ad}}$ and $\mathrm{RE_{free}}$.

User and query types. We capture variation in the generative engine's environment using user and query types. Each user has type $\gamma(x)$ capturing *ad sensitivity*, i.e., the user's responsiveness to ads—how strongly additional *ad exposure* lowers future engagement (Anderson & Coate, 2005; Cho & as, 2004; Edwards et al., 2002).

On the query side, $r_{\mathrm{RE_{ad}}}(x, q)$ measures *ad profitability* (the expected short-term revenue from showing ads on query q) (Ghose & Yang, 2009). Further, $\psi(x, q)$ measures the *AI experience gain* from showing an ad-free generative response, which accelerates user's experience accumulation and subscription conversion. Formally, these types enter the model through the retention and conversion functions (ρ, τ) and the experience updates δ_S, δ_C; see Appendix B.1 for detailed assumptions.

Proposition 4.3 *(Threshold structure over user and query features).* Under Assumptions 2 and 3, fix a pre-subscription state $(s, c, 0)$. Define the user–query type tuple as

$$t \equiv (\gamma, r, \psi) := \big(\gamma(x),\ r_{\mathrm{RE_{ad}}}(x,q),\ \psi(x,q)\big).$$

Define a partial order on types by

$$(\gamma, r, \psi) \preceq (\gamma', r', \psi') \iff \gamma \leq \gamma',\ \psi \leq \psi', \text{ and } r \geq r'.$$

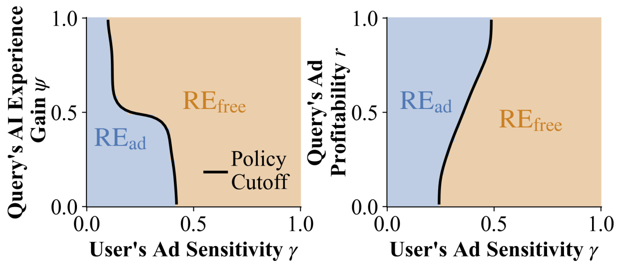

Figure 4: Illustration of the optimal policy cutoff across user and query types from Proposition 4.3. *Left:* User *ad sensitivity* γ vs. query *AI experience gain* ψ. *Right:* User *ad sensitivity* γ vs. query *ad profitability* r. The solid curve separates RE_{ad} (blue) and $\mathrm{RE}_{\text{free}}$ (orange) regions.

Then the following hold:

1. **(Monotone ordering across user–query types).** If $\mathrm{RE}_{\text{free}}$ is optimal at type t, then it is also optimal at any higher type $t' \succeq t$.
2. **(Type thresholds along a single index).** Holding (γ, r) fixed, there exists a threshold $\psi^*(s, c; \gamma, r) \in [-\infty, \infty]$ such that

$$\mathrm{RE}_{\text{ad}} \text{ is optimal } \iff \psi \le \psi^*(s, c; \gamma, r).$$

Holding (γ, ψ) fixed, there exists a threshold $r^*(s, c; \gamma, \psi) \in [-\infty, \infty]$ such that

$$\mathrm{RE}_{\text{ad}} \text{ is optimal } \iff r \ge r^*(s, c; \gamma, \psi).$$

Remark 4.4 *(Interpretation: when to show RE_{free} and RE_{ad}).* Proposition 4.3 delivers a key ordering result: the generative engine decides toward $\mathrm{RE}_{\text{free}}$ for users who are more *ad-sensitive* and for queries where an ad-free response generates larger user experience gains, and shifts toward RE_{ad} when the query's immediate *ad revenue* is higher. Moreover, this ordering implies a clear cutoff: as a user or query becomes more favorable to $\mathrm{RE}_{\text{free}}$ along these types, there is a point at which the generative engine switches from showing RE_{ad} to $\mathrm{RE}_{\text{free}}$.

Figure 4 presents the type-ordered structure in Proposition 4.3 by overlaying the optimal display decision on the population density of user-type indices. Holding the other parameters fixed, the $\mathrm{RE}_{\text{free}}$ region expands monotonically with ad sensitivity γ, yielding a clean cutoff that splits the population into RE_{ad}-preferred and $\mathrm{RE}_{\text{free}}$-preferred regions. See the illustration over query types $(r_{\mathrm{RE}_{\text{ad}}}, \psi)$ in Appendix B.8.2. We also show the illustration of Figure 3 under user heterogeneity at Appendix B.8.1.

Generative engine's weight on future outcomes. We also show that a more forward-looking generative engine (higher discount factor β) displays $\mathrm{RE}_{\text{free}}$ more often, as $\mathrm{RE}_{\text{free}}$ improves long-term value through retention and experience accumulation (Appendix B.4).

4.3. Implications of AI Economics: Inference Cost and Subscription Price

Inference costs and subscription pricing jointly influence the allocation between RE_{ad} and $\mathrm{RE}_{\text{free}}$. Inference costs affect the relative value of keeping a user in the free tier versus moving that user into the paid tier, while subscription pricing changes the revenue from conversion.

Inference cost. As defined in Section 3.2, both RE_{ad} and $\mathrm{RE}_{\text{free}}$ incur the same free-tier inference cost κ per query, while the paid subscription tier carries cost $\kappa_{\text{sub}} \ge \kappa$. To analyze how these costs affect the design policy, consider a user close to subscribing, i.e., whose AI experience s is near the subscription threshold $\tau(x, p, c)$, where inference costs have the most direct effect on the design choice (see Assumption 6 for the formal setup).

Proposition 4.5 *(Free-tier and paid-tier inference costs have opposite effects near subscription).* Under Assumptions 2, 5 and 6, recall the RE_{ad}'s value edge $\Delta_x(s, c, q)$ from Proposition 4.1; here we index it by inference costs and write $\Delta_x(s, c, q; \kappa, \kappa_{\text{sub}})$.

1. **Higher free-tier cost κ raises incentives to display $\mathrm{RE}_{\text{free}}$ in the near-threshold region.** $\Delta_x(s, c, q; \kappa, \kappa_{\text{sub}})$ is weakly decreasing in κ. That is, at near-threshold states where displaying $\mathrm{RE}_{\text{free}}$ triggers subscription while displaying RE_{ad} leaves the user pre-subscription, increasing κ pushes the decision toward $\mathrm{RE}_{\text{free}}$.
2. **Higher paid-tier cost κ_{sub} reduces incentives to display $\mathrm{RE}_{\text{free}}$ in the near-threshold region.** Under Assumption 6, $\Delta_x(s, c, q; \kappa, \kappa_{\text{sub}})$ is weakly increasing in κ_{sub}. That is, at near-threshold states where displaying $\mathrm{RE}_{\text{free}}$ would trigger subscription, increasing κ_{sub} pushes the decision further toward RE_{ad}.

Near the subscription threshold, the two costs work in opposite directions. A higher free-tier cost makes it more attractive to move the user out of the free tier, while a higher paid-tier cost makes immediate conversion less attractive.

Subscription price. Another key economic factor is the subscription price p. As defined in Section 3.3, p enters through both the subscription conversion threshold $\tau(x, p)$ and the post-conversion revenue $r_{\text{sub}}(p)$, so it affects the generative engine's incentive to display $\mathrm{RE}_{\text{free}}$. We study the price effect in the same one-step conversion region (as in Proposition 4.5), where displaying $\mathrm{RE}_{\text{free}}$ triggers immediate subscription.

Proposition 4.6 *(Higher subscription price shifts the design toward RE_{free} near the threshold.).* Under Assumptions 2, 5 and 7, fix a pre-subscription decision point (x, s, c, q) and consider any price interval I such that for all $p \in I$, the one-step conversion conditions in Assumption 6 hold at (x, p, s, c, q). Then the RE_{ad}'s value edge $\Delta_x(s, c, q;\, p)$ is

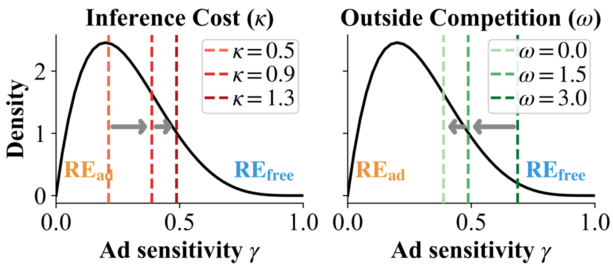

Figure 5: Policy cutoff shift illustration under inference cost and outside competition, at a fixed pre-subscription state (s, c). Ad sensitivity follows $\gamma \sim \text{Beta}$. The solid curve shows the user-type density; users to the left of the cutoff are served RE_{ad}, while those to the right are served RE_{free}. Dashed lines indicate the cutoff γ^* at each parameter level, and arrows indicate the direction of the cutoff shift. *(a)* Varying inference cost κ; *(b)* varying outside competition ω.

weakly decreasing in p on I. Equivalently, in this near-threshold region, a higher subscription price weakly shifts the optimal design toward RE_{free}.

These propositions share a common insight: AI economics has the sharpest implications for users who are close to subscribing. For these near-threshold users, showing an ad-free response can easily push them over the subscription threshold. In this region, a higher free-tier cost κ, a lower paid-tier cost κ_{sub}, and a higher subscription price p all make it more attractive to move the user into the paid tier by displaying RE_{free}. Away from the near-threshold region, price can also shift the conversion threshold $\tau(x, p)$ but the overall effect of p can be more nuanced.

4.4. Impact of Outside Competition

When outside alternatives strengthen, users engage less with displayed responses, reducing the generative engine's short-term ad revenue and shifting the long-term comparison toward RE_{free}.

Competitive-pressure channel. In our framework (Section 3), outside competition enters through the outside-option utility $v_0(x, q)$ in the per-query choice model, capturing factors such as advances in rival LLMs, ad-free policies by competitors (Chmielewski et al., 2026), or the growth of alternative GEs (Malik, 2025).

To formalize how outcomes change as the outside option improves, we index a uniform increase in outside-option attractiveness by $\omega \geq 0$, shifting $v_0(x, q) \to v_0(x, q) + \omega$ while holding all other parameters fixed. The uniform shift is a tractable baseline; heterogeneous shifts across user–query pairs are a natural extension. As outside competition intensifies, users disengage more from the generative engine, directly reducing RE_{ad}'s short-term ad revenue and shifting the long-term value advantage toward RE_{free}. To capture this full effect, we allow the continuation values to depend on ω and impose a sufficient condition on the long-term value advantage of RE_{free} over RE_{ad}. Under this condition, RE_{ad}'s total value edge shrinks as outside competition grows, shifting the optimal policy toward RE_{free} (see Assumption 8 in Appendix B.1 for details).

Proposition 4.7 *(Intensified competition shifts the policy toward* $\text{RE}_{\textit{free}}$*).* Under Assumptions 2 and 8, for any fixed pre-subscription decision point (s, c, q), the RE_{ad}'s edge $\Delta_x(s, c, q; \omega)$ is weakly decreasing in outside-option intensity ω. Equivalently, stronger outside competition makes the optimal policy weakly more likely to select RE_{free}.

To understand the mechanism, recall the short-term vs. long-term decomposition of Δ_x from Remark 4.2. Indexing both components by ω, competition weakens RE_{ad} through two channels: it directly reduces the share of users who engage with responses that include ads (shrinking the short-term ad-revenue component), and it widens RE_{free}'s long-term value advantage (the long-term component), because ad-free experience accumulation becomes relatively more valuable when users are harder to retain. Both effects are formalized in Assumption 8 and proved in Appendix B.3.

Proposition 4.7 tells us that competition shifts the *overall* policy toward RE_{free}, but a natural follow-up question is: *which* users and queries are affected first? Combining Proposition 4.7 with the type-threshold structure from Proposition 4.3 — under the additional requirement that Assumption 3 holds at each competition level ω (see Appendix B.1) — we can trace how the cutoffs from Section 4.2 move as competition intensifies:

Proposition 4.8 *(Competition shifts the type cutoffs toward* $\text{RE}_{\textit{free}}$*).* Under Assumptions 2, 3 and 8, with Assumption 3 holding for each fixed ω, fix a pre-subscription state $(s, c, 0)$ and index the type cutoffs from Proposition 4.3 by ω:

1. **Experience-gain cutoff decreases.** $\psi^*(s, c; \gamma, r, \omega)$ is weakly decreasing in ω: under stronger competition, RE_{free} becomes optimal even for queries whose ad-free responses generate smaller experience gains.
2. **Ad-revenue cutoff increases.** $r^*(s, c; \gamma, \psi, \omega)$ is weakly increasing in ω: under stronger competition, showing ads is justified only when the query's immediate ad revenue is higher.
3. **Distributional consequence.** For any fixed distribution μ over user–query types (γ, r, ψ), the share of user–query pairs for which RE_{ad} is optimal is weakly decreasing in ω.

In words, as competitors improve, the set of users and queries where showing ads is worthwhile shrinks: the generative engine can justify ads only for queries with high immediate revenue and low experience-gain potential, and

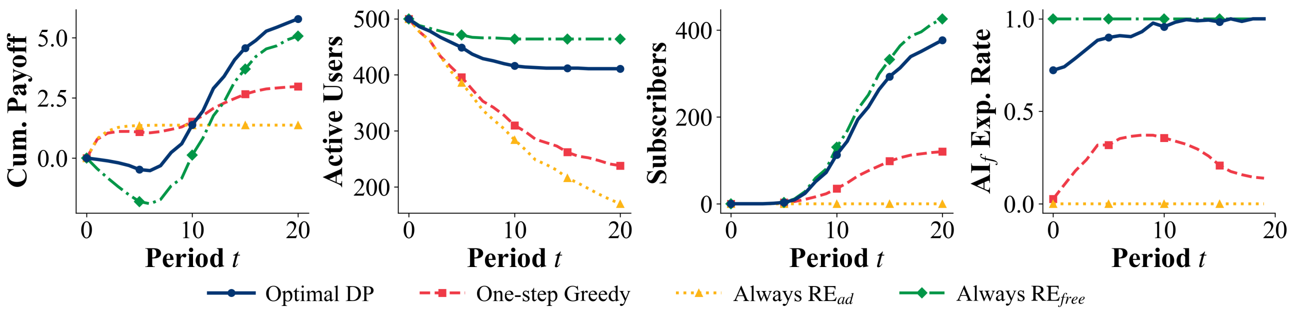

Figure 6: Comparison of four policies: *Optimal DP*, *One-step greedy*, *Always* RE_{ad}, and *Always* RE_{free}; $N = 500$, $T = 20$, $\beta = 0.95$. *(1)* Discounted cumulative generative engine payoff up to t. *(2)* Total active users (non-subscribers plus subscribers). *(3)* Cumulative subscribers by t. *(4)* AI exposure rate: the share of *active non-subscribers* shown $\mathrm{RE}_{\text{free}}$ at time t.

the overall share of interactions served with ads decreases.

Remark 4.9 *(Subscription insulation).* Under the permanent-subscription benchmark, the subscribed-state value $V_x(s, c, 1) = r_{\text{sub}}(p)/(1 - \beta)$ is independent of ω, since subscribers consume $\mathrm{RE}_{\text{paid}}$ directly without facing the outside option. By contrast, the pre-subscription phase is fully exposed to competition through the per-query engagement channel. This asymmetry suggests that competition selectively weakens free-tier value, helping explain why the generative engine's incentive to use $\mathrm{RE}_{\text{free}}$ as a bridge toward subscription grows stronger—particularly for users near the conversion threshold $\tau(x, p)$.

Figure 5 illustrates this cutoff shift numerically: as outside competition strengthens from $\omega = 0$ to $\omega = 1$, the policy threshold γ^* moves steadily to the right, shrinking the share of users served ads.

Implication for competitive markets. These results indicate that ad-heavy strategies are most viable when outside competition is weak but become less sustainable as rival GEs improve. The generative engine is pushed toward ad-free display to retain users who would otherwise switch, and toward subscription as a revenue source insulated from competition—consistent with the observed trend of generative engines investing in premium tiers as the market grows more competitive (Einav et al., 2025). This competitive-pressure effect complements the discount-factor result in Appendix B.4: both outside competition and forward-looking orientation independently push the policy toward $\mathrm{RE}_{\text{free}}$, but through distinct mechanisms.

5. Experiments

We conduct a simulation study to illustrate the dynamic trade-off between short-term monetization and long-term user development in design policy. The analysis traces how different allocations between GRs with ads and ad-free GRs shape provider payoff, retention, subscription conversion, and AI exposure over time.

5.1. Experimental Setup

We simulate a GEs market with $N = 500$ heterogeneous users over $T = 20$ periods and discount factor $\beta = 0.95$. In each period, every active user issues a query; the generative engine chooses RE_{ad} or $\mathrm{RE}_{\text{free}}$, and the user probabilistically decides whether to engage, subscribe, and remain active in the next period. Per-query engagement yields an immediate payoff and updates two states – ad exposure c and AI experience s – which govern future retention and subscription conversion. We follow the model specification and calibration in Sections 3 and 4; full details are in Appendix C.1.

Monetization Design Policies. We compare four design policies deployed by the generative engine:

1. *Optimal DP*: computed from the Bellman equation in Sections 3 and 4.
2. *One-step greedy*: maximizes immediate expected payoff plus next-period value.
3. *Always* RE_{ad}, which always displays GRs with ads.
4. *Always* RE_{free}, which always displays ad-free GRs.

5.2. Results

The simulations highlight three implications of the dynamic design problem: why myopic monetization underperforms, what mechanism drives the gap, and how the optimal policy balances the trade-off.

Cumulative payoff. Figure 6 (panel 1) highlights the payoff trade-off: policies that over-monetize early leave long-term value on the table. Policies that shift weight toward RE_{ad} front-load revenue but weaken continuation value, whereas those that shift toward $\mathrm{RE}_{\text{free}}$ slow early payoff but support stronger downstream gains through retention and conversion. At the extremes, *Always* RE_{ad} concentrates early revenue with limited continuation gains, while *Always* RE_{free} bears

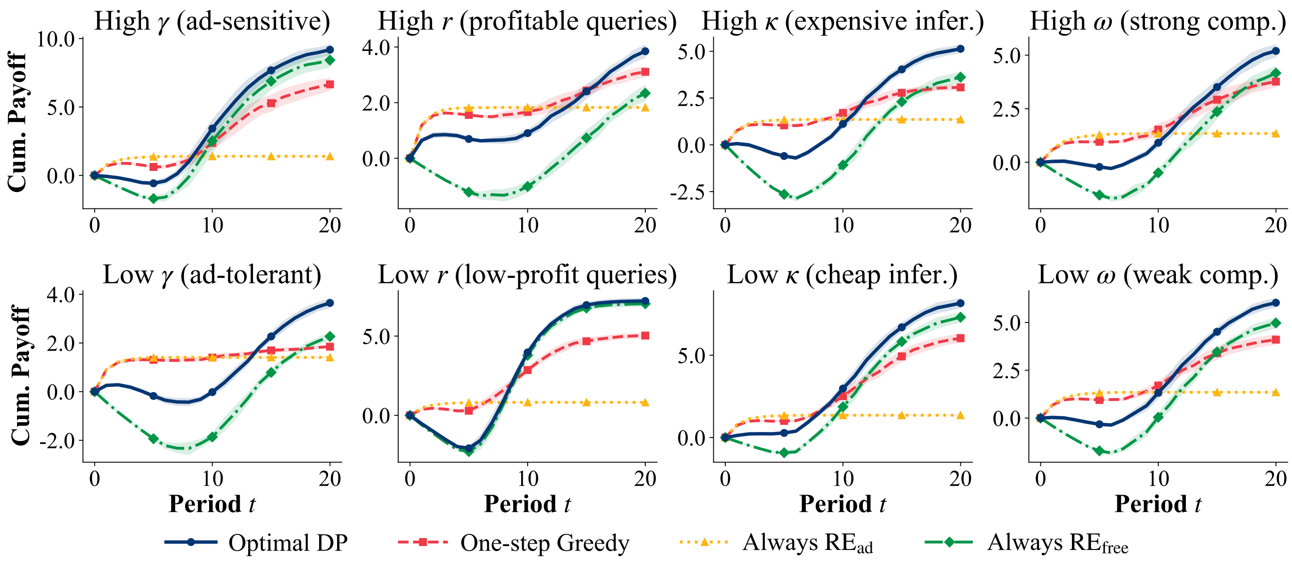

Figure 7: Cumulative payoff trajectories under eight market conditions. Columns correspond to the four market dimensions (γ, r, κ, ω); rows compare high and low levels. Shaded bands show $\pm 1\sigma$ across five random seeds (N=300 each).

the largest short-term cost but produces the strongest retention and subscription outcomes. The *Optimal DP* balances these forces by internalizing both the current-query payoff and the effect of today's design choice on future states. By contrast, *One-step greedy* remains too sensitive to short-term payoffs and therefore underweights these continuation benefits.

Active users and subscriptions. Figure 6 (panels 2–3) reveal the mechanism behind the cumulative payoffs. Policies that allocate more $\mathrm{RE}_{\mathrm{free}}$ build AI experience s faster. This lifts both retention and subscription rates. Policies that lean on $\mathrm{RE}_{\mathrm{ad}}$ slow that accumulation process and shrink the population of future active users. This is why short-sighted over-monetization underperforms: it sacrifices the user base that generates future revenue. At one extreme, *Always* RE_{ad} slows upgrading and accelerates churn; at the other, *Always* RE_{free} maximizes conversion at the cost of short-term revenue. The dynamic policies lie between these extremes by adjusting $\mathrm{RE}_{\mathrm{free}}$ exposure over time rather than committing to a fixed display rule.

AI exposure rate. We report the $\mathrm{RE}_{\mathrm{free}}$ exposure rate, defined as the share of active non-subscribers shown $\mathrm{RE}_{\mathrm{free}}$ by the policy. Figure 6 (panel 4) captures how aggressively each policy allocates $\mathrm{RE}_{\mathrm{free}}$ to active non-subscribers over time. Sustained $\mathrm{RE}_{\mathrm{free}}$ allocation supports retention and conversion over time (e.g., *Optimal DP*, *Always* RE_{free}), whereas short-term $\mathrm{RE}_{\mathrm{free}}$ pushes do not sustain these gains to the same extent (e.g., *One-step greedy*). This contrast shows that sustained allocation of ad-free responses matters more than temporary spikes in $\mathrm{RE}_{\mathrm{free}}$ exposure.

Strategic implications for GEs. Together, these results show that monetization design is not a static choice but a state-dependent allocation rule that balances short-term margins against long-term engagement returns. In practice, the policy depends on the economics of ad-free investment, the generative engine's weight on future outcomes, and the surrounding market environment. Inference cost governs how expensive it is to sustain ad-free exposure, forward-looking orientation determines how much the generative engine values continuation gains, and market conditions shape the return to building retention and subscription demand (see Section 4) (Dastin & Nellis, 2023; Cai & Sophia, 2025).

5.3. Sensitivity to Market Conditions

We test the robustness of the baseline findings through a paired sensitivity analysis. For each of four market dimensions—ad sensitivity (γ), query profitability (r), inference cost (κ), and outside-option strength (ω)—we simulate a "high" and "low" variant while holding all others at baseline values, yielding eight market conditions evaluated under all four policies. Full specifications and numerical summaries are in Appendix C.2. Figure 7 shows the cumulative payoff trajectories under the eight conditions, while Appendix Table Table 1 reports the detailed results.

Two economic patterns emerge from the sensitivity analysis.

Pattern 1: The value of dynamic adaptation is robust across market conditions. Across all eight conditions, allowing the generative engine to adjust its display policy to the market environment yields cumulative payoff that is at least as high as under any fixed rule. The difference is largest when a static rule is especially misaligned with the environment: Always $\mathrm{RE}_{\mathrm{ad}}$ misses subscription-building gains, whereas Always $\mathrm{RE}_{\text{free}}$ performs worse when free-tier inference is costly or when ad insertion is especially profitable.

Pattern 2: Market conditions reshape the return to ad-free investment. Each market dimension shifts the short-term cost vs. long-term benefit trade-off. High ad sensitivity strengthens the subscription channel, high inference cost raises the cost of ad-free display, and stronger outside competition compresses overall payoffs while making retention harder. Query profitability is particularly informative: total payoff can be higher under low r than under high r, because lower ad revenue reduces the opportunity cost of ad-free exposure and makes subscription-building relatively more attractive. Across conditions, the dynamic policy adjusts its $\mathrm{RE}_{\text{free}}$ exposure accordingly, consistent with the comparative statics in Proposition 4.3: it shifts toward $\mathrm{RE}_{\text{free}}$ for ad-sensitive users and low-profit queries, and toward RE_{ad} when immediate ad revenue is high.

Detailed trajectory plots for all four metrics (cumulative payoff, active users, subscribers, and AI exposure rate) under each condition are provided in Appendix C.2.

6. Conclusion

We develop a dynamic game-theoretic framework for GEs' monetization problems in which a generative engine chooses between advertising and subscription, accounting for how current choices shape future user engagement and subscription adoption. Our analysis highlights the core trade-off: ad-free responses build AI experience, which supports retention and subscription conversion over time, while ad-heavy policies raise short-term revenue at the cost of weaker long-term value. The optimal design shifts toward ad-free responses when the generative engine is more forward-looking, when subscription conversion becomes more valuable, and when outside competition weakens engagement-dependent ad revenue. The appendix extends the analysis to a social-welfare objective and to learning the model parameters from logged data; see Sections B.6 and B.7. Future work could extend the monopoly setting to a multi-player game among competing generative engines and incorporate additional operational constraints such as capacity or latency. As GEs scale, these results clarify when ad-free responses create long-term value and when ad-heavy policies mainly trade future retention and conversion for short-term revenue.

References

Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K., and Deshpande, A. Geo: Generative engine optimization. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 5–16, 2024.

Agichtein, E., Brill, E., and Dumais, S. Improving web search ranking by incorporating user behavior information. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 19–26, 2006.

Anderson, S. P. and Coate, S. Market provision of broadcasting: A welfare analysis. *The review of Economic studies*, 72(4):947–972, 2005.

Ascarza, E. and Hardie, B. G. A joint model of usage and churn in contractual settings. *Marketing Science*, 32(4): 570–590, 2013.

Athey, S. and Ellison, G. Position auctions with consumer search. *The Quarterly Journal of Economics*, 126(3): 1213–1270, 2011.

Babu, J. and Kachwala, Z. Openai to test ads in chatgpt in bid to boost revenue. *Reuters*, January 2026.

Bajari, P., Benkard, C. L., and Levin, J. Estimating dynamic models of imperfect competition. *Econometrica*, 75(5): 1331–1370, 2007.

Balseiro, S., Bhawalkar, K., Deng, Y., Feng, Z., Mao, J., Mehta, A., Mirrokni, V., Leme, R. P., Wang, D., and Zuo, S. Position auctions in ai-generated content. *arXiv preprint arXiv:2506.03309*, 2025.

Ben-Akiva, M. E. and Lerman, S. R. *Discrete choice analysis: theory and application to travel demand*, volume 9. MIT press, 1985.

Bhawalkar, K., Psomas, A., and Wang, D. Sponsored questions and how to auction them. *arXiv preprint arXiv:2512.03975*, 2025.

Cai, K. and Sophia, D. M. Alphabet plans massive capex hike, reports cloud revenue growth slowed. *Reuters*, February 2025.

Chapekis, A. and Lieb, A. Google users are less likely to click on links when an ai summary appears in the results, 07 2025. Accessed 2026-01-14.

Chapelle, O. and Zhang, Y. A dynamic bayesian network click model for web search ranking. In *Proceedings of the 18th international conference on World wide web*, pp. 1–10, 2009.

Chmielewski, D., Seetharaman, D., and Cherney, M. A. Anthropic buys super bowl ads to slap openai for selling ads in chatgpt. *Reuters*, February 2026.

Cho, C.-H. and as, U. o. T. a. A. i. a. Why do people avoid advertising on the internet? *Journal of advertising*, 33(4): 89–97, 2004.

Cox, D. R. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2):187–202, 1972.

Coyne, P. Case study: Impact of ai search on user behavior & ctr in 2026, 12 2025.

Craswell, N., Zoeter, O., Taylor, M., and Ramsey, B. An experimental comparison of click position-bias models. In *Proceedings of the 2008 international conference on web search and data mining*, pp. 87–94, 2008.

Dastin, J. and Nellis, S. Focus: For tech giants, ai like bing and bard poses billion-dollar search problem. *Reuters*, February 2023.

Datta, H., Foubert, B., and Van Heerde, H. J. The challenge of retaining customers acquired with free trials. *Journal of Marketing Research*, 52(2):217–234, 2015.

Della Vigna, S. and Malmendier, U. Paying not to go to the gym. *american economic Review*, 96(3):694–719, 2006.

Dubé, J.-P., Hitsch, G. J., and Rossi, P. E. State dependence and alternative explanations for consumer inertia. *The RAND Journal of Economics*, 41(3):417–445, 2010.

Dubey, A., Feng, Z., Kidambi, R., Mehta, A., and Wang, D. Auctions with llm summaries. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 713–722, 2024.

Duetting, P., Mirrokni, V., Paes Leme, R., Xu, H., and Zuo, S. Mechanism design for large language models. In *Proceedings of the ACM Web Conference 2024*, pp. 144–155, 2024.

Dupret, G. E. and Piwowarski, B. A user browsing model to predict search engine click data from past observations. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 331–338, 2008.

Edwards, S. M., Li, H., and Lee, J.-H. Forced exposure and psychological reactance: Antecedents and consequences of the perceived intrusiveness of pop-up ads. *Journal of advertising*, 31(3):83–95, 2002.

Einav, L., Klopack, B., and Mahoney, N. Selling subscriptions. *American Economic Review*, 115(5):1650–1671, 2025.

Evans, D. S. The economics of the online advertising industry. *Review of network economics*, 7(3), 2008.

Fader, P. S. and Hardie, B. G. Probability models for customer-base analysis. *Journal of interactive marketing*, 23(1):61–69, 2009.

Fader, P. S., Hardie, B. G., and Shang, J. Customer-base analysis in a discrete-time noncontractual setting. *Marketing Science*, 29(6):1086–1108, 2010.

Farrell, J. and Klemperer, P. Coordination and lock-in: Competition with switching costs and network effects. *Handbook of industrial organization*, 3:1967–2072, 2007.

Ghose, A. and Yang, S. An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management science*, 55(10):1605–1622, 2009.

Ghose, A., Iperotis, P. G., and Li, B. Towards a theory model for product search. In *Proceedings of the 20th The Web Conference (WWW)*, 2011.

Ghose, A., Iperotis, P. G., and Li, B. Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content. *Marketing Science*, 31:493–520, 2012.

Gleason, J. et al. Google the gatekeeper: How search components affect clicks and attention. In *Proceedings of ICWSM*, 2023.

Google. How ai max for search campaigns works. Google Ads Help, 2025.

Gupta, S., Hanssens, D., Hardie, B., Kahn, W., Kumar, V., Lin, N., Ravishanker, N., and Sriram, S. Modeling customer lifetime value. *Journal of service research*, 9 (2):139–155, 2006.

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

Jerath, K., Ma, L., and Park, Y.-H. Consumer click behavior at a search engine: The role of keyword popularity. *Journal of Marketing Research*, 51(4):480–486, 2014.

Jeziorski, P. and Segal, I. What makes them click: Empirical analysis of consumer demand for search advertising. *American Economic Journal: Microeconomics*, 7(3):24–53, 2015.

Joachims, T., Granka, L., Pan, B., Hembrooke, H., and Gay, G. Accurately interpreting clickthrough data as implicit feedback. In *Acm Sigir Forum*, volume 51, pp. 4–11. Acm New York, NY, USA, 2017.

Kaiser, C., Kaiser, J., Schallner, R., and Schneider, S. A new era of online search? a large-scale study of user behavior and personal preferences during practical search tasks with generative ai versus traditional search engines. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, CHI EA '25, New York, NY, USA, 2025a. Association for Computing Machinery. ISBN 9798400713958. doi: 10.1145/3706599.3720123. URL https://doi.org/10.1145/3706599.3720123.

Kaiser, C., Kaiser, J., Schallner, R., and Schneider, S. A new era of online search? a large-scale study of user behavior and personal preferences during practical search tasks with generative ai versus traditional search engines. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, pp. 1–7, 2025b.

Kirsten, E., Perdekamp, J. G., Upadhyay, M., Gummadi, K. P., and Zafar, M. B. Characterizing web search in the age of generative ai. *arXiv preprint arXiv:2510.11560*, 2025.

Kwon, W., Li, Z., Zhuang, S., Sheng, Y., Zheng, L., Yu, C. H., Gonzalez, J., Zhang, H., and Stoica, I. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th symposium on operating systems principles*, pp. 611–626, 2023.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474, 2020.

Li, J., Huffman, S., and Tokuda, A. Good abandonment in mobile and pc internet search. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, pp. 43–50, 2009.

Luccioni, A. S., Viguier, S., and Ligozat, A.-L. Estimating the carbon footprint of bloom, a 176b parameter language model. *Journal of machine learning research*, 24(253): 1–15, 2023.

Luce, R. D. et al. *Individual choice behavior*, volume 4. Wiley New York, 1959.

Ma, Y., Qian, R., Wang, B., Qi, Q., Liu, W., Tang, Q., Shen, Z., Zhong, W., Shen, B., Su, Y., et al. Genauction: A generative auction for online advertising. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 12372–12380, 2025.

Malik, A. Perplexity received 780 million queries last month, ceo says, June 2025.

Mcdonald, T. Aio impact on google ctr: September 2025 update, 2025.

McFadden, D. Conditional logit analysis of qualitative choice behavior. 1972.

Miller, K. M., Sahni, N. S., and Strulov-Shlain, A. Sophisticated consumers with inertia: Long-term implications from a large-scale field experiment. *Available at SSRN 4065098*, 2023.

Mordo, T., Tennenholtz, M., and Kurland, O. Sponsored question answering. In *Proceedings of the 2024 ACM SIGIR International Conference on Theory of Information Retrieval*, pp. 167–173, 2024.

Ning, J., Coelho, J., Kong, Y., Long, Y., Martins, B., Magalhães, J., Callan, J., and Xiong, C. Agentic search in the wild: Intents and trajectory dynamics from 14m+ real search requests. *arXiv preprint arXiv:2601.17617*, 2026.

OpenAI. Introducing chatgpt search, October 2024. Updated Feb. 5, 2025.

OpenAI. Introducing chatgpt go, now available worldwide, 2026.

Puterman, M. L. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.

Rust, J. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society*, pp. 999–1033, 1987.

Scarcella, M. Google defends ai search summaries in rolling stone publisher's lawsuit. *Reuters*, January 2026.

Shi, Z., Zhang, K., and Srinivasan, K. Freemium as an optimal strategy for market dominant firms. *Marketing Science*, 38(1):150–169, 2019.

Taylor, G. Search quality and revenue cannibalization by competing search engines. *Journal of Economics & Management Strategy*, 22(3):445–467, 2013. Foundational paper on why better search might hurt ad revenue.

Varian, H. R. Position auctions. *international Journal of industrial Organization*, 25(6):1163–1178, 2007.

Venkatachary, S. C. Ai overviews in search are coming to more places around the world, October 2024.

Yao, S. and Mela, C. F. A dynamic model of sponsored search advertising. *Marketing Science*, 30(3):447–468, 2011.

Zhao, C., Hu, Q., Song, S., Chen, D., Zhu, H., Xu, J., and Zheng, B. Llm-auction: Generative auction towards llm-native advertising. *arXiv preprint arXiv:2512.10551*, 2025.

Zhou, T. and Li, S. Understanding user switch of information seeking: From search engines to generative ai. *Journal of librarianship and information science*, pp. 09610006241244800, 2024.

This appendix collects all extensions, regularity conditions, proofs, and supplementary experiments. Appendix A extends the framework setup (probabilistic adoption, post-subscription churn, and query-dependent costs). Section B states regularity conditions, collects all proofs, and develops two extensions: a social-welfare objective (Section B.6) and learning DP primitives from data (Section B.7). Section C provides additional experimental details.

A. Notes for Framework Setups

A.1. Problem Setup and Preliminaries

Remark A.1 *(Beyond i.i.d. queries).* We assume $q_t \sim \mathcal{P}(\cdot \mid x)$ i.i.d. only to keep the pre-subscription state minimal. All subsequent policy comparisons are pointwise in the realized query q. More generally, if queries follow any stationary exogenous process – for example, $q_{t+1} \sim \mathcal{P}(\cdot \mid x, q_t)$ – the DP remains valid after augmenting the state with the relevant query-process variables (e.g., the current q_t or a latent topic state), with no change to the per-query policy differential.

A.1.1. Probabilistic Adoption Hazard

This appendix relaxes the deterministic subscription threshold. After the post-query state update to (s^+, c^+) under action $a \in \{\mathrm{RE}_{\text{ad}}, \mathrm{RE}_{\text{free}}\}$, the user adopts the subscription according to a Bernoulli hazard

$$Z_{t+1} \;\sim\; \mathrm{Bernoulli}\big(\pi(x, p, s^+, c^+)\big), \tag{11}$$

where $\pi(x, p, s, c)$ is weakly increasing in s, weakly decreasing in c, and takes values in $[0, 1]$. The deterministic threshold rule is nested as the special case $\pi(x, p, s, c) = \mathbf{1}\{s \geq \tau(x, p, c)\}$.

Adoption-weighted continuation value. Define the adoption-weighted continuation value

$$\widetilde{V}_x(s, c) \;:=\; \pi(x, p, s)\, V_x(s, c, 1) \;+\; \big(1 - \pi(x, p, s)\big)\, V_x(s, c, 0). \tag{12}$$

Then the pre-subscription Bellman operator becomes

$$V_x(s, c, 0) \;=\; \mathbb{E}\left[\max_{a \in \{\mathrm{RE}_{\text{ad}}, \mathrm{RE}_{\text{free}}\}} \Big\{ r_a(x, q) + \beta\, \rho_{z_a^+(q)}(x, s_a^+(q), c_a^+(q))\, \widetilde{V}_x\big(s_a^+(q), c_a^+(q)\big) \Big\}\right], \tag{13}$$

where the expectation is taken over $q \sim \mathcal{P}(\cdot \mid x)$ (and any within-query randomness, if present).

Implication for structural results. Equation (13) shows that all subsequent policy comparisons continue to take the same "flow vs. discounted continuation" form; relative to the deterministic threshold, the only change is the replacement of $\mathbf{1}\{\cdot\}$ by $\pi(\cdot)$ inside $\widetilde{V}_x$. In particular, under the same regularity conditions used in the main text (monotonicity of $V_x(\cdot, \cdot, z)$ and $\rho_{z_a^+(q)}$ in (s, c)), the single-crossing / threshold characterizations of the optimal design policy extend directly.

A.1.2. Experience-Dependent Churn After Subscription

We relax the absorbing-subscription benchmark by allowing churn that depends on the experience state. When $Z_t = 1$, the user remains subscribed into period $t + 1$ with probability

$$\rho_1(x, s, c) \in (0, 1], \tag{14}$$

and otherwise churns back to the pre-subscription regime ($Z_{t+1} = 0$) with probability $1 - \rho_1(x, s, c)$. We assume $\rho_1(x, s, c)$ is weakly increasing in s and weakly decreasing in c.

Subscribed-state Bellman equation. Let the platform serve the subscribed experience (e.g., $\mathrm{RE}^{\text{sub}}_{\text{free}}$) when $Z_t = 1$, yielding flow payoff $r_{\text{sub}}(q, x)$ and post-query updates $(s^+_{\text{sub}}(q), c^+_{\text{sub}}(q))$. Then the subscribed-state value satisfies

$$V_x(s, c, 1) \;=\; \mathbb{E}\Big[r_{\text{sub}}(q, x) + \beta\Big(\rho_1\big(x, s^+, c^+\big)\, V_x(s^+, c^+, 1) + \big(1 - \rho_1\big(x, s^+, c^+\big)\big)\, V_x(s^+, c^+, 0)\Big)\Big], \tag{15}$$

where $(s^+, c^+) = (s^+_{\text{sub}}(q), c^+_{\text{sub}}(q))$ and the expectation is over $q \sim \mathcal{P}(\cdot \mid x)$.

Pre-subscription Bellman equation. The pre-subscription Bellman equation remains identical to the main text after replacing the continuation term $V_x(\cdot, \cdot, 1)$ by the churn-adjusted subscribed value defined in (15).

Lemma A.2 *(Monotonicity of subscribed value under churn).* Suppose $\rho_1(x, s, c)$ is weakly increasing in s and weakly decreasing in c. If $r_{\text{sub}}(q, x)$ and $(s^+_{\text{sub}}(q), c^+_{\text{sub}}(q))$ do not reverse these orders (e.g., s^+_{sub} is nondecreasing in s and c^+_{sub} is nondecreasing in c), then $V_x(s, c, 1)$ is weakly increasing in s and weakly decreasing in c.

Implication for the optimal design policy. Lemma A.2 ensures that the continuation-value channel that drives the optimal design policy (via the prospect of reaching and remaining in $Z = 1$) preserves the same comparative statics as in the no-churn benchmark. Consequently, all policy characterizations in the main text (e.g., single-crossing / threshold structures in s conditional on (q, c)) extend under the same sufficient conditions, with the only modification that the subscribed continuation value now solves (15).

A.2. Context- and Query-Dependent Inference Costs

We now allow the inference cost to depend on user context and query. Fix a measurable function $\kappa : \mathcal{X} \times \mathcal{Q} \to \mathbb{R}_+$ and replace the constant cost in the pre-subscription flow payoffs by $-\kappa(x, q)$. All other primitives (choice probabilities, state updates, retention, and conversion) are unchanged.

Assumption 1 *(Contextual inference cost enters symmetrically).* In the pre-subscription regime ($z = 0$), the per-period payoffs take the form

$$r^{\kappa}_{\text{RE}_{\text{ad}}}(x, q, s, c) = \bar{r}_{\text{RE}_{\text{ad}}}(x, q, s, c) - \kappa(x, q), \qquad r^{\kappa}_{\text{RE}_{\text{free}}}(x, q, s, c) = \bar{r}_{\text{RE}_{\text{free}}}(x, q, s, c) - \kappa(x, q),$$

for some baseline payoffs $(\bar{r}_{\text{RE}_{\text{ad}}}, \bar{r}_{\text{RE}_{\text{free}}})$ that do not depend on κ. All transition primitives (including the within-query choice model and cross-period state updates) are independent of κ.

Remark A.3 *(Direct versus continuation effects of $\kappa(x, q)$).* Under Assumption 1, the common current-period cost term $-\kappa(x, q)$ enters both action values symmetrically, so it cancels in the immediate payoff comparison. Any effect of $\kappa(x, q)$ on the optimal design policy must therefore operate through continuation values, because future pre-subscription periods also incur the cost and the two actions can lead to different future states.

Proposition A.4 *(Value monotonicity under pointwise cost increases).* Under Assumption 1 and Appendix B.1, let κ_1, κ_2 be two cost functions such that $\kappa_2(x, q) \geq \kappa_1(x, q)$ for all (x, q). Let $V^{\kappa_i}_x$ denote the optimal value function under κ_i. Then

$$V^{\kappa_2}_x(s, c, z) \leq V^{\kappa_1}_x(s, c, z) \qquad \text{for all } (s, c, z).$$

Proof. Let $\mathcal{T}^{\kappa}_x$ be the Bellman operator under cost function κ. Under Assumption 1, for any bounded V and any pre-subscription state $(s, c, 0)$, each feasible action-value is shifted down by $\kappa(x, q)$ for the realized query q. Hence the pointwise maximand is shifted down by the same amount, and after integrating over q we have

$$\mathcal{T}^{\kappa_2}_x V(s, c, 0) \leq \mathcal{T}^{\kappa_1}_x V(s, c, 0),$$

with equality for subscribed states $z = 1$ (where no free-tier inference cost is incurred). By the contraction property from Appendix B.1, $\mathcal{T}^{\kappa}_x$ has a unique fixed point V^{κ}_x, and monotonicity of the operator implies the fixed points satisfy $V^{\kappa_2}_x \leq V^{\kappa_1}_x$ pointwise. □

Remark A.5 *(Uniform cost shifts).* A common comparative-statics experiment corresponds to the one-parameter family $\kappa_{\Delta}(x, q) := \kappa(x, q) + \Delta$. Proposition A.4 implies $V^{\kappa_{\Delta}}_x$ is weakly decreasing in Δ pointwise. Whether such uniform shifts also move the optimal design policy depends on how the two actions change the expected time spent in the pre-subscription regime.

B. Notes for Theoretical Results

B.1. Standing Regularity Conditions and Assumptions

Assumption 2 *(Standing regularity conditions).* Fix a user context x and price p, and let $\beta \in (0, 1)$. Throughout, queries satisfy $q \sim \mathcal{P}(\cdot \mid x)$ on a measurable space $(\mathcal{Q}, \mathcal{F})$.

1. **Measurability.** All primitives are measurable with respect to the underlying Borel σ-algebras: $r_a(x, q)$ is measurable in q, $\delta_S(x, q, s)$ and $\delta_C(x, q, c)$ are measurable in (q, s) and (q, c) , and $\rho_x(s, c, z)$ is measurable in (s, c) for $z \in \{0, 1\}$. Consequently, for each action $a \in \{\text{RE}_{\text{ad}}, \text{RE}_{\text{free}}\}$, the post-decision map $(s, c, q) \mapsto (S^+_{a,Y}, C^+_{a,Y}, Z^+_{a,Y})$ is measurable. $\Pr(Y_t = 1 \mid x, q, a)$ is measurable.

2. **Bounded one-period payoffs.** There exists $M < \infty$ such that

$$\sup_{q\in\mathcal{Q},\, c\geq 0} \left|r_{\mathrm{RE}_{\mathrm{ad}}}(x,q)\right| \leq M, \qquad 0 \leq \kappa < \infty, \qquad |r_{\mathrm{sub}}(p)| < \infty,$$

and we set $r_{\mathrm{RE}_{\mathrm{free}}}(x,q) \equiv -\kappa$.

3. **Well-defined retention.** The pre-subscription retention probability satisfies

$$\rho_x(s,c,z) \in [0,1] \qquad \text{for all } s \geq 0,\ c \geq 0.$$

4. **Nonnegative state increments.** The experience increments are nonnegative and finite:

$$\delta_S(x,q,s) \geq 0, \qquad \delta_C(x,q,c) \geq 0 \qquad \text{for all } q \in \mathcal{Q}.$$

5. **Immediate adoption convention (if used).** The subscription status at the start of a period is deterministic: $z = 1$ if and only if $s \geq \tau(x,p,c)$, and no exit once subscribed. Hence, in the pre-subscription regime $z = 0$ we restrict attention to states with $s < \tau(x,p,c)$.

Assumption 3 *(Type sufficiency and monotone type effects).* Fix a pre-subscription state $(s,c,0)$. Let $Q_x^a(s,c,q)$ be the action value for $a \in \{\mathrm{RE}_{\mathrm{ad}}, \mathrm{RE}_{\mathrm{free}}\}$ as defined in Proposition 4.1. Define the AI advantage

$$D_x(s,c,q) := Q_x^{\mathrm{RE}_{\mathrm{free}}}(s,c,q) - Q_x^{\mathrm{RE}_{\mathrm{ad}}}(s,c,q).$$

Note that $D_x = -\Delta_x$, where Δ_x is the $\mathrm{RE}_{\mathrm{ad}}$'s edge defined in Proposition 4.1.

Assume there exists a function $\widetilde{D}$ such that for all (x,q),

$$D_x(s,c,q) = \widetilde{D}\big(s,c;\gamma(x), r_{\mathrm{RE}_{\mathrm{ad}}}(x,q), \psi(x,q)\big).$$

Moreover, for each fixed (s,c), the map $(\gamma,r,\psi) \mapsto \widetilde{D}(s,c;\gamma,r,\psi)$ is weakly increasing in γ and ψ, and weakly decreasing in r. Moreover, for each fixed (s,c,γ,r), the function $\psi \mapsto \widetilde{D}(s,c;\gamma,r,\psi)$ is left-continuous in ψ, and for each fixed (s,c,γ,ψ), the function $r \mapsto \widetilde{D}(s,c;\gamma,r,\psi)$ is right-continuous in r.

When this assumption is invoked jointly with Assumption 8, we require that the above conditions hold for the ω-indexed action values $Q_x^a(s,c,q;\omega)$ and the corresponding edge $D_x(s,c,q;\omega) = \widetilde{D}^\omega(s,c;\gamma,r,\psi)$ at each fixed $\omega \geq 0$.

Assumption 4 *(Horizon monotonicity and continuation dominance).* Fix (x,p) and consider the family of problems indexed by $\beta \in (0,1)$. Let $V_x(\cdot;\beta)$ denote the unique bounded value function solving the Bellman equation at discount factor β, and let $Q_x^a(s,c,q;\beta)$ be the associated action value for $a \in \{\mathrm{RE}_{\mathrm{ad}}, \mathrm{RE}_{\mathrm{free}}\}$.

Define the traffic-weighted continuation value under action a by

$$W_x^a(s,c,q;\beta) := \mathbb{E}_{Y\sim\Pr(\cdot|x,q,a)}\Big[\rho_{Z^+_{a,Y}}\big(x, S^+_{a,Y}, C^+_{a,Y}\big)\, V_x\big(S^+_{a,Y}, C^+_{a,Y}, Z^+_{a,Y};\beta\big)\Big],$$

and define the continuation advantage of $\mathrm{RE}_{\mathrm{free}}$ over $\mathrm{RE}_{\mathrm{ad}}$ as

$$\Phi_x(s,c,q;\beta) := W_x^{\mathrm{RE}_{\mathrm{free}}}(s,c,q;\beta) - W_x^{\mathrm{RE}_{\mathrm{ad}}}(s,c,q;\beta).$$

Assume:

1. **Uniform continuation dominance of $\mathrm{RE}_{\mathrm{free}}$.** For all (s,c,q) and all $\beta \in (0,1)$,

$$\Phi_x(s,c,q;\beta) \geq 0.$$

2. **Continuation advantage monotonicity in β.** For any $\beta' > \beta$ and all (s,c,q),

$$\Phi_x(s,c,q;\beta') \geq \Phi_x(s,c,q;\beta).$$

Assumption 5 *(Common inference cost and subscription-tier margin).* In the pre-subscription regime ($z = 0$), both display options generate an AI answer and incur the same per-query inference cost $\kappa \geq 0$. Thus the flow payoffs are

$$r_{\mathrm{RE}_{\mathrm{ad}}}(x, q) - \kappa \quad \text{under } \mathrm{RE}_{\mathrm{ad}}, \qquad -\kappa \quad \text{under } \mathrm{RE}_{\mathrm{free}},$$

where $r_{\mathrm{RE}_{\mathrm{ad}}}(x, q)$ depends only on (x, q). After subscription ($z = 1$), subscription is absorbing and the provider earns per-period net payoff

$$r_{\mathrm{sub}}(p) := p - \kappa_{\mathrm{sub}},$$

with $\kappa_{\mathrm{sub}} \geq \kappa \geq 0$. All other primitives $(\rho, \delta_S, \delta_C, \tau)$ are independent of $(\kappa, \kappa_{\mathrm{sub}})$. We break ties in the design policy in favor of $\mathrm{RE}_{\mathrm{ad}}$.

Assumption 6 *(One-step conversion region).* Fix (x, p) and a pre-subscription decision point (s, c, q) with $s < \tau(x, p, c)$. Let $u := \delta_S(x, q, s)$. Assume:

1. **One-step conversion under $\mathrm{RE}_{\mathrm{free}}$.** $s + u \geq \tau(x, p, c)$.
2. **No one-step conversion under $\mathrm{RE}_{\mathrm{ad}}$.** Choosing $\mathrm{RE}_{\mathrm{ad}}$ does not increase s in one step, so the next-period subscription status remains $z = 0$.
3. **Engagement under $\mathrm{RE}_{\mathrm{free}}$ at this decision point.** $\Pr(Y = 1 \mid x, q, \mathrm{RE}_{\mathrm{free}}) = 1$.

Assumption 7 *(Price affects subscription margin and (possibly) the conversion threshold).* Fix x and hold $(\kappa, \kappa_{\mathrm{sub}})$ fixed. The subscription margin $r_{\mathrm{sub}}(p) = p - \kappa_{\mathrm{sub}}$ is weakly increasing in p. The conversion threshold $\tau(x, p, c)$ is weakly nondecreasing in p and weakly decreasing in c. All pre-subscription primitives $(r_{\mathrm{RE}_{\mathrm{ad}}}, \rho, \delta_S, \delta_C)$ are independent of p.

Assumption 8 *(Outside-competition comparative statics).* Fix (x, p) and index outside competition by $\omega \geq 0$. The within-query outside-option utility is shifted by

$$v_0(x, q; \omega) = v_0(x, q) + \omega,$$

while $v_a(x, q)$ for $a \in \{\mathrm{RE}_{\mathrm{ad}}, \mathrm{RE}_{\mathrm{free}}\}$ and all cross-period primitives $(\rho, \delta_S, \delta_C, \tau)$ are independent of ω.

Let V_x^ω denote the optimal value function under outside-option intensity ω, and define

$$W_x^\omega(s, c, z) := \rho_x(s, c, z) V_x^\omega(s, c, z).$$

For a fixed decision point (s, c, q), define

$$p_a(\omega) := \Pr(Y = 1 \mid x, q, a; \omega), \qquad a \in \{\mathrm{RE}_{\mathrm{ad}}, \mathrm{RE}_{\mathrm{free}}\},$$

and the continuation-value advantage of $\mathrm{RE}_{\mathrm{free}}$ over $\mathrm{RE}_{\mathrm{ad}}$ by

$$\begin{aligned} \Phi_x(s, c, q; \omega) := {} & \mathbb{E}_{Y \sim \Pr(\cdot \mid x, q, \mathrm{RE}_{\mathrm{free}}; \omega)} \Big[W_x^\omega \big(S^+_{\mathrm{RE}_{\mathrm{free}}, Y}, C^+_{\mathrm{RE}_{\mathrm{free}}, Y}, Z^+_{\mathrm{RE}_{\mathrm{free}}, Y} \big) \Big] \\ & - \mathbb{E}_{Y \sim \Pr(\cdot \mid x, q, \mathrm{RE}_{\mathrm{ad}}; \omega)} \Big[W_x^\omega \big(S^+_{\mathrm{RE}_{\mathrm{ad}}, Y}, C^+_{\mathrm{RE}_{\mathrm{ad}}, Y}, Z^+_{\mathrm{RE}_{\mathrm{ad}}, Y} \big) \Big]. \end{aligned} \tag{16}$$

Assume:

1. **Nonnegative ad revenue.** $R(x, q) \geq 0$.
2. **Continuation advantage strengthens with competition.** For the fixed decision point (s, c, q), the function

$$\omega \mapsto \Phi_x(s, c, q; \omega)$$

is weakly increasing.

B.2. Extra Theoretical Results

B.3. Proofs

Proof of Proposition 4.1. Fix (x, p) and suppress these arguments when notationally convenient.

Let the state space be

$$\mathcal{S} := [0, \infty) \times [0, \infty) \times \{0, 1\}$$

equipped with the product Borel σ-algebra. Let $\mathcal{B}$ denote the space of bounded *measurable* functions $V : \mathcal{S} \to \mathbb{R}$, equipped with the sup norm $\|V\|_\infty := \sup_{(s,c,z)\in\mathcal{S}} |V(s,c,z)|$.

Step 1: Bellman operator. First let $Y \sim \Pr(\cdot \mid x, q, a)$ denote the within-query choice, and write $(S^+, C^+, Z^+) := (S^+_{a,Y}(x,q,s),\ C^+_{a,Y}(x,q,c),\ Z^+_{a,Y}(x,q,s))$. Define the Bellman operator $\mathcal{T}_x : \mathcal{B} \to \mathcal{B}$ by

$$(\mathcal{T}_x V)(s,c,0) := \int_{\mathcal{Q}} \max_{a\in\{\mathrm{RE}_{\mathrm{ad}},\mathrm{RE}_{\mathrm{free}}\}} \Big\{ r_a(x,q) + \beta\, \mathbb{E}_{Y\sim\Pr(\cdot|x,q,a)} \big[\rho\big(S^+,C^+,Z^+\big)\, V\big(S^+,C^+,Z^+\big)\big] \Big\}\, \mathcal{P}(dq \mid x), \quad (17)$$

$$(\mathcal{T}_x V)(s,c,1) := r_{\mathrm{sub}}(p) + \beta\, V(s,c,1). \quad (18)$$

Under Assumption 2, the integrand in (17) is measurable in q: $r_{\mathrm{RE}_{\mathrm{ad}}}(x,q)$ is measurable in (q,c), $\delta_S(x,q,s)$ and $\delta_C(x,q,c)$ are measurable in q, and $V(\cdot)$ is measurable in (s,c,z); hence the composition $q \mapsto V(s^+_a(q), c^+_a(q), z^+_a(q))$ is measurable. Since the max of finitely many measurable functions is measurable, the integral is well-defined. Moreover, boundedness of one-period payoffs and V implies $\mathcal{T}_x V$ is bounded.

Step 2: Contraction. Let $V, W \in \mathcal{B}$. For any (s,c) and q, define

$$G_V(a;q) := r_a(x,q) + \beta\, \mathbb{E}_{Y\sim\Pr(\cdot|x,q,a)} \big[\rho\big(S^+,C^+,Z^+\big)\, V\big(S^+,C^+,Z^+\big)\big]$$

Then for each (s,c),

$$\begin{aligned} |(\mathcal{T}_x V)(s,c,0) - (\mathcal{T}_x W)(s,c,0)| &= \left| \int \Big(\max_a G_V(a;q) - \max_a G_W(a;q) \Big)\, \mathcal{P}(dq \mid x) \right| \\ &\le \int \max_a |G_V(a;q) - G_W(a;q)|\, \mathcal{P}(dq \mid x) \\ &\le \beta\, \mathbb{E}_{Y\sim\Pr(\cdot|x,q,a)} \big[\rho(\cdot)\, \|V - W\|_\infty\big] \ \le\ \beta\, \|V - W\|_\infty, \end{aligned}$$

where we used that $\rho_{z^+_a(q)} \in [0,1]$. For $z = 1$, (18) gives

$$|(\mathcal{T}_x V)(s,c,1) - (\mathcal{T}_x W)(s,c,1)| = \beta\, |V(s,c,1) - W(s,c,1)| \le \beta\, \|V - W\|_\infty.$$

Taking the supremum over (s,c,z) yields

$$\|\mathcal{T}_x V - \mathcal{T}_x W\|_\infty \le \beta\, \|V - W\|_\infty.$$

Hence $\mathcal{T}_x$ is a contraction on $(\mathcal{B}, \|\cdot\|_\infty)$.

Step 3: Existence and uniqueness of the value function. By the Banach fixed-point theorem, $\mathcal{T}_x$ admits a unique fixed point $V_x \in \mathcal{B}$ satisfying $V_x = \mathcal{T}_x V_x$. In particular, (17) implies

$$V_x(s,c,0) = \int_{\mathcal{Q}} \max_{a\in\{\mathrm{RE}_{\mathrm{ad}},\mathrm{RE}_{\mathrm{free}}\}} Q^a_x(s,c,q)\, \mathcal{P}(dq \mid x),$$

where Q^a_x is the action-value function defined in Proposition 4.1. Equation (18) yields $V_x(s,c,1) = r_{\mathrm{sub}}(p)/(1-\beta)$.

Step 4: Optimality and the sign characterization. Consider a period in which the platform is in the pre-subscription regime with state (s,c) and observes a realized query q. Conditional on (s,c,q), the only choice is $a \in \{\mathrm{RE}_{\mathrm{ad}}, \mathrm{RE}_{\mathrm{free}}\}$, and the continuation value under any admissible policy equals the corresponding action value $Q^a_x(s,c,q)$ computed using the fixed point V_x. Therefore an optimal decision rule at (s,c,q) selects any maximizer of $Q^a_x(s,c,q)$.

Since the action set is binary, define

$$\Delta_x(s,c,q) := Q^{\mathrm{RE}_{\mathrm{ad}}}_x(s,c,q) - Q^{\mathrm{RE}_{\mathrm{free}}}_x(s,c,q).$$

Then $Q^{\mathrm{RE}_{\mathrm{ad}}}_x(s,c,q) \ge Q^{\mathrm{RE}_{\mathrm{free}}}_x(s,c,q)$ if and only if $\Delta_x(s,c,q) \ge 0$. With ties broken toward $\mathrm{RE}_{\mathrm{ad}}$, an optimal design decision is

$$a^\star(s,c,q) = \begin{cases} \mathrm{RE}_{\mathrm{ad}}, & \Delta_x(s,c,q) \ge 0, \\ \mathrm{RE}_{\mathrm{free}}, & \Delta_x(s,c,q) < 0, \end{cases}$$

which proves the claim. □

Proof of Proposition 4.3. Fix (s, c) in the pre-subscription regime. By Proposition 4.1, conditional on (x, q, s, c), $\mathrm{RE}_{\mathrm{ad}}$ is optimal if and only if $Q_x^{\mathrm{RE}_{\mathrm{ad}}}(s, c, q) \geq Q_x^{\mathrm{RE}_{\mathrm{free}}}(s, c, q)$, equivalently $D_x(s, c, q) \leq 0$.

Under Assumption 3, we can write

$$D_x(s, c, q) = \widetilde{D}(s, c; \gamma, r, \psi), \quad \text{where } (\gamma, r, \psi) = \big(\gamma(x), r_{\mathrm{RE}_{\mathrm{ad}}}(x, q), \psi(x, q)\big).$$

Define the type acceptance set

$$\mathcal{A}(s, c) := \{(\gamma, r, \psi) : \ \widetilde{D}(s, c; \gamma, r, \psi) \geq 0\}.$$

This is exactly the set of types for which $\mathrm{RE}_{\mathrm{free}}$ is optimal (since $\mathrm{RE}_{\mathrm{free}}$ is optimal iff $D_x \geq 0$).

Step 1: $\mathcal{A}(s, c)$ is an upper set under $\preceq$. Take any $t \preceq t'$, i.e., $\gamma' \geq \gamma$, $\psi' \geq \psi$, and $r' \leq r$. By Assumption 3, $\widetilde{D}(s, c; \cdot)$ is weakly increasing in γ, ψ and weakly decreasing in r, hence

$$\widetilde{D}(s, c; t') \ \geq \ \widetilde{D}(s, c; t).$$

Therefore, if $t \in \mathcal{A}(s, c)$ then $\widetilde{D}(s, c; t) \geq 0$, implying $\widetilde{D}(s, c; t') \geq 0$, so $t' \in \mathcal{A}(s, c)$. This proves the monotone ordering claim.

Step 2: threshold in ψ holding (γ, r) fixed. Fix (γ, r) and define $f(\psi) := \widetilde{D}(s, c; \gamma, r, \psi)$. By Assumption 3, f is weakly increasing. Let $S := \{\psi \in \mathbb{R} : \ f(\psi) \leq 0\}$. Then S is a (possibly empty) lower interval. Define $\psi^*(s, c; \gamma, r) := \sup S$ with the convention $\sup \emptyset = -\infty$. Then $S = (-\infty, \psi^*(s, c; \gamma, r)]$, hence

$$\mathrm{RE}_{\mathrm{ad}} \text{ is optimal } \iff \psi \leq \psi^*(s, c; \gamma, r),$$

which proves the stated threshold characterization in ψ.

Step 3: threshold in r holding (γ, ψ) fixed. Fix (γ, ψ) and define $g(r) := \widetilde{D}(s, c; \gamma, r, \psi)$. By Assumption 3, g is weakly decreasing. Let $R := \{r \in \mathbb{R} : \ g(r) \leq 0\}$. Then R is a (possibly empty) upper interval. Define $r^*(s, c; \gamma, \psi) := \inf R$ with the convention $\inf \emptyset = +\infty$. Then $R = [r^*(s, c; \gamma, \psi), \infty)$, hence

$$\mathrm{RE}_{\mathrm{ad}} \text{ is optimal } \iff r \geq r^*(s, c; \gamma, \psi),$$

which proves the stated threshold characterization in r.

□

Proof of Proposition B.1. Fix (x, p) and (s, c, q). For each β, expanding the action values yields

$$\Delta_x(s, c, q; \beta) = r_{\mathrm{RE}_{\mathrm{ad}}}(x, q) - \beta \, \Phi_x(s, c, q; \beta), \tag{19}$$

where Φ_x is defined in Assumption 4. Let $\beta' > \beta$. Subtracting (19) at β from the same identity at β' gives

$$\Delta_x(s, c, q; \beta') - \Delta_x(s, c, q; \beta) = -\Big(\beta' \Phi_x(s, c, q; \beta') - \beta \Phi_x(s, c, q; \beta)\Big). \tag{20}$$

Rewrite the bracketed term as

$$\beta' \Phi(\beta') - \beta \Phi(\beta) = (\beta' - \beta)\Phi(\beta) + \beta'\big(\Phi(\beta') - \Phi(\beta)\big),$$

where $\Phi(\beta)$ is shorthand for $\Phi_x(s, c, q; \beta)$. By Assumption 4(i), $\Phi(\beta) \geq 0$, so $(\beta' - \beta)\Phi(\beta) \geq 0$. By Assumption 4(ii), $\Phi(\beta') - \Phi(\beta) \geq 0$, so the second term is also nonnegative. Hence $\beta' \Phi(\beta') - \beta \Phi(\beta) \geq 0$, and (20) implies $\Delta_x(s, c, q; \beta') \leq \Delta_x(s, c, q; \beta)$, proving Δ is weakly decreasing in β.

For set inclusion, if $\Delta_x(s, c, q; \beta) < 0$, then $\Delta_x(s, c, q; \beta') \leq \Delta_x(s, c, q; \beta) < 0$, so (s, c, q) remains in the $\mathrm{RE}_{\mathrm{free}}$ region at β'. □

Proof for Proposition 4.5. Fix (x, p, β) and suppress these arguments for convenience.

Part (i). Take $\kappa_2 \ge \kappa_1$ and write $\Delta\kappa := \kappa_2 - \kappa_1$. Under Assumption 6, choosing $\mathrm{RE}_{\text{free}}$ at (s, c, q) yields one-step conversion, so the next-period state is subscribed. Since subscription is absorbing, for $i \in \{1, 2\}$ we have

$$Q_x^{\mathrm{RE}_{\text{free}}}(s, c, q; \kappa_i, \kappa_{\text{sub}}) = (-\kappa_i) + \beta \, \frac{r_{\text{sub}}(p)}{1-\beta},$$

hence

$$Q_x^{\mathrm{RE}_{\text{free}}}(s, c, q; \kappa_2, \kappa_{\text{sub}}) - Q_x^{\mathrm{RE}_{\text{free}}}(s, c, q; \kappa_1, \kappa_{\text{sub}}) = -\Delta\kappa.$$

Now consider choosing RE_{ad} at (s, c, q). By Assumption 6, the next-period subscription status remains $z = 0$. Let $V_{x,i}$ denote the optimal value function under $(\kappa_i, \kappa_{\text{sub}})$. By Proposition A.4, specialized to constant cost functions κ_1 and κ_2 while holding κ_{sub} fixed, the value from any pre-subscription state is weakly lower under κ_2:

$$V_{x,2}(s', c', 0) \le V_{x,1}(s', c', 0) \qquad \text{for all } (s', c').$$

Therefore

$$\begin{aligned}
&Q_x^{\mathrm{RE}_{\text{ad}}}(s, c, q; \kappa_2, \kappa_{\text{sub}}) - Q_x^{\mathrm{RE}_{\text{ad}}}(s, c, q; \kappa_1, \kappa_{\text{sub}}) \\
&= \big(r_{\mathrm{RE}_{\text{ad}}}(x, q) - \kappa_2\big) - \big(r_{\mathrm{RE}_{\text{ad}}}(x, q) - \kappa_1\big) \\
&\quad + \beta \, \mathbb{E}_{Y \sim \Pr(\cdot | x, q, \mathrm{RE}_{\text{ad}})} \Big[\rho_x\big(S^+_{\mathrm{RE}_{\text{ad}}, Y}, C^+_{\mathrm{RE}_{\text{ad}}, Y}, 0\big) \\
&\qquad\qquad \times \Big(V_{x,2}\big(S^+_{\mathrm{RE}_{\text{ad}}, Y}, C^+_{\mathrm{RE}_{\text{ad}}, Y}, 0\big) - V_{x,1}\big(S^+_{\mathrm{RE}_{\text{ad}}, Y}, C^+_{\mathrm{RE}_{\text{ad}}, Y}, 0\big) \Big) \Big] \\
&\le -\Delta\kappa,
\end{aligned}$$

because $\rho_x(\cdot) \in [0, 1]$ and the bracketed value difference is nonpositive. Consequently,

$$\begin{aligned}
\Delta_x(s, c, q; \kappa_2, \kappa_{\text{sub}}) - \Delta_x(s, c, q; \kappa_1, \kappa_{\text{sub}}) &= \big(Q_x^{\mathrm{RE}_{\text{ad}}}(\kappa_2) - Q_x^{\mathrm{RE}_{\text{ad}}}(\kappa_1)\big) \\
&\quad - \big(Q_x^{\mathrm{RE}_{\text{free}}}(\kappa_2) - Q_x^{\mathrm{RE}_{\text{free}}}(\kappa_1)\big) \\
&\le 0.
\end{aligned}$$

Thus $\Delta_x(s, c, q; \kappa, \kappa_{\text{sub}})$ is weakly decreasing in κ, so a higher pre-subscription cost pushes the decision toward $\mathrm{RE}_{\text{free}}$ in the near-threshold region.

Part (ii). Consider $\kappa_{\text{sub},2} \ge \kappa_{\text{sub},1}$ and write $r_{\text{sub},i} := p - \kappa_{\text{sub},i}$ for $i \in \{1, 2\}$, so $r_{\text{sub},2} = r_{\text{sub},1} - (\kappa_{\text{sub},2} - \kappa_{\text{sub},1})$.

Under the one-step conversion condition $s < \tau(x, p, c) \le s + u$, choosing $\mathrm{RE}_{\text{free}}$ moves the experience state to at least $\tau(x, p, c)$ in one step, so the next-period state is subscribed. Since subscription is absorbing, the subscribed-state value satisfies $V_x(s', c', 1) = r_{\text{sub}}(p)/(1-\beta)$ for any (s', c'). Therefore,

$$Q_x^{\mathrm{RE}_{\text{free}}}(s, c, q; \kappa, \kappa_{\text{sub},i}) = (-\kappa) + \beta \, \frac{r_{\text{sub},i}}{1-\beta},$$

and hence

$$Q_x^{\mathrm{RE}_{\text{free}}}(s, c, q; \kappa, \kappa_{\text{sub},2}) - Q_x^{\mathrm{RE}_{\text{free}}}(s, c, q; \kappa, \kappa_{\text{sub},1}) = -\frac{\beta}{1-\beta}\big(\kappa_{\text{sub},2} - \kappa_{\text{sub},1}\big).$$

Now consider choosing RE_{ad} at (s, c, q) with $s < \tau(x, p, c)$ and RE_{ad} not changing s in one step. Then the next-period status remains pre-subscription ($z = 0$), so subscription payoffs that depend on κ_{sub} cannot begin before period $t + 2$. Even in the most extreme case where the user becomes subscribed from $t + 2$ onward with probability one, the total discounted exposure to the subscribed payoff stream from period $t + 2$ is at most $\sum_{k=2}^{\infty} \beta^k = \beta^2/(1-\beta)$. Because changing κ_{sub} shifts the subscribed per-period payoff by exactly $-(\kappa_{\text{sub},2} - \kappa_{\text{sub},1})$, this implies the action value under RE_{ad} satisfies the bound

$$Q_x^{\mathrm{RE}_{\text{ad}}}(s, c, q; \kappa, \kappa_{\text{sub},2}) - Q_x^{\mathrm{RE}_{\text{ad}}}(s, c, q; \kappa, \kappa_{\text{sub},1}) \;\ge\; -\frac{\beta^2}{1-\beta}\big(\kappa_{\text{sub},2} - \kappa_{\text{sub},1}\big).$$

Combining the two differences yields

$$\Delta_x(s,c,q;\kappa,\kappa_{\text{sub},2}) - \Delta_x(s,c,q;\kappa,\kappa_{\text{sub},1}) = \big(Q_x^{\text{RE}_{\text{ad}}}(\kappa_{\text{sub},2}) - Q_x^{\text{RE}_{\text{ad}}}(\kappa_{\text{sub},1})\big) - \big(Q_x^{\text{RE}_{\text{free}}}(\kappa_{\text{sub},2}) - Q_x^{\text{RE}_{\text{free}}}(\kappa_{\text{sub},1})\big)$$
$$\geq -\frac{\beta^2}{1-\beta}\Delta\kappa_{\text{sub}} + \frac{\beta}{1-\beta}\Delta\kappa_{\text{sub}} = \beta\,\Delta\kappa_{\text{sub}} \;\geq\; 0,$$

where $\Delta\kappa_{\text{sub}} := \kappa_{\text{sub},2} - \kappa_{\text{sub},1}$. This proves Δ_x is weakly increasing in κ_{sub} under the stated one-step condition, and the set inclusion follows immediately. □

Proof for Proposition 4.6. Take any $p_2 > p_1$ in I and write $\Delta p := p_2 - p_1 > 0$.

Under Assumption 6, choosing RE_{free} at (x,s,c,q) yields one-step conversion for both prices. Under Assumption 5, the current-period flow payoff under RE_{free} is $-\kappa$, and since subscription is absorbing,

$$Q_x^{\text{RE}_{\text{free}}}(s,c,q;\,p_i) = -\kappa + \beta\,\frac{p_i - \kappa_{\text{sub}}}{1-\beta}, \qquad i \in \{1,2\}.$$

Hence

$$Q_x^{\text{RE}_{\text{free}}}(s,c,q;\,p_2) - Q_x^{\text{RE}_{\text{free}}}(s,c,q;\,p_1) = \frac{\beta}{1-\beta}\Delta p.$$

Now consider choosing RE_{ad} at (x,s,c,q). By Assumption 6(ii), the next-period subscription status remains $z = 0$ for both prices. By Assumption 7, all pre-subscription primitives are independent of p, so the only effect of a higher price comes through the subscribed payoff stream. Moreover, $\tau(x,p,c)$ is weakly nondecreasing in p, so raising p cannot make subscription occur earlier; it can only weakly delay or reduce future conversion.

Therefore, even in the most favorable case for RE_{ad} under p_2, the additional gain from raising the price from p_1 to p_2 is bounded above by the discounted value of receiving the higher subscribed per-period margin Δp from period $t+2$ onward with probability one:

$$Q_x^{\text{RE}_{\text{ad}}}(s,c,q;\,p_2) - Q_x^{\text{RE}_{\text{ad}}}(s,c,q;\,p_1) \leq \sum_{k=2}^{\infty} \beta^k\,\Delta p = \frac{\beta^2}{1-\beta}\Delta p.$$

Combining the two action-value differences yields

$$\Delta_x(s,c,q;\,p_2) - \Delta_x(s,c,q;\,p_1) = \big(Q_x^{\text{RE}_{\text{ad}}}(p_2) - Q_x^{\text{RE}_{\text{ad}}}(p_1)\big)$$
$$- \big(Q_x^{\text{RE}_{\text{free}}}(p_2) - Q_x^{\text{RE}_{\text{free}}}(p_1)\big)$$
$$\leq \frac{\beta^2}{1-\beta}\Delta p - \frac{\beta}{1-\beta}\Delta p = -\beta\,\Delta p \leq 0.$$

Hence $\Delta_x(s,c,q;\,p)$ is weakly decreasing in p on I, proving that a higher subscription price weakly shifts the optimal design toward RE_{free} in this near-threshold region. □

Proof for Proposition 4.7. Fix (x,p,s,c,q) and suppress fixed arguments in notation. Under Assumption 8, for each action $a \in \{\text{RE}_{\text{ad}}, \text{RE}_{\text{free}}\}$,

$$p_a(\omega) = \frac{\exp(v_a)}{\exp(v_a) + \exp(v_0 + \omega)}, \qquad \frac{d}{d\omega}p_a(\omega) = -p_a(\omega)\big(1 - p_a(\omega)\big).$$

Hence $p_a(\omega)$ is weakly decreasing in ω for both actions, and in particular $p_{\text{RE}_{\text{ad}}}(\omega)$ is weakly decreasing in ω.

For each ω, define

$$C_{\text{RE}_{\text{ad}}}(\omega) := \mathbb{E}_{Y \sim \Pr(\cdot|x,q,\text{RE}_{\text{ad}};\omega)}\Big[W_x^{\omega}\big(S^+_{\text{RE}_{\text{ad}},Y}, C^+_{\text{RE}_{\text{ad}},Y}, Z^+_{\text{RE}_{\text{ad}},Y}\big)\Big],$$
$$C_{\text{RE}_{\text{free}}}(\omega) := \mathbb{E}_{Y \sim \Pr(\cdot|x,q,\text{RE}_{\text{free}};\omega)}\Big[W_x^{\omega}\big(S^+_{\text{RE}_{\text{free}},Y}, C^+_{\text{RE}_{\text{free}},Y}, Z^+_{\text{RE}_{\text{free}},Y}\big)\Big].$$

Then the action values at the fixed decision point are

$$Q_x^{\text{RE}_{\text{ad}}}(s,c,q;\omega) = -\kappa + R(x,q)p_{\text{RE}_{\text{ad}}}(\omega) + \beta C_{\text{RE}_{\text{ad}}}(\omega),$$

$$Q_x^{\mathrm{RE_{free}}}(s,c,q;\omega) = -\kappa + \beta C_{\mathrm{RE_{free}}}(\omega).$$

Therefore, using $\Delta_x = Q_x^{\mathrm{RE_{ad}}} - Q_x^{\mathrm{RE_{free}}}$ and the definition $\Phi_x(s,c,q;\omega) = C_{\mathrm{RE_{free}}}(\omega) - C_{\mathrm{RE_{ad}}}(\omega)$ from Equation (16),

$$\Delta_x(s,c,q;\omega) = R(x,q)p_{\mathrm{RE_{ad}}}(\omega) \; - \; \beta\Phi_x(s,c,q;\omega).$$

The first term is weakly decreasing in ω because $R(x,q) \geq 0$ and $p_{\mathrm{RE_{ad}}}(\omega)$ is weakly decreasing. The second term is also weakly decreasing because $\Phi_x(s,c,q;\omega)$ is weakly increasing by Assumption 8(ii). Hence $\Delta_x(s,c,q;\omega)$ is weakly decreasing in ω.

Therefore, if $\Delta_x(s,c,q;\omega) < 0$ at some ω, then for any $\omega' > \omega$, $\Delta_x(s,c,q;\omega') \leq \Delta_x(s,c,q;\omega) < 0$. This gives the stated set inclusion for the $\mathrm{RE_{free}}$-optimal region. □

Proof of Proposition 4.8. Fix (s,c) in the pre-subscription regime. By Proposition 4.7, for every (x,q) and $\omega' > \omega$, $\Delta_x(s,c,q;\omega') \leq \Delta_x(s,c,q;\omega)$. Under Assumption 3 (assumed to hold for each fixed ω), write $\Delta_x(s,c,q;\omega) = \widetilde{\Delta}^{\omega}(s,c;\gamma,r,\psi)$ where $\widetilde{\Delta}^{\omega}$ is weakly decreasing in γ and ψ, and weakly increasing in r.

Part (i). Fix (γ,r) and define $f^{\omega}(\psi) := \widetilde{\Delta}^{\omega}(s,c;\gamma,r,\psi)$. Since f^{ω} is weakly decreasing in ψ and $f^{\omega'}(\psi) \leq f^{\omega}(\psi)$ pointwise, the threshold $\psi^*(\omega) := \sup\{\psi : f^{\omega}(\psi) \geq 0\}$ satisfies $\psi^*(\omega') \leq \psi^*(\omega)$.

Part (ii). Fix (γ,ψ) and define $g^{\omega}(r) := \widetilde{\Delta}^{\omega}(s,c;\gamma,r,\psi)$. Since g^{ω} is weakly increasing in r and $g^{\omega'}(r) \leq g^{\omega}(r)$ pointwise, the threshold $r^*(\omega) := \inf\{r : g^{\omega}(r) \geq 0\}$ satisfies $r^*(\omega') \geq r^*(\omega)$.

Part (iii). The $\mathrm{RE_{ad}}$-optimal region at ω is $\mathcal{A}(\omega) := \{(\gamma,r,\psi) : \widetilde{\Delta}^{\omega} \geq 0\}$. Since $\widetilde{\Delta}^{\omega'} \leq \widetilde{\Delta}^{\omega}$ pointwise, $\mathcal{A}(\omega') \subseteq \mathcal{A}(\omega)$. For any measure μ, $\mu(\mathcal{A}(\omega')) \leq \mu(\mathcal{A}(\omega))$. □

B.4. Implications of the GE Provider's Weight on Future Outcomes

How much the generative engine values future outcomes relative to immediate value directly shapes the design policy.

In our framework as introduced in Section 3, the generative engine's weight on future outcomes is summarized by the discount factor $\beta \in (0,1)$. A higher β corresponds to a more forward-looking generative engine that places greater weight on future retention and subscription conversion relative to near-term value (Bajari et al., 2007). Accordingly, we treat the previously defined value functions (and the associated other functions) as depending on β, and study how the optimal design region changes as β varies.

Indexing the weight. We index the weight of value function V_x and action values Q_x by β, and write the edge function as $\Delta_x(s,c,q;\beta)$ (the same object as in Section 3.4, now viewed as a function of β). Because one-period payoffs (first term of (10)) do not depend on β, the weight on future outcomes effect affects entirely through the long-term value (second term of (10)).

Directional weight effects. To make the weight effect unambiguous, we impose two sufficient conditions (formalized in Appendix B.1). First, a larger weight on future outcomes increases the value of future outcomes from any state. Second, at any decision point, showing $\mathrm{RE_{free}}$ leads to weakly better future outcomes than showing $\mathrm{RE_{ad}}$. Together, these conditions imply that a more forward-looking generative engine selects $\mathrm{RE_{free}}$ more often.

Proposition B.1 *(More forward-looking generative engines show* RE_{free} *more often).* Under Assumptions 2 and 4, for any fixed (s,c,q), the $\mathrm{RE_{ad}}$'s edge of value over $\mathrm{RE_{free}}$ $\Delta_x(s,c,q;\beta)$ is weakly decreasing in β. Consequently, for any $\beta' > \beta$,

$$\begin{aligned}\big\{(s,c,q) : \Delta_x(s,c,q;\beta) < 0\big\} \subseteq \\ \big\{(s,c,q) : \Delta_x(s,c,q;\beta') < 0\big\}.\end{aligned} \tag{21}$$

Equivalently, the decision set of $\mathrm{RE_{free}}$ weakly expands with larger β.

Economic interpretation. A larger weight on future outcomes (i.e., a larger β) increases the generative engine's weight on users' future engagement and conversion to subscription relative to immediate ad revenue. Since $\mathrm{RE_{free}}$ improves continuation value V_x through retention ρ_x and the user's experience accumulation, a more forward-looking generative engine displays $\mathrm{RE_{free}}$ to a larger population of users (and queries).

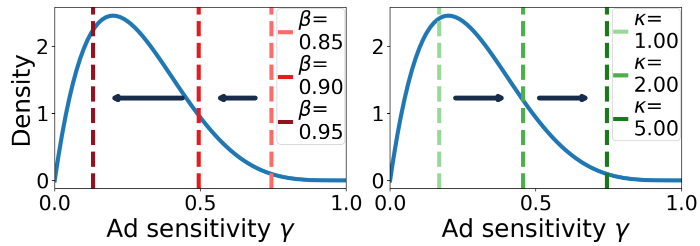

Figure 8: User-type policy cutoff shift under the generative engine's weight on future outcomes (indexed by β) and inference cost κ. Fix user context x and query type, and draw *ad sensitivity* $\gamma \sim \mathrm{Beta}$. Vertical dashed lines mark the optimal policy cutoff between $\mathrm{RE}_{\text{free}}$ and RE_{ad} for each parameter value. *Left:* varying discount factor β. *Right:* varying AI inference cost κ.

Figure 8 (left) illustrates Proposition B.1 at a fixed state and query environment. As β increases, the policy cutoff in *ad sensitivity* shifts left: $\mathrm{RE}_{\text{free}}$ becomes optimal for a larger share of users (i.e., the region to the right of the vertical cutoff expands.)

Section B.5 discusses finite-horizon formulations and how changes in the time horizon affect the policy shift.

B.5. Finite-Horizon Case

Proposition B.2 *(Finite-horizon value converges geometrically).* Fix (x, p) and $\beta \in (0, 1)$. Maintain Assumption 2. Let $\mathcal{T}_x$ be the Bellman operator at discount factor β (as in Proof 1), and define the finite-horizon value-iteration sequence by $V_x^{(0)} \equiv 0$ and

$$V_x^{(T)} := \mathcal{T}_x V_x^{(T-1)} \qquad \text{for } T \geq 1.$$

Let V_x^∞ denote the unique bounded fixed point of $\mathcal{T}_x$. Assume the one-period reward is uniformly bounded: there exists $R_{\max} < \infty$ such that for all (s, c, z) and all q,

$$\max_{a \in \{\mathrm{RE}_{\text{ad}}, \mathrm{RE}_{\text{free}}\}} \big| r_a(x, q) \big| \leq R_{\max}, \qquad |r_{\text{sub}}(p)| \leq R_{\max}.$$

Then $V_x^{(T)}$ converges uniformly to V_x^∞ and

$$\big\| V_x^\infty - V_x^{(T)} \big\|_\infty \;\leq\; \frac{R_{\max}}{1-\beta}\, \beta^T. \tag{22}$$

Proof. Under Assumption 2, $\mathcal{T}_x$ maps bounded measurable functions to bounded measurable functions and is a β-contraction under the sup norm:

$$\|\mathcal{T}_x V - \mathcal{T}_x W\|_\infty \leq \beta \, \|V - W\|_\infty \qquad \text{for all bounded } V, W.$$

Hence $\mathcal{T}_x$ admits a unique bounded fixed point V_x^∞ and, for any initial $V_x^{(0)}$, value iteration converges:

$$\|V_x^\infty - V_x^{(T)}\|_\infty \leq \beta^T \|V_x^\infty - V_x^{(0)}\|_\infty.$$

With $V_x^{(0)} \equiv 0$, it remains to bound $\|V_x^\infty\|_\infty$. For any bounded V and any state, by the reward bound and $\rho_z \in [0, 1]$,

$$|(\mathcal{T}_x V)(s, c, 0)| \leq \mathbb{E}\Big[\max_a |r_a(x, q)|\Big] + \beta \, \|V\|_\infty \leq R_{\max} + \beta \, \|V\|_\infty,$$

and similarly $|(\mathcal{T}_x V)(s, c, 1)| \leq R_{\max} + \beta \, \|V\|_\infty$. Applying this inequality to the fixed point $V_x^\infty = \mathcal{T}_x V_x^\infty$ yields $\|V_x^\infty\|_\infty \leq R_{\max}/(1-\beta)$. Therefore

$$\|V_x^\infty - V_x^{(T)}\|_\infty \leq \beta^T \|V_x^\infty\|_\infty \leq \frac{R_{\max}}{1-\beta}\, \beta^T,$$

which proves (22). □

Proposition B.3 *(Policy stabilizes away from an indifference band).* Fix (x, p) and $\beta \in (0, 1)$. Maintain Assumption 2 and the reward bound of Proposition B.2. For $a \in \{\mathrm{RE}_{\text{ad}}, \mathrm{RE}_{\text{free}}\}$, define the infinite-horizon action value

$$Q_x^{a,\infty}(s, c, q) := r_a(x, q) + \beta\, \rho_{z_a^+(q)}\big(x, s_a^+(q), c_a^+(q)\big)\, V_x^\infty\big(s_a^+(q), c_a^+(q), z_a^+(q)\big),$$

and the T-step (finite-horizon) action value

$$Q_x^{a,(T)}(s, c, q) := r_a(x, q) + \beta\, \rho_{z_a^+(q)}\big(x, s_a^+(q), c_a^+(q)\big)\, V_x^{(T-1)}\big(s_a^+(q), c_a^+(q), z_a^+(q)\big), \qquad T \geq 1.$$

Let the corresponding advantages be

$$D_x^\infty(s, c, q) := Q_x^{\mathrm{RE}_{\text{free}},\infty}(s, c, q) - Q_x^{\mathrm{RE}_{\text{ad}},\infty}(s, c, q), \qquad D_x^{(T)}(s, c, q) := Q_x^{\mathrm{RE}_{\text{free}},(T)}(s, c, q) - Q_x^{\mathrm{RE}_{\text{ad}},(T)}(s, c, q).$$

Then

$$\sup_{s,c,q} \big|D_x^\infty(s, c, q) - D_x^{(T)}(s, c, q)\big| \;\leq\; 2\beta\, \big\|V_x^\infty - V_x^{(T-1)}\big\|_\infty \;\leq\; \frac{2\beta R_{\max}}{1-\beta}\, \beta^{T-1}. \tag{23}$$

Consequently, if at some (s, c, q),

$$\big|D_x^\infty(s, c, q)\big| > \frac{2\beta R_{\max}}{1-\beta}\, \beta^{T-1}, \tag{24}$$

then $\mathrm{sign}\big(D_x^{(T)}(s, c, q)\big) = \mathrm{sign}\big(D_x^\infty(s, c, q)\big)$, and thus the finite-horizon and infinite-horizon optimal display decisions coincide at (s, c, q) (under the same tie-breaking rule).

Proof. Fix (s, c, q). For any $a \in \{\mathrm{RE}_{\text{ad}}, \mathrm{RE}_{\text{free}}\}$, the only difference between $Q_x^{a,\infty}$ and $Q_x^{a,(T)}$ is the continuation value, hence

$$\begin{aligned}\big|Q_x^{a,\infty}(s, c, q) - Q_x^{a,(T)}(s, c, q)\big| &= \beta\, \rho_{z_a^+(q)}\big(x, s_a^+(q), c_a^+(q)\big)\, \big|V_x^\infty(\cdot) - V_x^{(T-1)}(\cdot)\big| \\ &\leq \beta\, \big\|V_x^\infty - V_x^{(T-1)}\big\|_\infty,\end{aligned}$$

since $\rho_{z_a^+(q)} \in [0, 1]$. Therefore,

$$\begin{aligned}\big|D_x^\infty(s, c, q) - D_x^{(T)}(s, c, q)\big| &= \big|\big(Q_x^{\mathrm{RE}_{\text{free}},\infty} - Q_x^{\mathrm{RE}_{\text{ad}},\infty}\big) - \big(Q_x^{\mathrm{RE}_{\text{free}},(T)} - Q_x^{\mathrm{RE}_{\text{ad}},(T)}\big)\big| \\ &\leq \big|Q_x^{\mathrm{RE}_{\text{free}},\infty} - Q_x^{\mathrm{RE}_{\text{free}},(T)}\big| + \big|Q_x^{\mathrm{RE}_{\text{ad}},\infty} - Q_x^{\mathrm{RE}_{\text{ad}},(T)}\big| \\ &\leq 2\beta\, \big\|V_x^\infty - V_x^{(T-1)}\big\|_\infty,\end{aligned}$$

which proves the first inequality in (23). The second inequality in (23) follows by combining with Proposition B.2. Finally, if (24) holds, then

$$|D_x^{(T)} - D_x^\infty| < |D_x^\infty|,$$

which implies $D_x^{(T)}$ has the same sign as D_x^∞. Under the same tie-breaking rule, the induced optimal display decision therefore coincides at (s, c, q). □

B.6. A Social-Welfare Extension

We extend the framework to study when the generative engine's revenue-maximizing design policy agrees with a welfare-maximizing design policy. The user's engage-versus-outside choice, the state dynamics, the retention rule, and subscription conversion all remain exactly as in the main text. The only change is the objective: welfare counts both the generative engine's payoff and the user's benefit from the same per-query choice problem.

User benefit from the same per-query choice. For each displayed response $a \in \{\mathrm{RE}_{\text{ad}}, \mathrm{RE}_{\text{free}}\}$, let

$$u_{\text{user}}(x, q, a)$$

denote the user's per-query benefit from facing the same binary choice between engaging with the displayed response and taking the outside option. This is an extra primitive for the welfare extension only; it does not change the user's behavior, the state dynamics, or the revenue objective.

Welfare flow and value function. For pre-subscription users, define the per-query welfare flow under action a by

$$w_a(x,q) := r_a(x,q) + u_{\text{user}}(x,q,a). \tag{25}$$

After subscription, let

$$w_{\text{sub}}(x,p)$$

denote the per-period welfare flow in the paid tier, combining the generative engine's subscription payoff and the user's value from the paid experience.

Fix x. Let $W_x(s,c,z)$ denote the maximal expected discounted welfare starting from state (s,c,z). In the subscribed state,

$$W_x(s,c,1) = w_{\text{sub}}(x,p) + \beta W_x(s,c,1) = \frac{w_{\text{sub}}(x,p)}{1-\beta}. \tag{26}$$

In the pre-subscription state,

$$\begin{aligned} W_x(s,c,0) = \mathbb{E}_{q\sim\mathcal{P}(\cdot|x)}\Bigg[\max_{a\in\{\text{RE}_{\text{ad}},\text{RE}_{\text{free}}\}} \Big\{ w_a(x,q) \\ + \beta\, \mathbb{E}_{Y\sim\Pr(\cdot|x,q,a)}\Big[\rho_x\big(S^+_{a,Y}, C^+_{a,Y}, Z^+_{a,Y}\big)\, W_x\big(S^+_{a,Y}, C^+_{a,Y}, Z^+_{a,Y}\big)\Big]\Big\}\Bigg]. \end{aligned} \tag{27}$$

Relative to the revenue Bellman equation, the only change is that the current-period flow $r_a(x,q)$ is replaced by $w_a(x,q)$ and the subscribed-state flow is replaced by $w_{\text{sub}}(x,p)$.

Assumption 9 *(Bounded welfare flows).* There exist constants $U_{\max}, \Delta^{\max}_{\text{sub}} < \infty$ such that

$$|u_{\text{user}}(x,q,a)| \le U_{\max} \qquad \text{for all } (x,q,a),$$

and

$$|w_{\text{sub}}(x,p) - r_{\text{sub}}(p)| \le \Delta^{\max}_{\text{sub}} \qquad \text{for all } x.$$

Under Assumptions 2 and 9, the same contraction argument as in Proof 1 yields a unique bounded welfare value function W_x.

Welfare action values and the welfare edge. For $a \in \{\text{RE}_{\text{ad}}, \text{RE}_{\text{free}}\}$, define

$$Q^{a,\text{sw}}_x(s,c,q) := w_a(x,q) + \beta\, \mathbb{E}_{Y\sim\Pr(\cdot|x,q,a)}\Big[\rho_x\big(S^+_{a,Y}, C^+_{a,Y}, Z^+_{a,Y}\big)\, W_x\big(S^+_{a,Y}, C^+_{a,Y}, Z^+_{a,Y}\big)\Big]. \tag{28}$$

Define the welfare edge of RE_{ad} over RE_{free} by

$$\Delta^{\text{sw}}_x(s,c,q) := Q^{\text{RE}_{\text{ad}},\text{sw}}_x(s,c,q) - Q^{\text{RE}_{\text{free}},\text{sw}}_x(s,c,q). \tag{29}$$

As in the main text, we break ties in favor of RE_{ad}.

Proposition B.4 *(The welfare design policy follows the same sign rule).* Under Assumptions 2 and 9, the welfare-maximizing design policy chooses RE_{ad} at (s,c,q) if and only if

$$\Delta^{\text{sw}}_x(s,c,q) \ge 0.$$

Equivalently, it chooses RE_{free} if and only if

$$\Delta^{\text{sw}}_x(s,c,q) < 0.$$

Proof. By definition, the welfare-maximizing design policy chooses an action that maximizes the welfare action value at (s,c,q). Since the action set is binary and ties are broken toward RE_{ad}, the policy chooses RE_{ad} if and only if

$$Q^{\text{RE}_{\text{ad}},\text{sw}}_x(s,c,q) \ge Q^{\text{RE}_{\text{free}},\text{sw}}_x(s,c,q),$$

which is equivalent to $\Delta^{\text{sw}}_x(s,c,q) \ge 0$. □

This result shows that the welfare objective keeps the same point-wise comparison as the revenue objective. Welfare can change the preferred action only by changing the size of the $\mathrm{RE}_{\mathrm{ad}}$ edge, not by changing the basic decision rule.

Proposition B.5 *(Welfare and revenue agree away from small revenue margins).* Under Assumptions 2 and 9, let

$$\epsilon_{\mathrm{sw}} := \max\{U_{\max}, \Delta_{\mathrm{sub}}^{\max}\}.$$

Then

$$\|W_x - V_x\|_\infty \le \frac{\epsilon_{\mathrm{sw}}}{1-\beta}. \tag{30}$$

Moreover,

$$\sup_{s,c,q} \left|\Delta_x^{\mathrm{sw}}(s,c,q) - \Delta_x(s,c,q)\right| \le 2U_{\max} + \frac{2\beta}{1-\beta}\,\epsilon_{\mathrm{sw}}. \tag{31}$$

In particular, if at some (s, c, q),

$$\left|\Delta_x(s,c,q)\right| > 2U_{\max} + \frac{2\beta}{1-\beta}\,\epsilon_{\mathrm{sw}}, \tag{32}$$

then the welfare-maximizing and revenue-maximizing design policies choose the same action at (s, c, q).

Proof. Let $\mathcal{T}_x$ denote the revenue Bellman operator from Proof 1, and let $\mathcal{T}_x^{\mathrm{sw}}$ denote the welfare Bellman operator defined by (27).

Fix any bounded function F. For pre-subscription states,

$$\begin{aligned}
&\left|(\mathcal{T}_x^{\mathrm{sw}}F)(s,c,0) - (\mathcal{T}_xF)(s,c,0)\right| \\
&\le \mathbb{E}_{q\sim\mathcal{P}(\cdot|x)}\Big[\max_{a\in\{\mathrm{RE}_{\mathrm{ad}},\mathrm{RE}_{\mathrm{free}}\}} \left|u_{\mathrm{user}}(x,q,a)\right|\Big] \\
&\le U_{\max},
\end{aligned}$$

because the two operators differ only through the current-period user-benefit term. For subscribed states,

$$\left|(\mathcal{T}_x^{\mathrm{sw}}F)(s,c,1) - (\mathcal{T}_xF)(s,c,1)\right| = \left|w_{\mathrm{sub}}(x,p) - r_{\mathrm{sub}}(p)\right| \le \Delta_{\mathrm{sub}}^{\max}.$$

Hence

$$\|\mathcal{T}_x^{\mathrm{sw}}F - \mathcal{T}_xF\|_\infty \le \epsilon_{\mathrm{sw}}.$$

Using the fixed-point identities $W_x = \mathcal{T}_x^{\mathrm{sw}}W_x$ and $V_x = \mathcal{T}_xV_x$, together with the β-contraction property of $\mathcal{T}_x^{\mathrm{sw}}$, we obtain

$$\begin{aligned}
\|W_x - V_x\|_\infty &= \|\mathcal{T}_x^{\mathrm{sw}}W_x - \mathcal{T}_xV_x\|_\infty \\
&\le \|\mathcal{T}_x^{\mathrm{sw}}W_x - \mathcal{T}_x^{\mathrm{sw}}V_x\|_\infty + \|\mathcal{T}_x^{\mathrm{sw}}V_x - \mathcal{T}_xV_x\|_\infty \\
&\le \beta\|W_x - V_x\|_\infty + \epsilon_{\mathrm{sw}},
\end{aligned}$$

which yields (30).

Next fix (s, c, q) and $a \in \{\mathrm{RE}_{\mathrm{ad}}, \mathrm{RE}_{\mathrm{free}}\}$. Since the welfare and revenue problems share the same transition law,

$$\begin{aligned}
&Q_x^{a,\mathrm{sw}}(s,c,q) - Q_x^a(s,c,q) \\
&= u_{\mathrm{user}}(x,q,a) + \beta\,\mathbb{E}_{Y\sim\Pr(\cdot|x,q,a)}\Big[\rho_x\big(S_{a,Y}^+, C_{a,Y}^+, Z_{a,Y}^+\big) \\
&\qquad \times \big(W_x\big(S_{a,Y}^+, C_{a,Y}^+, Z_{a,Y}^+\big) - V_x\big(S_{a,Y}^+, C_{a,Y}^+, Z_{a,Y}^+\big)\big)\Big].
\end{aligned}$$

Therefore,

$$\left|Q_x^{a,\mathrm{sw}}(s,c,q) - Q_x^a(s,c,q)\right| \le U_{\max} + \beta\|W_x - V_x\|_\infty.$$

Applying this bound to both actions gives

$$
\begin{aligned}
&\left|\Delta_x^{\mathrm{sw}}(s,c,q) - \Delta_x(s,c,q)\right| \\
&\le \left|Q_x^{\mathrm{RE_{ad}},\mathrm{sw}}(s,c,q) - Q_x^{\mathrm{RE_{ad}}}(s,c,q)\right| + \left|Q_x^{\mathrm{RE_{free}},\mathrm{sw}}(s,c,q) - Q_x^{\mathrm{RE_{free}}}(s,c,q)\right| \\
&\le 2U_{\max} + 2\beta\left\|W_x - V_x\right\|_\infty \\
&\le 2U_{\max} + \frac{2\beta}{1-\beta}\,\epsilon_{\mathrm{sw}},
\end{aligned}
$$

which proves (31).

Finally, if (32) holds, then (31) implies that $\Delta_x^{\mathrm{sw}}(s,c,q)$ has the same sign as $\Delta_x(s,c,q)$. Because both policies break ties toward $\mathrm{RE_{ad}}$, they choose the same action at (s,c,q). □

This result localizes where welfare can matter. If the revenue edge already strongly favors one action, adding user-side value does not overturn the decision. Disagreement can arise only in states where the generative engine is already close to indifferent between $\mathrm{RE_{ad}}$ and $\mathrm{RE_{free}}$.

B.7. Learning DP Primitives and Robustness of the Plug-in Design Policy

This subsection shows how the dynamic program can be implemented with logged data while staying fully aligned with the main framework. We keep the same state variables, the same engage-versus-outside choice, and the same design policy problem. The only goal here is to clarify which primitives must be learned and how estimation error affects the learned design policy.

Logged data and overlap. Fix a user context x and focus on pre-subscription decisions. Suppose we observe logged tuples

$$
(x_t,\ q_t,\ s_t,\ c_t,\ a_t,\ Y_t,\ \Pi_t,\ S_{t+1},\ C_{t+1},\ Z_{t+1},\ R_{t+1}),
$$

where $a_t \in \{\mathrm{RE_{ad}}, \mathrm{RE_{free}}\}$ is the displayed response, $Y_t \in \{0,1\}$ is the user's engage-versus-outside decision, Π_t is the realized one-period payoff to the generative engine, and $R_{t+1} \in \{0,1\}$ indicates whether the user returns in period $t+1$. The post-update state $(S_{t+1}, C_{t+1}, Z_{t+1})$ can be observed directly or reconstructed from $(q_t, s_t, c_t, a_t, Y_t)$ using the model's state-update rules.

Assumption 10 *(Exploration/overlap).* There exists $\eta \in (0, 1/2]$ such that for all (x,q,s,c) in the support,

$$
\begin{aligned}
&\mathbb{P}(a_t = \mathrm{RE_{ad}} \mid x_t = x, q_t = q, s_t = s, c_t = c) \ge \eta, \\
&\mathbb{P}(a_t = \mathrm{RE_{free}} \mid x_t = x, q_t = q, s_t = s, c_t = c) \ge \eta.
\end{aligned} \tag{33}
$$

Primitives to estimate. The Bellman equation uses three one-step objects:

$$
m_a(x,q) := \mathbb{P}(Y_t = 1 \mid x_t = x,\ q_t = q,\ a_t = a), \tag{34}
$$

$$
r_a(x,q) := \mathbb{E}[\Pi_t \mid x_t = x,\ q_t = q,\ a_t = a]\,, \tag{35}
$$

$$
\rho_x(s,c,z) := \mathbb{P}(R_{t+1} = 1 \mid x_t = x,\ S_{t+1} = s,\ C_{t+1} = c,\ Z_{t+1} = z)\,. \tag{36}
$$

Here $m_a(x,q)$ is the engage probability from the main text, $r_a(x,q)$ is the generative engine's expected one-period payoff, and $\rho_x(s,c,z)$ is the retention rule evaluated at the post-update state. We treat the query distribution $\mathcal{P}(\cdot \mid x)$, the state-update mappings in (6), and the subscription rule as known model primitives. If those objects are also estimated, the same arguments below continue to hold after adding their error terms.

Let $\widehat{m}_a$, $\widehat{r}_a$, and $\widehat{\rho}_x$ denote fitted estimates of (34)–(36). Define the fitted engage probabilities

$$
\widehat{\pi}_a(1 \mid x,q) := \widehat{m}_a(x,q), \qquad \widehat{\pi}_a(0 \mid x,q) := 1 - \widehat{m}_a(x,q),
$$

and analogously $\pi_a(1 \mid x,q) := m_a(x,q)$ and $\pi_a(0 \mid x,q) := 1 - m_a(x,q)$.

Plug-in Bellman equation and learned design policy. For any bounded function F, define the estimated Bellman operator by

$$\begin{aligned}(\widehat{\mathcal{T}}_x F)(s,c,1) &:= r_{\mathrm{sub}}(p) + \beta F(s,c,1),\\ (\widehat{\mathcal{T}}_x F)(s,c,0) &:= \mathbb{E}_{q\sim\mathcal{P}(\cdot|x)}\Bigg[\max_{a\in\{\mathrm{RE}_{\mathrm{ad}},\mathrm{RE}_{\mathrm{free}}\}}\Big\{\widehat{r}_a(x,q)\\ &\qquad + \beta\sum_{y\in\{0,1\}}\widehat{\pi}_a(y\mid x,q)\,\widehat{\rho}_x\big(S^+_{a,y},C^+_{a,y},Z^+_{a,y}\big)\,F\big(S^+_{a,y},C^+_{a,y},Z^+_{a,y}\big)\Big\}\Bigg]. \qquad (37)\end{aligned}$$

Under admissible fitted primitives, $\widehat{\mathcal{T}}_x$ is a β-contraction and therefore has a unique bounded fixed point, denoted by $\widehat{V}_x$.

Define the plug-in action values by

$$\begin{aligned}\widehat{Q}^a_x(s,c,q) &:= \widehat{r}_a(x,q)\\ &\quad + \beta\sum_{y\in\{0,1\}}\widehat{\pi}_a(y\mid x,q)\,\widehat{\rho}_x\big(S^+_{a,y},C^+_{a,y},Z^+_{a,y}\big)\,\widehat{V}_x\big(S^+_{a,y},C^+_{a,y},Z^+_{a,y}\big), \qquad (38)\end{aligned}$$

and the plug-in $\mathrm{RE}_{\mathrm{ad}}$ edge by

$$\widehat{\Delta}_x(s,c,q) := \widehat{Q}^{\mathrm{RE}_{\mathrm{ad}}}_x(s,c,q) - \widehat{Q}^{\mathrm{RE}_{\mathrm{free}}}_x(s,c,q). \qquad (39)$$

The learned design policy uses the same tie-breaking rule as the main text:

$$\widehat{g}_x(s,c,q) = \mathrm{RE}_{\mathrm{ad}} \iff \widehat{\Delta}_x(s,c,q) \ge 0. \qquad (40)$$

Proposition B.6 *(Plug-in estimation error yields small value and policy error).* Fix x. Suppose the fitted primitives satisfy

$$0 \le \widehat{m}_a(x,q) \le 1, \qquad 0 \le \widehat{\rho}_x(s,c,z) \le 1,$$

and that there exists $R_{\max} < \infty$ such that

$$\sup_{a,q}\big|r_a(x,q)\big| \le R_{\max}, \qquad \sup_{a,q}\big|\widehat{r}_a(x,q)\big| \le R_{\max}, \qquad \big|r_{\mathrm{sub}}(p)\big| \le R_{\max}.$$

Define

$$\begin{aligned}\epsilon_m &:= \sup_{a,q}\big|\widehat{m}_a(x,q) - m_a(x,q)\big|,\\ \epsilon_r &:= \sup_{a,q}\big|\widehat{r}_a(x,q) - r_a(x,q)\big|,\\ \epsilon_\rho &:= \sup_{s,c,z}\big|\widehat{\rho}_x(s,c,z) - \rho_x(s,c,z)\big|, \qquad (41)\end{aligned}$$

and let

$$V_{\max} := \frac{R_{\max}}{1-\beta}, \qquad B_x := \epsilon_r + \beta\big(2\epsilon_m + \epsilon_\rho\big)V_{\max}.$$

Then the following hold:

1. $$\big\|\widehat{V}_x - V_x\big\|_\infty \le \frac{B_x}{1-\beta}.$$
2. $$\sup_{s,c,q}\big|\widehat{\Delta}_x(s,c,q) - \Delta_x(s,c,q)\big| \le \frac{2B_x}{1-\beta}.$$
3. If at some (s,c,q),
   $$\big|\Delta_x(s,c,q)\big| > \frac{2B_x}{1-\beta},$$
   then the learned and optimal design policies agree at (s,c,q):
   $$\widehat{g}_x(s,c,q) = g^*_x(s,c,q).$$

4. Let $V_x^{\widehat{g}}$ denote the value induced by following the learned design policy $\widehat{g}_x$ thereafter. Then for all pre-subscription states $(s, c, 0)$,

$$0 \le V_x(s,c,0) - V_x^{\widehat{g}}(s,c,0) \le \frac{2B_x}{(1-\beta)^2}.$$

Proof. Let $\mathcal{T}_x$ denote the true Bellman operator. Because both the true and fitted retention probabilities lie in $[0,1]$, both $\mathcal{T}_x$ and $\widehat{\mathcal{T}}_x$ are β-contractions on bounded functions.

First, by the bounded payoff assumption,

$$\|V_x\|_\infty \le V_{\max}.$$

Now fix any bounded F with $\|F\|_\infty \le V_{\max}$. For a pre-subscription state $(s, c, 0)$, the difference between the fitted and true one-step maximands is bounded by the largest action-level discrepancy. For any $a \in \{\mathrm{RE}_{\text{ad}}, \mathrm{RE}_{\text{free}}\}$ and query q,

$$\begin{aligned}
&\Big|\widehat{r}_a(x,q) + \beta \sum_y \widehat{\pi}_a(y \mid x,q)\, \widehat{\rho}_x\big(S^+_{a,y}, C^+_{a,y}, Z^+_{a,y}\big)\, F\big(S^+_{a,y}, C^+_{a,y}, Z^+_{a,y}\big) \\
&\qquad - r_a(x,q) - \beta \sum_y \pi_a(y \mid x,q)\, \rho_x\big(S^+_{a,y}, C^+_{a,y}, Z^+_{a,y}\big)\, F\big(S^+_{a,y}, C^+_{a,y}, Z^+_{a,y}\big)\Big| \\
&\le \epsilon_r + \beta \sum_y \widehat{\pi}_a(y \mid x,q)\, \Big|\widehat{\rho}_x\big(S^+_{a,y}, C^+_{a,y}, Z^+_{a,y}\big) - \rho_x\big(S^+_{a,y}, C^+_{a,y}, Z^+_{a,y}\big)\Big| \cdot \|F\|_\infty \\
&\qquad + \beta \Big| \sum_y \big(\widehat{\pi}_a(y \mid x,q) - \pi_a(y \mid x,q)\big)\, \rho_x\big(S^+_{a,y}, C^+_{a,y}, Z^+_{a,y}\big)\, F\big(S^+_{a,y}, C^+_{a,y}, Z^+_{a,y}\big)\Big| \\
&\le \epsilon_r + \beta \epsilon_\rho V_{\max} + \beta \big|\widehat{m}_a(x,q) - m_a(x,q)\big| \cdot \big|g_{a,1}(q) - g_{a,0}(q)\big|,
\end{aligned}$$

where

$$g_{a,y}(q) := \rho_x\big(S^+_{a,y}, C^+_{a,y}, Z^+_{a,y}\big)\, F\big(S^+_{a,y}, C^+_{a,y}, Z^+_{a,y}\big).$$

Since $|g_{a,y}(q)| \le \|F\|_\infty \le V_{\max}$, we have

$$\big|g_{a,1}(q) - g_{a,0}(q)\big| \le 2V_{\max},$$

so the last display is bounded by

$$\epsilon_r + \beta\big(2\epsilon_m + \epsilon_\rho\big)V_{\max} = B_x.$$

Taking the maximum over a and the expectation over q yields

$$\big\|\widehat{\mathcal{T}}_x F - \mathcal{T}_x F\big\|_\infty \le B_x.$$

Using the fixed-point identities $\widehat{V}_x = \widehat{\mathcal{T}}_x \widehat{V}_x$ and $V_x = \mathcal{T}_x V_x$, together with the contraction property of $\widehat{\mathcal{T}}_x$, we obtain

$$\begin{aligned}
\big\|\widehat{V}_x - V_x\big\|_\infty &= \big\|\widehat{\mathcal{T}}_x \widehat{V}_x - \mathcal{T}_x V_x\big\|_\infty \\
&\le \big\|\widehat{\mathcal{T}}_x \widehat{V}_x - \widehat{\mathcal{T}}_x V_x\big\|_\infty + \big\|\widehat{\mathcal{T}}_x V_x - \mathcal{T}_x V_x\big\|_\infty \\
&\le \beta \big\|\widehat{V}_x - V_x\big\|_\infty + B_x,
\end{aligned}$$

which proves (1).

Next fix (s, c, q) and $a \in \{\mathrm{RE}_{\text{ad}}, \mathrm{RE}_{\text{free}}\}$. By adding and subtracting terms exactly as above,

$$\begin{aligned}
\big|\widehat{Q}_x^a(s,c,q) - Q_x^a(s,c,q)\big| &\le \epsilon_r + \beta \big\|\widehat{V}_x - V_x\big\|_\infty + \beta\big(2\epsilon_m + \epsilon_\rho\big)V_{\max} \\
&= B_x + \beta \big\|\widehat{V}_x - V_x\big\|_\infty \\
&\le \frac{B_x}{1-\beta}.
\end{aligned}$$

Applying this bound to both actions gives

$$\begin{aligned}\left|\widehat{\Delta}_x(s,c,q)-\Delta_x(s,c,q)\right| &\le \left|\widehat{Q}_x^{\mathrm{RE}_{\mathrm{ad}}}(s,c,q)-Q_x^{\mathrm{RE}_{\mathrm{ad}}}(s,c,q)\right|+\left|\widehat{Q}_x^{\mathrm{RE}_{\mathrm{free}}}(s,c,q)-Q_x^{\mathrm{RE}_{\mathrm{free}}}(s,c,q)\right| \\ &\le \frac{2B_x}{1-\beta},\end{aligned}$$

which proves (2).

If (3) holds, then (2) implies that $\widehat{\Delta}_x(s,c,q)$ has the same sign as $\Delta_x(s,c,q)$. Because the learned and optimal design policies use the same tie-breaking rule in favor of $\mathrm{RE}_{\mathrm{ad}}$, they choose the same action, proving (3).

Finally, let

$$\epsilon_\Delta := \sup_{s,c,q}\left|\widehat{\Delta}_x(s,c,q)-\Delta_x(s,c,q)\right|.$$

At any query-state point where the learned and optimal policies disagree, the same-sign argument above implies

$$\left|\Delta_x(s,c,q)\right| \le \epsilon_\Delta.$$

Because the action set is binary, the loss from choosing $\widehat{g}_x$ instead of g_x^* at that point is at most ϵ_Δ in true action value. Therefore, for every pre-subscription state,

$$V_x(s,c,0) \le (\mathcal{T}_x^{\widehat{g}}V_x)(s,c,0)+\epsilon_\Delta,$$

where $\mathcal{T}_x^{\widehat{g}}$ is the Bellman operator induced by the learned design policy. Since $V_x^{\widehat{g}}=\mathcal{T}_x^{\widehat{g}}V_x^{\widehat{g}}$ and $\mathcal{T}_x^{\widehat{g}}$ is also a β-contraction,

$$\left\|V_x-V_x^{\widehat{g}}\right\|_\infty \le \epsilon_\Delta+\beta\left\|V_x-V_x^{\widehat{g}}\right\|_\infty,$$

which yields

$$\left\|V_x-V_x^{\widehat{g}}\right\|_\infty \le \frac{\epsilon_\Delta}{1-\beta} \le \frac{2B_x}{(1-\beta)^2}.$$

This proves (4). □

This proposition shows that learning matters most near indifference. If the true $\mathrm{RE}_{\mathrm{ad}}$ edge is already large in magnitude, moderate estimation error does not change the design choice. Estimation error can overturn the policy only in states where the generative engine is already close to indifferent between $\mathrm{RE}_{\mathrm{ad}}$ and $\mathrm{RE}_{\mathrm{free}}$.

Corollary B.7 *(Consistency of the learned design policy).* Fix x and consider a sequence of fitted primitives

$$\left(\widehat{m}_{a,n},\widehat{r}_{a,n},\widehat{\rho}_{x,n}\right)_{n\ge 1}$$

satisfying the admissibility conditions in Proposition B.6. If

$$\epsilon_{m,n}\to 0,\qquad \epsilon_{r,n}\to 0,\qquad \epsilon_{\rho,n}\to 0,$$

then the associated plug-in value functions, action edges, and learned design policies satisfy

$$\left\|\widehat{V}_{x,n}-V_x\right\|_\infty \to 0, \tag{42}$$

$$\sup_{s,c,q}\left|\widehat{\Delta}_{x,n}(s,c,q)-\Delta_x(s,c,q)\right| \to 0, \tag{43}$$

$$\sup_{s,c}\left(V_x(s,c,0)-V_x^{\widehat{g}_n}(s,c,0)\right) \to 0. \tag{44}$$

Moreover, for any fixed (s,c,q) such that

$$\Delta_x(s,c,q)\neq 0,$$

the learned and optimal design policies eventually agree at that point:

$$\widehat{g}_{x,n}(s,c,q)=g_x^*(s,c,q) \qquad \text{for all sufficiently large } n.$$

Proof. Let

$$B_{x,n} := \epsilon_{r,n} + \beta\big(2\epsilon_{m,n} + \epsilon_{\rho,n}\big)V_{\max}.$$

By assumption, $B_{x,n} \to 0$. Therefore Proposition B.6 implies

$$\big\|\widehat{V}_{x,n} - V_x\big\|_\infty \le \frac{B_{x,n}}{1-\beta} \to 0,$$

which proves (42), and

$$\sup_{s,c,q}\big|\widehat{\Delta}_{x,n}(s,c,q) - \Delta_x(s,c,q)\big| \le \frac{2B_{x,n}}{1-\beta} \to 0,$$

which proves (43). The value bound (4) from Proposition B.6 then yields (44).

Finally, fix (s,c,q) with $\Delta_x(s,c,q) \neq 0$. Since (43) gives

$$\widehat{\Delta}_{x,n}(s,c,q) \to \Delta_x(s,c,q),$$

the sign of $\widehat{\Delta}_{x,n}(s,c,q)$ must eventually match the sign of $\Delta_x(s,c,q)$. Because both policies break ties toward $\mathrm{RE}_{\mathrm{ad}}$, they eventually choose the same action, proving (B.7). □

This corollary gives the clean asymptotic implication. As the estimated primitives become uniformly accurate, the plug-in value function converges to the true value function, the learned $\mathrm{RE}_{\mathrm{ad}}$ edge converges to the true edge, and the learned design policy becomes value-optimal. Exact policy agreement is immediate away from knife-edge states where $\Delta_x(s,c,q) = 0$.

B.8. Extra Simulation Plots

B.8.1. Optimal Policy Characterization

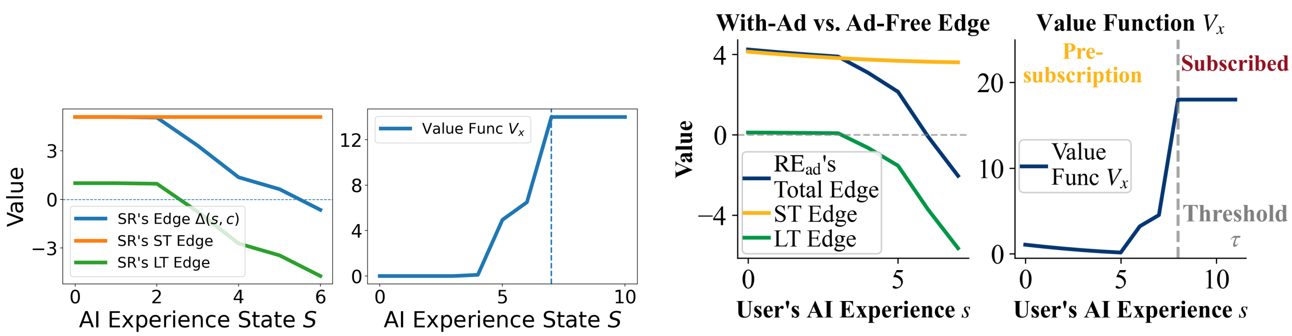

(a) $\gamma = 0.2$ (b) $\gamma = 0.8$

Figure 9: Simulation illustration of the threshold rule in Proposition 4.1 for a fixed context x and $\mathrm{RE}_{\mathrm{ad}}$ *experience state* c, under different user heterogeneity. *Left: Ad sensitivity* $\gamma = 0.2$; *Right: AI sensitivity* $\gamma = 0.8$. Inside each subplot: *Left:* Query-averaged $\mathrm{RE}_{\mathrm{ad}}$ edge $\mathbb{E}_q\Delta_x(s,c,q)$ and its short-term (ST) and continuation-value (LT) components from Remark 4.2 as *AI experience* s varies. *Right:* pre-subscription value $V_x(s,c,0)$ as a function of s; the dashed line marks the subscription threshold $\tau(x,p,c)$, after which the value equals the subscribed level.

B.8.2. Threshold Structure Under User-Query Types

C. Notes for Experiment

C.1. Additional Experimental Setups

Population, horizon, and shared randomness. We simulate one platform interacting with $N = 500$ heterogeneous users over a horizon $T = 20$ with discount factor $\beta = 0.95$ (unless otherwise noted). Each user i draws a fixed latent type γ_i i.i.d. from Beta. In each period t, the environment draws a query profitability signal r_{it} and an AI-quality signal ψ_{it} i.i.d. from $\mathrm{Unif}[0,1]^2$. To make policy comparisons apples-to-apples, we evaluate all policies on the *same* sampled population and the same realizations of $\{(r_{it}, \psi_{it})\}$, as well as the same uniform random draws used to realize consumption and retention outcomes.

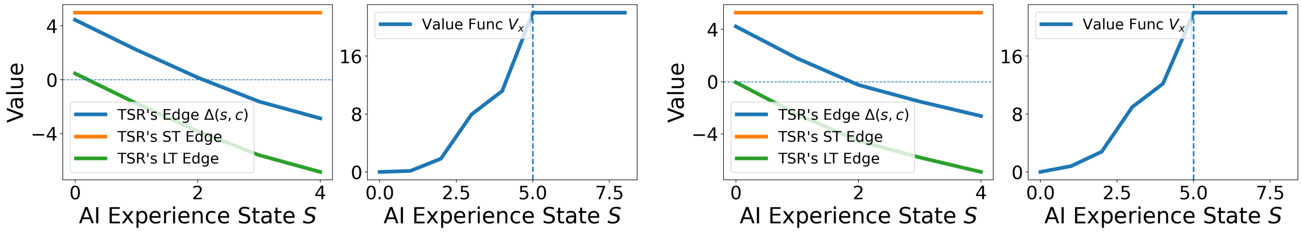

(a) $\theta = 0.2$ (b) $\theta = 0.8$

Figure 10: Simulation illustration of the threshold rule in Proposition 4.1 for a fixed context x and $\mathrm{RE_{ad}}$ *experience state* c, under different user heterogeneity. *Left: Ad reliance* $\theta = 0.2$; *Right: AI reliance* $\theta = 0.8$. Inside each subplot: *Left:* Query-averaged $\mathrm{RE_{ad}}$ edge $\mathbb{E}_q \Delta_x(s, c, q)$ and its short-term (ST) and continuation-value (LT) components from Remark 4.2 as *AI experience* s varies. *Right:* pre-subscription value $V_x(s, c, 0)$ as a function of s; the dashed line marks the subscription threshold $\tau(x, p, c)$, after which the value equals the subscribed level.

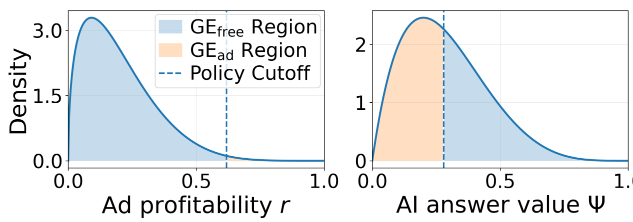

Figure 11: Query-type illustration with fixing user context x and drawing query types r, ψ from beta distributions. Shaded regions indicate which search outcome is optimal: $\mathrm{RE_{free}}$ vs. $\mathrm{RE_{ad}}$.

User utilities and consumption. Given the platform's displayed outcome $a_{it} \in \{\mathrm{RE_{ad}}, \mathrm{RE_{free}}\}$, the user chooses whether to consume the displayed content versus an outside option (outside option utility normalized to 0). Consumption follows a logit rule with utility indices

$$v_{\mathrm{RE_{free}}}(\gamma; r, \psi) = w_\psi\, \psi + w_{r,\mathrm{RE_{free}}}\, r, \qquad v_{\mathrm{RE_{ad}}}(\gamma; r, c) = u_0^{\mathrm{RE_{ad}}} + u_r^{\mathrm{RE_{ad}}}\, r - w_\gamma\, \gamma - u_c^{\mathrm{RE_{ad}}}\, c,$$

where c is the $\mathrm{RE_{ad}}$ fatigue state defined below. (All coefficients are fixed environment parameters; we use the values in our code/config unless explicitly overridden.)

State dynamics and platform payoffs. Each pre-subscription user maintains two integer states (s_{it}, c_{it}) initialized at $(0, 0)$ and truncated to caps $(S_{\max}, C_{\max})$. If the user consumes $\mathrm{RE_{free}}$, the AI experience state increases by a discrete increment $s_{i,t+1} \leftarrow \min\{s_{it} + \Delta_s(\psi_{it}), S_{\max}\}$ with $\Delta_s(\psi) = 1 + \mathbf{1}\{\psi \geq \psi_{\mathrm{cut}}\}$. If the user consumes $\mathrm{RE_{ad}}$, the fatigue state increases by one: $c_{i,t+1} \leftarrow \min\{c_{it} + 1, C_{\max}\}$. If the user does not consume, the state does not change. Platform per-period payoff is earned only upon consumption: consuming $\mathrm{RE_{free}}$ yields a serving cost $-\kappa_{\mathrm{free}}$, while consuming $\mathrm{RE_{ad}}$ yields ad revenue $r_{\mathrm{RE_{ad}}}(r, c, s) = \max\{b_0 + b_r\, r - b_c\, c - b_s\, s, 0\}$ (allowing SR monetization to deteriorate with fatigue and, optionally, with accumulated AI experience).

Retention, conversion, and subscription value. Pre-subscription retention is state-dependent with a floor:

$$\rho(s, c) = \rho_{\min} + (1 - \rho_{\min}) \cdot \sigma\big(\alpha_s\, s - \alpha_c\, c\big),$$

where $\sigma(\cdot)$ is the logistic link. Conversion is triggered deterministically when accumulated experience reaches a type-dependent cutoff: a user converts once $s_{it} \geq \lceil \tau \rceil$, with $\tau(\theta, p, c) = \tau_0 + \tau_p \cdot p - \tau_\theta \cdot \theta - \tau_c \cdot c$. Upon conversion, the user enters an absorbing subscribed regime; we add the discounted continuation value starting next period, $\beta^{t+1} V_{\mathrm{sub}}$ with $V_{\mathrm{sub}} = (p - \kappa_{\mathrm{paid}})/(1 - \beta)$, and stop simulating further pre-subscription dynamics for that user. In our "total active users" metric, we count both retained pre-subscription users and already-converted (subscribed) users as active.

DP computation and type binning. The Optimal DP policy is computed by value iteration on the integer grid (s, c) in the pre-subscription regime. Bellman expectations over query signals are approximated by Monte Carlo with n_q i.i.d. draws per state using a fixed random seed. To reduce the number of DP solves, we bin γ onto a uniform grid over $[0, 1]^2$ with n_{bins}

points per dimension and solve the DP only for the unique binned type pairs appearing in the simulated population; each user is mapped to the nearest grid point for policy lookup.

C.2. Market-Condition Sensitivity Analysis

Market conditions. We conduct a paired high/low sensitivity analysis across four market dimensions. For each dimension, we simulate both a "high" and "low" variant while holding all other parameters at baseline values:

1. **Ad sensitivity** (γ): High $\gamma \sim \mathrm{Beta}(5, 2)$ concentrates mass on ad-sensitive users; Low $\gamma \sim \mathrm{Beta}(2, 5)$ concentrates mass on ad-tolerant users. Baseline: $\gamma \sim \mathrm{Unif}[0, 1]$.
2. **Query profitability** (r): High $r \sim \mathrm{Beta}(5, 2)$ shifts queries toward high ad revenue; Low $r \sim \mathrm{Beta}(2, 5)$ reduces ad profitability. Baseline: $r \sim \mathrm{Unif}[0, 1]$.
3. **Inference cost** (κ_{free}): High $\kappa = 1.3$ makes $\mathrm{RE}_{\text{free}}$ expensive to serve; Low $\kappa = 0.5$ makes $\mathrm{RE}_{\text{free}}$ cheap. Baseline: $\kappa = 0.9$.
4. **Outside competition** (ω): High $\omega = 0.5$ raises the outside option; Low $\omega = 0.2$ weakens competition. Baseline: $\omega = 0$.

This yields $2 \times 4 = 8$ market conditions, each simulated under all four policies ($N = 300$ users, $T = 20$ periods, $\beta = 0.95$, $n_{\text{bins}} = 5$, $n_q = 60$). To ensure statistical reliability, each configuration is run across five independent random seeds (seeds $42, 1042, 2042, 3042, 4042$), for a total of $8 \times 4 \times 5 = 160$ simulation runs. All results report mean $\pm$ standard deviation across seeds.

Detailed results. Table 1 reports the final cumulative payoff (mean±std), $\mathrm{RE}_{\text{free}}$ exposure rate, and subscriber count under each condition.

Table 1: Full sensitivity analysis results (mean±std over 5 seeds, T=20). *DP* $\mathrm{RE}_{\textit{free}}$ *%* reports the average share of active non-subscribers shown $\mathrm{RE}_{\text{free}}$ by the Optimal DP policy. *Subs* reports the total subscribers under Optimal DP and Always $\mathrm{RE}_{\text{free}}$ policies.

Condition	Opt. DP	Greedy	Always RE_{ad}	Always $\mathrm{RE}_{\text{free}}$	$\mathrm{RE}_{\text{free}}$ %	Subs(DP)	Subs($\mathrm{RE}_{\text{free}}$)
High γ	**9.19**±0.35	6.66±0.44	1.40±0.02	8.42±0.50	92%	260±4	276±4
Low γ	**3.65**±0.18	1.85±0.05	1.41±0.04	2.27±0.17	87%	217±6	252±6
High r	**3.85**±0.22	3.11±0.17	1.83±0.02	2.34±0.27	87%	130±8	156±8
Low r	**7.22**±0.15	5.03±0.14	0.82±0.02	7.04±0.14	92%	285±2	292±2
High κ	**5.13**±0.20	3.07±0.12	1.36±0.02	3.62±0.28	64%	196±10	265±6
Low κ	**8.16**±0.28	6.03±0.27	1.36±0.02	7.30±0.29	93%	243±6	265±6
High ω	**5.20**±0.28	3.77±0.25	1.34±0.01	4.16±0.30	84%	196±7	226±6
Low ω	**6.03**±0.22	4.10±0.26	1.35±0.02	4.97±0.26	86%	223±6	251±5

Detailed trajectory plots. Figures 12 to 15 show the full trajectory plots for all four metrics—cumulative payoff, active users, subscribers, and AI exposure rate—under each market condition.

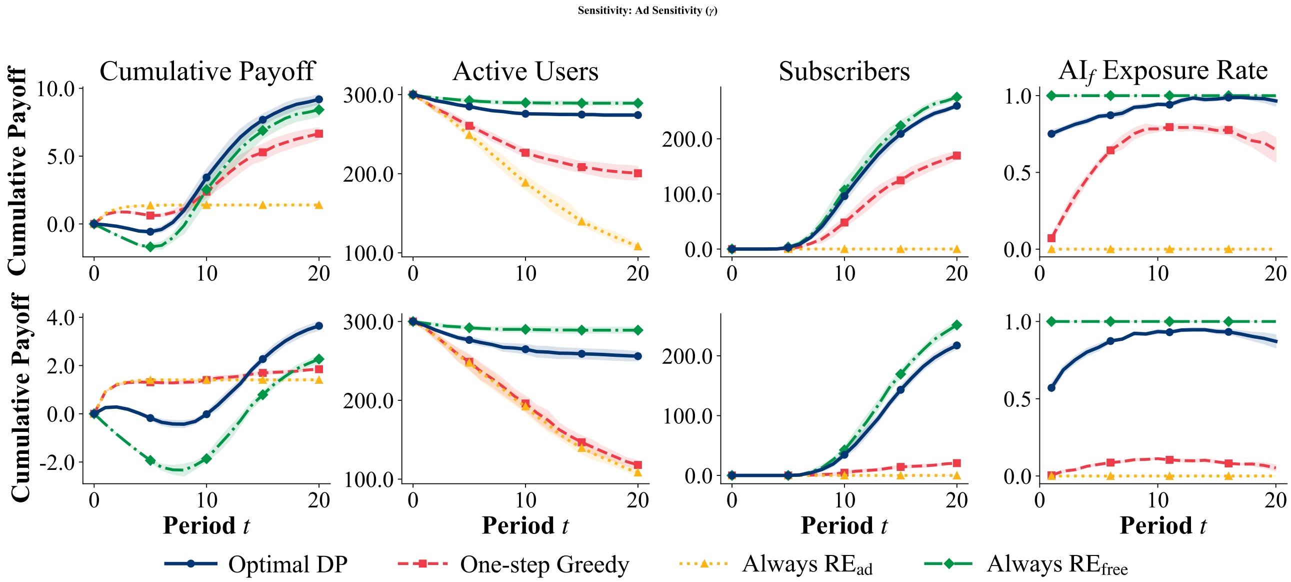

Figure 12: Sensitivity to ad sensitivity (γ): all four metrics under high-γ (top) and low-γ (bottom) conditions. High ad sensitivity accelerates subscription conversion and increases the DP's payoff advantage.

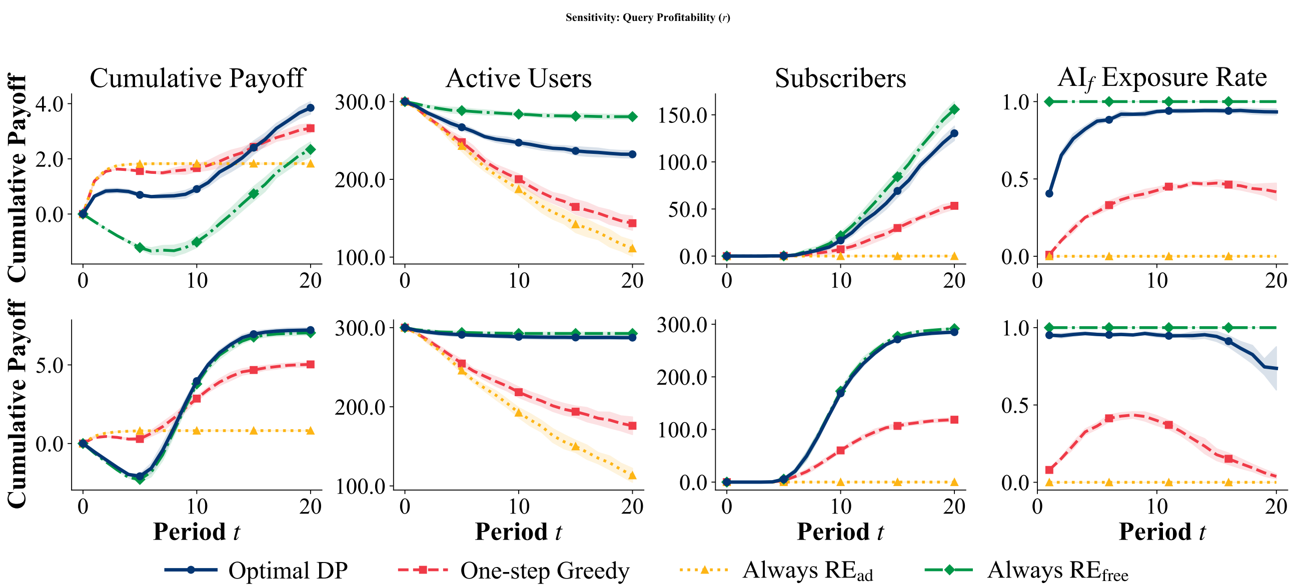

Figure 13: Sensitivity to query profitability (r): all four metrics under high-r (top) and low-r (bottom) conditions. Low-profit queries reduce the opportunity cost of $\mathrm{RE}_{\mathrm{free}}$, enabling the DP to invest more aggressively in subscription conversion.

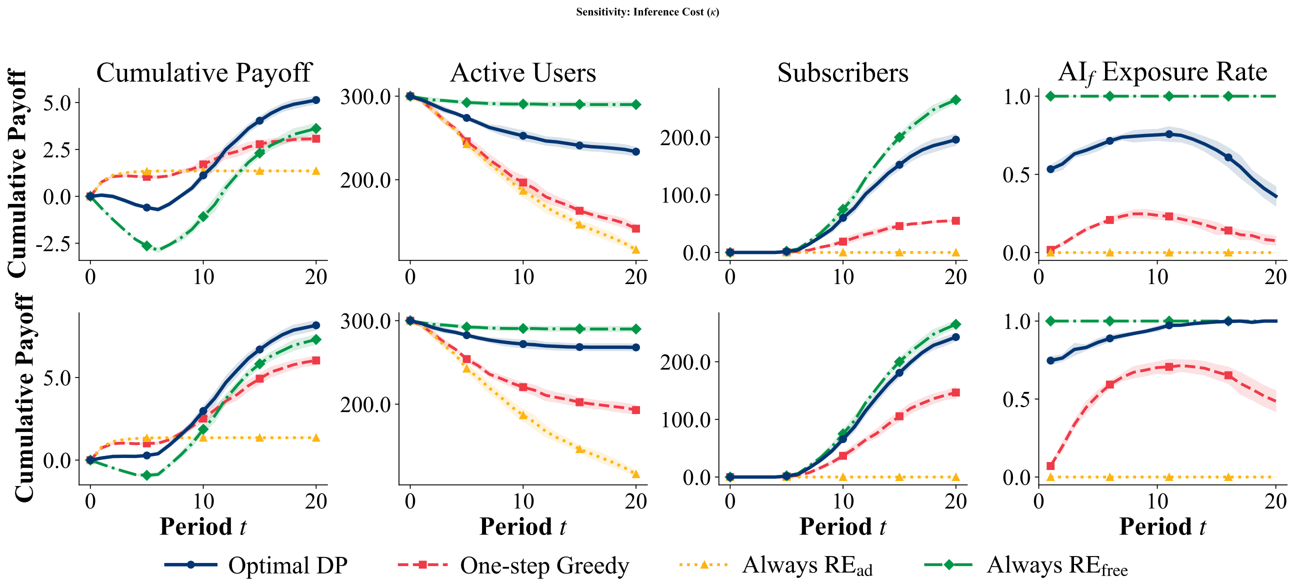

Figure 14: Sensitivity to inference cost (κ): all four metrics under high-κ (top) and low-κ (bottom) conditions. Under high inference cost, the DP curtails $\mathrm{RE}_{\text{free}}$ exposure, yet still outperforms by timing its $\mathrm{RE}_{\text{free}}$ investments. Note the negative early payoffs under high κ for Always $\mathrm{RE}_{\text{free}}$, reflecting the cost of building the subscription base.

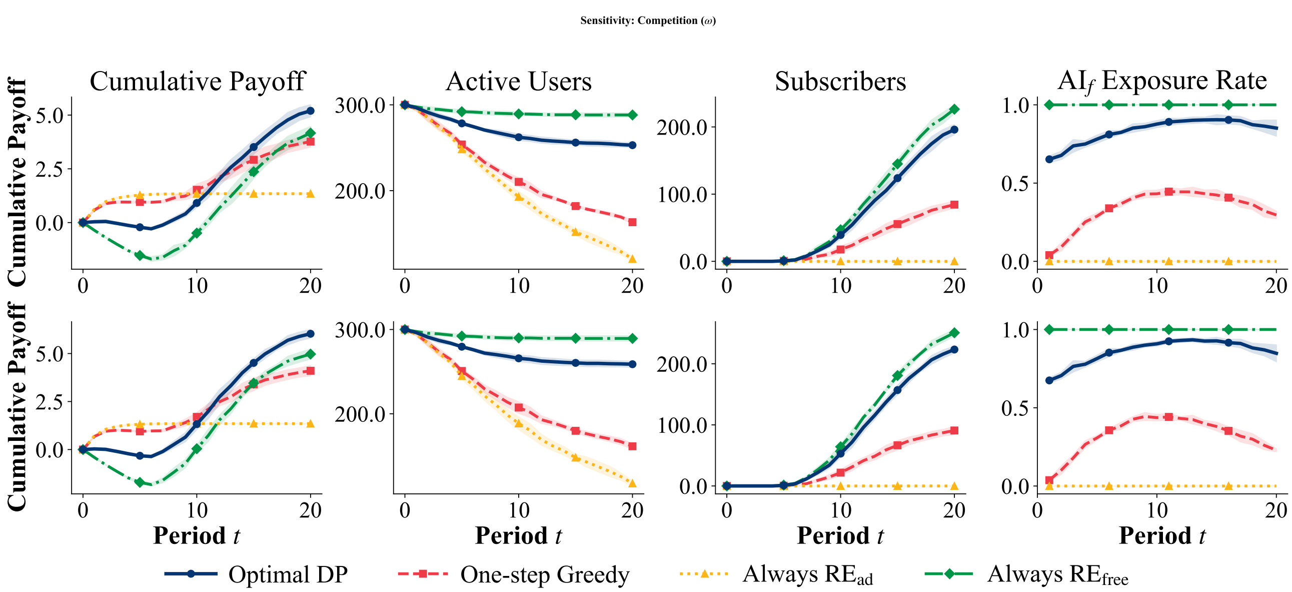

Figure 15: Sensitivity to outside competition (ω): all four metrics under high-ω (top) and low-ω (bottom) conditions. Stronger competition compresses margins but preserves the DP's structural advantage, consistent with Proposition 4.8.

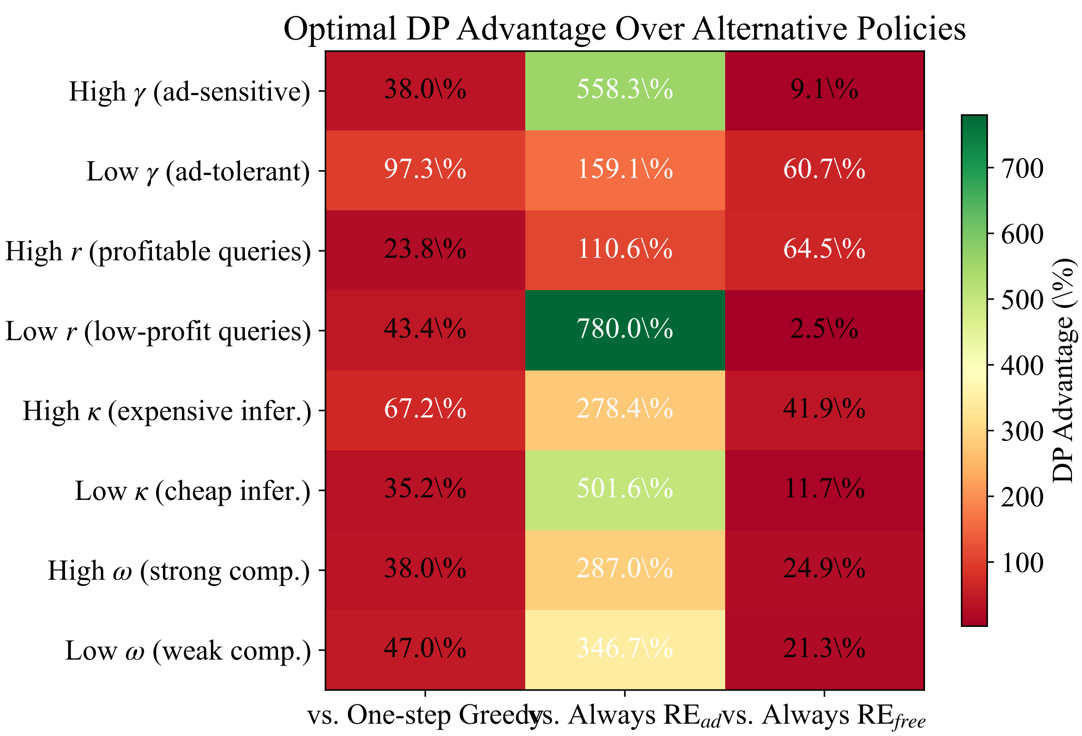

Figure 16: Heatmap of the Optimal DP's percentage advantage over each alternative policy across all eight market conditions. The advantage over Always $\mathrm{RE}_{\mathrm{ad}}$ is dominant across all conditions.