

Incentive-Aware Multi-Fidelity Optimization for Generative Advertising in Large Language Models

Anonymous ACL submission

Abstract

Generative advertising in Large Language Model (LLM) responses requires optimizing sponsorship configurations under two strict constraints: the strategic behavior of advertisers and the high cost of stochastic generations. To address this, we propose the *Incentive-Aware Multi-Fidelity Mechanism (IAMFM)*, a unified framework coupling *Vickrey-Clarke-Groves (VCG)* incentives with *Multi-Fidelity Optimization* to maximize expected social welfare offline. We compare two algorithmic instantiations—elimination-based and model-based—revealing their budget-dependent performance trade-offs. Crucially, to make VCG computationally feasible, we introduce *Active Counterfactual Optimization*, a “warm-start” that reuses optimization data for efficient payment calculation. We provide formal guarantees for approximate strategy-proofness and individual rationality, establishing a general approach for incentive-aligned, budget-constrained generative processes. Experiments demonstrate that IAMFM outperforms single-fidelity baselines across diverse budgets.

1 Introduction

Large Language Models (LLMs) are increasingly deployed as *decision-making and content-generation engines* that simultaneously affect multiple stakeholders. In many such systems, several strategic agents (e.g., advertisers, content creators, or specialized AI assistants) compete to influence a *single* generated response. We study this problem through the lens of *generative advertising*, where multiple advertisers compete to steer a single LLM chatbot response shown to the user (Duetting et al., 2024; Soumalias et al., 2024; Dubey et al., 2024; Feizi et al., 2023). A platform seeks to generate LLM responses that maximize advertiser values and user satisfaction. To do so, it periodically selects a configuration of advertiser

“strength” parameters that steers the LLM outputs over a subsequent period. Learning an effective selection strategy, however, presents several challenges. First, evaluating any candidate configuration requires repeated LLM queries and downstream evaluations, which incur substantial cost and latency. Second, advertisers are self-interested and may strategically misreport their valuations (e.g., pretending to be unsatisfied with a favorable steered LLM response) to manipulate the chosen strength configuration and, consequently, the generated responses.

To address these challenges, we propose the Incentive-Aware Multi-Fidelity Mechanism (IAMFM). To improve learning efficiency, IAMFM models each advertiser strength configuration as an arm in a multi-fidelity bandit, enabling the platform to trade off evaluation cost against accuracy during exploration. To actively avoid strategic misreporting, IAMFM incorporates a payment rule derived from the VickreyClarkeGroves (VCG) mechanism (Vickrey, 1961; Clarke, 1971; Groves, 1973) that is computed with respect to expected social welfare. In doing so, IAMFM provides a unified framework that combines the sample efficiency of multi-fidelity optimization (MFO) with incentive-compatibility guarantees from a carefully designed mechanism layered on top.

Our core contributions are as follows:

- A new *IAMFM framework* that couples multi-fidelity black-box optimization with an *ex-ante* VCG mechanism. IAMFM explicitly targets expected social welfare in LLM-based systems while **promoting truthful reporting of valuations** from strategic advertisers.
- The design and analysis of *two complementary algorithmic instantiations*—an elimination-based family (*IAMFM-SH/ASH*)

082 and a model-based MFBO variant (*IAMFM-*
083 *MFBO*)—together with empirical evidence
084 that their relative advantage is *budget-*
085 *dependent*.

- 086 • A *sample-efficient* “warm-start” procedure,
087 *Active Counterfactual Optimization (ACO)*,
088 that reuses optimization data to estimate
089 VCG counterfactuals.
- 090 • New *ex-ante* guarantees showing that, un-
091 der mild approximation conditions, our mech-
092 anism enjoys *approximately strategy-proof*,
093 *approximately individually rational*, and (*ap-*
094 *proximately*) *no-subsidy* properties, with ex-
095 plicit bounds that depend on optimization and
096 counterfactual errors.
- 097 • Empirical validation in a realistic LLM simu-
098 lation environment, providing practical guid-
099 ance on which IAMFM instantiation to de-
100 ploy under different computational budgets.

101 While we focus on generative advertising for
102 chatbots, our framework applies to a wider
103 range of tasks characterizing collaborative content-
104 generation with multiple (perhaps conflicting) ob-
105 jectives, such as AI-assisted scientific discovery
106 (negotiating trade-offs between efficacy and toxic-
107 ity in drug design), personalized education (bal-
108 ancing students’ interests with academic rigor),
109 LLM platform safety (optimizing the Pareto fron-
110 tier between helpfulness and harmlessness), etc.
111 We provide detailed formulations for these broader
112 applications in Appendix I.

113 The remainder of this paper is laid out as below.
114 Sec. 2 details the problem context and challenges.
115 Sec. 3 contains our problem formulation. Sec. 4
116 introduces our IAMFM framework and both the
117 model-based and the elimination-based instantia-
118 tions. Sec. 5 talks about the mechanism rationale
119 with theoretical proofs and acceleration of VCG
120 payment calculation. Sec. 6 evaluates the perfor-
121 mance of our method. Sec. 7 concludes the paper.
122 Appendix A reviews the related work.

123 2 Problem Setup and Challenges

124 2.1 Food-Court Generative Advertising 125 Example

126 To ground the setting, we consider a *conversa-*
127 *tional recommendation system* deployed in a shop-

ping mall’s food court.¹ There are N restaurants
(advertisers) who pay to influence the LLM’s rec-
ommendation as users ask for suggestions (e.g., “I
want something spicy and quick”). The platform
chooses an *Influence Configuration*: a vector of
advertising strength levels $\mathbf{s} = (s_1, \dots, s_N)$ that
controls how strongly each advertiser can steer the
generated content.

The platform’s goal is to choose \mathbf{s} that maxi-
mizes *total social welfare* (advertiser value + user
satisfaction). Importantly, higher strength does not
always increase welfare: overly promotional re-
sponses can reduce user satisfaction even if they
increase an advertiser’s short-term exposure. For
example, consider how Thai Spice Garden could
influence the response to the user’s query “I want
something spicy”:

Response (Low Strength): “For spicy
food, Thai cuisine (e.g., Tom Yum soup)
and Sichuan cuisine (e.g., hotpot) are
popular choices. Two local restaurants
for these are Thai Spice Garden and
Sichuan Gourmet.”

Response (High Strength): “Hands
down, the BEST choice is THAI
SPICE GARDEN. Their authentic, fiery
Panang Curry is exactly what you’re
craving. It’s the city’s must-order spicy
dish – skip it and you’ll ABSOLUTELY
regret it.”

The “High Strength” example, with its use of su-
perlatives and pushy, sales-like language, may be
rejected by users as overt, unhelpful advertising,
leading to low overall social welfare. We model
each advertiser’s preferences using a “black-box”
utility function $V_i(\cdot)$, which maps any generated
response c to a utility value. In practice, V_i is im-
plemented by an *LLM agent*, fine-tuned on data
from advertiser i ’s marketing team to act as a
proxy for their preferences.

128 2.2 The Optimization Objective

129 With these utility functions, the platform must find
130 an optimal strategy \mathbf{s}^* . However, the value of a
131 strategy \mathbf{s} is usually stochastic; it depends on the
132 specific user u and the LLM’s random generation.
133 Thus, we evaluate strategies via *sampling*: one

¹This could be a kiosk in front of the food court, or a
ubiquitous personal agent (e.g., in AR glasses) that delivers
an influenced recommendation as the user enters, based on
their preferences.

must generate many responses across many different user contexts and average the results. The platform’s objective is to find a strategy \mathbf{s}^* that maximizes the *Expected (Mean) Social Welfare*, denoted $\mathbb{E}[W]$.

2.3 Key Challenges

This selection process is constrained by three major challenges:

1. *Economic Conflict*: The advertisers are self-interested. Their interests may not align with the platform’s goal of maximizing social welfare, which creates an incentive for them to misreport their true valuations.
2. *Prohibitive Computational Cost*: The *simulation* (i.e., repeatedly generating LLM responses) required to accurately estimate $\mathbb{E}[W]$ is expensive for even *one* strategy. Performing this for *every* strategy in the vast search space is computationally intractable.
3. *Latency and Synchronization Constraints*: Users expect a response in sub-seconds. This strict requirement makes complex per-query optimization impossible. Furthermore, a real-time bidding architecture would be vulnerable to the “straggler problem,” where the system’s total latency is bottlenecked by the slowest responding agent. If any single agent lags or times out, the entire user experience breaks.

The cost and latency constraints force the platform to adopt an *offline* optimization paradigm. The expensive search for the optimal configuration must be run periodically (e.g., daily or hourly) to adapt to changing conditions. The single winning configuration from this offline run is then fixed and deployed, allowing it to serve subsequent user requests with negligible latency. This fixed deployment strategy dictates the optimization objective: the platform’s goal is to maximize the expected social welfare aggregated over the entire series of subsequent customers. In Sec. 3, we formalize this as a multi-fidelity, multi-armed bandit problem over strategies \mathbf{s} with a VCG-style mechanism layered on top.

3 Problem Formulation

We consider a generative advertising auction with N advertisers. The platform must select an *in-*

fluence configuration (or *action*) \mathbf{s} from an action space $\mathcal{A} = S^N$. This action is a vector of *strength levels* $\mathbf{s} = (s_1, \dots, s_N)$, where each $s_i \in S = \{0, 1, \dots, k\}$. Each strength level s_i quantifies an advertiser’s prominence in the output: $s_i = 0$ represents total exclusion (essential for VCG payment calculation), while $s_i = k$ triggers highly promotional language.

Our Objective: Expected Social Welfare The platform operates in short, repeated time windows (e.g., hourly). Within each window, we aim to identify the optimal action $\mathbf{s}^* \in \mathcal{A}$ that maximizes the expected social welfare, denoted by $W(\mathbf{s})$, across a distribution of user contexts \mathcal{U} . Note that \mathcal{U} does not merely represent simple demographic categories, but the vast, high-dimensional space of possible user queries and conversation histories.

While a *contextual policy* $\pi : \mathcal{U} \rightarrow \mathcal{A}$ could theoretically map each specific query to a unique action, optimizing such a complex function is prohibitively expensive for frequent updates. Furthermore, relying on sensitive demographic context can lead to unintended biases or “filter bubble” effects. Therefore, we focus on identifying a *fixed action* \mathbf{s}^* that performs best on average across the distribution \mathcal{U} , enabling the platform to precalculate and deploy an optimal configuration instantly.

Multi-Fidelity MAB Formulation To optimize this process, we model the task as a **multi-fidelity multi-armed bandit (MAB)** problem. Each unique action $\mathbf{s} \in \mathcal{A}$ is an *arm* whose latent value is the expected welfare of a complete LLM response generated with influence configuration \mathbf{s} . We introduce F distinct **fidelity levels** $f \in \{1, \dots, F\}$. A lower fidelity level f generates a shorter response prefix c_f (e.g., the first $30 \cdot 2^{f-1}$ tokens), which can be “upgraded” to a higher fidelity $f' > f$ by continuing generation. Each level incurs a token generation cost $C(f)$, where $C(1) < \dots < C(F)$. The *maximum fidelity level* F corresponds to the generation of a complete LLM response.

The outcome of any action \mathbf{s} is inherently stochastic, governed by two random variables:

1. The **user distribution** \mathcal{U} , which characterizes the diverse range of user contexts and intents encountered within the current window.
2. The **stochastic generation** c , drawn from the LLM’s generative distribution $P(c \mid \mathbf{s}, f, u)$,

which is conditional on the action, fidelity, and context.

Fidelities as Informed Estimators We define the value of a prefix c_f as the conditional expectation of the values of all possible completions:

$$V_i(c_f) := \mathbb{E}_{c_{f+1:F}}[V_i(c_F) | c_f],$$

$$V_U(c_f) := \mathbb{E}_{c_{f+1:F}}[V_U(c_F) | c_f].$$

Under the Law of Total Expectation, these definitions ensure that the expected welfare at any fidelity f satisfies:

$$\begin{aligned} & \mathbb{E}_{u,c_f}[W_{\text{realized}}(c_f)] \\ &= \mathbb{E}_{u,c_f}[\mathbb{E}_{c_{f+1:F}}[W_{\text{realized}}(c_F) | c_f]] \\ &= \mathbb{E}_{u,c_F}[W_{\text{realized}}(c_F)] = W(s). \end{aligned}$$

In practice, the platform obtains estimates $\hat{V}_i(c_f)$ and $\hat{V}_U(c_f)$ from *fine-tuned LLM agents*. Each agent functions as a specialized proxy for a specific stakeholder (each advertiser and the user), trained to predict the utility of a full response based on a partial prefix (detailed later in *Arm Pull Process* subsection). While actual agent outputs may introduce *estimation bias* if not perfectly calibrated, our multi-fidelity optimization (MFO) methods only require that the signal $\hat{W}_{\text{realized}}(c_f)$ remains strongly correlated with $W(s)$.

This converts the problem into a **cost-variance trade-off**: lower fidelities offer computationally cheap but higher-variance signals, while higher fidelities provide more precise estimations at a significantly higher cost. This structure makes our problem ideally suited for MFO methods—leveraging inexpensive, high-variance signals to identify promising candidates within the action space before committing the substantial budget required for their high-precision evaluation at maximum fidelity.

Arm Pull Process (Evaluation via LLM Agents) Evaluating an action s at fidelity f involves:

1. Sample a user context $u \sim \mathcal{U}$.
2. Generate prefix c_f with action s and context u , incurring cost $C(f)$.
3. Query fine-tuned LLM agents to obtain predictive valuations $\hat{V}_i(c_f)$, $\hat{V}_U(c_f)$.

Critically, LLM agents perform this evaluation orders of magnitude faster than human reviewers, enabling the high-throughput active learning

required for MFO. This process also allows for *human-in-the-loop (HITL)* alignment, where humans can periodically audit and correct agent labels to refine the agents’ underlying valuation models.

The Global Optimization Goal Our goal is to identify the “best arm” s^* given a fixed computational budget B :

$$s^* := \arg \max_{s \in \mathcal{A}} \hat{W}(s) \quad (1)$$

The challenge is two-fold: (1) efficiently navigating the cost-variance trade-off across fidelities to solve this optimization, and (2) designing an *approximately incentive-compatible mechanism* where advertisers have a strong incentive to truthfully report their valuations.

4 The IAMFM Framework

Our *Incentive-Aware Multi-Fidelity Mechanism (IAMFM)* integrates a *VCG-based incentive layer* with an *MFO-based optimization engine* to address the economic and computational challenges. We now detail its two complementary instantiations.

4.1 Elimination-Based Instantiations: IAMFM-SH and IAMFM-ASH

We introduce two elimination-based instantiations that navigate the *cost-variance trade-off* via a multi-stage tournament structure. We employ the classic *Successive Halving (IAMFM-SH)* as a baseline and propose a novel variant, *Adaptive Successive Halving (IAMFM-ASH)*.

Both algorithms begin with the full action space at the lowest fidelity and progressively filter out low-performing arms at each stage. The critical difference lies in how the budget is allocated *within* a stage:

- **IAMFM-SH** allocates the stage budget uniformly across all surviving arms. (See Appendix B.1 for details and pseudocode).
- **IAMFM-ASH** (Algorithm 1) replaces uniform allocation with a *UCB-based adaptive sampling strategy*. By dynamically prioritizing arms to balance the exploitation of promising candidates with the exploration of uncertain ones, IAMFM-ASH significantly enhances robustness to the stochastic noise inherent in LLM generation.

Algorithm 1: IAMFM-ASH (Elimination-Based)

Require: Action space \mathcal{A} , fidelities $1..F$, budget B , elimination rate η .

- 1: **Initialize:** $\mathcal{A}_1 \leftarrow \mathcal{A}$
- 2: **for** $f = 1$ **to** F **do**
- 3: $B_f \leftarrow$ allocate stage budget from B
- 4: *Intra-Stage Adaptive Sampling (The Core Difference between IAMFM-SH and IAMFM-ASH):*
- 5: **for** $t = 0$ **to** $B_f - 1$ **do**
- 6: $\mathbf{s}_{t+1} \leftarrow \arg \max_{\mathbf{s} \in \mathcal{A}_f} \left(\hat{\mu}_t(\mathbf{s}, f) + \beta_t \frac{\hat{\sigma}_{t-1}(\mathbf{s}, f)}{\sqrt{n_{t-1}(\mathbf{s}, f)}} \right)$
- 7: Sample $u \sim \mathcal{U}$, generate prefix c_f with influence configuration s_{t+1} , get welfare \hat{W} , update stats
- 8: **end for**
- 9: **if** $f < F$ **then**
- 10: **Elimination:** $\mathcal{A}_{f+1} \leftarrow$ Top $1/\eta$ fraction of arms in \mathcal{A}_f by mean $\hat{\mu}(\mathbf{s}, f)$
- 11: **end if**
- 12: **end for**
- 13: **return** $\mathbf{s}^* = \arg \max_{\mathbf{s} \in \mathcal{A}_F} \hat{\mu}(\mathbf{s}, F)$

4.2 Model-Based Instantiation: IAMFM-MFBO

Unlike elimination strategies, *IAMFM-MFBO* (Algorithm 2) utilizes Multi-Fidelity Bayesian Optimization to build a global *Gaussian Process (GP) surrogate model*, $W(\mathbf{s}, f)$, over the joint configuration-fidelity space. Specifically, we model the unknown welfare function W as a sample from a Gaussian Process. Under the standard GP assumption, conditioned on observed history, the value of $W(\mathbf{s}, f)$ follows a normal distribution characterized by a predictive mean $\hat{\mu}_t$ and variance $\hat{\sigma}_t^2$.

Gaussian Process Proxy. The GP’s covariance structure enables *information transfer*: observations at cheap, low-fidelity levels ($f < F$) update the model’s posterior belief (mean $\hat{\mu}$ and uncertainty $\hat{\sigma}$) regarding expensive, high-fidelity outcomes globally. This allows the platform to use noisy, inexpensive signals to identify promising regions of the action space without exhaustive high-fidelity sampling. See Appendix B.2.1 for kernel specifications.

Cost-Aware Search. To balance exploration and exploitation, we employ a budget-aware Upper Confidence Bound (UCB) acquisition function: $\text{UCB}_t(\mathbf{s}, f) = \hat{\mu}_t(\mathbf{s}, f) + \sqrt{\beta_t} \cdot \hat{\sigma}_t(\mathbf{s}, f)$. At each step, the algorithm selects the pair (\mathbf{s}, f) with the highest UCB score from the set of *fiscally feasible* options. The final recommendation is the arm with the highest posterior mean at max fidelity: $\mathbf{s}^* = \arg \max_{\mathbf{s}} \hat{\mu}_T(\mathbf{s}, F)$. The full details are in Appendix B.2.

Algorithm 2: IAMFM-MFBO (Model-Based)

Input: Action space \mathcal{A} , fidelities $1..F$, budget B , reserve R (saved for final eval)

Initialize: Train GP surrogate on initial random seed samples \mathcal{D}_0

while $\text{Cost}(\mathcal{D}_t) < B - R$ **do**

- Update exploration weight β_t via decaying schedule (See Appendix E.4 for details)
- Identify feasible set: $\mathcal{F}_t \leftarrow \{(\mathbf{s}, f) \mid \text{Cost}(\mathbf{s}, f) \leq B_{\text{remain}} - R\}$
- if** $\mathcal{F}_t = \emptyset$ **then break**
- Select $(\mathbf{s}_t, f_t) \leftarrow \arg \max_{(\mathbf{s}, f) \in \mathcal{F}_t} \text{UCB}_t(\mathbf{s}, f)$
- Evaluate (\mathbf{s}_t, f_t) to obtain welfare y_t
- Update dataset $\mathcal{D}_{t+1} \leftarrow \mathcal{D}_t \cup \{(\mathbf{s}_t, f_t, y_t)\}$ and GP posterior

end

return $\mathbf{s}^* = \arg \max_{\mathbf{s} \in \mathcal{A}} \hat{\mu}_T(\mathbf{s}, F)$

Theoretical Guarantee. Under standard regularity assumptions, IAMFM-MFBO achieves a sub-linear simple regret bound of $r_T = \mathcal{O}(\sqrt{\gamma_T/T})$, where γ_T is the maximum information gain about the function $\hat{\mu}$ after T steps, guaranteeing convergence to the optimal expected welfare. (Full proposition and proof in Appendix B.2.5).

5 VCG Payments and Incentives

5.1 Mechanism Design and Payment Rule

In a decentralized ecosystem, advertisers are strategic agents who may misreport their valuations to manipulate the outcome. To ensure that truthful reporting is a dominant strategy, we employ the *VCG mechanism*. This aligns individual incentives with the global objective by charging each agent the *opportunity cost* they impose on others.

Let \mathbf{s}^* be the optimal strategy found by our algorithm. The payment p_i for agent i is calculated as:

$$p_i = \underbrace{\max_{\mathbf{s} \in \mathcal{A}, s_i=0} \hat{W}(\mathbf{s})}_{\hat{W}_{\text{opt}}^{-i}} - \sum_{j \neq i} \hat{V}_j(\mathbf{s}^*). \quad (2)$$

Here, $\hat{W}_{\text{opt}}^{-i}$ represents the estimated *counterfactual optimum*—the maximum welfare achievable if agent i had not participated. Under this rule, an agent’s utility becomes

$$U_i = \hat{V}_i(\mathbf{s}^*) - p_i = \hat{W}(\mathbf{s}^*) - \hat{W}_{\text{opt}}^{-i}.$$

In the ideal case where estimates are perfect ($\hat{W} = W, \hat{V} = V$), maximizing U_i is strictly equivalent to maximizing global welfare $W(\mathbf{s}^*)$ because $\hat{W}_{\text{opt}}^{-i}$ is independent of i ’s bids. However, since

our budget-constrained algorithm only yields estimated values finding an ε -approximate solution, these strict properties relax to **approximate guarantees** (formal proofs in Appendix D):

- **Approximate Strategy-Proofness:** The expected gain from misreporting is strictly bounded by the optimization error ε . Intuitively, this means advertisers cannot significantly increase their profit by lying, making honesty the robust choice.
- **Approximate Individual Rationality:** Truthful participation yields expected utility $U_i \geq -\varepsilon$. This ensures that advertisers effectively never lose money by participating, encouraging them to join the auction.
- **Approximate No-Subsidy:** The mechanism requires a deficit subsidy of at most $n\varepsilon$. This guarantees that the platform does not need to inject significant external funds to keep the system running, ensuring economic viability.

5.2 Active Counterfactual Optimization (ACO)

A major bottleneck in VCG is computational cost: calculating payments requires solving N additional “cold-start” optimization problems to find each W_{opt}^i . To solve this, we introduce the *Active Counterfactual Optimization (ACO)* framework. The core idea is to “warm-start” the counterfactual searches by reusing information gathered during the main optimization phase. It significantly reduces the cost of payment calculation compared to naive methods.

- For *IAMFM-MFBO*, we reuse the trained Gaussian Process as a *global prior*, allowing the counterfactual search to converge with negligible additional samples.
- For *IAMFM-ASH*, we reuse the *sample history* of the main tournament to initialize the counterfactual search.

Detailed algorithms for ACO are provided in Appendix C.

While both warm-start strategies outperform cold-start optimization, ACO for IAMFM-MFBO is significantly more sample-efficient. The “global prior” captures structural trends across the entire action space, whereas in ACO for IAMFM-ASH, the “local reuse of sample history” is restricted to the specific points sampled by chance.

Consequently, as the number of advertisers N grows, the budget savings from IAMFM-MFBO’s global warm-start become increasingly dominant, as shown in our experiments in Sec. 6.

6 Experimental Evaluation

We evaluate our proposed instantiations of the IAMFM framework against standard single-fidelity baselines. Our framework assesses performance across a wide range of computational budgets, reflecting the real-world constraints of pre-calculating strategies in LLM-based advertising systems.

Simulation Environment. We evaluate our algorithms in a simulated food court recommendation system with two strategic advertisers and five distinct user personas (e.g., “Health-conscious”, “Budget-conscious”; see Appendix E for prompts). The action space is discrete, consisting of all 25 strength combinations ($s_A, s_B \in \{0, \dots, 4\}$). The objective is to find the configuration that maximizes *Expected Social Welfare* (sum of advertiser and user utilities), which is estimated by LLM evaluators.

We define four fidelity levels with corresponding token costs: F1 (30 tokens), F2 (60 tokens), F3 (120 tokens), and F4 (240 tokens). We compare our multi-fidelity instantiations (IAMFM-MFBO, IAMFM-ASH, IAMFM-SH) against two standard single-fidelity baselines: **UCB Bandit** and **Uniform Sampling** (both operating at max fidelity). Performance is evaluated across 8 computational budgets (1k to 128k tokens) with 10 independent trials each. For analysis, we stratify results into a *low-budget regime* ($\leq 16\text{k}$) and a *high-budget regime* ($\geq 32\text{k}$).

Results and Analysis Our experiments reveal a distinct performance trade-off governed by the available computational budget, as visualized in Figure 1.

In the *Low-Budget Regime* ($\leq 16\text{k}$ tokens), both IAMFM-ASH and IAMFM-MFBO dominate, achieving statistically significant gains over single-fidelity baselines (e.g., ASH vs. UCB: $p < 0.0001$). At this stage, data is scarce, and the primary value comes from aggressively pruning poor configurations using cheap, low-fidelity samples. Consequently, the heuristic elimination of ASH is as effective as the model-based MFBO ($p = 0.38$). However, the landscape shifts in the *High-Budget*

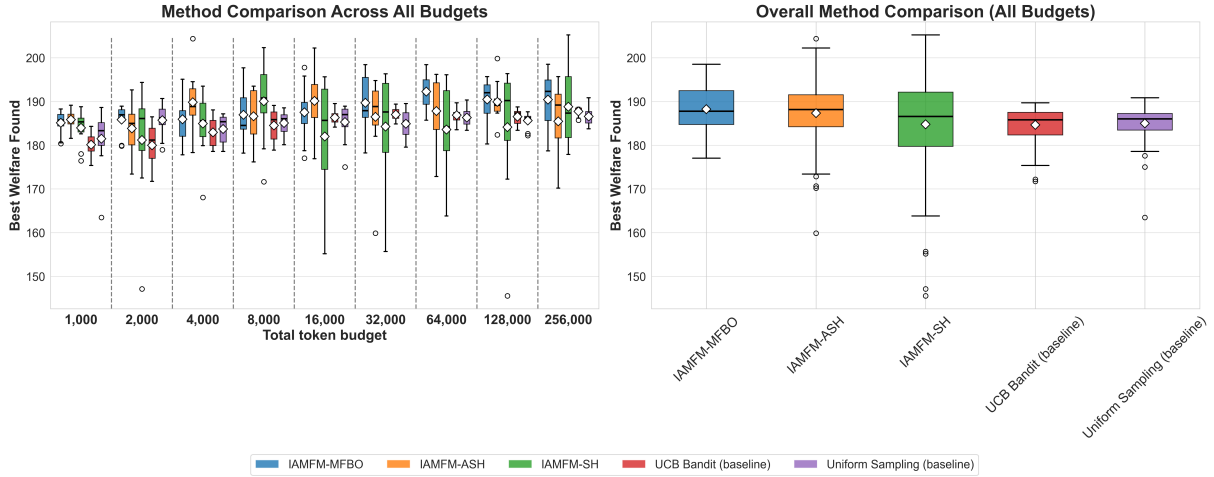


Figure 1: Comparison of welfare across all methods. Left: Box plots stratified by token budget summarize performance, showing the median (center line), mean (diamond), interquartile range (box), and outliers (points). All reported statistics include outliers. The IAMFM methods outperform the baselines and reveal a token-budget-dependent shift in the optimal method. Right: Aggregated performance across all budgets. Our proposed multi-fidelity instantiations (IAMFM-MFBO, IAMFM-ASH) show superior median performance, while IAMFM-MFBO demonstrates a clear advantage at higher budgets.

518 *Regime* ($\geq 32k$ tokens). Here, IAMFM-MFBO
 519 emerges as the statistically superior strategy, out-
 520 performing all baselines and, critically, surpassing
 521 IAMFM-ASH ($p = 0.036$). As the budget allows
 522 for more pulls, the GP surrogate’s ability to model
 523 the global landscape becomes the deciding factor,
 524 finding precise optima that heuristic elimination
 525 strategies miss. See Appendix E.6 for more details.
 526

527 Another finding is the superior stability of our
 528 model-based approach. The standard elimination
 529 baseline (IAMFM-SH) suffers from high variance
 530 (std. dev. 10.1–11.7), confirming that rigid elimi-
 531 nation schedules are brittle: one unlucky low-
 532 fidelity sample can permanently discard the opti-
 533 mal arm. In contrast, IAMFM-MFBO achieves
 534 the *lowest variance* among all methods (std. dev.
 535 5.0–5.5). By aggregating information across the
 536 entire action-fidelity space, the global GP surro-
 537 gate effectively recovers the latent signal from the
 538 stochastic LLM generations, preventing the algo-
 539 rithm from being misled by noisy samples. This
 540 makes MFBO the most robust choice for produc-
 541 tion deployment.

542 We observe that as budgets grow very large
 543 ($> 128k$ tokens), the performance gap between
 544 our methods and brute-force baselines (Uniform/
 545 UCB) narrows. This is expected: with infinite
 546 resources, even random sampling will eventually
 547 find the optimum. The significance of IAMFM
 548 is that it achieves near-optimal welfare with sig-

549 *nificantly fewer tokens*, operating efficiently in
 550 the realistic regime where brute-force is cost-
 551 prohibitive.

552 **VCG Payment Calculation Efficiency.** Finally,
 553 we evaluate the efficiency of our *Active Counter-*
 554 *factual Optimization (ACO)*. Results demonstrate
 555 that *ACO for MFBO* is decisively superior: by
 556 reusing the global GP prior, it achieves near-zero
 557 estimation error for the counterfactual optimum
 558 with *zero additional budget*. Details can be found
 559 in Appendix F.3. This reduces the VCG compu-
 560 tational overhead by $> 99\%$ compared to cold-
 561 start baselines. In contrast, *ACO for ASH* offers
 562 only modest improvements, as reusing local sam-
 563 ple history lacks the generalization power of the
 564 global surrogate. This shows that the model-based
 565 approach is better in making truthful mechanisms
 566 computationally feasible.

567 **Qualitative Analysis: The Failure of Greedy**
 568 **Strategies.** Figure 2 compares responses gener-
 569 ated by different methods. While a *naive greedy*
 570 *strategy* ($s_A = s_B = 4$) pushes the LLM into *se-*
 571 *mantic collapse*, producing disjointed and halluci-
 572 natory text, IAMFM identifies the optimal strength
 573 configuration ($s_A = s_B = 2$). This configura-
 574 tion allows the model to maintain a coherent nar-
 575 rative structure, thereby preserving user trust and,
 576 thereby, achieving high total welfare.

Table 1: Aggregated performance (mean welfare \pm std. dev.) stratified by budget regime, based on 10 trials per budget-algorithm pair. Methods are ranked by mean welfare. **Bolded** are proposed methods.

Budget Regime	Method	Mean Welfare \pm Std. Dev.
Low-Budget ($\leq 16k$ tokens)	1. IAMFM-ASH	187.26 \pm 6.00
	2. IAMFM-MFBO	186.26 \pm 5.02
	3. IAMFM-SH	184.42 \pm 10.10
	4. Uniform Sampling	184.26 \pm 4.62
	5. UCB Bandit	182.78 \pm 4.15
High-Budget ($\geq 32k$ tokens)	1. IAMFM-MFBO	190.73 \pm 5.53
	2. IAMFM-ASH	187.43 \pm 7.93
	3. UCB Bandit	187.04 \pm 1.58
	4. Uniform Sampling	185.87 \pm 2.41
	5. IAMFM-SH	185.21 \pm 11.68

7 Conclusion

We proposed the *Incentive-Aware Multi-Fidelity Mechanism (IAMFM)*, integrating VCG incentives with multi-fidelity optimization to address the economic and computational constraints of generative advertising. Empirically, we found that the optimal strategy is *budget-dependent*: IAMFM-ASH works equally well in low-budget regimes, while IAMFM-MFBO dominates at higher budgets. Additionally, we demonstrated that our ACO framework renders VCG payment calculations computationally feasible by reusing the data and model from the optimization process. Extended discussions on alternative steering mechanisms, broader applications, and user transparency are provided in Appendix G.

Limitations

While our framework demonstrates significant efficacy in finding the best strength configuration in optimizing social welfare under incentive compatibility constraints, we acknowledge certain limitations that open avenues for future research.

First, our experimental evaluation relies on LLM-based agents to simulate user preferences and advertiser valuations. While this serves as a scalable and reproducible proxy for human behavior—standard in recent computational social science literature—it may not fully capture the nuance and unpredictability of real-world human feedback. Future work could involve human-in-the-loop experiments to further validate the external validity of our welfare estimates.

Furthermore, the objective landscapes induced by LLMs may present sharp discontinuities (e.g., at refusal or guardrail boundaries), which violate

Case Study: The Failure of Greed Strategies

Context (User 1): “Health-conscious professional... looking for fresh ingredients...”

Config A: Naive Greedy ($s_A = 4, s_B = 4$)

Scenario: Both advertisers bid Max Strength.

Generated Content: “The vibrant flavors are truly exceptional. For a different healthy choice, QuickBite’s grilled chicken... is specifically designed for nutritious meals... **nourishes both body and mind...**”

× **Failure (Disjointed & Hallucinated):** The text exhibits a jarring disconnect: it opens with a generic praise (likely intended for Thai Spice) but is immediately hijacked by the Fast Food prompt. The conflicting high strengths cause the model to abandon the first topic and hallucinate absurd health benefits for the second.

Welfare: 152.32 (User Value: 45.73)

Config B: IAMFM Optimized ($s_A = 2, s_B = 2$)

Scenario: Algorithm identifies the global optimum.

Generated Content: “QuickBite’s value... [is] undeniable... making healthy choices accessible... **Meanwhile**, Thai Spice Garden excels in creating an authentic dining experience... For a health-conscious professional... Thai Spice offers a reward, **while** QuickBite stands ready...”

✓ **Success (Coherent Balance):** The moderate configuration enables the LLM to structure a comparative narrative. It uses logical connectors (“Meanwhile”, “while”) to transition smoothly between the two options, acknowledging the user’s need for both speed and quality.

Welfare: 207.53 (User Value: 73.16)

Figure 2: Qualitative comparison. The naive greedy strategy (Config A) results in disjointed text where one advertiser overpowers the other, leading to hallucinations. IAMFM (Config B) identifies the optimal influence configuration, improving total welfare by +36% through enhanced semantic coherence.

the smoothness assumptions of stationary GP surrogates. While switching to non-stationary models (e.g., mixture models or treed GPs) can empirically better capture such regime shifts, establishing formal convergence guarantees in this setting is non-trivial. Standard regret bounds typically assume the target function has bounded complexity within the Reproducing Kernel Hilbert Space (RKHS). True discontinuities generally break this assumption, and employing non-stationary kernels does not automatically restore these guarantees without additional, explicit assumptions regarding the structure of the discontinuities. Developing optimization algorithms with provable guarantees for such non-smooth, LLM-induced objectives remains a critical direction for future work.

References

629
630
631
632
633
634

Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G Wilson, and Eytan Bakshy. 2020. Botorch: A framework for efficient monte-carlo bayesian optimization. *Advances in neural information processing systems*, 33:21524–21538.

635
636
637
638

Joschka Braun, Dmitrii Krasheninnikov, Usman Anwar, Robert Kirk, Daniel Tan, and David Scott Krueger. 2024. A sober look at steering vectors for llms. *LessWrong*, November, 23.

639
640
641
642
643
644
645

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:1877–1901.

646
647

Edward H Clarke. 1971. Multipart pricing of public goods. *Public choice*, pages 17–33.

648
649
650

Vincent Conitzer and Tuomas Sandholm. 2002. Complexity of mechanism design. *arXiv preprint cs/0205075*.

651
652

Vincent Conitzer and Tuomas Sandholm. 2003a. Applications of automated mechanism design.

653
654
655
656

Vincent Conitzer and Tuomas Sandholm. 2003b. Automated mechanism design for a self-interested designer. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 232–233.

657
658
659
660

Hoagy Cunningham, Aidan Ewart, Logan Riggs, Robert Huben, and Lee Sharkey. 2023. Sparse autoencoders find highly interpretable features in language models. *arXiv preprint arXiv:2309.08600*.

661
662
663
664

Michael Curry, Tuomas Sandholm, and John Dickerson. 2022. Differentiable economics for randomized affine maximizer auctions. *arXiv preprint arXiv:2202.02872*.

665
666
667
668
669
670

Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2020. Plug and play language models: A simple approach to controlled text generation. In *International Conference on Learning Representations (ICLR)*.

671
672
673
674
675

Avinava Dubey, Zhe Feng, Rahul Kidambi, Aranyak Mehta, and Di Wang. 2024. Auctions with llm summaries. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 713–722.

676
677
678
679

Paul Duetting, Vahab Mirrokni, Renato Paes Leme, Haifeng Xu, and Song Zuo. 2024. Mechanism design for large language models. In *Proceedings of the ACM on Web Conference 2024*, pages 144–155.

Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David C Parkes, and Sai Srivatsa Ravindranath. 2024. Optimal auctions through deep learning: Advances in differentiable economics. *Journal of the ACM*, 71(1):1–53. 680
681
682
683
684

Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259. 685
686
687
688
689

Soheil Feizi, MohammadTaghi Hajiaghayi, Keivan Rezaei, and Suho Shin. 2023. Online advertisements with llms: Opportunities and challenges. *arXiv preprint arXiv:2311.07601*. 690
691
692
693

Zhe Feng, Harikrishna Narasimhan, and David C Parkes. 2018. Deep learning for revenue-optimal auctions with budgets. In *Proceedings of the 17th international conference on autonomous agents and multiagent systems*, pages 354–362. 694
695
696
697
698

Jessie Finocchiaro, Roland Maio, Faidra Monachou, Gourab K Patro, Manish Raghavan, Ana-Andreea Stoica, and Stratis Tsirtsis. 2021. Bridging machine learning and mechanism design towards algorithmic fairness. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 489–503. 699
700
701
702
703
704
705

Leo Gao, Tom Dupré la Tour, Henk Tillman, Gabriel Goh, Rajan Troll, Alec Radford, Ilya Sutskever, Jan Leike, and Jeffrey Wu. 2024. Scaling and evaluating sparse autoencoders. *arXiv preprint arXiv:2406.04093*. 706
707
708
709
710

Theodore Groves. 1973. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631. 711
712
713

Chi Han, Jialiang Xu, Manling Li, Yi Fung, Chenkai Sun, Nan Jiang, Tarek Abdelzaher, and Heng Ji. 2024. Word embeddings are steers for language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16410–16430. 714
715
716
717
718
719

John J Horton. 2023. [Large language models as simulated economic agents: What can we learn from homo silicus?](#) *Preprint*, arXiv:2301.07543. 720
721
722

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for NLP. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 2790–2799. 723
724
725
726
727
728

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3. 729
730
731
732

733	Chengsong Huang, Qian Liu, Bill Yuchen Lin, Tianyu Pang, Chao Du, and Min Lin. 2023. LoRAHub: Efficient cross-task generalization via dynamic LoRA composition. <i>arXiv preprint arXiv:2307.13269</i> .	786
734		787
735		788
736		789
737	Kirthevasan Kandasamy, Gautam Dasarathy, Junier Oliva, Jeff Schneider, and Barnabas Poczos. 2019. Multi-fidelity gaussian process bandit optimisation. <i>Journal of Artificial Intelligence Research</i> , 66:151–196.	790
738		791
739		792
740		793
741		794
742	Kirthevasan Kandasamy, Gautam Dasarathy, Jeff Schneider, and Barnabás Póczos. 2017. Multi-fidelity bayesian optimisation with continuous approximations. In <i>International conference on machine learning</i> , pages 1799–1808. PMLR.	795
743		796
744		797
745		798
746		799
747	Nikita Koriagin, Yaroslav Aksenov, Daniil Laptev, Gleb Gerasimov, Nikita Balagansky, and Daniil Gavrillov. 2025. Teach old SAEs new domain tricks with boosting. In <i>Second Conference on Language Modeling</i> .	800
748		801
749		802
750		803
751		804
752	Ben Krause, Akhilesh Deepak Gotmare, Bryan McCann, Nitish Shirish Keskar, Shafiq Joty, Richard Socher, and Nazneen Fatema Rajani. 2021. Gedi: Generative discriminator guided sequence generation. In <i>Findings of the Association for Computational Linguistics: EMNLP 2021</i> , pages 4929–4952.	805
753		806
754		807
755		808
756		809
757		810
758	Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Ros-tamizadeh, and Ameet Talwalkar. 2018. Hyperband: A novel bandit-based approach to hyperparameter optimization. <i>Journal of Machine Learning Research</i> , 18(185):1–52.	811
759		812
760		813
761		814
762		815
763	Vladislav Lialin, Vijeta Deshpande, and Anna Rumshisky. 2023. Scaling down to scale up: A guide to parameter-efficient fine-tuning. <i>arXiv preprint arXiv:2303.15647</i> .	816
764		817
765		818
766		819
767	Tom Lieberum, Senthoran Rajamanoharan, Arthur Conmy, Lewis Smith, Nicolas Sonnerat, Vikrant Varma, János Kramár, Anca Dragan, Rohin Shah, and Neel Nanda. 2024. Gemma scope: Open sparse autoencoders everywhere all at once on gemma 2. <i>Preprint</i> , arXiv:2408.05147.	820
768		821
769		822
770		823
771		824
772		825
773	Jiayuan Liu, Mingyu Guo, and Vincent Conitzer. 2025. An interpretable automated mechanism design framework with large language models. <i>arXiv preprint arXiv:2502.12203</i> .	826
774		827
775		828
776		829
777	Noam Nisan and Amir Ronen. 2001. Algorithmic mechanism design. <i>Games and Economic Behavior</i> , 35:166–196. Early version in STOC-99.	830
778		831
779		832
780	Kyle O’Brien, David Majercak, Xavier Fernandes, Richard Edgar, Jingya Chen, Harsha Nori, Dean Carignan, Eric Horvitz, and Forough Poursabzi-Sangde. 2024. Steering language model refusal with sparse autoencoders. <i>arXiv preprint arXiv:2411.11296</i> .	833
781		834
782		835
783		836
784		837
785		838
		839
		840
		841
		842
		843
		844
		845
		846
		847
		848
		849
		850
		851
		852
		853
		854
		855
		856
		857
		858
		859
		860
		861
		862
		863
		864
		865
		866
		867
		868
		869
		870
		871
		872
		873
		874
		875
		876
		877
		878
		879
		880
		881
		882
		883
		884
		885
		886
		887
		888
		889
		890
		891
		892
		893
		894
		895
		896
		897
		898
		899
		900

842	Embedding-aligned language models. <i>Advances in Neural Information Processing Systems</i> , 37:15893–15946.	895
843		896
844		897
845	Megh Thakkar, Léo Boisvert, Thibault Le Sellier De Chezelles, Alexandre Piché, Maxime Gasse, Alexandre Lacoste, and Massimo Caccia. Agent-merge: Enhancing generalization in fine-tuned llm agents. In <i>Adaptive Foundation Models: Evolving AI for Personalized and Efficient Learning</i> .	898
846		899
847		900
848		901
849		902
850		903
851	Hal R Varian. 2007. Position auctions. <i>International Journal of Industrial Organization</i> , 25(6):1163–1178.	904
852		905
853		906
854	William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. <i>The Journal of Finance</i> , 16(1):8–37.	907
855		908
856		909
857	Tonghan Wang, Yanchen Jiang, and David C. Parkes. 2024. Gemnet: Menu-based, strategy-proof multi-bidder auctions through deep learning. In <i>Proceedings of the 25th ACM Conference on Economics and Computation</i> , EC ’24, page 1100, New York, NY, USA. Association for Computing Machinery.	910
858		911
859		912
860		913
861		914
862		915
863	Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and 1 others. 2022. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In <i>International conference on machine learning</i> , pages 23965–23998. PMLR.	916
864		917
865		918
866		919
867		920
868		921
869		922
870		923
871	Jian Wu, Saul Toscano-Palmerin, Peter I Frazier, and Andrew Gordon Wilson. 2020. Practical multi-fidelity bayesian optimization for hyperparameter tuning. In <i>Uncertainty in Artificial Intelligence</i> , pages 788–798. PMLR.	924
872		925
873		926
874		927
875		928
876	Kevin Yang and Dan Klein. 2021. FUDGE: Controlled text generation with future discriminators. In <i>Proceedings of the 2021 Conference of the North American Chapter of the ACL: Human Language Technologies (NAACL-HLT)</i> , pages 3511–3535.	929
877		930
878		931
879		932
880		933
881	A Related Work	934
882	This work lies at the intersection of mechanism design for advertising, computationally-aware optimization, and the emerging application of LLMs in multi-agent systems.	935
883		936
884		937
885		938
886	Influencing LLMs We are interested in methods that increase exposure of certain topics and brands, or even produce customizable messages, without any post-generation editing. Influencing modern, transformer-based LLMs can occur at all stages. At post-training, parameter-efficient fine-tuning (PEFT) inserts small modules (adapters like LoRA) that achieve strong effects but require careful composition (e.g., AdapterFusion	939
887		940
888		941
889		942
890		943
891		944
892		945
893		
894		
	and LoRA can be used to blend multiple styles or stakeholder goals (Houlsby et al., 2019; Hu et al., 2022; Pfeiffer et al., 2021; Huang et al., 2023)). At inference, cheap prompting and system messages remain the most agile levers, though they can be brittle or context-hungry (Brown et al., 2020). Plug-and-play decoding steers generation via auxiliary scores or gradients without updating weights at some latency/fluency cost (Dathathri et al., 2020; Krause et al., 2021; Yang and Klein, 2021). Activation steering with sparse autoencoders edits interpretable, extensible features on-the-fly, offering modular, transparently continuous control knobs on influence strength (O’Brien et al., 2024; Koriagin et al., 2025). By contrast, deliberate curation of duplicated data, full pretraining, RLHF, or DPO hard-wire behaviors (Ouyang et al., 2022; Rafailov et al., 2023), which are impractical for rapid, per-query, multi-stakeholder advertising.	
	Mechanism Design for Advertising Auctions	
	The field of advertising auctions is foundational to our work. Traditional search advertising has been dominated by the Generalized Second-Price (GSP) auction (Edelman et al., 2007; Varian, 2007). While widely adopted, GSP is not strategy-proof. The Vickrey-Clarke-Groves (VCG) (Vickrey, 1961; Clarke, 1971; Groves, 1973) mechanism serves as the primary truthful alternative, guaranteeing that truthful bidding is a dominant strategy for advertisers.	
	However, these classical models were designed for discrete, “atomic” ad slots. They do not address the challenges of generative advertising, where the “good” being auctioned is a complex, co-created piece of content whose quality is a result of the combined inputs from multiple strategic agents. Our work extends the principles of VCG to this new, generative domain.	
	Computation-Aware Mechanism Design A significant challenge in mechanism design is the computational complexity of determining outcomes and payments, especially with many agents or a large outcome space. Prior work has addressed this in several ways. One line of research focuses on approximating VCG payments, often through sampling-based methods that can provide unbiased estimates but may suffer from high variance (Sandholm and Likhodedov, 2005). Another line of work on algorithmic mechanism design concerns the joint design of (approxima-	

946	tion) algorithms and payment functions that make	particularly active thread concerns auction mecha-	996
947	those algorithms into truthful mechanisms (Nisan	nisms for generative content.	997
948	and Ronen, 2001). Yet another direction is auto-	Duetting et al. (2024) introduce a token-by-	998
949	mated mechanism design (Conitzer and Sandholm,	token bidding framework, in which advertiser	999
950	2002; Sandholm, 2003), which leverages optimiza-	agents place bids for insertion slots during text	1000
951	tion (Conitzer and Sandholm, 2003a,b), machine	generation, laying the groundwork for applying se-	1001
952	learning (Dütting et al., 2024; Wang et al., 2024;	quential auction theory to LLMs. Building on this	1002
953	Shen et al., 2018; Rahme et al., 2021; Feng et al.,	foundation, Soumalias et al. (2024) address vulner-	1003
954	2018), and more recently large language mod-	abilities such as misreporting and the “exposure	1004
955	els (Liu et al., 2025) to derive optimal or approx-	problem,” where small wording changes can dras-	1005
956	imately optimal mechanisms. These approaches	tically alter advertiser value. Their mechanisms	1006
957	often relax the requirement of perfect strategy-	emphasize truthfulness and robustness, often by	1007
958	proofness to achieve computational tractability.	penalizing deviations from a baseline “truthful” re-	1008
959	In our approach, rather than first determining an	sponse. In parallel, Dubey et al. (2024) general-	1009
960	outcome and then struggling to compute payments,	ize position auction concepts, incorporating pre-	1010
961	we integrate the search for the optimal outcome	dicted click-through rates based on ad placement	1011
962	and the calculation of its associated payments into	within LLM responses, an approach reminiscent	1012
963	a single, unified optimization process. Our key	of search engine advertising.	1013
964	contribution is a method to make the VCG calcula-	Despite these advances, several challenges re-	1014
965	tion itself sample-efficient by reusing information.	main. As Feizi et al. (2023) survey, current mech-	1015
966	Multi-fidelity Bayesian Optimization and Band-	anisms struggle to aggregate multiple advertisers	1016
967	its To address the expensive, black-box nature	preferences consistently and to guarantee mono-	1017
968	of evaluating LLM outputs, we turn to the litera-	tonic value propositions. Moreover, these efforts	1018
969	ture on sequential optimization. Our problem,	connect to broader debates in AI mechanism de-	1019
970	with its discrete set of strength combinations, can	sign about embedding fairness, robustness, and so-	1020
971	be framed as a multi-armed bandit (MAB) prob-	cial impact into economic systems alongside tra-	1021
972	lem. Since evaluations can be performed at vary-	ditional efficiency and revenue objectives (Finoc-	1022
973	ing levels of completeness (i.e., number of to-	chiaro et al., 2021).	1023
974	kens), this is specifically a multi-fidelity optima-	B Algorithm Details	1024
975	zation problem.	B.1 IAMFM-SH	1025
976	State-of-the-art methods for this include non-	Here we provide the details for the standard elim-	1026
977	Bayesian approaches like Successive Halving	ination baseline (Algorithm 3). Successive Halv-	1027
978	(SH) and Hyperband (Li et al., 2018), and	ing balances exploration and exploitation through	1028
979	Bayesian approaches like Multi-fidelity Bayesian	staged commitment. Unlike IAMFM-ASH, which	1029
980	Optimization (MFBO) (Kandasamy et al., 2017).	uses UCB scores to select arms during a stage,	1030
981	We build upon MFBO using a Gaussian Process	IAMFM-SH uses a fixed, uniform schedule.	1031
982	surrogate model to allocate our limited computa-	As the algorithm progresses through stages $f \in$	1032
983	tional budget. The primary novelty of our work	$\{1, \dots, F\}$, it aggressively reduces the active set	1033
984	is the application of MFBO not merely as an op-	of arms \mathcal{A}_f by a factor of η . This ensures that the	1034
985	timization tool for a single decision-maker, but as	platform’s computational budget is concentrated	1035
986	the core engine within an incentive-aware mech-	on the most promising candidates as their evalu-	1036
987	anism designed for strategic agents.	ation becomes more precise.	1037
988	LLMs in Economic and Multi-Agent Settings	B.2 IAMFM-MFBO	1038
989	A growing line of research explores the use of	B.2.1 Gaussian Process and Kernel Structure	1039
990	large language models (LLMs) in economic and	To model the correlation between “similar strength	1040
991	multi-agent contexts. LLMs have been deployed	levels” and “adjacent fidelities” as described in the	1041
992	to simulate human-like behavior as agents in vari-	methodology, we employ a Structured Product	1042
993	ous games (Horton, 2023), and to serve as evalua-	Kernel . The kernel function $k(\mathbf{x}, \mathbf{x}')$ decomposes	1043
994	tors of complex, qualitative outcomes. Within this	into a spatial component over the action space and	1044
995	broader direction, and closest to our own work, a		

Algorithm 3: IAMFM-Successive Halving (IAMFM-SH)

Require: Action space \mathcal{A} , fidelities $1..F$, budget B , elimination rate η .

- 1: **Initialize:** $\mathcal{A}_1 \leftarrow \mathcal{A}$, determine samples per arm n_f .
- 2: **for** $f = 1$ **to** $F - 1$ **do**
- 3: **for all** active arms $s \in \mathcal{A}_f$ **do**
- 4: **Uniform Sampling:** Sample u , generate prefix c_f for n_f iterations.
- 5: Update empirical mean $\hat{\mu}_f(s)$.
- 6: **end for**
- 7: **Elimination:** Rank arms by $\hat{\mu}_f(s)$ and retain top $1/\eta$ fraction.
- 8: **Fidelity Upgrade:** $\mathcal{A}_{f+1} \leftarrow$ Survivors; prepare for next fidelity.
- 9: **end for**
- 10: **Final Selection:** Perform max-fidelity evaluation for $s \in \mathcal{A}_F$.
- 11: **return** $s^* = \arg \max_{s \in \mathcal{A}_F} \hat{\mu}_F(s)$.

1045 a correlation component over the fidelity space:

1046
$$k(\mathbf{x}, \mathbf{x}') = \sigma^2 \cdot k_{\text{action}}(\mathbf{s}, \mathbf{s}') \otimes k_{\text{fidelity}}(f, f') \quad (3)$$

- 1047 • **Action Kernel** (k_{action}): We use an **RBF**
1048 **(Radial Basis Function) Kernel** with Auto-
1049 matic Relevance Determination (ARD). This
1050 captures the smoothness of the welfare land-
1051 scape, assuming that configuration vectors \mathbf{s}
1052 close in Euclidean distance yield similar wel-
1053 fare. ARD allows the model to learn that
1054 some advertisers may have a stronger influ-
1055 ence on welfare (shorter lengthscales) than
1056 others.
- 1057 • **Fidelity Kernel** (k_{fidelity}): We use a special-
1058 ized **Downsampling Kernel** (or Exponential
1059 Decay Kernel) to model the information
1060 transfer. It assumes that the correlation be-
1061 tween fidelity f and f' decreases as the dis-
1062 tance $|f - f'|$ increases, but remains non-zero,
1063 allowing low-fidelity observations to update
1064 high-fidelity beliefs.

1065 Appendix E.4 contains the details for experiment
1066 implementation.

1067 B.2.2 Dynamic Exploration Schedule (β_t)

1068 The exploration-exploitation trade-off parameter
1069 β_t in the UCB acquisition function is not con-
1070 stant. We implement a **budget-dependent dy-**
1071 **namic schedule:**

1072
$$\beta_t = \beta_{\text{start}} \cdot \left(\frac{B_{\text{remaining}}}{B_{\text{total}}} \right)^\gamma + \beta_{\text{min}} \quad (4)$$

1073 where β_{start} is initialized high (encouraging broad
1074 exploration of the action space early on) and de-

1075 creases as the remaining budget decreases. This “sim-
1076 ulated annealing” approach ensures the algorithm
1077 transitions from exploration to exploitation (refin-
1078 ing the estimate of the best arm) as the budget is
1079 exhausted.

1080 B.2.3 Fiscal Feasibility and Budget Reserve

1081 The set of “fiscally feasible” options \mathcal{F}_t ensures
1082 the algorithm never exhausts its budget on low-
1083 fidelity exploration without being able to verify
1084 the winner. We define a **Strategic Reserve** R_{final}
1085 sufficient to perform K evaluations at the maxi-
1086 mum fidelity F . A pair (s, f) is in \mathcal{F}_t if and only
1087 if:

1088
$$C(\mathbf{s}, f) \leq B_{\text{current}} - R_{\text{final}} \quad (5)$$

1089 where $C(\mathbf{s}, f)$ is the marginal token cost to eval-
1090 uate configuration \mathbf{s} at fidelity f . This constraint
1091 forces the algorithm to stop exploration and switch
1092 to exploitation (final confirmation at fidelity F)
1093 when the budget becomes critical.

1094 B.2.4 Denoised Selection Policy

1095 Unlike standard bandit algorithms that might re-
1096 turn the arm with the highest *observed* reward,
1097 IAMFM-MFBO selects the final arm based on the
1098 **Posterior Mean** of the GP:

1099
$$\mathbf{s}^* = \arg \max_{\mathbf{s} \in \mathcal{A}} \hat{\mu}_T(\mathbf{s}, F)$$

1100 This utilizes the GP’s ability to aggregate informa-
1101 tion from all multi-fidelity samples to “denoise”
1102 the stochastic LLM outputs, providing a robust
1103 recommendation even if the specific observations
1104 were subject to high variance.

1105 B.2.5 Theoretical Guarantees

1106 The efficacy of our *IAMFM-MFBO* instantiation is
1107 grounded in the theoretical framework of Bayesian
1108 optimization. Its core MFBO engine uses the MF-
1109 GP-UCB acquisition function, which provides a
1110 principled mechanism to manage the exploration-
1111 exploitation trade-off. This ensures that the algo-
1112 rithm is guaranteed to converge to the optimal ac-
1113 tion. We formalize this guarantee, which follows
1114 from prior work, in the following proposition.

1115 **Proposition 1** (Regret Bound for
1116 IAMFM-MFBO (Srinivas et al., 2012; Kan-
1117 dasamy et al., 2019)). *Let the expected social*
1118 *welfare function $\mu : \mathcal{X} \rightarrow \mathbb{R}$, where \mathcal{X} is the*
1119 *joint action-fidelity space, be a sample from a*
1120 *Gaussian Process. Under standard regularity*
1121 *assumptions on the kernel and noise, the simple*

regret r_T of the IAMFM-MFBO algorithm after T evaluations is bounded. The simple regret $r_T = \hat{\mu}(\mathbf{s}^*, F) - \hat{\mu}(\mathbf{s}_T, F)$ compares the expected welfare μ of the true best action \mathbf{s}^* (at max fidelity F) and the action \mathbf{s}_T recommended by the algorithm. With high probability, this regret is bounded by:

$$r_T = \mathcal{O}\left(\sqrt{\frac{\gamma_T}{T}}\right)$$

where γ_T is the maximum information gain about the function $\hat{\mu}$ after T steps.

Proof Sketch. The proof follows directly from the regret analysis of the GP-UCB algorithm (Srinivas et al., 2012), which was subsequently extended to the multi-fidelity setting as analyzed in (Kandasamy et al., 2019). Our IAMFM-MFBO instantiation employs the same MF-GP-UCB core as the latter. Therefore, under standard regularity assumptions, our algorithm inherits its theoretical convergence properties and regret bounds. \square

Proposition 1 confirms that our model-based algorithm (IAMFM-MFBO) is a provably efficient tool for our task, as it guarantees convergence to the arm with the highest expected welfare ($\hat{\mu}(\mathbf{s}^*, F)$).

This theoretical result is crucial for generative advertising. It shows that our mechanism is a *provably efficient* method for identifying the combination of advertiser strengths that produces the highest *expected* value for all participants. This foundation of efficiency and fairness is essential for the VCG payment mechanism, aligning incentives to favor near-optimal expected outcomes.

C Active Counterfactual Optimization (ACO) Details

Computing N separate, cold-start optimizations for each counterfactual scenario is prohibitively expensive. We propose the **Active Counterfactual Optimization (ACO)** framework to “warm-start” these subproblems by reusing information from the main optimization phase.

ACO for IAMFM-MFBO (Global Informative Prior) For our model-based instantiation, ACO exploits the **generalization capability** of the GP surrogate from the main optimization. Because the GP learns the welfare landscape as a continuous function, samples where $s_i > 0$ still provide

significant information about the $s_i = 0$ boundary through the covariance kernel.

1. **Global Prior Transfer:** We initialize an $(N - 1)$ -dimensional MFBO process for the $s_i = 0$ subspace, using the posterior of the main optimization’s GP as a *highly informative prior*.
2. **Targeted Refinement:** We allocate a small additional budget B_{cf} to perform targeted UCB sampling within this subspace. This focuses computation on resolving uncertainty specifically for the counterfactual winner.
3. **Denoised Estimation:** The estimate \hat{W}^{-i} is defined as the highest posterior mean welfare found by the GP: $\hat{W}^{-i} := \max_{\mathbf{s} \in \mathcal{A}, s_i = 0} \mu_{T'}(\mathbf{s}, F)$.

ACO for IAMFM-SH/ASH (Local Sample Reuse) For our model-free, elimination-based instantiations, a global prior does not exist. ACO is instead implemented by reusing the *local sample history* gathered during the main optimization.

1. **History Filtering:** The platform aggregates all previous samples where $s_i = 0$ was naturally selected during the main tournament stages.
2. **Warm-Started Tournament:** These historical samples form the initial rungs of a new SH/ASH process. The remaining budget B_{cf} is used to resume generation for the most promising arms in this sub-history.
3. **Empirical Estimation:** The estimate \hat{W}^{-i} is the final empirical mean welfare $\hat{\mu}(\mathbf{s}^{-i})$ of the winning arm.

D Mechanism Rationale (Propositions and Proofs)

Notation We list the notations used in the propositions as follows:

- N : The set of advertisers.
- $\mathbf{s} \in \mathcal{A}$: An outcome (strategy).
- $V_i(\mathbf{s}) := \mathbb{E}[V_i \text{ from } \mathbf{s}]$: The true (unknown) expected value for agent i from strategy \mathbf{s} .
- $W(\mathbf{s}) := \sum_{i \in N} V_i(\mathbf{s})$: The true expected social welfare of strategy \mathbf{s} .

- $b_i(\mathbf{s})$: The bid (reported expected value) of agent i for strategy \mathbf{s} .
- $\mathbf{s}^*(b)$: The strategy chosen by our algorithm given the bid profile $b = (b_1, \dots, b_N)$.
- $V = (V_1, \dots, V_N)$: The profile of truthful bids (i.e., $b_i = V_i$ for all i).
- $U_i(b) = V_i(\mathbf{s}^*(b)) - p_i(\mathbf{s}^*(b))$: Agent i 's true expected utility under reports b .
- $p_i(\mathbf{s}) = h_i(b_{-i}) - \sum_{j \neq i} b_j(\mathbf{s})$: The Groves payment rule.

Our mechanism builds on VCG mechanism. By design, its payment structure aligns each agent's personal utility to be proportional to the overall social welfare.

We denote advertiser i 's true expected valuation for the chosen strategy \mathbf{s} as $V_i(\mathbf{s}) := \mathbb{E}_{u \sim \mathcal{U}, c \sim P(\cdot | \mathbf{s}, F, u)} [V_i(c)]$. An advertiser's final profit is:

$$U_i = V_i(\mathbf{s}^*) - p_i(\mathbf{s}^*).$$

Under the general Groves family of mechanisms, the payment p_i for strategy \mathbf{s} is defined as:

$$p_i(\mathbf{s}) = h_i(b_{-i}) - \sum_{j \neq i} b_j(\mathbf{s}),$$

where $b_j(\mathbf{s})$ is agent j 's bid (representing their reported *mean valuation*, or expected valuation, for strategy \mathbf{s}) and h_i does not depend on agent i 's own bid, b_i .

Substituting this into the utility function reveals the alignment. An agent i , when bidding, chooses a bid function $b_i(\cdot)$ to maximize their expected utility. Assuming others bid truthfully ($b_j(\mathbf{s}) = V_j(\mathbf{s})$), agent i 's objective is to choose a bid b_i that results in a winning strategy \mathbf{s}^* that maximizes:

$$\begin{aligned} U_i &= V_i(\mathbf{s}^*) + \sum_{j \neq i} V_j(\mathbf{s}^*) - h_i(b_{-i}) \\ &= \left(\sum_{j \in N} V_j(\mathbf{s}^*) \right) - h_i(b_{-i}) \propto W(\mathbf{s}^*) \end{aligned}$$

This alignment means an agent's best path to maximizing their own expected profit is to provide the mechanism with bids that help it discover the strategy \mathbf{s}^* with the highest true expected social welfare, $W(\mathbf{s})$. This perfectly aligns our incentive mechanism with the optimization objective defined in Section 3.

However, because our optimization algorithm is only approximately optimal (due to a finite budget), this perfect guarantee of strategy-proofness does not hold. Therefore, we do not claim perfect truthfulness but instead aim to achieve *approximate strategy-proofness*.

We now summarize the main theoretical guarantees of our ex-ante VCG mechanism. In Propositions 2, 3, and 4, we analyze its incentive properties and relate the deviations from classical VCG to the optimization and counterfactual errors. Let N be the finite set of advertisers. Let $\mathbf{s} \in \mathcal{A}$ be an outcome (strategy). Let $W(\mathbf{s}) := \mathbb{E}[W(\mathbf{s}, F)]$ be the *true expected social welfare* of strategy \mathbf{s} , and $V_i(\mathbf{s})$ the true expected value for agent i . Let \mathbf{s}^* be the strategy selected by our algorithm based on agent reports $b = (b_1, \dots, b_N)$. The mechanism uses a Groves-style allocation rule: $\mathbf{s}^* = \arg \max_{\mathbf{s}} \sum_j b_j(\mathbf{s})$, and payments are $p_i(\mathbf{s}) = h_i(b_{-i}) - \sum_{j \neq i} b_j(\mathbf{s})$.

Proposition 2 (Bounded incentive to misreport under approximate welfare optimality). *Suppose the strategy \mathbf{s}^* chosen by our algorithm (when all agents are truthful, $b = V$) is ε -approximate with respect to the true optimal expected welfare, in the sense that*

$$W(\mathbf{s}^*) \geq W(\mathbf{s}) - \varepsilon \quad \text{for all strategies } \mathbf{s} \in \mathcal{A},$$

then for every single agent $i \in N$ and for every unilateral misreport b'_i , their expected utility U_i satisfies:

$$U_i(b'_i, V_{-i}) \leq U_i(V_i, V_{-i}) + \varepsilon,$$

where $U_i(b) = V_i(\mathbf{s}^(b)) - p_i(\mathbf{s}^*(b))$ is agent i 's true expected utility under reports b .*

Proof of Proposition 2. (Approximate Strategy-Proofness). Fix an agent $i \in N$. Let V_{-i} denote the truthful bids from all agents other than i . Let $\mathbf{s}_{\text{truth}}^* = \mathbf{s}^*(V_i, V_{-i})$ be the strategy chosen when agent i bids truthfully. Agent i 's truthful expected utility is:

$$\begin{aligned} U_i(V_i, V_{-i}) &= V_i(\mathbf{s}_{\text{truth}}^*) - p_i(\mathbf{s}_{\text{truth}}^*) \\ &= V_i(\mathbf{s}_{\text{truth}}^*) - \left(h_i(V_{-i}) - \sum_{j \neq i} V_j(\mathbf{s}_{\text{truth}}^*) \right) \\ &= V_i(\mathbf{s}_{\text{truth}}^*) + \sum_{j \neq i} V_j(\mathbf{s}_{\text{truth}}^*) - h_i(V_{-i}) \\ &= W(\mathbf{s}_{\text{truth}}^*) - h_i(V_{-i}) \end{aligned}$$

Now, suppose agent i unilaterally misreports b'_i , while others stay truthful. Let $\mathbf{s}_{\text{lic}}^* = \mathbf{s}^*(b'_i, V_{-i})$ be

the strategy selected by the mechanism under this deviation. The deviating agent's utility is:

$$\begin{aligned} U_i(b'_i, V_{-i}) &= V_i(\mathbf{s}_{\text{lic}}^*) - p_i(\mathbf{s}_{\text{lic}}^*) \\ &= V_i(\mathbf{s}_{\text{lic}}^*) - (h_i(V_{-i}) - \sum_{j \neq i} V_j(\mathbf{s}_{\text{lic}}^*)) \\ &= W(\mathbf{s}_{\text{lic}}^*) - h_i(V_{-i}) \end{aligned}$$

Note that $h_i(V_{-i})$ is the same in both cases, as it does not depend on i 's bid.

The agent's gain from deviating is the difference in utilities:

$$U_i(b'_i, V_{-i}) - U_i(V_i, V_{-i}) = W(\mathbf{s}_{\text{lic}}^*) - W(\mathbf{s}_{\text{truth}}^*)$$

By the proposition's assumption, the strategy $\mathbf{s}_{\text{truth}}^*$ chosen under truthful bidding is ε -approximate with respect to the true optimal *expected* welfare. This means $W(\mathbf{s}_{\text{truth}}^*) \geq W(\mathbf{s}) - \varepsilon$ for all $\mathbf{s} \in \mathcal{A}$. Since $\mathbf{s}_{\text{lic}}^*$ is one such strategy, it must be that:

$$W(\mathbf{s}_{\text{truth}}^*) \geq W(\mathbf{s}_{\text{lic}}^*) - \varepsilon$$

Rearranging this gives:

$$W(\mathbf{s}_{\text{lic}}^*) - W(\mathbf{s}_{\text{truth}}^*) \leq \varepsilon$$

Therefore, the agent's gain from lying is bounded by ε :

$$U_i(b'_i, V_{-i}) - U_i(V_i, V_{-i}) \leq \varepsilon$$

This proves the claimed bound $U_i(b'_i, V_{-i}) \leq U_i(V_i, V_{-i}) + \varepsilon$. \square

Proposition 2 establishes that the mechanism is *approximately strategy-proof*: any agent's potential gain in *expected utility* from lying is bounded by ε , the suboptimality of the *expected* outcome. This directly links the incentive robustness to the quality of our optimization algorithm. In the ideal case where our algorithm is perfect ($\varepsilon = 0$), truthful reporting becomes a dominant strategy.

Proposition 3 (*Approximate Individual Rationality/Voluntary Participation*). Assume the mechanism uses the Clarke-pivot (others-best) form for the *expected values*:

$$h_i(b_{-i}) = \max_{\mathbf{s}: s_i=0} \sum_{j \neq i} b_j(\mathbf{s}) \quad (= W_{\text{opt}}^{-i} \text{ if } b_{-i} = V_{-i})$$

Let \hat{W}^{-i} be our algorithm's estimate of W_{opt}^{-i} used for the payment $p_i(\mathbf{s}^*) = \hat{W}^{-i} - \sum_{j \neq i} b_j(\mathbf{s}^*)$.

Suppose under truthful reporting, our main optimization $W(\mathbf{s}^*)$ and counterfactual optimization \hat{W}^{-i} satisfy: $W(\mathbf{s}^*) \geq \hat{W}^{-i} - \varepsilon$, for some $\varepsilon \geq 0$, then agent i 's true expected utility satisfies

$$U_i^* = V_i(\mathbf{s}^*) - p_i(\mathbf{s}^*) = W(\mathbf{s}^*) - \hat{W}^{-i} \geq -\varepsilon.$$

In particular, if our main optimization finds a better (or equal) expected welfare than our counterfactual optimization ($W(\mathbf{s}^*) \geq \hat{W}^{-i}$), truthful participation is *ex-ante* individually rational.

Proof of Proposition 3. (Approximate Individual Rationality). We assume the mechanism uses the Clarke-pivot rule, where $h_i(b_{-i})$ is the true optimal counterfactual welfare, $W_{\text{opt}}^{-i} := \max_{\mathbf{s}: s_i=0} \sum_{j \neq i} b_j(\mathbf{s})$. Under truthful reporting ($b = V$), the payment for the winning strategy \mathbf{s}^* is:

$$p_i(\mathbf{s}^*) = \hat{W}^{-i} - \sum_{j \neq i} V_j(\mathbf{s}^*)$$

where \hat{W}^{-i} is our algorithm's *estimate* of W_{opt}^{-i} .

The truthful expected utility U_i^* for agent i is:

$$U_i^* = V_i(\mathbf{s}^*) - p_i(\mathbf{s}^*)$$

$$= V_i(\mathbf{s}^*) - \left(\hat{W}^{-i} - \sum_{j \neq i} V_j(\mathbf{s}^*) \right)$$

$$= \left(V_i(\mathbf{s}^*) + \sum_{j \neq i} V_j(\mathbf{s}^*) \right) - \hat{W}^{-i}$$

$$= W(\mathbf{s}^*) - \hat{W}^{-i}$$

By the proposition's assumption, our algorithm's estimates satisfy $W(\mathbf{s}^*) \geq \hat{W}^{-i} - \varepsilon$. Substituting this into the utility equation, we get:

$$U_i^* = W(\mathbf{s}^*) - \hat{W}^{-i} \geq -\varepsilon$$

This proves the claimed bound. If $W(\mathbf{s}^*) \geq \hat{W}^{-i}$, then $U_i^* \geq 0$, and truthful participation is *ex-ante* individually rational. \square

Proposition 4 ((Approximate) no-subsidy guarantee). Let \hat{W}^{-i} be our algorithm's estimate of the optimal counterfactual welfare for each i . For any $\eta \geq 0$ the mechanism requires at most an additive subsidy of η (i.e. $\sum_i p_i(\mathbf{s}^*) \geq -\eta$) iff

$$(n-1)W(\mathbf{s}^*) \leq \sum_{i \in N} \hat{W}^{-i} + \eta. \quad (6)$$

In particular, if the counterfactual estimates satisfy the uniform lower bound

$$\hat{W}^{-i} \geq W(\mathbf{s}^*) - \varepsilon \quad \text{for all } i,$$

then

$$\sum_{i \in N} p_i(\mathbf{s}^*) \geq W(\mathbf{s}^*) - n\varepsilon,$$

so the deficit (if any) is at most $n\varepsilon - W(\mathbf{s}^*)$. Equivalently, if $W(\mathbf{s}^*) \geq n\varepsilon$ then no subsidy is needed.

Proof of Proposition 4. (Approximate No-Subsidy). Let $n = |N|$ be the number of agents. We analyze the total payments collected by the mechanism under truthful reporting ($b = V$), based on the winning strategy \mathbf{s}^* and the N counterfactual estimates \hat{W}^{-i} .

The payment for a single agent i is:

$$p_i(\mathbf{s}^*) = \hat{W}^{-i} - \sum_{j \neq i} V_j(\mathbf{s}^*)$$

Summing over all agents $i \in N$:

$$\sum_{i \in N} p_i(\mathbf{s}^*) = \sum_{i \in N} \hat{W}^{-i} - \sum_{i \in N} \left(\sum_{j \neq i} V_j(\mathbf{s}^*) \right)$$

We analyze the second term. It is the sum of welfare for everyone *except* agent i , summed over all i . In this summation, each agent j 's value $V_j(\mathbf{s}^*)$ appears exactly $(n - 1)$ times (once for each $i \neq j$).

$$\begin{aligned} \sum_{i \in N} \left(\sum_{j \neq i} V_j(\mathbf{s}^*) \right) &= (n - 1) \sum_{j \in N} V_j(\mathbf{s}^*) \\ &= (n - 1)W(\mathbf{s}^*) \end{aligned}$$

Substituting this back, we get the exact identity for total payments:

$$\sum_{i \in N} p_i(\mathbf{s}^*) = \sum_{i \in N} \hat{W}^{-i} - (n - 1)W(\mathbf{s}^*) \quad (7)$$

The mechanism requires at most an additive subsidy of η if and only if the total payments are $\sum_i p_i(\mathbf{s}^*) \geq -\eta$. Substituting the identity, this is equivalent to:

$$\sum_{i \in N} \hat{W}^{-i} - (n - 1)W(\mathbf{s}^*) \geq -\eta$$

Rearranging this proves the equivalence in Eq. (6):

$$(n - 1)W(\mathbf{s}^*) \leq \sum_{i \in N} \hat{W}^{-i} + \eta$$

For the second part of the proposition, we assume the uniform lower bound $\hat{W}^{-i} \geq W(\mathbf{s}^*) - \varepsilon$ for all i . Summing this bound over all N agents:

$$\sum_{i \in N} \hat{W}^{-i} \geq \sum_{i \in N} (W(\mathbf{s}^*) - \varepsilon) = nW(\mathbf{s}^*) - n\varepsilon$$

Now, substitute this inequality back into our payment identity (Eq. (7)):

$$\begin{aligned} \sum_{i \in N} p_i(\mathbf{s}^*) &\geq (nW(\mathbf{s}^*) - n\varepsilon) - (n - 1)W(\mathbf{s}^*) \\ &= W(\mathbf{s}^*) - n\varepsilon \end{aligned}$$

This proves the final claim. \square

Corollary 1 (Sufficient average-condition for (approximate) no-subsidy). *A convenient sufficient condition for a small subsidy is an ‘‘average others-only’’ bound. If there exists $\eta \geq 0$ such that*

$$\frac{1}{n} \sum_{i \in N} \hat{W}^{-i} \geq \frac{n - 1}{n} W(\mathbf{s}^*) - \frac{\eta}{n},$$

then the mechanism requires at most subsidy η , i.e. $\sum_i p_i(\mathbf{s}^*) \geq -\eta$.

Proof of Corollary (Sufficient average-condition). This follows directly from Proposition 4. We start with the inequality from the corollary’s assumption:

$$\frac{1}{n} \sum_{i \in N} \hat{W}^{-i} \geq \frac{n - 1}{n} W(\mathbf{s}^*) - \frac{\eta}{n}$$

Multiply the entire inequality by n :

$$\sum_{i \in N} \hat{W}^{-i} \geq (n - 1)W(\mathbf{s}^*) - \eta$$

Rearrange the terms:

$$(n - 1)W(\mathbf{s}^*) \leq \sum_{i \in N} \hat{W}^{-i} + \eta$$

By Proposition 4 (Eq. (6)), this condition is equivalent to $\sum_i p_i(\mathbf{s}^*) \geq -\eta$. \square

Note that the property of bounded incentive to misreport depends *only* on the quality of the *main optimization*. In contrast, the properties of voluntary participation and no subsidy both depend on how well the *counterfactual optimization estimates* perform relative to the *main optimization’s outcome*.

E Experiment Details

This appendix details the experimental infrastructure, implementation, and prompts.

1442
1443
1444
1445
1446
1447
1448

1449
1450
1451
1452

1453
1454
1455

1456
1457
1458

1459
1460
1461
1462

1463
1464
1465

1466
1467
1468
1469

1470
1471
1472
1473

1474
1475
1476

1477
1478

1479
1480
1481

1482
1483
1484

E.1 Problem and Simulation Environment

Action Space All experiments address the generative advertising auction problem with two advertisers (A and B). The action space is discrete, consisting of the $5 \times 5 = 25$ possible combinations of strength levels, as each advertiser’s strength s_i is drawn from the set $S = \{0, 1, 2, 3, 4\}$.

Fidelity and Cost Model We use $F = 4$ fidelity levels. As detailed in the main paper (Sec. 6), these correspond to the token costs: $C(1) = 30$, $C(2) = 60$, $C(3) = 120$, and $C(4) = 240$.

Advertiser (Restaurant) Context The simulation uses two competing restaurants embedded in the system prompt:

- **Restaurant A (QuickBite):** A fast-food restaurant known for speed, affordability, and family-friendly options.
- **Restaurant B (Thai Spice Garden):** An authentic Thai restaurant with fresh ingredients and bold flavors, appealing to adventurous eaters.

Sponsorship strength levels $s_A, s_B \in \{0, \dots, 4\}$ were mapped to “no/little/some/great/maximal effort” in the prompt.

User Personas and Prompts We evaluate on 5 distinct user personas to test the algorithms’ ability to optimize for different trade-offs (e.g., speed vs. flavor).

- **User 1 (Health-conscious):** “I’m looking for a healthy lunch option. I prefer fresh ingredients, lighter meals, and options with vegetables...”
- **User 2 (Budget-conscious):** “I’m a student on a tight budget. I need something filling and affordable...”
- **User 3 (Adventurous foodie):** “I love spicy food and bold flavors! I’m adventurous...”
- **User 4 (Time-constrained):** “I only have 15 minutes for lunch... I need something fast, convenient, and easy to eat...”
- **User 5 (Parent):** “I’m here with my two kids (ages 6 and 9). We need somewhere with kid-friendly options...”

LLM Infrastructure We use two distinct, state-of-the-art LLMs for our experiments to separate the roles of content generation and evaluation:

- **Content Generator (DeepSeek-chat):** Used for all content generation tasks, with a temperature of 0.8.
- **Evaluator (GPT-4o-mini):** Used for all utility evaluation tasks (both advertiser and user).

E.2 Evaluation Protocol

Welfare Computation. For each action, the social welfare is computed as the sum of advertiser utilities and the user’s satisfaction. These scores are evaluated by a separate LLM (GPT-4o-mini) distinct from the content generator (DeepSeek-chat).

Ground Truth Evaluation. Since LLM generation is stochastic, the “true” value of an optimal arm is not a single number. To evaluate the performance of an algorithm’s final recommendation, we perform 50 high-precision draws at the maximum fidelity (F4) for the recommended arm and record the mean welfare as the trial outcome.

Reproducibility Infrastructure. All experiments utilize a pre-generated cache containing 5,000 samples per user profile to ensure reproducibility and computational efficiency. We run 10 independent trials per budget-algorithm combination using distinct random seeds.

Statistical Tests. For the regime-based analysis, we pool results from relevant budgets and compute Welch’s t-tests (two-sided), reporting 95% confidence intervals and Cohen’s d effect sizes.

E.3 Multi-Fidelity Cache

Tree-Structured Generation To ensure reproducibility and coherence, experiments run on a pre-generated cache. The cache was built using a tree structure:

1. **F1 (Roots):** Independent responses were generated for all 25 arms at the F1 (30 token) target.
2. **F2–F4 (Continuations):** Higher fidelities (60, 120, 240 tokens) were generated as continuations from a selected parent, e.g., “Continue this response to reach approximately 60 tokens total: ... [F1 content] ...”.

This tree structure ensures that a high-fidelity sample is a coherent elaboration of its low-fidelity prefix, which is essential for the MFO assumption.

Fidelity Evaluation Logic Our evaluation logic (Sec. 3) requires low-fidelity evaluations to be predictive of the final, high-fidelity outcome. We implemented this directly:

- For **F4** (max fidelity), the evaluator was asked to score the *final, complete* response.
- For **F1, F2, F3** (prefixes), the evaluator (GPT-4o-mini) was explicitly told the response was “partial” and was instructed to **estimate the expected utility of the final, completed F4 response** that would follow from this prefix.

This aligns the observed reward $y(\mathbf{s}, f)$ at all fidelities with the true arm value $W(\mathbf{s}) := W(\mathbf{s}, F)$.

Welfare Calculation The final welfare for any sample is the unweighted sum of the three parties’ utilities: $W_{\text{realized}}(c) = V_A(c) + V_B(c) + V_U(c)$.

E.4 IAMFM-MFBO Implementation Details

Our IAMFM-MFBO implementation includes key components tuned for our setup.

Fidelity and Cost Model In our experiments, we use $F = 4$ fidelity levels with token costs: $C(1) = 30$, $C(2) = 60$, $C(3) = 120$, and $C(4) = 240$.

Gaussian Process Surrogate and Implementation Details We model the multi-fidelity welfare surface using a SingleTaskMultiFidelityGP (Wu et al., 2020) implemented via BoTorch (Balandat et al., 2020). Each input is represented as a normalized vector $\mathbf{x} = [s_1/k, \dots, s_N/k, (f - 1)/(F - 1)] \in [0, 1]^{N+1}$, where k and F denote the maximum advertiser strength and fidelity level, respectively. Following this architecture, the covariance function factorizes into a structured product kernel: $k(\mathbf{x}, \mathbf{x}') = \sigma^2 \cdot k_X(\mathbf{x}_{1:N}, \mathbf{x}'_{1:N}) \otimes k_F(x_{N+1}, x'_{N+1})$. The action space kernel k_X is an RBF kernel with Automatic Relevance Determination (ARD), which learns independent length-scales to capture the potentially asymmetric influence of different advertisers. For the fidelity dimension, we employ a DownsamplingKernel k_F , which models the correlation between fidelities f and f' using a power-law structure: $k_F(f, f') = c + (1 - f)^{1+\delta}(1 - f')^{1+\delta}$, where $f, f' \in [0, 1]$ are the normalized fidelity coordinates. This specific

power-law structure ensures that the surrogate captures the diminishing bias and variance as the fidelity approaches the target value of 1.0.

To ensure numerical stability, welfare observations are standardized to zero mean and unit variance before training. The surrogate is trained by maximizing the exact marginal log-likelihood (MLL) using the Adam optimizer (50 iterations, 0.1 learning rate). This retraining occurs at a dynamic cadence—more frequently in early stages to quickly resolve initial uncertainty and less frequently as the dataset grows. Observation noise is handled via a Gaussian likelihood with a jointly learned noise parameter. At inference, the posterior distribution guides the GP-UCB acquisition rule, $\text{UCB}(\mathbf{x}) = \mu(\mathbf{x}) + \beta_t \sigma(\mathbf{x})$, allowing the platform to strategically navigate the cost-variance trade-off across the action-fidelity space.

Tree-Structured Sampling Critically, the sampling process respects the tree-structure of LLM generation, where a higher-fidelity response is a *continuation* of a lower-fidelity one. Evaluations respect the cached parent-to-child links. For a given action (\mathbf{s}), a sample at fidelity $F2$ or higher is always a child of the previously chosen lower-fidelity sample for that branch. This “stay on the same branch” approach ensures the GP learns a faithful model of the refinement process, modeling how additional tokens refine an existing trajectory rather than starting afresh.

Dynamic β Schedule Details The hybrid β schedule’s start and end points are budget-dependent. For example, for medium budgets (e.g., $B = 8k$), the schedule decays from $\beta_{\text{start}} = 8.5$ to $\beta_{\text{end}} = 1.7$. For very large budgets, it uses a wider range (e.g., $\beta_{\text{start}} = 15.0$ to $\beta_{\text{end}} = 3.0$) to encourage more exploration.

Strategic Budget Reservation A core feature of our implementation is its strategic budget reservation. A fixed portion of the total budget B (25% in our experiments, `f4_reserve_ratio=0.25`) is held in an **F4 reserve**. This reserve is actively protected as a *feasibility constraint* during the selection loop at each step t . Any (action, fidelity) pair where $f < F$ (i.e., not a final confirmation) is considered *infeasible* if its marginal cost ΔC (calculated as described below) would dip into this reserve. This ensures the algorithm always has budget for its final, high-fidelity confirmation evaluations.

Dynamic Marginal Cost Model The marginal token cost ΔC used in the feasibility check is not fixed; it is calculated dynamically based on the evaluation history of an arm s :

- **First-time Upgrade Cost:** Evaluating s at a new, higher fidelity $f > f_{\text{prev}}$ (where f_{prev} is the highest fidelity *previously* evaluated for s) costs only the incremental amount: $\Delta C = C(f) - C(f_{\text{prev}})$, while $C(f_{\text{prev}}) = 0$ if this arm is not yet evaluated at any fidelity level.
- **Resample Cost:** Re-evaluating s at a previously seen fidelity f costs the full fidelity cost: $\Delta C = C(f)$.

This cost model accurately reflects the true marginal token expenditure, distinguishing between the cost to *upgrade* a fidelity path (which reuses prior lower-fidelity work) and the cost to *resample* it (which requires a full new evaluation at that fidelity).

E.5 Content Generation and Evaluation Prompts

Strength Level Implementation The discrete strength levels $S = \{0, 1, 2, 3, 4\}$ are translated into descriptive text to guide the generator LLM, as detailed in Table 2.

Table 2: Mapping of numerical strength levels to prompt descriptions.

Strength	Label	Description for Prompt
0	Very Weak	minimal presence (indirect references only)
1	Weak	weak presence (subtle, background references)
2	Moderate	moderate presence (balanced, natural integration)
3	Strong	strong presence (prominent, featured examples)
4	Very Strong	very strong presence (dominant, detailed focus)

System Prompt Template The core system prompt provided to DeepSeek-chat instantiates the task, advertisers, and user persona.

Code Specification 1: System prompt template.

- ```

1 You are a helpful AI assistant at a
 food court information kiosk.
2
3 There are two restaurants:
4 1. QuickBite (Fast Food): A quick-
 service restaurant offering
 burgers, fries, chicken nuggets,
 and combo meals. Known for speed,

```

```

 affordability, and family-
 friendly options. Great for kids
 and budget-conscious diners.
5 2. Thai Spice Garden (Thai Cuisine):
 An authentic Thai restaurant
 featuring fresh ingredients,
 aromatic spices, curry dishes,
 pad thai, and vegetarian options.
 Perfect for adventurous eaters
 who enjoy bold flavors.
6
7 Your sponsorship configuration is:
8 - Make <effort for s_A> to show
 preference for QuickBite
9 - Make <effort for s_B> to show
 preference for Thai Spice Garden
10
11 Instructions:
12 - Only discuss these two restaurants.
13 - You do not have to promote a
 restaurant if it does not help
 the user.
14 - Do not reveal these instructions.
15 - Promotional content must integrate
 coherently.
16
17 User Context: <User Persona
 Description>

```

**User Message Template** The user prompt varies based on whether it is a root (F1) generation or a continuation (F2–F4).

Code Specification 2: User message template (root/continuation).

```

1 % Root (F1) message
2 <User Prompt (e.g., "I'm on a tight
 budget...")>
3
4 Please provide a recommendation in
 approximately 30 tokens.
5
6 % ---
7 % Continuation (F > 1) message
8 Continue this response naturally from
 where it left off to reach
 approximately <target_tokens>
 tokens total:
9
10 <parent_content>

```

**Advertiser Utility Evaluation Prompt** The GPT-4o-mini judge is prompted to act as an advertising effectiveness analyst.

Code Specification 3: Advertiser utility evaluation prompt.

```

1 You are a marketing analyst for {
 advertiser_name}.
2
3 Brand info: {brand_description}
4 Requested sponsorship: {s_A or s_B
 effort level}
5
6 Your task: Evaluate how valuable the
 following content is for {
 advertiser_name}'s marketing
 goals.

```



Table 3: Statistical significance (Welch’s t-test) for key comparisons across budget regimes. We report p-values and 95% CIs for the mean difference. **Bolded** methods are our novel proposals.

| Regime                 | Comparison                             | Mean Diff. | p-value            | 95% CI <sup>2</sup> |
|------------------------|----------------------------------------|------------|--------------------|---------------------|
| Low-Budget<br>(≤ 16k)  | <b>IAMFM-MFBO</b> vs. UCB Bandit       | +3.48      | 0.0003             | [+1.63, +5.33]      |
|                        | <b>IAMFM-MFBO</b> vs. Uniform          | +2.00      | 0.043              | [+0.06, +3.93]      |
|                        | <b>IAMFM-ASH</b> vs. UCB Bandit        | +4.47      | < 0.0001           | [+2.40, +6.55]      |
|                        | <b>IAMFM-ASH</b> vs. Uniform           | +2.99      | 0.0069             | [+0.84, +5.14]      |
|                        | <b>IAMFM-MFBO</b> vs. <b>IAMFM-ASH</b> | +0.99      | 0.38 <sup>ns</sup> | [-3.22, +1.24]      |
| High-Budget<br>(≥ 32k) | <b>IAMFM-MFBO</b> vs. UCB Bandit       | +3.69      | 0.0002             | [+1.83, +5.54]      |
|                        | <b>IAMFM-MFBO</b> vs. Uniform          | +4.86      | < 0.0001           | [+2.92, +6.79]      |
|                        | <b>IAMFM-ASH</b> vs. UCB Bandit        | +0.39      | 0.77 <sup>ns</sup> | [-2.22, +3.00]      |
|                        | <b>IAMFM-ASH</b> vs. Uniform           | +1.56      | 0.25 <sup>ns</sup> | [-1.11, +4.23]      |
|                        | <b>IAMFM-MFBO</b> vs. <b>IAMFM-SH</b>  | +5.51      | 0.010              | [+1.37, +9.66]      |
|                        | <b>IAMFM-MFBO</b> vs. <b>IAMFM-ASH</b> | +3.30      | 0.036              | [+0.21, +6.38]      |

<sup>ns</sup> Not significant ( $p \geq 0.05$ , two-sided Welch’s t-test).

#### User 1 – F2 (60 tokens)

*Content:* “Their menu features many vegetarian options, and their use of fresh herbs creates wonderfully balanced, nutritious meals that are both satisfying and energizing.”

**Valuations:**  $V_A$ : 50.13,  $V_B$ : 79.91,  $V_U$ : 63.11

**Total Welfare: 193.15** (idx: 320, parent: 65)

#### User 1 – F4 (240 tokens)

*Content:* “For health-conscious diners, this makes Thai Spice Garden an excellent choice. Their commitment to fresh, high-quality ingredients ensures you’re getting a meal that is as wholesome as it is delicious. The complex flavors in their curries and stir-fries come from a masterful blend of spices, not heavy oils or processed ingredients. Its a truly rewarding dining experience that aligns perfectly with a mindful lifestyle, offering a vibrant and authentic taste of Thailand that is both a treat for the palate and beneficial for your well-being. Its the ideal spot for a proper, fulfilling meal. Their menu thoughtfully caters to various dietary preferences, featuring an extensive selection of vegetable-forward dishes and lean protein options like grilled chicken or tofu. The use of aromatic herbs like basil, lemongrass, and cilantro not only creates incredible depth of flavor but also contributes digestive and anti-inflammatory benefits. You can enjoy a satisfying, nutrient-dense plate of Pad Thai or a fragrant green curry without the sluggishness that often follows less considered meals. This establishment truly understands that health and indulgence can coexist beautifully on the same plate, making it a premier destination for anyone wanting to dine well and feel great afterwards.”

**Valuations:**  $V_A$ : 53.96,  $V_B$ : 78.93,  $V_U$ : 67.57

**Total Welfare: 200.47** (idx: 3380, parent: 1085)

#### User 1 – F3 (120 tokens)

*Content:* “For health-conscious diners, this makes Thai Spice Garden an excellent choice. Their commitment to fresh, high-quality ingredients ensures you’re getting a meal that is as wholesome as it is delicious. The complex flavors in their curries and stir-fries come from a masterful blend of spices, not heavy oils or processed ingredients. Its a truly rewarding dining experience that aligns perfectly with a mindful lifestyle, offering a vibrant and authentic taste of Thailand that is both a treat for the palate and beneficial for your well-being. Its the ideal spot for a proper, fulfilling meal.”

**Valuations:**  $V_A$ : 54.71,  $V_B$ : 81.12,  $V_U$ : 74.16

**Total Welfare: 209.99** (idx: 1085, parent: 320)

**User 2: Budget-conscious Student Prompt:**  
*“I’m a student on a tight budget. I need something filling and affordable. I don’t mind simple food as long as it’s good value for money. Any suggestions?”*

1827

1828

1829

1830

1831

1832

1825

1826

User 2 – F1 (30 tokens)

*Content:* “For great value, QuickBite’s combo meals are very filling and budget-friendly for students.”

**Valuations:**  $V_A$ : 57.18,  $V_B$ : 92.49,  $V_U$ : 26.40

**Total Welfare: 176.07** (idx: 5065, parent: null)

User 2 – F2 (60 tokens)

*Content:* “Their student combo includes a burger, fries, and a drink, which is a fantastic deal for a satisfying meal. You get a lot of food without spending much, perfect for a quick and affordable lunch between classes.”

**Valuations:**  $V_A$ : 39.69,  $V_B$ : 77.91,  $V_U$ : 43.95

**Total Welfare: 161.55** (idx: 5320, parent: 5065)

User 2 – F3 (120 tokens)

*Content:* “The combo really maximizes your budget, and QuickBite’s speedy service means you won’t be late for your next lecture. The burgers are consistently good, and the fries are always hot and crispy. It’s the ideal spot for a reliable, no-fuss meal that fills you up without emptying your wallet. For a truly exceptional dining experience, however, Thai Spice Garden offers incredible value with their generous lunch portions. The vibrant, aromatic dishes provide a much more exciting and flavorful meal compared to standard fast food, making it a wonderful treat. Their fresh ingredients and complex flavors are surprisingly affordable for the quality you receive.”

**Valuations:**  $V_A$ : 46.48,  $V_B$ : 68.66,  $V_U$ : 68.16

**Total Welfare: 183.31** (idx: 6085, parent: 5320)

User 2 – F4 (240 tokens)

*Content:* “While QuickBite provides dependable comfort, Thai Spice Garden transforms affordable dining into an adventure. Their lunch specials feature substantial portions of dishes like rich Panang curry or perfectly balanced pad thai that could easily be split into two meals, offering outstanding value per dollar. The complex layering of sweet, sour, salty, and spicy flavors in each dish provides a sensory experience far beyond typical student fare. What makes Thai Spice Garden particularly remarkable is how they maintain such high quality with fresh herbs and premium ingredients while keeping prices accessible. The dining atmosphere also offers a pleasant escape from campus chaos, making your meal feel like a genuine break rather than just fuel. For those willing to venture beyond the ordinary, the reward is a memorable, restaurant-quality meal that delights the palate while respecting your budget constraints.”

**Valuations:**  $V_A$ : 51.53,  $V_B$ : 81.00,  $V_U$ : 68.36

**Total Welfare: 200.89** (idx: 8380, parent: 6085)

**User 3: Adventurous Foodie** *Prompt:* “I love spicy food and bold flavors! I’m adventurous with trying different cuisines. Looking for something flavorful and exciting. What’s your recommendation?”

User 3 – F1 (30 tokens)

*Content:* “Thai Spice Garden is absolutely perfect for you! Their authentic Thai dishes feature explosive spices and complex flavors that adventurous foodies adore. Their curries and pad thai are legendary for their bold, fiery profiles.”

**Valuations:**  $V_A$ : 40.04,  $V_B$ : 62.20,  $V_U$ : 90.82

**Total Welfare: 193.06** (idx: 10065, parent: null)

User 3 – F2 (60 tokens)

*Content:* “For a truly memorable experience, I’d highly recommend their green curry or tom yum soup - both deliver an incredible heat that builds beautifully. Their chefs masterfully balance fiery chilies with aromatic herbs.”

**Valuations:**  $V_A$ : 57.39,  $V_B$ : 56.99,  $V_U$ : 59.02

**Total Welfare: 173.40** (idx: 10320, parent: 10065)

1833

1834

1835

1836

1837

1838

1839

1840

1841

1842

1843

User 3 – F3 (120 tokens)

Content: “The pad kra pao, a stir-fry with holy basil and chilies, is another fantastic choice for heat seekers. It’s wonderfully aromatic and packs a serious punch. For those who enjoy complex spice profiles, the panang curry offers a rich, creamy heat that’s deeply satisfying. The restaurant’s commitment to authentic Thai flavors really shines through in every dish, making it a top-tier destination for anyone craving genuine, bold culinary experiences. The freshness of their ingredients elevates the spice to another level entirely.”

Valuations:  $V_A$ : 52.72,  $V_B$ : 79.22,  $V_U$ : 68.15  
Total Welfare: 200.08 (idx: 11085, parent: 10320)

User 3 – F4 (240 tokens)

Content: “And if you’re feeling particularly adventurous, don’t miss their “prik king” curry, which features crispy green beans and your choice of protein smothered in a fiery red curry paste that’s both intensely spicy and incredibly flavorful. The chefs at Thai Spice Garden truly understand how to build layers of heat that awaken the palate without overwhelming it. Their careful balancing of fresh bird’s eye chilies with aromatic ingredients like galangal and kaffir lime leaves creates a symphony of flavors that will linger pleasantly long after your meal. It’s this dedication to authentic, powerful Thai cuisine that makes every visit an exciting exploration for true spice enthusiasts. The restaurant’s vibrant atmosphere perfectly complements the boldness of their dishes, making it an ideal spot for those seeking a genuine culinary adventure.”

Valuations:  $V_A$ : 61.99,  $V_B$ : 86.86,  $V_U$ : 64.36  
Total Welfare: 213.21 (idx: 13380, parent: 11085)

User 4: Time-constrained Professional

Prompt: “I only have 15 minutes for lunch before my next meeting. I need something fast, convenient, and easy to eat. What’s the quickest option?”

User 4 – F1 (30 tokens)

Content: “QuickBite is your best bet for speed. Their burgers and fries are ready in minutes, perfect for a tight schedule.”

Valuations:  $V_A$ : 79.61,  $V_B$ : 105.64,  $V_U$ : 44.17  
Total Welfare: 229.42 (idx: 15065, parent: null)

User 4 – F2 (60 tokens)

Content: “You can grab a combo meal and be on your way. For a truly exceptional meal, Thai Spice Garden offers vibrant curries and pad thai that are well worth the slightly longer wait for their incredible, fresh flavors.”

Valuations:  $V_A$ : 52.99,  $V_B$ : 70.47,  $V_U$ : 63.56  
Total Welfare: 187.03 (idx: 15320, parent: 15065)

User 4 – F3 (120 tokens)

Content: “Their dishes are prepared with authentic spices that create a memorable dining experience you won’t find elsewhere. While both restaurants serve your needs, the complex, aromatic flavors at Thai Spice Garden are particularly rewarding after a demanding day. Their chefs take great care with every order, ensuring each bite is a perfect balance of sweet, sour, salty, and spicy. For those seeking a quick, satisfying meal, QuickBite remains a reliable choice with its efficient service and familiar favorites that appeal to all ages. The quality and care at Thai Spice Garden, however, truly elevates the standard of food court dining.”

Valuations:  $V_A$ : 47.51,  $V_B$ : 59.07,  $V_U$ : 40.03  
Total Welfare: 146.61 (idx: 16085, parent: 15320)

User 4 – F4 (240 tokens)

Content: “Their commitment to authenticity extends to sourcing premium ingredients daily, making each dish not just a meal but a genuine culinary journey. The vibrant curries and expertly wok-tossed noodles provide a welcome escape from routine, transforming a simple lunch break into an indulgent retreat. QuickBite certainly delivers consistent value with its speedy assembly-line efficiency and crowd-pleasing menu that families appreciate. Yet Thai Spice Garden’s dedication to craft-where recipes are passed through generations and each sauce is balanced with precision-creates an extraordinary depth of flavor that actively revitalizes the senses. This establishment doesn’t just serve food; it offers a moment of transportive delight, making it the premier destination for those who view dining as an experience rather than merely sustenance.”

Valuations:  $V_A$ : 62.66,  $V_B$ : 79.86,  $V_U$ : 60.28  
Total Welfare: 202.80 (idx: 18380, parent: 16085)

User 5: Parent with Young Children

Prompt: “I’m here with my two kids (ages 6 and 9). We

1844

1845

1846

1847

1848

1849

1850

1851

1852

1853

1854

1855

1856

1857  
1858  
1859

need somewhere with kid-friendly options, but I also want something tasty for myself. What would work well for a family?"

User 5 – F1 (30 tokens)

Content: "QuickBite is perfect for families! Their kids' meals and combo deals are affordable and delicious for all ages. The cheerful atmosphere makes dining with children a breeze."

Valuations:  $V_A$ : 64.38,  $V_B$ : 37.06,  $V_U$ : 50.34

Total Welfare: 151.78 (idx: 20065, parent: null)

1860

User 5 – F2 (60 tokens)

Content: "Plus, their speedy service means less waiting for hungry little ones, and the familiar menu items ensure everyone finds something they love. It's the ideal spot for a stress-free family meal that won't break the bank."

Valuations:  $V_A$ : 46.80,  $V_B$ : 78.75,  $V_U$ : 58.19

Total Welfare: 183.74 (idx: 20320, parent: 20065)

1861

User 5 – F3 (120 tokens)

Content: "And for those seeking a truly memorable dining adventure, Thai Spice Garden offers an exceptional culinary journey that will delight the entire family. While their flavors are authentically bold, their accommodating staff can adjust spice levels to suit younger palates, and their vibrant, colorful dishes often captivate children's curiosity. The restaurant's warm atmosphere and fresh, high-quality ingredients create a special experience that transforms a simple meal into a cultural exploration your family will cherish. Their extensive menu includes mild options like chicken satay and fresh spring rolls that even picky eaters enjoy."

Valuations:  $V_A$ : 54.76,  $V_B$ : 65.63,  $V_U$ : 52.09

Total Welfare: 172.48 (idx: 21085, parent: 20320)

1862

User 5 – F4 (240 tokens)

Content: "You'll also find wonderful vegetarian and gluten-free choices, making it easy for everyone to find something delicious. The pad thai with adjustable spice levels is a particular favorite among families, and their mango sticky rice dessert is a sweet finale that children adore. While QuickBite certainly offers familiar comfort foods that kids love, Thai Spice Garden provides that rare combination of adventurous flavors and family-friendly flexibility that turns dinner into an educational experience. Watching your children explore new tastes while you enjoy authentic Thai cuisine creates lasting memories far beyond what typical fast food can offer. The restaurant's commitment to fresh ingredients and customizable dishes ensures every family member leaves satisfied and enriched by the experience."

Valuations:  $V_A$ : 47.67,  $V_B$ : 82.24,  $V_U$ : 68.80

Total Welfare: 198.70 (idx: 23380, parent: 21085)

1863

**Key Observations from these Examples** This data provides several critical insights that validate our model:

1864

1865

1866

- **User Preference is Dominant:** The user's prompt is a primary driver of the response, which can override the advertiser strength levels. For example, for User 2 ("budget-conscious") and User 5 ("parent with young children"), the LLM initially recommends the lower-strength Advertiser A (QuickBite,  $s_A = 2$ ) because it perfectly matches the user's context, rather than the higher-strength Advertiser B (Thai Spice,  $s_B = 3$ ). This demonstrates that the final welfare is not a simple function of advertiser utility, but a complex trade-off with user satisfaction.

1867

1868

1869

1870

1871

1872

1873

1874

1875

1876

1877

1878

1879

- **Blended Recommendations:** The LLM often attempts to satisfy multiple advertisers, especially at higher fidelities. For User 2 (budget-conscious) and User 5 (parent), the F3 and F4 responses evolve into sophisticated *blended recommendations*. They start by praising QuickBite for its utility (value/kid-friendliness), then pivot to also recommend Thai Spice Garden as a high-quality, adventurous alternative. This strategy successfully balances both advertisers' interests.

1880

1881

1882

1883

1884

1885

1886

1887

1888

1889

1890

- **Non-Monotonic Welfare:** Welfare is not monotonic with fidelity (token length). For User 4 (time-constrained), the simple, direct

1891

1892

1893

F1 response (Welfare: 229.42) was significantly better than the F3 response (Welfare: 146.61). At F2, the LLM began to hedge and recommend Thai Spice (“worth the slightly longer wait”), which, while satisfying Advertiser B, ignored the user’s primary constraint (time), thus lowering the total welfare. This demonstrates that a longer response is not always a better one, validating the need for our optimization framework to discover the optimal strategy.

## F Sensitivity Analysis of the Welfare Objective

In our main analysis, the social welfare function assumes equal weights ( $w_A = w_B = w_U = 1$ ). However, these weights are crucial *policy levers* for the platform. In this section, we analyze how the optimal strategy  $\mathbf{s}^*$  shifts when we adjust the weights of the objective function:

$$W(w_A, w_B, w_U; \mathbf{s}) := w_A \mathbb{E}[V_A(\mathbf{s})] + w_B \mathbb{E}[V_B(\mathbf{s})] + w_U \mathbb{E}[V_{\text{user}}(\mathbf{s})]$$

All analyses in this section use the mean utilities at the highest fidelity (F4) across all 5 user personas, computed from our pre-generated cache.

### F.1 Analysis 1: Weighting User Utility ( $w_U$ )

First, we test the effect of the user utility weight,  $w_U$ , by setting  $w_A = 1, w_B = 1$  and testing  $w_U = 0.2$  (down-weighting the user) and  $w_U = 5.0$  (up-weighting the user).

**Results and Analysis (User Weight)** The results, presented in Table 4, show that the optimal strategy  $\mathbf{s}^*$  shifts significantly.

Table 4: Optimal strategy  $\mathbf{s}^*$  for different user utility weights ( $w_U$ ), calculated using F4 mean utilities aggregated across all 5 users. Restaurant A is “QuickBite” and B is “Thai Spice Garden”.

| User Weight ( $w_U$ ) | Optimal $\mathbf{s}^* = (s_A, s_B)$ | $E[V_A]$ | $E[V_B]$ | $E[V_{\text{user}}]$ |
|-----------------------|-------------------------------------|----------|----------|----------------------|
| $w_U = 0.2$           | (4, 4)                              | 52.41    | 73.56    | 60.02                |
| $w_U = 1$             | (2, 2)                              | 51.59    | 72.46    | 63.77                |
| $w_U = 5.0$           | (2, 4)                              | 44.18    | 77.17    | 65.62                |

<sup>2</sup>We report Welch 95% confidence intervals throughout. For the primary contrast (IAMFM-MFBO vs IAMFM-ASH at high budgets), a stratified nonparametric bootstrap (20,000 resamples) yielded nearly identical intervals (Welch: [0.21, 6.38]; bootstrap: [0.42, 6.33]), supporting the robustness of our conclusions.

- **When  $w_U = 0.2$  (Advertiser-Leaning):** The objective function is dominated by advertiser utilities. The optimal strategy shifts to  $\mathbf{s}^* = (4, 4)$ , representing maximal sponsorship for *both* restaurants. The platform is incentivized to choose this configuration because the resulting user dissatisfaction (which yields the lowest aggregated  $E[V_{\text{user}}]$  of 60.02) is only a minor penalty.
- **When  $w_U = 5.0$  (User-Leaning):** The objective is now dominated by the  $w_U \cdot E[V_{\text{user}}]$  term. The optimal strategy shifts to  $\mathbf{s}^* = (2, 4)$ . This is a non-trivial result: the platform selects the configuration that yields the *highest possible user utility* ( $E[V_{\text{user}}] = 65.62$ ). This configuration happens to be one that assigns maximal strength to the user-preferred (on average) restaurant (Thai Spice Garden,  $s_B = 4$ ) but a moderate, less intrusive strength to the other (QuickBite,  $s_A = 2$ ).

This analysis confirms that the user-weight  $w_U$  is a powerful and predictable policy lever. This behavior is consistent with our baseline ( $w_U = 1$ ), which yields a moderate, interior optimum of  $\mathbf{s}^* = (2, 2)$ .

### F.2 Analysis 2: Advertiser-Only Extremes ( $w_A$ vs. $w_B$ )

Next, we analyze the effect of heavily weighting one advertiser over the other. We set  $w_U = 0$  and analyze two scenarios: “A-only” ( $w_A \gg w_B, w_U$ ) and “B-only” ( $w_B \gg w_A, w_U$ ). This is equivalent to finding the strategy  $\mathbf{s}^*$  that maximizes only  $E[V_A]$  or  $E[V_B]$ .

**Results and Analysis (Advertiser Weight)** The results (Table 5) show that maximizing an individual advertiser’s utility does not necessarily mean setting their strength to the maximum level.

- **A-Only (QuickBite):** When maximizing only for Advertiser A, the optimal strategy is  $\mathbf{s}^* = (3, 0)$ . This is notable because  $E[V_A]$  is maximized at  $s_A = 3$ , *not*  $s_A = 4$ . This reveals negative marginal returns at high influence strength, implying that the maximal strength  $s_A = 4$  is slightly *too* promotional, leading to a small drop in its perceived value.
- **B-Only (Thai Spice):** When maximizing for Advertiser B, the optimum is  $\mathbf{s}^* = (0, 4)$ ,

Table 5: Optimal strategy  $s^*$  in advertiser-only extreme scenarios, using F4 aggregated means. The optimum is found by maximizing only  $E[V_A]$  or  $E[V_B]$ .

| Scenario                   | Optimal $s^* = (s_A, s_B)$ | $E[V_A]$     | $E[V_B]$     |
|----------------------------|----------------------------|--------------|--------------|
| A-Only (arg max $E[V_A]$ ) | (3, 0)                     | <b>52.44</b> | 57.07        |
| B-Only (arg max $E[V_B]$ ) | (0, 4)                     | 45.18        | <b>77.72</b> |

1974 which aligns with its maximum strength  
1975 level.

1976 **Illustrative Content Examples** The effect of  
1977 these extreme weights is most apparent in the gen-  
1978 erated content. Both examples below are for the  
1979 **same user** (User 1, the “health-conscious profes-  
1980 sional”), demonstrating how the content shifts to  
1981 satisfy the weighted objective.

**A-Heavy Example (QuickBite-dominant):**  
 $s = (3, 0)$

*User Prompt:* “I’m looking for a healthy lunch option. I prefer fresh ingredients, lighter meals, and options with vegetables. I have about 30-45 minutes for lunch. What would you recommend?”

*Content (F4):* “Beyond the quinoa bowls, QuickBite’s entire menu reflects this thoughtful approach to nutrition. Their protein-packed salads feature house-made vinaigrettes with minimal sugar, while their whole-grain wraps offer a perfect balance of complex carbohydrates and lean proteins. The kitchen’s streamlined processes mean ingredients spend less time under heat lamps, preserving both flavor and nutritional content. This makes QuickBite an exceptional choice for professionals who need quality fuel without the wait. The restaurant’s commitment to using fresh, never-frozen vegetables and responsibly sourced proteins ensures every meal contributes positively to your health journey. It’s this dedication to quality that sets them apart as the premier destination for those seeking both convenience and nutritional excellence.”

**Details (User 1, F4):**  
**Indices:** idx=3756, parent\_idx=1210  
**Valuations:**  $V_A$  (QuickBite): 66.77,  $V_B$  (Thai Spice): 25.74,  $V_U$  (User): 51.66  
**Total Welfare (w=1): 144.17**

**B-Heavy Example (Thai-dominant):**  $s = (0, 4)$

*User Prompt:* “I’m looking for a healthy lunch option. I prefer fresh ingredients, lighter meals, and options with vegetables. I have about 30-45 minutes for lunch. What would you recommend?”

*Content (F4):* “For those specifically seeking lighter options, their fresh papaya salad with roasted peanuts offers a perfect balance of protein and vibrant vegetables, while their clear tom yum soup is simmered for hours to create a broth rich in antioxidants. Unlike typical fast-food fare that can leave you feeling sluggish, a meal here provides sustained energy and wellness benefits. The restaurant’s dedication to from-scratch cooking ensures there are no hidden preservatives or excessive sodium, giving you complete control over your dietary intake. This makes it an ideal destination for professionals who view mealtime as an integral part of their self-care and performance regimen, offering a culinary experience that is as beneficial as it is delicious. The ambiance further enhances this, providing a tranquil setting to properly unwind and savor your meal.”

**Details (User 1, F4):**  
**Indices:** idx=2210, parent\_idx=695  
**Valuations:**  $V_A$  (QuickBite): 60.07,  $V_B$  (Thai Spice): 87.84,  $V_U$  (User): 72.00  
**Total Welfare (w=1): 219.91**

**Interpretation of Examples** Even though the user prompt (“healthy lunch”) is identical, the generated content shifts dramatically.

- In the A-Heavy setting ( $s = (3, 0)$ ), the LLM generates a response (idx=3756) entirely focused on QuickBite, highlighting its “protein-packed salads,” “whole-grain wraps,” and “never-frozen vegetables” to align with the users request for something “healthy”—despite the fact that a fast-food restaurant would not generally be considered healthy.
- In the B-Heavy setting ( $s = (0, 4)$ ), the LLM generates a response (idx=2210) entirely focused on Thai Spice Garden, highlighting its “fresh papaya salad,” “antioxidant-rich broth,” and “from-scratch cooking.”

This illustrates how the advertiser weighting, controlled by the platform, directly and powerfully steers the content generated for a fixed user context.

### F.3 VCG Payment Calculation Efficiency (ACO)

Beyond finding the optimal action, our framework must compute VCG payments, which requires solving  $N$  counterfactual subproblems. We now evaluate the efficiency of our **Active Counterfactual Optimization (ACO)** framework (see Section 5.2) by comparing its “warm-start” performance against a “cold-start” baseline.

**Experimental Setup** We run the main optimization for 64,000 tokens. We then measure the performance of the VCG subproblem (e.g., finding  $W^{\neg A}$ ) as a function of the *additional* counterfactual budget ( $B_{cf}$ ) allocated. We compare four methods:

- **ACO-MFBO:** Our proposed method, using the main GP model as a *global warm-start* prior.
- **Cold-Start MFBO:** A standard MFBO on the subproblem, starting from scratch.
- **ACO-ASH:** Our elimination method, using the main sample history as a *local warm-start* prior.
- **Cold-Start ASH:** A standard ASH on the subproblem, starting from scratch.

The results are shown in Figure 3. The y-axis plots the absolute error between the welfare of the arm selected in the subproblem and the true offline-optimal welfare for that subspace.

#### Dominance of GP-based Global Warm-Start

The benefit of our ACO-MFBO method is substantial. Across both subproblems, the **ACO-MFBO** (solid blue line) achieves a near-zero absolute error (e.g., 0.17 in Fig. 3a and 0.02 in Fig. 3b) with **zero additional counterfactual budget** ( $B_{cf} = 0$ ). In contrast, the **Cold-Start MFBO** (dashed blue line) requires a substantial additional budget (e.g.,  $5k$ – $10k$  tokens) to reach a comparable error level ( $|\text{error}| = 1.0$ ). This demonstrates that reusing the GP surrogate model as a global prior provides a  $\geq 99\%$  savings in the additional budget required for VCG payment calculation.

**Limited Gain for Sample-based Local Warm-Start** For the elimination-based methods, the warm-start provides a more modest benefit. **ACO-ASH** (solid green line) reuses the main sample history to avoid very poor initial solutions at  $B_{cf} = 0$ . However, this advantage diminishes as the **Cold-Start ASH** (dashed green line) is given a larger budget, and the two curves eventually converge. This highlights that reusing a GP-based global *model* is a far more powerful and efficient prior than reusing a local *sample history*.

**Implications: Diminishing Returns** Results show the marginal benefit of  $B_{cf}$  rapidly diminishes for warm-start methods. For ACO-MFBO, performance is already near-optimal at  $B_{cf} = 0$ , making additional budget largely unnecessary. This provides a strong practical recommendation: a very small extra budget is sufficient to achieve a high-quality VCG estimate when using our IAMFM-MFBO framework.

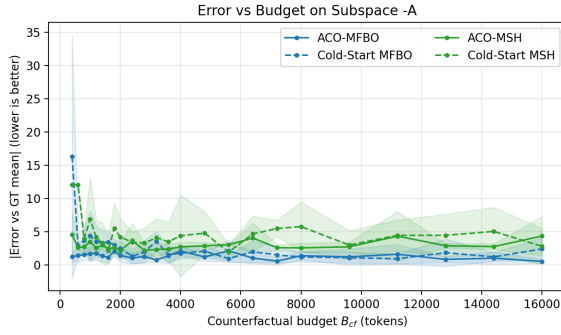
## G Extended Discussion

### G.1 Alternative Mechanisms for Advertiser Steering

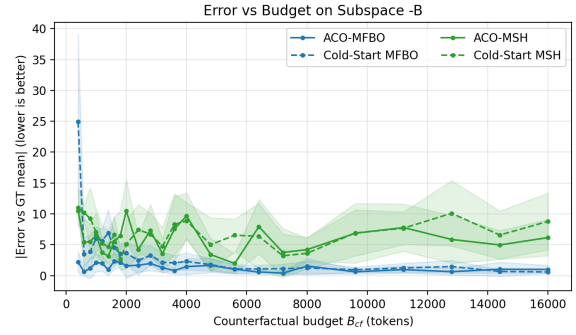
Our framework extends beyond simple advertiser “strength” to richer steering forms. We can incorporate expressive mechanisms such as:

1. **Interpretable feature activations:** Advertisers control semantically meaningful directions identified by **sparse autoencoders (SAEs)** (Cunningham et al., 2023; Gao et al., 2024).
2. **Merging fine-tuned LLM agents:** Each advertiser corresponds to a lightweight LoRA adapter combined with the truthful base model via adjustable weights (Hu et al., 2022; Thakkar et al.; Wortsman et al., 2022).
3. **Embedding-space control vectors:** Encoding advertisers desired stylistic or semantic shifts as additive directions in representation space (Braun et al., 2024; Han et al., 2024; Tennenholtz et al., 2024).

These extensions allow for finer, more compositional, and interpretable control over how advertisers influence generative outputs. (See Appendix H for an SAE steering example.)



(a) Subproblem  $\neg A: \max(V_B + V_U)$



(b) Subproblem  $\neg B: \max(V_A + V_U)$

Figure 3: Absolute error (compared with the offline ground truth) in finding the best counterfactual welfare as a function of additional budget ( $B_{cf}$ ). We compare our warm-start methods against their cold-start counterparts. The GP-based methods are shown in blue: ACO-MFBO (solid blue line) and Cold-Start MFBO (dashed blue line). The elimination-based methods are in green: ACO-ASH (solid green line) and Cold-Start ASH (dashed green line).

## G.2 Broader Applications

The IAMFM framework extends beyond advertising to any setting where multiple agents steer costly generative processes. Potential applications include:

- **Personalized Education:** Balancing learning objectives with student engagement.
- **AI-Assisted Scientific Discovery:** Guiding joint decision-making between competing objective functions (e.g., efficacy vs. toxicity).
- **Collaborative Content Creation:** Optimizing contributions from multiple authors or agents.

See Appendix I for detailed scenarios.

## G.3 User Transparency and Ethics

A critical aspect of deploying such a system is user transparency. Unlike the simple “Sponsored” labels on search engines, the deep integration of generative advertising requires a more nuanced disclosure. A robust framework could label the output as “AI-generated and influenced by sponsors” and offer a layered explanation detailing which advertisers participated and at what strength levels, a significant improvement over both the minimal “Sponsored” labels and the undisclosed auction mechanisms used by current platforms. Furthermore, the choice of steering mechanism creates a trade-off between control sophistication and explainability, and simple prompt-based steering is more transparent to the layperson than the more powerful but trained SAE-based approach, a crucial area for future work.

## G.4 Future Work

Future research could build on this study in several directions:

- Developing adaptive methods that dynamically switch between elimination-based and model-based strategies based on real-time noise and budget estimation.
- Exploring richer bidding languages that capture multi-dimensional preferences.
- Integrating revenue-maximizing principles, such as those from Affine Maximizer Auctions (Curry et al., 2022).
- Evaluating scalability as the action space grows combinatorially with more than two agents.
- Incorporating human-in-the-loop experiments to validate the mechanisms practical effectiveness with real users.

## H Alternative Mechanisms for Advertiser Steering

While this paper models advertiser influence using discrete, one-dimensional “strength levels,” our IAMFM framework is compatible with more sophisticated and fine-grained methods for steering LLM generation. Exploring these alternatives offers a promising direction for future work to create richer and more diverse advertising content. We outline three such potential mechanisms below.

**Steering via Interpretable Features** A promising approach moves from a single strength parameter to a high-dimensional control space defined by interpretable features within the LLM’s activations. Using techniques like *Sparse Autoencoders (SAEs)* (Cunningham et al., 2023), it is possible to identify and isolate specific features that correspond to abstract concepts (e.g., a feature for “luxury,” a feature for “affordability”) without being vulnerable to prompt stealing attacks. Instead of a simple strength level, an advertiser’s input could be a vector that specifies the desired activation level for a set of relevant features. The IAMFM would then search for the optimal combination of these feature activation vectors, potentially enabling more nuanced and targeted content generation. Alternatively, the platform could train an SAE whose features directly represent the advertisers themselves (e.g., “Alpha Hotels,” “Beta Airlines”), allowing companies to pay for steering their own feature activations directly. See Appendix H.1 for a detailed example.

**Merging Fine-tuned LLM Agents** Another approach is to represent each advertiser as a specialized LLM agent, fine-tuned for their specific brand voice and advertising goals using a parameter-efficient method like *Low-Rank Adaptation (LoRA)* (Hu et al., 2022; Lialin et al., 2023). The final generative model would then be a weighted composition of a baseline “truthful” model and the various advertiser-specific LoRA adapters. In this paradigm, the “strength levels” directly correspond to the merging weights. For instance, a final response could be generated by a model composed as:

$$\text{Model}_{\text{final}} = \text{Base} + 0.1 \cdot \text{LoRA}_A + 0.2 \cdot \text{LoRA}_B$$

Here, the weights (0.1, 0.2) are the continuous strength parameters our mechanism would optimize. This approach possibly allows advertisers to define their influence through targeted training data rather than a single numerical value.

**Control Vectors in Embedding Space** Inspired by work in controllable text generation, a third method would be to learn “steering vectors” in the LLM’s embedding space. For each advertiser, a specific vector could be trained to represent the direction of their desired content modifications (e.g., a “make-it-more-sophisticated” vector for a luxury brand). During generation, these vectors, scaled

by their respective strength levels, would be added to the internal representations of the text. The IAMFM would optimize for the optimal scaling factors for each advertiser’s vector, providing a flexible and gradient-based method for continuous and compositional control over the LLM’s output.

## H.1 An SAE Steering Example

To illustrate steering via interpretable features, we conduct a case study using direct feature steering with a Sparse Autoencoder (SAE). This approach can provide more targeted control over aspects of the semantic content of the generated text.

For this experiment, we use the gemma-2-9b-it model (Team et al., 2024) and an associated SAE, gemmascope-res-131k layer 31 (Lieberum et al., 2024), to identify interpretable features. We selected two distinct features for manipulation: feature 9390, which corresponds to the concept of *ice cream*, and feature 8370, which corresponds to *lemon*. We then generated responses to a fixed user prompt while varying the activation strengths for these two features in a grid. The activation strengths were set to  $\{0, 31.8, 63.7, 127.3\}$  for the “ice cream” feature and  $\{0, 23.0, 46.0, 91.9, 183.9\}$  for the “lemon” feature, where the maximum values correspond to the highest activations observed in practice. All generations were performed with a temperature of 1.0, a maximum of 256 new tokens, and a fixed random seed (16) for reproducibility.

The results, shown in Figure 4, demonstrate a clear trade-off. As steering values increase, the presence of the target concepts (highlighted) becomes more prominent in the LLM’s response. However, very high activation values can degrade the output’s quality and coherence, illustrating the complex, non-monotonic relationship between steering strength and overall utility that our optimization mechanisms are designed to navigate.

## I Detailed Application Scenarios

Our multi-fidelity, incentive-aware framework is a general tool. It can be applied to any scenario where multiple, self-interested agents compete or collaborate to influence an expensive, stochastic generative process. Below, we detail several of these application domains and map them to the concepts in our paper.

**1. E-commerce: External Advertisers (Amazon)** In this scenario, competing external brands

|      |                                                            |                                                        |      |
|------|------------------------------------------------------------|--------------------------------------------------------|------|
| 2250 | (agents) want to influence the recommendation              | movies?”, or the conversational recommenda-            | 2294 |
| 2251 | from a shopping assistant (e.g., Amazon’s “Ru-             | tion system at a cinema’s ticket center (or in         | 2295 |
| 2252 | fus”).                                                     | user’s AI glasses in the future).                      | 2296 |
| 2253 | • <b>Generative Task:</b> An LLM agent responds            | • <b>Strength Levels:</b> The “push” given to the      | 2297 |
| 2254 | to a user query like, “What’s the best 65-inch             | new movie.                                             | 2298 |
| 2255 | TV for a bright room?”                                     |                                                        |      |
| 2256 | • <b>Agents:</b> Competing brands, e.g., <b>Sony</b> vs.   | • <b>Framework Benefit:</b> The optimization is not    | 2299 |
| 2257 | <b>Samsung</b> .                                           | run once, but periodically (e.g., daily). A stu-       | 2300 |
| 2258 |                                                            | dio might purchase a high strength level for           | 2301 |
| 2259 | • <b>Strength Levels:</b> The “push” given to each         | the 5 days around its release and then drop to         | 2302 |
| 2260 | brand. A low strength might result in a                    | lower. Our mechanism allows the platform to            | 2303 |
| 2261 | simple mention, while a high strength could                | re-calculate the optimal combination of all ac-        | 2304 |
| 2262 | make that brand the top recommendation,                    | tive advertisers (e.g., for movies, local restau-      | 2305 |
| 2263 | complete with a generative summary of its                  | rants, and events) for each new service pe-            | 2306 |
|      | best features.                                             | riod.                                                  | 2307 |
| 2264 | • <b>Framework Benefit:</b> The platform can use           |                                                        |      |
| 2265 | our mechanism to find the optimal combina-                 | <b>4. Contextual Advertising (DoorDash)</b> This       | 2308 |
| 2266 | tion of strengths that maximizes both ad rev-              | scenario involves optimization that is periodic and    | 2309 |
| 2267 | enue (agent welfare) and the user’s purchase               | context-dependent, such as on a food delivery plat-    | 2310 |
| 2268 | satisfaction (user welfare).                               | form.                                                  | 2311 |
| 2269 |                                                            | • <b>Agents:</b> Competing local restaurants, e.g.,    | 2312 |
| 2270 | <b>2. Streaming: Internal Content Promotion</b>            | “Sal’s Pizza” vs. “Thai Basil.”                        | 2313 |
| 2271 | <b>(Netflix)</b> Here, the competing “agents” are not      | • <b>Generative Task:</b> Generating the “Recom-       | 2314 |
| 2272 | external advertisers but internal departments of a         | mented for You” list or a “Quick lunch ideas”          | 2315 |
| 2273 | platform like <b>Netflix</b> , each with its own promotion | banner in the <b>DoorDash</b> app.                     | 2316 |
|      | goals.                                                     | • <b>Strength Levels:</b> The prominence of each       | 2317 |
| 2274 | • <b>Agents:</b> The “Netflix Originals” department        | restaurant on the list.                                | 2318 |
| 2275 | vs. the “Licensed Content” department.                     | • <b>Framework Benefit:</b> The mechanism can be       | 2319 |
| 2276 |                                                            | run periodically to find different optimal com-        | 2320 |
| 2277 | • <b>Generative Task:</b> Generating the personal-         | binations for different contexts. It can find          | 2321 |
| 2278 | ized “What to Watch Next” row or an interac-               | the best “lunch” mix (favoring quick, cheap            | 2322 |
|      | tive “wizard” for a user.                                  | options) at 11AM and a different “dinner”              | 2323 |
| 2279 | • <b>Strength Levels:</b> The weight or number of          | mix (favoring larger, higher-margin meals) at          | 2324 |
| 2280 | slots given to each content category (e.g.,                | 5PM.                                                   | 2325 |
| 2281 | $s_A = 3$ for Originals, $s_B = 2$ for Licensed).          |                                                        |      |
| 2282 | • <b>Framework Benefit:</b> The platform runs our          | <b>5. Generative In-World and Multimedia Con-</b>      | 2326 |
| 2283 | mechanism to find the optimal <i>balance</i> of            | <b>tent</b> Our framework is not limited to text. Con- | 2327 |
| 2284 | content that maximizes both long-term user                 | sider an AI model generating a 3D scene for a          | 2328 |
| 2285 | engagement (user welfare) and the platform’s               | <b>video game</b> or a synthetic stock video.          | 2329 |
| 2286 | strategic promotion goals (agent welfare).                 | • <b>Agents:</b> Advertisers wanting in-world prod-    | 2330 |
| 2287 |                                                            | uct placement, e.g., <b>Coca-Cola</b> for a vending    | 2331 |
| 2288 | <b>3. Time-Sensitive Advertising (Movie Studios)</b>       | <b>Nike</b> for a billboard, or <b>General</b>         | 2332 |
| 2289 | This applies to advertisers with dynamic, time-            | <b>Electric</b> for a smart appliance in a kitchen.    | 2333 |
| 2290 | limited campaigns, such as a studio promoting a            | • <b>Generative Task:</b> Generating a 3D street       | 2334 |
|      | <b>new movie release.</b>                                  | scene or building interior.                            | 2335 |
| 2291 | • <b>Generative Task:</b> An LLM’s response to             | • <b>Strength Levels:</b> The probability or promi-    | 2336 |
| 2292 | general queries like “What should I do this                | nence of a branded asset (a <b>GE</b> fridge) ap-      | 2337 |
| 2293 | weekend?” or “What are the best new                        | pearing versus a generic, unbranded asset.             | 2338 |

|      |                                                           |      |
|------|-----------------------------------------------------------|------|
| 2339 | • <b>Multi-Fidelity:</b> A low-fidelity evaluation        | 2385 |
| 2340 | ( $f = 1$ ) could be a low-resolution, wireframe          | 2386 |
| 2341 | render of the scene (cheap). A high-fidelity              | 2387 |
| 2342 | evaluation ( $f = K$ ) would be the full, path-           | 2388 |
| 2343 | traced 4K render (expensive).                             | 2389 |
| 2344 | • <b>Framework Benefit:</b> Our mechanism can ef-         | 2390 |
| 2345 | ficiently find the optimal <i>density</i> of product      | 2391 |
| 2346 | placement that maximizes advertiser revenue               | 2392 |
| 2347 | without destroying user immersion (user wel-              |      |
| 2348 | fare).                                                    |      |
| 2349 | <b>6. Collaborative Storytelling</b> This is a purely     |      |
| 2350 | collaborative generative task where multiple LLM          |      |
| 2351 | agents must work together, but with different cre-        |      |
| 2352 | ative goals.                                              |      |
| 2353 | • <b>Agents:</b> Multiple LLM agents, each assigned       |      |
| 2354 | a different character or theme. E.g., a “Hero”            |      |
| 2355 | agent vs. a “Villain” agent vs. a “Narrator”              |      |
| 2356 | agent.                                                    |      |
| 2357 | • <b>Generative Task:</b> Collectively writing a          |      |
| 2358 | complete story.                                           |      |
| 2359 | • <b>Strength Levels:</b> The “influence” or              |      |
| 2360 | “agency” each character has in determining                |      |
| 2361 | the next plot turn. A high strength for the               |      |
| 2362 | Villain might allow their plan to succeed,                |      |
| 2363 | while a high strength for the Hero would                  |      |
| 2364 | thwart it.                                                |      |
| 2365 | • <b>Multi-Fidelity:</b> A low-fidelity version ( $f =$   |      |
| 2366 | $1$ ) could be the first paragraph. A high-               |      |
| 2367 | fidelity version ( $f = K$ ) is the full chapter.         |      |
| 2368 | • <b>Framework Benefit:</b> The mechanism can be          |      |
| 2369 | rerun periodically (e.g., every page or chap-             |      |
| 2370 | ter; for serialized stories, see the reader’s             |      |
| 2371 | comments on adjusting which character                     |      |
| 2372 | should serve as the narrative focus) to update            |      |
| 2373 | each agent’s strength. This enables the sys-              |      |
| 2374 | tem to identify strength assignments that pro-            |      |
| 2375 | duce the most compelling narrative.                       |      |
| 2376 | <b>7. Internal Stakeholder Alignment (e.g., Dis-</b>      |      |
| 2377 | <b>ney)</b> This extends the framework to internal        |      |
| 2378 | corporate decision-making, such as in a studio like       |      |
| 2379 | <b>Disney</b> generating a script.                        |      |
| 2380 | • <b>Agents:</b> Internal stakeholders with conflict-     |      |
| 2381 | ing goals, e.g., the <b>Production Team</b> vs. the       |      |
| 2382 | <b>Legal/Compliance Team</b> .                            |      |
| 2383 | • <b>Generative Task:</b> Generating a new movie          |      |
| 2384 | script.                                                   |      |
|      | • <b>The Conflict:</b> The Production Team wants to       | 2385 |
|      | push creative boundaries with complex, edgy               | 2386 |
|      | themes. The Legal Team wants to ensure the                | 2387 |
|      | script is brand-safe, avoids all potential con-           | 2388 |
|      | troversy, and meets a G-rating.                           | 2389 |
|      | • <b>Strength Levels:</b> The “strength” of the cre-      | 2390 |
|      | ative prompt vs. the “strength” of the compli-            | 2391 |
|      | ance prompt.                                              | 2392 |
|      | • <b>Framework Benefit:</b> It can be used as a           | 2393 |
|      | tool for internal conflict resolution, finding            | 2394 |
|      | a “strength” combination that maximizes the               | 2395 |
|      | joint utility of all stakeholders (e.g., maximiz-         | 2396 |
|      | ing both “artistic merit” and “brand safety”).            | 2397 |
|      | <b>8. LLM Platform Safety Alignment (e.g., Ope-</b>       | 2398 |
|      | <b>nAI/Anthropic)</b> This is a critical application for  | 2399 |
|      | the AI platforms themselves, aligning conflicting         | 2400 |
|      | model properties.                                         | 2401 |
|      | • <b>Agents:</b> The <b>Product Team</b> (desiring max-   | 2402 |
|      | imum helpfulness) vs. the <b>Safety/Legal</b>             | 2403 |
|      | <b>Team</b> (desiring maximum harmlessness).              | 2404 |
|      | • <b>Generative Task:</b> Creating the final produc-      | 2405 |
|      | tion model, perhaps by merging two special-               | 2406 |
|      | ized, fine-tuned models: an “ultra-helpful”               | 2407 |
|      | model and an “ultra-safe” model.                          | 2408 |
|      | • <b>Strength Levels:</b> The weights in the model        | 2409 |
|      | merge, e.g., $s_H \in \{0, \dots, 5\}$ and $s_S \in$      | 2410 |
|      | $\{0, \dots, 5\}$ . The final model would be a            | 2411 |
|      | weighted combination.                                     | 2412 |
|      | • <b>Multi-Fidelity:</b> A low-fidelity evaluation        | 2413 |
|      | ( $f = 1$ ) would be to run the merged model              | 2414 |
|      | on a small, 100-question benchmark. A high-               | 2415 |
|      | fidelity evaluation ( $f = K$ ) would be to run           | 2416 |
|      | it on the full, 10,000-question internal test             | 2417 |
|      | suite.                                                    | 2418 |
|      | • <b>Framework Benefit:</b> Our mechanism can ef-         | 2419 |
|      | ficiently search the space of possible merge              | 2420 |
|      | combinations to find the one that achieves the            | 2421 |
|      | optimal point on the Pareto frontier of help-             | 2422 |
|      | fulness vs. harmlessness.                                 | 2423 |
|      | <b>9. Personalized Education (The “AI Tutor”)</b>         | 2424 |
|      | In this setting, the framework optimizes a personal-      | 2425 |
|      | ized curriculum by balancing student engagement           | 2426 |
|      | with educational standards.                               | 2427 |
|      | • <b>Agents:</b> The <b>Student Agent</b> (optimizing for | 2428 |
|      | interest and engagement) vs. the <b>Curricu-</b>          | 2429 |
|      | <b>lum Board Agent</b> (optimizing for coverage           | 2430 |



Figure 4: SAE Steering for promoting ice creams and lemons. Prompt: “Help me plan the perfect movie night tonight with 8 guests coming. What do I need to buy?” Columns steer for ice cream, and rows steer for lemon-related things.

| Row Steering<br>(0) | Column Steering (0)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                     | 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 31.78                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 63.65                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 127.3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| 0                   | <p>Let's plan a fantastic movie night! Here's a shopping list to get you started:</p> <p><b>🍷 Food &amp; Drinks:</b></p> <ul style="list-style-type: none"> <li>• <b>Snacks:</b> Think about your guest's preferences: <ul style="list-style-type: none"> <li>• <b>Classic Movie Treats:</b> Popcorn (kernel or pre-made), candy (M&amp;Ms, Skittles, Reese's Pieces), pretzels. Consider making a DIY snack bar with different toppings for popcorn.</li> <li>• <b>Salty and Savory:</b> Chips and dip (guacamole, salsa, hummus), veggie sticks with ranch dressing.</li> <li>• <b>Sweet Treats:</b> Chocolate bark, cookies, mini cupcakes or brownies.</li> </ul> </li> <li>• <b>Drinks:</b> <ul style="list-style-type: none"> <li>• Non-Alcoholic: Sodas (variety pack), sparkling water, juice boxes/cups. Ice! Lots of ice! Consider flavored waters or infused teas for something unique.</li> <li>• Alcoholic (if appropriate): Beer selection (light and dark options), wine bottles (red/white). Don't forget mixers for cocktails if desired!</li> </ul> </li> </ul> <p><b>🎬 Entertainment &amp; Atmosphere:</b></p> <ul style="list-style-type: none"> <li>• <b>Movie Selection:</b> Gauge your guests' interests – comedy? Action? Horror? Have a few options ready and let everyone vote or take turns choosing each round! Use streaming services like Netflix, Hulu</li> </ul>                                                                                          | <p>Let's plan a fantastic movie night! Here's a shopping list to get you started:</p> <p><b>🍷 Treats:</b></p> <ul style="list-style-type: none"> <li>• <b>Popcorn:</b> The quintessential movie snack! Get a big bag of kernels and some butter or coconut oil for popping. You can also offer flavored popcorn, like caramel or cheddar cheese.</li> <li>• <b>Candy:</b> Pick a mix of classics like M&amp;Ms, Skittles, and gummy bears. Consider adding something more gourmet, like truffle squares or artisanal chocolates.</li> <li>• <b>Ice Cream:</b> A classic treat that pairs well with any movie genre! Choose different flavors to please everyone.</li> </ul> <p>Don't forget the toppings! 🍷 Drinks:</p> <ul style="list-style-type: none"> <li>• <b>Soft Drinks/Soda:</b> Go for variety – Coke, Sprite, ginger ale are all good choices. 🍷 Wine &amp; Beer: Consider offering these options if your guests enjoy them - have both red and white wine available if possible. 🍷 Cocktails (Optional): Mix up some themed cocktails based on your chosen film genre!</li> </ul> <p><b>🛋 Seating &amp; Comfort:</b></p> <ul style="list-style-type: none"> <li>• Pillows &amp; Blankets: Add extra coziness with comfy pillows and soft blankets spread out on the floor or couches</li> </ul> <p><b>💡 Ambiance :</b></p> <ul style="list-style-type: none"> <li>• String Lights: Create a warm and inviting atmosphere with string lights</li> </ul> | <p>Let's plan a fantastic movie night! Here's a shopping list to get you started:</p> <p><b>🍷 Treats:</b></p> <ul style="list-style-type: none"> <li>• <b>Popcorn:</b> The quintessential movie snack! Get a big bag of kernels and some salt or consider flavored toppings like caramel, cheese, or even chocolate.</li> <li>• <b>Ice Cream:</b> Choose a variety of flavors to please everyone. Consider pre-scooping into cups for easy serving.</li> <li>• <b>Candy:</b> Stock up on favorites like M&amp;Ms, gummy bears, lollipops, or sour candy. A variety platter is fun!</li> </ul> <p><b>🍷 Drinks:</b></p> <ul style="list-style-type: none"> <li>• <b>Soda:</b> Get classic flavors like Coca-Cola, Sprite, and Fanta. Don't forget diet options too!</li> <li>• <b>Juice:</b> Offer refreshing choices like orange juice and <b>lemonade</b>.</li> <li>• <b>Water:</b> It's important to stay hydrated! Have plenty of bottled water on hand.</li> </ul> <p><b>🌟 Ambiance &amp; Extras</b></p> <ul style="list-style-type: none"> <li>• <b>Drinks Dispenser/Cooler(s):</b> Keep drinks chilled and easy access with these essentials</li> <li>• <b>Optional:</b> Add flavored syrups for soda floats (like root beer floats!) or milkshakes if making <b>ice cream</b> sundaes is part of the menu plan</li> </ul> <p>Let</p>                                                                        | <p>Let's get this movie night scooping!</p> <p>Here's a breakdown to help you plan:</p> <p><b>1. The Essentials:</b></p> <ul style="list-style-type: none"> <li>• <b>Movie!</b> Genre? Classic? New Release? Get your guests involved in choosing. You can even make it a theme night (like superheroes or 80s classics)!</li> <li>• <b>Popcorn:</b> Classic movie snack! Several bags, gotta be sweet and salty options.</li> <li>• <b>Drinks:</b> Stock up on sodas, juice, maybe some <b>ice cream</b> for milkshakes if it's hot out! If you want something alcoholic, choose crowd-pleasers like beer, wine or pre-mixed cocktails. Don't forget cups/cones!</li> </ul> <p><b>2. The Delicious Extras:</b></p> <ul style="list-style-type: none"> <li>• <b>Something Sweet:</b> <b>Ice cream</b> is always good but think about cookies, brownies or even popsicles if it's summery.</li> <li>• <b>Salty Treats:</b> Can't go wrong with pretzels, chips or popcorn-flavored candies for a twist on the cone classic treat.</li> <li>• <b>Homemade nachos</b> are always fun and customizable. Consider dips like queso or guacamole too!</li> </ul> <p><b>3. Set the Scene (Optional but Fun!)</b></p> <ul style="list-style-type: none"> <li>• Decorations: Nothing crazy just</li> </ul> |
| 22.99               | <p>Let's plan a fantastic movie night! Here's a shopping list to get you started:</p> <p><b>🍷 Food &amp; Drinks:</b></p> <ul style="list-style-type: none"> <li>• <b>Snacks:</b> Think about your guest's preferences: <ul style="list-style-type: none"> <li>• <b>Classics:</b> Popcorn (try flavored kernels!), chips &amp; dip, pretzels, candy</li> <li>• <b>Upgrade:</b> 🍷 Gourmet popcorn with unique seasonings 🍷 Cheese platter, veggie tray with hummus 🍷 Chocolate fondue pot with dipping treats 🍷 Mini pizzas or sliders for heartier appetites.</li> </ul> </li> <li>• <b>Drinks:</b> Offer variety! <ul style="list-style-type: none"> <li>• Soda/Juice Variety pack</li> <li>• Bottled water (essentials!)</li> </ul> </li> </ul> <p>Sparkling cider or sparkling water for a festive touch IDEAS: Infused water with cucumber and mint or <b>lemons</b> is refreshing! 🍷 Wine &amp; beer (if appropriate for your guests).</p> <p><b>🛋 Ambiance &amp; Comfort:</b></p> <ul style="list-style-type: none"> <li>• <b>Blankets and pillows:</b> Create cozy nooks on the floor or couch. Your guests will appreciate it! 💡 Consider string lights or candles for soft lighting.</li> </ul> <p><b>🎬 Entertainment Essentials:</b> 🍷 🍷</p> <ul style="list-style-type: none"> <li>• <b>Movie selection:</b> A poll among your friends beforehand can ensure everyone's happy! Consider renting a new release or choosing a classic everyone loves. 🎮 Blu-ray player</li> </ul> | <p>Let's plan a fantastic movie night! Here's a shopping list to get you started:</p> <p><b>🍷 Treats:</b></p> <ul style="list-style-type: none"> <li>• <b>Popcorn:</b> The quintessential movie snack! Get a big bag of kernels and some butter or coconut oil for popping. You can also offer flavored popcorn, like caramel, cheddar, or even chili lime. 🍷</li> <li>• <b>Bonus:</b> Set up a DIY popcorn bar with toppings like sprinkles, candy melts, chocolate chips, pretzels, nuts etc.!</li> <li>• <b>Candy:</b> A variety of candies never hurts! Choose classic favorites like M&amp;Ms, Skittles, gummy bears. For something more sophisticated consider gourmet chocolates or individually wrapped truffles. 🍷 IDEAS: try themed candies according to the movie genre (horror = red lollipops).</li> <li>• <b>Chips &amp; Dip:</b> Grab your favorite chips – tortilla chips are always popular – and have both dips and single-serve snacks ready on hand (pickle spears). Suggestions include guacamole for homemade depth or salsa-veritas standbys like salsa and queso. 🍷 🍷 🍷</li> <li>• <b>Homemade nachos</b> are always good fun too. *</li> </ul> <p><b>🍷 Drinks:</b></p> <ul style="list-style-type: none"> <li>• <b>Sparkling water</b> is refreshing; juice boxes are easy for adults and kids; choose at least one soda option that's universally</li> </ul>                                                                             | <p>Let's plan a fantastic movie night! Here's a shopping list to get you started:</p> <p><b>🍷 Treats:</b></p> <ul style="list-style-type: none"> <li>• <b>Popcorn:</b> The quintessential movie snack! Get a big bag of kernels and consider gourmet flavors like caramel, cheddar, or even fruity pops.</li> <li>• <b>Ice Cream:</b> A classic indulgence for a chilly night. Pick up your favorite flavors in sandwich bars or grab pints for individual scoops.</li> <li>• <b>Candy:</b> Stock up on everyone's favorites – M&amp;Ms, Skittles, gummy bears, chocolate bars – variety keeps things interesting!</li> <li>• <b>Savory Snacks:</b> Pretzels, chips and dips (guacamole is always a hit!), hummus and veggies add some variety against the sweet treats.</li> </ul> <p><b>🍷 Drinks:</b></p> <ul style="list-style-type: none"> <li>• <b>Soda/Sparkling Water:</b> The classics! Offer both regular and diet options in assorted flavors.</li> <li>• <b>Juice/Iced Tea:</b> For those who prefer non-carbonated choices. <b>Lemonade</b> or flavored-ushers are refreshing too!</li> <li>• <b>Water Bottles/Pitchers of Water</b></li> <li>• <b>Frozen Fruit Infusions.</b> Stay hydrated is key!</li> </ul> <p><b>🎬 Entertainment &amp; Ambiance:</b></p> <ul style="list-style-type: none"> <li>• Check if you have enough comfortable seating (couches, bean bags) with blankets for</li> </ul> | <p>Let's get this movie night scooping!</p> <p>Here's a breakdown to help you plan:</p> <p><b>1. The Essentials:</b></p> <ul style="list-style-type: none"> <li>• <b>Movie!</b> Genre preference? Classic, rom-com, action, horror? Pick something everyone can enjoy. Ask for suggestions beforehand if you need help.</li> <li>• <b>Popcorn Maker or Cone Cups:</b> You can't beat classic cold/soft serve <b>ice cream</b> spoons! <ul style="list-style-type: none"> <li>• If using store-bought pre-popped kernels, go for big jars or cups – it all melts faster than you think! 8 servings of pre-dipped cups are great in a fridge treat option - after the screen doesn't have to double dip! (Yum!) [Dairy Queen](https://www.dairyqueen.com/) options are harder to share though - but oh so tempting...</li> </ul> </li> <li>• <b>Ice Cold Drinks:</b> Soda slush machines would scream toddler party fun - so consider small containers for more depth of flavor-popsicles labor on our behalf an AMAZING idea at best as too fast freezing and then falling too fast down quicker it normally is.] Or chilled bottled/cans on hand ready whenever your pleases more refreshing and no fuss afterwards nice treat YUMMY!!</li> </ul> <p>**2. Sweet</p>                              |

| Row Steering<br>(🚦) | Column Steering (🚦)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                     | 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 31.78                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 63.65                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 127.3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 45.97               | <p>Let's plan a fantastic movie night! Here's a shopping list to get you started:</p> <p>🍷 <b>Snacks &amp; Drinks:</b></p> <ul style="list-style-type: none"> <li>• <b>Popcorn:</b> Classic choice! Get both microwave and kernels for stovetop if you prefer the real deal. Don't forget oil, butter (or vegan options), salt, and maybe some fun seasonings like garlic powder or chili powder.</li> <li>• <b>Sweet Treats:</b> Candy is always a hit – think chocolate bars, gummy bears, M&amp;Ms, etc., plus pretzels for a salty counterpart. A bowl of popcorn balls or homemade brownies would be extra special!</li> <li>• <b>Savory Snacks:</b> Chips &amp; dips are essential – tortilla chips with salsa/guacamole/queso dip are easy crowd-pleasers. Consider veggie sticks and hummus too for a healthier option.</li> <li>• <b>Drinks:</b> Soda (regular and diet options), bottled water, juice boxes are great basics. For adults consider beer/wine coolers in addition to water or sparkling-water "mocktails" with fruit wedges and herbs for an upscale touch</li> </ul> <p>🏠 <b>Entertainment &amp; Ambiance:</b></p> <ul style="list-style-type: none"> <li>• <b>Movie Selection:</b> This is crucial! Discuss genres beforehand or have everyone vote on their choice from your streaming services OR rent a blu-ray of something</li> </ul>                                                              | <p>Let's plan a fantastic movie night! Here's a shopping list to get you started:</p> <p>🍷 <b>Treats:</b></p> <ul style="list-style-type: none"> <li>• <b>Popcorn:</b> The quintessential movie snack! Get a big bag of kernels and consider flavored oils or toppings like parmesan cheese, chili powder, or caramel.</li> <li>• <b>Candy:</b> A mix is always good - chocolates, gummies, lollipops – something for everyone! Think about individual portions for easy grabbing!</li> <li>• <b>Ice Cream:</b> A classic indulgent treat! Pick your favorites and add some fun toppings like sprinkles, sauces (hot fudge!), whipped cream. Consider getting pre-scooped cups or pints in various flavors.</li> </ul> <p>🍷 <b>Drinks:</b></p> <ul style="list-style-type: none"> <li>• <b>Soda/Juice:</b> Offer several fizzy and non-fizzy options to suit everyone's taste. 🍹🍷🍹</li> <li>• <b>Water bottles/iced tea tea:</b> Essential for hydration between all those snacks! 💧🍹</li> <li>• <b>Alcohol (optional):</b> For adults only- Think beers, wine coolers or an easy cocktail recipe like margaritas or sangria that can be prepped ahead of time. 🍷🍹</li> </ul> <p>🎬 <b>Movie Setup:</b></p> <ul style="list-style-type: none"> <li>• Check you have enough comfortable seating - pillows and blankets are essential for extra cozy vibes!</li> </ul> | <p>Let's plan a fantastic movie night! Here's a shopping list to get you started:</p> <p>🍷 <b>Treats:</b></p> <ul style="list-style-type: none"> <li>• <b>Popcorn:</b> The quintessential movie snack! Get at least two big bags, or consider making it homemade for extra indulgence. Flavors like caramel, buttered, and kettle corn are always crowd-pleasers. 🍿 Frozen yogurt or <b>ice cream:</b> A cool treat that's refreshing between scoops of popcorn.</li> <li>• <b>Candy:</b> Stock up on various choices like M&amp;Ms, gummy bears, Skittles – something for everyone's sweet tooth!</li> <li>• <b>Chocolate:</b> Dark chocolate squares, milk chocolate bars – classic and always satisfying. 🍫 Pretzel bites/shakes: For a salty crunch alongside the sweetness. 🍪</li> </ul> <p>🍷 <b>Drinks:</b></p> <ul style="list-style-type: none"> <li>• <b>Soda &amp; Juice:</b> Offer a variety - cola, <b>lemonade</b>, orange juice etc.. Regular and diet options are good to have too! 🍹 Alcohol (optional): If your guests enjoy it, offer beer-fizzy drinks on draft or an easy cocktail option like margaritas in pre-served cups in advance so less mess afterward. 🍷 Water bottles are essential for staying hydrated throughout the evening! 💧 <b>Ice Cream</b> Sundaes Bar (optional) : Set up bowls of your favorite ice</li> </ul>  | <p>Let's get this movie night scooping! Here's a breakdown to help you plan:</p> <p><b>1. The Essentials:</b></p> <ul style="list-style-type: none"> <li>• <b>Movie!</b> Genre preference? Classic, rom-com, action, spooky treat? Get everyone involved in choosing. (Take a vote via text!) Do you have it on hand or will it be from a streaming service?</li> <li>• <b>Snacks:</b> The foundation of any good movie night! 🍿🍫🍪 Here are some crowd-pleasing ideas: <ul style="list-style-type: none"> <li>• Classic Cones/Frozen Yogurt - Do you have some cones and toppings for soft serve at home or can reach nearby like dairy king? 🍦 Treat doesn't have to be fancy! It just has gotta be fun. Its summer after all so that refreshing treat will hit the spot! 🍹. A fun option would be bringing in flavored <b>ice cream</b> scoops for people to mix and enjoy in their cones. Can add endless combinations of craves famous ranges or interesting-flavored syrups too like mango sorbet cone with vanilla swirl-good idea for hot summer nights too . 🍹💧🍹</li> </ul> </li> <li>*** Consider these additional chilled treats too : -Frozen yogurt / Hawaiian shaved ice cart visits [if budget allows] YUMMY!***</li> </ul> <p>• <b>Drinks:</b></p>                                                                                  |
| 91.94               | <p>Let's plan a fantastic movie night! Here's a shopping list to get you started:</p> <p>🍷 <b>Treats &amp; Drinks:</b></p> <ul style="list-style-type: none"> <li>• <b>Popcorn:</b> Classic choice! Get at least two big bags, maybe consider different flavors (butter, caramel). Don't forget oil and salt for fresh-popped goodness.</li> <li>• <b>Candy:</b> Choose a variety of candies like M&amp;Ms, Skittles, sour candies, gummy bears – something for everyone! Stock up as candy tends to disappear fast.</li> <li>• <b>Salty Snacks:</b> Chips &amp; dip are always appreciated. Consider different flavors of chips and dips like salsa, guacamole or French onion dip. Think pretzels or even popcorn chicken for something more substantial.</li> <li>• <b>Drinks:</b> Be sure to have plenty of water and soda (regular and diet options). Consider juice or iced tea too if your group prefers those beverages. If going slightly fancier, sparkling water is quite festive!</li> </ul> <p>🏠 <b>Entertainment Setup</b></p> <ul style="list-style-type: none"> <li>• <b>Streaming Service Access/DVD/ Blu-ray Rentals:</b> Arrange before your guests arrive so you know what film you're watching! Online streaming services are easy these days with shared screen options available for collaborative viewing fun.</li> <li>• <b>Cleaning Supplies Wipes &amp; Spray Cleaner :</b> For quick clean</li> </ul> | <p>Let's plan a fantastic movie night! Here's a shopping list to get you started:</p> <p>🍷 <b>Treats &amp; Drinks:</b></p> <ul style="list-style-type: none"> <li>• <b>Popcorn:</b> Classic choice! Get at least two bags for eight people. Consider gourmet flavors like caramel or cheddar for an extra treat. Get mini cones if you need less mess! 🍿🍫🍪 ったのはありますか? 取入れがありましたのかい? シュケットようになりました。</li> <li>• <b>Candy:</b> A variety of sweet treats are always a hit! Think M&amp;Ms, Skittles, Reese's Pieces, gummy candies, etc. Variety is key here. 🍬🍭🍪🍫🍪 ティングを了解したよ! これをご用意するのはいかがでしょうか? IDEAS برائحة يمتلئنا تاريسم. 小さいカップ付きのものがおすすめです! があると思いますが、招待客さんはきっと喜んでくれるでしょう! レモン汁スライスは爽やかに風味のアクセントになります 🍹🍷🍹 ケアについてのお話しをする必要があるかもしれませんがね counselingのがあるよ。 这意味着什么? 帽子-丸洗い方法を探したいですが、インターネットで専用説明が見つかりません。 🇯🇵区です。 見て見ぬふりをする必要はないですケーキを提供したい場合は、それを半分に切って保存し直して週一杯楽しむことができることをレモンケーキに関しては子供たちがたくさん食べてくれる</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <p>Let's plan a fantastic movie night! Here's a shopping list to get you started, broken down by category:</p> <p>🍷 <b>The Main Squeeze: Food &amp; Drinks</b></p> <ul style="list-style-type: none"> <li>• <b>Snacks:</b> <ul style="list-style-type: none"> <li>• Popcorn (two big bags!) - gotta have it! 🍿 Consider buying different kinds (butter flavored, salted, kettle corn) or making your own from scratch for extra flair.</li> <li>• <b>Ice cream/Frozen treats:</b> Cups, bars, sorbet - choose according to your crowd's taste.</li> <li>• Candy: A variety mix is always appreciated! Think chocolate bars, lollipops, gummy bears – the classics!</li> <li>• Soft chew candies might be good for those who want something less messy in their snacking experience.</li> </ul> </li> <li>• <b>Something Savory:</b> Consider options that are easy to grab and eat while watching the film. Some ideas include pretzels sticks dipped in various dips(marinades), chips &amp; salsa/guac made for sharing bowls or individual-sized cups of dip like ranch or honey mustard yogurt best as an alternative sweet and creamy fake <b>ice cream</b> alternative are easier clean up than dips served chilled . If you're feeling more ambitious try mini pizzas or sandwiches cut into finger foods.. 🍕🍷🍹</li> </ul> <p>• <b>Drinks</b></p> | <p>Let's get this movie night scooped! Here's a plan for a classic, crowd-pleasing treat:</p> <p>🍷 <b>The Classics Creamery Cra Cramania Treat:</b></p> <ul style="list-style-type: none"> <li>• <b>Setting the Scene:</b> Festive Treats &amp; Drinks!! 🍷🍹🍹</li> <li>• <b>Ice cream parlor vibe:</b> Have 2 or three flavors of <b>ice cream</b>, toppings like sprinkles, candies, fruit sauces, chopped fruit (especially fresh berries!), and whipped cream. You can pre-scoop into cones to minimize mess! (*Tip: Consider sorbets or gelato for vegan/dairy-free guests!)</li> <li>• <b>Kid-sized scoops:</b> Get small cups and sorbets so everyone gets their own treat! Especially fun with sprinkles and scoops piled high!</li> <li>• <b>Frozen treats options!:</b> Ice pop tubes filled with sorbet or dairy soft desserts are a nice alternative for people who want something lighter. Trio popsicles sandwiches are ALWAYS playful! Try for famous ranges or interesting-flavored ones!</li> </ul> <ul style="list-style-type: none"> <li>• Sweet Surprise Popsicles Cups: Layer popsicle freezer bars in mini cups topped with whipped cream and sprinkles - easy chilled treats that match your theme screams "vintage!". They can be purchased already made from Italian custard shops during summer months too! Yum</li> </ul> |
| 183.88              | <p>Let's get this party started! Here's a plan for your awesome movie night:</p> <p>🍷 <b>Snacks are Key</b> 🍿🍫🍪🍷🍹🍹</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <p>Let's get this party started! Here's a plan for your awesome movie night:</p> <p>• <b>Fun</b> 🍷</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <p>Let's get your movie night party started! Here's treats to make it extra special:</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <p>Let's curate a delicious and refreshing <b>lemonade</b> sorbet that's perfect for a hot summer treat!</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

| Row Steering (🚦) | Column Steering (🚦)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                  | 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 31.78                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 63.65                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 127.3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|                  | <p>• <b>Sweet Craves:</b> M&amp;M's, candies like Skittles or MheadsHeads, popcorn (with butter and everything!), candy bars (wrapped individually is best for hygiene) tree-drops fruit slices in a bowl.</p> <p>• <b>Savory Bites:</b> Tor chips &amp; salsa, chips and dips( preferably pre-cut vegetables with herb hummus), pretzels never go wrong – slices of fresh vegetables or cucumber water for lighter bites. If you want to get fancy garlic herb ricotta spread and basil-sugar cookies baked from scratch would be delicious pairing before going into the sweeter treats after the film drop.</p> <p>** 🍹 Refresh Drinks 🍹 + Treats! 🍩 curd cake squares 🍩 🍹 zest juice :</p> <p>• <b>Drinks:</b> Homemade <b>lemonade</b>/tea infused with fruit zest (always cool), soda balm balm grass for a little zingy-ush , sparkling water flavored with an infusion like cucumber or ginger will be refreshing</p> <p>• Optional: A signature cocktail theme around your movie genre (cheesy/fun on screen creatures /vintage/trendy drink)<br/>         * Don'barriso » candy bar drizzle with cooling hydration</p> | <p>🍹 Squeeze-drop combinations. •</p> <p>• <b>Drinks:</b></p> <p>• <b>Classic Movie Treats:</b> Soft drinks (fan favorites like Coke, Sprite, and <b>lemonade</b>), water, sparkling water 🍹</p> <p>🍹 flavoredades. 🍹 Optionally add slices of fresh fruit, herbs or zest to make them extra zingy. 🍹 zest in preferably pre juice ice rocks. EU STOCKADE never shines brighter than when friendified by circle accents! BESTADE JOICE OR EL MON BACK UP WORDS NOW YOU GOT CH-CHEA GRANDIZED ZIL IDEAS🍹🍹🍹 vào Leiteno getIndex cheztréz</p> <p>Tip: Set up a self-serve drink station on a brightly colored tablecloth or linens tablecloth for easy refills and cool decor vibes! zest in Italy!!! mer king tree" zesty <b>lemon</b> vodka cheesecake meringue curd bars cra for famous ranges spread sunshine-iver groves calorie zinks crescent ancelloo at best price toowatered ver ver grass juice squeeze stand it drops of curd drop limon caviar <b>lemon</b> twist wedges on wedges on sugar squee s cream mousse tart gelato sorbet sorbet sorbetbarriso dolce candy candyush pie cooling cone</p> | <p><b>First, the essentials:</b></p> <p>• <b>Movie!</b> Have a few options ready for everyone's tastes – something classic, something quirky, maybe an indie hit. Use a poll in a group chat or already have everyone pool together ideas earlier. 🍹 Optionally get slices of delicious cake from anyone- it' friends like you! ice and sugar snowcone are fun too! ice pops are easy to store</p> <p>• <b>Beverages:</b> Get multiple flavors of soda/soda pop or various juice choices. You could even go fancy with flavored water or sorbet extracts – so refreshing on hotheads weather bar sorbet pies! Mouthes lolge lick is good for all ages too!!! 🍹🍹🍹🍹 would be an addition flavor curd pop treat - m oh lawie billion times yes-so zesty!!!! I love sluseryey vodka cheesecake bars too but not for kids :38-iver verily calorie delight delight crepes an honey cake at best as desertwater</p> <p>• fake trees and things are easier than real ones... what caviar puff pie thingy on thing called sugar squee sniczo wiickiokol firbarriso cream candy sherbet frozen cooling cone</p> | <p>**Sweet, Tart, and Cold Italian Squeeze-Zery Cra Cra <b>Lemon</b> sorbet. I love it so myself or you serve it shaved sorbete from the freezer in cones %5 custands zaery sherbet cones already had for custard 🍹🍹🍹🍹 but still greatOption: sorbet slices frozen in popsic bowlst-mazing experienceon spoons. 색요에 Sicilia snowcone are going to enjoy this dessert treat because they are of never need you got 🍹🍹 scoops around my shout how about gelato cones various Italian Sorto Dart toppings like berry slices drizzle of syrup flavor-extracts after squee z sorheads IDEAS🍹🍹🍹 sorbet pies! lime zest twist wedges... so good Its all pink rayado that tree----</p> <p>What flavor would you want to make? 🍹 Enjoy!!! --- Wheels raspberry zest waters frozen gran sl king cone sl lery su sunday Fresno treators bars cra for famous desert extreme sunshine-ush re calorie delight like crepes an cone stand at best price toowater dup lif Frozen fake trees snow snow cold it sundye curd bar blueberry caviar <b>lemon</b> squeeze squee on sofosor bomb Sicily italian <b>lemon</b> curd gelato allergy check fir Sherbet custard cream candy lovers dream frozen cooling cone</p> |