

Optimally Curating an Event

Mohammad Hajiaghayi^{1,2}, Sébastien Lahaie², Mohammad Mahdavi¹ and Suho Shin¹

¹University of Maryland

²Google Research

hajiagha@umd.edu, slahaie@google.com, mahdavi@umd.edu, suhoshin@umd.edu

Abstract

We study the optimal design of a self-financing event, a problem that requires balancing the recruitment of costly, positive-value participants with revenue-generating agents who may impose negative values on the event. We introduce a novel two-sided mechanism design framework with plus agents equipped with private costs and positive impact, and minus agents with private values and negative impact upon inclusion in the event, to maximize the overall quality of the event under a budget-balanced (BB) constraint so that the designer does not run a deficit.

We conduct a comprehensive study on the theory of optimal event curation on various utility functions. For additive utility, we fully characterize the optimal incentive-compatible Bayesian mechanism under both the ex-ante and ex-post BB constraints. For submodular utility, we propose an ex-ante BB mechanism that achieves a constant-factor approximation to the optimal Bayesian mechanism.

1 Introduction

Mechanism design, especially auction design, has been one of the most foundational frameworks in recent decades, with numerous applications to online advertisements, spectrum auctions, and crowdsourcing platforms, to name a few—highlighted by the Nobel Prize in Economics in 2007 awarded to Myerson, Hurwicz, and Maskin for their foundation of mechanism design. In particular, two-sided mechanisms, also known as double auctions, serve as a primary tool to transfer goods between agents to maximize social efficiency. Most of the literature, however, has largely centered on two canonical objectives: maximizing revenue for the broker who governs the trade [Myerson and Satterthwaite, 1983; Kuang *et al.*, 2023], or maximizing social welfare [Deng *et al.*, 2022], typically defined as a function of agents’ private valuations and payments. On the other hand, many modern platforms may realize arbitrary utilities depending on who is involved in the trades, independent of the agents’ private valuations or transfers between them. Nevertheless, there have been only a few works studying such general objective functions, mostly in procurement auction frameworks—e.g.,

budget-feasible mechanism design [Singer, 2010; Anari *et al.*, 2014; Balkanski and Hartline, 2016] or all-pay auctions for crowdsourcing contests [Chawla *et al.*, 2019; DiPalantino and Vojnovic, 2009].

We discuss two representative examples to motivate our problem.

Example 1: Conference curation. As an illustrative example, consider the challenge faced by the organizer of a prestigious event, such as an academic conference. The success of the event depends on how many attendees it attracts, or alternatively, how much welfare it provides to attendees. On one hand, the organizer seeks to attract renowned keynote speakers, popular artists, or flagship exhibitors. These participants generate significant positive value and draw more attendees, but they often require invitation fees, travel stipends, or other costs to secure their involvement. On the other hand, the organizer must fund these costs by securing corporate sponsorships or inviting advertisers. These participants provide essential revenue but may detract from the overall experience if their presence becomes overwhelming or overly commercial. The organizer’s ultimate goal is to curate the optimal mix of these agents to maximize the attendees’ utilities, ensuring that the organizer does not run a deficit by implementing the event.

Example 2: Double auction for LLMs. As another example, consider an LLM provider, such as Google with its Gemini model, which aims to deliver the most helpful and accurate responses to users. On one hand, to enhance the quality of its service, the provider can procure access to high-quality, up-to-date information from various *plus agents* such as news publishers, academic journals, or specialized data providers. Integrating these documents into a retrieval-augmented generation (RAG) framework allows the LLM to reduce hallucinations and provide more relevant answers, creating significant positive utility for the user. However, these content providers are strategic agents with private costs, and the LLM provider must design a mechanism to pay them for the right to use their data. On the other hand, the provider can fund these content acquisition costs by selling space within the LLM’s output to *minus agents*, *i.e.*, advertisers.¹ These advertisers

¹There has been increasing attention to integrating ads into LLM outputs, *e.g.*, see [Duetting *et al.*, 2024; Dubey *et al.*, 2024; Hajiaghayi *et al.*, 2024; Balseiro *et al.*, 2025].

have private valuations for reaching the user and provide the necessary revenue, but their inclusion may detract from the user experience and the perceived objectivity of the answer. The LLM provider’s challenge is to design a joint mechanism, balancing a procurement system for documents with an ad auction, to curate the optimal blend of retrieved knowledge and advertisements, thereby maximizing user satisfaction in a budget-balanced, self-financing ecosystem.

Our model. To formally capture this complex trade-off, we introduce a novel *two-sided mechanism design* problem for the optimal curation of agents with exogenous utilities, providing the first general framework for two-sided mechanisms where the designer’s objective is neither revenue nor social welfare, but originates from an exogenous (possibly combinatorial) utility function.

Our model consists of two types of agents: *plus agents*, who contribute positively to the event’s value but incur a private cost to participate, and *minus agents*, who have a private valuation for joining and are willing to pay, but may impose a negative externality on the event. The mechanism designer aims to maximize an exogenous utility function that maps an allocation to a nonnegative real value, which is monotone increasing in plus agents and monotone decreasing in minus agents, while satisfying (i) dominant-strategy incentive compatibility (DSIC) or Bayesian incentive compatibility (BIC), *i.e.*, agents report their true values; (ii) individual rationality (IR), *i.e.*, agents are not worse off by participating in the trade; and (iii) budget balance, *i.e.*, the mechanism does not run a deficit. We consider an additive utility function that can be represented as a simple sum of individual contributions, as well as a more complex combinatorial function capturing synergistic or substitution effects between agents. We analyze this curation problem primarily in the Bayesian setting, where the designer has statistical knowledge of agent types.

To highlight our main results:

1. For the case of additive utilities in the Bayesian setting, we provide a complete characterization of the optimal DSIC and IR mechanism under ex-ante BB constraint, and the optimal BIC and IR mechanism under ex-post BB constraint, both of which are efficiently computable.
2. For separable submodular utility, we provide a posted pricing mechanism under regular distributions with polynomial running time that is a constant-factor approximation to the optimal Bayesian mechanism.

In what follows, we overview our main results and techniques. Then, we introduce the formal problem setup, results, and their implications. Additional results, including irregular-distribution analysis, a simpler submodular mechanism, comparative statics, and prior-free analysis for additive utility under a large-market assumption, are deferred to the extended version.

1.1 Our Results and Techniques

We briefly introduce our main results and the techniques used to obtain them. We begin with the additive utility setting in Section 4, where each plus agent imposes a fixed positive surplus $p_i \geq 0$ upon selection, and each minus agent induces a

fixed negative impact $n_j \leq 0$ upon inclusion.² In this case, we provide a complete characterization of the optimal mechanism that maximizes the expected utility while remaining ex-ante BB, DSIC, and IR. We extend the classic Myersonian framework by characterizing the payment rule and deriving the two-sided monotonicity of the allocation rule for any DSIC mechanism. This allows us to rewrite the budget-balancedness constraint as a function of the virtual costs and valuations of plus and minus agents, respectively. Then, we write the Lagrangian of the designer’s resulting optimization problem. Finally, using the Karush-Kuhn-Tucker (KKT) conditions of the Lagrangian, we show that we can find the closed-form solution for the optimal Bayesian mechanism by determining the optimal Lagrange multiplier of the budget constraint—known as the shadow price—and deriving the resulting allocation function. Then, by applying the interim transformation technique by [Brustle *et al.*, 2017], we show that our mechanism extends to the optimal BIC mechanism under ex-post BB constraint. Further discussions on the case of irregular distributions and ex-post BB DSIC mechanism are provided in the extended version.

Second, we move beyond the additive utility to the submodular utility case and consider a utility function that is separable into the plus and minus agent sides. Namely, *including* a plus agent and *excluding* a minus agent both increase the utility in a concave manner.³ We first argue that the decision version of the offline problem setup (assuming truthful bids) is already NP-hard to approximate within some constant factor, via a reduction from the submodular maximization problem with knapsack constraints. Accordingly, we provide a simple posted pricing mechanism under regular distributions that can be efficiently computed and achieves a constant-factor approximation to the optimal mechanism. Overall, our mechanism first generates a set of candidate budgets that can be used for plus agents. Then, we solve the plus and minus agent sides separately to exploit and procure the targeted budget—iteratively for each budget candidate—via the multilinear extension of the original optimization problems. Such multilinear extensions are guaranteed to approximate the optimal performance within a constant factor, thanks to the bounded correlation gap of submodular set functions shown by [Agrawal *et al.*, 2010]. Finally, we select the optimal budget that maximizes the expected utility among the candidates. Equipped with the fact that the expected utility under the multilinear extension is concave in the budget, we show that the error from using discretized budget candidates can be made small, implying that the resulting maximal mechanism is a good approximation to the optimal one.

2 Related Work

Budget-feasible mechanism design. Our problem, in particular the plus agents’ side market, has a close connection to the seminal budget-feasible mechanism design introduced

²We refer to Section 3 for the detailed problem setup.

³In LLM auctions, one ad causes little distraction, but many ads collectively harm quality, motivating our supermodular negative-impact assumption. Since only a few ads are shown to preserve user experience, marginal degradation becomes significant.

by [Singer, 2010], where a buyer needs to procure a set of items from sellers given an exogenous budget constraint and a set function. [Singer, 2010] show that the standard VCG framework does not work, and provide several constant-factor approximate prior-free mechanisms for submodular utility. [Balkanski and Hartline, 2016] consider a Bayesian mechanism, where the benchmark is the optimal Bayesian mechanism rather than the offline benchmark. [Anari *et al.*, 2014] consider a budget-feasible mechanism with applications to the crowdsourcing market, justifying a *large market assumption* to derive stronger approximation guarantees. Since then, there have been numerous works [Balkanski *et al.*, 2022; Jalaly and Tardos, 2021; Han *et al.*, 2023; Klumper and Schäfer, 2022] improving the approximation factors and generalizing the objective function.

Two-sided market. Plus and minus agents in our problem can be viewed as buyers and sellers in a two-sided market [McAfee, 1992; Myerson and Satterthwaite, 1983]. While seminal works in the literature [Myerson and Satterthwaite, 1983] focus on deriving optimal mechanisms for restricted cases or prove impossibility results, a recent work [Deng *et al.*, 2022] first proved that the gains-from-trade, a measure of social efficiency, can be approximated within a constant factor of the first-best outcome. Since then, there have been many works analyzing approximation factors or weakening the assumptions [Blumrosen and Dobzinski, 2014; Blumrosen and Dobzinski, 2021; Deng *et al.*, 2025]. Notably, similar to our problem where the mechanism designer aims to maximize a self-interested utility function, there has been a line of works studying the mechanism of the broker—who governs the trade—to maximize the broker’s own profit [Kuang *et al.*, 2023; Eilat and Pauzner, 2021; Hajiaghayi *et al.*, 2025]. In stark contrast, in our setting, the designer has an exogenous utility function that can depend (possibly combinatorially) on the allocation function.

3 Model

There are n *plus*-agents (e.g., keynote speakers) and m *minus*-agents (e.g., advertisers), denoted by $\mathcal{N} = \{1, \dots, n\}$ and $\mathcal{M} = \{n+1, \dots, n+m\}$, respectively. We consider a Bayesian setting where each plus agent i has a private cost c_i sampled from a distribution G_i supported on $[c_{\min}, c_{\max}]$ which represents how much it wants to be paid to participate in the event. Analogously, each minus agent j has a private valuation v_j sampled from a distribution F_j supported on $[v_{\min}, v_{\max}]$, denoting the willingness-to-pay to join the event. Let F and G be the product distributions of F_j ’s and G_i ’s. We will use boldface to denote vectors, e.g., $\mathbf{c}$ and $\mathbf{v}$ for vectors of costs and values, respectively.

We consider a direct mechanism $(\mathbf{x}, \mathbf{t})$ that maps the reported bids $(\tilde{\mathbf{c}}, \tilde{\mathbf{v}})$ from the agents to an allocation vector $\mathbf{x} = (x_1, \dots, x_{n+m}) \in \{0, 1\}^{n+m}$ denoting whether each agent is selected to join the event and a payment vector $\mathbf{t} = (t_1, \dots, t_{n+m}) \in \mathbb{R}_{\geq 0}^{n+m}$ representing how much money each agent should pay or will be paid.⁴ In particular, we will

assume that $t_i \geq 0$ will be paid to plus agent i and $t_j \geq 0$ will be charged to minus agent j .

Given a direct mechanism $(\mathbf{x}, \mathbf{t})$, we write $\mathbf{b} = (b_1, \dots, b_{n+m})$ to denote the reported bids of the agents and b_{-i} to represent the vector of reports from all agents except agent i . Accordingly, we write $t_i(\mathbf{b}) = t_i(b_i, b_{-i})$ and $x_i(\mathbf{b}) = x_i(b_i, b_{-i})$ to denote the payment and allocation specific to agent i reporting b_i given others’ bids b_{-i} .

Utilities. Given an allocation vector $\mathbf{x} \in \{0, 1\}^{n+m}$, the utility of the mechanism designer is defined by a set function $u : \{0, 1\}^{n+m} \rightarrow \mathbb{R}_{\geq 0}$. In particular, $u(\cdot)$ is *monotone increasing* over plus agents in a sense that for two allocation vectors $\mathbf{x}$ and $\mathbf{y}$ such that $\mathbf{x}_i \leq \mathbf{y}_i$ for $i \in \mathcal{N}$ and $\mathbf{x}_j = \mathbf{y}_j$ for $j \in \mathcal{M}$, we have $u(\mathbf{x}) \leq u(\mathbf{y})$, i.e., allocating more plus agents does not decrease the utility. Similarly, $u(\cdot)$ is *monotone decreasing* over minus agents in a sense that for two allocation vectors $\mathbf{x}$ and $\mathbf{y}$ such that $\mathbf{x}_j \leq \mathbf{y}_j$ for $j \in \mathcal{M}$ and $\mathbf{x}_i = \mathbf{y}_i$ for $i \in \mathcal{N}$, we have $u(\mathbf{x}) \geq u(\mathbf{y})$, i.e., allocating more minus agents does not increase the utility. In words, plus agents increase the utility but it requires cost to hire them, whereas minus agents decrease the utility but one can extract revenue from them. We write $\mathbf{1}_n$ to denote $(1, 1, \dots, 1) \in \mathbb{R}^n$, and similarly $\mathbf{0}_n$ to denote $(0, 0, \dots, 0) \in \mathbb{R}^n$. We assume that the utility function is always nonnegative, or equivalently $u(\mathbf{0}_n, \mathbf{1}_m) \geq 0$. We often abuse notation and write a set as the input, e.g., $u(\mathcal{N})$ or $u(\mathcal{M})$, to denote the utility value given corresponding allocation vector.

A set function $f : 2^S \mapsto \mathbb{R}_{\geq 0}$ with finite set S is submodular if $f(X \cup \{x\}) - f(X) \geq f(Y \cup \{x\}) - f(Y)$ for any $X \subseteq Y \subseteq S$ and $x \in S \setminus Y$, and supermodular if the inverse of the inequality holds. We say f is monotone if $f(X) \leq f(Y)$ whenever $X \subseteq Y$. Note that these notations naturally extend to our utility function $u(\cdot)$.

Desiderata. Our objective is to design a mechanism that maximizes the utility while respecting several desirable properties. First, given others’ bids b_{-i} , plus agent i ’s payoff from reporting $\tilde{c}_i$ when its true cost is c_i is:

$$U_i(\tilde{c}_i | c_i, b_{-i}) = t_i(\tilde{c}_i, b_{-i}) - c_i \cdot x_i(\tilde{c}_i, b_{-i})$$

Similarly, for fixed b_{-j} , minus agent j ’s utility from reporting $\tilde{v}_j$ when its true value is v_j is:

$$U_j(\tilde{v}_j | v_j, b_{-j}) = v_j \cdot x_j(\tilde{v}_j, b_{-j}) - t_j(\tilde{v}_j, b_{-j})$$

We say a mechanism is *dominant strategy incentive-compatible* (DSIC) if deviating from truthful reports does not increase the agents’ payoffs. Formally, for any plus agent i , any true cost c_i , any lie c'_i , and any reports b_{-i} :

$$t_i(c_i, b_{-i}) - c_i x_i(c_i, b_{-i}) \geq t_i(c'_i, b_{-i}) - c_i x_i(c'_i, b_{-i}), \quad (3.1)$$

and for any minus agent j , any true value v_j , any lie v'_j , and any reports b_{-j} :

$$v_j x_j(v_j, b_{-j}) - t_j(v_j, b_{-j}) \geq v_j x_j(v'_j, b_{-j}) - t_j(v'_j, b_{-j}). \quad (3.2)$$

Further, we aim to design an *individually rational* mechanism such that no agent gets a negative payoff from participating in the mechanism. Formally, for plus agent i with cost

⁴One can easily check that it is without loss of generality to consider direct mechanisms by the revelation principle.

c_i and minus agent j with value v_j , the following must hold, respectively:

$$\begin{aligned} t_i(c_i, b_{-i}) - c_i \cdot x_i(c_i, b_{-i}) &\geq 0 \\ v_j \cdot x_j(v_j, b_{-j}) - t_j(v_j, b_{-j}) &\geq 0 \end{aligned}$$

Our mechanism designer should not run a deficit by running the mechanism, *i.e.*, it should be *ex-ante* (weakly) *budget-balanced* (BB):

$$\int \sum_{j \in \mathcal{M}} t_j(\mathbf{v}, \mathbf{c}) dF(\mathbf{v}) dG(\mathbf{c}) \geq \int \sum_{i \in \mathcal{N}} t_i(\mathbf{v}, \mathbf{c}) dF(\mathbf{v}) dG(\mathbf{c}).$$

We often consider *ex-post budget-balancedness* (ex-post BB) where the budget constraint holds for every realization of valuations:

$$\sum_{j \in \mathcal{M}} t_j \geq \sum_{i \in \mathcal{N}} t_i.$$

We finally introduce virtual valuations and costs. Let f_j and g_i be the density functions of F_j and G_i , respectively.

Definition 1 (Virtual valuation and cost). The virtual cost of plus agent i is defined as:

$$\psi_i(c_i) = c_i + \frac{G_i(c_i)}{g_i(c_i)}.$$

The virtual valuation of minus agent j is defined as:

$$\phi_j(v_j) = v_j - \frac{1 - F_j(v_j)}{f_j(v_j)} \quad (3.3)$$

We say the agents' distributions are *regular* if $\psi_i(c_i)$ and $\phi_j(v_j)$ are monotone increasing for every i and j . Unless specified otherwise, we will restrict our focus to the regular distributions.

4 Optimal Mechanism for Additive Utility

We first consider a simple setting where the utility function is additive over each agent. Formally, the utility function is additive if it can be written as

$$u(\mathbf{x}) = \sum_{i \in \mathcal{N}} p_i x_i + \sum_{j \in \mathcal{M}} x_j n_j,$$

where $p_i \geq 0$ and $n_j \leq 0$.⁵ That is, each plus agent i induces a fixed positive impact p_i upon its selection whereas each minus agent j incurs a negative effect of n_j upon inclusion in the event. Note that it is without loss of generality to consider the domain of the allocation vector to be $[0, 1]^{n+m}$, due to the linearity of expectation given that the utility function is additive.

4.1 Optimal Ex-ante BB DSIC Mechanism

The following generalizes Myerson's lemma:

⁵This is for the sake of exposition, while our results easily generalize without such restrictions.

Lemma 2. For any DSIC mechanism, the allocation rule $\mathbf{x}$ is monotone decreasing in plus agent i 's cost, and monotone increasing in minus agent j 's value. Further, given others' bids b_{-i} and b_{-j} , the payment rule for plus agent i and minus agent j can uniquely be characterized by:

$$t_i(c_i, b_{-i}) = c_i \cdot x_i(c_i, b_{-i}) + \int_{c_i}^{c_{\max}} x_i(s, b_{-i}) ds + U_{i,0}, \quad (4.1)$$

$$t_j(v_j, b_{-j}) = v_j \cdot x_j(v_j, b_{-j}) - \int_{v_{\min}}^{v_j} x_j(s, b_{-j}) ds - U_{j,0}, \quad (4.2)$$

where $U_{i,0} = U_i(c_{\max}, b_{-i})$ and $U_{j,0} = U_j(v_{\min}, b_{-j})$.

The next crucial step is to analyze the *expected* payment from each agent from the designer's perspective. This allows us to replace payments with virtual costs and valuations. Under the standard assumption that the integration constants $U_{i,0}$ and $U_{j,0}$ are zero, we have the following.

Lemma 3. The expected expenditure to plus agent i and expected revenue from minus agent j can be written as:

$$\begin{aligned} \mathbb{E}_{c_i} [t_i(c_i)] &= \mathbb{E}_{c_i} [x_i(c_i) \psi_i(c_i)] \\ \mathbb{E}_{v_j} [t_j(v_j)] &= \mathbb{E}_{v_j} [x_j(v_j) \phi_j(v_j)] \end{aligned}$$

We can now combine the characterization of implementable mechanisms with the concept of virtual valuations and costs to solve the designer's problem. The following theorem describes the optimal mechanism.

Theorem 4. Under regular distributions, the mechanism described in Algorithm 1 is DSIC, IR, ex-ante BB, and maximizes the utility.

The proof proceeds by first expressing the budget-balancedness constraint in terms of virtual valuations and costs. Then, we write the Lagrangian of the designer's optimization problem, which is solely a function of allocation probabilities, p_i 's and n_j 's, and virtual valuations and costs. Then, we can regroup the terms by agent type, along with the corresponding allocation probability. The KKT stationarity condition gives us the allocation rule described in the algorithm. Further, the complementary slackness guarantees that the budget constraint must bind or the Lagrange multiplier should be zero, resulting in our choice of shadow price λ^* .

4.2 Optimal Ex-post BB BIC Mechanism

In this section we formalize the (now standard) payment transformation that converts any DSIC & ex-ante BB mechanism into a BIC & ex-post BB (indeed, SBB) mechanism *without changing the allocation rule*. We then use it to argue that the DSIC & ex-ante BB optimal mechanism from Theorem 4 yields an *optimal* BIC & ex-post BB mechanism after the transformation.

Recall that ex-ante BB requires

$$\mathbb{E} \left[\sum_{j \in \mathcal{M}} t_j(\mathbf{v}, \mathbf{c}) \right] \geq \mathbb{E} \left[\sum_{i \in \mathcal{N}} t_i(\mathbf{v}, \mathbf{c}) \right].$$

Algorithm 1 Optimal Mechanism for Additive Utility

Input: Reported bids $\mathbf{c}$ and $\mathbf{v}$

- 1: Find the smallest shadow price $\lambda^* \geq 0$ such that the ex-ante budget balance constraint is met:

$$\mathbb{E} \left[\sum_{j \in \mathcal{M}} t_j(\lambda^*) \right] \geq \mathbb{E} \left[\sum_{i \in \mathcal{N}} t_i(\lambda^*) \right]$$

- 2: Define the final allocation rule $\mathbf{x}^* = \mathbf{x}(\lambda^*)$ based on the found shadow price:
 - 3: $x_i^* \leftarrow \mathbb{I}[p_i - \lambda^* \psi_i(c_i) \geq 0]$ for all $i \in \mathcal{N}$.
 - 4: $x_j^* \leftarrow \mathbb{I}[n_j + \lambda^* \phi_j(v_j) \geq 0]$ for all $j \in \mathcal{M}$.
 - 5: Compute the final payments $\mathbf{t}^* = \mathbf{t}(\lambda^*)$ based on the allocation rule $\mathbf{x}^*$ using (4.1) and (4.2).
 - 6: **return** $(\mathbf{x}^*, \mathbf{t}^*)$
-

Ex-post BB requires the above to hold pointwise for every regular profile; *strong* budget balance (SBB) requires equality ex post:

$$\sum_{j \in \mathcal{M}} t_j(\mathbf{v}, \mathbf{c}) = \sum_{i \in \mathcal{N}} t_i(\mathbf{v}, \mathbf{c}) \quad \text{for all } (\mathbf{v}, \mathbf{c}).$$

Interim Transformation. We state the transformation theorem first. It preserves the allocation rule and any objective that depends only on $\mathbf{x}$.

Lemma 5 (Payment transformation by [Brustle *et al.*, 2017]). *Let $(\mathbf{x}, \mathbf{t})$ be a mechanism that is BIC, interim IR, and ex-ante BB with a nonnegative payment rule (i.e., $t_i(\cdot) \geq 0$ for $i \in \mathcal{N}$ and $t_j(\cdot) \geq 0$ for $j \in \mathcal{M}$). Then there exists another payment rule $\mathbf{t}'$ such that:*

1. Allocation preserved: $(\mathbf{x}, \mathbf{t}')$ uses the same allocation rule $\mathbf{x}$.
2. BIC & interim IR preserved: $(\mathbf{x}, \mathbf{t}')$ remains BIC and interim IR, and each agent's interim expected payment is unchanged.
3. Ex-post SBB: For every profile $(\mathbf{v}, \mathbf{c})$,

$$\sum_{j \in \mathcal{M}} t'_j(\mathbf{v}, \mathbf{c}) = \sum_{i \in \mathcal{N}} t'_i(\mathbf{v}, \mathbf{c}).$$

Moreover, the converse holds: any BIC, interim IR, SBB mechanism with nonnegative payments can be converted to a BIC, interim IR, ex-ante BB mechanism with the same allocation by an inverse linear transformation of payments.

Remark 6 (On “same allocation”). Lemma 5 preserves the ex-post allocation rule $\mathbf{x}(\cdot)$ profile-by-profile; in particular, interim allocation probabilities remain identical. Since all of our objectives in this paper depend only on the allocation (transfers do not enter the objective), the designer's value is unchanged by the transformation.

Optimality. We now show that the DSIC & ex-ante BB optimum characterized in Section 4 yields an *optimal* BIC & ex-post BB mechanism after the transformation in Lemma 5.

Theorem 7 (Optimal BIC & ex-post BB). *In the additive setting under regular distributions, let $(\mathbf{x}^*, \mathbf{t}^*)$ be the DSIC, IR, ex-ante BB mechanism from Theorem 4. Then there exists a payment rule $\hat{\mathbf{t}}$ such that $(\mathbf{x}^*, \hat{\mathbf{t}})$ is BIC, IR, and ex-post BB (indeed SBB), achieves the same expected utility as $(\mathbf{x}^*, \mathbf{t}^*)$, and is optimal among all BIC, IR, ex-post BB mechanisms.*

Corollary 8 (Equality of optimal values). *In the additive setting with regular distributions,*

$$OPT_{\text{BIC, ex-ante BB}} = OPT_{\text{BIC, ex-post BB}} = \mathbb{E}[u(\mathbf{x}^*)],$$

where $\mathbf{x}^*$ is the allocation from Theorem 4.

5 Beyond Additive Utility

When the utility function $u(\cdot)$ is not additive but combinatorial, e.g., submodular, the previous framework does not work. In particular, even the offline optimization problem, assuming truthful bids, is well-known to be NP-hard to solve for many classes of functions. Thus, we aim to derive a mechanism that approximates the optimal mechanism.

Let S be the set of selected agents given an allocation vector $\mathbf{x}$. Accordingly, we consider the following class of separable functions

$$u(S) = v_p(S_p) - v_m(S_m),$$

where S_p and S_m are the sets of *selected* plus and minus agents, respectively. Further, we assume v_p is monotone increasing submodular and v_m is monotone decreasing supermodular. This is practical since the marginal decrease in the audience experience can grow significantly as we assign more advertisements, even though a few of them might be fine. From theoretical perspectives, this is reasonable since if we alternatively assume v_m is submodular, then the resulting function is the sum of a submodular v_p and a supermodular $-v_m$, which admits no polynomial approximation algorithm.⁶

Note that maximizing u is equivalent to maximizing

$$u(S) + v_m(\mathcal{M}) = v_p(S_p) + (v_m(\mathcal{M}) - v_m(S_m)).$$

In particular, we can write $u_m(S_m) = v_m(\mathcal{M}) - v_m(S_m)$ as a monotone increasing submodular function defined on the set of excluded minus agents S_m . Thus, redefining S_m as the set of excluded minus agents for notational convenience, we will consider the utility function $u(S) = v_p(S_p) + u_m(S_m)$ for monotone increasing nonnegative submodular functions v_p and u_m where S_p denotes the set of selected plus agents, and S_m the set of excluded minus agents.

5.1 Budget-enumeration Posted Pricing

We first introduce our main mechanism, the *budget-enumeration posted pricing* (BEPP) mechanism, presented in Algorithm 2. It takes a discretization parameter $\varepsilon > 0$ which controls the approximation factor as well as its time complexity. The following theorem formalizes the performance of the BEPP mechanism.

⁶In general, maximizing a monotone supermodular function always requires an exponential number of queries.

Algorithm 2 BEPP Mechanism

Input: Parameter $\varepsilon > 0$

- 1: Compute $B_{\max}$, the maximum expected revenue extractable from the minus agents via Myerson's optimal auction.
- 2: Find minimal M such that $(1 + \varepsilon)^M > B_{\max}$.
- 3: Define $\mathcal{B} \leftarrow \{0\} \cup \{(1 + \varepsilon)^i : 0 \leq i \leq M - 1\}$.
- 4: **for** each budget target $B_k \in \mathcal{B}$ **do**
- 5: Find $\mathbf{q}(B_k)$ from multilinear version of (5.1).
- 6: Find $\mathbf{r}(B_k)$ from multilinear version of (5.2).
- 7: $W^{\text{ml}}(B_k) \leftarrow u_p(\mathbf{q}(B_k)) + u_m(\mathbf{r}(B_k))$.
- 8: **end for**
- 9: Find the best budget $B^* \leftarrow \arg \max_{B_k \in \mathcal{B}} W^{\text{ml}}(B_k)$.
- 10: Set $t_i^p \leftarrow G_i^{-1}(q_i(B^*))$ for each $i \in \mathcal{N}$ and $t_j^m \leftarrow F_j^{-1}(r_j(B^*))$ for each $j \in \mathcal{M}$.
- 11: **return** posted prices $(\mathbf{t}^p, \mathbf{t}^m)$: pay t_i^p to plus agent i upon acceptance, and charge t_j^m from minus agent j upon acceptance.

Theorem 9. For any $\varepsilon > 0$, the BEPP mechanism is DSIC, IR, and ex-ante BB, runs in $\text{poly}(n, m, 1/\varepsilon, \log B_{\max})$ time, and achieves a $(1 - 1/e - O(\varepsilon))$ -approximation to the optimal Bayesian mechanism for separable submodular utility and regular distributions.

Proof overview. We now provide the overall proof steps to show that the BEPP mechanism is a constant-factor approximation of the optimal Bayesian mechanism for separable submodular utilities. The proof proceeds by defining a series of relaxations for the plus-agent and minus-agent subproblems, establishing an upper bound on the optimal welfare, and then showing that the expected welfare of our posted-price mechanism is within a constant factor of this bound.

We first formally define the two key relaxations of a set function, following the approaches used by [Balkanski and Hartline, 2016]. Let $u : 2^{\mathcal{O}} \rightarrow \mathbb{R}_{\geq 0}$ be a set function on a ground set $\mathcal{O}$. Let $\mathbf{p} = (p_k)_{k \in \mathcal{O}}$ be a vector of marginal probabilities. We now define the multilinear extension of u .

Definition 10 (Multilinear Extension). The **multilinear extension** of u , denoted $U(\mathbf{p})$, is the expected value of the function when each element $k \in \mathcal{O}$ is included in a set S independently with probability p_k .

$$U(\mathbf{p}) = \mathbb{E}_{S \sim \mathbf{p}} [u(S)] = \sum_{S \subseteq \mathcal{O}} u(S) \prod_{k \in S} p_k \prod_{k \notin S} (1 - p_k).$$

The key feature in analyzing approximation factors derived from multilinear extension is its relation to the following concave relaxation of u .

Definition 11 (Concave Closure). The **concave closure** of u , denoted $U^+(\mathbf{p})$, is the maximum expected value of the function over any *correlated* distribution $\mathcal{D}$ on subsets of $\mathcal{O}$ with marginal probabilities $\mathbf{p}$.

$$U^+(\mathbf{p}) = \max_{\mathcal{D} \text{ with marginals } \mathbf{p}} \mathbb{E}_{S \sim \mathcal{D}} [u(S)]$$

Importantly, note here that the probability vectors can be correlated unlike the multilinear extension. The connection

between these two relaxations is the following notion of correlation gap.

Theorem 12 ([Agrawal et al., 2010]). For any non-negative, monotone submodular function $u(\cdot)$, its multilinear extension and concave closure satisfy:

$$\min_{\mathbf{p}} \frac{U(\mathbf{p})}{U^+(\mathbf{p})} \geq 1 - \frac{1}{e}$$

We structure the proof as a series of lemmas. First, we write concave closures of our original optimization problems for each side market. Define $\mathbf{q} = (q_1, \dots, q_n)$ as a vector that represents the marginal probability that plus agent i is selected, and similarly $\mathbf{r} = (r_1, \dots, r_m)$ as a vector that represents the probability that minus agent j is excluded. In particular, we consider the following program for the plus agent side for each B_k :

$$\begin{aligned} (\text{PLUS}) \quad & \max_{\mathbf{q}} \quad u_p(\mathbf{q}) \\ \text{s.t.} \quad & \sum_{i \in \mathcal{N}} q_i G_i^{-1}(q_i) \leq B_k, \end{aligned} \quad (5.1)$$

and the following for the minus agent side:

$$\begin{aligned} (\text{MINUS}) \quad & \max_{\mathbf{r}} \quad u_m(\mathbf{r}) \\ \text{s.t.} \quad & \sum_{j \in \mathcal{M}} (1 - r_j) F_j^{-1}(r_j) \geq B_k. \end{aligned} \quad (5.2)$$

Multilinear versions of the above programs require the probabilities to be independent, whereas concave closure allows them to be correlated.

We will compare each optimization problem with the expected utilities obtainable from each side market. Formally, we will consider a class of mechanisms $\text{MEC}(B)$ parameterized by B such that its total expected payment to plus agents is at most B , and its total expected revenue from minus agents is at least B . Given a budget B , we use the following notation:

- $W_p(B)$ and $W_m(B)$: the optimal expected utilities achievable by any DSIC and IR mechanism in $\text{MEC}(B)$, on the plus and minus side respectively. We write $W(B) = W_p(B) + W_m(B)$.
- $W_p^+(B)$ and $W_m^+(B)$: the optimal values of (PLUS) and (MINUS) under the concave closure.
- $W_p^{\text{ml}}(B)$ and $W_m^{\text{ml}}(B)$: the optimal values of (PLUS) and (MINUS) under the multilinear extension, with $\mathbf{q}(B)$ and $\mathbf{r}(B)$ denoting the corresponding optimal marginal probabilities.

Throughout, we assume without loss of generality (by rescaling costs and values) that $W(x) \leq (1 + \varepsilon)W(0)$ for all $x \in (0, 1]$.

The performance of any mechanism given a budget, including the optimal one, will be upper-bounded by the concave closure as shown by the following lemma:

Lemma 13 (Concave relaxation). Given $B > 0$ and a mechanism $\text{MEC}(B)$ that is DSIC, IR and ex-ante BB, its expected utility u_p from plus agents is upper bounded by $W_p^+(B)$ and u_m (that from minus agents) is upper bounded by $W_m^+(B)$.

Then, we consider multilinear versions of the programs (PLUS) and (MINUS), and solving them gives us prices to offer as depicted in Algorithm 2. Thanks to Theorem 12, these prices guarantee a constant-factor approximation to the concave closure given B .

Lemma 14 (Posted pricing given a budget). *For any budget B , offering prices $\mathbf{t}^p = G^{-1}(\mathbf{q}(B))$ to plus agents guarantees an expected utility of at least $(1 - 1/e) \cdot W_p^+(B)$ from plus agents, and offering $\mathbf{t}^m = F^{-1}(\mathbf{r}(B))$ to minus agents guarantees an expected utility of at least $(1 - 1/e) \cdot W_m^+(B)$ from minus agents.*

Therefore, one can conclude that if we know the expected budget used by the optimal mechanism, the corresponding posted pricing mechanism gives us the desired result. On the other hand, this is not known in advance, so we generate a finite list of candidate optimal budgets and select the candidate budget that maximizes the expected utility. Since this list cannot always include the optimal budget due to the continuity of the entire budget space, we require the following concavity of the optimal expected utility with respect to the budget.

Lemma 15 (Concavity of the optimal utility). *Under regular agent distributions, the plus-agent's optimal utility $W_p(B)$ and the minus-agent's optimal utility $W_m(B)$ are both concave functions of B .*

The above lemma results in the following lemma, implying that the error from the finiteness of the budget candidates can be made sufficiently small.

Lemma 16 (Error from budget discretization). *Let $W(B) = W_p(B) + W_m(B)$ be the total utility function given B . Let B_k and B_{k+1} be consecutive grid points in $\mathcal{B}$ such that $B_k \leq B^* \leq B_{k+1}$. Then, we have $W(B_k) \geq (1 - \varepsilon)W(B^*)$.*

Combining the lemmas, let B_{OPT} be the expected budget used by the optimal Bayesian mechanism, and let B_k, B_{k+1} be consecutive grid points bracketing it ($B_k \leq B_{\text{OPT}} \leq B_{k+1}$). Then

$$\begin{aligned} W^{\text{ml}}(B^*) &\geq W^{\text{ml}}(B_k) \\ &\geq (1 - 1/e) \cdot W^+(B_k) && \text{(Theorem 12)} \\ &\geq (1 - 1/e) \cdot W(B_k) && \text{(Lemma 13)} \\ &\geq (1 - 1/e)(1 - \varepsilon) \cdot W(B_{\text{OPT}}) && \text{(Lemma 16)} \\ &= (1 - 1/e)(1 - \varepsilon) \cdot \text{OPT}, \end{aligned}$$

which yields the claimed $(1 - 1/e - O(\varepsilon))$ -approximation.

6 Conclusion

In this paper, we introduced and analyzed a novel two-sided mechanism design framework for self-financing events, where a designer's goal is to maximize an exogenous utility function by curating a set of participants. This problem requires balancing the recruitment of costly, high-impact plus agents with revenue-generating minus agents who may detract from the event's quality, all under a budget-balanced constraint.

Our analysis addresses this design problem primarily within the Bayesian setting. For the case of additive utilities, we fully characterize the optimal mechanisms that are

incentive-compatible, IR, and BB. Extending our results to more general objectives, we then provide constant-factor approximation mechanisms for utility functions that are separable and submodular.

Key avenues for future work include extending our analysis beyond submodular separable utility functions to model the complex interplay between participants. Broadening the scope of prior-free results for this framework, by providing robust guarantees for submodular utilities or by relaxing the large market assumption, could also enhance its practical applicability.

Acknowledgements

This work is partially supported by DARPA expMath, NSF AF:Small #2218678, NSF AF:Small #2114269, Army-Research Laboratory (ARL) #W911NF2410052, and MURI on Algorithms, Learning and Game Theory.

References

- [Agrawal *et al.*, 2010] Shipra Agrawal, Yichuan Ding, Amin Saberi, and Yinyu Ye. Correlation robust stochastic optimization. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 1087–1096. SIAM, 2010.
- [Anari *et al.*, 2014] Nima Anari, Gagan Goel, and Afshin Nikzad. Mechanism design for crowdsourcing: An optimal $1-1/e$ competitive budget-feasible mechanism for large markets. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 266–275. IEEE, 2014.
- [Balkanski and Hartline, 2016] Eric Balkanski and Jason D Hartline. Bayesian budget feasibility with posted pricing. In *Proceedings of the 25th International Conference on World Wide Web*, pages 189–203, 2016.
- [Balkanski *et al.*, 2022] Eric Balkanski, Pranav Garimidi, Vasilis Gkatzelis, Daniel Schoepflin, and Xizhi Tan. Deterministic budget-feasible clock auctions. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2940–2963. SIAM, 2022.
- [Balseiro *et al.*, 2025] Santiago Balseiro, Kshipra Bhawalkar, Yuan Deng, Zhe Feng, Jieming Mao, Aranyak Mehta, Vahab Mirrokni, Renato Paes Leme, Di Wang, and Song Zuo. Position auctions in ai-generated content. *arXiv preprint arXiv:2506.03309*, 2025.
- [Blumrosen and Dobzinski, 2014] Liad Blumrosen and Shahr Dobzinski. Reallocation mechanisms. *arXiv preprint arXiv:1404.6786*, 2014.
- [Blumrosen and Dobzinski, 2021] Liad Blumrosen and Shahr Dobzinski. (almost) efficient mechanisms for bilateral trading. *Games and Economic Behavior*, 130:369–383, 2021.
- [Brustle *et al.*, 2017] Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 589–590, 2017.

- [Chawla *et al.*, 2019] Shuchi Chawla, Jason D Hartline, and Balasubramanian Sivan. Optimal crowdsourcing contests. *Games and Economic Behavior*, 113:80–96, 2019.
- [Deng *et al.*, 2022] Yuan Deng, Jieming Mao, Balasubramanian Sivan, and Kangning Wang. Approximately efficient bilateral trade. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 718–721, 2022.
- [Deng *et al.*, 2025] Yuan Deng, Jieming Mao, Balasubramanian Sivan, Kangning Wang, and Jinzhao Wu. Approximately efficient bilateral trade with samples. In *Proceedings of the 26th ACM Conference on Economics and Computation*, pages 206–223, 2025.
- [DiPalantino and Vojnovic, 2009] Dominic DiPalantino and Milan Vojnovic. Crowdsourcing and all-pay auctions. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 119–128, 2009.
- [Dubey *et al.*, 2024] Avinava Dubey, Zhe Feng, Rahul Kidambi, Aranyak Mehta, and Di Wang. Auctions with llm summaries. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 713–722, 2024.
- [Duetting *et al.*, 2024] Paul Duetting, Vahab Mirrokni, Renato Paes Leme, Haifeng Xu, and Song Zuo. Mechanism design for large language models. In *Proceedings of the ACM on Web Conference 2024*, pages 144–155, 2024.
- [Eilat and Pauzner, 2021] Ran Eilat and Ady Pauzner. Bilateral trade with a benevolent intermediary. *Theoretical Economics*, 16(4):1655–1714, 2021.
- [Hajiaghayi *et al.*, 2024] MohammadTaghi Hajiaghayi, Sébastien Lahaie, Keivan Rezaei, and Suho Shin. Ad auctions for llms via retrieval augmented generation. *arXiv preprint arXiv:2406.09459*, 2024.
- [Hajiaghayi *et al.*, 2025] Ilya Hajiaghayi, MohammadTaghi Hajiaghayi, Gary Peng, and Suho Shin. Gains-from-trade in bilateral trade with a broker. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4827–4860. SIAM, 2025.
- [Han *et al.*, 2023] Kai Han, You Wu, He Huang, and Shuang Cui. Triple eagle: Simple, fast and practical budget-feasible mechanisms. *Advances in Neural Information Processing Systems*, 36:33894–33911, 2023.
- [Jalaly and Tardos, 2021] Pooya Jalaly and Éva Tardos. Simple and efficient budget feasible mechanisms for monotone submodular valuations. *ACM Transactions on Economics and Computation (TEAC)*, 9(1):1–20, 2021.
- [Klumper and Schäfer, 2022] Sophie Klumper and Guido Schäfer. Budget feasible mechanisms for procurement auctions with divisible agents. In *International Symposium on Algorithmic Game Theory*, pages 78–93. Springer, 2022.
- [Kuang *et al.*, 2023] Zhonghong Kuang, Weiran Shen, and Fan Wu. Profit-maximizing mechanism in bilateral trade with interdependent valuations. *Available at SSRN 4474002*, 2023.
- [McAfee, 1992] R Preston McAfee. A dominant strategy double auction. *Journal of economic Theory*, 56(2):434–450, 1992.
- [Myerson and Satterthwaite, 1983] Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of economic theory*, 29(2):265–281, 1983.
- [Singer, 2010] Yaron Singer. Budget feasible mechanisms. In *2010 IEEE 51st Annual Symposium on foundations of computer science*, pages 765–774. IEEE, 2010.