Mathematics of Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

From Monetary to Nonmonetary Mechanism Design via Artificial Currencies

Artur Gorokh, Siddhartha Banerjee, Krishnamurthy Iyer

MATHEMATICS OF OPERATIONS RESEARCH
Articles in Advance, pp. 1–21
ISSN 0364-765X (print), ISSN 1526-5471 (online)
http://pubsonline.informs.org/journal/moor

From Monetary to Nonmonetary Mechanism Design via Artificial Currencies

Artur Gorokh,[a] Siddhartha Banerjee,[b] Krishnamurthy Iyer[c]

[a] Center for Applied Mathematics, Cornell University, Ithaca, New York 14850; [b] Operations Research and Information Engineering, Cornell University, Ithaca, New York 14850; [c] Industrial and Systems Engineering, University of Minnesota, Minneapolis, Minnesota 55455

Contact: ag2282@cornell.edu (AG); sbanerjee@cornell.edu, https://orcid.org/0000-0002-8954-4578 (SB); kriyer@umn.edu, https://orcid.org/0000-0002-5538-1432 (KI)

Received: July 31, 2017
Revised: December 8, 2018
Accepted: January 25, 2020
Published Online in Articles in Advance: February 11, 2021

MSC2000 Subject Classification: Primary: 91A06, 91A10
OR/MS Subject Classification: Primary: bidding/auctions

https://doi.org/10.1287/moor.2020.1098

Abstract. Nonmonetary mechanisms for repeated allocation and decision making are gaining widespread use in many real-world settings. Our aim in this work is to study the performance and incentive properties of simple mechanisms based on artificial currencies in such settings. To this end, we make the following contributions: For a general allocation setting, we provide two black-box approaches to convert any one-shot monetary mechanism to a dynamic nonmonetary mechanism using an artificial currency that simultaneously guarantees vanishing gains from nontruthful reporting over time and vanishing losses in performance. The two mechanisms trade off between their applicability and their computational and informational requirements. Furthermore, for settings with two agents, we show that a particular artificial currency mechanism also results in a vanishing price of anarchy.

Funding: This work was supported by the National Science Foundation Division of Electrical, Communications and Cyber Systems [Grant ECCS-1847393], Division of Mathematical Sciences [Grant DMS-1839346], and Division of Civil, Mechanical and Manufacturing Innovation [Grants CMMI-1462592 and CMMI-1633920], and the Army Research Laboratory [Grant W911NF-17-1-0094].

Keywords: dynamic mechanism design • artificial currencies • budget constraints

1. Introduction

Today, many real-world settings require nonmonetary mechanisms for allocating scarce resources or making collective decisions over a period of time, among agents with time-varying preferences. As an example, consider the problem of allocating computing resources in a university cluster: over time, the limited processors must be divided between different users with diverse memory/processing/software requirements, different urgency levels, etc. Other examples include distributing food among food banks, course seats among students, parking spots and vacation days among employees, collective decision making within organizations, sharing cooperatives, etc. All these settings share two characteristics: (1) resources are allocated among a set of agents repeatedly over time, and (2) the use of money as a means to elicit the agents' preferences is undesirable and/or even prohibited.

The lack of monetary payments as a means to incentivize agents makes mechanism design in such settings more challenging. For instance, in the absence of payments, agents naturally seek to inflate their reported value for resources. In repeated settings, however, future allocations provide a means to incentivize agent behavior. For the case of single-item allocation, this idea has been explored theoretically in a set of recent works (Balseiro et al. [5], Guo et al. [25]). However, the resulting mechanisms are complex and are hard to extend to more general settings.

In contrast, a class of simple nonmonetary mechanisms, namely, mechanisms based on *artificial currencies*, has gained recent attention because of their successful use in practice (e.g., in food banks; Prendergast [37]). Such mechanisms involve endowing each agent with a budget of an artificial currency and organizing a static monetary mechanism in each period with payments in the artificial currency. Despite their success in practice, not much is known about the incentive properties of such mechanisms. *The main contribution of our work is to provide the theoretical foundations for the use of such simple mechanisms in practice.*

Specifically, our goal in this work is twofold: first, we seek to connect the incentive and utility properties of nonmonetary mechanisms for repeated allocation settings to those of monetary mechanisms in static settings. Given the large body of literature on static monetary mechanism design, this connection provides a valuable benchmark for the performance of nonmonetary mechanisms. Second, we seek to exploit this connection to design simple mechanisms that can be employed in practice. In particular, our results give operational insights into implementing simple mechanisms based on artificial currencies.

1.1. Main Results and Road Map of This Paper

We consider a general repeated allocation setting over T periods and with n agents. In each period, each agent s has a private type drawn from some underlying distribution. Agent types can be multidimensional—for example, in the setting of combinatorial auctions, the type specifies an agent's value for any bundle of items. For most of the work, we assume that the type of agent s in each period is drawn independently (across time and agents) and identically (across time) from a distribution F_s (we relax this in Section 6).

We focus on direct-revelation mechanisms, which ask each agent to report her type at each period, and then choose an allocation from among a feasible set of allocations. This allocation decision can be based on the agents' current and past reports. Crucially, the mechanism cannot make monetary transfers to the agents or receive transfers from them.

Our objective is to construct a nonmonetary mechanism M that mirrors the properties of a given monetary mechanism M_{money} for the static setting (i.e., when the allocation is to be done once). In particular, we want M to provide the same incentives and ensure the same utility to the agents as M_{money}. However, doing this perfectly is in general impossible: for $T = 1$ and general type distributions F_s, no nonmonetary mechanism for single-item allocation can achieve the efficiency as guaranteed by the second-price auction (as a concrete example, consider a single-item setting and two agents with independent and identically distributed (i.i.d.) Uniform$[0, 1]$ values).

Faced with this difficulty, we relax our objective and instead seek a nonmonetary mechanism that approximately preserves the incentive and utility guarantees of a monetary mechanism. Formally, we say a nonmonetary mechanism M is an (α, β)-approximation of a given *Bayesian incentive compatible* (*BIC*) mechanism M_{money} if it simultaneously guarantees the following two properties (see Definition 1):

1. For any agent, the gains from unilaterally deviating from truthful reporting under M is at most αT. Because M_{money} is BIC, this ensures that M approximately matches the incentive properties of M_{money}.
2. Assuming truthful reporting, the mechanism M guarantees the same expected utility for every agent as M_{money} (excluding payments), up to an additive loss βT.

With this definition in place and given the general impossibility of a $(0, 0)$-approximate mechanism, one question is whether there exists a mechanism that is (α, β)-approximate, where both α and β are *vanishing*, that is, $o(1)$, with respect to T. This question is interesting from two perspectives. First, it suggests that as the time horizon T increases, such a mechanism becomes a better approximation of M_{money}. Second, and more importantly, under such a guarantee (specifically $\alpha = o(1)$), there are strong behavioral justifications for why an agent would not unilaterally deviate from truthful reporting despite the associated (small) gain (see Section 2.4 for more details).

Our work answers this question in the affirmative. In particular, we construct two *black-box* approaches (NonMonetary Black-box Reduction (NMBER) and Repeated Endowed All-Pay (REAP)) that each take any monetary mechanism M_{money} and return a nonmonetary mechanism M that is (α, β)-approximate for $\alpha = O(\sqrt{\log T/T})$ and $\beta = O(1/T)$. The two approaches differ in the assumptions on the settings, as well as their informational and computational requirements.

In more detail, our main results are as follows:

1. Under the assumption of *excludability* (see Assumption 1), where the mechanism is able to exclude an agent without affecting other agents, we provide a simple nonmonetary black-box reduction (which we refer to as the NMBR mechanism; see Section 3). Informally, our recipe comprises replacing money with an artificial currency for payments in the monetary mechanism, coupled with tractable procedures for setting initial endowments of the artificial currency and simulating bids of budget-depleted agents. More specifically, we provide a black-box reduction that, given a monetary mechanism, produces a nonmonetary mechanism, for which we prove the desired (α, β)-approximation guarantee (Theorems 1 and 2). Furthermore, we show that this guarantee persists even when the principal has access only to a finite number of type reports from the agent, and that the computational burden of running the NMBR mechanism is comparable with that of running M_{money} T times.
2. In Section 4, we drop the assumption of excludability and construct an alternative black-box mechanism, the repeated endowed all-pay, or REAP, mechanism. The idea in REAP is again endowing agents with artificial currency, but instead of directly running M_{money} for allocations and payments, we make payments of agents depend only on their reports and not the outcome of the allocation. We show that REAP recovers the

incentive and performance guarantees of NMBR (Theorem 3); however, the mechanism requires exact knowledge or agents' type distributions and may be computationally intractable.

3. In addition, we show that the incentive guarantees we provide for the REAP mechanism can be strengthened for the case when there are two agents and the Vickrey-Clarke-Groves (VCG) mechanism is employed as M_{money}. More precisely, in this setting, we show that utility profile at any equilibrium is close to that of M_{money} for the case of two agents (Theorem 4). We also leverage this result to show the *existence* of a (potentially complex) $(0, o(1))$-approximate mechanism (Corollary 3).

1.1.1. Technical Contributions. We prove our results by showing a connection between incentives of an agent in the monetary setting and those of an agent constrained by a budget of artificial currency. We do this via three steps: (i) We consider the perspective of agent s playing against truthful opponents, and for this agent construct an auxiliary problem that can be viewed as allowing the agent to violate the budget constraint on some sample paths, but still satisfy it in expectation. (ii) Via concentration arguments, we show that the performance of any strategy in the original mechanism is close to that in the auxiliary problem, and thus whenever truthful reporting is approximately optimal in the auxiliary problem, it is approximately optimal in the mechanism as well. (iii) Finally, we connect incentives in the auxiliary problem to those in the monetary setting; in particular, we show that incentive compatibility (IC) of M_{money} in the monetary setting implies approximate optimality of truthful reporting in the auxiliary problem.

The auxiliary problem also proves useful for proving our price of anarchy bounds, in combination with standard arguments that are reminiscent of the smooth games framework (Roughgarden [39]). Moreover, the modular nature of our proof allows us to discuss applicability of our technique to settings more general than the one considered in this paper. We discuss several such potential extensions in Section 6.

1.2. Related Work

Our setting sits at the intersection of work on mechanisms without money and dynamic mechanism design; both topics have attracted significant interest in recent years. We briefly summarize work that is closest to our setting.

1.2.1. Dynamic Mechanism Design. Dynamic mechanism design focuses on extending the theory of mechanisms for single-period settings (Milgrom [33], Myerson [34]) to repeated allocation settings. The difficulty in doing so arises because of factors that couple auctions across time, for example, incomplete information and learning over time (Devanur et al. [15], Iyer et al. [28], Kanoria and Nazerzadeh [30], Nekipelov et al. [36]), cross-period combinatorial constraints including limits imposed by budgets (Balseiro et al. [4], Gummadi et al. [23], Leme et al. [32], Nazerzadeh et al. [35]), stochastic fluctuations in the underlying setting (Bergemann and Said [6], Bergemann and Välimäki [7], Gershkov and Moldovanu [19]), etc. In our setting, the cross-period coupling arises essentially because of budget constraints. Similar repeated auctions with budget constraints have been considered by Gummadi et al. [23] and Balseiro et al. [4], wherein the authors use mean-field approaches to circumvent the difficulty in deriving equilibrium behavior of agents. Our results are similar in spirit in that we eschew exact IC for approximate IC; however, our technique of approximating the dynamic setting via an auxiliary static game and then proving closeness between the two as T scales is novel compared with existing approaches.

1.2.2. Mechanism Design Without Money. This is a broad area of study, encompassing diverse topics ranging from matching theory to social choice, that broadly considers strategic allocation in settings where money is not permitted for various reasons. Most of this literature deals primarily with single-period settings and typically involves working with alternate notions of equilibria. An approximation-based approach to such settings was proposed by Procaccia and Tennenholtz [38], who used the specific example of a facility location game; this approach was subsequently explored by many others (Dughmi and Ghosh [17], Guo and Conitzer [24]). Several alternate approaches have also been proposed, including ones based on verifiability (Brânzei and Procaccia [8]), proof-of-work ("money burning"; Hartline and Roughgarden [26]), and two-tiered resource redistribution (Cavallo [11]).

In the case of repeated allocation without money, a notable line of work is that by Guo et al. [25] and its subsequent refinement by Balseiro et al. [5]. These works consider a model identical to ours but with a single item and symmetric agents, and under discounted infinite-horizon settings. The former uses variations of the AGV mechanism (d'Aspremont and Gérard-Varet [14]) to achieve perfect IC at the cost of efficiency. The latter work builds on this model to develop a novel BIC mechanism that achieves vanishing welfare loss as discount factor goes to one. In contrast, we consider much more general allocation settings under the finite horizon and focus on using simple mechanisms that utilize existing monetary mechanisms with artificial currency. In this context, we obtain welfare and incentive guarantees that vanish with the horizon length.

1.2.3. Artificial Currency Mechanisms. These are a particular subset of nonmonetary mechanisms that have attracted a lot of recent interest, in part because of recent successful implementations in real-world settings (Budish et al. [10], Prendergast [37], Walsh [40]). Our work follows in the line of several recent papers in attempting to establish a theoretical foundation for such mechanisms. Among these, the closest to ours are the works of Budish [9] and Jackson and Sonnenschein [29]; we describe these now in more detail.

Budish [9] studies the use of artificial currency mechanisms in the context of static combinatorial assignment problems under (arbitrary) ordinal preferences. As in our work, the incentive compatibility constraint is relaxed to satisfy other design objectives, namely, Pareto optimality, approximate market clearing, and envy-freeness; this is analogous to our use of approximate efficiency guarantees under approximate IC constraints, as in Definition 1. Similar guarantees are also established in static settings for additive valuations via maximizing Nash welfare by Cole and Gkatzelis [13]. In contrast to these static mechanisms, our approach applies primarily for dynamic settings by exploiting future allocations to ensure approximate incentive compatibility.

Jackson and Sonnenschein [29] (henceforth, JS) are also concerned with one-shot allocation problems, but unlike Budish [9], they consider settings where agents simultaneously participate in multiple resource allocation problems, with i.i.d. types in each instance. For this setting, JS provide a mechanism that guarantees near-optimal welfare at any equilibrium in the mechanism; this is done by essentially endowing agents with separate budgets for reporting each possible type across instances. Our approach also shares this technique of linking separate problem instances to enforce incentives; however, our work differs from JS's in three significant dimensions: dynamics, scalability, and incentive guarantees. The primary difference is that we consider repeated allocation settings, rather than simultaneous allocations as in the work of JS. This temporal aspect makes the analysis of incentives more challenging in our setting. Next, because JS make use of a separate budget for every possible report, the number of linked instances required for nontrivial efficiency guarantees is prohibitively large for multidimensional type spaces (e.g., combinatorial auctions). In contrast, our mechanisms endow every agent with just a single artificial currency budget, and, as a consequence, the number of periods needed to guarantee near-optimal efficiency is independent of size of the type space. In particular, our mechanism performs well even if there are not enough periods for an agent to sample all of her types. Such an advantage comes at the cost of admitting weaker incentive guarantees—the approximate efficiency is achieved under the particular ϵ-equilibrium of truthful reporting (with the per-round gains from nontruthful bidding going to zero for large T), whereas the guarantees given by JS hold under any Bayes–Nash equilibrium (BNE). We note though that in the case of two agents, we are able to match the incentive guarantees of JS using our simpler single-currency mechanism.

1.2.4. Ex Ante Relaxation. A technique we use extensively throughout our paper is to first solve a relaxed version of the problem, in which agents are to satisfy their budget constraints in expectation, and then use this result to establish guarantees for the original problem with ex post constraints. A similar approach was used by Alaei [1] in a monetary setting to establish a black-box reduction from a single-agent to a multiagent mechanism with an item supply constraint. The same technique was later adopted by Chawla and Miller [12] in the context approximating revenue for selling multiple items to several heterogeneous buyers. Another related technique used in revenue maximization problems is that of a correlation gap; see Yan [41].

1.2.5. Dynamic Bidding Under Budget Constraints. Finally, another related line of literature analyzes bidding in repeated auctions with budget constraints, in the context of advertising markets. In this setting, Balseiro and Gur [3] consider the problem of regret minimization from the bidders' perspective and demonstrate strategies that constitute approximate equilibrium. Karande et al. [31] design a system to optimize bids in large, repeated ad auctions with budgets. Finally, Goel et al. [20] model various mechanisms to throttle the bids of budgeted agents across time to maximize revenue.

2. Model

2.1. Setting

We consider a repeated allocation setting with n agents. In each period[1] $t \in [T]$, each agent s has a type $\theta_s^t \in \Theta_s$, drawn from a distribution F_s. These draws are i.i.d. across periods and independent across agents. We let

$\theta^t = (\theta_s^t : s \in [n])$ denote the type profile at time t, and let $\theta_{-s}^t = (\theta_q^t : q \in [n], q \neq s)$. Let $\Theta = \times_{s\in[n]}\Theta_s$ denote the set of type profiles.

At each time $t \in [T]$, a principal chooses an allocation $X^t \in \mathcal{X}$, where $\mathcal{X}$ denote the set of feasible allocations. (We describe the principal's choice of allocation below.) Let $v_s(\theta_s^t, X^t)$ denote the utility that the agent s receives from allocation X^t when having type θ_s^t. We allow for v_s to take negative values, but assume that under any allocation and any type, the absolute value of utility of agents is uniformly bounded by $v_{\max}$.

For our results in Section 3, we require that the allocation setting satisfy the following *excludability* assumption.

Assumption 1 (Excludability). *For any feasible allocation X and any $S \subseteq [n]$, there exists a feasible allocation $X|_S$ such that for all θ, we have $v_s(\theta_s, X|_S) = v_s(\theta_s, X)$ for $s \in S$ and $v_s(\theta_s, X|_S) = 0$ for $s \notin S$.*

The excludability assumption holds in many centralized allocation problems, in particular, combinatorial assignment settings (where it is sometimes referred to as the *downward-closure* property). Nevertheless, this is a restrictive assumption that many settings of interest do not satisfy: for example, excludability need not hold in bilateral trade settings, or in the case of provision of public nonrival goods. (See Appendix B for a more detailed description of settings that do or do not satisfy excludability.)

2.2. Nonmonetary Mechanisms and Agents' Strategies and Utilities

We focus on settings where the principal employs a nonmonetary, *direct-revelation* mechanism to select the allocation X^t at each time t. Formally, a nonmonetary direct-revelation mechanism M requires each agent s to submit a report $\hat{\theta}_s^t \in \Theta_s$ at each time t. Let $\mathcal{H}^t = \{(X^\tau, \hat{\theta}_1^\tau, \ldots, \hat{\theta}_n^\tau)\}_{\tau<t}$ denote the public history, that is, the set of past allocations and reports. Subsequent to obtaining the reports $\hat{\theta}^t = (\hat{\theta}_s^t : s \in [n])$, the mechanism M selects an allocation $X^t = X^t(\hat{\theta}^t, \mathcal{H}^t; M)$ at each time t. (Note that there are no monetary payments from the agents.)

A strategy A_s of an agent s specifies in each period t a (possibly random) report $\hat{\theta}_s^t$ based on her current type θ_s^t, her types $\{\theta_s^\tau\}_{\tau<t}$ in periods $\tau < t$, and the public history $\mathcal{H}^t$. We denote the strategy of truthful reporting by Tr; that is, when agent s follows the strategy Tr, her reports are $\hat{\theta}_s^t = \theta_s^t$ for all $t \in [T]$. Let Tr_{-s} denote the case where all agents other than agent s are reporting their types truthfully.

For a strategy profile $A = \{A_s : s \in [n]\}$, let $U_s(A; M)$ denote the total utility obtained by agent s over T periods:

$$U_s(A; M) = \sum_{t=1}^{T} v_s\big(\theta_s^t, X^t\big),$$

where $X^t = X^t(\hat{\theta}^t, \mathcal{H}^t; M)$. Observe that $U_s(A; M)$ is a random variable that depends on the specific realization of the agents' types (we omit θ from arguments to simplify our notation), their reports, and the (possibly random) allocation. We let $u_s(A; M) = \mathbb{E}[U_s(A; M)]$ denote the expected utility of agent s under strategy profile A and mechanism M, where the expectation is over all the aforementioned randomness. Abusing notation, we write $U_s(A_s; M)$ and $u_s(A_s; M)$ for $U_s(A_s, \mathsf{Tr}_{-s}; M)$ and $u_s(A_s, \mathsf{Tr}_{-s}; M)$, respectively.

2.3. Static Setting and Monetary Mechanisms

To design and benchmark nonmonetary mechanisms for the setting described above, we consider monetary mechanisms for a related one-shot allocation setting. We briefly describe this setting below.

Formally, consider a one-shot allocation setting with n agents, where each agent s has a private type θ_s drawn independently from the distribution F_s. A direct-revelation mechanism M_{money} is defined in terms of its allocation rule $X(\hat{\theta}_1, \ldots, \hat{\theta}_n; M_{\text{money}}) \in \mathcal{X}$ and payment rules $\{C_s(\hat{\theta}_1, \ldots, \hat{\theta}_n; M_{\text{money}}) : s \in [n]\}$. Specifically, given the agents' report $\hat{\theta} = (\hat{\theta}_1, \ldots, \hat{\theta}_n)$, the mechanism chooses an allocation $X(\hat{\theta}; M_{\text{money}})$ and charges each agent s a monetary payment of $C_s(\hat{\theta}; M_{\text{money}})$. We assume that the agents have quasi-linear utilities, that is, the utility of the agent s is given by $v_s(\theta_s, X(\hat{\theta}; M_{\text{money}})) - C_s(\hat{\theta}; M_{\text{money}})$. We also assume that all monetary mechanisms employed in this paper have nonnegative payments $C_s(\hat{\theta}; M_{\text{money}}) \geq 0$ for all $\hat{\theta}$.

We say a direct-revelation mechanism M_{money} is BIC if reporting truthfully is a Bayes–Nash equilibrium; that is, for each agent s and for all θ_s, we have

$$\theta_s \in \arg\max_{\hat{\theta}_s} \; \mathbb{E}\Big[v_s\big(\theta_s, X\big(\hat{\theta}_s, \theta_{-s}; M_{\text{money}}\big)\big) - C_s\big(\hat{\theta}_s, \theta_{-s}; M_{\text{money}}\big) | \theta_s\Big].$$

Here, the expectation is over the types θ_{-s} of all agents other than s. For a BIC mechanism M_{money}, we let $u_s(\mathsf{Tr};M)$ denote the expected ex ante utility (excluding payment) obtained by an agent s in the truthful equilibrium, that is, $u_s(\mathsf{Tr};M_{\text{money}}) = \mathbb{E}[v_s(\theta_s, X(\theta;M_{\text{money}}))]$, where the expectation is over $\theta = (\theta_s : s \in [n])$.

Throughout, we focus on mechanisms M_{money} that satisfy the following *opt-out* condition, meant to capture voluntary participation.

Assumption 2 (Opt-Out Report). *Under mechanism M_{money}, every agent s has a report $\emptyset$ that guarantees zero payment $C_s(\emptyset, \hat{\theta}_{-s}; M_{\text{money}}) = 0$ and an allocation $X(\emptyset, \hat{\theta}_{-s}; M_{\text{money}})$ such that $v_s(\theta_s, X(\emptyset, \hat{\theta}_{-s}; M_{\text{money}})) = 0$ for any type θ_s and reports $\hat{\theta}_{-s}$.*

Finally, for a given monetary BIC mechanism M_{money}, we define $\bar{c}_s = \mathbb{E}_{\theta\sim F}[C_s(\theta, M_{\text{money}})]$ to be the expected payment charged to s in M_{money} in truthful equilibrium, and $c_s^{\max} = \max\{C_s(\theta; M_{\text{money}}) : \theta \in \Theta\}$ to be the maximum possible payment for s. We also define $R_s = \bar{c}_s / c_s^{\max}$.

2.4. Design Requirements

Returning to the repeated allocation setting, we seek to design nonmonetary mechanisms that approximate the utility and incentive characteristics of a given monetary BIC mechanism M_{money}. Formally, we consider the following definition for how well a nonmonetary mechanism M captures the incentives and utility profile of a given monetary mechanism M_{money}.

Definition 1. A mechanism M is an (α, β)-approximation of a monetary BIC mechanism M_{money} if it simultaneously guarantees the following:

1. Truthful reporting Tr is an α-equilibrium for M: For any agent s, assuming all other agents play truthfully, we have

$$\frac{u_s(\mathsf{Tr};M)}{T} \geq \sup_{A_s}\left(\frac{u_s(A_s;M)}{T}\right) - \alpha.$$

2. M guarantees the same utility profile as M_{money} up to an additive loss of βT under truthful reporting:

$$\frac{u_s(\mathsf{Tr};M)}{T} \geq u_s(\mathsf{Tr};M_{\text{money}}) - \beta.$$

Note that under the above definition, a $(0, 0)$-approximate mechanism guarantees exact incentive compatibility and attains the same utility profile (and, consequently, welfare) as M_{money}; however, as we mention above, this is too stringent, especially for small values of T. Instead, we seek to find an (α, β)-approximate mechanism where both α and β are vanishing (i.e., $o(1)$ with respect to T).

Such a guarantee captures a natural desideratum for a behavior model—that an agent is willing to deviate from truthful reporting only if the gains from doing so are significant when compared with the total utility the agent obtains. One reason for this could be that there is a cognitive burden of having to figure out a good deviation. More formally, suppose that to deviate from truth-telling in a profitable way (and/or to find out what the gains of such deviation would be), an agent needs to expend some computational effort c per turn. In this case, for large enough T, the computational overhead of deviating from truthful reporting would overwhelm the benefit that it could yield. A similar argument has been used in other works (see, e.g., Azevedo and Budish [2], Nekipelov et al. [36]).

3. Black-Box Reduction from Monetary Mechanisms

In this section, we present our black-box reduction that takes a monetary mechanism as an input and produces a nonmonetary mechanism that approximately matches its efficiency and incentive compatibility guarantees.

3.1. The NMBR Mechanism

We now describe our black-box reduction technique, which converts any chosen monetary BIC mechanism M_{money} to a nonmonetary mechanism with the desired approximation guarantees.

Formally, our mechanism is described in Algorithm 1. Informally, the mechanism proceeds as follows: it takes a monetary one-shot mechanism M_{money} as an input and endows each agent s with a budget of artificial currency. This endowment includes a small mutiplicative surplus δ_s over the agent's expected spending under the truthful equilibrium of M_{money} (i.e., $\bar{c}_s T$). We also initialize the set of active agents, $\mathsf{ACT} = [n]$. In each period, agents participate in the original mechanism M_{money} using the artificial currency instead of actual money for payments.

We then use the following procedure for handling bankrupt agents: when an agent s runs out of currency, we declare her inactive, $\mathsf{ACT} \leftarrow \mathsf{ACT} \setminus s$. From then on, we disregard reports of this agent and use independent samples for the agent's type when running M_{money}. Furthermore, we exclude inactive agents from the allocations produced by M_{money} (note that without the excludability assumption (Assumption 1), we would not be able to carry out this step).

The intuition behind introducing the budget surplus δ_s is that this ensures that, with high probability, truthful agents do not run out of credits before the final round T. Furthermore, the procedure of simulating the bankrupt agent removes the incentive agents might have to deplete their opponents, as upon depletion, an agent gets replaced with a replica; from the strategic perspectives of other agents, this replica is equivalent to the original agent.

Algorithm 1 (NMBR Mechanism)

Require: Static, monetary BIC mechanism M_{money}, expected agent payments under truthful reporting $\bar{c}_s$, sample access to agent type distributions F_s, surplus parameter δ_s

1: Allocate endowment of $B_s = (1+\delta_s)\bar{c}_s T$ credits to each agent s. Let $B_s^1 = B_s$ and initialize the set of active agents $\mathsf{ACT} = [n]$.
2: **for all** $t = 1$ to T **do**
3: For each agent $s \in \mathsf{ACT}$, get her report $\hat{\theta}_s^t$.
4: For each agent $s \notin \mathsf{ACT}$, sample a report $\hat{\theta}_s^t \sim F_s$ from her type distribution.
5: For each agent $s \in \mathsf{ACT}$, if $B_s^t \geq C_s(\hat{\theta}^t; M_{\text{money}})$, charge her the payment $C_s(\hat{\theta}^t; M_{\text{money}})$, and set $B_s^{t+1} = B_s^t - C_s(\hat{\theta}^t; M_{\text{money}})$. If $B_s^t < C_s(\hat{\theta}^t; M_{\text{money}})$, update $\mathsf{ACT} \leftarrow \mathsf{ACT} \setminus \{s\}$ and $B_s^{t+1} = 0$.
6: Let $X^t = X(\hat{\theta}^t; M_{\text{money}})$. Implement $X^t|_{\mathsf{ACT}}$ (refer Assumption 1)
7: **end for**

Our main result states that the above mechanism is approximately incentive compatible and guarantees a utility profile close to that under the M_{money} mechanism. Formally, we have the following theorem.

Theorem 1. *For a given BIC monetary mechanism M_{money}, consider the corresponding NMBR mechanism with $\delta_s = \sqrt{\frac{3\log T}{R_s T}}$. Then, for any strategy A_s, assuming all other agents play truthfully, we have*

$$\frac{u_s(\mathsf{Tr}; \text{NMBR})}{T} \geq \frac{u_s(A_s; \text{NMBR})}{T} - \sqrt{\frac{3\bar{c}_s c_{\max} \log T}{T}} - 2v_{\max} T^{-1}. \tag{1}$$

Additionally, the utility of each agent s satisfies

$$\frac{u_s(\mathsf{Tr}; \text{NMBR})}{T} \geq u_s(\mathsf{Tr}; M_{\text{money}}) - 2v_{\max} T^{-1}. \tag{2}$$

In other words, the NMBR mechanism is (α, β)-approximate, with $\alpha = O(\sqrt{\bar{c}_s \log T / T})$ and $\beta = O(T^{-1})$.

Note that although the number of agents n does not explicitly enter the expression for the bounds, there is an implicit dependency associated with the constant $\bar{c}_s$: as the number of the agents n increases, the average payment $\bar{c}_s$ decreases. In particular, this implies that, as n increases, the incentive guarantee α of the NMBR mechanism improves in the absolute (additive) sense. However, we note that as n increases, an agent's expected utility itself decreases, and the resulting bound on α may not be strong in a relative (multiplicative) sense.

We also note that although Theorem 1 establishes truthful reporting to be only approximately optimal ex ante, it is possible to demonstrate that similar guarantees hold in later rounds and degrade toward the end of the game. For more detail on this, see Appendix C.

To illustrate the mechanism and the main result, we provide a simple example, in which an item is allocated between n symmetric agents.

Example 1 (Single-Item Allocation Among n Agents)**.** Suppose there are n agents with uniformly distributed values $v_{it} \sim \mathsf{Unif}[0,1]$. Let M_{money} be the second-price auction. The expected payment of an agent s in M_{money} is given by $\bar{c}_s = \frac{(n-1)}{n(n+1)}$, with $c_s^{\max} = 1$, and $R_s = \frac{\bar{c}_s}{c_s^{\max}} = \frac{(n-1)}{n(n+1)}$.

The corresponding NMBR mechanism then proceeds as follows. Both agents are endowed with $(1+\sqrt{\frac{18\log T}{T}})\frac{(n-1)}{n(n+1)}T$ credits, and they participate in a second-price auction (with payments in credits) in each time period. If an agent s runs out of sufficient credits, then the mechanism excludes that agent, but in each subsequent time period includes an independent sample from $\mathsf{Unif}[0,1]$ as the agent's report. If at any time the winning report is from an agent with insufficient credits, then the mechanism does not allocate the item. With this, Theorem 1 guarantees the following:

$$\frac{u_s(\mathsf{Tr};\mathrm{NMBR})}{T} \geq \frac{u_s(A_s;\mathrm{NMBR})}{T} - \sqrt{\frac{3\log T}{nT}} - \frac{2}{T},$$
$$\frac{u_s(\mathsf{Tr};\mathrm{NMBR})}{T} \geq u_s(\mathsf{Tr};M_{\mathrm{money}}) - \frac{2}{T}.$$

Although gains from deviations decrease with n, because the agent's utility decreases with n as $O(1/n)$, the ratio of possible gains to total utility actually grows with n. Thus, for the asymptotical guarantees to stay meaningful, we need $T = O(n\log n)$.

Algorithm 1 along with Theorem 1 thus gives a general recipe for converting any one-shot monetary mechanism into a nonmonetary mechanism for repeated allocation with desired guarantees. From a computational viewpoint, the NMBR mechanism is equivalent to executing the original mechanism M_{money} with the caveat that we also need $\bar{c}_s$ as an input for each agent s; computing this, however, involves taking expectation over payments of M_{money} and may be costly. In Section 3.3, we circumvent this by showing that the incentive and welfare guarantees are preserved even if we instead use a *sample average* $\bar{c}_s^m$ over m simulated instances of the mechanism as input to NMBR (in particular, we show that $m = T$ samples are sufficient).

3.2. Proof of Theorem 1

In this section, we outline the overall strategy for proving Theorem 1 and establish some key lemmas for our proof. For brevity, we highlight the main ideas of our argument and defer some of the proof details to Appendix A.

The proof of Theorem 1 involves two parts: (1) proving the approximate incentive compatibility and (2) proving that NMBR achieves sublinear loss in welfare assuming agents report truthfully. The main challenge is in showing the former; the latter then follows from a simple concentration argument. We begin with the first part next.

To establish the approximate incentive compatibility of the NMBR mechanism for an agent s, we must compare the agent's utility under truthful reporting against her utility under the optimal strategy (assuming all other agents report truthfully). In the NMBR mechanism, an agent stops receiving any utility once she runs out of budget. Thus, in reasoning about her optimal strategy, an agent effectively has to satisfy a budget constraint on every sample path. This is a challenging decision problem to analyze. To circumvent this challenge, we first consider an *auxiliary problem* in which agent s is playing against truthful opponents with her budget constraint relaxed to be met only in expectation. We then show that (a) in the auxiliary problem, truthful reporting is approximately optimal, and that (b) expected utilities of truthful reporting in original game and the auxiliary problem are close.

Formally, suppose agent s participates repeatedly in the M_{money} mechanism for T rounds, with the other agents reporting their types truthfully. Given a strategy A_s, the expected utility of agent s in auxiliary problem is defined as

$$\widehat{u}_s(A_s) \triangleq \mathbb{E}\left[\sum_{t=1}^{T} v_s\left(\theta_s^t, X\left(\hat{\theta}_s^t, \theta_{-s}^t; M_{\mathrm{money}}\right)\right)\right], \tag{3}$$

Where, on the right-hand side, the reports $\{\hat{\theta}_s^t : t \in [T]\}$ are determined according to the strategy A_s based on past history. Here, the expectation is taken over truthful reports of other agents, types of agent s, and any randomness in the strategy A_s. Similarly, under the strategy A_s, the expected spending of agent s is defined as

$$\widehat{c}_s(A_s) \triangleq \mathbb{E}\left[\sum_{t=1}^{T} C_s\left(\hat{\theta}_s^t, \theta_{-s}^t; M_{\mathrm{money}}\right)\right]. \tag{4}$$

Given these definitions, the auxiliary problem for agent s is defined as

$$\begin{aligned} \max_{A_s} \quad & \widehat{u}_s(A_s) \\ \text{s.t.} \quad & \widehat{c}_s(A_s) \leq B_s. \end{aligned} \tag{5}$$

Note that the budget constraint in the preceding problem is required to hold only in expectation, and not sample-path-wise, unlike in the NMBR mechanism.

Next, we argue that (5) is indeed a relaxation of agent s's decision problem in the NMBR mechanism. To see this, first, for any strategy A_s in the NMBR mechanism, define $A_s^{\emptyset}$ as the strategy that mimics A_s until the agent's budget runs out and reports $\hat{\theta}_s^t = \emptyset$ thereafter. It is immediate that $u_s(A_s; \text{NMBR}) = u_s(A_s^{\emptyset}; \text{NMBR})$ and $c_s(A_s; \text{NMBR}) = c_s(A_s^{\emptyset}; \text{NMBR})$. Second, Assumption 2 implies that $u_s(A_s^{\emptyset}; \text{NMBR}) = \widehat{u}_s(A_s^{\emptyset})$ and $c_s(A_s^{\emptyset}; \text{NMBR}) = \widehat{c}_s(A_s^{\emptyset})$. Finally, because the budget constraint holds sample-path-wise in the NMBR mechanism, the expected budget constraint in (5) is satisfied by $A_s^{\emptyset}$. Taken together, we obtain that for any strategy A_s in the NMBR mechanism, the strategy $A_s^{\emptyset}$ is feasible for the auxiliary problem and achieves the same expected utility.

We next compare the performance of a feasible strategy A_s for (5) to that of the truthful report Tr. First, note that by definition of $\bar{c}_s$, because $B_s > \bar{c}_s T$, truthful reporting Tr is feasible for (5). The next lemma establishes a sensitivity result, which states that the gain in expected utility for the agent upon deviating from truthful reporting is bounded above by her budget surplus.

Lemma 1 (Sensitivity)**.** *Suppose the budget of agent s is $B_s = \bar{c}_s T + \Delta$, for some $\Delta > 0$. Then, for any feasible A_s for (5), we have*

$$\widehat{u}(A_s) \le \widehat{u}(\mathsf{Tr}) + \Delta.$$

Proof. First, note that in the monetary setting, running a BIC mechanism M_{money} repeatedly in each round is overall a BIC mechanism for all T rounds taken together. This follows from the fact that decisions of agents in some round t do not affect their quasi-linear utility in another round t'. This implies that for any strategy A_s feasible in the auxiliary game, we have

$$\widehat{u}_s(A_s) - \widehat{c}_s(A_s) \le \widehat{u}_s(\mathsf{Tr}) - \widehat{c}_s(\mathsf{Tr}).$$

Indeed, if this inequality did not hold for some strategy $\tilde{A}_s$, we could use this strategy to invalidate the assumption that M_{money} is a BIC mechanism.

Now, the budget constraint (4) implies $\widehat{c}_s(A_s) \le B_s = \bar{c}_s T + \Delta$, whereas, by definition, we have $\widehat{c}_s(\mathsf{Tr}) = \bar{c}_s T$. Combining these inequalities yields the needed bound. □

Note that the above sensitivity result is unidirectional, in that it holds in the above form only for $\Delta \ge 0$.

Next, coming back to the NMBR mechanism, to show its approximate incentive compatibility, we seek to bound $u_s(A_s; \text{NMBR}) - u_s(\mathsf{Tr}; \text{NMBR})$ for any strategy A_s. We begin by writing this difference as follows:

$$\begin{aligned}
&u_s(A_s; \text{NMBR}) - u_s(\mathsf{Tr}; \text{NMBR}) \\
&\quad = u_s\big(A_s^{\emptyset}; \text{NMBR}\big) - u_s(\mathsf{Tr}; \text{NMBR}) \\
&\quad = u_s\big(A_s^{\emptyset}; \text{NMBR}\big) - \widehat{u}_s\big(A_s^{\emptyset}\big) + \widehat{u}_s\big(A_s^{\emptyset}\big) - \widehat{u}_s(\mathsf{Tr}) + \widehat{u}_s(\mathsf{Tr}) - u_s(\mathsf{Tr}; \text{NMBR}) \\
&\quad = \underbrace{\widehat{u}_s\big(A_s^{\emptyset}\big) - \widehat{u}_s(\mathsf{Tr})}_{\dagger} + \underbrace{\widehat{u}_s(\mathsf{Tr}) - u_s(\mathsf{Tr}; \text{NMBR})}_{\star},
\end{aligned} \tag{6}$$

where the first equality follows from the fact that $u_s(A_s; \text{NMBR}) = u_s(A_s^{\emptyset}; \text{NMBR})$, and the last equality follows from the fact that $u_s(A_s^{\emptyset}; \text{NMBR}) = \widehat{u}_s(A_s^{\emptyset})$. Now, Lemma 1 gives a bound on the term (†), because $A_s^{\emptyset}$ is feasible for (5). Thus, to complete our proof, we seek to bound the term (⋆), that is, compare the utility under truthful reporting in the auxiliary problem and the NMBR mechanism.

To do this, we begin by defining the following *budget depletion event* $\mathcal{E}_s$:

$$\mathcal{E}_s \triangleq \left\{ \sum_{1}^{T} C_s\big(\theta_s^t, \theta_{-s}^t; M_{\text{money}}\big) \ge B_s \right\}. \tag{7}$$

The event $\mathcal{E}_s$ captures all the sample paths under which agent s's spending exceeds her budget B_s in the auxiliary problem, under truthful reporting. Alternatively, one can envisage $\mathcal{E}_s$ as the event under which the agent s becomes inactive before the end of the NMBR mechanism under truthful reporting. We have the following lemma.

Lemma 2.

$$\widehat{u}_s(\mathsf{Tr}) \le u_s(\mathsf{Tr};\mathrm{NMBR}) + 2\mathbb{P}(\mathcal{E}_s)Tv_{\max}.$$

Proof. Define the utility received by the agent s in the auxiliary problem under truthful reporting as

$$\widehat{U}_s(\mathsf{Tr}) = \sum_{t=1}^{T} v_s\big(\theta_s^t, X\big(\theta_s^t, \theta_{-s}^t; M_{\mathrm{money}}\big)\big).$$

We can write the expected utility in the NMBR mechanism and the auxiliary problem respectively as

$$u_s(\mathsf{Tr};\mathrm{NMBR}) = \mathbb{P}(\neg\mathcal{E}_s)\mathbb{E}\big[U_s(\mathsf{Tr};\mathrm{NMBR})|\neg\mathcal{E}_s\big] + \mathbb{P}(\mathcal{E}_s)\mathbb{E}\big[U_s(\mathsf{Tr};\mathrm{NMBR})|\mathcal{E}_s\big],$$
$$\widehat{u}_s(\mathsf{Tr}) = \mathbb{P}(\neg\mathcal{E}_s)\mathbb{E}\Big[\widehat{U}_s(\mathsf{Tr})|\neg\mathcal{E}_s\Big] + \mathbb{P}(\mathcal{E}_s)\mathbb{E}\Big[\widehat{U}_s(\mathsf{Tr})|\mathcal{E}_s\Big].$$

Note that we have $U_s(\mathsf{Tr};\mathrm{NMBR}) = \widehat{U}_s(\mathsf{Tr})$ on the event $\neg\mathcal{E}_s$. This is because, on this event, in the NMBR mechanism, the agent remains active until time T and receives the same sequence of values as in the M_{money} mechanism. On the event $\mathcal{E}_s$, we have the trivial bound $|U_s(\mathsf{Tr};\mathrm{NMBR}) - \widehat{U}_s(\mathsf{Tr})| \le 2Tv_{\max}$. Substituting this into the decomposition above yields the needed bound. □

Putting Lemma 1 and Lemma 2 together, along with the fact that $\Delta = \delta_s\bar{c}_s T$ for the NMBR mechanism, we obtain from (6) that

$$u_s(A_s;\mathrm{NMBR}) - u_s(\mathsf{Tr};\mathrm{NMBR}) \le \delta_s\bar{c}_s T + 2Tv_{\max}\mathbb{P}(\mathcal{E}_s). \tag{8}$$

As a final step, we bound the probability of the budget depletion event $\mathcal{E}_s$. In particular, we show that with high probability, an agent s remains active until the end of the NMBR mechanism under truthful reporting. We establish this by showing that the expected payment of the agent in the auxiliary problem concentrates around its mean. We have the following lemma.

Lemma 3 (Concentration of Spending)**.** *For every agent s, we have*

$$\mathbb{P}[\mathcal{E}_s] \le \exp\left(-\frac{\delta_s^2 R_s T}{3}\right).$$

Proof. Recall from the NMBR mechanism that $B_S = \bar{c}_s T(1+\delta_s)$. The bound follows from direct application of the following standard Chernoff bounds (for more details, see Dubhashi and Panconesi [16, chapter 1]): For X_i independent random variables with $0 \le X_i \le 1$ and $X = \sum_i X_i$, we have

$$\mathbb{P}\big[X \ge \mathbb{E}[X](1+\epsilon)\big] \le \exp\left(-\frac{\epsilon^2\mathbb{E}[X]}{3}\right). \quad \square$$

Thus, we obtain from Lemma 3 and (8) that

$$u_s(A_s;\mathrm{NMBR}) - u_s(\mathsf{Tr};\mathrm{NMBR}) \le \delta_s\bar{c}_s T + 2Tv_{\max}\exp\left(-\frac{\delta_s^2 R_s T}{3}\right). \tag{9}$$

As we increase δ_s, the gains from potential deviations in the auxiliary game increase (as in Lemma 1), but the gap in the performance of truthful reporting between the original and auxiliary games decreases (via Lemmas 2 and 3). Choosing $\delta_s = \sqrt{\frac{3\log T}{R_s T}}$ balances this trade-off, thereby establishing approximate incentive compatibility guarantee (1). Similarly, from Lemma 2 and (8), after substituting the aforementioned value of δ_s, we have

$$u_s(\mathsf{Tr};\mathrm{NMBR}) \ge \widehat{u}_s(\mathsf{Tr}) - 2v_{\max}T^{-1}. \tag{10}$$

Observe that $\widehat{u}_s(\mathsf{Tr}) = Tu_s(\mathsf{Tr};M_{\mathrm{money}})$, because the auxiliary problem involves repeating M_{money} for T time periods. Upon dividing by T, we obtain the utility guarantee (2).

3.3. Tractable Black-Box Reduction via Sample-Averaged Budgets

Note that the NMBR mechanism requires the expected payment $\bar{c}_s$ of each agent s under M_{money} as an input to compute the initial budgets. However, for general M_{money} and Θ_s, computing $\bar{c}_s$ exactly may be computationally hard. One way to resolve this is to compute $\bar{c}_s$ approximately using a finite number of type samples. In this section, we show how one can preserve our approximation guarantees while using such sample-averaged budget estimates.

Before stating and proving our main results in this section, we need to introduce some additional notation. For each agent s, suppose we have m independent samples $\{\theta_s^{(1)}, \theta_s^{(2)}, \ldots, \theta_s^{(m)}\}$ from the agent's type distribution F_s. Let $\bar{c}_s{}^m$ be the *sample-average cost* over m rounds of the mechanism M_{money} for agent s, where simulated round k uses sampled types $\{\theta_s^{(k)} : s \in [n]\}$. We define $\tilde{B}_s = \bar{c}_s{}^m(1+\delta_s)T$ to be the *sample-average budget* for agent s.

The main result of this section states that when the budgets of agents are set using sample averages $\tilde{B}_s$ with $m = T$, the (α, β)-approximation guarantee of the NMBR mechanism is preserved with high probability.

Theorem 2. *Consider the NMBR mechanism with budgets set to $\tilde{B}_s = T\bar{c}_s{}^m(1+\delta_s)$ (with the same choice of δ_s as in Theorem 1) via sampling the type of every agent $m = T$ times. Then, with probability at least $1 - 2nT^{-1}$, the following bound holds for any agent s and any strategy A_s:*

$$\frac{u_s(\mathsf{Tr};\text{NMBR})}{T} \geq \frac{u_s(A_s;\text{NMBR})}{T} - 3\sqrt{\frac{27\bar{c}_s{}^m c_{\max}\log T}{T}} - 4v_{\max}T^{-1},$$
$$\frac{u_s(\mathsf{Tr};\text{NMBR})}{T} \geq u_s\big(\mathsf{Tr}; M_{\text{money}}\big) - 2v_{\max}T^{-1}.$$

The proof of this theorem relies on showing a concentration bound for $\tilde{B}_s$ and then repeating the argument employed for proving Theorem 1. For brevity, we defer the proof to Appendix A.

The bounds provided in this theorem are asymptotically equivalent to those in Theorem 1; however, the constants are larger, as we need to account for the errors imposed by the sampling procedure. The proof of this theorem can be found in Appendix A. Using this, we get the following corollary that establishes that the computational complexity of running the NMBR mechanism is similar to that of running the monetary mechanism M_{money} over T rounds.

Corollary 1 (Computational Cost of the NMBR Mechanism)**.** *Let C be the computational cost of running monetary mechanism M_{money} over one period. Then, implementing the NMBR mechanism corresponding to M_{money} over T periods requires $2T$ samples for every type distribution F_s and has a computational cost of $2TC_{\text{money}} + O(Tn)$.*

Proof. Given a mechanism M_{money}, we first run it T times with sampled types θ_s^t as input and use this to compute the average costs $\bar{c}_s^m$ for every agent; this requires $O(nT)$ operations. Moreover, during the actual execution of the mechanism, we run the mechanism M_{money} for T times (once per period) with the actual reports of the agents, and in the worst case, the NMBR mechanism simulates every agent's bid T times. Thus, the NMBR mechanism needs at most $2T$ sampled types per agent and $2T$ executions of the mechanism M_{money}. □

4. Beyond Excludability: The REAP Mechanism

Though the NMBR mechanism in the previous section provides a tractable black-box mechanism for a wide range of repeated allocation settings, it is crucially dependent on the excludability assumption (Assumption 1). This assumption may not hold in certain allocation settings, such as in exchange economies and social choice settings. In this section, we propose an alternate mechanism, which we call REAP (for repeated endowed all-pay), which satisfies guarantees similar to those of Theorem 1 without requiring the excludability assumption. This extension, however, comes at a cost, namely, that we lose the tractability of the NMBR mechanism in settings with rich type spaces.

Before presenting the algorithm, we define some notation. Given a monetary mechanism M_{money}, for any agent s, we define a *personalized pricing rule* $\{c_s(\cdot) : \theta \in \Theta\}$, where $c_s(\theta)$ is a price agent s pays to report type θ and is defined as follows:

$$c_s(\theta_s) \stackrel{\Delta}{=} \mathbb{E}\big[C_s\big(\theta_s, \theta_{-s}; M_{\text{money}}\big)\big]. \tag{11}$$

Here, the expectation is over θ_{-s} (drawn from the product measure $F_{-s} = \prod_{q\neq s} F_q$).

Example 2 (Single-Item Allocation). Recall the setting we used to illustrate the NMBR mechanism in the previous section: there are n agents with values $v_{it} \sim \text{Unif}[0,1]$, and M_{money} is the second-price auction. In this case, the personalized price (11) for reporting some value $\hat{v}_s$ in the REAP mechanism is given by

$$c_s(\hat{v}_s) = \mathbb{E}\left[\mathbf{1}\left\{\max_{r \neq s} v_r \leq \hat{v}_s\right\} \cdot \max_{r \neq s} v_r\right] = \frac{n-1}{n}\hat{v}_s^n.$$

Similarly to the previous section, we define $\bar{c}_s = \mathbb{E}_{\theta \sim F_s}[c_s(\theta)]$ to be the expected price charged to s under $c_s(\cdot)$ in a single period under truthful reporting and $c_s^{\max} = \max\{c_s(\theta_s) : \theta_s \in \Theta_s\}$ to be the maximum price charged to agent s, and finally, $R_s = \bar{c}_s / c_s^{\max}$.

We can now define the REAP mechanism. Formally, it is described in Algorithm 2. Informally, the REAP mechanism proceeds as follows: similarly to NMBR, it takes a monetary mechanism M_{money} as an input and initializes by endowing each agent with a budget of credits. Then, in each period, agents report their types, and the resulting allocation is computed via M_{money}. In contrast with NMBR, the payment of a given agent depends only on the report of that agent and is computed as the expected payment for the agent's report in M_{money} when other agents report truthfully.

Algorithm 2 (Repeated Endowed All-Pay Mechanism)

Require: Static monetary BIC mechanism M_{money}, type distributions $\{F_s\}$, surplus parameter δ_s

1: Compute $\{c_s(\cdot) : \theta_s \in \Theta_s\}$, the pricing rule, for each agent s as described in (11); also compute $\bar{c}_s$, the expected per-period payment under truthful reporting.
2: Allocate endowment of $B_s = (1 + \delta_s)\bar{c}_s T$ credits to each agent s. Let $B_s^1 = B_s$.
3: **for all** $t = 1$ to T **do**
4: Get report $\hat{\theta}_s^t$ from each agent s. If $B_s^t - c_s(\hat{\theta}_s^t) < 0$, update $\hat{\theta}_s^t = \emptyset$. Charge each agent s a price of $c_s(\hat{\theta}_s^t)$ credits, and let $B_s^{t+1} = B_s^t - c_s(\hat{\theta}_s^t)$.
5: Implement the allocation $X(\hat{\theta}^t; M_{\text{money}})$
6: **end for**

For this mechanism, we establish a vanishing (α, β)-approximation guarantee resembling that of Theorem 1. Formally, we have the following theorem.

Theorem 3. *Consider the REAP mechanism (Algorithm 2) with $\delta_s = \sqrt{6(\log n + \log T)/R_s T}$. Then, for each agent s and for any strategy A_s, assuming all other agents report truthfully, we have*

$$\frac{u_s(\text{Tr}; \text{REAP})}{T} \geq \frac{u_s(A_s; \text{REAP})}{T} - \sqrt{\frac{6\bar{c}_s c_{\max}(\log n + \log T)}{T}} - 4nv_{\max}T^{-1}.$$

Moreover, assuming all agents play truthfully, we have

$$\frac{u_s(\text{Tr}; \text{REAP})}{T} \geq u_s(\text{Tr}; M_{\text{money}}) - 2nv_{\max}T^{-1}.$$

Comparing this result to that of Theorem 1, we notice that the number of agents n now enters the incentive guarantee explicitly (in addition to entering it implicitly through $\bar{c}_s$). This comes from the fact that in the absence of the simulation step of the NMBR mechanism (Algorithm 1, Step 4), meaningful guarantees hold only on the sample paths on which not a single truthful agent runs out of credits. This can be viewed as a relatively low cost for disposing of this step in the REAP mechanism.

Finally, we make a note that running the REAP mechanism might not be computationally tractable. The reason is that to execute REAP, we need to compute $c_s(\theta_s)$ according to (11), and doing so exactly might be intractable for large type domains. An interesting direction for future research is whether it is possible to overcome this intractability in the general allocation setting.

4.1. Proof of Theorem 3

To prove Theorem 3, we adopt an approach analogous to the one we used in Section 3. Specifically, we establish three lemmas that are analogous to Lemmas 1, 2, and 3. The main difference is that instead of defining a separate auxiliary problem for each agent, here we define a single auxiliary game played by all agents simultaneously and show that the agents' utility profiles under the REAP mechanism and the auxiliary game are close. In addition to its use in this proof, the auxiliary game also plays a role in proving the *price of anarchy* bound in the next section.

Formally, in the auxiliary game, a strategy A_s for an agent s maps the history at any time $t \in [T]$ and her type θ_s^t to her report $\hat{\theta}_s^t$. Given a strategy profile (A_s, A_{-s}), the expected utility of agent s in the auxiliary game is defined as

$$\widehat{u}_s(A_s, A_{-s}) = \mathbb{E}\left[\sum_{t=1}^{T} v_s\left(\theta_i^t, X\left(\hat{\theta}_s^t, \hat{\theta}_{-s}^t; M_{\text{money}}\right)\right)\right], \tag{12}$$

where in the right-hand side, for each agent $q \in [n]$, the reports $\{\hat{\theta}_q^t : t \in [T]\}$ are determined according to the strategy A_q. Similarly, under a strategy A_s, the expected spending of agent s is defined as

$$\widehat{c}_s(A_s) = \mathbb{E}\left[\sum_{t=1}^{T} c_s\left(\hat{\theta}_s^t\right)\right]. \tag{13}$$

Note that, in contrast to analysis of the NMBR mechanism as in Section 3, here the expected spending of agent s is independent of the strategies of the other agents, as an agent's payment depends only on her own reports. We say a strategy A_s is *feasible* for an agent s if it satisfies the expected budget constraint, $c_s(A_s) \le B_s$.

Because $B_s > \overline{c}_s T$, it follows that the strategy Tr is feasible for each agent s. As in the previous section, for an agent s and for any strategy A_s in the REAP mechanism, define the strategy $A_s^{\emptyset}$ as one that mimics A_s until the first time t for which $\sum_{\tau=1}^{t} c_s(\hat{\theta}_s^\tau) > B_s$ holds and reports $\hat{\theta}_s^t = \emptyset$ thereafter. It is straightforward to verify that $A_s^{\emptyset}$ is a feasible strategy in the auxiliary game, satisfying $u_s(A_s, A_{-s}; \text{REAP}) = \widehat{u}_s(A_s^{\emptyset}, A_{-s}^{\emptyset})$ and $c_s(A_s; \text{REAP}) = \widehat{c}_s(A_s^{\emptyset})$ for all s.

Having defined the auxiliary game, we start with the following lemma, an analog of Lemma 1. The proof is also analogous and is omitted for the sake of brevity.

Lemma 4 (Sensitivity)**.** *In the auxiliary game, suppose the budget of agent s is $B_s = \overline{c}_s T + \Delta$ for some $\Delta > 0$. Then, for any feasible strategy A_s of agent s, we have*

$$\widehat{u}_s(A_s, \mathsf{Tr}_{-s}) \le \widehat{u}_s(\mathsf{Tr}_s, \mathsf{Tr}_{-s}) + \Delta.$$

The next lemma, analogous to Lemma 2, establishes that an agent's utility in the REAP mechanism is close to that in the auxiliary game, when other agents report truthfully. Here, we prove this closeness holds for arbitrary strategies of the agent, not just when she herself makes truthful reports. To prove this lemma, we consider the budget depletion event $\mathcal{E} = \cup_{s \in [n]} \mathcal{E}_s$, where $\mathcal{E}_s$ is defined as

$$\mathcal{E}_s = \left\{\sum_{t=1}^{T} c_s\left(\theta_s^t\right) \ge B_s\right\}. \tag{14}$$

We have the following lemma, whose proof is similar to that of Lemma 2 and is deferred to Appendix A.

Lemma 5 (Closeness of Auxiliary and Original Games)**.** *For any strategy A_s employed by agent s under the REAP mechanism, assuming all other agents are playing truthfully, the following inequality holds:*

$$u_s(A_s, \mathsf{Tr}_{-s}; \text{REAP}) \le \widehat{u}_s\left(A_s^{\emptyset}, \mathsf{Tr}_{-s}\right) + 2\mathbb{P}[\mathcal{E}] T v_{\max}.$$

Furthermore, we also have

$$\widehat{u}_s(\mathsf{Tr}_s, \mathsf{Tr}_{-s}) \le u_s(\mathsf{Tr}_s, \mathsf{Tr}_{-s}; \text{REAP}) + 2\mathbb{P}(\mathcal{E}) T v_{\max}.$$

Finally, we have the third lemma, analogous to Lemma 3, which bounds the probability of the budget depletion event $\mathcal{E}$. Once again, the proof is analogous and is omitted. Recall that $R_s = \overline{c}_s / c_s^{\max}$, where $c_s^{\max} = \max\{c_s(\theta_s) : \theta_s \in \Theta_s\}$, and $\overline{c}_s = \mathbb{E}[c_s(\theta_s)]$.

Lemma 6.

$$\mathbb{P}[\mathcal{E}] \leq n \exp\left(-\frac{T \min_s \delta_s^2 R_s}{3}\right). \tag{15}$$

Using these three lemmas, we obtain

$$\begin{aligned} &u_s(A_s, \mathsf{Tr}_{-s}; \text{REAP}) - u_s(\mathsf{Tr}_s, \mathsf{Tr}_{-s}; \text{REAP}) \\ &\quad = \left(u_s(A_s, \mathsf{Tr}_{-s}; \text{REAP}) - \widehat{u}_s(A_s^{\emptyset}, \mathsf{Tr}_{-s})\right) + \left(\widehat{u}_s(A_s^{\emptyset}, \mathsf{Tr}_{-s}) - \widehat{u}_s(\mathsf{Tr}_s, \mathsf{Tr}_{-s})\right) \\ &\qquad + \left(\widehat{u}_s(\mathsf{Tr}_s, \mathsf{Tr}_{-s}) - u_s(\mathsf{Tr}_s, \mathsf{Tr}_{-s}; \text{REAP})\right) \\ &\quad \leq \bar{c}_s T \delta_s + 4\mathbb{P}[\mathcal{E}] T v_{\max}, \end{aligned}$$

where we have used the fact that $B_s = \bar{c}_s T(1 + \delta_s)$ in the inequality. As in the previous section, the appropriate choice of δ_s yields the theorem statement. For more details, see Appendix A.

5. Stronger Equilibrium Guarantees for Two Agents

Theorem 1 and Theorem 3 establish utility guarantees under truthful reporting while simultaneously ensuring that truthful reporting is an ϵ-equilibrium. In particular, as the number of time periods T increases, the expected per-round gain of an agent from a unilateral deviation goes to zero. This raises two related questions:

1. Do NMBR and REAP mechanisms have a vanishing *price of anarchy*; that is, does the per-round welfare[2] under any equilibrium (if one exists) approach the welfare achieved by M_{money} as T increases?
2. Is there a nonmonetary BIC mechanism that achieves vanishing welfare loss (i.e., guarantees $\alpha = 0, \beta = o(1)$ in the formalism of Definition 1)?

Given the generality of our setting, answering these two questions is challenging. Nevertheless, in this section, we show that we can address these questions in a restricted setting, namely, when there are exactly two agents and the REAP mechanism is employed with the VCG mechanism as M_{money}. We refer to this setting as the REAP-VCG mechanism, or RVCG for short.

The rest of this section proceeds as follows. We first establish a price of anarchy for the RVCG mechanism using the auxiliary game approach developed in previous section. We then leverage this result to prove the existence of a BIC mechanism via a revelation principle argument.

5.1. REAP-VCG Mechanism

For the case of two agents, we prove price of anarchy bounds for the REAP mechanism when VCG mechanism is used as input M_{money} (we refer to this setting as the REAP-VCG mechanism, and RVCG for short, throughout). Specifically, we have the following allocation and payment rules:

$$X^*(\theta) \in \arg\max_{X \in \mathcal{X}} \sum_{s \in [n]} v_s(\theta_s, X),$$

and

$$c_s(\theta_s) = L_s - \mathbb{E}_{F_{-s}}\left[\sum_{q \neq s} v_q\left(\theta_q, X^*(\theta)\right)\right], \tag{16}$$

where L_s is defined as

$$L_s = \sup_{\theta_s} \mathbb{E}_{F_{-s}}\left[\sum_{q \neq s} v_q\left(\theta_q, X^*(\theta)\right)\right].$$

Here $F_{-s} = \Pi_{q \neq s} F_q$ denotes the product measure over the types of all agents other than agent s. Note that this choice of L_s ensures that the prices are nonnegative. We also restrict our attention to only those settings where the VCG mechanism satisfies the opt-out assumption (Assumption 2). This implies that for every agent s, there exists a report θ_s such that $c_s(\theta_s) = 0$. For example, in the single-item setting we discussed earlier, the payment rule (16) yields the expected externality payments as in Example 2, and an agent can report $\hat{v}_s = 0$ to guarantee zero payment.

5.2. Price of Anarchy of the REAP-VCG Mechanism

Throughout this section, we state our results from the perspective of agent 1; however, all results admit a corresponding version for agent 2.

Our approach to proving price of anarchy guarantee for the REAP-VCG mechanism is via a technique reminiscent of smoothness (see Roughgarden [39]). In particular, given any strategy profile (A_1, A_2), it is enough for us to show that individual deviations from that profile to truth-telling can guarantee agents their share of welfare independently of what strategy the other agent is playing. Recall that $u_1(A_1, A_2; \mathrm{RVCG})$ denotes the expected utility of agent 1 in the RVCG mechanism, when the strategy profile is (A_1, A_2). The main result of this section is as follows.

Theorem 4 (Smoothness of REAP-VCG)**.** *Given a two-player setting under the RVCG mechanism, for any arbitrary strategy profile* (A_1, A_2)*, we have*

$$\frac{u_1(\mathsf{Tr}, A_2; \mathrm{RVCG})}{T} \geq \frac{u_1(\mathsf{Tr}, \mathsf{Tr}; \mathrm{RVCG})}{T} - O\left(\sqrt{\frac{\log T}{T}}\right). \tag{17}$$

Note that (17) does not exactly fit the definition of (λ, μ) smoothness (see Roughgarden [39]), because we characterize the difference between "entangled" and "truthful" utilities via an additive rather than multiplicative factor. This difference, however, does not change the nature of the argument. In particular, this immediately yields the following corollary.

Corollary 2 (Price of Anarchy)**.** *Let* (A_1, A_2) *be any equilibrium profile under the REAP-VCG mechanism. Then* $u_1(A_1, A_2; \mathrm{RVCG}) + u_2(A_2, A_1; \mathrm{RVCG}) \geq u_1(\mathsf{Tr}, \mathsf{Tr}; \mathrm{RVCG}) + u_2(\mathsf{Tr}, \mathsf{Tr}; \mathrm{RVCG}) - o(T)$.

Proof. Because (A_1, A_2) is an equilibrium profile, we have $u_1(A_1, A_2; \mathrm{RVCG}) \geq u_1(\mathsf{Tr}, A_2; \mathrm{RVCG})$ (and similarly for agent 2). Moreover, from (17), we have $u_1(\mathsf{Tr}, A_2; \mathrm{RVCG}) \geq u_1(\mathsf{Tr}, \mathsf{Tr}; \mathrm{RVCG}) - o(T)$. Combining and summing over both agents, we get our result. □

Thus, the rest of this subsection is dedicated to proving Theorem 4. We do so by again employing the auxiliary game we constructed earlier (see Section 3.2), in which agents are required to satisfy their budget constraints only in expectation. We start by proving a smoothness property for the auxiliary game; we then use the closeness of the auxiliary and original game to port the smoothness result over to the original game.

We first need some additional notation. We use $\widehat{U}_1^t(\theta_1^t, \tilde{\theta}_2^t)$ to denote the utility of agent 1 in period t in the auxiliary game if she reports type θ_1^t and agent 2 reports type $\tilde{\theta}_2^t$. The following lemma, the main driver of the proof of Theorem 4, establishes a connection between utility of the first agent when she is reporting truthfully and the payments of the second agent. This proof follows by a direct rearrangement of the terms in the pricing rule (16) and is omitted for brevity.

Lemma 7. *In the auxiliary game, suppose agent* 2*'s report at time t is* $\tilde{\theta}_2^t$*. Then the expected utility of agent* 1 *under truthful reporting (i.e.,* $\theta_1^t \sim F_1$*) satisfies*

$$\mathbb{E}\left[\widehat{U}_1^t\left(\theta_1^t, \tilde{\theta}_2^t\right) \middle| \tilde{\theta}_2^t\right] = L_2 - c_2\left(\tilde{\theta}_2^t\right). \tag{18}$$

Using this characterization, we can now demonstrate smoothness for the auxiliary game.

Lemma 8 (Smoothness of the Auxiliary Game)**.** *Suppose the budget of agent* 2 *is* $B_2 = (1 + \delta_2)\bar{c}_2 T$*. Then, for any feasible strategy* A_2 *of agent* 2*, we have*

$$\widehat{u}_1(\mathsf{Tr}, A_2) = \widehat{u}_1(\mathsf{Tr}, \mathsf{Tr}) - \delta_2 \bar{c}_2 T.$$

Proof. Let $\{\tilde{\theta}_2^t : t \in [1, 2, \ldots, T]\}$ denote a sequence of reported types produced by agent 2's strategy A_2 in a sample path of the mechanism. Note that these reports may be correlated with each other, or with the reports of the first agent. Now we have

$$\begin{aligned}
\widehat{u}_1(\mathsf{Tr}, A_2) &= \mathbb{E}\left[\sum_{t=1}^{T} \hat{U}_1\left(\theta_1^t, \tilde{\theta}_2^t\right)\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{T} \mathbb{E}\left[\hat{U}_1\left(\theta_1^t, \tilde{\theta}_2^t\right) \middle| \tilde{\theta}_2^t\right]\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{T} \left(L_2 - c_2\left(\tilde{\theta}_2^t\right)\right)\right] \qquad \text{(by Lemma 7)} \\
&= T L_2 - \widehat{c}_2(A_2).
\end{aligned}$$

Now, because the strategy A_2 is feasible, we have $\widehat{c}_2(A_2) \le B_2$, and hence

$$\widehat{u}_1(\mathsf{Tr}, A_2) \ge TL_2 - B_2 = TL_2 - \bar{c}_2 T - \delta_2 \bar{c}_2 T = \widehat{u}_1(\mathsf{Tr}, \mathsf{Tr}) - \delta_2 \bar{c}_2 T,$$

where the last equality follows from observing $\bar{c}_s[2]T = TL_2 - \widehat{u}_1(\mathsf{Tr}, \mathsf{Tr})$, which can be derived from (18) by taking expectation over truthful reports $\theta_2^t \sim F_2$. □

Next, to extend this smoothness result to the original setting, we need a generalization of Lemma 5 for settings where agent s reports truthfully, while other agents are playing some strategy A_{-s} (which is feasible in the original game). As before, we let $\mathcal{E}_s$ denote the event that agent s depletes her budget before time T while reporting truthfully under the REAP-VCG mechanism. We have the following lemma, whose proof is analogous to that of Lemma 2 and is omitted for brevity.

Lemma 9. *For any strategy A_2 of agent 2 in the REAP-VCG mechanism, we have*

$$\left|u_1(\mathsf{Tr}, A_2) - \widehat{u}_1(\mathsf{Tr}, A_2)\right| \le T v_{\max} \mathbb{P}(\mathcal{E}_1).$$

Finally, these lemmas allow us to prove Theorem 4 in a manner analogous to that of the proof of Theorem 1.

Proof of Theorem 4. We can write the entangled utility term (left-hand side of (17)) as

$$\begin{aligned}
&u_1(\mathsf{Tr}, A_2; \mathrm{RVCG}) - u_1(\mathsf{Tr}, \mathsf{Tr}; \mathrm{RVCG}) \\
&= \big(u_1(\mathsf{Tr}, A_2; \mathrm{RVCG}) - \widehat{u}_1(\mathsf{Tr}, A_2)\big) + \big(\widehat{u}_1(\mathsf{Tr}, A_2) - \widehat{u}_1(\mathsf{Tr}, \mathsf{Tr})\big) \\
&\quad + \big(\widehat{u}_1(\mathsf{Tr}, \mathsf{Tr}) - u_1(\mathsf{Tr}, \mathsf{Tr}; \mathrm{RVCG})\big) \\
&\ge -T v_{\max} \mathbb{P}(\mathcal{E}_1) + \delta_2 \bar{c}_2 T - T v_{\max}(\mathbb{P}(\mathcal{E}_1) + \mathbb{P}(\mathcal{E}_2)),
\end{aligned}$$

where the first bound follows from Lemma 5, the second from Lemma 8, and the last from Lemma 9. Substituting our choice of $\delta_1 = \sqrt{\frac{6(\log 2 + \log T)}{R_1 T}}$, we get the result. □

Unfortunately, the technique of proving vanishing price of anarchy via analyzing the deviation to truthfulness does not extend to the case of more than two agents; namely, our result relies on the connection between agent 1's utility and agent 2's payment as stated in Lemma 7, which holds even if agent 2 is misreporting. Such a connection does not persist when there are multiple opponents, because externality imposed by a group of agents does not equal to the sum of externalities of individual agents. Whether the price of anarchy result persists in the case of more than two agents is an open question.

5.3. Existence of a Near-Efficient BIC Mechanism

Finally, we return to the question of whether there exists a nonmonetary BIC mechanism with vanishing inefficiency ($\alpha = 0$, $\beta = o(1)$, according to Definition 1). We now show how we can leverage the price of anarchy result from the previous section to prove the existence of such mechanism, under the assumption that the type spaces are finite.

Theorem 5. *If the type spaces Θ_s are finite, there exists a Bayes–Nash equilibrium in the REAP-VCG mechanism.*

The proof of this theorem is based on the fact that REAP-VCG is a finite state space game and can be found in Appendix A.

Corollary 3 (Existence of a BIC Mechanism). *Under the conditions in Theorem 5, there exists a BIC mechanism for two agents with expected welfare loss of $o(T)$.*

Proof. In Theorem 5, we establish that a Bayes–Nash equilibrium exists for the REAP-VCG mechanism; moreover, Corollary 2 guarantees that expected welfare loss at this equilibrium is at most $o(T)$. Thus, given such a Bayes–Nash equilibrium strategy, one can use the revelation principle to construct a mechanism that is Bayesian incentive compatible and has the same welfare guarantees. □

6. Discussion

We have presented a general black-box reduction technique that allows us to convert any static monetary mechanism to a repeated nonmonetary mechanism, while approximately preserving the incentive, welfare, and tractability guarantees of the original mechanism. The modular nature of the proof of our main results allows us to study how it extends to settings more general than the one we have analyzed. Below, we briefly discuss two such extensions.

6.1. Nonstationary Independent Type Distributions

A simple extension to our model allows agents' type distributions $F_s^t(\cdot)$ that depend on the time period. It is straightforward to show that our results extend to this setting. We briefly describe some details here.

Both NMBR and REAP mechanisms have a natural extension to nonstationary distributions: in the NMBR mechanism, type sampling (Algorithm 1, Step 4) would be performed with respect to the appropriate distribution $F_s^t(\theta_s)$, and in the REAP mechanism, the pricing rule (11) would depend on time t. We update the definition of $\bar{c}_s$ to $\bar{c}_s = \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}[c_s^t(\theta^t; M_{\text{money}})]$, where $c_s^t(\theta^t; M_{\text{money}})$ denotes the payment of agent s in the static BIC mechanism M_{money}, when types are distributed as F^t. Similarly, we define $c_s^{\max} = \max_{t,\theta} c_s^t(\theta; M_{\text{money}})$.

With these modifications, our proofs go through with no change; we illustrate this in the proof for the NMBR mechanism. The definition of the auxiliary problem remains untouched, and the sensitivity (Lemma 1) and the closeness (Lemma 2) results still follow, as their proofs do not use stationarity. We can also extend the concentration result (Lemma 3), as the Chernoff bound holds for bounded, nonidentical, independent distributions.

6.2 Conditionally Independent Types

Our results are less amenable to the case when types of agents are correlated, either across agents or across time. The primary reason for this is that there is no straightforward way to extend the NMBR or REAP mechanisms to these cases. For instance, when types are independent, the NMBR mechanism uses samples from an agent's distribution once her budget runs out (see Algorithm 1, Step 4). When types are correlated, the (conditional) distribution of an agent's type may be a priori unknown to the principal, and it is unclear how agents with no remaining budgets should be handled. Similarly, when types are independent, the REAP mechanism charges each agent her expected payment in the M_{money} mechanism given her report (see (11)). When types are correlated, an agent's payment in the M_{money} mechanism will be correlated with her type; without access to an agent's type, in general, the principal cannot compute the expected payment in M_{money}.

Despite the above concerns, our (α, β)-approximation results can be extended under a mild form of correlation in agents' types, namely, when the agents' types are *conditionally independent*. Formally, this corresponds to the case where there exists an underlying random process $\{\omega^t : t \in [T]\}$ such that at each time t, conditional on ω^t, the agents' types θ^t are independent. We further require that for any t, given $\{\omega^r : r < t\}$, the distribution of ω^t is independent of $\{\theta^r : r < t\}$. Finally, we assume that the principal has access to ω_t prior to making the allocation at each time t. This assumption captures resource allocation settings where agents' value for the resource depends on an (observable) quality of the resource (e.g., weather conditions in the case of allocation of parking spots or vacation days). (Note that nonstationary independent type distributions form a special case of conditionally independent distributions, where $\omega_t = t$ for each $t \in [T]$.)

Under conditional independence, both NMBR and REAP mechanisms allow a straightforward extension. In particular, in the simulation step (Algorithm 1, Step 4) of the NMBR mechanism, for any agent s with no remaining budget, the mechanism samples a report $\hat{\theta}_s^t$ from the conditional distribution given ω^t. Similarly, for the REAP mechanism, each agent s is charged her (conditional) expected payment in the M_{money} mechanism, given ω^t. With these modifications, Lemma 1 and Lemma 2 continue to hold. Thus, our (α, β)-approximation results survive as long as the process $\{\omega^t : t \in [T]\}$ satisfies a concentration result analogous to Lemma 3. This holds, for example, when ω^t is independent across time.

Acknowledgments

The majority of this work was carried out while the third author was with Cornell University. Preliminary versions of this work were presented at the 2016 International Workshop on Internet and Network Economics and the 2017 ACM Conference on Economics and Computation, and accompanied by one-page extended abstracts (Gorokh et al. [21, 22]). The authors thank the anonymous reviewers from these conferences for their helpful feedback.

Appendix A. Additional Proofs of Results

A.1. Proofs for the NMBR Mechanism

Proof of Theorem 2. Let $\mathcal{G}_s$ be the event that $\tilde{B}_s$ differs from its expected value (the endowment B_s in the NMBR mechanism) by more than $T\bar{c}_s\delta_s/2$, that is,

$$\mathcal{G}_s = \{|\tilde{B}_s - B_s| \geq T\bar{c}_s\delta_s/2\}.$$

We now have the following concentration bound.

Lemma A.1. *Given sample-average budget $\tilde{B}_s = \bar{c}_s{}^m(1+\delta_s)$ obtained from m samples (for some $0 < \delta_s < 1/2$), we have*

$$\mathbb{P}[\mathcal{G}_s] \le 2\exp\left(-\frac{\delta_s^2 R_s m}{27}\right).$$

Proof. *The first bound follows from standard Chernoff bounds (see Lemma 6):*

$$\mathbb{P}\Big[\tilde{B}_s \ge T\bar{c}_s(1+\delta+\delta/2)\Big] \le \mathbb{P}\left[T\bar{c}_s{}^m(1+\delta) \ge T\bar{c}_s(1+\delta)\frac{1+3\delta/2}{1+\delta}\right]$$
$$\le \mathbb{P}\big[T\bar{c}_s{}^m(1+\delta) \ge T\bar{c}_s(1+\delta)(1+\delta/3)\big] \le \exp\left(-\frac{\delta_s^2 R_s m}{27}\right).$$

Analogous derivation for the lower bound of $\tilde{B}_s$ yields the result. □

We denote by $u_s(A_s;\mathcal{G}_s)$ the expected utility of agent s playing strategy A_s conditioned on the event $\mathcal{G}_s$. When the event $\mathcal{G}_s$ happens, we can repeat the argument of Theorem 1 to obtain

$$u_s(A_s;\mathcal{G}_s) - u_s(\mathsf{Tr};\mathcal{G}_s) \le \frac{3}{2}T\bar{c}_s\delta_s + 4\exp\left(-\frac{\delta_s^2 R_s T}{12}\right)Tv_{\max}. \tag{A.1}$$

The expected utility profile bound is obtained analogously to the proof of Theorem 1, by taking a union bound of events $\mathcal{G}_s$ and $\mathcal{E}_s$:

$$u_s(\mathsf{Tr};\mathcal{G}_s) \ge Tu_s\big(\mathsf{Tr};M_{\text{money}}\big) - 2\exp\left(-\frac{\delta^2 R_s T}{27}\right)Tv_{\max}. \tag{A.2}$$

Now choose $\delta_s = \sqrt{\frac{27\log T}{T\tilde{R}_s}}$, where $\tilde{R}_s = \frac{\bar{c}_s{}^m}{c_{\max}}$. Note that here we had to use $\tilde{R}_s$ instead of $R_s = \frac{\bar{c}_s}{c_{\max}}$, which may be unknown to the mechanism designer. This, however, does not affect the bound: it follows from Lemma A.1 that $\mathbb{P}[|c_s^m(\mathsf{Tr}) - \bar{c}_s| \le \bar{c}_s/2] \le 2\exp(-\frac{R_s m}{27}) = o(T^{1-k})$ for any k. This, together with the choice of δ_s and the bounds (A.1) and (A.2), yields the needed result. □

A.2. Proofs for the REAP Mechanism

Proof of Lemma 5. Define the utility received by the agent s in the auxiliary problem when playing strategy A_s against truthful opponents as

$$\widehat{U}_s(A) = \sum_{t=1}^{T} v_s\big(\theta_s^t, X\big(\hat{\theta}_s^t, \theta_{-s}^t; M_{\text{money}}\big)\big).$$

Observe that, conditioned on the event $\neg\mathcal{E}_{-s}$ (i.e., restricting to sample paths where no agent apart from s runs out of credits), expected utilities in original and auxiliary games are equal; formally, we have $\mathbb{E}\big[U_s(A_s)|\neg\mathcal{E}_{-s}\big] = \mathbb{E}[\widehat{U}_s(A_s^{\emptyset})|\neg\mathcal{E}_{-s}]$. This gives us

$$u_s(A_s) = \mathbb{P}\Big[\neg\mathcal{E}_{-s}\Big]\mathbb{E}[U_s(A_s)|\neg\mathcal{E}_{-s}] + \mathbb{P}\Big[\mathcal{E}_{-s}\Big]\mathbb{E}[U_s(A_s)|\mathcal{E}_{-s}] \le \mathbb{P}[\neg\mathcal{E}_{-s}]\mathbb{E}\Big[\widehat{U}_s\big(A_s^{\emptyset}\big)|\neg\mathcal{E}_{-s}\Big] + 2\mathbb{P}[\mathcal{E}_{-s}]Tv_{\max} \le \widehat{u}_s\big(A_s^{\emptyset}\big) + 2\mathbb{P}[\mathcal{E}_{-s}]Tv_{\max}. \tag{A.3}$$

For the second inequality, note that, conditioned on $\neg\mathcal{E}_{-s} \cap \neg\mathcal{E}_s$ (i.e., no agent runs out of credits), the utility of a truthful agent in the auxiliary game is equal to her utility in the original game, that is, $\mathbb{E}[U_s(\mathsf{Tr}_s)|\neg\mathcal{E}_{-s} \cap \neg\mathcal{E}_s] = \mathbb{E}[U_s(\mathsf{Tr}_s)|\neg\mathcal{E}_{-s} \cap \neg\mathcal{E}_s]$. The bound is then derived similarly to Equation (A.3). □

We now turn back to the REAP mechanism and prove the promised (α,β)-approximation guarantee.

Proof of Theorem 3. Let A_s denote an arbitrary strategy in REAP. We have

$$u_s(A_s) - u_s(\mathsf{Tr}) = u_s(A_s) - \widehat{u}_s\big(A_s^{\emptyset}\big) + \widehat{u}_s(A_s) - \widehat{u}_s(\mathsf{Tr}) + \widehat{u}_s(\mathsf{Tr}) - u_s(\mathsf{Tr}). \tag{A.4}$$

We can now use bounds from Lemma 5 for $u_s(A_s) - \widehat{u}_s(A_s^{\emptyset})$ and $\widehat{u}_s(\mathsf{Tr}) - \widehat{u}_s(\mathsf{Tr})$, and Lemma 4 for $\widehat{u}_s(A_s) - \widehat{u}_s(\mathsf{Tr})$. Substituting them into (A.4) gives

$$u_s(A_s) - u_s(\mathsf{Tr}) \le T\bar{c}_s\delta + 2(2n+1)Tv_{\max}\exp\left(-\frac{\delta_s^2 R_s T}{3}\right).$$

Substituting $\delta_s = \sqrt{\frac{6(\log n + \log T)}{R_s T}}$ in the above equation, we get the first promised bound (for the gains from deviation).

To prove the β-approximation, we first need some additional notation. Let $\mathcal{E} = \cup_{s\in[n]}\mathcal{E}_s$ denote the event where at least one agent runs out of money before the end of T rounds. Note also that under our choice of δ_s, we have, for all s, $\delta_s^2 R_s = \frac{6(\log n + \log T)}{T} \triangleq \delta^2 R$. Finally, from Lemma 6, we have $\mathbb{P}[\mathcal{E}] \le n\exp\left(-\frac{\delta^2 RT}{3}\right)$. Now we have

$$\mathbb{E}\left[u_s(\mathsf{Tr})\right] \ge \mathbb{P}[\neg\mathcal{E}]\mathbb{E}\left[U_s(\mathsf{Tr}_s)|\neg\mathcal{E}\right] \ge \widehat{u}_s(\mathsf{Tr}) - nv_{\max}T\mathbb{P}[\mathcal{E}] \ge u_s\left(\mathsf{Tr}; M_{\text{money}}\right) - 2nv_{\max}T\mathbb{P}[\mathcal{E}]\,.$$

Substituting $\delta^2 R = \frac{6(\log n + \log T)}{T}$ yields the result. □

Proof of Theorem 5. Note that in the REAP mechanism, the budget of each agent at any time t is completely determined by the previous reports. Because the type space Θ is finite and we focus on direct-revelation mechanisms, the set of possible reports is finite. Thus, the set of possible budgets of each agent at any time t is finite.

Now, in the REAP mechanism, the budget of an agent at any time t denotes her state, and the type space denotes her actions. Because this is a finite game, the existence of a BNE follows using standard arguments (Fudenberg and Tirole [18, chapter 14]). □

Appendix B. Applicability of the REAP Mechanism: Examples

The main difference between the mechanisms described in Sections 3 and 4 is the dependence on excludability assumption (Assumption 1); namely, the NMBR mechanism makes use of it and the REAP mechanism does not. This can be viewed as a trade-off between generality and tractability, but how restricting is Assumption 1? And are there important problems that can be solved with REAP but not NMBR?

In this section, we aim to answer this question by providing several examples.

B.1. Assumptions on the Allocation Setting

We start with the example of central allocation of items with combinatorial preferences, for which we argue both mechanisms are applicable.

Example B.1 (Combinatorial Auction)**.** There are n agents and m items, and agents have combinatorial preferences for the allocated set of items that result from independently and randomly drawn types θ_s.

Assumption 1 clearly holds for this setting. For the NMBR mechanism, we can use VCG with the Clarke pivot rule as an input monetary mechanism M_{money}. If more is known about utility functions, other results from monetary mechanism design can be employed for tractability purposes, for example, using the work of Hartline et al. [27].

Example B.2 (Mutually Beneficial Exchange of Goods)**.** There are two agents, and at each round, an item is given to one of them at random, and agents sample their values $v_s^t \sim F$ for the item independently from the same distribution F at every round. A mechanism is to choose whether to reallocate the item at every round.

It is easy to see that, whatever the value distribution is, the expected optimal allocation in this case Pareto dominates the default one, as both players have greater utility under the mechanism compared with initial allocation, and so (as we argue in Section 2) our individual rationality criterion is satisfied. Thus, REAP is an approximately IC, approximately efficient, and ex ante individually rational mechanism.

Example B.3 (Voting)**.** There are n agents and m options. In each round, agents have value $v_{sk}^t \sim F_{sk}$ for the option k (drawn i.i.d.). A mechanism is to choose a single option in every round, and utilities of agents are their values for the option chosen.

Again, Assumption 1 does not apply here, as it is impossible to prevent an agent from deriving utility from whatever option is chosen by the mechanism. However, there is nothing that prevents us from applying the REAP mechanism to derive near-optimal welfare.

Appendix C. Interim Guarantees for the NMBR Mechanism

Although throughout this paper we focus on ex ante guarantees, our results also imply that after a constant fraction γT (for any $\gamma < 1$) of the time periods has passed, with high probability (but decreasing in γ), the agent will continue to find truthful reporting to be approximately optimal (with approximation factor decreasing in γ). More formally, one can show the following result.

Theorem C.1. *Consider the NMBR mechanism and any constant $\gamma \in (0,1)$, and let $\tilde{T} = (1-\gamma)T$. Then, with probability $p \ge 1 - 2e^{-\frac{\delta_s^2 R_s \gamma T}{3}}$, the following hold simultaneously for all agents s at round γT, if all agents report truthfully up to this turn*:

$$\frac{u_s^{\tilde{T}}(A_s;\text{NMBR}) - u_s^{\tilde{T}}(\mathsf{Tr};\text{NMBR})}{\tilde{T}} \le \bar{c}_s\sqrt{\frac{3(1+\gamma)\log\tilde{T}}{(1-\gamma)R_s\tilde{T}}} + 2v_{\max}\tilde{T}^{-1}$$

$$\textit{and}\qquad u_s\left(\mathsf{Tr}; M_{\text{money}}\right) - \frac{u_s^{\tilde{T}}(\mathsf{Tr};\text{NMBR})}{\tilde{T}} \le 2nv_{\max}\tilde{T}^{-1},$$

where $u_s^{\tilde{T}}(A_s;\mathrm{NMBR}) = \sum_{t=\gamma T}^{T} v_s(\hat{\theta}_s,, X(\hat{\theta}_s, \theta_{-s};\mathrm{NMBR}))$ *is the utility agent s derives from playing strategy* A_s *against truthful opponents in the NMBR mechanism, starting with round* $t = \gamma T$.

Proof. Note that after γT periods, the remaining game is equivalent to the original one, with budgets $B_s^{\gamma T} = B_s^0 - \sum_{t=1}^{\gamma T} C_s(\theta; M_{\mathrm{money}})$ and $\tilde{T}$ rounds remaining. Thus, we can reduce the statement of the theorem to that of the Theorem 1, if we can show that, with high probability, $B_s^{\gamma t} = \bar{c}_s \tilde{T}(1 + \tilde{\delta}_s)$ for some $\tilde{\delta}_s$.

In particular, we can use the Chernoff bound to see that

$$\mathbb{P}\left[\sum_{t=1}^{\gamma T} C_s(\theta^t; M_{\mathrm{money}}) \in [\gamma T \bar{c}_s(1-\delta_s), \gamma T \bar{c}_s(1+\delta_s)]\right] \geq 1 - 2e^{-\frac{\delta_s^2 R_s \gamma T}{3}}.$$

Also, because the initial budget is set to $B_s = \bar{c}_s T(1 + \delta_s)$, the remaining budget at turn $t = \gamma T$ is given by

$$B_s^{\gamma T} = \bar{c}_s T(1+\delta_s) - \sum_{t=1}^{\gamma T} C_s(\theta; M_{\mathrm{money}}).$$

Combining the Chernoff bound with this, we get that $B_s^{\gamma T} = \bar{c}_s \tilde{T}(1 + \tilde{\delta}_s)$, where

$$\tilde{\delta}_s \in \left[\delta_s, \left(1 + \frac{2\gamma}{1-\gamma}\right)\delta_s\right] \quad \text{with probability at least } 1 - 2e^{-\frac{\delta_s^2 R_s \gamma T}{3}}.$$

We can now reduce the proof of the theorem to that of Theorem 1. To do this, we just need to notice that, to upper bound $\mathbb{P}[\mathcal{E}_s]$, we can use the lower bound for the value of $\tilde{\delta}_s$ (because the probability to run out of credits is highest when the budget is the smallest). To upper bound gains from deviations, we use the upper bound for $\tilde{\delta}_s$, as the bound given by Lemma 4 is widest when the budget surplus is large. We can now substitute the correct values of $\tilde{\delta}_s$ in place of δ_s in the statement of Theorem 1 to obtain the needed result. □

Endnotes

[1] Throughout, we use the notation $[m]$ to denote the set $\{1, 2, \ldots, m\}$ for any positive integer m.

[2] Here, we define welfare as the sum of the agents' utilities.

References

[1] Alaei S (2014) Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM J. Comput.* 43(2):930–972.
[2] Azevedo EM, Budish E (2017) Strategy-proofness in the large. Technical report, National Bureau of Economic Research, Cambridge, MA.
[3] Balseiro SR, Gur Y (2019) Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Sci.* 65(9): 3952–3968.
[4] Balseiro SR, Besbes O, Weintraub GY (2015) Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Sci.* 61(4):864–884.
[5] Balseiro S, Gurkan H, Sun P (2019) Multi-agent mechanism design without money. *Oper. Res.* 67(5):1417–1436.
[6] Bergemann D, Said M (2011) Dynamic auctions. *Wiley Encyclopedia of Operations Research and Management Science* (Wiley, Hoboken, NJ).
[7] Bergemann D, Välimäki J (2010) The dynamic pivot mechanism. *Econometrica* 78(2):771–789.
[8] Brânzei S, Procaccia AD (2015) Verifiably truthful mechanisms. *Proc. 2015 Conf. Innovations Theoret. Comput. Sci.* (Association for Computing Machinery, New York), 297–306.
[9] Budish E (2011) The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *J. Political Econom.* 119(6):1061–1103.
[10] Budish EB, Cachon G, Kessler JB, Othman A (2017) Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Oper. Res.* 65(2):314–336.
[11] Cavallo R (2014) Incentive compatible two-tiered resource allocation without money. *Proc.* 2014 *Internat. Conf. Autonomous Agents Multiagent Systems* (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC), 1313–1320.
[12] Chawla S, Miller JB (2016) Mechanism design for subadditive agents via an ex ante relaxation. *Proc. 2016 ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 579–596.
[13] Cole R, Gkatzelis V (2015) Approximating the Nash social welfare with indivisible items. *Proc. 47th Annual ACM Sympos. Theory Comput.* (Association for Computing Machinery, New York), 371–380.
[14] d'Aspremont C, Gérard-Varet LA (1979) Incentives and incomplete information. *J. Public Econom.* 11(1):25–45.
[15] Devanur NR, Peres Y, Sivan B (2015) Perfect Bayesian equilibria in repeated sales. *Proc. 26th Annual* ACM-*SIAM Sympos. Discrete Algorithms* (Society for Industrial and Applied Mathematics, Philadelphia), 983–1002.
[16] Dubhashi DP, Panconesi A (2009) *Concentration of Measure for the Analysis of Randomized Algorithms* (Cambridge University Press, Cambridge, UK).
[17] Dughmi S, Ghosh A (2010) Truthful assignment without money. *Proc. 11th ACM Conf. Electronic Commerce* (Association for Computing Machinery, New York), 325–334.
[18] Fudenberg D, Tirole J (1991) *Game Theory* (MIT Press, Cambridge, MA).
[19] Gershkov A, Moldovanu B (2009) Dynamic revenue maximization with heterogeneous objects: A mechanism design approach. *Amer. Econom. J. Microeconom.* 1(2):168–198.

[20] Goel A, Mahdian M, Nazerzadeh H, Saberi A (2010) Advertisement allocation for generalized second-pricing schemes. *Oper. Res. Lett.* 38(6):571–576.

[21] Gorokh A, Banerjee S, Iyer K (2016) Near-efficient allocation using artificial currency in repeated settings. Preprint, submitted October 15, http://dx.doi.org/10.2139/ssrn.2852895.

[22] Gorokh A, Banerjee S, Iyer K (2017) From monetary to non-monetary mechanism design via artificial currencies. *Proc. 2017 ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 563–564.

[23] Gummadi R, Key P, Proutiere A (2012) Repeated auctions under budget constraints: Optimal bidding strategies and equilibria. Preprint, https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.1046.6525&rep=rep1&type=pdf.

[24] Guo M, Conitzer V (2010) Strategy-proof allocation of multiple items between two agents without payments or priors. *Proc. Ninth Internat. Conf. Autonomous Agents Multiagent Systems*, vol. 1 (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC), 881–888.

[25] Guo M, Conitzer V, Reeves DM (2009) Competitive repeated allocation without payments. Leonardi S, ed. *Internet and Network Economics.* Lecture Notes in Computer Science, vol. 5929 (Springer, Berlin), 244–255.

[26] Hartline JD, Roughgarden T (2008) Optimal mechanism design and money burning. *Proc. 40th Annual ACM Sympos. Theory Comput.* (Association for Computing Machinery, New York), 75–84.

[27] Hartline JD, Kleinberg R, Malekian A (2011) Bayesian incentive compatibility via matchings. *Proc.* 22*nd Annual ACM-SIAM Sympos. Discrete Algorithms* (Society for Industrial and Applied Mathematics, Philadelphia).

[28] Iyer K, Johari R, Sundararajan M (2014) Mean field equilibria of dynamic auctions with learning. *Management Sci.* 60(12):2949–2970.

[29] Jackson MO, Sonnenschein HF (2007) Overcoming incentive constraints by linking decisions. *Econometrica* 75(1):241–257.

[30] Kanoria Y, Nazerzadeh H (2014) Dynamic reserve prices for repeated auctions: Learning from bids. *WINE 2014* (Springer, Cham, Switzerland), 232.

[31] Karande C, Mehta A, Srikant R (2013) Optimizing budget constrained spend in search advertising. *Proc. Sixth ACM Internat. Conf. Web Search Data Mining* (Association for Computing Machinery, New York), 697–706.

[32] Leme RP, Syrgkanis V, Tardos É (2012) Sequential auctions and externalities. *Proc. 23rd Annual ACM-SIAM Sympos. Discrete Algorithms* (Society for Industrial and Applied Mathematics, Philadelphia), 869–886.

[33] Milgrom PR (2004) *Putting Auction Theory to Work* (Cambridge University Press, Cambridge, UK).

[34] Myerson RB (1981) Optimal auction design. *Math. Oper. Res.* 6(1):58–73.

[35] Nazerzadeh H, Saberi A, Vohra R (2008) Dynamic cost-per-action mechanisms and applications to online advertising. *Proc.* 17*th Internat. Conf. World Wide Web* (Association for Computing Machinery, New York), 179–188.

[36] Nekipelov D, Syrgkanis V, Tardos E (2015) Econometrics for learning agents. *Proc. 16th ACM Conf. Econom. Comput*. (Association for Computing Machinery, New York), 1–18.

[37] Prendergast C (2016) The allocation of food to food banks. *EAI Endorsed Trans. Serious Games* 3(10):e4.

[38] Procaccia AD, Tennenholtz M (2013) Approximate mechanism design without money. *ACM Trans. Econom. Comput.* 1(4):18.

[39] Roughgarden T (2012) The price of anarchy in games of incomplete information. *Proc. 13th ACM Conf. Electronic Commerce* (Association for Computing Machinery, New York), 862–879.

[40] Walsh T (2014) Allocation in practice. *Joint German/Austrian Conf. Artificial Intelligence* (Springer, Berlin), 13–24.

[41] Yan Q (2011) Mechanism design via correlation gap. *Proc.* 22*nd Annual ACM-SIAM Sympos. Discrete Algorithms* (Society for Industrial and Applied Mathematics, Philadelphia), 710–719.