FAIRNESS AND EFFICIENCY IN ONLINE ALLOCATION OF GOODS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Artur Gorokh

May 2020

FAIRNESS AND EFFICIENCY IN ONLINE ALLOCATION OF GOODS

Artur Gorokh, Ph.D.

Cornell University 2020

The subject of this thesis is the problem of allocating goods to people without using monetary payments. Specifically, I concentrate on the case of repeated allocation of goods, a common problem in practice, as many resources are allocated to recipients on monthly or daily basis. Throughout this work, we show how repeated nature of allocation can be leveraged to achieve strong efficiency, fairness and incentive guarantees via well-chosen mechanisms. The work presented here is based on three papers on the topic written by me and my collaborators. When taken together, these works also outline the theoretical trade-off between various assumptions on the allocation setting and the strength of the resulting guarantees.

This document is dedicated to Maria Antonovna Gorokh.

ACKNOWLEDGEMENTS

This work would be impossible without the guidance and friendship of my advisors. Collaborating with Sid and Kris for the last 4 years has been a great joy, most of the time. It was taking Sid's course that made me fall in love with the field of mechanism design which I at the time did not even know existed. His mentorship helped me understand and navigate the world of academia, and his unremitting optimism gave me the confidence and the will to keep going when I would have otherwise given up. I owe an equal amount of gratitude to Kris, who apart from becoming a good friend, taught me rigor and patience I desperately lack. It is a rare joy to collaborate with someone who fully understands your ideas without you saying as much as a few words, and I feel lucky to have experienced this. Finally, a significant chunk of this dissertation is work I have done with Vasilis Gkatzelis, who introduced me to the problem and since proved to be a superb collaborator.

I also want to thank numerous friends at the Center for Applied Mathematics, who throughout the years supported me, heard me vent about my work, and have given me great advice. PhD research can be lonely business, and it is thanks to CAM that for me it wasn't. I would especially like to thank Matt Davidow, who practically always knew the exact thing I was working on, occasionally helped me attack research problems, and has been a great company throughout the years.

Finally, I want to thank my parents, Anzhelika and Anatoly Gorokh, who have supported me in every way since as long as I can remember, and always encouraged me to pursue my passions, even when these passions took me across the globe.

BIOGRAPHICAL SKETCH

Artur Gorokh was born in 1992 in Moscow, Russia. He completed his Specialist degree in Mathematical Physics at the Physics Department of Lomonosov Moscow State University in 2015. In the same year he started working on his PhD in the Center for Applied Mathematics at Cornell University, conducting research in game theory and mechanism design with his advisors Sid Banerjee and Krishnamurthy Iyer.

TABLE OF CONTENTS

CHAPTER 1

INTRODUCTION

Who gets what, and why? This question has been the heart and soul of mechanism design, the field concerned with making decisions in presence of strategic participants. In decades since its founding, the field has seen many successes. Some of these are theoretical, showing the often surprising boundaries of what is possible in the world of selfish agents. Other ideas have directly influenced the economy and the lives of real people by bringing carefully thought-out designs to settings where ad hoc procedures used to rule.

One particular area in which mechanism design has seen a lot of theoretical and practical success is the allocation of goods without money, that is, situations where a principal has to distribute resources without charging participants monetary payments. The practical examples of such successes include residency and school matching systems, allocation of food to food banks or course seats to students. Whenever brought to scale, these innovations have demonstrated their potential to greatly improve human welfare, and there are probably many more problems where further improvements are possible with appropriate mechanisms.

The problem of allocating goods to people without using monetary payments is the subject of this thesis. The work documented here is based on three papers on the topic written by me and my collaborators. All three are concerned with a particular yet very common scenario, dynamic allocation of goods, i.e. allocating resources repeatedly through time. Repeated allocation is frequent in practice, be that an non-profit organization distributing resources between people on a monthly basis, or a tech corporation distributing computing resources to employees every minute.

The reason repeated allocation of goods deserves separate treatment is that, in comparison with one-shot allocation, repeated allocation gives the designer more degrees of freedom which in turn can allow for stronger guarantees than those one can derive by repeatedly employing one-shot allocation results. To illustrate, when allocating a single item only once, the goals of efficiency (giving the item to the agent with highest value) and fairness (making sure agents are treated equally) are in direct conflict. When the item is allocated many times, this conflict between efficiency and fairness goes away, as different agents end up having highest value on different rounds, and goals of efficiency and fairness are entirely compatible (assuming agents are equally likely to have highest value on any round). Similarly, repeated allocation can allow for incentive compatible mechanisms that would not work in a one shot setting (this was first noticed in [50]).

All the results presented here are theoretical, and as with any theoretical work, the assumptions matter: assuming more allows for stronger results, but potentially renders them less practically relevant. In the context of mechanism design, these assumptions are typically either about the information availability (what does the designer and the participants know?), the preferences of the participants (how are preferences of participants distributed?), or the allocation space (what is being allocated?). It is instructive to think of the three chapters that follow as taking different choices in picking the assumptions along dimensions described above and consequentially obtaining results of varying theoretical and practical value.

1.1 Overview

Below we provide a brief summary of the chapters in this thesis, discuss their main findings and the different trade-offs in assumptions they take.

To outline our results, it is helpful to first present a simple model, which can serve as an example and make the results we state below more concrete. Consider the following allocation problem: there are n agents and T rounds with a single item available on each round, and on every round agent i has value v_{it} for the item. Agent's i utility is given by $u_i = \sum_t x_{it} v_{it}$, where x_{it} is the indicator agent i was awarded the item on round t. On every round, a mechanism in this setting first communicates with the agents and then picks x_{it}, without necessarily knowing values v_{it} (as the agents may not want to disclose those truthfully).

The goal of allocating efficiently in this setting can be formalized in several ways. First, one might ask the mechanism to implement a Pareto optimal outcome, i.e., an allocation such that no agent's utility can be improved without decreasing utility of some other agent. This requirement leaves many degrees of freedom, as there are typically many Pareto optimal allocations to choose from. In this work we consider two primary ways of specifying a particular Pareto optimal outcome of interest. Maximizing utilitarian social welfare, defined as sum of agent's utilities $W = \sum_i u_i$, leads to a Pareto optimal outcome where, in the setting described above, the item always goes to the agent with highest value. Alternatively, one can allocate as to maximize Nash Social Welfare (NSW) defined as geometric mean of agents' utilities $NSW = \left(\prod_i u_i\right)^{1/n}$ (this would also produce to a Pareto optimal allocation). In contrast to utilitarian social welfare, maximizing this objective leads to more egalitarian outcomes, as an allocation that fails to allocate enough to even a single agent can suffer significant NSW loss compared to the optimal allocation.

In Chapter 2(based on [39]), we make strong assumptions about both information availability as well as preferences of the participants. In particular, we assume that values of agents for the allocated resources are independently and randomly drawn in each period, and that the designer knows these distributions of values(or at least has a certain number of samples from them). On the other hand, we make virtually no assumptions on the allocation problem, so the results are applicable to a variety of allocation settings, such as general combinatorial assignment problems or public choice settings.

Under these assumptions we prove a rather strong result: we construct a black-box reduction that converts any incentive compatible one-shot mechanism with money to a dynamic non-monetary mechanism which is approximately incentive compatible and approximately as a efficient as the monetary mechanism. More precisely, we show that under the constructed reduction per-round incentives to deviate as well as per-round utilitarian welfare loss go to zero as the number of rounds is increased.

The assumption of full information about value distributions may not be practical in most settings, and the result is mostly of theoretical value. What it tells us is that there is little gap between monetary and non-monetary mechanisms in a setting where such information is available.

In Chapter 3(based on [40]), we explore whether results similar to the ones described above are possible under more practical informational assumptions and with the use of simple practical mechanisms. We still assume that distribution for values of agents is known to the agent herself, but not necessarily to the designer or the other agents.

We consider the following mechanism: the participants are allocated budgets of artificial currency in the beginning of the game, and the items are repeatedly allocated with first price auction. This artificial currency mechanism is simple to explain and practical to implement, and in fact has been occasionally employed in real life: a mechanism nearly identical to the one we outlined has been used by Feeding America, a non-profit organization managing the allocation of donations to food banks [70]. However, to the best of our knowledge, no prior work has addressed the efficiency properties of this or similar mechanisms while accounting for the incentives of the participants; the aim of this work is to fill this gap.

For the mechanism described above, we show that, independently of the strategies adopted by the others, every agent can secure a 0.5-approximation of their expected utility at a Pareto optimal outcome chosen by the designer [1]. To compare it with Chapter 2, the efficiency guarantee is weaker, as here we only guarantee half of possibly achievable efficiency. In return for weaker efficiency guarantees we gain two features: first, we employ considerably weaker and more practical informational assumptions, and secondly, a much stronger incentive guarantee. In particular, we assume the designer does not necessarily know anything about the value distributions of the agents, and the utility an agent can secure does not depend on other agents converging to an equilibrium or being rational at all. These two features make the results much more robust which in turn should make the guarantees more convincing to practitioners considering adoption of such a mechanism. The results also suggest how the principal can set budgets in order to pick a desired Pareto-optimal outcome to approximate.

Finally, in Chapter 4(based on [41]),we tackle the same problem of dynamically allocating items to agents, and take an even further step in relaxing assumptions on information available to the principal – we now assume that the values agents have for the items are arbitrary (i.e., adversarial). Under such weak assumption on the

[1] In fact, instead of Pareto optimality we use a somewhat stronger notion of ideal utilities, which we introduce in the corresponding chapter.

values, the problem of allocating efficiently is complex even in a single item setting and without the consideration of incentives, so this work assumes the participants report their values truthfully and concentrates on the algorithmic aspect of the problem.

In this setting, we set out to construct an online algorithm that allocates items to maximize Nash Social Welfare (NSW), a popular efficiency and fairness metric defined as geometric mean of agents' utilities. We then show two complementary results. First, we construct a simple online algorithm that guarantees logarithmic competitive ratio for NSW, when parametrizing the bound either in the number of agents or the number of rounds. We also show a matching impossibility result, i.e., we prove that an asymptotically better competitive ratio cannot be achieved by any online algorithm. We complement these results with several generalizations. In particular, we show that the competitive ratio is strictly better when the values of agents are not much bigger than the mean, and we demonstrate how our results can be generalized to settings with indivisible items and several items per round.

CHAPTER 2

FROM MONETARY TO NON-MONETARY MECHANISMS VIA ARTIFICIAL CURRENCIES

2.1 Introduction

Today, many real-world settings require non-monetary 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 ([44, 6]). However, the resulting mechanisms are complex and are hard to extend to more general settings.

In contrast, a class of simple non-monetary mechanisms, namely mechanisms based on *artificial currencies*, has gained recent attention due to their successful use in practice (e.g., in food banks ([69])). 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 non-monetary 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 non-monetary 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.

2.1.1 Main results and roadmap

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 multi-dimensional – 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 4.6).

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 non-monetary mechanism M that mirrors the properties of a given monetary mechanism $M_{\textsf{money}}$ 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 $M_{\textsf{money}}$. However, doing this perfectly is in general impossible: for $T = 1$ and general type distributions F_s, no non-monetary 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 2 agents with i.i.d. Uniform$[0, 1]$ values).

Faced with this difficulty, we relax our objective and instead seek a non-monetary mechanism that approximately preserves the incentive and utility guarantees of a monetary mechanism. Formally, we say a non-monetary mechanism M is an (α, β)-approximation of a given *Bayesian incentive compatible* (BIC) mechanism $M_{\textsf{money}}$ if it simultaneously guarantees the following two properties (cf. Definition 2.2.3):

1. For any agent, the gains from unilaterally deviating from truthful reporting under M is at most αT. Since $M_{\textsf{money}}$ is BIC, this ensures that M approximately matches the incentive properties of $M_{\textsf{money}}$.
2. Assuming truthful reporting, the mechanism M guarantees the same expected utility for every agent as $M_{\textsf{money}}$ (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*, i.e. $o(1)$, 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 2.2.4 for more details).

Our work answers this question in the affirmative. In particular, we construct two *blackbox* approaches (NMBR and REAP) that each take any monetary mechanism M_{money}, and return a non-monetary 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* (cf. Assumption 2.2.1), where the mechanism is able to exclude an agent without affecting other agents, we provide a simple Non-Monetary Blackbox Reduction (which we refer to as the NMBR mechanism; see Section 2.3). Informally, our recipe comprises of 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 blackbox reduction that, given a monetary mechanism, produces a non-monetary mechanism, for which we prove the desired

α, β approximation guarantee (Theorems 2.3.1 and 2.3.6). 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 2.4 we drop the assumption of excludability and construct alternative blackbox mechanism, 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 2.4.2); 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 REAP mechanism can be strengthened for the case when there are 2 agents, and the 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 2 agents (Theorem 2.5.1). We also leverage this result to show the *existence* of a (potentially complex) $(0, o(1))$-approximate mechanism (Corollary 2.5.7).

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 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 ([72]). Moreover, the modular nature of our proof allows us to discuss applicability of our technique to settings more general than the one considered in the paper. We discuss several such potential extensions in Section 4.6.

2.1.2 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 which is closest to our setting.

Dynamic mechanism design focuses on extending the theory of mechanisms for single period settings ([64, 61]) to repeated allocation settings. The difficulty in doing so arises due factors which couple auctions across time; for example, incomplete information and learning over time ([29, 53, 67, 49]), cross-period combinatorial constraints including limits imposed by budgets ([66, 7, 42, 58]), stochastic fluctuations in the underlying setting ([37, 13, 14]), etc. In our setting, the cross-period coupling arises essentially due to budget constraints. Similar repeated auctions with budget constraints have been considered by [42] and [7], 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 to existing approaches.

Mechanism design without money: This is a broad area of study, encompassing diverse topics ranging from matching theory to social choice, which broadly considers strategic allocation in settings where money is not permitted due to 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 [71], who used the specific example of a facility location game; this approach was subsequently explored by many others ([43, 31]). Several alternate approaches have also been proposed, including ones based on verifiability [18], proof-of-work ('money burning') ([47]), and two-tiered resource redistribution ([23]).

In the case of repeated allocation without money, a notable line of work is that by [44], and its subsequent refinement by [6]. These works consider an identical model as ours, but with a single-item and symmetric agents, and under discounted infinite-horizon settings. The former uses variations of the AGV mechanism [28] 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 1. 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.

Artificial currency mechanisms: These are a particular subset of non-monetary mechanisms which have attracted a lot of recent interest, in part due to recent successful implementations in real-world settings ([21, 74, 69]). Our work follows in the line of several recent papers in attempting to establish a theoretical foundations for such mechanisms. Among these, the closest to us is the work of [19] and [50]; we describe these now in more detail.

[19] studies the use of artificial currency mechanisms in the context of static combinatorial assignment problems under (arbitrary) ordinal preferences. As in our work, incentive compatibility constraint is relaxed in order 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 2.2.3. Similar guarantees are also established in static settings for additive valuations via maximizing Nash Welfare by [26]. In contrast to these static mechanisms, our approach applies primarily for dynamic settings by exploiting future allocations to ensure approximate incentive compatibility.

[50] (henceforth JS) are also concerned with a one-shot allocation problems, but unlike Budish, they consider settings where agents simultaneously participate in multiple resource allocation problems, with iid 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 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 JS; this temporal aspect makes the analysis of incentives more challenging in our setting. Next, since JS make use of a separate budget for every possible report, the number of linked instances required for non-trivial 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 non-truthful bidding going to 0 for large T), while the guarantees in JS hold under any BNE. We note though that in the case of 2 agents, we are able to match the incentive guarantees of JS using our simpler single-currency mechanism.

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 [1] in monetary setting to establish a blackbox reduction from single-agent to multi-agent mechanism with item supply constraint. The same technique was later adopted by [25] in the context approximating revenue for selling multiple items to several heterogeneous buyers. Another related technique used in revenue maximization problems is that of correlation gap, see [75].

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, [8] consider the problem of regret minimization from the bidders' perspective and demonstrate strategies that constitute approximate equilibrium. [54] design a system to optimize bids in large repeated ad auctions with budgets. Lastly, [38] model various mechanisms to throttle the bids of budgeted agents across time to maximize revenue.

2.2 Model

2.2.1 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 iid 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

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

uniformly bounded by $v_{\max}$.

For our results in Section 2.3, we require that the allocation setting satisfies the following *excludability* assumption.

Assumption 2.2.1 (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 case of provision of public non-rival goods. (See Appendix 2.7 for a more detailed description of settings that do or do not satisfy excludability.)

2.2.2 Non-monetary mechanisms, agents' strategies and utilities

We focus on settings where the principal employs a non-monetary, *direct-revelation* mechanism to select the allocation X^t at each time t. Formally, a non-monetary direct-revelation mechanism M requires each agent s to submit a report $\hat{\theta}_s^t \in \Theta_s$ at each time t. Let $\mathcal{H}^{\sqcup} = \{(\mathcal{X}^\tau, \hat{\theta}^\tau_\infty, \cdots, \hat{\theta}^\tau_\backslash)\}_{\tau<\sqcup}$ denote the public history, i.e., 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 $X^t = X^t(\hat{\theta}^t, \mathcal{H}^{\sqcup}; \mathcal{M})$ 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 $\mathcal{H}^{\sqcup}$. We denote the strategy of truthful reporting as Tr, i.e., 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:

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

where $X^t = X^t(\hat{\theta}^t, \mathcal{H}^{\sqcup}; \mathcal{M})$. Observe $U_s(A; M)$ 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.

2.2.3 Static setting and monetary mechanisms

To design and benchmark non-monetary 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

$X(\hat{\theta}_1, \ldots, \hat{\theta}_n; M_{\mathsf{money}}) \in \mathcal{X}$, and payment rules $\left\{C_s(\hat{\theta}_1, \ldots, \hat{\theta}_n; M_{\mathsf{money}}) : s \in [n]\right\}$. Specifically, given the agents' report $\hat{\theta} = (\hat{\theta}_1, \cdots, \hat{\theta}_n)$, the mechanism chooses an allocation $X(\hat{\theta}; M_{\mathsf{money}})$ and charges each agent s a monetary payment of $C_s(\hat{\theta}; M_{\mathsf{money}})$. We assume that the agents have quasilinear utilities, i.e., the utility of the agent s is given by $v_s(\theta_s, X(\hat{\theta}; M_{\mathsf{money}})) - C_s(\hat{\theta}; M_{\mathsf{money}})$. We also assume that all monetary mechanisms employed in this paper have non-negative payments $C_s(\hat{\theta}; M_{\mathsf{money}}) \geq 0$ for all $\hat{\theta}$.

We say a direct revelation mechanism M_{money} is *Bayesian-incentive compatible* (BIC) if reporting truthfully is a Bayes-Nash equilibrium. That is, for each agent s and for all θ_s, we have

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

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, i.e., $u_s(\mathsf{Tr}; M_{\mathsf{money}}) = \mathbb{E}[v_s(\theta_s, X(\theta; M_{\mathsf{money}}))]$ 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 2.2.2 (Opt-out report)**.** *Under mechanism M_{money}, every agent s has a report $\emptyset$ that guarantees zero payment $C_s(\emptyset, \hat{\theta}_{-s}; M_{\mathsf{money}}) = 0$ and an allocation $X(\emptyset, \hat{\theta}_{-s}; M_{\mathsf{money}})$ such that $v_s(\theta_s, X(\emptyset, \hat{\theta}_{-s}; M_{\mathsf{money}})) = 0$ for any type θ_s and reports $\hat{\theta}_{-s}$.*

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

2.2.4 Design requirements

Returning to the repeated allocation setting, we seek to design non-monetary 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 non-monetary mechanism M captures the incentives and utility profile of a given monetary mechanism M_{money}.

Definition 2.2.3. *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:*
$$\frac{u_s(\mathsf{Tr}; M)}{T} \geq u_s(\mathsf{Tr}; M_{\mathsf{money}}) - \beta.$$

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 only willing to deviate from truthful reporting only if the gains from doing so are significant when compared to 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 in order 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 can yield. A similar argument has been used in other works (see for example [5, 67]).

2.3 Blackbox reduction from monetary mechanisms

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

2.3.1 The Non-monetary Blackbox Reduction (NMBR) mechanism

We now describe our blackbox reduction technique, that converts any chosen monetary BIC mechanism M_{money} to a non-monetary 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 of 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 2.2.1) 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.

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 2.3.1. *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, as-*

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.

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 $\mathsf{ACT} = [n]$.

for all $t = 1$ to T **do**

  For each agent $s \in \mathsf{ACT}$, get her report $\hat{\theta}_s^t$.

  For each agent $s \notin \mathsf{ACT}$, sample a report $\hat{\theta}_s^t \sim F_s$ from her type distribution.

  For each agent $s \in \mathsf{ACT}$, if $B_s^t \geq C_s(\hat{\theta}^t; M_{\mathsf{money}})$, charge her the payment $C_s(\hat{\theta}^t; M_{\mathsf{money}})$, and set $B_s^{t+1} = B_s^t - C_s(\hat{\theta}^t; M_{\mathsf{money}})$.
  If $B_s^t < C_s(\hat{\theta}^t; M_{\mathsf{money}})$, update $\mathsf{ACT} \leftarrow \mathsf{ACT} \setminus \{s\}$ and $B_s^{t+1} = 0$.

  Let $X^t = X(\hat{\theta}^t; M_{\mathsf{money}})$. Implement $X^t|_{\mathsf{ACT}}$ (refer Assumption 2.2.1).

end for

suming all other agents play truthfully, we have,

$$\frac{u_s(\mathsf{Tr}; \mathrm{NMBR})}{T} \geq \frac{u_s(A_s; \mathrm{NMBR})}{T} - \sqrt{\frac{3\bar{c}_s c_{\max} \log T}{T}} - 2v_{\max} T^{-1}. \tag{2.1}$$

Additionally, the utility of each agent s satisfies,

$$\frac{u_s(\mathsf{Tr}; \mathrm{NMBR})}{T} \geq u_s(\mathsf{Tr}; M_{\mathit{money}}) - 2v_{\max} T^{-1}. \tag{2.2}$$

In other words, 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 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 2.3.1 only establishes truthful reporting to be approximately optimal ex ante, it is possible to demonstrate that similar guarantees hold in later rounds, and degrade towards the end of the game. For more detail on this see Section 2.8.

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

Example 2.3.2. 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 2.3.1 guarantees* *the following:*

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

Although gains from deviations decrease with n, because 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 2.3.1 thus gives a general recipe for converting any one-shot monetary mechanism into a non-monetary 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 2.3.3, 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).

2.3.2 Proof of Theorem 2.3.1

In this section, we outline the overall strategy for proving Theorem 2.3.1, 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 Section 3.7.2.

The proof of Theorem 2.3.1 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 (cf. Lemma 2.3.5 below). 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 budgets. 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

$$\widehat{u}_s(A_s) \triangleq \mathbb{E}\left[\sum_{t=1}^{T} v_s\left(\theta_s^t, X(\hat{\theta}_s^t, \theta_{-s}^t; M_{\mathsf{money}})\right)\right], \tag{2.3}$$

where in 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 the any randomness in the strategy A_s. Similarly, under the strategy A_s, the expected spending of agent s is defined as

$$\widehat{c}_s(A_s) \triangleq \mathbb{E}\left[\sum_{t=1}^{T} C_s(\hat{\theta}_s^t, \theta_{-s}^t; M_{\mathsf{money}})\right]. \tag{2.4}$$

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

$$\begin{aligned} \max_{A_s} \quad & \widehat{u}_s(A_s) \\ \text{s.t. } & \widehat{c}_s(A_s) \leq B_s. \end{aligned} \tag{2.5}$$

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 (2.5) is indeed a relaxation of the 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 2.2.2 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, since the budget constraint holds sample-path wise in the NMBR mechanism, the expected budget constraint in (2.5) 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 (2.5) to that of the truthful report Tr. First, note that by definition of $\bar{c}_s$, since $B_s > \bar{c}_s T$, truthful reporting Tr is feasible for (2.5). 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 2.3.3 (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 (2.5), we have*

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

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

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

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 (2.4) implies $\widehat{c}_s(A_s) \leq B_s = \bar{c}_s T + \Delta$, 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 $\Delta \geq 0$.

Next, coming back to the NMBR mechanism, in order 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:

$$\begin{aligned}
&u_s(A_s; \text{NMBR}) - u_s(\mathsf{Tr}; \text{NMBR}) \\
&\quad = u_s(A_s^{\emptyset}; \text{NMBR}) - u_s(\mathsf{Tr}; \text{NMBR}) \\
&\quad = u_s(A_s^{\emptyset}; \text{NMBR}) - \widehat{u}_s(A_s^{\emptyset}) + \widehat{u}_s(A_s^{\emptyset}) - \widehat{u}_s(\mathsf{Tr}) + \widehat{u}_s(\mathsf{Tr}) - u_s(\mathsf{Tr}; \text{NMBR}) \\
&\quad = \underbrace{\widehat{u}_s(A_s^{\emptyset}) - \widehat{u}_s(\mathsf{Tr})}_{\dagger} + \underbrace{\widehat{u}_s(\mathsf{Tr}) - u_s(\mathsf{Tr}; \text{NMBR})}_{\star},
\end{aligned} \tag{2.6}$$

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 2.3.3 gives a bound on the term (†), since $A_s^{\emptyset}$ is feasible for (2.5). Thus, to complete our proof, we seek to bound the term ($\star$), i.e., 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* $\mathcal{E}_s$:

$$\mathcal{E}_s \triangleq \left\{ \sum_{1}^{T} C_s(\theta_s^t, \theta_{-s}^t; M_{\mathsf{money}}) \geq B_s \right\}. \tag{2.7}$$

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 2.3.4.

$$\widehat{u}_s(\mathsf{Tr}) \leq u_s(\mathsf{Tr}; \text{NMBR}) + 2\mathbb{P}(\mathcal{E}_s) T v_{\max}$$

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

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

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

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

Note that we have $U_s(\mathsf{Tr}; \text{NMBR}) = \widehat{U}_s(\mathsf{Tr})$ 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 $|U_s(\mathsf{Tr}; \text{NMBR}) - \widehat{U}_s(\mathsf{Tr})| \leq 2Tv_{\max}$. Substituting this into the decomposition above yields the needed bound. □

Putting Lemma 2.3.3 and Lemma 2.3.4 together, along with the fact that $\Delta = \delta_s \bar{c}_s T$ for the NMBR mechanism, we obtain from (2.6) that

$$u_s(A_s; \text{NMBR}) - u_s(\mathsf{Tr}; \text{NMBR}) \leq \delta_s \bar{c}_s T + 2Tv_{\max}\mathbb{P}(\mathcal{E}_s). \tag{2.8}$$

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 2.3.5 (Concentration of spending)**.** *For every agent s, we have*

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

Proof. 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 (cf. [30], Ch 1 for more details): For X_i independent r.v.s with $0 \leq X_i \leq 1$, and $X = \sum_i X_i$, we have

$$\mathbb{P}[X \geq \mathbb{E}[X](1 + \epsilon)] \leq \exp\left(-\frac{\epsilon^2 \mathbb{E}[X]}{3}\right)$$

□

Thus, we obtain from Lemma 2.3.5 and (2.8) that

$$u_s(A_s; \text{NMBR}) - u_s(\mathsf{Tr}; \text{NMBR}) \leq \delta_s \bar{c}_s T + 2Tv_{\max} \exp\left(-\frac{\delta_s^2 R_s T}{3}\right). \tag{2.9}$$

As we increase δ_s, the gains from potential deviations in the auxiliary game increase (as in Lemma 2.3.3), but the gap in the performance of truthful reporting between the original and auxiliary games decreases (via Lemmas 2.3.4 and 2.3.5). Choosing $\delta_s = \sqrt{\frac{3 \log T}{R_s T}}$ balances this trade-off, thereby establishing approximate incentive compatibility guarantee (2.1). Similarly, from Lemma 2.3.4 and (2.8), after substituting the aforementioned value of δ_s, we have

$$u_s(\mathsf{Tr}; \mathrm{NMBR}) \geq \widehat{u}_s(\mathsf{Tr}) - 2v_{\max} T^{-1}. \tag{2.10}$$

Observe that $\widehat{u}_s(\mathsf{Tr}) = T u_s(\mathsf{Tr}; M_{\mathsf{money}})$, because the auxiliary problem involves repeating M_{money} for T time periods. Upon dividing by T, we obtain the utility guarantee (2.2).

2.3.3 Tractable blackbox 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 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 $\left\{\theta_s^{(1)}, \theta_s^{(2)}, \ldots, \theta_s^{(m)}\right\}$ 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 $\left\{\theta_s^{(k)} : s \in [n]\right\}$. 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 NMBR mechanism is preserved with high probability.

Theorem 2.3.6. *Consider* 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 2.3.1) 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*:*

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

The proof of this theorem relies on showing a concentration bound for $\tilde{B}_s$ and then repeating the argument employed for proving Theorem 2.3.1. For brevity, we defer the proof to the appendix.

The bounds provided in this theorem are asymptotically equivalent to those in Theorem 2.3.1, 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 Section 3.7.2. 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 2.3.7 (Computational cost of 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 re-*

quires $2T$ samples for every type distribution F_s, and has a computational cost of $2TC_{\textsf{money}} + O(Tn)$.

Proof. Proof. Given a mechanism $M_{\textsf{money}}$, 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 $M_{\textsf{money}}$ for T times (once per period) with the actual reports of the agents, and in the worst case 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 $M_{\textsf{money}}$. □

2.4 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 2.2.1); 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 2.3.1 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 $M_{\textsf{money}}$, 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:

$$c_s(\theta_s) \triangleq \mathbb{E}[C_s(\theta_s, \theta_{-s}; M_{\textsf{money}})]. \tag{2.11}$$

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

Example 2.4.1 (Single item allocation)**.** *Recall the setting we used to illustrate the* NMBR *mechanism in the previous section: there are n agents with values $v_{it} \sim$ $\textsf{Unif}[0,1]$, and $M_{\textsf{money}}$ is the second-price auction. In this case, the personalized price* (2.11) *for reporting some value $\hat{v}_s$ in the REAP mechanism is given by*

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

Similarly to 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; 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: similar to NMBR, it takes a monetary mechanism $M_{\textsf{money}}$ 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 $M_{\textsf{money}}$. In contrast with NMBR, payment of a given agent only depends on report of this agent, and is computed as expected payment for agent's report in $M_{\textsf{money}}$ when other agents report truthfully.

For this mechanism we are able to establish vanishing (α, β)-approximation guarantee resembling that of Theorem 2.3.1. Formally, we have the following theorem.

Theorem 2.4.2. *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,*

ALGORITHM 2: Repeated Endowed All-Pay (REAP) 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 (2.11); 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 $X(\hat{\theta}^t; M_{\mathsf{money}})$

6: **end for**

assuming all other agents report truthfully, we have

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

Moreover, assuming all agents play truthfully, we have

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

Comparing this result to that of Theorem 2.3.1, 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 simulation step of 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 REAP mechanism.

Finally, we make a note that running REAP mechanism might not be computationally tractable. The reason is that, in order to execute REAP we need to compute $c_s(\theta_s)$ according to (2.11), 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.

2.4.1 Proof of Theorem 2.4.2.

To prove Theorem 2.4.2, we adopt an approach analogous to the one we used in Section 2.3. Specifically, we establish three lemmas that are analogous to Lemmas 2.3.3, 2.3.4 and 2.3.5. 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

$$\widehat{u}_s(A_s, A_{-s}) = \mathbb{E}\left[\sum_{t=1}^{T} v_s(\theta_i^t, X(\hat{\theta}_s^t, \hat{\theta}_{-s}^t; M_{\mathsf{money}}))\right], \tag{2.12}$$

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

$$\widehat{c}_s(A_s) = \mathbb{E}\left[\sum_{t=1}^{T} c_s(\hat{\theta}_s^t)\right]. \tag{2.13}$$

Note that, in contrast to analysis of the NMBR mechanism as in Section 2.3, here the expected spending of agent s is independent of the strategies of the other agents, as an agent's payment only depends on her own reports. We say a strategy A_s is *feasible* for an agent s if it satisfies the expected budget constraint, $c_s(A_s) \leq B_s$.

Since $B_s > \bar{c}_s T$, 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 which 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 2.3.3. The proof is also analogous and is omitted for the sake of brevity.

Lemma 2.4.3 (Sensitivity). *In the auxiliary game, suppose the budget of agent s is $B_s = \bar{c}_s T + \Delta$ for some $\Delta > 0$. Then, for any feasible strategy A_s of agent s, we have*

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

The second lemma, analogous to Lemma 2.3.4, 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

$$\mathcal{E}_s = \left\{ \sum_{t=1}^{T} c_s(\theta_s^t) \geq B_s \right\}. \tag{2.14}$$

We have the following lemma, whose proof is similar to that of Lemma 2.3.4 and is deferred to Appendix 3.7.2.

Lemma 2.4.4 (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*

$$u_s(A_s, \mathsf{Tr}_{-s}; \mathrm{REAP}) \leq \widehat{u}_s(A_s^\emptyset, \mathsf{Tr}_{-s}) + 2\mathbb{P}[\mathcal{E}]Tv_{\max}.$$

Furthermore, we also have

$$\widehat{u}_s(\mathsf{Tr}_s, \mathsf{Tr}_{-s}) \leq u_s(\mathsf{Tr}_s, \mathsf{Tr}_{-s}; \mathrm{REAP}) + 2\mathbb{P}(\mathcal{E})Tv_{\max}.$$

Finally, we have the third lemma, analogous to Lemma 2.3.5, which bounds the probability of the budget depletion event $\mathcal{E}$. Once again, the proof is analogous and is omitted. Recall that $R_s = \bar{c}_s/c_s^{\max}$, where $c_s^{\max} = \max\{c_s(\theta_s) : \theta_s \in \Theta_s\}$, and $\bar{c}_s = \mathbb{E}[c_s(\theta_s)]$.

Lemma 2.4.5.

$$\mathbb{P}[\mathcal{E}] \leq n \exp\left(-\frac{T \min_s \delta_s^2 R_s}{3} \right). \tag{2.15}$$

Using these three lemmas, we obtain

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

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 3.7.2.

2.5 Stronger equilibrium guarantees for two agents

Theorem 2.3.1 and Theorem 2.4.2 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*, i.e., 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 non-monetary BIC mechanism that achieves vanishing welfare loss (i.e., guarantees $\alpha = 0, \beta = o(1)$ in the formalism of Definition 2.2.3)?

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

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

section; we then leverage this result to prove the existence of a BIC mechanism via a revelation principle argument.

2.5.1 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:

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

and

$$c_s(\theta_s) = L_s - \mathbb{E}_{F_{-s}} \left[\sum_{q \neq s} v_q(\theta_q, X^*(\theta)) \right], \tag{2.16}$$

where L_s is defined as

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

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 non-negative. We also restrict our attention to only those settings where the VCG mechanism satisfies the opt-out assumption (Assumption 2.2.2). 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 (2.16) yields the expected externality payments as in Example 2.4.1, and an agent can report $\hat{v}_s = 0$ to guarantee zero payment.

2.5.2 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 (cf. [72]). 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 2.5.1 (Smoothness of REAP-VCG)**.** *Given a two-player setting under the* RVCG *mechanism, for any arbitrary strategy profile* (A_1, A_2)*, we have*

$$\frac{u_1(\mathsf{Tr}, A_2; \mathrm{RVCG})}{T} \geq \frac{u_1(\mathsf{Tr}, \mathsf{Tr}; \mathrm{RVCG})}{T} - O\left(\sqrt{\frac{\log T}{T}}\right). \tag{2.17}$$

Note that (2.17) does not exactly fit the definition of (λ, μ) smoothness (cf. [72]), since 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 2.5.2 (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. Proof. Since (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 (2.17), 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 2.5.1. We do so by again by employing the auxiliary game we constructed earlier (cf. Section 2.3.2), 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 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 2.5.1, 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 (2.16), and is omitted for brevity.

Lemma 2.5.3. *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*

$$\mathbb{E}\Big[\widehat{U}_1^t(\theta_1^t, \tilde{\theta}_2^t)\big|\tilde{\theta}_2^t\Big] = L_2 - c_2(\tilde{\theta}_2^t). \tag{2.18}$$

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

Lemma 2.5.4 (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. Proof. Let $\left\{\tilde{\theta}_2^t : t \in [1, 2, \dots T]\right\}$ 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

$$\begin{aligned}
\widehat{u}_1(\mathsf{Tr}, A_2) &= \mathbb{E}\left[\sum_{t=1}^{T} \hat{U}_1(\theta_1^t, \tilde{\theta}_2^t)\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{T} \mathbb{E}\left[\hat{U}_1(\theta_1^t, \tilde{\theta}_2^t) \middle| \tilde{\theta}_2^t\right]\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{T} \left(L_2 - c_2(\tilde{\theta}_2^t)\right)\right] \qquad \text{(By Lemma 2.5.3)} \\
&= TL_2 - \widehat{c}_2(A_2).
\end{aligned}$$

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

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

where the last equality follows from observing $\bar{c}_2 T = TL_2 - \widehat{u}_1(\mathsf{Tr}, \mathsf{Tr})$, which can be derived from (2.18) 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 2.4.4 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 2.3.4 and is omitted for brevity.

Lemma 2.5.5. *For any strategy A_2 of agent 2 in the REAP-VCG mechanism, we have*

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

Finally, these lemmas allow us to prove Theorem 2.5.1 in a manner that is analogous to Theorem 2.3.1.

Proof. Proof of Theorem 2.5.1. We can write the entangled utility term (LHS of (2.17)) as

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

where the first bound follows from Lemma 2.4.4, the second from Lemma 2.5.4 and the last from Lemma 2.5.5. Substituting our choice of $\delta_1 = \sqrt{\frac{6(\log 2 + logT)}{R_1 T}}$, 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 2 agents. Namely, our result relies on the connection between agent 1 utility and agent's 2 payment as stated in Lemma 2.5.3, which holds even if agent 2 is misreporting. Such connection does not persist when there are more than one opponent, 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 2 agents is an open question.

2.5.3 Existence of near-efficient BIC mechanism

Finally, we return to the question of whether there exists a non-monetary BIC mechanism with vanishing inefficiency ($\alpha = 0$, $\beta = o(1)$, according to Definition

2.2.3). We now show how we can leverage the POA result from the previous section to prove the existence of such mechanism, under the assumption that the type spaces are finite.

Theorem 2.5.6. *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 is Section 3.7.2.

Corollary 2.5.7 (Existence of BIC mechanism)**.** *Under the conditions in Theorem 2.5.6, there exists a BIC mechanism for 2 agents with expected welfare loss of $o(T)$.*

Proof. Proof. In Theorem 2.5.6, we establish that a Bayes-Nash equilibrium exists for the REAP-VCG mechanism; moreover, Corollary 2.5.2 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. □

2.6 Discussion

We have presented a general blackbox reduction technique that allows us to convert any static monetary mechanism to a repeated non-monetary 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. **Non-stationary 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 non-stationary distributions: in NMBR, type sampling (Algorithm 1, Step 4) would be performed with respect to the appropriate distribution $F_s^t(\theta_s)$, and in REAP mechanism the pricing rule 2.11 would depend on time t. We update the definition of $\bar{c}_s$ to $\bar{c}_s = \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}[c_s^t(\theta^t; M_{\mathsf{money}})]$, where $c_s^t(\theta^t; M_{\mathsf{money}})$ denotes the payment of agent s in the static BIC mechanism M_{money}, when types are distributed as F^t. Similarly, we define $c_s^{\max} = \max_{t,\theta} c_s^t(\theta; M_{\mathsf{money}})$.

With these modifications, our proofs go through with no change; we illustrate this on the proof for NMBR mechanism. The definition of auxiliary problem remains untouched, and the sensitivity (Lemma 2.3.3) and the closeness (Lemma 2.3.4) results still follow as their proofs do not use stationarity. We can also extend the concentration result (Lemma 2.3.5), as Chernoff bound holds for bounded, non-identical, independent distributions.

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 (2.11)). 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 (for example, weather conditions in the case of allocation of parking spots or vacation days). (Note that non-stationary 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 the simulation step (Algorithm 1, Step 4) of the NMBR mechanism, for any agent s with no remaining budget, the mechanism samples a report $\hat{\theta}^t_s$ 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 2.3.3 and Lemma 2.3.4 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 2.3.5. This holds, for example, when ω^t is independent across time.

2.6.1 Proofs for NMBR mechanism

Proof. Proof of Theorem 2.3.6

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 $T\bar{c}_s\delta_s/2$, i.e.,

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

We now have the following concentration bound

Lemma 2.6.1. *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*

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

Proof. The first bound follows from standard Chernoff bounds (cf. Lemma 2.4.5):

$$\begin{aligned}
\mathbb{P}\Big[\tilde{B}_s \geq T\bar{c}_s(1+\delta+\delta/2)\Big] &\leq \mathbb{P}\bigg[T\bar{c}_s^m(1+\delta) \geq T\bar{c}_s(1+\delta)\frac{1+3\delta/2}{1+\delta}\bigg] \\
&\leq \mathbb{P}[T\bar{c}_s^m(1+\delta) \geq T\bar{c}_s(1+\delta)(1+\delta/3)] \\
&\leq \exp\bigg(-\frac{\delta_s^2 R_s m}{27}\bigg)
\end{aligned}$$

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

We denote $u_s(A_s;\mathcal{G}_s)$ to be 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 2.3.1 to obtain:

$$u_s(A_s;\mathcal{G}_s) - u_s(\mathsf{Tr};\mathcal{G}_s) \le \frac{3}{2}T\bar{c}_s\delta_s + 4\exp\left(-\frac{\delta_s^2 R_s T}{12}\right)Tv_{\max}. \tag{2.19}$$

The expected utility profile bound is obtained analogously to the proof of Theorem 2.3.1, by taking a union bound of events $\mathcal{G}_s$ and $\mathcal{E}_s$:

$$u_s(\mathsf{Tr};\mathcal{G}_s) \ge Tu_s(\mathsf{Tr};M_{\mathsf{money}}) - 2\exp\left(-\frac{\delta^2 R_s T}{27}\right)Tv_{\max} \tag{2.20}$$

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 2.6.1 that $\mathbb{P}[|c_s^m(\mathsf{Tr}) - \bar{c}_s| \le \bar{c}_s/2] \le 2\exp\left(-\frac{R_s m}{27}\right) = o(T^{1-k})$ for any k. This, together with the choice of δ_s and the bounds (2.19) and (2.20) yields the needed result. □

2.6.2 Proofs for REAP mechanism

Proof. Proof of Lemma 2.4.4. Define the utility received by the agent s in the auxiliary problem under when playing strategy A_s against truthful opponents

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

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 $\mathbb{E}[U_s(A_s)|\neg\mathcal{E}_{-s}] = \mathbb{E}\left[\widehat{U}_s(A_s^{\emptyset})|\neg\mathcal{E}_{-s}\right]$. This gives us

$$\begin{aligned} u_s(A_s) &= \mathbb{P}[\neg\mathcal{E}_{-s}]\mathbb{E}[U_s(A_s)|\neg\mathcal{E}_{-s}] + \mathbb{P}[\mathcal{E}_{-s}]\mathbb{E}[U_s(A_s)|\mathcal{E}_{-s}] \\ &\le \mathbb{P}[\neg\mathcal{E}_{-s}]\mathbb{E}\left[\widehat{U}_s(A_s^{\emptyset})|\neg\mathcal{E}_{-s}\right] + 2\mathbb{P}[\mathcal{E}_{-s}]Tv_{\max} \\ &\le \widehat{u}_s(A_s^{\emptyset}) + 2\mathbb{P}[\mathcal{E}_{-s}]Tv_{\max} \end{aligned} \tag{2.21}$$

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, i.e., $\mathbb{E}[U_s(\mathsf{Tr}_s)|\neg\mathcal{E}_{-s}\cap\neg\mathcal{E}_s] = \mathbb{E}[U_s(\mathsf{Tr}_s)|\neg\mathcal{E}_{-s}\cap\neg\mathcal{E}_s]$. The bound is then derived similarly to Eqn. (2.21). □

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

Proof. Proof of Theorem 2.4.2. Let A_s denote an arbitrary strategy in REAP. We have

$$u_s(A_s) - u_s(\mathsf{Tr}) = u_s(A_s) - \widehat{u}_s(A_s^{\emptyset}) + \widehat{u}_s(A_s) - \widehat{u}_s(\mathsf{Tr}) + \widehat{u}_s(\mathsf{Tr}) - u_s(\mathsf{Tr}) \quad (2.22)$$

We can now use bounds from Lemma 2.4.4 for $u_s(A_s) - \widehat{u}_s(A_s^{\emptyset})$ and $\widehat{u}_s(\mathsf{Tr}) - \widehat{u}_s(\mathsf{Tr})$, and Lemma 2.4.3 for $\widehat{u}_s(A_s) - \widehat{u}_s(\mathsf{Tr})$. Substituting them into (2.22) gives

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

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 2.4.5, we have $\mathbb{P}[\mathcal{E}] \leq n\exp\left(-\frac{\delta^2 RT}{3}\right)$. Now we have

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

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

Proof. Proof of Theorem 2.5.6. Note that in the REAP mechanism, the budget of each agent at any time t is completely determined by the previous reports. Since 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. Since this is a finite game, the existence of a BNE follows using standard arguments [36]. □

2.7 Applicability of REAP mechanism: examples.

The main difference between mechanisms described in Sections 2.3 and 2.4 is the dependence on excludability assumption (Assumption 2.2.1); namely, NMBR mechanism makes use of it and REAP mechanism does not. This can be viewed as a trade-off between generality and tractability, but how restricting is Assumption 2.2.1 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.

2.7.1 Assumptions on allocation setting

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

Example 2.7.1 (Combinatorial auction)**.** *There are n agents and m items, agents* *have combinatorial preferences for the allocated set of items that result from independently and randomly drawn types θ_s.*

Assumption 2.2.1 clearly holds for this setting. For NMBR mechanism, we can use VCG with Clarke Pivot as an input monetary mechanism M_{money}. If more is known about utility functions, other results from monetary mechanism design can be employed for tractibility purpuse, for example using the work of [46].

Example 2.7.2 (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 player has a greater utility under the mechanism when compared to initial allocation, and so (as we argue in Section 3.3) our IR criteria is satisfied. Thus REAP is an approximately IC, approximately efficient ex ante IR mechanism.

Example 2.7.3 (Voting)**.** *There are n agents and m options. On each round agents have value $v_{sk}^t \sim F_{sk}$ for the option k (drawn iid). A mechanism is to choose a single option on every round, and utilities of agents are their values for the option chosen.*

Again, Assumption 2.2.1 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 REAP mechanism to derive near-optimal welfare.

2.8 Interim guarantees for NMBR mechanism

Although throughout the 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 2.8.1. *Consider the* NMBR *mechanism and any constant* $\gamma \in (0,1)$, *denote* $\tilde{T} = (1-\gamma)T$. *Then, with probability* $p \geq 1 - 2e^{-\frac{\delta_s^2 R_s \gamma T}{3}}$, *the following hold simultaneously* for all agents s *at round* γT, *if all agents report truthfully up to this turn:*

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

$$\text{and} \qquad u_s(\mathsf{Tr}; M_{\textsf{money}}) - \frac{u_s^{\tilde{T}}(\mathsf{Tr}; \text{NMBR})}{\tilde{T}} \leq 2 n v_{\max} \tilde{T}^{-1},$$

where $u_s^{\tilde{T}}(A_s; \text{NMBR}) = \sum_{t=\gamma T}^{T} v_s(\hat{\theta}_s, , X(\hat{\theta}_s, \theta_{-s}; \text{NMBR}))$ *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. Proof. Note that, after γT periods, the remaining game is equivalent to the original one, with budgets $B_s^{\gamma T} = B_s^0 - \sum_{t=1}^{\gamma T} C_s(\theta; M_{\textsf{money}})$ and $\tilde{T}$ rounds remaining. Thus, we can reduce the statement of the theorem to that of the Theorem 2.3.1, 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 Chernoff bound to see that

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

Also, since 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:

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

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

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

We can now reduce the proof of the theorem to that of Theorem 2.3.1. To do this, we just need to notice that, in order to upper bound $\mathbb{P}[\mathcal{E}_s]$, we can use the lower bound for the value of $\tilde{\delta}_s$ (since 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 2.4.3 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 2.3.1 to obtain the needed result. □

CHAPTER 3

EFFICIENCY OF SIMPLE SCRIP MECHANISMS

3.1 Introduction

There are several employees at a firm who all make use of some scarce common resource, for example a cluster, conference rooms, or company cars. How should the company decide whom to allocate the resource on a given day? The principal wants the allocation to be efficient, that is, to allocate the resources to the people who would make most use of it. The use of money, however, is disallowed, as the salaried employees may refuse to pay out of pocket in order to be more productive.

One simple solution is for the principal to issue *scrip* (i.e., an artificial currency) among the employees; then the principal can repeatedly run a first-price auction using scrip for payments in order to decide who gets what. This intuitive approach is sometimes successfully employed in practice, one prominent example being the allocation system used by Feeding America [69].

In this paper we argue that the intuitive attractiveness of this approach can be justified game-theoretically. In particular, we show that every agent is able to secure a constant fraction of their ideal utility (a benchmark we define below), independently of the strategies adopted by other agents. We also argue that the budgets of agents can be interpreted as setting the probabilities of getting a given item.

3.1.1 Repeated Allocation without Money: Model and Objectives

We consider a model where there are T time periods, and the principal has m different items to allocate in each period. We assume that each agent has a random valuation function, drawn independently from some underlying distribution in each period; moreover, our results hold under the assumption that the value each agent derive from a bundle of goods in any given period is fractionally subadditive (XOS) in the set of items allocated.

In the absence of monetary transfers, defining an appropriate notion of a desirable allocation outcome can present a challenge. One useful constraint is *Pareto efficiency*, that is, requiring that no agent can be made better-off without hurting another agent. However, Pareto efficiency alone seems insufficient to fully capture the notion of a desirable outcome – for example, a dictatorial allocation that gives all resources to a single agent is Pareto efficient, but arguably is not an acceptable solution to the problem. To circumvent this, one needs to impose additional 'fairness' constraints on the outcome of the allocation.

To specify a particular desirable utility profile, we assume that each agent is guaranteed a certain 'status quo' allocation. In particular, we assume that, on average, agent i is guaranteed to get item j with probability α_{ij} . Every agent has an ideal allocation – an allocation they like most subject to having satisfy the status quo allocation for the remaining agents. Throughout the paper we strive to allocate in a way that ensures that expected utilities of agents are approximately their ideal expected utilities.

3.1.2 Incentives

Given the above desired allocation rule, the challenge in implementing it is that the principal is unaware of the preferences of participating agents. The aim is to construct a mechanism in which agents are incentivized to behave in a way that leads to a desired outcome.

Throughout the paper, we employ the notion of *maxmin value* for reasoning about incentive guarantees. In particular we introduce the property of β-utopia, that states that the maxmin value of every agent is at least β-approximation of her ideal utility. In other words, this property implies that, no matter what other agents choose to do in the mechanism, any agent can secure an expected utility that is close to her ideal utility.

Such a guarantee implies (and indeed, is strictly stronger than) a Price of Anarchy guarantee for our setting, i.e., establish equivalent properties at all Nash equilibria. However, in a β-utopian mechanism, even in situations where the population as a whole fails to converge to an equilibrium, individual agents can still secure a β fraction of their ideal utility by unilaterally adopting a certain bidding strategy.

3.1.3 Outline of Results

We now give a brief outline of our results.

We start with considering the setting of allocating a single item in Section 3.4 – doing so allows us to clearly communicate the analytic approach we adopt in proving a more technical result in Section 3.6, and the theorems we prove about multiple item allocation in Section 3.5 are proved by reducing them to the established single item results.

For allocating of a single item, we consider a simple mechanism (Algorithm 3), in which agents are endowed with budgets of artificial currency and then repeatedly play a first price auction, using the provided currency for payments.

Our main result in this setting is establishing a $\left(\frac{1}{2} - o(1)\right)$-Utopia for this mechanism (Theorem 3.4.1), where $o(1)$ is vanishing for large T. To do so we consider a specific class of strategies we call Bernoulli strategies (see (3.4)), and for this class we prove a core result, the Bang for Buck Lemma (Lemma 3.4.2), that guarantees a linear bound between agent's expected utility and expected spending. Having established this result, we can then reduce proving 1/2-Utopia to proving that every agent can spend their entire budget while playing the appropriate Bernoulli strategy.

We then switch to the setting of allocating multiple items, and consider a generalization of the simple mechanism we describe above. In the Chromatic Auction (Algorithm 4), agents are endowed with m budgets of currencies (one for each item), and simultaneous first-price auctions are executed in every period for each individual items, using corresponding currencies for payments. For arbitrary endowments α_{ij} we demonstrate a $\left(\frac{1}{2} - o(1)\right)$-Utopia property for this mechanism (Theorem 3.5.2)using a reduction to results from Section 3.4.

Because employing a separate currency for each item type may be burdensome in real life problems, in Section 3.6 we consider a simpler mechanism where a single currency is employed for all payments. We argue that the guarantees we provide for Chromatic Auction are impossible here for arbitrary endowments α_{ij}, but establish a $\left(\frac{1}{2} - o(1)\right)$-Utopia result for when the endowments are equal across items $\alpha_{ij} = \alpha_i$. In doing so we adopt an approach similar to the one employed in Section 3.4, first proving a Generalized Bang-per-Buck Lemma (3.6.2) and then reducing the proof of 1/2-Utopia to proving that any agent can spend their entire budget while playing a specific Bernoulli strategy.

Finally, in Appendix 3.7.1 we consider the question of whether individual rationality constraints are satisfied for the mechanisms we considered. In particular we demonstrate how, for particular lottery allocation rules, agents in the mechanisms we have presented can ensure the expected utility that is at least their expected utility under these lottery allocation rules.

3.1.4 Discussion

Somewhat surprisingly, the informational assumptions needed to demonstrate our results are very frugal: the agents need not know the preferences of other agents, nor the equilibrium strategy profile. The principal need not know anything about the distributions of the participants. Moreover, the mechanism is robust to individual agents deviating from rational behavior. Even if the system as a whole fails to converge to a Nash equilibrium, individual agents can still obtain good expected utility by bidding in a rather intuitive way.

A consequence of this generality is that the scrip mechanisms we consider, as well as the results we prove, are in line with *Wilson's doctrine*, that suggests the use of mechanisms that do not depend on the primitives of the model (such as type distributions). Thus, we hope our work provides additional justification for the success of artificial currency mechanisms in practice [70]. Moreover, our ideas also provide guidance for improving the design of such systems; in particular, we draw attention to the use of multiple currencies for implementing different ideal utility profiles (cf. 3.5), and the benefits of employing the first-price over second-price auction in scrip economies (cf. discussion before Lemma 3.4.2).

3.2 Related Literature

Our work is broadly situated in the area of mechanism design without money. Much of the literature in this area deals with single-period settings – in particular, there is a large body of work devoted to finding one-shot allocations that satisfy various combinations fairness and efficiency constraints [63, 71, 22]. These works typically disregard strategic considerations and incentives, as achieving fairness, efficiency and incentive compatibility is known to be impossible in most settings (Budish [19] provides a good overview of the related trade-offs). Another approach is to relax some of the constraints to be met approximately in order to provide simultaneous fairness, efficiency and incentive guarantees. Two notable successes in this line are by Budish [19], who constructs an approximately fair and approximately incentive compatible mechanism for one-shot combinatorial allocations, and Brânzei et al. [17], who present a mechanism that approximates Nash Welfare (maximizing which can be viewed as a combination of fairness and efficiency constraints) under any Nash equilibrium.

An alternate approach to achieving a desired allocation without sacrificing incentives is to couple together multiple allocation problems. This approach was pioneered in the work of Jackson and Sonnenschein [50], who consider the problem of choosing many simultaneous allocations, and provide a mechanism that guarantees near-optimal welfare at any equilibrium in the mechanism. They achieve this by essentially endowing agents with separate budgets for reporting each possible type across instances. This result is of theoretical importance, but is not practically useful, since the mechanism makes use of a separate budget for every possible *type*, which may be prohibitively large for multidimensional type-spaces (e.g. combinatorial auctions). We note that our Chromatic auction (Algorithm 4) has a similar flavor, but crucially, requires separate currency for each *item*.

The idea of linking decisions can also be used for allocating in multiple periods over time. This idea is explored for single-item allocation and symmetric agents in a line of work starting with [44], and refined by [6] (this idea has also been extended to general allocation settings [39]; more on this below). These works assume an endogenously provided way of combining preferences of agents into a welfare function and strive to maximize welfare subject to incentive compatibility. There are several major differences that set us apart from this work: 1) in line with literature in economics, we avoid agglomerating utilities of agents into a welfare function in the absence of transfers, as it can be challenging to rigorously justify doing so; 2) the allocation setting we consider here is considerably more general; 3) our guarantees hold for the maxmin value, rather than using the Nash equilibrium as a solution concept; and 4) we concentrate on the performance of simple, practically implementable mechanisms.

Another closely related line of work explores the properties of artificial currency (i.e., *scrip*) economies [55, 45]. These are widely used in practice, with recent successes in course-allocation [20] and the distribution of food-bank donations [70]. With respect to the latter, Prendergast [69] considers a model for allocating food to foodbanks using auctions with artificial currency; however he eschews reasoning about the incentives of the mechanism. Gorokh et al. [39] consider the allocation setting, and construct a blackbox reduction from single-shot monetary mechanisms into their artificial currency-based counterparts. They demonstrate that the resulting mechanism is approximately efficient and approximately incentive compatible. The differences with our work are similar as outlined above, namely, the use of welfare function and Nash equilibrium as the solution concept.

Finally, we note that our maxmin bounds directly imply equivalent Price of Anarchy guarantees. There is an extensive literature on price-of-anarchy guarantees for various settings [72], including for simultaneous auctions for items with XOS valuations [73]. Most of this literature is for static settings, or dynamic settings with no-regret strategies [60]. One exception is a result of Gorokh et al. [39], who demonstrate that for 2-agents, the artificial currency all-pay auction has vanishing price-of-anarchy; they however mention challenges in extending this result. We note that one consequence of our results is the *first general price-of-anarchy guarantees for repeated non-monetary allocation settings*, with any number of agents and XOS valuations.

3.3 Model

We consider a repeated allocation setting with n agents. At each time[1] $t \in [T]$, a principal has a fresh set of m indivisible items $[m] = \{1, 2, \dots, m\}$ to allocate in that round, and must choose an allocation X^t of items to agents, where X^t_{ij} is the indicator that agent i gets item j (we sometimes use X^t_i as a vector of allocations for agent i). We denote $\mathcal{X} = \{X | X_{ij} \in \{0, 1\}, \sum_{i \in [n]} X_{ij} \leq 1 \, \forall \, j \in [m]\}$ to be the

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

set of feasible allocations in each round.

3.3.1 Preferences of agents.

Agents have preferences over lotteries of allocations, which we model with the use of cardinal utilities (we assume that preferences of agents satisfy von Neumann-Morgenstern axioms). We furthermore assume that these preferences depend on an agent's *type*, where we denote an agent i's type at time t by θ_i^t. We assume that the types $\theta_i^t \in \Theta_i$ are drawn from a distribution F_i, and are i.i.d across periods, and independent across agents. We let $\theta^t = (\theta_i^t : i \in [n])$ denote the type profile at time t, and let $\theta_{-i}^t = (\theta_q^t : q \in [n], q \neq i)$. Let $\Theta = \times_{i\in[n]}\Theta_i$ denote the set of type profiles.

Let $v_i(\theta_i^t, X_i^t)$ denote the utility that the agent i receives from allocation X_i^t when having type θ_i^t. Throughout the paper, we focus on the setting where agents have *fractionally subadditive (XOS) values* [57, 33] for the items, i.e. the utility an agent derives from a set of items is a maximum over k linear valuations. Formally, there exists matrices $D(\theta_i) \in \mathbb{R}^{m\times k}$ for each $\theta_i \in \Theta_i$ such that $v_i(\theta_i, X) = \max_{k\in K} \sum_{j\in[m]} D_{jk}(\theta_i) \cdot X_{ij}$. Note that this covers the case of single item allocation, with $D(\theta_i) = \theta_i \in \mathbb{R}$, and $\mathcal{X} = \{0, 1\}^n$.

3.3.2 Mechanisms, Strategies and Maxmin Value

In the mechanisms we consider, each agent i submits a vector of bids b_i^t at each time $t \in [T]$, subsequent to which an allocation $X^t \in \mathcal{X}$ is chosen. At time t, the history h_i^t of agent i is defined as the set of all past types of agent i, all past bids, and past allocations:

$$h_i^t = \{(\theta_i^{t'}, b^{t'}, X^{t'}), t' < t\}$$

Similarly, the *common history* H^t at time t is defined as the set of all past bids and allocations: $H^t = \{(b^{t'}, X^{t'}), t' < t\}$. Finally, a mechanism M specifies for each time t, common history H^t and bids b^t, an allocation[2] $X^t(M) \triangleq X^t(M, H^t, b^t) \in \mathcal{X}$.

For each agent i, a strategy $\mathcal{S}_i$ is a mapping that specifies, for each history h_i^t and current type θ_i^t, a distribution over bids b_i^t in the current period; a strategy profile $\mathcal{S} = \{\mathcal{S}_i\}_{i \in [n]}$ is the collection of individual agent strategies. We define $u_i(M, \mathcal{S}) = \mathbb{E}\left[\sum_t v_i(\theta_i^t, X_i^t(M))\right]$ as the expected utility of agent i in mechanism M under strategy profile $\mathcal{S}$, where the expectation is taken over types θ^t as well as over the randomization in the agents' strategies. We abuse the notation and write $u_i(\mathcal{S})$ when the mechanism M is clear from context. Now, for reasoning about the incentives and strategic behavior of agents, we use the notion of the *maxmin value*, defined as follows:

Definition 3.3.1 (Maxmin value). *For any given mechanism M, the maxmin value of any agent i is the highest expected utility she is guaranteed to get under M, without knowing the strategies chosen by other agents. Formally:*

$$\bar{V}_i(M) = \max_{S_i} \min_{\mathcal{S}_{-i}} u_i(M, \mathcal{S})$$

The standard definition of the maxmin value requires agents to select independent strategies; our results however hold under weaker conditions, wherein we permit other agents to collude in their strategies. Indeed, it is convenient to view all agents other than i as acting as a single adversary, playing some strategy

[2]Though our notation implies a deterministic mechanism, our results carry over for randomized mechanisms, where the allocation is chosen at random. To keep the notation simple, we avoid defining it formally.

$\mathcal{S}_a = \{S_{i'}\}_{i' \neq i}$. Since such an adversary can adopt strategies that are infeasible for uncoordinated agents, our bounds are strictly stronger then the usual notion of the maxmin value.

3.3.3 Solution concept

Since there are no monetary transfers in our setting, in general, there can be a large number of Pareto efficient outcomes (E.g., give all items to player i). To define a desirable outcome, we adopt the following concepts from bargaining theory.

We define a status quo outcome as one where agent i gets items according to some allocation rule $X^{sq}_{ij}(\theta_i)$. Throughout the paper we consider X^{sq} to be various weighted lottery allocation rules, where agent i is endowed with α_{ij} probability of getting item j. We refer to the probabilities $\{\alpha_{ij}\}_j$ as agent i's endowments.

We assume the endowments are exogenously specified, perhaps arising from some underlying rights structure among the agents; for example, if the principal considers agents to be equal in their rights for the resource, a logical choice is to set them equal.

Given any endowments α_{ij}, we define the notion of the *ideal utility*, which we use as a benchmark throughout the paper.

Definition 3.3.2 (**Ideal utilities**). *Given endowments $\{\alpha_{ij}\}$, $\sum_i \alpha_{ij} = 1$ we define the (per-period expected) ideal utility of agent i as the maximum utility she can get while ensuring every other agent i' get at least her endowed share $\alpha_{i'j}$ of* *each item j:*

$$v_i^*(\alpha) = \max_{X_i(\theta_i)} \mathbb{E}\left[v_i\left(\theta_i, X_i(\theta_i)\right)\right]$$
$$s.t.\ \mathbb{E}\left[X_{ij}(\theta_i)\right] \leq \alpha_{ij}. \tag{3.1}$$

We note that this is a rather strong benchmark, as it allows direct access to types of agents, which are private. We also briefly note that the above definition can be related to the Kalai-Smorodinsky bargaining solution [51]; in particular, if we consider the zero allocation as the disagreement point, and restrict the space of allocations to ones where each agent i obtains each item j at least α_{ij} fraction of the time, $v^*(\alpha)$ is the ideal point in KS solution.

The following example illustrates the notion of ideal utilities for single-item settings.

Example 3.3.3 (Ideal utilities in symmetric single-item setting)**.** *Consider a setting with a single item per round, and n agent, with i.i.d. valuations in reach round distributed as:*

$$v_i = \begin{cases} 1 & w.p.\ \frac{1}{n} \\ 0 & otherwise \end{cases}$$

Suppose now the designer chooses equal endowments for the agents, i.e., $\alpha_i = \frac{1}{n}$. Then each agent has an average value of $1/n$ in each round, and hence the utility each agent gets from the uniform lottery is $v_i^{lottery} = \frac{1}{n^2}$. In contrast, the ideal utility of each agent is $v_i^*(\alpha) = \frac{1}{n}$, which corresponds to utility the agent obtains from their favorite $1/n$ share of the items. Note though that this utility profile $v_i^*(\alpha)$ is infeasible – in particular, under the symmetric Pareto-optimal allocation (that maximizes the welfare function $W(v) = \sum_i X_i v_i$), each agent obtains an expected utility of $v_i^{PO} = \frac{1}{n}\left(1 - \left(1 - \frac{1}{n}\right)^n\right) \approx \left(1 - \frac{1}{e}\right) v_i^*(\alpha)$.

The above example shows that, in general, there may exist no allocation rule $X(\cdot)$ that guarantees v_i^* for all agents simultaneously. Another subtle point is that the expected utility profile v^* is only Pareto optimal for some choices of α_{ij} - we address this point in detail in Section 3.5.

The main guarantee we desire of our mechanisms in this paper is of the following form:

Definition 3.3.4 (**β-Utopia:**). *We say a mechanism M is β-utopic if, for any given endowment vector α, the maxmin value for each agent i in the mechanism $\bar{V}_i(M)$ is a $\beta-$approximation of her ideal utility, i.e.,*

$$\bar{V}_i(M) \geq \beta T v_i^*(\alpha) \qquad \forall\, i \tag{3.2}$$

This property states that any participant can get a β-fraction of their ideal outcome utility, no matter how the rest of the participants behave. Note that such a guarantee also implies an equivalent guarantee for the utilities of agents at any Nash equilibrium (i.e., a β-*Price of Anarchy* (PoA) guarantee [72, 73]). However, it is more robust than typical PoA guarantees: even if for some reason other agents are not playing according to an equilibrium, any agent can obtain a β fraction of the optimal utility by bidding in a certain way. Moreover, as we will see later, in order to do so the agent does not need to know the strategies of other agents, nor their type distributions.

Another property one can ask for a mechanism to satisfy is individual rationality – given that there is some status quo allocation rule (e.g. a weighted lottery), are all the agents incentivized to switch to the proposed mechanism? We can formalize this requirement by asking that $\bar{V}_i \geq T v_i^{sq}$. Depending on the type distributions, guaranteeing β-Utopia might already imply this bound (intuitively, this is the case when agent values have high variability); when this is not the case, we demonstrate how individual rationality constraint can be explicitly satisfied (see Appendix 3.7.1).

3.4 Allocation of a single item

For the allocation of a single item in each round, we consider a simple mechanism (Algorithm 3) where we provide each agent i with a budget B_i of artificial credits, and then run a first price auction on every round. Since the credits are scale invariant, we assume W.L.O.G. that $\sum_i B_i = T$.

ALGORITHM 3: Artificial Currency First-Price Auction

Require: Budgets B_i s.t. $\sum_i B_i = T$

Let $B_i^1 = B_i$

for all $t = 1$ to T **do**

  Agents submit bids b_i^t

  Budgets are enforced $b_i^t \leftarrow \min(B_i^t, b_i^t)$

  Item is allocated to a highest bidder $i^\star \in argmax\{b_i^t\}$, breaking ties arbitrarily

  Budgets are updated $B_i^{t+1} = B_i^t - b_i^t \mathbb{I}[i = i^\star]$.

end for

Recall that in case of the single item allocation, we have $v_i(\theta_i, X) = \theta_i X_i$. Before stating the main result of this section, we first note in the single item setting, the per-round ideal expected utility for agent i with endowment α_i is given by the following expressions:

$$v_i^*(\alpha) = \max_X \; \mathbb{E}\left[\theta_i X_i(\theta_i)\right] \quad \text{s.t., } \mathbb{E}\left[X_i(\theta_i)\right] \leq \alpha_i \tag{3.3}$$

The main result of this section is as follows.

Theorem 3.4.1. *The Artificial Currency First-Price Auction M_{FP} with $B_i = \alpha_i T$ satisfies $\frac{1}{2}$-Utopia:*

$$\bar{V}_i(M_{FP}) \geq \frac{1}{2} T v_i^*(\alpha) \left[1 - O\left(\frac{\log T}{\sqrt{T}}\right)\right].$$

We prove this result using a 2-stage approach, which we also employ in later sections. First we construct a strategy that guarantees a certain *bang-per-buck*, i.e., the expected utility under this strategy is at least a guaranteed fraction of the expected spending. We then prove a *spending lemma*, showing that the constructed strategy, with high probability, results in the agent entirely depleting her budget. Combining these two results together yields a lower bound for the maxmin value $\bar{V}_i(M_{FP})$.

For the rest of the paper, for any chosen agent i with budget B_i, it is convenient to view the other agents as a single adversary with a budget of $B_a = \sum_{j \neq i} B_j$. Now given some *fictitious allocation* rule $\hat{X}_i : \Theta_i \to \{0, 1\}$, we define the expected value of this allocation rule $\nu_i(\hat{X}_i) = \mathbb{E}[v_i(\theta_i, \hat{X}_i(\theta_i))]$ (where the expectation is over type θ_i, and also any potential randomness in X_i).

We also define the following bidding strategy $\mathcal{S}(\hat{X}_i, r)$, parametrized by an allocation rule $\hat{X}_i(\theta_i)$ and a scalar $r \geq 0$, which we refer to as a *Bernoulli strategy*

:

$$b_i(\theta_i) = \begin{cases} r & if\ \hat{X}_i(\theta_i) = 1 \\ 0 & otherwise. \end{cases} \tag{3.4}$$

We briefly describe the intuition for picking this set of strategies. When agent i submits a bid b_i^t with some probability α_i her expected spending in the abscence of adversary's bids is only $\alpha_i b_i^t$. On the other hand, an adversary who wants to prevent the agent from obtaining any utility needs to bid $b_a = b_i$ and spend $C_a^t = b_i^t$. Thus an adversary needs to spend $1/\alpha_i$ more than agent i would in order to block her. The point of the Bernoulli strategy is that it allows agent i to compensate the asymmetry between her budget $B_i = \alpha_i T$ and the (potentially much larger) budget of the adversary $B_a = (1 - \alpha_i)T$. The choice of the first-price auction is critical for this technique, and one can easily verify that this bidding approach breaks when considering the second-price auction; in particular, it is hard for an agent to create the described spending asymmetry between herself and the adversary.

The powerhouse behind our 1/2-utopia guarantee is the following lemma, which relates the expected utility of an agent i with her expected spending, for *a* Bernoulli strategy parametrized by any fictitious allocation rule $\hat{X}_i(\theta_i)$.

Lemma 3.4.2 (Bang-for-Buck Lemma)**.** *In the First Price Auction mechanism, if agent i plays $\mathcal{S}(\hat{X}_i, r)$, then the expected total utility of agent i under any strategy profile $\mathcal{S}_{-i}$ of other agents, satisfies* [3]*:*

$$u_i(\mathcal{S}(\hat{X}_i, r), \mathcal{S}_{-i}) \geq \frac{\nu_i(\hat{X}_i)}{r\mathbb{E}[\hat{X}_i(\theta_i)]} c_i(\mathcal{S}(\hat{X}_i, r), \mathcal{S}_{-i}).$$

Proof. Consider any round t, and let C_i^t be the spending of agent i in this round. Since our goal is to lower-bound utility of the agent, for the remainder of the proof

[3] We note that when $\hat{X}_i(\theta_i) = 0$, then both sides of the inequality are 0.

we assume that the agent never gets the item when bidding zero (as getting the item in those instances can only improve agent's utility).

Suppose on this round agent i wins the item $X_i^t = 1$ and pays their bid $C_i^t = r$. This means the agent bid $b_i^t = r$ (which in turn means that $\hat{X}_i(\theta_i^t) = 1$) and the adversary bid $b_a^t \leq r$. Also note that adversary's bid b_a^t is independent of agent i's type θ_i^t (because the agent's type is private). With all this in mind, we have

$$\mathbb{E}[U_i^t | C_i^t = r] = \mathbb{E}[U_i^t | X_i^t = 1] = \mathbb{E}[\theta_i^t | \hat{X}_i(\theta_i) = 1 \; and \; b_a^t \leq r] =$$
$$= \mathbb{E}[\theta_i^t | \hat{X}_i(\theta_i) = 1] = \frac{\mathbb{E}[\theta_i \hat{X}_i(\theta_i)]}{\mathbb{E}[\hat{X}_i(\theta_i)]} = \frac{\nu_i(\hat{X}_i)}{\mathbb{E}[\hat{X}_i(\theta_i)]}.$$

Here the first equality uses the discussed above assumption that agent only wins when they bid above zero.

On the other hand, when the agent does not win the item $X_i^t = 0$ the payment is zero $C_i^t = 0$ and so is the obtained utility $U_i^t = 0$. Thus $\mathbb{E}[U_i^t | C_i^t = 0] = 0$.

Since 0 and r are the only possible payments for an agent playing a Bernoulli strategy, we thus have

$$\mathbb{E}[U_i^t | C_i^t] \geq \frac{\nu_i(\hat{X}_i)}{r\mathbb{E}[\hat{X}_i(\theta_i)]} C_i^t.$$

Summing up over all rounds and taking expectation over potential randomness from tie-breaking and the strategies of other agents, we obtain the statement of the lemma. □

Next, consider the fictitious allocation rule $\hat{X}_i(\theta_i) = X_i^*(\theta_i)$ to be the one that achieves the ideal value as defined in (3.3). Thus $\mathbb{E}[X_i^*(\theta_i)] = \alpha_i$ and $\nu_i(X_i^*) = v_i^*(\alpha)$ (as per notation outlined before Lemma 3.4.2). We now consider the Bernoulli strategy $\mathcal{S}(X_i^*, 2)$, that is, bidding $b_i^t = 2$ whenever $X_i^*(\theta_i) = 1$, and 0 otherwise. Note that the expected bid is then $2\alpha_i$, and thus it takes $T/2$ periods (if not blocked) on average for the agent i to run out of budget $B_i = \alpha_i T$. Our next lemma shows that the expected spending of agent i when playing this strategy is close to the total budget.

Lemma 3.4.3 (Agent spends most of her budget)**.** *If agent i employs $\mathcal{S}(X_i^*(\cdot), 2)$ then*

$$c_i(\mathcal{S}(X_i^*(\cdot), 2), \mathcal{S}_a) > \alpha_i T \left(1 - \sqrt{\frac{2 \log T}{\alpha_i T}} - O(1/T)\right)$$

for any strategy $\mathcal{S}_a$ adopted by the adversary.

The proof of the preceding lemma uses the following auxiliary lemma, which characterizes the optimal strategy of the adversary when the agent follows the Bernoulli strategy $\mathcal{S}(X_i^*, 2)$.

Lemma 3.4.4. *Suppose agent i employs strategy $\mathcal{S}(X_i^*, 2)$, and the adversary strives to minimize the expected spending $c_i(\mathcal{S}(X_i^*, 2), \mathcal{S}_{-i})$ of agent i. Then it is optimal for adversary to always bid $b_a^t = 2$ (until he runs out of budget).*

Proof Outline. A formal proof of this involves a somewhat technical dynamic programming argument; we defer this to the Appendix 3.7.3, and instead provide an informal argument here.

For simplicity, assume that the ties are always broken in favor of the adversary. Then the strategy of adversary reduces to deciding on which rounds to "block" the agent i by bidding $b_a^t = 2$ (bidding anything other than 2 or 0 is dominated). Since the adversary has a budget of $B_a = \sum_{j \neq i} B_i \leq T$, she can only block agent i on $T/2$ rounds.

Since the decision on agent i whether to bid $b_i^t = 2$ is independent on previous history, the adversary gains no advantage from deciding on which periods to block in a dynamic fashion, and thus any static policy that chooses $T/2$ rounds and blocks them is optimal (including the policy of blocking on the first $T/2$ rounds, which this lemma states to be optimal). □ □

We can now combine Lemmas 3.4.2 and 3.4.3 to obtain the statement of Theorem 3.4.1. According to Lemma 3.4.2 agent i obtains $2v_i^*(\alpha)/\alpha_i$, and Lemma 3.4.3 tells us that expected spending of agent i is at least $\alpha_i T\left(1 - O\left(\sqrt{\log T/T}\right)\right)$. Taken together, these bounds yield the theorem's statement.

Before proceeding to the next section, we briefly discuss the intuition between picking the parameter $r = 2$ for the Bernoulli strategy $\mathcal{S}(\hat{X}, r)$ we employed in our proof (and which led to the value of $\beta = \frac{1}{2}(1 - o(1))$). The trade-off in picking r is that, lower values of r allow a higher conversion ratio of spending into utility as given by the Bang-per-Buck lemma, but also decreases the spending of the agent (as it becomes possible for the adversary to block more periods). More precisely, the adversary can block up to T/r periods, and the utility-spending conversion ratio is given by v_i^*/r. Because the spending of agent i is at most her budget $B_i = \alpha_i T$, the total expected utility of the agent becomes

$$u_i(\mathcal{S}(X^*, r), \mathcal{S}_{-i}) \geq \frac{v_i^*}{2} \min(\alpha_i T, (T - T/r)\alpha_i r)$$

Now it is easy to see that the value $r = 2$ achieves optimal value of this bound.

3.5 Chromatic auction

We now switch to a more general setting, where there are m items and the agents have fractionally subadditive valuations. A straightforward way to generalize the mechanism we employed for the single item setting (Algorithm 3) is to run m First Price auctions on every round simultaneously, using a separate currency (of a corresponding "color") for each item j. The formal description of the mechanism is given by Algorithm 4.

ALGORITHM 4: Chromatic auction

Require: Budgets B_{ij}, $\sum_i B_{ij} = T\ \forall j$

1: Let $B^1_{ij} = B_{ij}$

2: **for all** $t = 1$ to T **do**

3: Every agent i submits bids $b^t_{i,j},\ b^t_{ij} \leq B^t_{ij}$

4: Each item is allocated to highest bidder, and the winner is charged their bid in the corresponding currency: $B^{t+1}_{ij} = B^t_{ij} - b^t_{ij}\mathbf{1}[b_{ij} \geq b_{-i,j}]$

5: **end for**

Recall that we defined the ideal per-period expected utility $v^*_i(\alpha)$ as

$$v^*_i(\alpha) = \max_{X_i(\theta_i)} \mathbb{E}\left[v_i\left(\theta_i, X_i(\theta_i)\right)\right]$$

$$s.t.\ \mathbb{E}\left[X^t_{ij}(\theta_i)\right] \leq \alpha_{ij}. \tag{3.5}$$

Before proceeding to our main result, we briefly discuss how the choice of parameters α affects the ideal utility profile $v^*(\alpha)$. Of particular interest is whether $v^*(\alpha)$ Pareto-dominates any feasible expected utility profile. It is clearly so in the case of a single item allocation, however it is no longer true for the multiple item case, as we demonstrate with the example below.

Example 3.5.1 (Ideal utility profile may not be Pareto optimal)**.** *There are n items and n agents with constant unit-demand preferences: $v_i(X_i) = \mathbf{1}[i \in X_i]$. Now, consider a potential choice of endowments $\hat{\alpha}_{ij} = 1/n$. Because agent i can only get $1/n$ fraction of the item they like, we have $v_i^*(\alpha) = 1/n$.*

However, there is a feasible allocation rule, $X_{ij} = \mathbb{I}[i = j]$, that gives utility of 1 to every agent. Thus the utility profile under this allocation rule Pareto dominates the ideal utility profile $v_i^*(\alpha)$

Thus, if a designer strives to achieve approximate Pareto efficiency, it is important to pick α_{ij} correctly. It is in fact the case that, for any Pareto Efficient allocation rule $X(\theta_i)$, there exist coefficients α_{ij} such that the ideal utilities weakly dominate the expected utilities under this allocation rule $v_i^*(\alpha) \geq \mathbb{E}v_i(\theta_i, X(\theta_i))$.

3.5.1 Maxmin Characterization of Chromatic Auctions

Our main result in this section is the following theorem, which establishes that the Chromatic Auction satisfies $\frac{1}{2}(1 - o(1))$-Utopia.

Theorem 3.5.2. *Chromatic Auction mechanism M_{CA} satisfies $\frac{1}{2}\left(1 - O\left(\sqrt{\frac{\log T}{T}}\right)\right)$-utopia with $B_{ij} = \alpha_{ij}$:*

$$\bar{V}_i(M_{CA}) \geq \frac{1}{2} T v_i^*(\alpha) \left[1 - O\left(\sqrt{\frac{\log T}{T}}\right)\right]$$

Before proving Theorem 3.5.2, we first recall a useful property of XOS valuations that we employ throughout the paper. From here onward, we sometimes slightly abuse the notation by writing $v_i(\theta_i, S) = v_i(\theta_i, X_i)$ where $S = \{j | X_{ij} = 1\}$.

Proposition 3.5.3. *For any given agent i with type θ_i, and any set of items $S \subset [m]$, there exists a price-vector $p_i(\theta_i, S)$ satisfying $v_i(\theta_i, S) = \sum_{j \in S} p_{ij}(\theta_i, S)$ and $v_i(\theta_i, S') \geq \sum_{j \in S'} p_{ij}(\theta_i, S)$ for any set $S' \subset [m]$.*

The ideal per-period expected utility $v_i^*(\alpha)$ and allocation rule $X_{ij}^*(\theta_i)$ are given by (3.5), and using the above lemma and denoting $S_i^*(\theta_i) = \{j | X_{ij}^*(\theta) = 1\}$ we can rewrite $v_i^*(\alpha)$ as

$$\begin{aligned} v_i^*(\alpha) = \mathbb{E} v_i(\theta_i, X_i^*(\theta_i)) &= \mathbb{E} \sum_{j=1}^{m} p_j(\theta_i, S_i^*) X_{ij}^*(\theta_i) \\ &= \sum_{j=1}^{m} \mathbb{E}[p_j(\theta_i, S_i^*) | X_{ij}^*(\theta_i) = 1] \mathbb{P}[X_{ij}^*(\theta_i) = 1] \\ &= \sum_{j=1}^{m} \alpha_{ij} \mathbb{E}[p_j(\theta_i, S_i^*) | X_{ij}^*(\theta_i) = 1]. \end{aligned} \tag{3.6}$$

Proof of Theorem 3.5.2. We prove the theorem using a reduction to the single item setting (Theorem 3.4.1).

Suppose on some rounds the agent wins some set of items S. Using Lemma 4.4.4 we get

$$U_i^t \geq \sum_{j \in S} p_j(\theta_i, S_i^*(\theta_i)) X_{ij}^*(\theta_i).$$

Thus the utility of the agent i can be lower-bounded with the utility of an agent with additive valuation function

$$\tilde{v}_i(\theta_i, S) = \sum_{j \in S} p_j(\theta_i, S_i^*(\theta_i)) X_{ij}^*(\theta_i) = \sum_{j} \tilde{v}_{ij}(\theta_i),$$

where $\tilde{v}_{ij}(\theta_i) = p_j(\theta_i, S_i^*(\theta_i)) X_{ij}^*(\theta_i)$.

Now, we concentrate on an item j and valuation $\tilde{v}_j(\theta_i)$ and apply Theorem 3.4.1 directly. Recall that the ideal utility $v_{ij}^*(\alpha_j)$ for the single item (the item j)

is given by:

$$\tilde{v}_{ij}^*(\alpha_j) = \max_X \ \mathbb{E}\left[\tilde{v}_j(\theta_i)X_{ij}(\theta_i)\right] \quad \text{s.t., } \mathbb{E}\left[X_{ij}(\theta_i)\right] = \alpha_{ij}$$

Since $\tilde{v}_j(\theta_i) = 0$ whenever $X_{ij}^*(\theta_i) = 0$ (which happens with probability α_{ij}), we get $\tilde{v}_j^* = \alpha_{ij}\mathbb{E}[\tilde{v}_j(\theta_i)|X_{ij}^*(\theta_i) = 1] = \alpha_{ij}\mathbb{E}[p_{ij}(\theta_i, S_i^*)|X_{ij}^*(\theta_i) = 1]$.

Thus the maxmin value of agent i for item j is, by Theorem 3.4.1 at least

$$\bar{V}_{ij} \geq \frac{1}{2}\mathbb{E}[p_{ij}(\theta_i, S_i^*)|X_{ij}^*(\theta_i) = 1]$$

We now argue that maxmin value of agent playing a chromatic auction is bounded by the sum of maxmin values $\bar{V}_{ij}$ for the utility she gets from the individual items j. First note that any set of strategies $\mathcal{S}_{ij}$ agent i can adopt for bidding on item j can always be combined together into a single strategy $\mathcal{S}_i$ for the chromatic auction, and because the payments and budgets for different auctions are fully decoupled, the expected utility in the chromatic auction is just sum of expected utilities agent i gets for the individual items. Thus, the strategy of playing single-item maxmin strategies under valuations $\tilde{v}_{ij}$ is a feasible strategy that yields an expected utility of

$$\bar{V}_i \geq \sum_j \bar{V}_{ij} \geq \frac{1}{2}\sum_j \alpha_{ij}\mathbb{E}[p_{ij}(\theta_i, S_i^*(\theta_i))|X_{ij}^*(\theta_i) = 1] = \frac{1}{2}Tv_i^*(\alpha).$$

Here for the last equality we employed the characterization (3.14) we derived earlier. Thus, the strategy of playing single-item maxmin strategy under valuations $\tilde{v}_{ij}$ guarantees expected utility of at least $u_i \geq \sum_j \bar{V}_{ij}$, which yields the theorem statement. □

3.6 Shared Currency Auctions

We have shown that Chromatic Auction mechanism (Algorithm 4) satisfies 1/2 - Utopia, however its one potential disadvantage is that its relative complexity: if there are many items j (for example different types of food in the case of foodbank allocations), having a separate currency for each item type might become unwieldy. A simpler mechanism (Algorithm 5) would employ a single currency for all items.

ALGORITHM 5: Shared Currency Auctions

Require: Budgets B_i, $\sum_i B_i = T$

Let $B_i^1 = B_i$

for all $t = 1$ to T **do**

  Every agent i submits bids $b_{i,j}^t$, $\sum_j b_{ij}^t \leq B_i^t$

  Each item is allocated to highest bidder, and the winner is charged their bid $B_i^{t+1} = B_i^t - \sum_j b_{i,j}^t \mathbf{1}[b_{i,j} \geq b_{-i,j}]$

end for

Consider the constrained ideal per-period utility profile:

$$v_i^*(\alpha_i) = \max_{X_i(\theta_i)} \mathbb{E}\left[v_i\left(\theta_i, X_i(\theta_i)\right)\right] \qquad (3.7)$$
$$s.t.\ \mathbb{E}\left[X_{ij}(\theta_i)\right] \leq \alpha_i.$$

This differs from our definition (3.5) in that we require that the fractions of items agent i is "endowed" with are all equal to a single constant α_i.

The main result for this section are as follows.

Theorem 3.6.1. *Simultaneous Auction mechanism M_{SC} satisfies $\frac{1}{2}$-Utopia with $B_i = \alpha_i$:*

$$\bar{V}_i(M_{SA}) \geq \frac{1}{2} T v_i^*(\alpha) \left[1 - O\left(\frac{\log T}{\sqrt{T}} \right) \right]$$

This result is weaker than Theorem 3.5.2, as it only allows to implement approximate utopia for some coefficients $\alpha_{ij} = \alpha_i$. In particular, as we discuss in the Section 3.5, for particular type distributions the coefficients α_{ij} implementable by the shared currency auctions may not be Pareto Efficient in the space of all possible allocation rules. (it however is Pareto efficient when the type distributions of agents are symmetric).

3.6.1 Generalized Bang for Buck Lemma

In this section we establish a core lemma, an analog of Lemma 3.4.2, that shows how agent i, given any fictitious allocation rule, can get approximate utility under this allocation rule by playing a certain Bernoulli bidding strategy.

First, we define some terms. Let $\hat{X}_i(\theta_i)$ be a fictitious allocation rule for agent i, i.e. $0 \leq \hat{X}_{ij}(\theta_i) \leq 1$ for any $\theta_i \in \Theta_i$. As in the previous section we also use notation $\hat{S}_i(\theta_i) = \{j \mid \hat{X}_{ij}(\theta_i) = 1\}$

Let $\bar{p}_{ij} = \mathbb{E}_{\theta_i}\left[p_{ij}(\theta_i, S_i(\theta_i)) | \hat{X}_{ij}(\theta_i) = 1\right]$ (as per Lemma 4.4.4,), and define the expected utility $v_{\hat{X}} = \mathbb{E} v_i(\theta_i, \hat{X}_i(\theta_i)) = \sum_j \alpha_i \bar{p}_{ij}$. We now define strategy $\mathcal{S}(X(\cdot), r)$ for agent i as a binary strategy:

$$b_{ij}^t(\theta_i) = \begin{cases} \rho_j & if\ \hat{X}_{ij}(\theta_i) = 1 \\ 0 & otherwise, \end{cases} \tag{3.8}$$

where $\rho_j = \frac{r\bar{p}_{ij}}{\sum_j \bar{p}_{ij}}$ and $r \geq 1$ is a constant (particularly useful will be the cases $r = 1$ and $r = 2$).

Lemma 3.6.2 (Bang-for-Buck lemma). *Fix any fictitious allocation rule* $\hat{X}_{ij}(\theta_i)$. *Agent* i *that adopts the strategy* $\mathcal{S}(\hat{X}, r)$ *(as described above) is then guaranteed to get an expected utility of*

$$u_i(\mathcal{S}(\hat{X}, r), \mathcal{S}_a) \geq \frac{\nu_i(\hat{X})}{r\mathbb{P}[\hat{X}_{ij}(\theta_i) = 1]} \cdot c_i(\mathcal{S}(\hat{X}, r), \mathcal{S}_a)$$

for any strategy $\mathcal{S}_a$ *adopted by the adversary.*

The proof of the lemma is similar to that of Lemma 3.4.2 and is deferred to Appendix.

3.6.2 β-Utopia

We now are ready to prove Theorem 3.6.1.

To specify a particular Bernoulli strategy, we fix fictitious allocation rule $\hat{X}_i(\theta_i) = X_i^*(\theta_i)$ to be the ideal allocation rule of agent i (3.7). We then have $\mathbb{P}[X_{ij}^*(\theta_i) = 1] = \alpha_i$, $v_{X^*} = v_i^*(\alpha)$ (where v_i^* is the ideal expected utility as defined in (3.7)).

Suppose agent i adopts the Bernoulli strategy $\mathcal{S}(X^*(\cdot), 2)$, i.e. bidding $b_i^t = \frac{2\bar{p}_{ij}}{\sum_j \bar{p}_{ij}}$ whenever $X_{ij}^*(\theta_i) = 1$, where $\bar{p}_{ij} = \mathbb{E}[p_{ij}(\theta_i, S^*(\theta_i))|X_{ij}^*(\theta_i) = 1]$.

Lemma 3.6.2 now implies that agent i obtains $\frac{v_i^*(\alpha)}{2}$ per unit of expected spending, and thus in order to prove Theorem 3.6.1 we only need to show that expected spending of agent i is close to $\alpha_i T$ when playing strategy $\mathcal{S}(X^*(\cdot), 2)$. Similarly to our proof of Theorem 3.4.1, we do this by first introducing an auxiliary lemma:

Lemma 3.6.3. *Suppose agent* i *employs strategy* $\mathcal{S}(X^*(\cdot), 2)$*, and the adversary strives to minimize the expected spending* $c_i(\mathcal{S}(X^*(\cdot), 2), \mathcal{S}_{-i})$ *of agent* i*. Then it* *is optimal for adversary to always bid* $b^t_{aj} = \rho_j$.

The intuition behind this lemma as well as its proof is analogous to the one employed in the proof of Lemma 3.4.4 and for brevity we defer the proof to the Appendix.

Using this lemma we can now prove that expected spending of agent i is in fact close to her total budget $B_i = \alpha_i T$.

Lemma 3.6.4. *If agent* i *employs* $\mathcal{S}(X^*(\cdot), 2)$ *then*

$$c_i(\mathcal{S}(X^*(\cdot), 2), \mathcal{S}_a) \geq B_i(1 - \sqrt{\frac{2l \log T}{\alpha_i T}} - O(1/T))$$

for any strategy $\mathcal{S}_a$ *adopted by the adversary.*

Proof. This proof closely mimics that of Lemma 3.4.3.

Proving this lemma for any strategy of the adversary is equivalent to proving it when the adversary adopts a strategy $\mathcal{S}^*_a$ that minimizes the expected spending $c_i(\mathcal{S}_2, \mathcal{S}^*_a)$. When adversary employs strategy S^*_a and always bids $b^t_{aj} = \rho_j$, the number of periods before he runs out of budget can be bounded by $t = \frac{B_a}{2} \leq T/2$ (as sum of all bids in a round is $\sum_{j=1}^m \rho_j = 2$). Thus in the remaining $t > T/2$ periods agent i bids $Y_t = \sum b^t_{ij} = \sum_{j=1^t} \rho_j X^*_{ij}(\theta_i)$. The spending of agent i in these periods is

$$\sum_{t=T/2}^{T} C^t_i \geq \min(B_i, \sum_{t=T/2}^{T} Y_t).$$

We now can apply Chernoff bound to $Y \triangleq \sum_{t=T/2}^T Y_t$ to show that, with high probability, it is close to B_i. Y is a sum of independent random variables, each bounded by 2, and

$$\mathbb{E}Y = \frac{T}{2}\sum_{j=1}^{m} \rho_j \mathbb{P}[b_{ij} = \rho_j] = \frac{T}{2}\alpha_i \sum_{j=1}^{m} \rho_j = \alpha_i T$$

By Chernoff bound, we now obtain

$$\mathbb{P}[Y/2 < \alpha_i T(1-\delta)/2] \leq e^{\frac{-\delta^2 B_i}{4}}$$

Now substituting $\delta = 2\log T/\sqrt{\alpha_i T}$ and $B_i = \alpha T$ yields the lemma's statement.

□

Now, we have shown that agent i's expected spending is at least $B_i(1-\sqrt{\frac{2l\log T}{\alpha_i T}})$, and substituting this value into Bang-for-Buck lemma (Lemma 3.6.2) yields the statement of the theorem.

3.7 Additional results and proofs

3.7.1 Individual rationality

An additional design constraint often considered in mechanism design literature is individual rationality. In the context of non-monetary allocation, one can impose this constraint by assuming that there is some status quo allocation rule (e.g. a lottery) and requiring that no agent is worse off in the new mechanism when compared to the status quo.

Such guarantees are not the focus of this paper, but in this section we show that, for the particular choices of the status quo allocation rule, artificial currency mechanisms we have presented are in fact individually rational.

Formally, we define a status quo outcome as one where agent i gets items according to some allocation rule X_{ij}^{sq}, where in the role of X^{sq} we would consider various lotteries. We also define status quo expected per-period utilities

$$v_i^{sq} = \mathbb{E}_{\theta_i, X_i^{sq}]} v_i(\theta_i, X_i^{sq}), \tag{3.9}$$

We can now formally define individual rationality.

Individual rationality (IR): All agents can secure expected utility that is (weakly) larger than their status quo utility v_i^{sq}

$$\bar{V}_i(M) \geq T v_i^{sq}(\alpha) \tag{3.10}$$

Single item setting

In the case of a single item allocation, a natural status quo allocation rule X_i^{sq} is a weighted lottery, with agent i getting the item with probability α_i. The per-period expected utility for agent i is then

$$v_i^{sq}(\alpha) = \alpha_i \cdot \mathbb{E}\left[\theta_i\right]$$

We now have the following result:

Theorem 3.7.1. *First Price Auction mechanism M_{FP} with $B_i = \alpha_i T$ satisfies*

$$\bar{V}_i(M_{FP}) \geq T v_i^{sq}(\alpha).$$

Consider a fictitious allocation rule $\hat{X}_i(\theta_i) = 1$, $\nu_i(\hat{X}) = \mathbb{E}\theta_i$ and the corresponding strategy $\mathcal{S}(\hat{X}_i(\cdot), 1)$ (all notation here is as defined in Section 3.4), that is bidding $b_i^t = 1$ on all periods.

With the following lemma we show that adopting this strategy ensures that agent i spends all of her budget.

Lemma 3.7.2 (Spending lemma)**.** *If agent i adopts constant bid strategy $\mathcal{S}(\hat{X}_i(\cdot), 1)$ her expected spending is $c_i = B_i = \alpha_i T$, for any strategy adopted by other agents $\mathcal{S}_{-i}$*

Proof. To prevent agent i from spending $C_i^t = 1$ on any round, some other agent needs to spend at least $C_{-i}^t = 1$. Since the budget of all agents but i is $B_a = (1 - \alpha_i)T$, there have to be at least $\alpha_i T$ periods when agent i wins the item and pay $C_i^t = 1$. [4] The statement of the lemma now follows from the fact that the agent cannot win more than $\alpha_i T$ rounds with the constant bid strategy. □ □

We can now combine this result with the Bang-per-Buck lemma to prove the Theorem. By Lemma 3.7.2 the expected spending of agent i is $c_i = \alpha_i T$, and by Lemma 3.4.2 the agent obtains v_i^{sq} of expected utility per unit of spending. Thus the total expected utility of the agent is $u_i(\mathcal{S}_1, \mathcal{S}_{-i}) \geq v_i^{sq} c_i = \alpha_i T v_{\hat{X}} = T v_i^{sq}(\alpha)$.

Shared Currency auctions

We now switch to the setting of allocating multiple items and equal marginal probabilities $\alpha_{ij} = \alpha_i$ (the setting of Section 3.6, in which we proposed using a Shared Currency Auctions mechanism (Algorithm 5). Here a status quo allocation rule X^{sq} we consider is weighted all-or-nothing lottery, that is a lottery in which agent i obtains the bundle of all items with probability α_i and nothing otherwise. . The per-period expected utility for agent i under this lottery is then

[4] $\alpha_i T$ may not be an integer, in which case agent is guaranteed to win the item on $\lfloor \alpha_i T \rfloor$ periods. However, the corresponding loss in utility does not scale with T and is omitted for readability.

$$v_i^{sq}(\alpha) = \alpha_i \mathbb{E} v_i(\theta_i, [m])$$

We note that this all-or-nothing lottery may not be a Pareto optimal one, in particular performing a lottery on each item independently might yield better expected utilities for all agents. Unfortunately, our proof technique (and, potentially, the use of maxmin value as the solution concept) does not allow to prove IR constraints for arbitrary lotteries.

Theorem 3.7.3. *Shared Currency Auctions mechanism M_{SC} satisfies IR with $\alpha_i = B_i$:*

$$\bar{V}_i(M_{SA}) \geq T v_i^{sq}(\alpha)$$

Consider a fictitious allocation rule $\hat{X}_{ij}(\theta) = 1$. We then get $\mathbb{P}[\hat{X}_{ij}(\theta_i) = 1] = 1$,

$$\nu_i(\hat{X}) = \mathbb{E} v_i(\theta_i, [m]) = \sum_{j=1}^{m} \mathbb{E} p_{ij}(\theta_i, [m]).$$

(all notation here is as defined in 3.6).

Consider the strategy $\mathcal{S}(\hat{X}, 1)$, i.e. always bidding $b_i^j = \frac{\bar{p}_{ij}}{\sum_j \bar{p}_{ij}}$. The Generalized Bang-for-Buck Lemma 3.6.2 then implies that, for any strategy $\mathcal{S}_a$ adopted by the adversary we have

$$u_i(\mathcal{S}(\hat{X}, 1), \mathcal{S}_a) \geq \nu_i(\hat{X}) c_i(\mathcal{S}(\hat{X}, 1), \mathcal{S}_a).$$

Since the status quo per-period expected utility $v_i^{sq} = \frac{\nu_i(\hat{X})}{\alpha_i}$, it is enough to show that expected spending $c_i(\mathcal{S}(\hat{X}, 1), \mathcal{S}_a) = B_i = \alpha_i T$ in order to prove Theorem 3.7.3.

We do this in the following lemma, which is a straightforward generalization of Lemma 3.7.2.

Lemma 3.7.4. *In the Shared Currency Auctions mechanism, if agent i adopts the strategy $\mathcal{S}(\hat{X}, 1)$, her expected spending is $c_i(\mathcal{S}(\hat{X}, 1), \mathcal{S}_a) = \alpha_i T$ for any strategy $\mathcal{S}_a$ chosen by the adversary.*

Proof. Let y_j^t be the indicator that adversary bids $b_a^t \geq \rho_j$ or more and wins the item j on period t. Let $I^t = \{j|y_j^t = 1\}$, $O^t = \{j|y_j^t = 0\}$.

For spending of adversary and agent i on period t we have

$$C_a^t \geq \sum_{j \in I^t} \rho_j$$

$$C_i^t = \sum_{j \in O^t} \rho_j$$

This implies that the sum of payments on any round is

$$C_i^t \geq \sum_{j=1}^{m} \rho_j - C_j^t = 1 - C_j^t$$

. Because the budget of the adversary is limited $\sum_t B_a = (1 - \alpha_i)T$, we have $\sum C_i^T \geq T - \sum_t C_j^t \geq T - B_a = \alpha_i T$. □

Chromatic Auction

We finally consider the setting of multiple items and arbitrary endowments α_{ij}, for which in section 3.5 we analyzed the Chromatic Auction mechanism (Algorithm 4).

When probabilities of getting different items α_{ij} differ between items, the all-or-nothing lottery is no longer well-defined, which presents an obstacle in generalizing our guarantee from the case of equal weights α_{ij} we considered above.

We assume that all agents are endowed with some probabilities α_{ij} of getting item j on round i. However, because value of agents depends on the set of the items allocated, the marginals α_{ij} alone are not enough to characterize the agent's expected utility. In this section we adopt an individual rationality constraint, that, roughly, requires that all agents are no worse off than under the worst case lottery that guarantees probabilities α_{ij} for getting the individual items.

Formally, let $X_i \sim G$ be a random allocation rule that is independent of θ_i (a lottery). We define status quo expected utility as follows:

$$u_i^{sq} = \min_G \mathbb{E}_{\theta_i, X_i \sim G} v_i(\theta_i, X_i) \tag{3.11}$$

$$s.t.\ \mathbb{E}[X_{ij}] \geq \alpha_{ij} \tag{3.12}$$

Using this value as a benchmark, we can now state our Individual Rationality result.

Theorem 3.7.5. *Chromatic Auction mechanism M_{CA} satisfies Individual Rationality for the worst-case lottery*

$$\bar{V}_i(M_{FP}) \geq T v_i^{sq}(\alpha)$$

Similarly to Theorem 3.5.2, we prove this result by using results from the single item setting (Section 3.4). Consider a simple constant bid strategy for agent i, $b_{ij}^t = 1$. According to Lemma 3.7.2, agent i is going to win item j on at least

$\alpha_{ij}T$ periods. This implies $\sum_t X_{ij}^t \geq \alpha_i j$, where X^t is the realized allocation on round t. Let $\tilde{G}$ be the empirical distribution of X^t, i.e. the empirical distribution of sets agent i has won on different periods. For any t have $\mathbb{P}_{\hat{X}\sim\tilde{G}}[\hat{X} = X^t] = 1/T$

$$\mathbb{E}_{\theta_i}[\sum_{t=1}^{T} U_i^t | \{X^t\}_{t\in[T]}] = T\sum_{t=1}^{T}\frac{1}{T}v_i(\theta_i, X^t) = T\mathbb{E}_{\theta_i,\hat{X}\sim\tilde{G}}v_i(\theta_i, \hat{X})$$

Since, as we argued above, the empirical distribution $\tilde{G}$ satisfies the marginal probability constraints $\mathbb{P}[X_{ij} = 1] \geq \alpha_{ij}$, this bound implies the worst-case bound (3.11) we are aiming to prove.

3.7.2 Proofs

Single item

Proof of Lemma 3.4.3. Proving this lemma for any strategy of the adversary is equivalent to proving it when the adversary adopts a strategy $\mathcal{S}_a^*$ that minimizes the expected spending $c_i(\mathcal{S}_2, \mathcal{S}_a^*)$. When adversary employs strategy S_a^* and always bids $b_a^t = 2$, the number of turns before he runs out of budget can be bounded by $t = \frac{B_a}{2} \leq T/2$. Thus in the remaining $t > T/2$ rounds agent i pays $C_i^t = 2$ whenever she bids $b_i^t = 2$, and thus we get

$$\sum_{t=1}^{T} C_i^t = \min(B_i, \sum_{t=T/2}^{T} b_i^t) \tag{3.13}$$

We now can apply a Chernoff bound to $Y = \sum_{t=T/2}^{T} b_i^t$ to show that, with high probability, it is close to B_i. Y is a sum of $T/2$ independent random variables, each bounded by 2, and

$$\mathbb{E}Y = 2\alpha_i\frac{T}{2} = \alpha_i T = B_i.$$

By Chernoff, we obtain

$$\mathbb{P}[Y/2 < B_i(1-\delta)/2] \leq e^{\frac{-\delta^2 B_i}{4}}$$

Since $C_i = \min(B_i, \hat{C}_i)$, $\hat{C}_i > B_i(1-\delta) \Rightarrow C_i > B_i(1-\delta)$. Now substituting $\delta = 2\log T/\sqrt{\alpha_i T}$ and $B_i = \alpha T$ yields the lemma's statement. □ □

Multiple items

Proof of Theorem 3.5.2. The ideal per-period expected utility $v_i^*(\alpha)$ and allocation rule $X_{ij}^*(\theta_i)$ are given by (3.5), and using the above lemma and denoting $S_i^*(\theta_i) = \{j | X_{ij}^*(\theta) = 1\}$ we can rewrite $v_i^*(\alpha)$ as

$$\begin{aligned} v_i^*(\alpha) = \mathbb{E}v_i(\theta_i, X_i^*(\theta_i)) &= \mathbb{E}\sum_{j=1}^{m} p_j(\theta_i, S_i^*)X_{ij}^*(\theta_i) \\ &= \sum_{j=1}^{m} \mathbb{E}[p_j(\theta_i, S_i^*)|X_{ij}^*(\theta_i) = 1]\mathbb{P}[X_{ij}^*(\theta_i) = 1] \\ &= \sum_{j=1}^{m} \alpha_{ij}\mathbb{E}[p_j(\theta_i, S_i^*)|X_{ij}^*(\theta_i) = 1]. \end{aligned} \tag{3.14}$$

We prove the theorem using a reduction to the single item setting (Theorem 3.4.1).

Suppose on some rounds the agent wins some set of items S. Using Lemma 4.4.4 we get

$$U_i^t \geq \sum_{j\in S} p_j(\theta_i, S_i^*(\theta_i))X_{ij}^*(\theta_i).$$

Thus the utility of the agent i can be lower-bounded with the utility of an agent with additive valuation function

$$\tilde{v}_i(\theta_i, S) = \sum_{j\in S} p_j(\theta_i, S_i^*(\theta_i))X_{ij}^*(\theta_i) = \sum_{j} \tilde{v}_{ij}(\theta_i),$$

where $\tilde{v}_{ij}(\theta_i) = p_j(\theta_i, S_i^*(\theta_i))X_{ij}^*(\theta_i)$.

Now, we concentrate on an item j and valuation $\tilde{v}_j(\theta_i)$ and apply Theorem 3.4.1 directly. Recall that the ideal utility $v_{ij}^*(\alpha_j)$ for the single item (the item j) is given by:

$$\tilde{v}_{ij}^*(\alpha_j) = \max_X \ \mathbb{E}\left[\tilde{v}_j(\theta_i)X_{ij}(\theta_i)\right] \quad \text{s.t., } \mathbb{E}\left[X_{ij}(\theta_i)\right] = \alpha_{ij}$$

Since $\tilde{v}_j(\theta_i) = 0$ whenever $X_{ij}^*(\theta_i) = 0$ (which happens with probability α_{ij}), we get $\tilde{v}_j^* = \alpha_{ij}\mathbb{E}[\tilde{v}_j(\theta_i)|X_{ij}^*(\theta_i) = 1] = \alpha_{ij}\mathbb{E}[p_{ij}(\theta_i, S_i^*)|X_{ij}^*(\theta_i) = 1]$.

Thus the maxmin value of agent i for item j is, by Theorem 3.4.1 at least

$$\bar{V}_{ij} \geq \frac{1}{2}\mathbb{E}[p_{ij}(\theta_i, S_i^*)|X_{ij}^*(\theta_i) = 1]$$

We now argue that maxmin value of agent playing a chromatic auction is bounded by the sum of maxmin values $\bar{V}_{ij}$ for the utility she gets from the individual items j. First note that any set of strategies $\mathcal{S}_{ij}$ agent i can adopt for bidding on item j can always be combined together into a single strategy $\mathcal{S}_i$ for the chromatic auction, and because the payments and budgets for different auctions are fully decoupled, the expected utility in the chromatic auction is just sum of expected utilities agent i gets for the individual items. Thus, the strategy of playing single-item maxmin strategies under valuations $\tilde{v}_{ij}$ is a feasible strategy that yields an expected utility of

$$\bar{V}_i \geq \sum_j \bar{V}_{ij} \geq \frac{1}{2}\sum_j \alpha_{ij}\mathbb{E}[p_{ij}(\theta_i, S_i^*(\theta_i))|X_{ij}^*(\theta_i) = 1] = \frac{1}{2}Tv_i^*(\alpha).$$

Here for the last equality we employed the characterization (3.14) we derived earlier. Thus, the strategy of playing single-item maxmin strategy under valuations

$\tilde{v}_{ij}$ guarantees expected utility of at least $u_i \geq \sum_j \bar{V}_{ij}$, which yields the theorem statement. □ □

Proof of Lemma 3.6.2. Fix a round t and assume the agent has sufficient budget left (i.e., $B_i^t r$). For adversary's bids b_a^t, let $\bar{S} = \{j \mid b_{aj}^t < \rho_j\}$ denote the set of items on which the adversary does not block the agent. Since agent i only bids on items in $\hat{S}_i(\theta_i)$, her expected spending in this round is given by

$$\begin{aligned}\mathbb{E}[C_i^t | \ b_a^t, B_i^t \geq r] &= \mathbb{E}\left[\sum_{j \in \bar{S}} \rho_j \mathbb{I}[\hat{X}_{ij}(\theta_i) = 1]\right] \\ &= \sum_{j \in \bar{S}} \rho_j \mathbb{P}\left(X_{ij}(\theta_i) = 1\right),\end{aligned}$$

Similarly, the agent's expected utility in this round is given by

$$\begin{aligned}\mathbb{E}\left[U_i^t | \ b_a^t, B_i^t \geq r\right] &= \mathbb{E}[v_i(\theta_i, \hat{S}_i(\theta_i) \cap \bar{S})] \\ &\geq \mathbb{E}\left[\sum_{j \in \bar{S}} p_{ij}(\theta_i, \hat{S}_i) \mathbb{I}\{\hat{X}_{ij}(\theta_i) = 1\}\right] \\ &= \sum_{j \in \bar{S}} \bar{p}_{ij} \cdot \mathbb{P}\left(\hat{X}_{ij}(\theta_i) = 1\right).\end{aligned}$$

Here, we have used Lemma 4.4.4 for the first inequality. Recalling the expression for ρ_j, we get

$$\mathbb{E}[U_i^t | \ b_a^t, B_i^t \geq r] \geq \frac{\sum_j \bar{p}_{ij}}{r} \mathbb{E}[C_i^t | \ b_a^t, B_i^t \geq r] = \frac{\nu_i(\hat{X})}{r\mathbb{P}[\hat{X}_{ij}(\theta_i) = 1]} \mathbb{E}[C_i^t | \ b_a^t, B_i^t \geq r]$$

Finally, note that if the agent's budget is already depleted on period t, then both the expected utility and the spending are zero on this period, and hence the bound above holds as well. Since the bound holds for any time t and arbitrary adversary bids b_a, we obtain the lemma statement by summing up the bound over all periods and taking the expectation over potential randomness in the strategy of the adversary. □ □

Proof of Lemma 3.6.4. This proof closely mimics that of Lemma 3.4.3.

Proving this lemma for any strategy of the adversary is equivalent to proving it when the adversary adopts a strategy $\mathcal{S}_a^*$ that minimizes the expected spending $c_i(\mathcal{S}_2, \mathcal{S}_a^*)$. When adversary employs strategy S_a^* and always bids $b_{aj}^t = \rho_j$, the number of periods before he runs out of budget can be bounded by $t = \frac{B_a}{2} \leq T/2$ (as sum of all bids in a round is $\sum_{j=1}^m \rho_j = 2$). Thus in the remaining $t > T/2$ periods agent i bids $Y_t = \sum b_{ij}^t = \sum_{j=1^t} \rho_j X_{ij}^*(\theta_i)$. The spending of agent i in these periods is

$$\sum_{t=T/2}^{T} C_i^t \geq \min(B_i, \sum_{t=T/2}^{T} Y_t).$$

We now can apply Chernoff bound to $Y \triangleq \sum_{t=T/2}^T Y_t$ to show that, with high probability, it is close to B_i. Y is a sum of independent random variables, each bounded by 2, and

$$\mathbb{E}Y = \frac{T}{2} \sum_{j=1}^{m} \rho_j \mathbb{P}[b_{ij} = \rho_j] = \frac{T}{2} \alpha_i \sum_{j=1}^{m} \rho_j = \alpha_i T$$

By Chernoff bound, we now obtain

$$\mathbb{P}[Y/2 < \alpha_i T(1-\delta)/2] \leq e^{\frac{-\delta^2 B_i}{4}}$$

Now substituting $\delta = 2 \log T / \sqrt{\alpha_i T}$ and $B_i = \alpha T$ yields the lemma's statement. □

3.7.3 Dynamic Programming Arguments

Single Item

In this section, we consider the adversary's best response when the agent follows a Bernoulli strategy. Specifically, we study the adversary's problem of minimizing the agent's expected payment in order to prove Lemma 3.4.4.

Before we analyze the decision problem, we note that since the agent follows a Bernoulli strategy, the adversary's decision in any round can be reduced to a binary action: whether or not to *block* the agent from winning that round. When the adversary decides not to block the agent, her optimal action is to submit a bid $b_a^t = 0$, whereas when the adversary chooses to block the agent, her optimal action is to submit a bid $b_a^t = 2$, where we assume that the tie is always broken in the adversary's favor. Thus, we can reduce the adversary's decision to a binary action $y_a^t \in \{0, 1\}$, where $y_a^t = 1$ implies the adversary blocks the agent, whereas $y_a^t = 0$ implies that the adversary does not block the agent.

Given the preceding discussion, we seek to show that the adversary's optimal strategy is to keep blocking the agent until the adversary's budget runs out. We show this using a dynamic programming argument. Specifically, let y_i^t denote the Bernoulli random variable that denotes whether or not the agent bids on the item at time t. We have the following expressions for the agent's spending until time t, and the left-over budgets:

$$C_i^t = 2y_i^t(1 - y_a^t)\mathbb{I}\{B_i^{t-1} \geq 2\} \quad , \quad B_i^t = B_i - \sum_{\tau=1}^{t-1} C_i^\tau \quad , \quad B_a^t = B_a - \sum_{\tau=1}^{t-1} 2y_a^t.$$

We define the following DP: the state is given by the tuple $x_t = (T - t + 1, B_a^t, B_i^t)$ (i.e., time-to-go and left-over budgets), where $t \in [T]$, and
$B_i^t \in \{0, 1, \cdots, \lceil \alpha_i T \rceil\}, B_a^T \in \{0, 1, \cdots, \lceil T(1 - \alpha_i) \rceil\}$. The adversary's available actions are $y_a^t \in \{0, 1\}$ if $B_a^t \geq 2$, and $y_a^t \in \{0\}$ otherwise. The transitions are

specified by

$$\text{if } y_a^t = 1, \text{ then } x_{t+1} = x_t - (1,2,0),$$

$$\text{if } y_a^t = 0, \text{ then } x_{t+1} = \begin{cases} x_t - (1,0,0) - (0,0,2)\mathbb{I}\{B_i^t \geq 2\} & \text{w.p. } \alpha_i, \\ x_t - (1,0,0) & \text{w.p. } 1-\alpha_i. \end{cases}$$

Now the adversary's decision problem is to choose a policy $\pi = \{y_a^t : t \geq 0\}$ that minimizes

$$\mathbb{E}^\pi \left[\sum_{t=1}^{t} 2\alpha_i \mathbb{I}\{B_i^t \geq 2\}(1 - y_a^t) \right].$$

Let $e_i, i \in \{1,2,3\}$ denote the standard unit vectors (i.e., $e_1 = (1,0,0)$, etc.). The Bellman equation can now be written as follows via the following cases:

1. If $t \leq T$, $B_i^t \geq 2$ and $B_a^t \geq 2$

   $$J(x_t) = \min\{J(x_t - e_1 - 2e_2), 2\alpha_i + \alpha_i J(x_t - e_1 - 2e_3) + (1-\alpha_i)J(x_t - e_1)\}$$

2. If $t \leq T$, $B_i^t < 2$ and $B_a^t \geq 2$

   $$J(x_t) = \min\{J(x_t - e_1 - 2e_2), J(x_t - e_1)\}$$

3. If $t \leq T$, $B_i^t \geq 2$ and $B_a^t < 2$

   $$J(x_t) = 2\alpha_i + \alpha_i J(x_t - e_1 - 2e_3) + (1-\alpha_i)J(x_t - e_1)$$

4. In all other cases

   $$J(x_t) = 0$$

Now we have the following structural characterization for the value function.

Lemma 3.7.6. *For any $t \leq T, B_a^t \geq 2$, the adversary's optimal action is $y_a^t = 1$ (i.e., block the agent).*

Proof. To prove the lemma, it is sufficient to show that in cases (1) and (2), the first term (corresponding to blocking) is the minimizer; thus, we need to show that $J(x_t) = J(x_t - e_1 - 2e_2)$ when $t \leq T$ and $B_a^t \geq 2$. For case (2), note that if $B_i^t < 2$, then $J(x_t) = J(x_t - e_1 - 2e_2) = 0$. To complete the proof, we need to show that $J(k, x + 2, y) = J(k - 1, x, y)$ for all $x, y \geq 2$.

We now show this via induction on k. For the base case, we have $J(1, x+2, y) = 0 = J(0, x, y)$ for any $x, y \geq 2$, since with one period left, the adversary is always better off blocking. Next, assume that for some $k \geq 1$, we have $J(k, x + 2, y) = J(k - 1, x, y)$ for all $x, y \geq 2$. Then, we have

$$
\begin{aligned}
&J(k + 1, x + 2, y) = \\
&= \min\{J(k, x, y), 2\alpha_i + \alpha_i J(k, x + 2, y - 2) + (1 - \alpha_i)J(k, x + 2, y)\} \\
&= \min\{J(k - 1, x - 2, y), 2\alpha_i + \alpha_i J(k - 1, x, y - 2) + (1 - \alpha_i)J(k - 1, x, y)\} \\
&= J(k, x, y)
\end{aligned}
$$

This completes the induction step. □

Shared Currency Auction

Here we give the proof of Lemma 3.6.3, that states that it is optimal for adversary to "block" all items $b_{aj}^t = \rho_j$ until he runs out of budget.

We assume that once the budget of adversary falls below 2 he stops blocking any items. (since the utility loss agent would experience as a result of adversary not doing so does not scale with T).

The proof of the lemma closely follows the one we provided above.

We define the dynamic program for the adversary as follows. The state is defined as before $x_t = (T-t+1, B_a^t, B_i^t)$. The set of available actions is $A = 2^{[m]}$ (which items to block on this period). Upon taking the action $S \subseteq [m]$, the transitions from the state $(T-t+1, B_a^t, B_i^t)$ to state $(T-t, B_a^t - \sum_{j\in S}\rho_j, B_i^t - \sum_{j\in A\subseteq S^c}\rho_j)$ with probability $\alpha_i^{|A|}(1-\alpha_i)^{|S^c\setminus A|}$ (the adversary pays for the items that he blocks). Given this, the bellman equation is

$$
\begin{aligned}
&J(T-t+1, B_a^t, B_i^t) = \\
&= \min_{S\in[m]}\Bigg\{\sum_{A\subseteq S^c}\alpha_i^{|A|}(1-\alpha_i)^{|S^c\setminus A|}\sum_{j\in A}\rho_j \\
&\quad + \sum_{A\subseteq S^c}\alpha_i^{|A|}(1-\alpha_i)^{|S^c\setminus A|}J(T-t, B_a^t - \sum_{j\in S}\rho_j, B_i^t - \sum_{j\in A\subseteq S^c}\rho_j)\Bigg\} \\
&= \min\{J(T-t, B_a^t - 2, B_i^t), \min_{S\in[m], S\neq[m]}\Bigg\{\sum_{A\subseteq S^c}\alpha_i^{|A|}(1-\alpha_i)^{|S^c\setminus A|}\sum_{j\in A}\rho_j \\
&\quad + \sum_{A\subseteq S^c}\alpha_i^{|A|}(1-\alpha_i)^{|S^c\setminus A|}J(T-t, B_a^t - \sum_{j\in S}\rho_j, B_i^t - \sum_{j\in A\subseteq S^c}\rho_j)\Bigg\}\}
\end{aligned}
$$

To prove the lemma, as before it is sufficient to show that $J(T-t+1, B_a^t, B_i^t) = J(T-t, B_a^T - 2, B_i^t)$ for any $T-t$, which we do by induction. Assume that the equality holds for $T-t = k$, then for $T-t+1 = k+1$ we have

$$
\begin{aligned}
&J(k+1, x+2, y) = \\
&= \min_{S \in [m], S \neq [m]} \Bigg\{ J(k, x, y), \sum_{A \subseteq S^c} \alpha_i^{|A|} (1-\alpha_i)^{|S^c \setminus A|} \sum_{j \in A} \rho_j \\
&\quad + \sum_{A \subseteq S^c} \alpha_i^{|A|} (1-\alpha_i)^{|S^c \setminus A|} J(k, x - \sum_{j \in S} \rho_j, y - \sum_{j \in A \subseteq S^c} \rho_j) \Bigg\} \\
&= \min_{S \in [m], S \neq [m]} \Bigg\{ J(k-1, x-2, y), \sum_{A \subseteq S^c} \alpha_i^{|A|} (1-\alpha_i)^{|S^c \setminus A|} \sum_{j \in A} \rho_j \\
&\quad + \sum_{A \subseteq S^c} \alpha_i^{|A|} (1-\alpha_i)^{|S^c \setminus A|} J(k-1, x - 2 - \sum_{j \in S} \rho_j, y - \sum_{j \in A \subseteq S^c} \rho_j) \Bigg\} \\
&= J(k, x, y)
\end{aligned}
$$

The second equality follows from the induction hypothesis, and the first and the third equality follows from Bellman equation. Thus, we obtain that the minimum in the Bellman equation is attained when the adversary blocks all the items. This proves the lemma statement.

CHAPTER 4

EFFICIENT ONLINE ALLOCATION UNDER ADVERSARIAL VALUATIONS

4.1 Introduction

We consider the following natural allocation problem: a principal has a single divisible resource each day (over a period of T days), and needs to divide it between n agents. Every day, each agent has a value for the resource – these values become known to the principal at the beginning of the day, but no earlier. The principal wants to allocate the resource as fairly and efficiently as possible, despite the lack of information about the future.

Such allocation problems are common in practice. For example, consider the problem of allocating the computational resources of a cluster among the employees of a firm or a university, where people might have very high value for gaining access to a cluster on some days (e.g., due to a conference deadline), and willing to pass up on their access on other days. A similar problem is faced by food banks that allocate food each day to soup kitchens and other local charities [70] – the number of people coming to each distribution facility varies from day to day, and the food bank wants to ensure some measure of equity and efficiency in its allocation. The online nature of this problem makes it quite challenging from an algorithmic perspective, and the main goal of our paper is to design algorithms that provide non-trivial fairness guarantees.

The goals of fairness and efficiency can be formalized in many ways, but the primary focus of this paper is on the objective of maximizing *the Nash social* *welfare (NSW)* – the geometric mean of the agents' utilities. This objective was initially proposed about 70 years ago [65, 52] and it is known to satisfy several desirable properties such as scale-independence, meaning that the scale of the agents' valuations does not affect the NSW maximizing solution [62]. Although computing a solution that maximizes the NSW is a well-understood and computationally tractable problem in an offline setting [68, Chapter 5], to the best of our knowledge this is the first work that studies the extent to which we can optimize this objective when the items arrive online.

To this end, we provide a near-complete characterization: on the positive side, we propose an online algorithm that guarantees a competitive ratio of $\log(\min\{T, n\})$, and on the negative side, we demonstrate a family of instances on which, for any chosen ϵ, no online algorithm can achieve competitive ratio of better than $\log^{1-\epsilon} n$ or $\log^{1-\epsilon} T$.

Going beyond the results themselves, the techniques we introduce may be of independent interest. From an algorithmic perspective, our result is based on using a novel yet intuitive idea of *promised utilities*, which in a certain sense, formalizes notions of an individual endowment in a collective settings. Moreover, we also provide examples which highlight why such a notion is in some sense essential for ensuring fairness in online allocation – in particular, we demonstrate how many simple heuristics can perform very poorly for this problem. Some of our theoretical ideas may thus prove useful for reasoning about practical algorithms for real-world applications.

On the other hand, our negative result is based on constructing an sequence of instances, which make it unavoidable for an algorithm to make mistakes over time. As with many other online decision-making settings, what makes online fair-allocation hard is that choices that appear locally symmetric can lead to vast asymmetry much later in the future. Imagine, on the first day all agents report having identical values. In this situation, an online algorithm can't do better than treating the agents symmetrically. In later days, however, some agents have plenty of opportunities to obtain good value without imposing externalities on others (i.e., when nobody else wants the item much), and an online algorithm may then regret wasting resources in earlier rounds on them. Forcing the algorithm into incurring this impossible-to-avoid regret forms the backbone of our impossibility results. We highlight the applicability of this technique by considering another popular fairness metric, egalitarian welfare, and showing that in the same setting it is impossible to achieve a practically useful competitive ratio for this objective.

4.1.1 Main Results and Outline

We consider a setting where divisible items arrive in an online fashion over a sequence of T rounds and we study the problem of designing online algorithms that decide how the item that becomes available in each round t should be distributed among a group of n agents. Although the decisions of the algorithm for the allocation of the item in each round t are irrevocable, our goal is to provide guarantees regarding the fairness of the final allocation after all of the items have been distributed. We provide a formal definition of the setting and our fairness measures in Section 4.2.

In Section 4.3, we first briefly consider the problem of maximizing the egalitarian social welfare objective, which is equal to the minimum utility across all agents, capturing the intuition that a fair allocation may seek to maximize the happiness of the least happy agent. For this objective we prove that no online algorithm can achieve a competitive ratio better than $\sqrt{n}/3$ or $\sqrt{T}/3$ (see Theorem 4.3.1), indicating that the quality of the final allocation with respect to this objective quickly degrades as the size of the instance grows. In light of this result, the rest of the paper focuses on the Nash social welfare objective, which is equal to the geometric mean of the agents' utilities.

We then consider the Nash social welfare objective, which is equal to the geometric mean of the agents' utilities and is well-known to capture a natural trade-off between fairness and efficiency (see e.g. [22]). In Section 4.4, we first discuss how some natural algorithmic solutions to this problem fail to achieve a good competitive ratio, and then we define the Set-Aside Greedy algorithm and prove that it achieves a competitive ratio of $O(\log n)$ and $O(\log T)$ (see Theorem 4.4.3).

This algorithm works by dividing each item in half and allocating each half using a different approach: the first half is allocated equally between the agents and the second half is allocated in a greedy manner aiming to myopically maximize the Nash social welfare. However, this combination of the two mechanisms would fail, were it not for a crucial adjustment of the agent utilities that are used in the greedy portion. In particular, in each round t, instead of using the actual utility that each agent has accrued up to that round, the greedy algorithm uses higher utility numbers, which we refer to as "promised utilities". The promised utility of each agent in round t corresponds to the actual utility that the agent has accrued from the greedy allocation up to round t, plus an additional utility of $1/(2n)$ which the agent is guaranteed to receive from the other half of the items after the completion of the algorithm. This modification circumvents the issue that the greedy algorithm would otherwise disproportionately prioritize agents with low accrued utility in early rounds, even if their value for the available items is low. At the same time, the algorithm never over-promises, i.e., it ensures that the agents will indeed eventually receive a utility at least as high as their promised utility in any round. In the analysis we leverage duality to design a (to our knowledge, novel) technique for constructing a certificate for the optimality gap of an arbitrary allocation, which is applicable and may be of interest outside of the online setting.

In Section 4.5 we complement the positive result of the Set-Aside Greedy algorithm with a matching inapproximability result. In Theorem 4.5.1, we show that no online algorithm can achieve a competitive ratio of $O(\log^{1-\epsilon} n)$ or $O(\log^{1-\epsilon} T)$ for a constant $\epsilon > 0$. The technique we use for this result is similar in spirit to the one used to prove the inapproximability of the egalitarian social welfare, however the result requires a much more elaborate and technical construction.

We conclude with Section 4.6, where we discuss some of the ways in which our results can be generalized, e.g., to settings with multiple items allocated in each round or settings where the items are indivisible, and we also provide additional motivation and context regarding our model.

4.1.2 Related Work

The work that is most relevant to ours is that of [12] and [76]. [12] consider a setting where a set of items arrives in an sequence of T rounds, and an algorithm needs to allocate them in an online manner. The main differences to our work is that they focus on items that are indivisible, i.e., can be allocated only to a single agent, and that they use the amount of envy among the agents as a measure of fairness. They show that allocating each item uniformly at random among the agents, without taking their values into consideration, leads to an expected maximum envy among the agents of $\tilde{O}(\sqrt{T/n})$ and they argue that this bound it tight up to polylogarithmic factors. They then also propose a more elaborate deterministic algorithm that depends on the agent values and achieves the same bound.

In more recent work, [76] revisit this setting and study the extent to which approximate envy-freeness can be combined with approximate Pareto efficiency. They consider a spectrum of increasingly powerful adversary models and they show that even for a non-adaptive adversary (which is weaker than the adaptive adversary model we consider in this work) there is no algorithm that can guarantee the aforementioned approximate envy while Pareto dominating the random allocation algorithm. In particular, they define an outcome to be α Pareto-efficient if improving every agent's utility by a factor α is infeasible, and they show that it is impossible to combine the envy bound with a $1/n$ approximation of Pareto efficiency, which is a trivial approximation that the random allocation algorithm satisfies.

Although some of the obstacles that we face in this paper are similar, in essence, to the ones that this prior work aims to address, our results are not directly comparable. In particular, we evaluate the quality of the outcomes using aggregate measures of fairness instead of pairwise fairness, so our bounds do not directly translate into interesting guarantees for envy-freeness, or vice-versa. On the other hand, our competitive ratio guarantees do directly translate into approximate Pareto efficiency bounds. In particular a competitive ratio of α for the egalitarian social welfare or the Nash social welfare directly implies an α approximation of Pareto efficiency. Another important difference is that in our setting the items are divisible and the agent values over all rounds are normalized. For a more detailed discussion regarding these assumptions, see Section 4.6.

More generally, our work can be viewed as a part of growing literature on online fair division. [48] study a setting similar to ours, but allow to reallocate some of the previously allocated items, and show that envy-freeness up to one item can be achieved using $O(T)$ reallocations. Other works considered the question of online fair division in a setting where values of agents are random, e.g. [16] consider objectives of fair-share and welfare, [76] concentrate on envy-freeness and Pareto-optimality, and [56] study the notion of maxmin share in such a setting. Another line of work on dynamic fair division, which is not very closely related to our work, considers settings where it is the agents, instead of the items, that arrive (and possibly depart) online, e.g., [35, 34, 59].

The egalitarian social welfare and the Nash social welfare objectives have also played a central role in the literature on the fair allocation of *indivisible* items. Much of this research has focused on the computational complexity of optimizing these objectives, which are both APX-hard. Maximizing the egalitarian social welfare objective is known as the Santa Claus problem [15, 9, 4, 32, 11, 24, 3, 2] and there is still a wide gap between the best known upper and lower bounds. On the other hand, recent work on the Nash social welfare objective has been more fruitful, leading to constant factor approximation algorithms [26, 10]. Also, a very influential paper by [22] showed that maximizing the NSW objective when allocating indivisible items leads to approximate envy-freeness, which provided additional motivation for studying this objective.

Finally, some of the prior work on the NSW objective has also focused on settings where the items are divisible, but values of the agents are private information. In this case, research has focused on the design of truthful mechanisms that approximate the NSW [27] or the analysis of non-truthful mechanisms with respect to their price of anarchy with respect to the NSW [17]. In contrast to this work, in this paper we assume that the agent valuations are public and focus on the complications introduced by the online nature of the problem.

4.2 Setting

There is a set N of n agents, and T rounds with a single divisible item arriving on each round. Each agent i's value for the item on round t is denoted by $v_{i,t} \geq 0$, and the values are normalized so that $\sum_{t=1}^{T} v_{i,t} = 1$[1]. We let $x_{i,t} \geq 0$ be the fraction of the item allocated to the agent i on round t, $\sum_i x_{i,t} = 1$. Utility of agent i is linear over rounds and is defined as $u_i(x) = \sum_{t=1}^{T} v_{i,t} x_{i,t}$.

In this setting, an online algorithm is a mapping from history of previous rounds and the agents' values on the current round $\{v_{i,t}\}_{i \in N}$ to an allocation $x_t = \{x_{i,t}\}_{i \in N}$ to be made on this round.

We evaluate the quality of the final allocation x of an online algorithm using two widely studied objectives: the egalitarian social welfare and the Nash social welfare. Given allocation x, the egalitarian social welfare of this allocation is equal to the minimum utility across all the agents for this allocation

$$ESW(x) = \min_i \{u_i(x)\}.$$

This objective captures the intuition that a fair allocation may be one that makes the least happy agent as happy as possible.

[1] In fair allocation literature, normalization of total values can often be assumed without loss of generality. However, in online setting it becomes a meaningful assumption about information available to the agents. See Section 4.6 for more discussion of this.

The second objective we consider is Nash social welfare (NSW), which is defined as the geometric mean of the agents' utilities

$$NSW(x) = \left(\prod_i u_i(x)\right)^{1/n}.$$

Similarly to the egalitarian social welfare, maximizing the Nash social welfare objective is well-aligned with the goal of ensuring a fair allocation. But, rather than looking to maximize the minimum utility irrespective of how much this costs to the other agents, the Nash social welfare does so only if the relative gains in utility outweigh the losses (for more information see [22]).

Following the standard approach from the online algorithms literature, we evaluate the performance of online algorithms using their *competitive ratio* with respect to each of these objectives. Let $\tilde{x}(v)$ denote the allocation that the algorithm outputs on instance v and let $x^{ESW}(v)$ be the optimal allocation for this instance with respect to the egalitarian social welfare objective. Then the competitive ratio of this algorithm is defined as

$$\max_v \frac{ESW(x^{ESW}(v))}{ESW(\hat{x}(v))}.$$

Similarly, if $x^{NSW}(v)$ denotes the optimal allocation for instance v with respect to the Nash social welfare objective, then the corresponding competitive ratio is

$$\max_v \frac{NSW(x^{NSW}(v))}{NSW(\hat{x}(v))}.$$

Our main results in this paper take the form of lower or upper bounds regarding the competitive ratio that is achievable for the two aforementioned objectives. Lower bounds indicate a competitive ratio guarantee that no online algorithm can achieve, and upper bounds provide a specific algorithm and prove that its competitive ratio is always at least as good as that upper bound.

To express these bounds as a function of the instances at hand, we consider different parameters of the instance, such as the number of agents, n, or the number of rounds, T. Note that the online problem becomes trivial in the extreme cases when $n = 1$ (since we can just allocate all the items to the single agent) or $T = 1$ (since this would correspond to a static fair division problem), but the problem becomes harder as n and T grow. Our goal is to understand how the competitive ratio that we can achieve degrades with these parameters.

4.3 Lower bound for the egalitarian social welfare

In this section we consider the problem of designing online algorithms with a good competitive ratio with respect to the egalitarian social welfare objective. Our main result here is a lower bound, showing that a practically useful competitive ratio is impossible to achieve for this objective. The main idea behind this proof is one that also provides some intuition for the much more demanding lower bound construction of Section 4.5. In particular, when facing an instance where the agents appear similar during the first rounds, an online algorithm cannot do better than treating them in a symmetric fashion. However, some of these agents may later on have a high value during a round with low competition, and are therefore easy to satisfy, while others may only value items during rounds with high competition, which makes them hard to satisfy. An algorithm with access to future valuations would therefore prioritize the latter agents during the first rounds, making the optimal solution significantly better than what any online algorithm can hope to achieve.

Theorem 4.3.1. *For the objective of egalitarian social welfare, there is no online algorithm that can guarantee a competitive ratio better than* $\frac{\sqrt{n}}{3}$*, nor is there an* *online algorithm that can guarantee a competitive ratio better than* $\frac{\sqrt{T}}{3}$.

Proof. To prove the desired bounds on the competitive ratio of any given online algorithm, we construct an instance adaptively depending on the decisions of the algorithm. Without loss of generality, we assume that n is even and that $\sqrt{n}$ is an integer. In the first round of this instance, the valuation of every agent i is equal to $v_{i1} = 1 - \frac{1}{\sqrt{n}}$. Irrespective of how the online algorithm allocates the item of the first round, the number of agents that receive at least a $2/n$ fraction of the good in this round is at most $n/2$ (since these fractions add up to at most 1), which implies that at least $n/2$ agents receive less than a $2/n$ fraction. Let A be a subset of any $\sqrt{n}$ of these agents, who receive less than $2/n$ of the item and hence a value less than $2/n$ from it.

In each of the next $n - \sqrt{n}$ rounds, the item is valued only by a distinct agent not in A for a value of $\frac{1}{\sqrt{n}}$ (i.e., all of the remaining value of that agent), and all other agents have a value of zero for it. Finally, in the last round all the agents in A have a value of $1/\sqrt{n}$ for the same item.

One feasible allocation for this instance would be to allocate the items of the first and last round equally among the $\sqrt{n}$ agents in A, providing them with a value of $\frac{1}{\sqrt{n}}$ each, and then allocating each of the remaining items fully to the distinct agent that values them, providing the remaining agents with a value of $\frac{1}{\sqrt{n}}$ as well. Therefore, the optimal egalitarian social welfare is at least $\frac{1}{\sqrt{n}}$.

On the other hand, since the outcome of the online algorithm allocated no more than a $2/n$ fraction of the first item to each of the agents in A, the average value these agents receive from the first and the last item is no more than $\frac{2}{n}\left(1 - \frac{1}{\sqrt{n}}\right) + \frac{1}{\sqrt{n}}\frac{1}{\sqrt{n}} \leq \frac{3}{n}$. Since the minimum value among the agents in A is less or equal to their average value, we conclude that the competitive ratio of the algorithm is at least $(1/\sqrt{n})/(3/n) = \sqrt{n}/3$.

Since the family of instances described above has a total number of rounds $T = n - \sqrt{n} + 2$, this directly implies that no online algorithm can guarantee a competitive ratio better than $\sqrt{T}/3$ for all values of T. □

In contrast to the negative result of this section for the egalitarian social welfare, the next section considers the Nash social welfare objective and provides an online algorithm that guarantees a logarithmic competitive ratio with respect to both n and T. Section 4.5 then shows that this is the best competitive ratio that any online algorithm can achieve.

4.4 Set-Aside Greedy algorithm for maximizing Nash Social Welfare

Before describing our proposed online algorithm for maximizing the Nash social welfare objective, it is instructive to consider some simpler approaches that fail to achieve a good competitive ratio. It is easy to verify that uniformly allocating every item, i.e., returning $x_{i,t} = 1/n$ for every agent i and every round t, performs quite poorly. For example, consider an instance with $T = n$ rounds where $v_{i,t} = 1$ if $i = t$ and $v_{i,t} = 0$ otherwise. The optimal Nash social welfare for this instance would be equal to 1 by, of course, allocating the item of each round t to the agent $i = t$, but the Nash social welfare of the uniform allocation is equal to $1/n$ for every instance, leading to a competitive ratio of n (and a competitive ratio of T if we use the number of rounds to parameterize the bound).

A natural alternative to the uniform allocation would be a proportional allocation, according to which in each round t the allocation of each agent i is proportional to her value, i.e., $x_{i,t} = \frac{v_{i,t}}{\sum_j v_{j,t}}$. As we show in Appendix 4.7.2, the proportional allocation Pareto dominates the uniform allocation and thus guarantees that every agent receives a utility of at least $1/n$. Nevertheless, a proportional allocation algorithm would also fail to achieve an appealing competitive ratio. We use the following example to illustrate this fact.

Example 4.4.1 (Proportional allocation fails to achieve high NSW)**.** *Consider an instance with $T = n$ rounds such that for each round t agent $i = t$ has a value of $v_{i,t} = 1/\sqrt{n}$, and every other agent has a value of $v_{i,t} = (1 - 1/\sqrt{n})/(n - 1)$. In each round t, the proportional allocation assigns to agent $i = t$ a fraction equal to*

$$\frac{1/\sqrt{n}}{(n-1)(1-1/\sqrt{n})/(n-1) + 1/\sqrt{n}} = \frac{1}{\sqrt{n}},$$

and it splits the remaining $1 - 1/\sqrt{n}$ portion of the item equally among the other $n - 1$ agents. As a result, the final utility of every agent i is equal to

$$\frac{1}{\sqrt{n}} \cdot \frac{1}{\sqrt{n}} + (n-1)\frac{1-1/\sqrt{n}}{n-1} \cdot \frac{1-1/\sqrt{n}}{n-1} \leq \frac{3}{n}.$$

Since all the agents receive the same utility, the Nash social welfare is also going to be equal to that utility. On the other hand, if the item of each round t was fully allocated to agent $i = t$, then the Nash social welfare would be equal to $1/\sqrt{n}$, which implies that the competitive ratio would be no better than $\sqrt{n}/3$ (or $\sqrt{T}/3$ if we use the number of rounds to parameterize the bound).

Finally, another natural approach to consider is the greedy algorithm. i.e. allocating the item on round t as to maximize the Nash social welfare at the end of the round, essentially assuming that no additional items will arrive in the future. As we show in Appendix 4.7.1, this greedy approach also fails to achieve a good competitive ratio. Specifically, this approach gives higher priority to agents whose utility up to the given round is low, since these are the agents that can yield the highest multiplicative increases, which in some instances can lead to the greedy algorithm allocating the item to agents with vanishingly small values while ignoring agents with substantially larger values.

4.4.1 The Set-Aside Greedy Algorithm

After this short introduction to the obstacles that natural online algorithms face in this context, we are ready to present the main result of this section, which is a polynomial time algorithm that achieves a logarithmic competitive ratio with respect to both n and T.

We begin by giving an informal description of the algorithm. The Set-Aside Greedy algorithm divides every item in half, and then uses a different strategy for allocating each half. One half of each item is set aside and distributed uniformly among the agents, so each agent receives a $1/(2n)$ fraction of every item. Since the values are normalized so that $\sum_t v_{i,t} = 1$, this ensures that after the completion of the algorithm, every agent will have received a utility of at least $1/(2n)$ through this half. The other half of each item is allocated greedily using this foresight: the algorithm pretends that every agent has *already* received a utility of $1/(2n)$ in advance, and this informs the greedy algorithm's allocation. As a result, since every agent appears to begin with a non-trivial amount of value (which in reality has not been distributed yet), the greedy allocation circumvents the issues it suffers from when applied by itself (as we discuss in previous subsection).

To define the algorithm formally, we first introduce some additional notation.

We let $y_{i,t}$ and $z_{i,t}$ to be the semi-allocations from first and second half of the item, i.e $x_{i,t} = y_{i,t} + z_{i,t}$, where $\sum_i y_{i,t} = \sum_i z_{i,t} = 1/2$.

Definition 4.4.2 (Promised utility)**.** *Given some round t, the semi-allocations $\tilde{z}_{i,t'}$ made on previous rounds $t' < t$, and the allocation $z_{i,t}$ made on the current round, the* promised utility *of agent i at round t is defined as*

$$\tilde{u}_{i,t}(\tilde{z}_1, \tilde{z}_2, \ldots, \tilde{z}_{t-1}, z_t) \;=\; \tilde{u}_{i,t}(z_t) \;=\; \frac{1}{2n} + z_{i,t}v_{i,t} + \sum_{t'=1}^{t-1} \left(\tilde{z}_{it'} v_{it'}\right). \tag{4.1}$$

Although the promised utility of an agent in round t depends on all the previous allocations $(\tilde{z}_{t'})_{t'<t}$, for notational simplicity throughout the paper we denote it by $\tilde{u}_{i,t}(z_t)$, dropping the dependence on previous allocations, which is implicit.

The reason we call $\tilde{u}_{i,t}(z_t)$ promised utility is that, as long as the first half of the item is always distributed uniformly $\tilde{y}_{i,t} = \frac{1}{2n}$, the final utility of an agent can be lower bounded with the promised utility $u_i(\tilde{x}) \geq \tilde{u}_{i,t}(z_t)$. After the final allocation is made, these two quantities coincide $u_i(\tilde{x}) = u_{i,T}(\tilde{z}_T)$, thus delivering on the promise.

We can now give the formal definition of the Set-Aside Greedy algorithm, which is provided as Algorithm 6 below. The allocation $\hat{z}_t$ on round t is computed by maximizing the Nash social welfare with respect to the *promised* utilities , $\prod_i \tilde{u}_{i,t}(z_t)^{1/n}$. This is equivalent to maximizing $\sum_i \log(\tilde{u}_{i,t}(z_t))$, so the solution can be computed by solving a convex program, but in Appendix 4.7.3 we also argue that there is an even more efficient way of doing this step, leading to an overall running time of $O(Tn)$.

ALGORITHM 6: Set-Aside Greedy algorithm

```
for all t = 1 to T do
    For each agent i, let ŷ_{i,t} = 1/(2n)
    Compute ẑ_t = arg max_{z_t} {∑ log(ũ_{i,t}(z_t))}   subject to   ∑_i z_{i,t} ≤ 1/2
    and   z_{i,t} ≥ 0
    Allocate x̂_{i,t} = ŷ_{i,t} + ẑ_{i,t}
end for
```

An interesting fact is that our algorithm does not need to know what the value of T is (however, the lower bound we provide in the next section holds even for algorithms that have this information).

We also note that one can modify the algorithm so that the half of each item that we set aside is actually allocated in a proportional manner instead of uniform by setting $\hat{y}_{i,t} = \frac{v_{i,t}}{2\sum_j v_{j,t}}$. In Appendix 4.7.2 we show that this approach would also guarantee that every agent i will get the utility of at least $\frac{1}{2n}$ and it actually weakly Pareto dominates the uniform allocation on all instances. This would arguably lead to a more practical algorithm without affecting any of our theoretical results. However, since we do not take advantage of such improved performance in our analysis, we assume that the set-aside half is allocated uniformly, for simplicity.

We are now ready to state the main result of this section. For a given instance v, we let $k(v) = T \max_{i,t}\{v_{i,t}\}$ (denoted by just k when the dependence on v is clear), i.e the ratio between the maximum and the average value across all agents. Since $v_{i,t} \leq 1$, we always have $k(v) \leq T$.

Theorem 4.4.3. *The competitive ratio of the Set-Aside Greedy algorithm is at* *most*

$$2\ln(\min\{k(v), n\}) + 2\ln(2) + 1.$$

In particular, for any instance v, the competitive ratio is at most

$$2\ln(\min\{T, n\}) + 2\ln(2) + 1.$$

One intuition we capture with our main result is that, when the values of agents are not much larger than their average value (across rounds), the performance of the online solution should improve. To see this, note that at the extreme, when values are equal across all rounds, allocation problem becomes trivial (as the optimal solution is to distribute all items equally).

The rest of the section is devoted to proving the result above. We start by describing a duality-based approach for constructing an upper bound on the competitive ratio of a given allocation. We then leverage this technique to prove the theorem above.

4.4.2 A Duality-Based Upper Bound for the Competitive Ratio

In proving the main result of this section, we employ a dual optimization problem that allows us to construct a certificate of approximation during the execution of the online algorithm. We also show that the Set-Aside Greedy algorithm can be interpreted as greedily minimizing *promised prices*, which are lower bounds on the dual solution.

Informally, the following lemma states that, after an allocation is made a vector of “prices” can be easily constructed and the sum of prices can be related to the sub-optimality of NSW under the chosen allocation.

Lemma 4.4.4. *Given an allocation $\tilde{x}$, if for every round there exists a number $p_t \geq 0$ such that $p_t \geq \frac{v_{i,t}}{u_i(\tilde{x})}$ for every every agent i, then*

$$NSW(x^*) \ \leq \ \frac{\sum_t^T p_t}{n} NSW(\tilde{x}),$$

where x^* is the Nash social welfare maximizing allocation.

Proof. Consider the value profile v and the allocation $\tilde{x}$, and let p be a vector of prices that satisfy the constraints in the statement of the lemma. Observe that it then is a feasible solution to the following ”dual” program:

$$\begin{aligned} \min_p \ & \sum_{t=1}^T p_t \\ \text{s.t. } & p_t \geq \frac{v_{i,t}}{u_i(\tilde{x})} \quad \text{for all } i, t \\ & p_t \geq 0 \end{aligned}$$

The ”primal” for this program is

$$\begin{aligned} \max_x \ & \sum_{i=1}^n \frac{u_i(x)}{u_i(\tilde{x})} \\ \text{s.t. } & \sum_i x_{i,t} \leq 1 \quad \text{for all } t \\ & x \geq 0 \end{aligned}$$

Here $u_i(x) = \sum_{t=1}^T x_{i,t} v_{i,t}$. We stress that only x (but not $\tilde{x}$) is a variable in this program.

Because a feasible solution for dual p gives an upper bound for the value of the primal program, we have for any feasible allocation x.

$$\frac{1}{n}\sum_{i=1}^{n}\frac{u_i(x)}{u_i(\tilde{x})} \leq \frac{\sum_t p_t}{n}$$

In particular, this implies that this holds for $x = x^*$, the optimal NSW solution. Then, since the arithmetic mean is larger than the geometric mean, we have

$$\left[\frac{u_i(x^*)}{u_i(\tilde{x})}\right]^{1/n} \leq \frac{\sum_t p_t}{n}$$

Thus we have shown that $\tilde{x}$ achieves the promised NSW approximation. □

Given this dual view of the problem, it is useful to re-interpret the described in previous section algorithm in dual terms.

Suppose, in light of the lemma above, one wants to achieve good approximation by choosing allocation x that minimizes prices p_t. The problem is, the right side of the inequality $p_t \geq \frac{v_{i,t}}{u_i(x)}$, namely total utility $u_i(x)$, depends on future values and allocations, which are unavailable to an online algorithm.

A natural fix to this problem is to construct a lower bound for $u_i(x)$, and a corresponding upper bound for the feasible price p_t. This is exactly what the earlier defined promised utility (4.1) does, because, as noted earlier, promised utility lower-bounds the final utility (if we commit to allocating half of the item uniformly $y_{i,t} = \frac{1}{2n}$).

Let's call the corresponding upper bound for the price $\tilde{p}_t(z_t) = \max_i \frac{v_{i,t}}{\tilde{u}_{i,t}(z_{i,t})}$ promised price. Since it does not depend on future rounds, one can choose semi-allocations z_t in order to minimize the promised price $\tilde{p}_t(z_t)$ on this round. In fact, doing so leads to exactly the same allocation as Set-Aside Greedy, as we show with the following lemma. [2]

[2]This is in fact the form of the algorithm that the authors discovered first

Lemma 4.4.5. *For the allocation $\hat{z}_t$ chosen by Set-Aside Greedy on round t, if $\hat{z}_{i,t} \neq 0$ for some agent i, then $\frac{v_{i,t}}{\tilde{u}_{i,t}(\hat{z}_{i,t})} = \max_{i'} \frac{v_{i',t}}{\tilde{u}_{i',t}(\hat{z}_{i,t})}$, i.e., $\frac{v_{i,t}}{\tilde{u}_{i,t}(\hat{z}_{i,t})} = \tilde{p}_t(\hat{z}_{i,t})$.*

Proof. We prove the lemma by reducing a step of online allocation to a static Fisher market and showing that the statement of the lemma is equivalent to the KKT conditions for this market.

Consider the following Fisher market: there are n agents and $n+1$ items. Item i is only valued by agent i at $v_{i,i} = \tilde{u}_{i,t}(0)$ (i.e. their promised utility excluding current round allocation). Agents value the last item at $v_{i,n+1} = v_{i,t}/2$. Then utility of agent i from a static allocation x is $x_{i,i} u_{i,t}(0) + \frac{x_{i,n+1} v_{i,t}}{2}$.

Clearly, in this Fisher market, it is optimal to allocate items 1 to n to the only agent who wants it ($x^*_{i,i} = 1$), and thus the problem of maximizing Nash Welfare for this market becomes

$$\max \sum_i \log\left(u_{i,t}(0) + \frac{x_{i,n+1} v_{i,t}}{2}\right)$$
$$s.t. \sum_i x_{i,n+1} \leq 1$$
$$x_{i,j} \geq 0$$

Changing variables $z_i = \frac{x_{i,n+1}}{2}$ leads to the optimization problem identical to the one in Set-Aside Greedy. This implies the equivalence between the online solution $\hat{z}_t$ and the optimal solution x^* for the constructed Fisher market: $\hat{z}_{i,t} = x^*_{i,n+1}/2$.

On the other hand, KKT conditions for this static Fisher market give ([68, Chapter 5]):

$$x^*_{i,n+1} \neq 0 \Rightarrow \frac{v_{i,t}}{\tilde{u}_{i,t}(0) + \frac{v_{i,t} x^*_{i,n+1}}{2}} = \max_{i'} \frac{v_{i',t}}{\tilde{u}_{i',t}(0) + \frac{v_{i',t} x^*_{i',n+1}}{2}}$$

Now, substituting $\hat{z}_{i,t} = \frac{x^*_{i,n+1}}{2}$ yields the statement of the lemma. □

4.4.3 Logarithmic Upper Bound for the Competitive Ratio of Set-Aside Greedy

Lemma 4.4.4 provides us with a way of proving an upper bound on the competitive ratio of our algorithm: our goal is to find a set of p_t values for each round t such that $p_t \geq \frac{v_{i,t}}{u_i(\hat{x})}$ for every agent i, where $\hat{x}$ is the allocation that our algorithm outputs and $\sum_t^T p_t$ is as small as possible. We will achieve this goal by using the promised prices of each round t:

$$\tilde{p}_t(\hat{z}_t) = max_i \frac{v_{i,t}}{\tilde{u}_{i,t}(\hat{z}_t)}.$$

It is easy to verify that these prices satisfy the condition of Lemma 4.4.4, since $\tilde{p}_t(\hat{z}_t) = max_i \frac{v_{i,t}}{\tilde{u}_{i,t}(\hat{z}_t)} \geq \frac{v_{i,t}}{u_i(\hat{x})}$, since for every agent i the utility $u_i(\hat{x})$ from the allocation at the end of the algorithm is at least as high as the promised utility $u_{i,t}(\hat{z}_t)$ at round t. What remains to be shown is that this set of prices also provide an appealing bound, i.e., that the $\sum_t^T \tilde{p}_t(\hat{z}_t)$ is small. In order to provide an upper bound for the promised prices we first need to define the following function.

Definition 4.4.6. *For each round t of the Set-Aside Greedy algorithm and each possible semi-allocation z_t of that round's item, we let*

$$U_t(z_t) = \sum_i \ln \left(\min \left\{ \frac{k}{n}, \; \tilde{u}_{i,t}(z_t) \right\} \right)$$

The following lemma provides an upper bound for the promised price on round t in terms of the change in the $U_t(\hat{z}_t)$ function.

Lemma 4.4.7. *Under Set-Aside Greedy, at every round t the following inequality holds*

$$\tilde{p}_t(\hat{z}_t) \leq 2(U_t(\hat{z}_t) - U_{t-1}(\hat{z}_{t-1})) + \frac{n}{T}$$

Proof. We first note that since $\tilde{u}_{i,t}(z_t)$ is weakly increasing in t, for every $t > 1$ we have $U_t(\hat{z}_t) - U_{t-1}(\hat{z}_{t-1}) \geq 0$. As a result, the lemma clearly holds for every round t when $\tilde{p}_t(\hat{z}_t) \leq \frac{n}{T}$. The rest of the proof shows that the lemma is also true for rounds t when $\tilde{p}_t(\hat{z}_t) > \frac{n}{T}$, i.e., $\max_i \frac{v_{i,t}}{\tilde{u}_{i,t}(\hat{z}_t)} > \frac{n}{T}$. This implies that for every agent i we have $\tilde{u}_{i,t}(\hat{z}_t) < \frac{Tv_{i,t}}{n} \leq \frac{k}{n}$, and thus, according to Definition 4.4.6 of $U_t(z_t)$, we get

$$\begin{aligned}
U_t(\hat{z}_t) - U_{t-1}(\hat{z}_{i,t-1}) &= \sum_i [\ln \tilde{u}_{i,t}(\hat{z}_{i,t}) - \ln \tilde{u}_{i,t-1}(\hat{z}_{i,t-1})] \\
&= \sum_i [\ln \tilde{u}_{i,t}(\hat{z}_{i,t}) - \ln \tilde{u}_{i,t}(0)] && (4.2) \\
&= \sum_i \int_{z_{i,t}=0}^{\hat{z}_{i,t}} \frac{d \ln \tilde{u}_{i,t}(z_{i,t})}{dz_{i,t}} dz_{i,t} \\
&= \sum_i \int_{z_{i,t}=0}^{\hat{z}_{i,t}} \frac{1}{\tilde{u}_{i,t}(z_{i,t})} \frac{d\tilde{u}_{i,t}(z_{i,t})}{dz_{i,t}} dz_{i,t} \\
&= \sum_i \int_{z_{i,t}=0}^{\hat{z}_{i,t}} \frac{v_{i,t}}{\tilde{u}_{i,t}(z_{i,t})} dz_{i,t} && (4.3) \\
&\geq \sum_i \hat{z}_{i,t} \frac{v_{i,t}}{\tilde{u}_{i,t}(\hat{z}_{i,t})} && (4.4) \\
&= \sum_i \hat{z}_{i,t} \tilde{p}_t(\hat{z}_t) && (4.5) \\
&= \frac{\tilde{p}_t(\hat{z}_t)}{2} && (4.6)
\end{aligned}$$

To get (4.2), we use the fact that for every agent i the utility $\tilde{u}_{i,t}(0)$ in the beginning of round t is equal to the utility $\tilde{u}_{i,t-1}(\hat{z}_{i,t-1})$ at the end of round $t-1$. We then reach (4.3) by taking the derivative of $\ln \tilde{u}_{i,t}(\hat{z}_{i,t})$ with respect to $z_{i,t}$, and

Inequality (4.4) uses the fact that $\tilde{u}_{i,t}(z_{i,t}) \le \tilde{u}_{i,t}(\hat{z}_{i,t})$ for all $z_{i,t} \le \hat{z}_{i,t}$. Equation (4.5) uses the fact that $\hat{z}_{i,t} \neq 0$ implies $\frac{v_{i,t}}{\tilde{u}_{i,t}(\hat{z}_{i,t})} = \max_{i'} \frac{v_{i',t}}{\tilde{u}_{i',t}(\hat{z}_{i,t})} = \tilde{p}_t(\hat{z}_t)$, according to Lemma 4.4.5. Finally, (4.6) uses the fact that $\sum_i \hat{z}_{i,t} = 1/2$ as specified in the description of Set-Aside Greedy. Therefore, since $n/T \ge 0$ this concludes the proof of the lemma. □

Using this lemma, we can now show that the sum of promised prices under the allocation of Set-Aside Greedy algorithm is indeed sufficiently small. The main result of this section, Theorem 4.4.3 is then a simple corollary of this bound.

Lemma 4.4.8. *For the allocation $\hat{z}_t$ made by Set-Aside Greedy algorithm, the following inequality holds for any instance*

$$\frac{\sum_{t=1}^{T} \tilde{p}_t(\hat{z}_t)}{n} \le 2\ln(k) + 2\ln(2) + 1,$$

Corollary 4.4.9. *The competitive ratio of the Set-Aside Greedy algorithm is at most*

$$2\ln(\min\{k(v), n\}) + 2\ln(2) + 1.$$

The corollary is just an application of the Lemma 4.4.4, as the promised prices $\tilde{p}_t(\hat{z}_t)$ satisfy the inequalities in the statement of the lemma (this in turn is because promised utility is always a lower bound for the total ex-post utility $\tilde{u}_{i,t}(\hat{z}_t) \le u_i(\hat{x})$).

Proof of Lemma 4.4.8. Using Lemma 4.4.7, we get

$$\sum_{t=1}^{T} \tilde{p}_t(\hat{z}_t) \le \sum_{t=1}^{T} \left(2(U_t(\hat{z}_t) - U_{t-1}(\hat{z}_{t-1})) + \frac{n}{T}\right) = 2(U_T(\hat{z}_T) - U_1(\hat{z}_1)) + n. \quad (4.7)$$

Note that, since for every round t and allocation z_t we have $\tilde{u}_{i,t}(\hat{z}_t) \le 1$, this implies that $U_t(\hat{z}_t) \le 0$. Also, note that since $k \ge 1$ and $\tilde{u}_{i,t}(\hat{z}_t) \ge 1/(2n)$ for every agent i and round t, we get that $U_t(\hat{z}_t) \geq n \ln(1/(2n))$. Therefore, in any round t we have $U_t(\hat{z}_t) \in [n \ln(1/(2n)), 0]$ which, combined with (4.7) above yields the desired bound with respect to n:

$$\begin{aligned}
\frac{\sum_{t=1}^{T} \tilde{p}_t(\hat{z}_t)}{n} &\leq \frac{2}{n}(U_T(\hat{z}_T) - U_0(\hat{z}_0)) + 1 \\
&\leq \frac{2}{n}(0 - n \ln(1/(2n))) + 1 \\
&= 2 \ln(n) + 2 \ln(2) + 1,
\end{aligned}$$

In fact, for instances with $k(v) < n$ we get a better upper bound of $U_t(\hat{z}_t) \leq n \ln(k/n)$, which leads to a stronger bound, parameterized by k:

$$\begin{aligned}
\frac{\sum_{t=1}^{T} \tilde{p}_t(\hat{z}_t)}{n} &\leq \frac{2}{n}(n \ln(k/n) - n \ln(1/(2n))) + 1 \\
&= 2 \ln(k) + 2 \ln(2) + 1,
\end{aligned}$$

□

Before concluding this section, we note that the choice of allocating half an item uniformly (rather than some other constant fraction) is somewhat arbitrary: however, the optimal choice for this fraction cannot be derived in closed form and does not improve the approximation factor asymptotically.

4.5 Lower bound for Nash Social Welfare

We now prove the following result, which complements the positive result of the previous section.

Theorem 4.5.1. *There exists no online algorithm that can achieve a competitive ratio of $O(\log^{1-\epsilon} n)$ or $O(\log^{1-\epsilon} T)$ with respect to the Nash social welfare.*

The main idea behind the construction that we provide for this theorem is similar in spirit to the one use in the proof of Theorem 4.3.1. First, at some time t we ask the algorithm allocate between agents with identical values. In later rounds, some agents are able to obtain large values without competition (i.e., there are rounds when only one agent has a non-zero value). Other agents, however, clash for a single item on some future round $t' > t$, and are thus unable to get much value from future rounds. This scheme exploits the fact that an online algorithm is unable to predict which agents are going to clash and which aren't.

4.5.1 Lower Bound Construction

We now give a formal description of the instance. Let's choose an integer $M = \log^{(1-\epsilon)} n$ (throughout the proof we assume that for large enough n and rational ϵ all choices of integer parameters are in fact integer), and also a constant (i.e., independent of n) integer L. Let $w^{m\ell}$ for $m \leq M, \ell \leq L$ be numbers s.t. $w^{m\ell} = \frac{1}{\log^{(L-\ell+1)(M-m+1)} n}$. Note that for any constant L, $w^{m\ell} = \Omega(1/n) = o(1/M)$, and also that for any fixed ℓ we have $\frac{w^{m\ell}}{w^{m+1,\ell}} = \frac{1}{\log n} = o(1/M)$.

Let n be large enough such that $M = \log^{(1-\epsilon)} n$ is an integer (throughout the proof we assume that for large enough n and rational ϵ all choices of integer parameters are in fact integer). Also, let L be an integer constant to be chosen later on. For each pair of natural numbers m, ℓ such that $m \leq M$ and $\ell \leq L$, we let $w^{m\ell} = \frac{1}{\log^{(L-\ell+1)(M-m+1)} n}$. Note that for any constants L, $w^{m\ell} = \Omega(1/n)$ and $w^{m\ell} = o(1/M)$. Also, for any fixed ℓ we have $\frac{w^{m\ell}}{w^{m+1,\ell}} = \frac{1}{\log n} = o(1/M)$.

To make our description more readable, it is useful to split up rounds in our instance into a hierarchy or repeatable cycles, we reserve the words month, year and era for this purpose.

There are 3 eras: Era of Banishment, Era of Plenty, and Era of Collapse.

Era of Banishment. This era lasts $n(1 - 1/2^L)$ rounds, and it consists of L years (indexed by ℓ) each lasting $n/2^\ell$ rounds. Years consists of months (indexed by m), every month lasts M rounds (thus, different years have different number of months in them, namely $\frac{n}{2^\ell M}$). Every year of the Era of Banishment agents are split into "banished" and "cleared", by banishing M agents and clearing M agents per month. Once banished, an agent never has non-zero valuation again until very last round (the Era of Collapse). The cleared agents lose their status at the end of the year (and can become cleared or banished in the next year).

At first round of month m of year ℓ there are M agents who are neither banished nor cleared who have value w^{1l} (everyone else has value 0). After the algorithm $\mathcal{A}$ makes an allocation x_t, we find an agent for whom $x_{i,t} \leq 1/M$ (such agent always exists by pigeonhole principle) and this agent becomes banished.

On round $z \in [2, M]$ of this month there are $M - 1$ agents from previous round plus one new agent (who isn't banished and hasn't been cleared this year) with values $w^{z\ell}$ (remaining agents, including the banished ones, have value 0). After the allocation is made, we repeat the banishing procedure described above. By the end of the month, we have banished M agents, and we declare all agents who had non-zero value this month but weren't banished (there are $M - 1$ such agents) to be cleared. To make the accounting easier, we additionally clear 1 more agent who hasn't been cleared or banished before in this year (thus equalizing number of banished and cleared agents in this month). Also note that all agents can have non-zero value during only one month in a year, since during this month they have to become either banished or cleared.

The number of months in the year is $\frac{n}{2^\ell M}$, as after this many months every agent is either banished or cleared. At the end of the year, the cleared agents lose their status (but banished agents remain such), and the next year begins.

Thus, every year of this era we banish half the agents that were not banished, and by the end of the era (which lasts L years) we banish $n(1 - 1/2^L)$ agents.

We also note that very little value is seen by the agents during this Era (compared to their total allowed value of $\sum_t v_{i,t} = 1$). Since during the Era of Banishment each agent participates in at most one month a year, the total value seen by an agent during the era of Banishment can be bound by $\sum_{m,\ell} w^{m\ell} \leq MLw^{ML} = \frac{L\log^{1-\epsilon} n}{\log n} = o(1)$ (since the number of years L is a constant).

Era of Plenty. This era begins at $t = n(1-1/2^L)+1$ and lasts $n/2^L$ rounds. At every round there is a single non-banished agent with the remainder of their value $v_{i,t} = 1 - \sum_{t'=1}^{n(1-1/2^L)} v_{i,t'} = 1 - o(1)$ (see the paragraph above for the justification of the last equality). All other agents have zero value on this round).

Era of Collapse. This era consists of just one round, $t = T = n + 1$. At this round, all banished agents have the remainder of their value $v_{i,t} = 1 - \sum_{t'=1}^{n} v_{i,t'} = 1 - o(1)$.

4.5.2 Proof of Theorem 4.5.1

Assume that there exists an online algorithm $\mathcal{A}$ that achieves an competitive ratio of $\log^{1-\epsilon} T$ or $\log^{1-\epsilon} n$ for some constant $\epsilon \geq 0$. We prove that this leads to a contradiction by showing that $\mathcal{A}$ would fail to satisfy these guarantees for the instance described above.

We first provide a lower bound for the optimal Nash social welfare in this instance. During the Era of Banishment we can give the item to the agent who is about to become banished, during the Era of Plenty we give the item to the only agent with non-zero value, during the Era of Collapse we split the item equally between the agents.

In this case, utility of banished agents u_i^* can be lower bounded with the value w^{ij} on the round when they got banished, and utility of non-banished agents is $1 - o(1)$, because $\sum_{m \leq M, \ell \leq L} w^{m\ell} = o(1)$ according to our definition of $w^{m\ell}$ above.

We now compute utilities under the allocation made by the online algorithm. Every banished agent got $1/M$ fraction of the item on the round that they got banished, and their values on previous rounds is negligible $v_{i,t'} = o(v_{i,t}/M)$, since $\frac{w^{m-1,\ell}}{w^{m\ell}} = o(1/M)$.

Finally, we argue that the last round does not affect the utilities of banished agents much. Intuitively, this is because there is only one item and $O(n)$ banished agents who have approximately equal value for it. We provide a more careful argument below.

We upper bound optimal utilities on the last round by assuming banished agents got $1/M$ of the item on which they got banished (by construction, maximum possible amount), and assuming that the algorithm makes optimal NSW allocation on the last round.
The optimal NSW solution, when choosing between agents with equal values, allocates the marginal fraction of the item to the agent with lowest utility. Note that, on the last round, there are $n/2M$ agents with value w^{11} (the smallest value

of w^{lm}). Thus, if the item is distributed between them, each agent gets the utility of $\frac{2M}{n} = \frac{2\log^{1-\epsilon}}{n} = o(w^{11}/M)$, thus not improving their utility by even a constant factor (and so it will distribute the entire item between the agents with w^{11}) .

Thus, overall, utility of banished agents can be upper-bounded by $u_i \leq \frac{u_i^*}{M} + o(\frac{u_i^*}{M})$. We bound utility of non-banished agents with $u_i \leq 1$.

Now we can bound the approximation factor achieved by the online algorithm:

$$\frac{NSW(x)}{NSW(x^*)} \leq \left[\frac{u_i}{u_i^*}\right]^{1/n} \leq \left[\left(\frac{1}{M}\right)^{n(1-1/2^L)} \left(\frac{1}{1-o(1)}\right)^{n/2^L}\right]^{1/n} \approx \log^{1-\epsilon'} n.$$

Now, we can pick the constant L (and appropriately large n) to make sure $\epsilon' \leq \delta$.

4.6 Discussion

We now conclude the paper with a brief discussion regarding the our model and possible extensions of our results.

Allocating multiple items per round. It is worth noting that all our results readily generalize to a setting where multiple items can arrive on a single round. The reduction is to split a single round with multiple items into multiple "imaginary" rounds with a single item in each. Under this reduction, our approximation guarantees stay intact, with number of rounds T replaced by the total number of items.

Allocating indivisible items using randomness. Also, note that even if the items are actually *indivisible*, i.e., each item can be allocated to at most one agent, then both our upper bounds and our lower bounds extend to this setting as well by considering randomized algorithms. In this case, we can evaluate the agents' *expected* utility based on the probability that they receive each item, and linearity of expectations reduces this setting to the one studied in this paper.

Normalization of valuations. Throughout the paper we assume agent valuations are normalized so that $\sum_{t=1}^{T} v_{i,t} = 1$. This normalization is the standard assumption in all of the cake-cutting literature, as well as in other work on the offline fair division of goods, where the values are used to indicate *relative* preferences of the agents among the goods rather than absolute values. For example, if we have three items and an agent values the second one three times as much as the first and the third one six times as much as the first, the valuation vectors $[1, 3, 6]$ and $[0.1, 0.3, 0.6]$ both convey the same information. Similarly, in our setting each agent i's $v_{i,t}$ value in round t effectively compares her value for the current item relative to her (anticipated) average value in the long run. Thus, in the online setting, this normalization can be interpreted as the assumption that agents have some information about the future that allows them to compare their current value to the long-term average. In fact, it is easy to see that without any information regarding the possible future values of the agents it is impossible to achieve any non-trivial competitive ratio (see Appendix 4.7.1 for a formal argument). In many real world applications agents implicitly express these normalized preferences through some form of tokens, or a budget of artificial currency, that they need to spend within a predetermined time horizon. This artificial budget essentially asks that the agents evaluate how much they need to use a resource today (e.g., using a cluster to process some task) compared to their anticipated future demand for this resource, which leads to exactly the same normalization.

It is worth noting that the actual constant that the values add up to in the normalization has no impact on our competitive ratio bounds. For example, if were to instead scale the valuations so that $\sum_{t=1}^{T} v_{i,t} = C$ for every agent i, the value of C would not alter the competitive ratio bounds.

4.7 Additional results and proofs

4.7.1 Counter-examples

Non-normalized valuations

Here we show that without the assumption $\sum_t v_{i,t} = 1$ there is no hope for a non-trivial competitive ratio for NSW in online setting. It is easy to see that splitting every item equally gives a $1/n$ approximation of NSW. We now prove the following result.

Theorem 4.7.1. *Assume the setting of the paper, excluding the assumption $\sum_t v_{i,t} = 1$. Then, for any $\delta > 0$ there is no online algorithm that guarantees a $\frac{e}{n} + \delta$ approximation of NSW.*

Proof. We assume an arbitrary online allocation algorithm M is employed and construct an instance on which it fails to provide a $\frac{e}{n} + \delta$ approximation to NSW.

There are $T = n$ rounds in the instance. At $t = 1$ all agents have the value of $v_{i,t} = 1$. We then pick an agent for whom $x_{i,t} \leq 1/n$ (who exists by pigeonhole principle) and that ensure this agent would never have non-zero values for the rest of the instance (we would refer to such agents as frozen).

At any round t we will have non-frozen agent to have values $v_{i,t} = \frac{1}{\epsilon^t}$ (we will pick a sufficiently small $\epsilon < 1$ later on), and freeze the agent whose allocation is $x_{i,t} \leq 1/(n-t)$.

NSW of the optimal allocation x^* can be lower bounded by, on round t, allocating all of to the agent who is about to get frozen, this yields $u_i(x^*) \geq 1/\epsilon^{i-1}$ (we rename indices in the order of being frozen).

On the other hand, by construction the algorithm M achieves $u_i \leq \frac{u_i(x^*)}{i} + O(\epsilon \frac{u_i(x^*)}{i})$. This yields a lower bound for the NSW approximation ratio $\frac{NSW}{NSW^*} \leq (\frac{1}{n!} + O(\epsilon))^{1/n} \approx \frac{e}{n} + O(\epsilon^{1/n})$.

Clearly, choosing a sufficiently small ϵ would yield the needed lower bound for any $\delta > 0$. □

Naive greedy algorithm

Recall that for our main result we reserved half of each item to be split equally among all agents. Here we show that a simpler algorithm that just greedily maximizes Nash Welfare fails to guarantee a non-trivial approximation factor.

Let $u_{i,t}(x_t) = \sum_{t'=1} x_{i,t'} v_{i,t'}$ be the utility of agent i for all rounds up to t (as in the main body, we consider allocations $x_{t'}$ to be fixed by the previous choices of the algorithm). The greedy algorithm is then formally defined with:

$$\hat{x}_t = \arg\max_{x_t} \sum_i \log(u_{i,t}(x_t))$$
$$s.t. \sum_i x_{i,t} \leq 1$$
$$x_{i,t} \geq 0$$

The following lemma is a dual interpretation of greedy algorithm as a myopic price-minimizer, and is proven analogously to lemma 4.4.5.

Lemma 4.7.2. *Consider allocation $\hat{x}_t$ made during execution of greedy algorithm on some round t. Then*

$$\hat{x}_{i,t} \neq 0 \Rightarrow \frac{v_{i,t}}{u_{i,t}(\hat{x}_t)} = \max_{i'} \frac{v_{i't}}{\tilde{u}_{i't}\hat{x}_{i,t}} = \hat{p}_t$$

We now demonstrate an instance on which greedy algorithm fails to do better than allocating all items equally between the agents (approximation factor of n). Consider the following instance with $T = n^2$.

On rounds $t \in [1, n]$ agent $i = 1$ is 'active',by which we mean the agent has value $v_{1t} = 1/n$. Agents other than 1 have values $v_{i,t} = 1/n^{n^2-t+1}$. Thus, on every round their value is asymptotically bigger than on the previous one, and yet is still vanishingly small compared to the active agent.

After n rounds agent 1 exhausted their total value of 1 (and so will have zero values for the rest of the instance), and we switch onto agent 2,who is active for another n rounds, i.e. has value $v_{2t} = \frac{1}{n} - \frac{\sum_{t'<t} v_{2t'}}{n} = \frac{1}{n} - o(\frac{1}{n})$. Other agents $i > 2$ continue to have value $v_{i,t} = 1/n^{(n^2-t+1)}$.

After n^2 rounds, every agent has been active for n rounds, and the instance reaches its end.

We now proceed to bounding the competitive ratio of greedy algorithm on this instance. The optimal allocation can be lower bounded by always giving the item to the active agent, which yields a bound for utilities in the NSW-optimal allocation $u_i(x^*) \geq 1 - o(1/n)$.

Now consider the greedy algorithm applied to the instance described above.

Lemma 4.7.2 implies that the greedy algorithm allocates the marginal item to the agents who currently maximize the ratio $\frac{v_{i,t}}{u_{i,t}(x_t)}$.

On the round t where agent i is active for the first time, we upper bound their utility by assuming they obtain the whole item (and thus the value of $k/T - o(1/T)$). On the following rounds the ratio $\frac{v_{i,t}}{u_{i,t}(x_t)} = 1 - o(1)$ for the active agents, but $\frac{v_{i,t}}{u_{i,t}(x_t)} \geq n - i$ for the other agents with non-zero value (the bound of $n - i$ comes from allocating all of the items equally between non-active agents with non-zero value). Thus the active agent is not allocated anything on any rounds except for the first one.

Thus for the utilities of agents in the greedy solution can be bounded with $u_i(\hat{x}) \leq \frac{1}{n} + o(\frac{1}{n})$.

This implies an approximation factor $A \geq n$, which is no better than just allocating all items equally between the agents.

4.7.2 Proportional allocation rule

In Section 4.4 we mentioned that proportional allocation Pareto-dominates uniform allocation in terms of agents' utilities. With the following theorem we prove this fact.

Theorem 4.7.3. *For any instance, let $\tilde{x}$ be the allocations made by proportional allocation $\tilde{x}_{i,t} = \frac{v_{i,t}}{\sum_j v_{j,t}}$. Then*

$$u_i(\tilde{x}) \geq \frac{1}{n}$$

Proof. We fix any agent i, the parameters T and n, and the values of this agent $v_{i,t}$.

We prove the statement of the theorem by optimizing the values of all other agents to make the utility of agent i as small as possible, and show that the minimum value is never less than $1/n$.

Notice that for the Proportional Sharing allocations $x_i t$ of agent i it does not matter what particular values everyone else reports, only their sum is important, as $x_{i,t} = \frac{v_{i,t}}{v_{i,t}+\sum_{j \neq i} v_{j,t}}$. So, in minimizing the utility of agent i, one can think of agents other than i as a single super-agent who can report values v'_t that sum up to no more than $n-1$. Then, the problem becomes one of choosing v'_t as to minimize the utility of agent i, which is the following convex program:

$$\min_{v'} \sum_t v_{i,t} \frac{v_{i,t}}{v_{i,t} + v'_t}$$
$$s.t. \sum_t v'_t \leq n-1.$$

Lagrangifying the constraint with λ variable, one can write down the first-order conditions for this program:

$$-\frac{v_{i,t}^2}{(v_{i,t} + v'_t)^2} + \lambda = 0 \quad for\ all\ t \tag{4.8}$$

$$\sum_t v'_t = n-1 \tag{4.9}$$

To solve it, we substitute $v'_t = \alpha v_{i,t}$, which solves the first equation for all t if $\lambda = \frac{1}{(1+\alpha)^2}$, and then find α from the second equation:

$$\sum_t v'_t = \alpha \sum_t v_{i,t} = n-1$$

Since $\sum_t v_{i,t} = 1$, we get $\alpha = n-1$ and $v'_t = (n-1)v_{i,t}$.

This means that on every round the utility-minimizing allocation $x_{i,t} = \frac{v_{i,t}}{v_{i,t}+v'_t} = \frac{1}{n}$ is the uniform allocation, and the statement of the theorem follows.

□

4.7.3 Computational Tractability of the Set-Aside Greedy Algorithm

We now briefly note that the Set-Aside Greedy algorithm is quite tractable from a computational standpoint. The most demanding operation of this algorithm involves the computation of $\hat{z}_{i,t}$ for each round t according to the following program:

$$\begin{aligned}\hat{z}_t = \arg\max_{z_t} & \sum \log(\tilde{u}_{i,t}(z_t))\\ s.t. & \sum_i z_{i,t} \le \frac{1}{2}\\ & z_{i,t} \ge 0.\end{aligned}$$

Although this is a convex program that can be solved in polynomial time, we can actually also provide a fast alternative process for computing this allocation. Lemma 4.4.5 implies that the problem above is equivalent to the problem of minimizing the current promised price:

$$\hat{z}_t = \arg\min_{z_t} \; p_t \tag{4.10}$$

$$s.t. \; \sum_i z_{i,t} \le \frac{1}{2}$$

$$z_{i,t} \ge 0$$

$$p_t \ge \frac{v_i}{\tilde{u}_i(z_t)} \tag{4.11}$$

We now provide an algorithm to solve this problem. The idea is to allocate the marginal fraction of the item to the agents for whom the constraint (4.11) is binding, as there is no other way to decrease the objective. We now demonstrate how to allocate the entire item this way n steps.

We initialize $z_{i,t} = 0$ for every agent, and initialize the unallocated fraction of the item at $B = 1/2$. We order the agents in decreasing order on their $\frac{v_{i,t}}{\tilde{u}_{i,t}(z_{i,t})}$ ratios using their promised utilities. Assuming that the agents are re-indexed according to this order, agent 1 has the highest ratio and agent n has the lowest ratio. For each agent $i < n$, we then compute the value δ_i which is the solution to the following equation:

$$\frac{v_{i,t}}{\tilde{u}_{i,t}(\delta_i)} = \frac{v_{i+1,t}}{\tilde{u}_{i+1,t}(0)}.$$

This value, δ_i captures how much allocation does agent i need to catch up to the next agent in terms of the $\frac{v_{i,t}}{\tilde{u}_{i,t}(\delta_i)}$ ratio.

Starting with $i = 1$, we then execute the following steps. If $i\delta_i \leq B$, we allocate δ_i to agent i and all agents before i: $z_{j,t} := z_{j,t} + \delta_i \ for \ j \leq i$, update the remaining allocation $B := B - i\delta_i$ and move on to the next step $i := i + 1$.

If $i\delta_i > B$, we allocate the remaining item fraction B equally between the agent i and agents before i, $z_{j,t} := z_{j,t} + \frac{B}{i} \ for \ j \leq i$ (terminating the algorithm on this round).

If the algorithm gets to $i = n$, we similarly allocate the remaining item fraction equally between everyone $z_{j,t} := z_{j,t} + \frac{B}{n} \ for \ all \ j \in N$.

Thus the number of operations on a single round is at most $O(n)$, giving us $O(Tn)$ complexity for making the allocations over all rounds.

We now briefly give an argument for correctness. First, as we mentioned before, it is always optimal to allocate the marginal fraction of the item to the agent with the smallest $\frac{v_{i,t}}{\tilde{u}_{i,t}(z_{i,t})}$ ratio, where $z_{i,t}$ is the allocation given to agent i so far. The second thing we need to note is that, when this ratio is tied between several agents

$$\frac{v_{i,t}}{\tilde{u}_{i,t}(z_{i,t})} = \frac{v_{j,t}}{\tilde{u}_{j,t}(z_{j,t})}, \tag{4.12}$$

adding equal values ϵ to their allocation preserves the tie

$$\frac{v_{i,t}}{\tilde{u}_{i,t}(z_{i,t} + \epsilon)} = \frac{v_{j,t}}{\tilde{u}_{j,t}(z_{j,t} + \epsilon)}$$

To see that this is true, recall that $\tilde{u}_{i,t}(z_{i,t} + \epsilon) = u_{i,t}(z_{i,t}) + v_{i,t}\epsilon$, and cancel out equal terms implied by (4.12).

This implies that the algorithm always allocates the marginal fraction of the item to the agents minimizing the $\frac{v_{i,t}}{\tilde{u}_{i,t}(z_{i,t})}$ ratio, and thus solves (4.10) exactly.

BIBLIOGRAPHY

[1] Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.

[2] Chidambaram Annamalai, Christos Kalaitzis, and Ola Svensson. Combinatorial algorithm for restricted max-min fair allocation. *ACM Trans. Algorithms*, 13(3):37:1–37:28, 2017.

[3] Arash Asadpour, Uriel Feige, and Amin Saberi. Santa claus meets hypergraph matchings. *ACM Transactions on Algorithms*, 8(3):24, 2012.

[4] Arash Asadpour and Amin Saberi. An approximation algorithm for max-min fair allocation of indivisible goods. *SIAM J. Comput.*, 39(7):2970–2989, 2010.

[5] Eduardo M Azevedo and Eric Budish. Strategy-proofness in the large. Technical report, National Bureau of Economic Research, 2017.

[6] Santiago Balseiro, Huseyin Gurkan, and Peng Sun. Multi-agent mechanism design without money. *Working paper*, 2017.

[7] Santiago R Balseiro, Omar Besbes, and Gabriel Y Weintraub. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science*, 61(4):864–884, 2015.

[8] Santiago R Balseiro and Yonatan Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 2019.

[9] Nikhil Bansal and Maxim Sviridenko. The Santa Claus problem. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 31–40, 2006.

[10] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In Éva Tardos, Edith Elkind, and Rakesh Vohra, editors, *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, pages 557–574. ACM, 2018.

[11] MohammadHossein Bateni, Moses Charikar, and Venkatesan Guruswami. Maxmin allocation via degree lower-bounded arborescences. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009*, pages 543–552, 2009.

[12] Gerdus Benade, Aleksandr M. Kazachkov, Ariel D. Procaccia, and Christos-Alexandros Psomas. How to make envy vanish over time. In Éva Tardos, Edith Elkind, and Rakesh Vohra, editors, *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, pages 593–610. ACM, 2018.

[13] Dirk Bergemann and Maher Said. Dynamic auctions. *Wiley Encyclopedia of Operations Research and Management Science*, 2011.

[14] Dirk Bergemann and Juuso Välimäki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010.

[15] Ivona Bezáková and Varsha Dani. Allocating indivisible goods. *SIGecom Exch.*, 5(3):11–18, April 2005.

[16] Anna Bogomolnaia, Hervé Moulin, and Fedor Sandomirskiy. A simple online fair division problem. *CoRR*, abs/1903.10361, 2019.

[17] Simina Brânzei, Vasilis Gkatzelis, and Ruta Mehta. Nash social welfare approximation for strategic agents. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC 2017*, pages 611–628, 2017.

[18] Simina Brânzei and Ariel D Procaccia. Verifiably truthful mechanisms. In *ACM ITCS 2015*, pages 297–306. ACM, 2015.

[19] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

[20] Eric Budish, Judd B Kessler, et al. *Bringing Real Market Participants' Real Preferences into the Lab: An Experiment that Changed the Course Allocation Mechanism at Wharton*. National Bureau of Economic Research, 2016.

[21] Eric B Budish, Gerard Cachon, Judd B Kessler, and Abraham Othman. Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Chicago Booth Research Paper*, (15-28), 2015.

[22] Ioannis Caragiannis, David Kurokawa, Hervé C. Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC 2016*, pages 305–322, 2016.

[23] Ruggiero Cavallo. Incentive compatible two-tiered resource allocation without money. In *AAMAS 2014*, pages 1313–1320, 2014.

[24] Deeparnab Chakrabarty, Julia Chuzhoy, and Sanjeev Khanna. On allocating goods to maximize fairness. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009*, pages 107–116, 2009.

[25] Shuchi Chawla and J Benjamin Miller. Mechanism design for subadditive agents via an ex ante relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 579–596. ACM, 2016.

[26] Richard Cole and Vasilis Gkatzelis. Approximating the Nash social welfare with indivisible items. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015*, pages 371–380, 2015.

[27] Richard Cole, Vasilis Gkatzelis, and Gagan Goel. Positive results for mechanism design without money. In *International conference on Autonomous Agents and Multi-Agent Systems, AAMAS '13*, pages 1165–1166, 2013.

[28] Claude d'Aspremont and Louis-André Gérard-Varet. Incentives and incomplete information. *Journal of Public economics*, 11(1):25–45, 1979.

[29] Nikhil R Devanur, Yuval Peres, and Balasubramanian Sivan. Perfect bayesian equilibria in repeated sales. In *ACM-SIAM SODA 2015*, pages 983–1002. SIAM, 2015.

[30] Devdatt P Dubhashi and Alessandro Panconesi. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.

[31] Shaddin Dughmi and Arpita Ghosh. Truthful assignment without money. In *ACM EC 2010*, pages 325–334. ACM, 2010.

[32] Uriel Feige. On allocations that maximize fairness. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 287–293, 2008.

[33] Uriel Feige. On maximizing welfare when utility functions are subadditive. *SIAM Journal on Computing*, 39(1):122–142, 2009.

[34] Eric Friedman, Christos-Alexandros Psomas, and Shai Vardi. Dynamic fair division with minimal disruptions. In *Proceedings of the sixteenth ACM conference on Economics and Computation*, pages 697–713, 2015.

[35] Eric J. Friedman, Christos-Alexandros Psomas, and Shai Vardi. Controlled dynamic fair division. In Constantinos Daskalakis, Moshe Babaioff, and Hervé Moulin, editors, *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, pages 461–478. ACM, 2017.

[36] Drew Fudenberg and Jean Tirole. Game theory, 1991. *Cambridge, Massachusetts*, 393:12, 1991.

[37] Alex Gershkov and Benny Moldovanu. Dynamic revenue maximization with heterogeneous objects: A mechanism design approach. *American economic Journal: microeconomics*, 1(2):168–198, 2009.

[38] Ashish Goel, Mohammad Mahdian, Hamid Nazerzadeh, and Amin Saberi. Advertisement allocation for generalized second-pricing schemes. *Operations Research Letters*, 38(6):571–576, 2010.

[39] Artur Gorokh, Siddhartha Banerjee, and Krishnamurthy Iyer. From monetary to non-monetary mechanism design via artificial currencies. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, pages 563–564, 2017.

[40] Artur Gorokh, Banerjee Siddhartha, and Krishnamurthy Iyer. On the effectiveness of first price auctions in artificial currency economies. *Submitted to the Twenty-First ACM Conference on Economics and Computation.*

[41] Artur Gorokh, Ghatzelis Vasilis, and Banerjee Siddhartha. Online fair division via promised utilities. *Submitted to the Twenty-First ACM Conference on Economics and Computation.*

[42] Ramakrishna Gummadi, Peter Key, and Alexandre Proutiere. Repeated auctions under budget constraints: Optimal bidding strategies and equilibria. *SSRN working draft*, 2012.

[43] Mingyu Guo and Vincent Conitzer. Strategy-proof allocation of multiple items between two agents without payments or priors. In *AAMAS 2010*, pages 881–888, 2010.

[44] Mingyu Guo, Vincent Conitzer, and Daniel M Reeves. Competitive repeated allocation without payments. In *WINE 2009*, pages 244–255. Springer, 2009.

[45] Joseph Y Halpern. Beyond nash equilibrium: Solution concepts for the 21st

century. In *Proceedings of the twenty-seventh ACM symposium on Principles of distributed computing*, pages 1–10. ACM, 2008.

[46] Jason D Hartline, Robert Kleinberg, and Azarakhsh Malekian. Bayesian incentive compatibility via matchings. In *Proceedings of the twenty-second annual ACM-SIAM SODA 2011*, 2011.

[47] Jason D Hartline and Tim Roughgarden. Optimal mechanism design and money burning. In *ACM STOC 2008*, pages 75–84. ACM, 2008.

[48] Jiafan He, Christos-Alexandros Psomas, Ariel D Procaccia, and David Zeng. Achieving a fairer future by changing the past. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019.

[49] Krishnamurthy Iyer, Ramesh Johari, and Mukund Sundararajan. Mean field equilibria of dynamic auctions with learning. *Management Science*, 60(12):2949–2970, 2014.

[50] Matthew O Jackson and Hugo F Sonnenschein. Overcoming incentive constraints by linking decisions. *Econometrica*, 75(1):241–257, 2007.

[51] Ehud Kalai and Meir Smorodinsky. Other solutions to nash's bargaining problem. *Econometrica: Journal of the Econometric Society*, pages 513–518, 1975.

[52] Mamoru Kaneko and Kenjiro Nakamura. The Nash social welfare function. *Econometrica*, 47(2):pp. 423–435, 1979.

[53] Yash Kanoria and Hamid Nazerzadeh. Dynamic reserve prices for repeated auctions: Learning from bids. *Available at SSRN 2444495*, 2014.

[54] Chinmay Karande, Aranyak Mehta, and Ramakrishnan Srikant. Optimizing budget constrained spend in search advertising. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 697–706. ACM, 2013.

[55] Ian A Kash, Eric J Friedman, and Joseph Y Halpern. Optimizing scrip systems: Efficiency, crashes, hoarders, and altruists. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 305–315. ACM, 2007.

[56] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. When can the maximin share guarantee be guaranteed? In *Proceedings of the Thirtieth*

AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA, pages 523–529, 2016.

[57] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006.

[58] Renato Paes Leme, Vasilis Syrgkanis, and Éva Tardos. Sequential auctions and externalities. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 869–886. SIAM, 2012.

[59] Bo Li, Wenyang Li, and Yingkai Li. Dynamic fair division problem with general valuations. *arXiv preprint arXiv:1802.05294*, 2018.

[60] Thodoris Lykouris, Vasilis Syrgkanis, and Éva Tardos. Learning and efficiency in games with dynamic population. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 120–129. Society for Industrial and Applied Mathematics, 2016.

[61] Paul Robert Milgrom. *Putting auction theory to work*. Cambridge University Press, 2004.

[62] H. Moulin. *Fair Division and Collective Welfare*. The MIT Press, 2003.

[63] Hervé Moulin. The proportional random allocation of indivisible units. *Social Choice and Welfare*, 19(2):381–413, 2002.

[64] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[65] J. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, April 1950.

[66] Hamid Nazerzadeh, Amin Saberi, and Rakesh Vohra. Dynamic cost-per-action mechanisms and applications to online advertising. In *WWW 2008*, pages 179–188. ACM, 2008.

[67] Denis Nekipelov, Vasilis Syrgkanis, and Eva Tardos. Econometrics for learning agents. In *ACM EC 2015*, pages 1–18. ACM, 2015.

[68] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

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

[70] Canice Prendergast. How food banks use markets to feed the poor. *Journal of Economic Perspectives*, 31(4):145–62, 2017.

[71] Ariel D Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. *ACM Transactions on Economics and Computation (TEAC)*, 1(4):18, 2013.

[72] Tim Roughgarden. The price of anarchy in games of incomplete information. In *ACM EC 2012*, pages 862–879. ACM, 2012.

[73] Vasilis Syrgkanis and Eva Tardos. Composable and efficient mechanisms. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 211–220. ACM, 2013.

[74] Toby Walsh. Allocation in practice. In *Joint German/Austrian Conference on Artificial Intelligence (Künstliche Intelligenz)*, pages 13–24. Springer, 2014.

[75] Qiqi Yan. Mechanism design via correlation gap. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 710–719. Society for Industrial and Applied Mathematics, 2011.

[76] David Zeng and Alexandros Psomas. Fairness-efficiency tradeoffs in dynamic fair division. *CoRR*, abs/1907.11672, 2019.