

Prior-free Dynamic Mechanism Design With Limited Liability

MARK BRAVERMAN, Department of Computer Science, Princeton University

JON SCHNEIDER, Google Research

S. MATTHEW WEINBERG, Department of Computer Science, Princeton University

We study the problem of repeatedly auctioning off an item to one of k bidders where: a) bidders have a per-round individual rationality constraint, b) bidders may leave the mechanism at any point, and c) the bidders' valuations are adversarially chosen (the prior-free setting). Without these constraints, the auctioneer can run a second-price auction to "sell the business" and receive the second highest total value for the entire stream of items. We show that under these constraints, the auctioneer can attain a constant fraction of the "sell the business" benchmark, but no more than $2/e$ of this benchmark.

In the course of doing so, we design mechanisms for a single bidder problem of independent interest: how should you repeatedly sell an item to a (per-round IR) buyer with adversarial valuations if you know their total value over all rounds is V but not how their value changes over time? We demonstrate a mechanism that achieves revenue V/e and show that this is tight.

CCS Concepts: • **Applied computing** → **Online auctions**.

Additional Key Words and Phrases: dynamic mechanisms, algorithmic mechanism design, auctions

ACM Reference Format:

Mark Braverman, Jon Schneider, and S. Matthew Weinberg. 2021. Prior-free Dynamic Mechanism Design With Limited Liability. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC '21), July 18–23, 2021, Budapest, Hungary*. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3465456.3467654>

1 INTRODUCTION

Classical dynamic mechanism design traditionally assumes that the values of the participants are generated *stochastically*: for example, drawn iid from some distribution every round, or generated according to a simple stochastic process known to the mechanism designer. These assumptions are often somewhat unrealistic; actual valuations drift over time and are subject to shocks in a way which is hard to predict or model. A mechanism designer may hope for a dynamic mechanism robust to these features of the problem.

The goal of this paper is to initiate the study of dynamic mechanism design in the presence of *adversarially* chosen valuations (the prior-free setting). Specifically: there are k bidders, and T rounds. Bidder i has value $v_{i,t}$ for the round- t item, and bidders are additive across rounds. The designer's goal is to maximize their *revenue*.

It is of course impossible to achieve any approximation guarantee with respect to the "first-best" benchmark of $\sum_t \max_i \{v_{i,t}\}$ (perhaps all but one bidder has a constant value of 0 — there is no way to guess the correct price to set for the remaining bidder). This challenge also arises in the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '21, July 18–23, 2021, Budapest, Hungary.

© 2021 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8554-1/21/07...\$15.00.

<https://doi.org/10.1145/3465456.3467654>

vast literature on prior-free auctions in other domains, and care is required to pick an insightful “second-best” benchmark [15, 16, 19, 22–24].

Let us quickly consider two potential benchmarks. First, one might target the second-best per-round: $\text{Rev}_{SPA} := \sum_t \text{Second-Highest}_i\{v_{i,t}\}$. There is also a simple, dominant-strategy truthful auction which achieves Rev_{SPA} : simply auction each item via a second-price auction. This benchmark, however, may be infinitely far from optimal if for each item there is a unique bidder with non-zero value. On the other extreme, one might target the aggregate second-best: $\text{Rev}_{STB} := \text{Second-Highest}_i\{\sum_t v_{i,t}\}$. There is also a simple, dominant-strategy truthful auction which achieves Rev_{STB} : simply “sell the business” and auction the entire stream of items at the very beginning.

Limited Liability and Per-Round Individual Rationality. In a static setting, Rev_{SPA} corresponds to selling the items separately (using a second-price auction), and Rev_{STB} corresponds to bundling the items together (for a second-price auction). Both mechanisms are simple to implement in the static setting, and in fact variants of (the better of) these are known to yield good approximation guarantees in static settings [4, 11, 12, 14, 32]. In a dynamic setting, however, bundling all time slots together requires the bidder to pay a large amount up front and trust that seller will honor their promise to follow the protocol in the rounds to come. For this reason, much literature on dynamic mechanism design studies what is possible under per-round “limited liability” or “individual rationality” constraints on the buyer, where they never pay more than their value per round [1, 2, 5, 28–31]. Recalling that the benchmark Rev_{STB} can be achieved in the static setting (or the dynamic setting without limited liability), the main question we ask is:

What approximation to Rev_{STB} can be guaranteed by a dominant-strategy truthful mechanism in the dynamic setting with limited liability?

Observe further that any positive resolution to the above question will also imply an approximation guarantee to $\max\{\text{Rev}_{STB}, \text{Rev}_{SPA}\}$, by simply flipping a fair coin to run either the designed mechanism or a per-round second-price auction (which already satisfies limited liability). In other words, the Rev_{STB} benchmark is key to understanding the revenue lost in prior-free settings *specifically due to limited liability constraints* (rather than other sources which harm revenue already in the static setting).

1.1 Main results and techniques

We design a dominant-strategy truthful, limited-liability mechanism which achieves a constant-factor approximation to Rev_{STB} . We also establish that it is impossible to achieve an arbitrarily close approximation to Rev_{STB} , let alone match it. Below are informal statements of our main results, along with pointers to the formal statements.

THEOREM 1 (INFORMAL RESTATEMENT OF THEOREM 12). *There exists a mechanism (for two bidders) which guarantees revenue at least $\frac{1}{2e} \text{Rev}_{STB} - O(1)$.*

It is also worth noting that our mechanism does *not* require the bidders to know their individual per round values ahead of time, but just their average value over all T rounds. Still, the mechanism is dominant-strategy truthful even if they have full knowledge of the future. For k bidders, a simple generalization of our 2-bidder mechanism leads to a competitive ratio of $1/(ke)$. We also show how to improve this to a constant competitive ratio independent of k , although the solution concept for our auction is implementation in undominated strategies (that is, it guarantees an approximation guarantee of α as long as every bidder plays an undominated strategy).

THEOREM 2 (INFORMAL RESTATEMENT OF THEOREM 14). *There exists a mechanism for k bidders which guarantees revenue at least $\alpha \cdot \text{Rev}_{STB} - O(k)$ for some constant α (independent of k).*

Finally, we show it is impossible to achieve revenue arbitrarily close to Rev_{STB} , even when there are only two bidders.

THEOREM 3 (INFORMAL RESTATEMENT OF THEOREM 10). *For any $\alpha > 2/e$, there is no mechanism (for two bidders) which always guarantees revenue at least $\alpha \cdot \text{Rev}_{STB} - O(1)$.*

Since the k bidder game contains the 2 bidder game as a subcase, it is still the case that no mechanism for k bidders can achieve an approximation ratio better than $2/e$ (Corollary 11).

Technical Overview: A Related Single-Bidder Problem. The main ingredient in our above results is actually a thorough understanding of a related single-buyer problem. Specifically, consider exactly the same setup as our model, except there is only a single bidder *and the designer knows* $V = \sum_t v_t$ (or equivalent, that the designer knows $v := V/T$). In the static setting, there is a trivial optimal solution (sell the bundle of all items for V , which is the first-best). Here we seek to answer the same questions as above: what fraction of this benchmark can the seller guarantee subject to limited liability?

Note that the prior-free nature of the valuations is an essential detail here. If the valuations were generated stochastically from a time-invariant distribution, the seller could run mechanisms of the form “continue giving the item to the buyer as long as the buyer has paid at least $(1 - \epsilon)v$ on average each round so far” and extract almost the entire expected welfare. When valuations are adversarial, such mechanisms are no longer possible – it could be the case that the value of the buyer is entirely back-loaded and is zero for the first half of the rounds. Nonetheless, we show that the seller can still (asymptotically) attain a $1/e$ fraction of the total welfare V , and that this is tight.

THEOREM 4 (INFORMAL RESTATEMENT OF THEOREMS 6 AND 8). *There is a mechanism for the single buyer game that achieves revenue $V/e - O(1)$. Moreover, this is tight: for any $\alpha > 1/e$, there is no mechanism for the single buyer game that can guarantee revenue $\alpha V - O(1)$.*

The mechanism we design for Theorem 6 is what we call a *pay-to-play mechanism*. Such mechanisms are parameterized by an increasing function $\rho : [0, 1] \rightarrow (0, 1]$ and work as follows: on round t , the seller allocates the item to the buyer with probability $\rho(X_t/V)$, where X_t is the total payment the buyer has made to the seller by time t . The buyer is free to make any payment they wish each round – since ρ is increasing, they are incentivized to pay early to increase their allocation probability later (and in particular, for whatever total payment they choose to make, they are incentivized to make that payment as soon as possible).

We show that the optimal such f (specifically, $f(w) = \min(e^{ew-1}, 1)$) leads to a pay-to-play mechanism with revenue V/e . Our impossibility result shows that this mechanism is optimal, and uses much of the same machinery. One highlight of this approach is that we reparameterize *any* truthful mechanism in terms of the allocation/payment *as a function of the maximum possible welfare the buyer could have had so far* (rather than as a function of the round t). Under this reparameterization, we find a pay-to-play interpretation of any mechanism, meaning that the optimal pay-to-play mechanism is in fact optimal among all mechanisms.

Technical Overview: Back to Multiple Bidders. We apply the following technique to extend these ideas to 2 (and $k > 2$) bidders. We begin by soliciting each bidder’s total value (e.g. $V_a = \sum a_t$ and $V_b = \sum b_t$). We then split each item in half and run two copies of the single bidder mechanism: for Alice, we run the single bidder mechanism on one of the halves, assuming her total value for these halves is $V_b/2$, and for Bob we run the single bidder mechanism on the other halves, assuming

his total value for these halves is $V_a/2$. If $V_a > V_b$, then $V_b/2$ is an underestimate for Alice’s value, and our single bidder mechanism guarantees we receive revenue at least $V_b/2e$, or equivalently, $\text{Rev}_{STB}/2e$.

This simple idea extends to k bidders, but causes the approximation guarantee to degrade to $1/(ke)$. The main issue with this approach is that we are allocating $1/k$ of the items to each bidder, even to bidders that cannot afford them. To fix this, we introduce an element of competition to this mechanism in the form of a first-price auction. Roughly speaking, each round t we allow each bidder i to submit a bid $b_{i,t}$ in an attempt to increase the rate at which they are allocated items. The bidder with the highest bid pays that amount, which in turn increases the rate at which they are allocated items (this description elides several details which can be found in Section 4.2). This mechanism is not truthful, but we show that *as long as all bidders play undominated strategies*, this mechanism achieves an α -approximation to Rev_{STB} for a constant α independent of k .

1.2 Related Work

The problem of selling a stream of items to one of several bidders is a central problem in the field of dynamic mechanism design (for a general introduction to the area, we recommend the survey [8]). Our treatment of this problem imposes two important constraints: the assumption of lack of long-term trust between the auctioneer and seller, and the adversarial prior-free nature of the bidders’ valuations.

Lack of long-term commitment. There is a major recent line of work which studies the design of dynamic mechanisms satisfying ex-post IR constraints (or “limited liability”) [1, 2, 5, 28–31]. In this language, our constraint that bidders never pay more than their value in a given round is a “per round ex-post IR” constraint. Most of these papers study a specific class of ex-post IR mechanisms called *bank account mechanisms*, where the seller maintains a virtual “bank account” for each bidder, and allows bidders to draw from and pay to this bank account so that they end up paying (for example) their average value per round. Such mechanisms work well and can be shown to be optimal for stochastic valuations (where the total liability of the seller – the size of the bank account – can be bounded sublinearly in T) but fail for adversarial valuations (where the liability of the seller can grow linearly in T , e.g. for a bidder whose values are back-loaded).

Prior-free valuations. There is a large spectrum of ways to model the valuations of the buyer, from weakest (being independently drawn each round from a time-invariant distribution) to strongest (the prior-free setting). The majority of the work mentioned above studies settings where each bidder’s value is drawn iid each round from some distribution known to the seller (or occasionally more generally, the value for round t is drawn from a distribution specific to round t but also known to the seller in advance). One notable exception is [30] (and the follow-up work [17]), which studies “non-clairvoyant dynamic mechanism design” for selling to a single buyer, where the buyer has a different value distribution for each round but the seller is only aware of the distribution for the current round. Another is [27], which studies the competition complexity of auctions in a non-clairvoyant setting.

One related line of work is the line of work on dynamic mechanisms for buyers with evolving values [3, 6, 7, 9, 13, 26]. In these works, the value of the buyer evolves according to some stochastic process over time (perhaps depending on the action of the seller or buyer). Again, the standard assumption here is that the seller has detailed knowledge of this stochastic process, and as such tend not to apply to adversarial valuations.

To the best of our knowledge, very few papers have studied dynamic mechanism design in a completely adversarial setting. The most closely related such work is [18], which studies the problem of selling to a buyer with adversarial valuations, but makes the additional assumption that

the buyer is running a low-regret learning algorithm. In contrast to this, prior-free *static* mechanism design has been extensively studied [15, 16, 19, 21–24]. See also Chapter 7 of [25].

Contract Theory. The primary benchmark we compare against is the revenue obtained when “selling the business”; i.e. selling the entire stream of items at the beginning of the protocol through a second-price auction. A similar strategy of “selling the firm” is a common concept in the contract theory literature. As in our setting, for the classical principal-agent problem with hidden actions, “selling the firm” is the optimal revenue strategy for the principal in the absence of limited liability. [20] studies the approximation ratio between the optimal contract and this first-best benchmark in this classical problem.

Indeed, our setting can be viewed as a peculiar multi-agent dynamic contract problem, where the principal is the seller and the buyers are agents (with actions corresponding to possible payments to the principal). A similar approach is taken in [10], which examines this model (in the stochastic values setting) when the principal uses a low-regret learning algorithm to allocate items.

2 MODEL

2.1 Multiple bidder game

We consider a setting with $k \geq 2$ bidders and one seller. There are T rounds, and in each round the seller has one item for sale. Bidder i has value $v_{i,t} \in [0, 1]$ for the item sold at time t . Unlike in much prior work, where $v_{i,t}$ are generated stochastically from some known prior, we consider the case where all $v_{i,t}$ are adversarially set at the beginning of the game. We assume the seller has no knowledge of the valuations $v_{i,t}$. We further assume that each bidder i has full knowledge of their own valuations $v_{i,t}$ over all times t . This assumption exists largely so that strategic play is well-defined for each bidder; in many of the mechanisms we present, we will see that it is possible to relax this assumption to bidders who only know their average valuation per round.

During each round t , the seller and the k bidders are allowed to participate in arbitrary communication. At the end of round t , the seller allocates the item among the k bidders by giving some fraction¹ $r_{i,t}$ of the item to bidder i (with $\sum_i r_{i,t} \leq 1$; in particular, the seller does not need to allocate all of the item). In return for this portion of the item, bidder i pays the seller some price $p \in [0, r_{i,t}v_{i,t}]$. Note that we do not include any mechanism for the seller to force the bidder to commit to paying a specific price p for the item in a given round; but as this is a repeated game, informal commitments of the form “if bidder i does not pay the agreed upon price in a round, the seller will never allocate any item to bidder i ever again” are possible².

Note also that we assume the bidder never pays the seller more than their value for the item in a given round. This can be seen as an “ex post individual rationality” or “limited liability” assumption, and rules out protocols where the seller simply sells the entire stream of items at the beginning of the protocol to the highest bidder. Nonetheless, we will see that our mechanisms have good performance even when some or all of the bidders are allowed to break this constraint and pay more than their value.

Finally, at the end of the game, we allow a single bonus round where the seller is allowed to reimburse each bidder an amount of the seller’s choosing. These reimbursements provide a convenient way to incentivize bidders to participate for all T rounds (otherwise e.g. there is never

¹For mathematical convenience, we assume the seller runs a deterministic strategy with the ability to fractionally allocate the item. We believe our results should carry over (with perhaps some loss in terms sublinear in T) to randomized settings where the seller awards the entire item to bidder i with probability $r_{i,t}$.

²In particular, all our results extend to an alternate model where the seller offers bidder i a price p for some fraction of the item, and the buyer either can either accept the price and take the item or reject the price and leave the item.

any incentive to pay the seller anything in round T), but are otherwise unimportant and are small (on the order of $O(1)$) in most of the mechanisms we present³.

Strategies and solution concepts. We are primarily interested in mechanisms with dominant strategy equilibria for the bidders. Specifically, let s_i denote a deterministic strategy for bidder i (i.e., a function from states of the mechanism and the valuations $v_{i,t}$ to allowed actions of bidder i at that state), and let $U_i(s_i, s_{-i})$ denote the expected utility received by bidder i (where the expectation is over potential randomness in the seller's actions) when i plays according to s_i and all other bidders play according to s_{-i} . Then the tuple $(s_1^*, s_2^*, \dots, s_k^*)$ is a dominant strategy equilibrium for this mechanism if for all s_i and s_{-i} ,

$$U_i(s_i^*, s_{-i}) \geq U_i(s_i, s_{-i}).$$

Any mechanism with a dominant strategy equilibrium can be transformed into a truthful direct-revelation mechanism via the revelation principle, where the bidders report their private types (their valuations $v_{i,t}$) to the seller at the beginning of the protocol and the seller uses this to simulate the bidders' dominant strategies s_i^* on behalf of the bidders. In our setting, we specify a direct-revelation mechanism as follows.

A *direct-revelation mechanism* for the multiple bidder game specifies the following information for each type profile $\mathbf{v} = \{v_{i,t}\}_{i \in [k], t \in [T]}$ of all k bidders:

- For each bidder i , the fraction $r_{\mathbf{v},i,t}$ of the item allocated to bidder i at time t . (These must satisfy $\sum_i r_{\mathbf{v},i,t} \leq 1$ for any $\mathbf{v}$ and t).
- For each bidder i , the price $x_{\mathbf{v},i,t}$ the seller expects bidder i to pay on round t .
- For each bidder i , the reimbursement $g_{\mathbf{v},i}$ paid to bidder i at the end of the game.

To run a direct-revelation mechanism, the seller proceeds as follows:

- (1) The seller begins by soliciting the type profile $\mathbf{v}$ from all of the bidders. The seller computes $r_{\mathbf{v},i,t}$, $p_{\mathbf{v},i,t}$, and $g_{\mathbf{v},i}$, and distributes this information to the bidders.
- (2) On round t , the seller allocates to bidder i a fraction $r_{\mathbf{v},i,t}$ of the item (note that there can be some fraction of the item that goes unallocated), unless bidder i has been eliminated (in which case bidder i receives nothing).
- (3) This bidder responds by paying the seller some price p . If $p \neq x_{\mathbf{v},i,t}$, then the seller eliminates bidder i from further consideration (i.e. will never allocate to or reimburse bidder i again).
- (4) Finally, at the end of the game, the seller gives a reimbursement of $g_{\mathbf{v},i}$ to each uneliminated bidder i .

In order for a direct-revelation mechanism to be *truthful*, these allocations, prices, and reimbursements must satisfy the following two constraints:

- (*Limited liability*) For any $\mathbf{v}$, i , and t , we must have $x_{\mathbf{v},i,t} \leq r_{\mathbf{v},i,t}v_{i,t}$ (i.e., a limited-liability bidder reporting truthfully must be able to pay for the item if they win in this round).
- (*Incentive compatibility*) Fix a type profile $\mathbf{v}$ and a bidder i , and let $\mathbf{v}'$ be the profile where bidder i 's type $v_{i,t}$ is replaced by a new type $v'_{i,t}$. Then, if it is the case that

$$x_{\mathbf{v}',i,t} \leq r_{\mathbf{v}',i,t}v'_{i,t}$$

for all $t \leq \tau$ (i.e. our bidder with limited liability can pretend to have type $\mathbf{v}'$ for τ rounds) we must have that, for any $\tau \leq T$,

³One notable exception is in Section 4.2, where we first demonstrate a mechanism for the k bidder case which uses reimbursements on the order of $O(T)$. We subsequently show how to reduce the size of these reimbursements to $O(1)$, but some may find the mechanism with large reimbursements more natural.

$$\left(\sum_{t=1}^T r_{v,i,t} v_{i,t} - x_{v,i,t} \right) + g_{v,i} \geq \left(\sum_{t=1}^{\tau} r_{v',i,t} v_{i,t} - x_{v',i,t} \right) + x_{v',i,\tau}.$$

That is, a bidder cannot increase their expected utility by misreporting their type and then following the resulting mechanism for τ steps before defecting. The additional term on the RHS represents the fact that the bidder does not need to pay the payment $x_{v',i,\tau}$ if they defect on round τ . Note that both $x_{v',i,\tau}$ and $g_{v,i}$ will typically be $O(1)$, in comparison to the two sums which will typically be $\Theta(T)$.

Revenue and benchmarks. We wish to design truthful mechanisms with high *revenue*: the total sum of payments paid to the seller, minus reimbursements. For a truthful direct-revelation mechanism and bidders with type profile $\mathbf{v}$, this can be written as

$$\text{Rev}(\mathbf{v}) = \sum_{i=1}^k \sum_{t=1}^T x_{v,i,t} - \sum_{i=1}^k g_{v,i}.$$

We will compare the revenue the seller achieves to a benchmark function of $\mathbf{v}$. In this paper, we primarily consider the benchmark of the revenue obtained when auctioning off all T items at once at the beginning of the game, i.e. “selling the business”. This is not possible with limited-liability bidders, but with unconstrained bidders this strategy obtains a revenue equal to

$$\text{Rev}_{STB}(\mathbf{v}) = \max_{(2)} \left(\left\{ \sum_{t=1}^T v_{i,t} \mid i \in [k] \right\} \right),$$

where we write $\max_{(2)}(S)$ to denote the second largest element in S . We say that a mechanism (more specifically, a family of mechanisms, one for each T) is α -competitive with selling the business if for all T and type profiles $\mathbf{v}$,

$$\text{Rev}(\mathbf{v}) \geq \alpha \text{Rev}_{STB}(\mathbf{v}) - O_k(1).$$

(Here we use the notation $O_k(1)$ to denote a constant possibly depending on k but not on T or any other parameter). It should be noted that Rev_{STB} is not the “optimal” policy for the seller when bidders are unconstrained, merely a natural choice. In fact, for any type profile $\mathbf{v}$, there is a mechanism that achieves revenue equal to $\text{Rev}(\mathbf{v}) = \sum_t \max(\{v_{i,t} \mid i \in [k]\})$ on this specific type profile (even if bidders are limited-liability), namely the mechanism which hopes that the type profile is exactly $\mathbf{v}$ and posts prices accordingly.

We will write $\mathbf{v}_i$ to refer to the specific type of player i (in the type profile $\mathbf{v}$), and $|\mathbf{v}_i| = \sum_{t=1}^T v_{i,t}$ to denote the total value of type $\mathbf{v}_i$. For instance, using this notation, we can write $\text{Rev}_{STB}(\mathbf{v}) = \max_{(2)}(\{|\mathbf{v}_i| \mid i \in [k]\})$.

3 OPTIMAL MECHANISMS FOR SELLING TO A SINGLE BUYER

In the course of designing mechanisms and proving lower bounds for the multiple bidder game described in the previous section, it will be useful to think about a different game between a single seller and a single buyer. This game is identical to the $k = 1$ variant of the multiple buyer game described in the previous section, with the following modifications:

- The seller knows (at the beginning of the protocol) the total value V of the buyer over all T rounds.

- Instead of trying to compete with the benchmark of selling the business (the sum of the second-largest value per round), the seller wishes to be α -competitive with respect to the total value V ; i.e., the seller's mechanism should satisfy

$$\text{Rev}(\mathbf{v}) \geq \alpha V - O(1)$$

for all T and types $\mathbf{v}$.

We can define truthful, direct-revelation mechanisms as before. Since there is only a single buyer in this case, we drop all subscripts i where relevant; e.g. a direct-revelation mechanism in the single buyer game is parametrized by sequences $r_{v,t}$, $x_{v,t}$, and $g_{v,t}$.

In this section, we will show that there exists a $(1/e)$ -competitive mechanism for the seller in the single buyer game. Moreover, we will show that this competitive ratio is optimal; no mechanism for the single buyer game can obtain more than a $1/e$ fraction of the welfare.

Our main technique for the both the upper and lower bound is understanding how the rate and allocation functions of the mechanism at time t depend on the total amount of revenue received from the buyer so far. Our $(1/e)$ -competitive mechanism will take the form of what we refer to as a “pay-to-play” mechanism, where the fraction of the item you are allocated at time t is equal to some fixed function of the total amount you have paid to the seller so far.

3.1 Continuous-time models

The analysis of the single buyer game is more clear in a continuous-time setting, where instead of being discretized into T rounds, the buyer's value (and the various parameters of his mechanism) are described by well-behaved functions over a continuous interval of time. We formalize this continuous-time model here. The analogous upper and lower bounds in the discrete-time model are formulated and proved in the appendices of the online supplemental material⁴ – they follow essentially the same logic, but are slightly messier.

In this continuous variant, our buyer has a Riemann-integrable value function $v(t) : [0, 1] \rightarrow [0, \infty)$ representing their value at time t . If the buyer has total value V (known to the seller), this function $v(t)$ should satisfy $\int_0^1 v(t)dt = V$.

A direct-revelation mechanism in this setting specifies for each type $v(t)$ a continuous rate function $r_v(t) : [0, 1] \rightarrow [0, 1]$ (the fraction of the infinitesimal item the buyer receives at time t) and a continuous payment function $x_v(t) : [0, 1] \rightarrow [0, \infty)$. In order for this direct-revelation mechanism to be truthful, it must satisfy the following analogues of the corresponding constraints in Section 2.1:

- For all types v and $t \in [0, 1]$,

$$r_v(t)v(t) - x_v(t) \geq 0.$$

- For any types v, v' and $\tau \in [0, 1]$ that satisfy

$$r_{v'}(t)v(t) - x_{v'}(t) \geq 0 \quad \forall 0 \leq t \leq \tau \quad (1)$$

(i.e., it is possible for v under limited liability constraints to imitate type v' up until time τ), it is the case that

$$\int_0^1 (r_v(t)v(t) - x_v(t))dt \geq \int_0^\tau (r_{v'}(t)v(t) - x_{v'}(t))dt. \quad (2)$$

⁴Also available at <https://arxiv.org/abs/2103.01868>

Note that for continuous mechanisms, we don't need to allow for a final reimbursement; this reflects the fact that only an infinitesimal amount of the product is up for grabs at any point, so we can guarantee a buyer cannot gain positive utility by defecting.

The total revenue this mechanism achieves on type v is given by $\int_0^1 x_v(t)dt$; we denote this quantity as X_v . Our goal is to maximize our worst-case revenue over all types, $M = \inf_v X_v$; if a mechanism attains a specific M , we say it is M -competitive.

Reparametrization by welfare. Define $w_v(t) = \int_0^t r_v(s)v(s)ds$ to be the total utility received by the bidder up until time t , not including payments (we call this quantity the *welfare* of the bidder up until time t). Note that for $t \in [0, 1]$, $w_v(t)$ is monotone increasing from 0 to $W_v = w_v(1)$.

This allows us to reparametrize our existing functions (value, rate, payment) in terms of welfare. For example, we (abusing notation) let $r_v(w) = r_v(w_v^{-1}(w))$ to be the rate at the time when bidder v has received welfare w . We similarly define $v(w)$ and $x_v(w)$.

3.2 A $(1/e)$ -competitive mechanism

In this section we demonstrate a mechanism for the single buyer game which is $(1/e)$ -competitive. As mentioned, this mechanism falls into the class of mechanisms that we call *pay-to-play mechanisms*.

A pay-to-play mechanism is specified by a continuous, weakly-increasing function $\rho(X) : [0, 1] \rightarrow (0, 1]$, denoting the fraction of the item we allocate to the buyer if the buyer has given the seller a fraction X of his welfare so far. The buyer is free to pay however much they wish to at time t , under the constraint that a limited-liability buyer cannot pay more than $\rho(X)$ of their value in any round. We claim that for any pay-to-play mechanism, the dominant strategy for any limited-liability buyer (regardless of their type v) is to pay as much as they can every round until they have paid an X_{opt} fraction (for some value X_{opt} depending only on ρ) of their total value V , and then never pay again.

More formally, we have the following lemma.

LEMMA 5. *Let $\rho(X) : [0, 1] \rightarrow (0, 1]$ be a weakly-increasing continuous function and let X_{opt} be a value of $x \in [0, 1]$ which maximizes the expression*

$$\rho(x) \left(1 - \int_0^x \frac{dx'}{\rho(x')} \right).$$

Consider the pay-to-play mechanism defined by $\rho(X)$. Fix a type v with total value V . Then there exists a dominant strategy for the buyer with type v where they pay at least $X_{opt}V$ to the seller. Moreover, if X_{opt} is a strict maximum, then this property holds for any dominant strategy.

PROOF. Begin by fixing any dominant strategy for the buyer $x_v(t)$ (since there is only one player, this is just any strategy which optimizes the buyer's utility). We first note that if $\int_t^1 x_v(t)dt > 0$ (the buyer pays some amount to the seller after time t), then $\int_0^t x_v(t)dt = \int_0^t r_v(t)v(t)dt = w_v(t)$ (the buyer passes along all reward at times $s < t$). In other words, the buyer's payments are front-loaded; i.e. any dominant strategy must pay as much as possible until it passes some threshold, then stop paying entirely. To see this, it suffices to note that since $\rho(X)$ is an increasing function of X , moving later payments earlier increases the value of all future rewards and hence is strictly optimal.

Any dominant strategy for the buyer can therefore be characterized by the total amount the buyer gives to the seller. We will show there is one dominant strategy where they give at least $X_{opt}V$. Assume that the buyer gives a total of $\bar{x}V$, for some $\bar{x} \in [0, 1]$. If the buyer pays the seller

until time τ , then τ must satisfy $\int_0^\tau x_v(t)dt = \bar{x}V$. But from the above argument, we know that $\int_0^\tau x_v(t)dt = w_v(\tau)$, so this means $\tau = w_v^{-1}(\bar{x}V)$.

Now, note that the eventual utility of the buyer is $\rho(\bar{x})(V - \int_0^\tau v(t)dt)$ (the buyer gets no net utility from the item until time τ , and from then on they get a constant fraction $\rho(\bar{x})$ of the remainder of the item).

Recall that $w_v(t) = \int_0^t r_v(s)v(s)ds$, so $\frac{dw_v}{dt} = r_v(t)v(t)$. This lets us reparametrize the integral $\int_0^\tau v(t)dt$ in terms of welfare via

$$\begin{aligned} \int_0^\tau v(t)dt &= \int_0^{w_v(\tau)} v(t) \cdot \frac{1}{r_v(t)v(t)} dw \\ &= \int_0^{\bar{x}V} \frac{dw}{r_v(w)} \\ &= \int_0^{\bar{x}V} \frac{dw}{\rho(w/V)} \\ &= V \int_0^{\bar{x}} \frac{dw'}{\rho(w')}. \end{aligned}$$

In the second-to-last equality, we have used the fact that while the buyer is passing along their full welfare to the seller, $r_v(w) = \rho(w/V)$. It follows that the net utility of the buyer is equal to

$$V\rho(\bar{x}) \left(1 - \int_0^{\bar{x}} \frac{dw}{\rho(w)} \right).$$

Since $\bar{x}$ was chosen to maximize this expression, this is the maximal possible utility possible for the buyer and hence this is a dominant strategy. If $\bar{x}$ is a strict maximizer, any dominant strategy must pay a total of $\bar{x}V$ to the seller (or it will be dominated by this strategy). $\square$

With Lemma 5 in hand, we can exhibit our $(1/e)$ -competitive mechanism.

THEOREM 6. *There exists a mechanism for the seller which obtains V/e total revenue.*

PROOF. Consider the pay-to-play mechanism defined by the function $\rho(w)$, where

$$\rho(w) = \begin{cases} e^{ew-1} & \text{if } w \in [0, 1/e] \\ 1 & \text{if } w \in [1/e, 1] \end{cases}$$

Note that for $x \leq 1/e$, we have that

$$\begin{aligned} \rho(w) \left(1 - \int_0^x \frac{dx'}{\rho(x')} \right) &= e^{ex-1} \left(1 - \int_0^x e^{-ex'+1} dx' \right) \\ &= e^{ex-1} e^{-ex} \\ &= e^{-1}. \end{aligned}$$

For $x > 1/e$, this expression is decreasing (since $\rho(x)$ is constant for $x \in [1/e, 1]$, but $\int_0^1 (1/\rho(x'))dx'$ is increasing). It follows that $X_{opt} = 1/e$ is a value of x which maximizes utility. By Lemma 5, it follows that a pay-to-play mechanism with this choice of ρ results in the seller receiving V/e total revenue, as desired. $\square$

Remark 1. Note that for the ρ in Theorem 6, the expression in Lemma 5 is only weakly maximized at $X_{opt} = 1/e$ (and in fact attains this maximum for all $X \in [0, 1/e]$). It is possible to perturb ρ slightly so that the expression is *strictly* maximized at some $X_{opt} \geq 1/e - \varepsilon$ for some arbitrarily small ε , thus satisfying the stronger conditions of Lemma 5. One option is to choose $\tilde{\rho}(w) = \rho(\min(w/(1-\varepsilon), 1))$. This can be thought of as allocating the item via the original pay-to-play mechanism but as if the buyer actually had value $(1-\varepsilon)V$ (to see why this works, see Remark 2). Alternatively, it suffices to give any positive reimbursement at the end of the game contingent on the buyer paying at least V/e .

Remark 2. The above analysis assumes that the total value of the buyer is exactly V , but in fact works as long as V is any lower bound on the total value of the buyer (that is, it still guarantees a revenue of at least V/e). Indeed, if the buyer's true value is $V' \geq V$, then by stopping at a payment of $\bar{x}V$ they receive a net utility of

$$\rho(\bar{x}) \left(V' - V \int_0^{\bar{x}} \frac{dw}{\rho(w)} \right).$$

For the ρ in Theorem 6, this expression is strictly increasing for $\bar{x} \in [0, 1/e]$ (in particular, it equals $(V' - V)\rho(\bar{x}) + e^{-1}V$) and strictly decreasing for $\bar{x} > 1/e$, so it is strictly dominant for the buyer to pay V/e .

Remark 3. Likewise, although the above analysis assumes that the buyer is subject to limited liability and does not pay more than his value, this too is not necessary. It is still true that a non-limited liable agent will front-load payments as much as possible, and thus we will have that $\tau \leq w_0^{-1}(\bar{x}V)$ (instead of in the limited-liable case, where they are equal). The rest of the proof proceeds as before.

For example, in the extreme case where the buyer can pay as much as they want at the very beginning, they will simply pay V/e at the very beginning and receive $(1 - 1/e)V$ total utility (in contrast to if they were limited-liable, in which case they receive V/e utility).

Remark 4. The discrete version of this mechanism is presented in the online appendix. As expected, it works essentially equivalently to the continuous mechanism of Theorem 6. Two small differences are: i) instead of guaranteeing a revenue of V/e , it only guarantees a revenue of $V/e - O(1)$ (this additive loss is inevitable due to the fact that in the discrete case we cannot guarantee a payment in the very last round), and ii) it employs a small $O(1)$ reimbursement at the end of the protocol to guarantee truthfulness.

3.3 Upper bound of $1/e$

In this section we show that no incentive compatible mechanism can extract more than $1/e$ of the welfare of the bidder. Throughout this section, we will assume without loss of generality that the total value V of the bidder is normalized to 1.

We will need the following auxiliary lemma.

LEMMA 7. Fix an $\alpha \in [0, 1]$, and let $f : [0, 1] \rightarrow [0, 1]$ be a function satisfying

$$\left(1 - \int_0^x \frac{dx'}{f(x')} \right) f(x) \leq \alpha$$

for all $x \in [0, 1]$. Then

$$\int_0^{\alpha \log(1/\alpha)} \frac{dx}{f(x)} \geq 1 - \alpha.$$

PROOF. Define $g(x) = 1/f(x)$, so $g(x)$ satisfies

$$g(x) \geq \frac{1}{\alpha} - \frac{1}{\alpha} \left(\int_0^x g(x') dx' \right)$$

Let $\lambda(x) = \alpha e^{x/\alpha}$. Multiplying the above inequality by $\lambda(x)$ and integrating it from 0 to $\alpha \log(1/\alpha)$, we get that

$$\begin{aligned} \int_0^{\alpha \log(1/\alpha)} \lambda(x) g(x) dx &\geq \frac{1}{\alpha} \int_0^{\alpha \log(1/\alpha)} \lambda(x) dx - \frac{1}{\alpha} \int_0^{\alpha \log(1/\alpha)} \lambda(x) \int_0^x g(x') dx' dx \\ &= \frac{1}{\alpha} \int_0^{\alpha \log(1/\alpha)} \lambda(x) dx - \frac{1}{\alpha} \int_0^{\alpha \log(1/\alpha)} g(x) \int_x^{\alpha \log(1/\alpha)} \lambda(x') dx' dx. \end{aligned}$$

From this, we have that

$$\begin{aligned} 1 - \alpha &= \frac{1}{\alpha} \int_0^{\alpha \log(1/\alpha)} \lambda(x) dx \\ &\leq \int_0^{\alpha \log(1/\alpha)} \lambda(x) g(x) dx + \frac{1}{\alpha} \int_0^{\alpha \log(1/\alpha)} g(x) \int_x^{\alpha \log(1/\alpha)} \lambda(x') dx' dx \\ &= \int_0^{\alpha \log(1/\alpha)} g(x) \left(\lambda(x) + \frac{1}{\alpha} \int_x^{\alpha \log(1/\alpha)} \lambda(x') dx' \right) dx \\ &= \int_0^{\alpha \log(1/\alpha)} g(x) dx, \end{aligned}$$

from which the desired result follows. $\square$

We can now prove our main theorem.

THEOREM 8. *The worst-case revenue of any incentive compatible mechanism is at most $1/e$.*

PROOF. In other words, we must show that for any mechanism $M \leq 1/e$.

For a type v let $U_v = W_v - X_v$ be the total utility this type ends up with at the end of the game if they act truthfully. We will need the following lemma, which shows that if all U_v are bounded by some constant and M is large enough, we can bound U_v by an even smaller constant.

LEMMA 9. *If for all types v , $U_v \leq \alpha$ (for some $\alpha \in [0, 1]$), then for all types v , $U_v \leq f(\alpha) - M$, where*

$$f(\alpha) = \alpha \log(1/\alpha) + \alpha$$

PROOF. Fix any type v and any $w \in [0, W_v]$; let $t_v(w) = w_v^{-1}(w)$. Fix a small $\varepsilon > 0$, and consider a new type, which we will call $\text{spike}(w, \varepsilon)$ with value equal to $v(t)$ for $t \leq t_v(w)$, value equal to 0 for $t \geq t_v(w) + \varepsilon$, and a constant value $v_{\text{spike}} = \frac{1}{\varepsilon} \int_{t_v(w)}^{t_v(w) + \varepsilon} v(t) dt$ within the interval $[t_v(w), t_v(w) + \varepsilon]$ (here v_{spike} is chosen so that the total value of this type equals 1).

This type can deviate as follows: it can pretend to be v until time $t_v(w) + \varepsilon$ and then abort. By performing this deviation, the type $\text{spike}(w, \varepsilon)$ receives utility at least the utility it gains during this spike, so

$$\begin{aligned}
 U'_{\text{spike}(w, \varepsilon)} &\geq v_{\text{spike}} \int_{t_v(w)}^{t_v(w)+\varepsilon} r_v(t) dt - \int_{t_v(w)}^{t_v(w)+\varepsilon} x_v(t) dt \\
 &= \left(\int_{t_v(w)}^1 v(t) dt \right) \left(\frac{1}{\varepsilon} \int_{t_v(w)}^{t_v(w)+\varepsilon} r_v(t) dt \right) - \int_{t_v(w)}^{t_v(w)+\varepsilon} x_v(t) dt.
 \end{aligned}$$

On the other hand, by dynamic incentive-compatibility it holds that the utility the spike receives if it plays truthfully, $U_{\text{spike}(w, \varepsilon)}$, is at least $U'_{\text{spike}(w, \varepsilon)}$. Since by assumption $U_{\text{spike}(w, \varepsilon)} \leq \alpha$, this means

$$\alpha \geq \left(\int_{t_v(w)}^1 v(t) dt \right) \left(\frac{1}{\varepsilon} \int_{t_v(w)}^{t_v(w)+\varepsilon} r_v(t) dt \right) - \int_{t_v(w)}^{t_v(w)+\varepsilon} x_v(t) dt.$$

Since $r_v(t)$ is continuous, taking the limit as $\varepsilon \rightarrow 0$, the term $\left(\frac{1}{\varepsilon} \int_{t_v(w)}^{t_v(w)+\varepsilon} r_v(t) dt \right)$ converges to $r_v(t_v(w)) = r_v(w)$, and the term $\int_{t_v(w)}^{t_v(w)+\varepsilon} x_v(t) dt$ converges to 0. It follows that

$$\alpha \geq \left(\int_{t_v(w)}^1 v(t) dt \right) r_v(w). \quad (3)$$

We'll now express the integral in (3) in terms of w . First, note that since $\int_0^1 v(t) dt = 1$,

$$\int_{t_v(w)}^1 v(t) dt = 1 - \int_0^{t_v(w)} v(t) dt.$$

We will perform a substitution to express $\int_0^{t_v(w)} v(t) dt$ as an integral in terms of w . Since $w_v(t) = \int_0^t r_v(s) v(s) ds$, $\frac{dw_v(t)}{dt} = r_v(t) v(t)$ and $dt = \frac{dw_v(t)}{r_v(t) v(t)}$. It follows then that

$$\int_0^{t_v(w)} v(t) dt = \int_0^w \frac{dw'}{r_v(w')}. \quad (4)$$

Substituting this into (3), we have that:

$$\alpha \geq \left(1 - \int_0^w \frac{dw'}{r_v(w')} \right) r_v(w). \quad (5)$$

Note that by Lemma 7, (5) implies that

$$\int_0^{\alpha \log(1/\alpha)} \frac{dw}{r_v(w)} \geq 1 - \alpha. \quad (6)$$

Now, since $\int_0^1 v(t) dt = 1$ and $w(1) = W_v$, substituting $w = W_v$ into (4) we have that

$$\int_0^{W_v} \frac{1}{r_v(w)} dw = 1. \quad (7)$$

If $W_v > \alpha \log(1/\alpha)$, we therefore have that

$$\begin{aligned}
 1 &= \int_0^{W_v} \frac{1}{r_v(w)} dw \\
 &= \int_0^{\alpha \log(1/\alpha)} \frac{1}{r_v(w)} dw + \int_{\alpha \log(1/\alpha)}^{W_v} \frac{1}{r_v(w)} dw \\
 &\geq (1 - \alpha) + (W_v - \alpha \log(1/\alpha)).
 \end{aligned}$$

where here we've applied (6) and the fact that $r_v(w) \leq 1$. For this to be true, we must have

$$W_v \leq \alpha \log(1/\alpha) + \alpha. \quad (8)$$

On the other hand, if $W_v \leq \alpha \log(1/\alpha)$, W_v also satisfies (8). Therefore in either case, W_v is at most $\alpha \log(1/\alpha) + \alpha = f(\alpha)$.

Since $U_v = W_v - X_v$, and since $X_v \geq M$ for all types v (by definition of M), it follows that $U_v \leq f(\alpha) - M$, as desired. $\square$

Now, note that if $f(\alpha) - M < \alpha$ for all $\alpha \in [0, 1]$, repeatedly applying Lemma 9 (starting from $\alpha = 1$) will eventually imply $U_v < 0$ for all types v . But this is clearly not possible in any incentive compatible mechanism (any type could improve their situation by immediately deviating and not paying).

We claim that $f(\alpha) - M < \alpha$ for all $\alpha \in [0, 1]$ if $M > 1/e$, thus completing the proof. One way to see this is to show that the maximum value of $f(\alpha) - \alpha$ on the interval $[0, 1]$ is at most $1/e$. But $f(\alpha) - \alpha = \alpha \log(1/\alpha)$, which is maximized at $\alpha = 1/e$ and has a maximum value of $1/e$, as desired. $\square$

4 SELLING TO MULTIPLE BIDDERS

In this section, we apply the results from the single buyer problem to the problem of selling to multiple bidders.

We begin (in Section 4.1) by considering the case when there are only two bidders. We show that even with just two bidders, it is impossible to get arbitrarily close to the benchmark of selling the business (in particular, we show that no algorithm can achieve a better competitive ratio than $2/e$). On the other hand we show that a constant competitive-ratio is achievable: by adapting the pay-to-play mechanism from the single buyer game, we can get a $(1/2e)$ -competitive mechanism.

For k bidders, this same strategy results in a $(1/ke)$ -competitive mechanism. We show how to improve this by adding competition in the form of a first-price auction. The resulting mechanism is no longer truthful, but has the guarantee that it is $\Theta(1)$ -competitive as long as all bidders play non-dominated strategies.

4.1 Two bidders

We begin by showing that there is no limited liability mechanism which is guaranteed to get more than a factor of $2/e$ of the benchmark of selling the business.

THEOREM 10. *There is no truthful mechanism for two bidders which is α -competitive against selling the business for $\alpha > 2/e$.*

PROOF SKETCH. Fix a $v > 0$ and a sufficiently large T . Let $V = vT$. From our upper bound (Theorem 8), we know that for any truthful mechanism for the single bidder game, there exists a value profile $\mathbf{v}$ with $\sum_t v_t = V$ such that our mechanism receives revenue at most V/e on this profile. By the minimax theorem, this means there exists a distribution μ over value profiles $\mathbf{v}$ with $|\mathbf{v}| = V$ such that no mechanism receives expected revenue more than V/e on a value profile sampled from this distribution. (There is a technical issue here in that von Neumann's minimax theorem requires finite action spaces, but here there are an uncountably infinite set of valuations and protocols. We address this by performing appropriate discretizations – the full detailed proof is in the online appendix.)

Consider the following distribution $\mathcal{D}$ of instances for the two bidder game. All instances will have $2T$ rounds. The valuation profile $\mathbf{v}_1$ of the first bidder will be sampled so that $(v_{1,1}, \dots, v_{1,T})$ is

sampled according to μ and $v_{1,t} = 0$ for $t > T$. Likewise, v_2 will be sampled so that $(v_{2,T+1}, \dots, v_{2,2T})$ is (independently from v_1) sampled according to μ and so that $v_{2,t} = 0$ for $t \leq T$.

Assume to the contrary that there exists a truthful mechanism $\mathcal{M}$ for two bidders which is α -competitive for some $\alpha > 2/e$. Since $|v_1| = |v_2| = V$ for all instances v in the support of $\mathcal{D}$, $\text{Rev}_{STB}(v) = V$, and this means that this mechanism receives expected revenue at least αV over instances sampled from $\mathcal{D}$.

Now, note that from the mechanism $\mathcal{M}$ we can construct two mechanisms $\mathcal{M}_1$ and $\mathcal{M}_2$ for the single-bidder game. To construct $\mathcal{M}_1$, after soliciting the bidder's valuation v , additionally sample a dummy valuation v' from μ . Consider running $\mathcal{M}$ on the valuation profile (v, v') : since v and v' have disjoint supports, $\mathcal{M}$ chooses some sequence of allocations / payments for the first bidder for the first T rounds and some sequence of allocations / payments for the second bidder for the second T rounds. Use this sequence of allocations / payments for the first T rounds as the allocations / payments for $\mathcal{M}_1$; since $\mathcal{M}$ is truthful, it follows that $\mathcal{M}_1$ is also truthful. We construct $\mathcal{M}_2$ symmetrically.

By construction, the expected revenue of $\mathcal{M}$ on instances drawn from $\mathcal{D}$ is the same as sum of the expected revenues of mechanisms $\mathcal{M}_1$ and $\mathcal{M}_2$ on instances drawn from μ . Since $\mathcal{M}$ receives expected revenue at least αV , this means that either $\mathcal{M}_1$ or $\mathcal{M}_2$ must receive expected revenue at least $(\alpha/2)V$ on instances from μ . But $(\alpha/2)V > V/e$, contradicting the fact that no mechanism can receive expected revenue larger than V/e from μ . The theorem statement follows. $\square$

Note that since we can always embed the 2 bidder game in the k bidder game for any $k > 2$, we have the immediate corollary that there is no α -competitive mechanism for k bidders with $\alpha \geq 2/e$.

COROLLARY 11. *For any $\varepsilon > 0$ and integer $k \geq 2$, there is no mechanism for k ε -truthful bidders which is α -competitive against selling the business for $\alpha > 2/e$.*

PROOF. See the online appendix. $\square$

On the other hand, we can reuse the pay-to-play mechanism of Section 3.2 to get a $(1/2e)$ -competitive mechanism. This mechanism works as follows:

- (1) At the beginning of the game, ask both players for their total values $V_1 = \sum_t v_{1,t}$ and $V_2 = \sum_t v_{2,t}$.
- (2) Define

$$\rho(w) = \begin{cases} e^{ew-1} & \text{if } w \in [0, 1/e] \\ 1 & \text{if } w \in [1/e, \infty) \end{cases}$$

(Recall that this is the rate function for the single buyer mechanism in Theorem 6).

- (3) In a given round, let X_1 and X_2 be the total payments of players 1 and 2 thus far. Allocate fraction $\rho(2X_1/V_2)/2$ of the item to player 1, and allocate fraction $\rho(2X_2/V_1)/2$ of the item to player 2. Note that since $\rho(x) \leq 1$, our allocation is valid.

THEOREM 12. *This mechanism for 2 bidders is $(1/2e)$ -competitive mechanism against selling the business.*

PROOF. Assume without loss of generality that $V_1 \geq V_2$. We therefore must show that this mechanism obtains revenue $V_2/(2e)$.

From the perspective of player 1, they are playing the single buyer mechanism for half of the item each round (for which their total valuation is $V_1/2$) with the lower bound on their value of $V_2/2$. Since $V_1/2 \geq V_2/2$, the single buyer mechanism guarantees that player 1 will pay at least $(1/e) \cdot (V_2/2) = V_2/(2e)$, as desired. $\square$

4.2 More than two bidders

What approximation factor can we achieve when we have $k > 2$ bidders? Adapting the mechanism of Theorem 12 (by splitting the item into k pieces and running the single buyer mechanism on each piece with the highest other reported value), we can obtain an approximation factor of $1/ek$. In this section, we will show that it is in fact possible to achieve a constant approximation factor independent of k . Unlike the other mechanisms in this paper, this mechanism is not a direct revelation mechanism – nonetheless, we will show that as long as bidders follow non-dominated strategies, we will receive expected revenue within a constant factor of $\text{Rev}_{STB}(\mathbf{v})$.

4.2.1 Selling shares via a first-price auction. Assume, to begin, that we know any lower bound V^* on the largest total value of any bidder. We will first show how to construct a mechanism that guarantees we obtain revenue at least a constant fraction of V^* . This mechanism will also have the property that it requires reimbursements of size $O(T)$ at the end of the protocol. Note that this is not inherently at odds with limited liability: our protocol will still never expect a bidder to pay more than their value in any given round. Nevertheless, we will show how to remove this constraint in Section 4.2.3.

Roughly, our mechanism proceeds as follows. Each round we will split the item into two halves, which we will allocate in different ways. One half of the item we will allocate via a first-price auction among all k bidders (each bidder will submit one bid per round in advance). We will allocate a fraction $\rho(X_i) = 10(X_i/V^*)$ of the other half of the item (i.e. $\rho(X_i)/2$ of the item) to bidder i , where X_i is the total payment of bidder i up to this round in the first-price auction.

Now, note that it is possible for $\sum_i \rho(X_i) > 1$; this would make it impossible to allocate the second half of the item as described above. To get around this, as soon as $\sum_i \rho(X_i) = 1$ we will freeze the allocation rates as is. That is, we will continue allocating the first half of the item to the bidder with the highest bid via a first-price auction, but stop updating the values X_i . Finally, at the end of the protocol, we will reimburse each bidder for the amount they paid in the first-price auction after the allocation rates have been fixed.

More formally, our mechanism operates as follows:

- (1) At the beginning of the mechanism, the seller asks each bidder i to (simultaneously) report their desired bid $b_{i,t}$ in the first price auction at time t .
- (2) In round t , the seller splits the item into two equal halves which they allocate separately. We will keep track of a value $X_i(t)$ for each bidder i representing the amount they have paid the seller up until round t before the allocation cap is hit. A fraction $\rho(X_i(t)) = 10(X_i(t)/V^*)$ of the first half (the *allocation half*) is allocated to bidder i . We will guarantee $\sum_i \rho(X_i(t)) \leq 1$ for all t .
- (3) The second half of the item (the *auction half*) is allocated to the bidder $i = \arg \max_i b_{i,t}$ (if multiple bidders share the same highest bid, split it evenly among them). This bidder is expected to pay $b_{i,t}$ this round (if they do not, we refuse to ever allocate to them in the future). If this causes $\sum_i \rho(X_i(t))$ to exceed 1, decrease the payment until it exactly equals 1. Note that if $\sum_i \rho(X_i(t)) = 1$, the payment in this round does not contribute towards $X_i(t)$.
- (4) Finally, at the end of the protocol, as long as they successfully paid their bids when they were chosen in step (3), we provide each bidder a reimbursement equal to the amount they paid the seller after the allocation cap was hit (i.e. after the first round where $\sum_i \rho(X_i(t)) = 1$). In addition, provide each bidder any $O(1)$ reimbursement $g > 1$ in case the other reimbursement is 0. This makes it dominant for each bidder to only report bids $b_{i,t}$ which they have the ability of paying (i.e. $b_{i,t}$ satisfying $b_{i,t} \leq v_{i,t}/2$) and also to participate for the entire protocol.

Note that (for this mechanism) a pure strategy for bidder i consists of their choice of $b_{i,t}$ to report at the beginning of the mechanism, along with a rule for when to defect from the mechanism. Recall that we write s_i to denote a pure strategy for bidder i , s_{-i} to denote a profile of pure strategies for the other $k - 1$ bidders, and $U_i(s_i, s_{-i})$ the utility received by bidder i when they play strategy s_i and the other bidders play s_{-i} . We call a pure strategy s_i weakly-dominated if there exists a (possibly mixed) strategy σ'_i such that $U_i(s_i, s_{-i}) \leq U_i(\sigma'_i, s_{-i})$ for all s_{-i} , and if there exists at least one choice of s_{-i} such that $U_i(s_i, s_{-i}) < U_i(\sigma'_i, s_{-i})$; a strategy is non-dominated if it is not weakly-dominated (and a mixed strategy is non-dominated if it is supported on non-dominated strategies).

For example, note that since we offer a reimbursement of at least $g > 1$ at the end of the protocol, any strategy which defects (doesn't pay $b_{i,t}$ when they are expected to) is dominated; by not paying $b_{i,t}$ they possibly gain $b_{i,t}$ but lose g (plus any other utility they may have obtained) for a net negative loss. In other words, all non-dominated strategies will never defect.

We show that, as long as all bidders are playing non-dominated strategies, the above mechanism extracts a constant fraction of V^* .

THEOREM 13. *Assume we are given a lower bound V^* on the largest total value belonging to any bidder. Then, as long as all bidders play non-dominated strategies, the above mechanism achieves revenue at least $\alpha V^* - O(1)$ for some constant $\alpha > 0$ (independent of the number of bidders k).*

The full proof of Theorem 13 can be found in the online appendix. The main intuition behind the proof is as follows. Without loss of generality, assume that bidder 1 is the bidder with the highest total value (and so in particular, $V_1 \geq V^*$). We argue that for “sufficiently early” rounds t , if bidder 1 ever bids less than $v_{1,t}/2$ (their value for the auction half of the item this round), then this strategy is dominated by the strategy where they bid $v_{1,t}/2$. This follows from the construction of the mechanism: increasing their bid this round costs bidder 1 a little bit more if they win, but this is more than paid back by the value they obtain from the allocation half of the item over the remaining rounds. In particular, we can think of this as the bidder paying one unit of cost to obtain a $5/V^*$ share of the remaining value of the item; if the bidder values the remaining rounds at $V^*/5$ or more, the bidder should be happy to pay this amount (and in fact, the bidders in the mechanism need to compete for the ability to buy these shares in a first-price auction).

Since $V_1 \geq V^*$, there are many rounds where the bidder values the remainder of the item (i.e. the total value of the item across the remaining rounds) at over $V^*/5$; in fact, this happens for at least $4V_1/5$ of the bidder's welfare. If the auction is active for all these rounds, then it generates revenue at least $2V_1/5$ if bidder 1 is playing a non-dominated strategy (and is thus bidding at least half their value per round). But if the auction ever stops, this means that $\sum_i \rho(X_i) = 1$, and therefore that $\sum_i X_i \geq V^*/10$, so we have received $V^*/10$ in revenue. In either case, we have received at least a constant fraction of V^* (since $2V_1/5 \geq 2V^*/5$).

Why do we reimburse each bidder for the amount they paid after the allocation rates have been fixed? Note that if we do not do this, then bidders may be incentivized to shade their bids for rounds where they think the cap has been hit, and we can no longer argue that they will bid their full value. Providing this reimbursement allows us to encourage bidders to bid their full value for the item and incentivizes them to hit the cap.

Remark 5. In fact, Theorem 13 does not require all k bidders to play non-dominated strategies – it merely requires the bidder with largest total value to play a non-dominated strategy.

Similarly, although we assume bidders cannot bid above their true value, note that the mechanism in Theorem 13 is still α -competitive even if some bidders can bid above their true value: it is still in a bidder's interest to bid at least their value in early rounds of the first price auction (and in fact, it might be in their interest to bid above their value, if they can afford to do so).

Remark 6. This mechanism has the additional caveat that, unlike our previous mechanisms (and the $O(1/k)$ -competitive generalization), it does require the bidders to know their individual per round values in advance (at least for Theorem 13 to hold). Since bidders' bids can affect the utility of other bidders, allowing bidders to bid dynamically can lead to undominated grim trigger strategies where bidders punish each other for not bidding according to a specific schedule. It is an interesting open question whether it is possible to adapt this mechanism to remove these constraints.

4.2.2 Soliciting bidders' values. The mechanism of Theorem 13 assumes we have a lower bound V^* on the highest total value of a bidder, and extracts a constant fraction of this bound as revenue. However, our ultimate goal is to construct a completely prior-free mechanism that extracts constant fraction of the second-highest total value $\text{Rev}_{STB}(\mathbf{v})$.

To do so, we will at the very beginning of the mechanism solicit the value of $\text{Rev}_{STB}(\mathbf{v})$ using a similar technique as in our construction of the $k = 2$ algorithm. Specifically, we will do the following at the very beginning of the protocol:

- At the beginning of the mechanism, the seller asks each bidder i to (simultaneously) report their total value V_i . (This should be done simultaneously with the reporting of bids in the mechanism of Theorem 13).
- The seller splits the bidders randomly into two subsets S_{price} and S_{play} . Each bidder is independently assigned to S_{price} with probability $1/2$ and to S_{play} with probability $1/2$. (If either set is empty the seller immediately stops the auction, never allocating the item).
- Define $V^* = \max_{i \in S_{\text{price}}} V_i$. The seller now removes all players in S_{price} from the auction (i.e., the seller will only allocate to players in S_{play} in the future).
- Run the mechanism from Theorem 13 on the bidders in S_{play} with this V^* .

Note that $1/4$ of the time, the bidder with the highest value belongs to S_{play} and the bidder with the second highest value belongs to S_{price} . In this case, $V^* = \text{Rev}_{STB}(\mathbf{v})$ and the conditions of Theorem 13 are satisfied with constant probability (assuming bidders report their values truthfully). Moreover, note that the value V that a bidder reports in this step is completely independent from their eventual utility (in particular, if they are selected to belong to S_{play} , their value V is never used). It is therefore never better for a bidder to report a value V' not equal to their true value V .

If bidders cannot pay more than their value each round, it is possible to make it dominant for each bidder to report their true value V . To do so, with some small probability $\delta > 0$, instead of running the rest of the mechanism after step (1), run the following procedure: allocate $\varepsilon \approx 1/kT$ of the item to each buyer each round. If buyer i has paid a total of εV_i at the end of the mechanism, reimburse them an additional $2\varepsilon V_i$.

If we do this, then bidder i receives a total utility of $\varepsilon V_i - \varepsilon V'_i + 2\varepsilon V'_i = \varepsilon(V_i + V'_i)$ by misreporting V'_i instead of one's true valuation V_i ; to maximize this, a bidder will want to report the largest V'_i they can possibly afford, which is V_i (they cannot report a larger V'_i while paying at most their value every round). It is therefore dominated for a bidder i to report any value $V'_i \neq V_i$.

Our main theorem for selling to k buyers therefore follows.

THEOREM 14. *As long as all bidders play non-dominated strategies, the above mechanism is an α -competitive mechanism for k bidders for some constant $\alpha > 0$ (independent of k).*

4.2.3 Removing large reimbursements. As written, Theorem 14 could require post-protocol reimbursements of up to $O(T)$, since we reimburse bidders for all payments after the allocation rates for the allocation half have been fixed. Although this is not at odds with limited-liability, ideally we could reduce these reimbursements to $O(1)$ as in our other protocols.

Since we reimburse bidders for their payments after the allocation cap has been hit, one simple idea is to simply remove these payments. In other words, after the allocation cap has been hit, still

allocate the item to the bidder with the highest bid, but charge them nothing. This largely works, but has the unfortunate side effect that now bidders might gamble and submit bids larger than their value for later rounds where they think the allocation cap has already been hit. Note that this was originally not an issue because in the first-price auction a bidder had to immediately pay their bid if they won the item, thus checking that they in fact could afford it. But now, it might be strategic for a bidder to report a higher bid than their value for a late round hoping that the cap has been met (and if the cap has not been met, defecting at that point).

Nonetheless, we can still address this in the following (admittedly contrived) way. The key idea is to spot check that bidders are capable of paying the bids they have reported. More specifically, for some uniformly randomly selected constant fraction (e.g. 50%) of the rounds, instead of running the mechanism described above, allocate a tiny ε of the item (for some $\varepsilon \approx 1/kT$) to each bidder and ask them to pay 2ε times their bid (since their bid was for half the item, this should be at most ε times their value). If they fail to do this, eject them from the protocol.

We then can show that with this change, it is dominated for a limited-liability bidder to bid higher than their value. In particular, the expected increase in utility they can possibly get from bidding above their value is outweighed by the expected loss in future utility (including the $O(1)$ reimbursement at the end) which they would lose from being forced to defect.

COROLLARY 15. *It is possible to implement Theorem 14 with reimbursements of size $O(1)$.*

Details and the full proof of Corollary 15 can be found in the online appendix.

5 ACKNOWLEDGMENTS

M.B. is supported in part by the NSF Alan T. Waterman Award, Grant No. 1933331, a Packard Fellowship in Science and Engineering, and the Simons Collaboration on Algorithms and Geometry. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author and do not necessarily reflect the views of the National Science Foundation.

S.M.W is supported by NSF CAREER Award CCF-1942497.

REFERENCES

- [1] Shipra Agrawal, Constantinos Daskalakis, Vahab S. Mirrokni, and Balasubramanian Sivan. 2018. Robust Repeated Auctions under Heterogeneous Buyer Behavior. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*. 171.
- [2] Itai Ashlagi, Constantinos Daskalakis, and Nima Haghpahan. 2016. Sequential Mechanisms with Ex-post Participation Guarantees. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*. 213–214.
- [3] Susan Athey and Kyle Bagwell. 2001. Optimal Collusion with Private Information. *RAND Journal of Economics* 32, 3 (2001), 428–65. <http://EconPapers.repec.org/RePEc:rje:randje:v:32:y:2001:i:3:p:428-65>
- [4] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. 2014. A Simple and Approximately Optimal Mechanism for an Additive Buyer. In *the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- [5] Santiago R. Balseiro, Vahab S. Mirrokni, and Renato Paes Leme. 2017. Dynamic Mechanisms with Martingale Utilities. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*. 165.
- [6] David P. Baron and David Besanko. 1984. Regulation and information in a continuing relationship. *Information Economics and Policy* 1, 3 (1984), 267 – 302.
- [7] Dirk Bergemann and Philipp Strack. 2015. Dynamic revenue maximization: A continuous time approach. *Journal of Economic Theory* 159, Part B (2015), 819 – 853. Symposium Issue on Dynamic Contracts and Mechanism Design.
- [8] Dirk Bergemann and Juuso Valimäki. 2018. *Dynamic Mechanism Design: An Introduction*. Cowles Foundation Discussion Papers 2102R. Cowles Foundation for Research in Economics, Yale University.
- [9] David Besanko. 1985. Multi-period contracts between principal and agent with adverse selection. *Economics Letters* 17, 1–2 (1985), 33 – 37.

- [10] Mark Braverman, Jieming Mao, Jon Schneider, and S Matthew Weinberg. 2019. Multi-armed Bandit Problems with Strategic Arms. In *Conference on Learning Theory*. 383–416.
- [11] Yang Cai, Nikhil Devanur, and S. Matthew Weinberg. 2016. A Duality Based Unified Approach to Bayesian Mechanism Design. In *Proceedings of the 48th ACM Conference on Theory of Computation (STOC)*.
- [12] Yang Cai and Mingfei Zhao. 2016. Simple Mechanisms for Subadditive Buyers via Duality. *Manuscript* (2016).
- [13] Shuchi Chawla, Nikhil R. Devanur, Anna R. Karlin, and Balasubramanian Sivan. 2016. Simple Pricing Schemes For Consumers With Evolving Values. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*. 1476–1490.
- [14] Shuchi Chawla and J. Benjamin Miller. 2016. Mechanism Design for Subadditive Agents via an Ex Ante Relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*. 579–596. <https://doi.org/10.1145/2940716.2940756>
- [15] Ning Chen, Nick Gravin, and Pinyan Lu. 2014. Optimal competitive auctions. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, David B. Shmoys (Ed.). ACM, 253–262. <https://doi.org/10.1145/2591796.2591855>
- [16] Ning Chen, Nikolai Gravin, and Pinyan Lu. 2015. Competitive Analysis via Benchmark Decomposition. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, Tim Roughgarden, Michal Feldman, and Michael Schwarz (Eds.). ACM, 363–376. <https://doi.org/10.1145/2764468.2764491>
- [17] Yuan Deng, Sebastien Lahaie, and Vahab Mirrokni. 2019. Robust pricing in non-clairvoyant dynamic mechanism design. *Available at SSRN* (2019).
- [18] Yuan Deng, Jon Schneider, and Balasubramanian Sivan. 2019. Prior-Free Dynamic Auctions with Low Regret Buyers. In *Advances in Neural Information Processing Systems*. 4804–4814.
- [19] Nikhil R. Devanur, Jason D. Hartline, and Qiqi Yan. 2015. Envy freedom and prior-free mechanism design. *J. Econ. Theory* 156 (2015), 103–143. <https://doi.org/10.1016/j.jet.2014.08.001>
- [20] Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. 2018. Simple versus optimal contracts. *arXiv preprint arXiv:1808.03713* (2018).
- [21] Amos Fiat, Andrew V. Goldberg, Jason D. Hartline, and Anna R. Karlin. 2002. Competitive generalized auctions. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, John H. Reif (Ed.). ACM, 72–81. <https://doi.org/10.1145/509907.509921>
- [22] Andrew V. Goldberg, Jason D. Hartline, Anna R. Karlin, and Michael E. Saks. 2004. A Lower Bound on the Competitive Ratio of Truthful Auctions. In *STACS 2004, 21st Annual Symposium on Theoretical Aspects of Computer Science, Montpellier, France, March 25-27, 2004, Proceedings (Lecture Notes in Computer Science)*, Volker Diekert and Michel Habib (Eds.), Vol. 2996. Springer, 644–655. https://doi.org/10.1007/978-3-540-24749-4_56
- [23] Andrew V. Goldberg, Jason D. Hartline, Anna R. Karlin, Michael E. Saks, and Andrew Wright. 2006. Competitive auctions. *Games Econ. Behav.* 55, 2 (2006), 242–269. <https://doi.org/10.1016/j.geb.2006.02.003>
- [24] Andrew V. Goldberg, Jason D. Hartline, and Andrew Wright. 2001. Competitive auctions and digital goods. In *Proceedings of the Twelfth Annual Symposium on Discrete Algorithms, January 7-9, 2001, Washington, DC, USA*, S. Rao Kosaraju (Ed.). ACM/SIAM, 735–744. <http://dl.acm.org/citation.cfm?id=365411.365768>
- [25] Jason D Hartline. 2013. Mechanism design and approximation. *Book draft. October* 122 (2013).
- [26] Sham M Kakade, Ilan Lobel, and Hamid Nazerzadeh. 2013. Optimal dynamic mechanism design and the virtual-pivot mechanism. *Operations Research* 61, 4 (2013), 837–854.
- [27] Siqi Liu and Christos-Alexandros Psomas. 2018. On the competition complexity of dynamic mechanism design. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2008–2025.
- [28] Vahab S. Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. 2016. Dynamic Auctions with Bank Accounts. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*. 387–393.
- [29] Vahab S. Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. 2016. Optimal dynamic mechanisms with ex-post IR via bank accounts. *CoRR abs/1605.08840* (2016). [arXiv:1605.08840](http://arxiv.org/abs/1605.08840) <http://arxiv.org/abs/1605.08840>
- [30] Vahab S. Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. 2018. Non-clairvoyant Dynamic Mechanism Design. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*. 169.
- [31] Christos Papadimitriou, George Pierrakos, Christos-Alexandros Psomas, and Aviad Rubinfeld. 2016. On the Complexity of Dynamic Mechanism Design. In *Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '16)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1458–1475. <http://dl.acm.org/citation.cfm?id=2884435.2884535>
- [32] Andrew Chi-Chih Yao. 2015. An n-to-1 bidder reduction for multi-item auctions and its applications. In *the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.