

Dynamic Auctions with Bank Accounts*

Vahab Mirrokni and Renato Paes Leme

Google Research, NY, USA
 {mirrokni, renatopl}@google.com

Pingzhong Tang and Song Zuo

Institute for Interdisciplinary Information Sciences,
 Tsinghua University, Beijing, China
 {kenshinping, songzuo.z}@gmail.com

Abstract

Consider a buyer with independent additive valuations for a set of goods, and a seller who is constrained to sell one item at a time in an online fashion. If the seller is constrained to run independent auctions for each item, then he would run Myerson's optimal auction for each item. If the seller is allowed to use the full power of dynamic mechanism design and have the auction for each item depend on the outcome of the previous auctions, he is able to perform much better. The main issues in implementing such strategies in online settings where items arrive over time are that the auction might be too complicated or it makes too strong assumptions on the buyer's rationality or seller's commitment over time. This motivates us to explore a restricted family of dynamic auctions that can be implemented in an online fashion and without too much commitment from the seller ahead of time. In particular, we study a set of auction in which the space of single-shot auctions is augmented with a structure that we call *bank account*, a real number for each node that summarizes the history so far. This structure allows the seller to store deficits or surpluses of buyer utility from each individual auction and even them out on the long run. This is akin to enforcing individual rationality constraint on average rather than per auction. We also study the effect of enforcing a maximum limit to the values that bank account might grow, which means that we enforce that besides the auction being individually rational on average it is also not far from being individually rational at any given interval. Interestingly, even with these restrictions, we can achieve significantly better revenue and social welfare compared to separate Myerson auctions.

*This work was supported by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the Natural Science Foundation of China Grant 61033001, 61361136003, 61303077, 61561146398, a Tsinghua Initiative Scientific Research Grant and a China Youth 1000-talent program. We acknowledge anonymous reviewers for helpful comments.

1 Introduction

The theory of dynamic mechanism design shows that if auctions are repeated over time and buyers and sellers care about the aggregate outcome of all auctions (as opposed to the outcome of each individual auction), then it is possible to design mechanisms that are superior to running the same static single-shot auction in terms of both revenue and social welfare. This is in contrast to the auction formats used by search engines to sell internet advertisement, which are mostly static single-shot auctions. The advantage of static formats is that they are easy to implement and the economic principles behind them are well understood. Nevertheless, by restricting to static mechanisms, exchanges are forgoing possibly very significant gains in revenue.

What are the barriers preventing dynamic mechanisms to become more prevalent in internet advertisement? If the seller is allowed to use the full power of dynamic mechanism design and have the auction for each item depend on the outcome of the previous auctions, such dependency induces a gigantic design space, making it difficult to solve for or even to describe such auctions. This creates problems not only for the mechanism designer but also for buyers participating in the mechanism, who need to play in a mechanism where every decision may lead to unintended repercussions.

Our goal in this paper is to study a restricted family of dynamic mechanisms, called *bank account mechanisms*, in which single-shot auctions are augmented with a structure we call *bank account*, a real number for each node that summarizes the history so far. This structure allows the seller to store deficits or surpluses of buyer utility from each individual auction and even them out on the long run, satisfying individual rationality on average.

In dynamic mechanism design the incentive compatibility and individual rationality constraints are satisfied on average over all queries and not per each individual query. One can think of bank account balance as how far each individual auction is from satisfying per-query incentive constraints. In other words, bank accounts quantify how far a certain dynamic mechanism is from a static mechanism.

This last observation motivates us to introduce the notion of a *bank account limit*, which is the maximum bank account balance allowed. A mechanism with zero limit corresponds to a static mechanism. We identify a trade-off between the bank account limit and the revenue extracted: the revenue extracted

per stage is upper bounded by the Myerson revenue plus the bank account limit (Theorem 3), and is lower bounded by constructed deterministic bank account mechanisms (Theorem 2).

We further identify a subset of bank account mechanisms, called *double-reserve auctions*, that can be easily implemented and can guarantee higher revenue than simply run separate Myerson auctions for each item (Theorem 2), while ensuring the same buyer utility as separate Myerson auction.

We also put forward an algorithm to compute the optimal double-reserve auction with any given bank account limit constraints via dynamic programming. Furthermore, the algorithm admits an FPTAS for any multiplicative ϵ -approximation (Theorem 4).

Finally, we empirically evaluate the revenue performance of a heuristically constructed double-reserve auction and the optimal one on an infinite sequence of i.i.d. items. Moreover, we prove a lower bound of its revenue as a function of the bank account limit for any regular distribution (Theorem 5).

In a parallel work [Zuo *et al.*, 2016], we consider the dynamic mechanism design under *ex-post IR* constraints. We also prove a characterization result concerning sufficiency of bank account mechanism in that setting and establish some simple bank account mechanisms that guarantee constant approximation ratios to the optimal revenue.

Related Work The area of *dynamic mechanism design* has attracted a large body of research work in the last decade [Cavallo, 2008; Pai and Vohra, 2008; Pavan *et al.*, 2008; Cavallo *et al.*, 2009; Gershkov and Moldovanu, 2009; Bergemann and Välimäki, 2010; Gershkov and Moldovanu, 2010; Athey and Segal, 2013; Kakade *et al.*, 2013; Papadimitriou *et al.*, 2014; Pavan *et al.*, 2014]. We refer to [Bergemann and Said, 2011] for a comprehensive survey on the topic and cite here a couple of representative papers: [Bergemann and Välimäki, 2010] study the problem of efficient mechanism design in a dynamic environment where agents receive private information over time, and generalize the idea of pivot mechanisms [Green and Laffont, 1977] to this dynamic setting. [Pavan *et al.*, 2014] provide characterization of dynamic local and global incentive compatibility constraints, and use the characterization to design optimal dynamic mechanisms in Markov environments. However, the optimal mechanism is individually rational only when the initial signal being observed, while in our independent valuation setting, it leads to full surplus extraction from all rest stages. [Papadimitriou *et al.*, 2014] work on the discrete (except for the last stage) and correlated valuation distribution setting, and prove hardness results to compute deterministic optimal mechanism subject to ex post individually rational constraint for the two-stage case.

The problem of optimal multidimensional mechanism design [Manelli and Vincent, 2006; 2007; Cai *et al.*, 2012a; 2012b; Hart and Nisan, 2012; Daskalakis *et al.*, 2013; Li and Yao, 2013; Wang and Tang, 2014; Daskalakis *et al.*, 2015; Cai *et al.*, 2016] is closely related to our approach, especially the *bundling* technique used to improve the expected revenue [Tang and Sandholm, 2012; Babaioff *et al.*, 2014].

2 Preliminaries and Model

Problem Consider the problem for selling a sequence of items, one at each stage, to a buyer who has independent additive valuations for the items. In each stage $t \in [T]$, the type (or valuation) $v_t \in \mathbb{R}_+$ of the buyer is privately drawn from a public distribution \mathcal{F}_t , i.e., $v_t \sim \mathcal{F}_t$. The prior distribution \mathcal{F}_t is independent stage-wise.

The outcome of each stage t is specified by a pair (x_t, p_t) , where $x_t \in [0, 1]$ denotes the probability of item the buyer obtains in stage t and $p_t \in \mathbb{R}_+$ denotes the corresponding payment. The buyer utility (or simply utility) of this stage is

$$u_t(\cdot; v_t) = x_t(\cdot) \cdot v_t - p_t(\cdot).$$

Then the seller's objective is to design some mechanism that maximizes the total revenue, $\sum_{\tau=1}^T p_\tau$. Let $v_{(\tau, \tau')} = v_\tau, \dots, v_{\tau'}$ be the vector of buyer's types from stage τ to τ' . Similar for $x_{(\tau, \tau')} = x_\tau, \dots, x_{\tau'}$, $p_{(\tau, \tau')} = p_\tau, \dots, p_{\tau'}$, etc.

2.1 Dynamic Mechanisms

By revelation principle, it is without loss of generality to restrict attention to *direct mechanisms*, where the agent reports its private type to the mechanism in each stage.

Definition 1 (Direct Mechanism). A *direct mechanism* is a pair of allocation rule and payment rule, i.e., $M = \langle x_{(1, T)}, p_{(1, T)} \rangle$, where $x_t : \mathbb{R}_+^t \rightarrow [0, 1]$, $p_t : \mathbb{R}_+^t \rightarrow \mathbb{R}_+$, $\forall t$.

Furthermore, the allocation $x_{(1, T)}$ and payment $p_{(1, T)}$ satisfy the following constraints,

- **Incentive compatible (IC).** *Incentive compatible in perfect Bayesian equilibrium: for any history, truth-telling is the best response for current and upcoming stages in expectation.*
- **Individually rational (IR).** *Stage-wise interim individually rational: for any history, the expected utility for current and upcoming stages is non-negative.*

Formally, we first define the following notation for the expected buyer utility after stage t , with given history $v_{(1, t)}$ and bidding strategy $b_{(t+1, T)}$, i.e.,

$$U_t^{b_{(t+1, T)}}(v_{(1, t)}) = \mathbb{E}_{v_{(t+1, T)}} \left[\sum_{\tau=t+1}^T u_\tau(v_{(1, \tau)}^*; v_\tau) \right], \quad (1)$$

where $v_{(1, \tau)}^* = v_{(1, t)}, b_{t+1}(v_{t+1}), \dots, b_\tau(v_\tau)$. We use U_t without superscript to denote the utility by truthful bidding, and the formal definitions are $\forall t, v_{(1, t)}, v'_t, b_{(t+1, T)}$,

$$\begin{aligned} \text{IC: } & u_t(v_{(1, t)}; v_t) + U_t(v_{(1, t)}) \\ & \geq u_t(v_{(1, t-1)}, v'_t; v_t) + U_t^{b_{(t+1, T)}}(v_{(1, t-1)}, v'_t) \end{aligned} \quad (2)$$

$$\text{IR: } u_t(v_{(1, t)}; v_t) + U_t(v_{(1, t)}) \geq 0 \quad (3)$$

For for any mechanism M , denote the overall expected revenue by $\text{REV}(M) = \mathbb{E}_{v_{(1, T)}} \left[\sum_{\tau=1}^T p_\tau(v_{(1, \tau)}) \right]$, overall expected buyer utility by $\text{UTL}(M) = U_0$, and overall expected efficiency by $\text{EFF}(M) = \text{UTL}(M) + \text{REV}(M)$.

The mechanism is *deterministic*, if for each stage t , and any history $v_{(1, t)}$, $x_t(v_{(1, t)}) \in \{0, 1\}$; the mechanism is *history independent*, if both x_t and p_t only depend on the type of current stage, v_t , but not $v_{(1, t-1)}$.




Figure 1: Cash flow in BAM. Seller's revenue comes from both Spending (Step 1) and Payment (Step 2). The balance in the bank account belongs to the buyer, but the seller can spend it according to the spending policy s_t .

Mechanism 1 ((Static) Myerson Auction). *Separate Myerson auction M^S is the auction that sells each item separately using single-stage Myerson auction.*

Let $\phi_t \in \operatorname{argmax}_u \rho_t(u)$ be the Myerson price of stage t , where $\rho_t(u) = (1 - \mathcal{F}(u)) \cdot u$. Hence $\operatorname{REV}(M^S) = \sum_{\tau=1}^T \rho_\tau(\phi_\tau)$.

Separate Myerson auction is deterministic and history independent. More generally, we called a mechanism *separate posted-price auction*, if the seller sells each item separately using single-stage posted-price.

2.2 Mechanisms with Bank Accounts

Intuitively, a bank account mechanism (BAM) is a separate auction augmented with a *bank account* to ensure IR on average. Instead of explicitly depending on the full history, a BAM depends on a real number, *bank account balance* (or simply balance).

Definition 2 (BAM). *A bank account mechanism $B = \langle z_{(1,T)}, q_{(1,T)}, d_{(1,T)}, s_{(1,T)}, \operatorname{bal}_{(1,T)} \rangle$, where for each t ,*

- *Allocation rule $z_t : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow [0, 1]$, maps balance and value to allocation.*
- *Payment rule $q_t : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$, maps balance and value to payment.*
- *Deposit policy $d_t : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$, maps balance and value to how much money the buyer needs to add into the bank account.*
- *Spend policy $s_t : \mathbb{R}_+ \rightarrow \mathbb{R}$, maps balance to how much money to spend from the bank account.*
- *Initial balance of stage t , bal_t , is a function of history $v_{(1,t-1)}$, satisfying the follows, where $\operatorname{bal}_1 = 0$.*

$$\forall v_t, \operatorname{bal}_{t+1} = \operatorname{bal}_t - s_t(\operatorname{bal}_t) + d_t(\operatorname{bal}_t, v_t). \quad (4)$$

At stage t , a BAM works as follows, (see Figure 1)

1. The seller spends the buyer's balance by $s_t(\operatorname{bal}_t)$. (So $s_t(\operatorname{bal}_t)$ becomes part of seller's revenue.)
2. Upon receiving buyer's report v_t , the seller allocates the item to the buyer with probability $z_t(\operatorname{bal}_t, v_t)$ and charges $q_t(\operatorname{bal}_t, v_t)$.

3. The buyer deposits an amount of $d_t(\operatorname{bal}_t, v_t)$ to the bank account.

By Definition 2, bank account mechanism is a subset of direct mechanisms. A BAM is *deterministic*, if in each stage $t, \forall v_t, \operatorname{bal}_t, z_t(\operatorname{bal}_t, v_t) \in \{0, 1\}$.

We impose a limit L on the maximum value that the balance can reach, i.e.,

$$L_t = \max_{v_{(1,t)}} \operatorname{bal}_t, \quad L = \max_{t \in [T]} L_t.$$

Intuitively, the larger the limit, the more trust needed from the buyer.

Consider the following example that illustrates the bank account mechanisms.

Example 1. *Consider a two-stage example, where v_1 and v_2 are two i.i.d. random variables which take value 1 or 2 with probability 1/2. A bank account mechanism B is defined as follows.*

$$\begin{aligned} z_1(\operatorname{bal}_1, v_1) &= 1; z_2(0, 1) = 0, z_2(0, 2) = z_2(0.5, v_2) = 1; \\ q_1(\operatorname{bal}_1, v_1) &= 1; q_2(0, 1) = 0, q_2(0, 2) = 2, q_2(0.5, v_2) = 1; \\ d_1(\operatorname{bal}_1, 1) &= 0, d_1(\operatorname{bal}_1, 2) = 0.5; d_2(\operatorname{bal}_2, v_2) = 0; \\ s_1(\operatorname{bal}_1) &= 0; s_2(\operatorname{bal}_2) = \operatorname{bal}_2. \end{aligned}$$

B operates as follows,

1. *The first item is sold at posted-price 1.*
2. *If $v_1 = 2$, the buyer offers a pre-payment of amount 0.5 to the seller before learning v_2 ($d_1(\operatorname{bal}_1, 1) = 0.5, s_2(\operatorname{bal}_2) = \operatorname{bal}_2$).*
3. *The seller sets a posted-price r for the second item. If $\operatorname{bal}_2 = 0, r = 2$; if $\operatorname{bal}_2 = 0.5, r = 1$.*
4. *The buyer learns v_2 , and buys the second item if and only if $v_2 \geq r$.*

The revenue comes from both the payment when the buyer buys the items (step 1 and 4) and the pre-payment (step 2), and is 2.25 in expectation.

2.3 Partial Characterization

The following lemma characterizes sufficient conditions for a BAM to satisfy IC and IR.

Lemma 1. *Let $\hat{u}_t(\operatorname{bal}_t, v'_t; v_t)$ denote the stage utility for bank account mechanism B , i.e.,*

$$\hat{u}_t(\operatorname{bal}_t, v'_t; v_t) = z_t(\operatorname{bal}_t, v'_t) \cdot v_t - q_t(\operatorname{bal}_t, v'_t). \quad (5)$$

B is IC and IR, if the following are satisfied, $\forall t, \operatorname{bal}_t, v_t, v'_t$.

$$\hat{u}_t(\operatorname{bal}_t, v_t; v_t) \geq \hat{u}_t(\operatorname{bal}_t, v'_t; v_t), \quad (6)$$

$$s_t(\operatorname{bal}_t) = \mathbb{E}_{v_t} [\hat{u}_t(\operatorname{bal}_t, v_t; v_t) - \hat{u}_t(0, v_t; v_t)], \quad (7)$$

$$\hat{u}_t(\operatorname{bal}_t, v_t; v_t) \geq 0. \quad (8)$$

In particular, (6)(7) \implies IC; (7)(8) \implies IR.

In the rest of the paper, BAM refers to bank account mechanism satisfying (6)(7)(8) and $q_t \geq 0$.

Proof. Throughout the proof, we use short notations as follows,

$$\begin{aligned} \text{bal}_t &= \text{bal}_t(v_{(1,t-1)}), \text{bal}'_{t+1} = \text{bal}_{t+1}(v_{(1,t-1)}, v'_t), \\ s_t &= s_t(\text{bal}_t), s'_t = s_t(\text{bal}'_t), \\ u_t &= u_t(v_{(1,t)}; v_t), u'_t = u_t(v_{(1,t-1)}, v'_t; v_t). \end{aligned}$$

We construct a direct mechanism M as follows, which has the same allocations and payments with B for any $v_{(1,T)}$.

$$x_t(v_{(1,t)}) = z_t(\text{bal}_t, v_t), p_t(v_{(1,t)}) = q_t(\text{bal}_t, v_t) + s_{t+1}$$

Firstly, we prove that M is IC. According to (5),

$$\hat{u}_t(\text{bal}_t, v'_t; v_t) = u_t(v_{(1,t-1)}, v'_t; v_t) + s'_{t+1}, \quad (9)$$

combining with (6), $\forall t, v_{(1,t-1)}, v_t, v'_t$,

$$u_t(v_{(1,t-1)}, v'_t; v_t) + s'_{t+1} \leq u_t(v_{(1,t)}; v_t) + s_{t+1}. \quad (10)$$

Let v'_t be $b_t(v_t)$. Take expectation to (10) over v_t , and subtract s_t from both sides,

$$\mathbb{E}_{v_t}[u'_t + s'_{t+1} - s_t] \leq \mathbb{E}_{v_t}[u_t + s_{t+1} - s_t]. \quad (11)$$

Denote $u''_{t+1} = u_{t+1}(v_{(1,t-1)}, v'_t, v'_{t+1}; v_{t+1})$, and $s''_{t+2} = s_{t+2}(\text{bal}_{t+2}(v_{(1,t-1)}, v'_t, v'_{t+1}))$, then

$$\begin{aligned} &\mathbb{E}_{v_t} [u'_t + s'_{t+1} - s_t + \mathbb{E}_{v_{t+1}}[u''_{t+1} + s''_{t+2} - s'_{t+1}]] \\ &= \mathbb{E}_{v_t, v_{t+1}} [u'_t + u''_{t+1} + s''_{t+2}] - s_t \end{aligned} \quad (12)$$

Sum over the LHS of (11), and recursively apply (12),

$$\begin{aligned} &\sum_{\tau=t+1}^T \text{LHS}(11) \\ &= U_t^{b^{(t+1,T)}}(v_{(1,t-1)}, v'_t) - s'_{t+1} + \mathbb{E}[s_{T+1}(\text{bal}_{T+1}^*)] \\ &= U_t^{b^{(t+1,T)}}(v_{(1,t-1)}, v'_t) - s'_{t+1}, \end{aligned} \quad (13)$$

where $s_{T+1} \equiv 0$. Similarly, sum over both sides of (11),

$$U_t^{b^{(t+1,T)}}(v_{(1,t-1)}, v'_t) - s'_{t+1} \leq U_t(v_{(1,t)}) - s_{t+1}. \quad (14)$$

Adding (14) to (10) implies (2), so M is IC.

Secondly, we prove that M is IR. By (9) and (7),

$$\begin{aligned} \mathbb{E}_{v_t}[u_t + s_{t+1} - s_t] &= \mathbb{E}_{v_t}[\hat{u}_t(\text{bal}_t, v_t; v_t)] - s_t \\ &= \mathbb{E}_{v_t}[\hat{u}_t(0, v_t; v_t)] \end{aligned}$$

Thus $\forall v_{(1,t)}$, $U_t(v_{(1,t)}) - s_{t+1}$ is a constant, i.e.,

$$U_t(v_{(1,t)}) - s_{t+1} = \sum_{\tau=t+1}^T \mathbb{E}_{v_\tau} [\hat{u}_\tau(0, v_\tau; v_\tau)]. \quad (15)$$

Add (15) to (9),

$$\begin{aligned} u_t + U_t(v_{(1,t)}) &= \hat{u}_t(\text{bal}_t, v_t; v_t) \\ &\quad + \sum_{\tau=t+1}^T \mathbb{E}_{v_\tau} [\hat{u}_\tau(0, v_\tau; v_\tau)], \end{aligned}$$

By (8), the RHS is non-negative, so M is IR by (3).

In summary, for any BAM, B , the equivalent direct mechanism M is IC or IR if the corresponding sufficient conditions are satisfied. \square

2.4 Bank Account Mechanism is WLOG

In this paper, we focus on bank account mechanisms instead of general direct mechanisms. This is formally justified by the following theorem.

Theorem 1. *For any direct mechanism M , there is a constructive BAM, B , such that,*

$$\text{REV}(B) \geq \text{REV}(M), \text{UTL}(B) = \text{UTL}(M).$$

In particular, if M is deterministic, B is also deterministic.

We omit the proof here and refer readers to the full version of our paper.

3 Double-Reserve Auction

By Theorem 1, the revenue optimal bank account mechanism is optimal among any mechanism. An extreme example is that a BAM could extract almost all valuation by allocating the t -th item and charging the buyer $\mathbb{E}[v_{t+1}]$ in the t -th stage for each $t = 1, \dots, T-1$.

However, there are a list of drawbacks of such mechanisms, such as, randomized allocation rule, large bank account, and pre-payment (positive payment even not get allocated).

In this paper, we introduce a subset of practical bank account mechanisms coined *double-reserve auction*, which in practice can be easily implemented as a special posted-price auction: for each stage t , there are a low reserve price and a high reserve price; if the $(t-1)$ -th item is sold to the buyer, sell the t -th item at the low reserve price, otherwise at the high price.

Mechanism 2 (Double-Reserve Auction). *A double-reserve auction B is defined upon a separate posted-price auction M and pre-specified bank account limits, $\{L_t\}_{t=1}^T$. Let $r_{(1,T)}$ be a series of auxiliary functions with $r_t(0)$ being the reserve price of M on stage t . Denote $\text{Val}_t(a, b) = \int_a^b (1 - \mathcal{F}_t(v)) dv$. Then B is formally defined as, where \mathcal{I} is indicator function.*

$$z_t(\text{bal}_t, v_t) = \mathcal{I}[v_t \geq r_t(s_t(\text{bal}_t))] \quad (16)$$

$$q_t(\text{bal}_t, v_t) = z_t(\text{bal}_t, v_t) \cdot r_t(s_t(\text{bal}_t)) \quad (17)$$

$$s_t(\text{bal}_t) = \text{bal}_t \quad (18)$$

$$d_t(\text{bal}_t, v_t) = z_t(\text{bal}_t, v_t) \cdot \bar{L}_{t+1} \quad (19)$$

$$\bar{L}_t = \min\{L_t, \text{Val}_t(0, r_t(0))\}$$

The function $r_t(\xi)$ above is defined to meet the requirements of Lemma 1, particularly, $r_t(\xi)$ equals to the value such that $\text{Val}_t(r_t(\xi), r_t(0)) = \xi$.¹

3.1 Revenue Properties of Bank Account Mechanism with Limits

We now study the revenue properties of double-reserve auction and general bank account mechanism with limits. Concretely, for double-reserve auction, we prove a revenue lower bound (Theorem 2); for general bank account mechanisms, we prove a revenue upper bound (Theorem 3).

Theorem 2. *For any separate posted-price mechanism M , the double-reserve auction B defined upon M with any given bank account limits, $\{L_t\}_{t=1}^T$, satisfies*

¹ $r_t(\xi)$ is unique, non-negative, and strictly decreasing.

- B is a BAM (satisfying Lemma 1);
- $\text{UTL}(B) = \text{UTL}(M)$;
- $\text{REV}(B) \geq \text{REV}(M)$;
- $\text{REV}(B)$ is increasing in $L_{(1,T)}$.

Proof. **1. B is a BAM (satisfying Lemma 1).**

By (16), B is deterministic. Meanwhile, by (18), balance is always completely spent at the beginning of each stage. Thus by (19), bal_t only has two possible values, 0 or \bar{L}_t ($\text{bal}_1 = 0$ by definition).

Since $\bar{L}_t \leq \text{Val}_t(0, r_t(0))$, $r_t(\text{bal}_t)$ is well-defined. Moreover, according to the proof of Lemma 1, B can be implemented by a direct mechanism with stage payment as follows,

$$\begin{aligned} p_t(v_{(1,t)}) &= q_t(\text{bal}_t, v_t) + s_{t+1}(\text{bal}_{t+1}) \\ &= z_t(\text{bal}_t, v_t) \cdot (r_t(s_t(\text{bal}_t)) + \bar{L}_{t+1}). \end{aligned}$$

In particular, when the item is sold, the payment can be two different values, depending on whether bal_t is 0.

Condition (6) and (8) are satisfied by (16) and (17), while (7) is ensured as follows.

$$\begin{aligned} \mathbb{E}_{v_t}[\hat{u}_t(\text{bal}_t, v_t; v_t)] &= \int_{r_t(\text{bal}_t)}^{\infty} (v - r_t(\text{bal}_t)) d\mathcal{F}_t(v) \\ &= \int_{r_t(\text{bal}_t)}^{\infty} (1 - \mathcal{F}_t(v)) dv = \text{Val}_t(r_t(\text{bal}_t), +\infty) \\ \implies \mathbb{E}_{v_t}[\hat{u}_t(\text{bal}_t, v_t; v_t) - \hat{u}_t(0, v_t; v_t)] \\ &= \text{Val}_t(r_t(\text{bal}_t), +\infty) - \text{Val}_t(r_t(0), +\infty) = s_t(\text{bal}_t) \end{aligned}$$

The second “=” is from integration by parts.

2. $\text{UTL}(B) = \text{UTL}(M)$.

$$\begin{aligned} \text{UTL}(B) &\stackrel{(15)}{=} \sum_{\tau=1}^T \mathbb{E}_{v_\tau} [\hat{u}_\tau(0, v_\tau; v_\tau)] \\ &= \sum_{\tau=1}^T \int_{r_\tau(0)}^{\infty} (v - r_\tau(0)) d\mathcal{F}_\tau(v) = \text{UTL}(M) \end{aligned}$$

3. $\text{REV}(B) \geq \text{REV}(M)$.

Denote the expected revenue improvement for stage t by κ_t ,

$$\begin{aligned} \kappa_t &= \mathbb{E}_{v_{(1,t)}} [q_t(\text{bal}_t, v_t) + s_t(\text{bal}_t)] - \rho_t(r_t(0)) \\ &= \mathbb{E}_{\text{bal}_t} [\rho_t(r_t(\text{bal}_t)) + \text{bal}_t] - \rho_t(r_t(0)) \\ &= \alpha_t (\rho_t(r_t(\bar{L}_t)) + \bar{L}_t) - \alpha_t \rho_t(r_t(0)) \\ &= \alpha_t (\rho_t(r_t(\bar{L}_t)) + \text{Val}_t(r_t(\bar{L}_t), r_t(0)) - \rho_t(r_t(0))) \\ &= \alpha_t \int_{r_t(\bar{L}_t)}^{r_t(0)} v d\mathcal{F}_t(v) \geq 0 \end{aligned}$$

where $\alpha_t = \Pr[\text{bal}_t = \bar{L}_t]$, and the last “=” is from integration by parts. Hence the overall revenue improvement is non-negative, i.e., $\text{REV}(B) - \text{REV}(M) = \sum_{\tau=1}^T \kappa_\tau \geq 0$.

4. $\text{REV}(B)$ is increasing in $L_{(1,T)}$.

Moreover, each κ_t is weakly increasing in bank account limit $L_{t'}$ for $t' \leq t$.² Hence the expected revenue increases as the bank account limit increases. \square

As we proposed, double-reserve auction is easy to implement, and can improve the revenue without hurting the buyer utility under any given bank account limit. A natural question

² α_t is increasing in $L_{t'}$. $r_t(\bar{L}_t)$ is decreasing (hence the integration is increasing) in $L_{t'}$ when $t = t'$. Therefore κ_t is increasing.

is: what is the limitation of such mechanisms comparing with the optimal mechanism under the same bank account limit? The following theorem provides an upper bound on the revenue for any mechanism under given limit.

Theorem 3. For any bank account mechanism B with limits $\{L_t\}_{t=1}^T$,

$$\text{REV}(B) \leq \sum_{\tau=1}^T L_\tau + \text{REV}(M^S) = \sum_{\tau=1}^T (L_\tau + \rho_\tau(\phi_\tau)),$$

where M^S denotes the separate Myerson auction. In other words, the revenue improvement comparing with separate Myerson auction is bounded by the sum of bank account limit on each stage.

Proof. Assume by contradiction that there exists a BAM, B , such that,

$$\text{REV}(B) > \sum_{\tau=1}^T (L_\tau + \rho_\tau(\phi_\tau)). \quad (20)$$

Then we construct a mechanism M that operates independently on each stage but generates strictly higher revenue than Myerson auction in at least one stage.

To ensure that the constructed mechanism is history independent and IC-IR within each single stage. We create a fake bank account to mimic the behavior of the real bank account in B . One can think of the balance bal_t^\dagger of the fake bank account as simply a number, satisfying that the distribution of bal_t^\dagger is the same as the distribution of the real balance (as a random-valued function of previous types).

Consider the following construction of M based on B .

$$x_t(v_t) = z_t(\text{bal}_t^\dagger, v_t), \quad p_t(v_t) = q_t(\text{bal}_t^\dagger, v_t),$$

where in each stage t , bal_t^\dagger is a random variable drawn by the following process for τ from 1 to $t - 1$.

$$\text{bal}_{\tau+1}^\dagger = \text{bal}_\tau^\dagger - s_\tau(\text{bal}_\tau^\dagger) + d_\tau(\text{bal}_\tau^\dagger, v_\tau^\dagger).$$

Thus bal_t^\dagger is independent of the buyer's real types, and has the same distribution with bal_t . Note that v_t is independent of the history, hence independent of bal_t as well. Then

$$\begin{aligned} \mathbb{E}_{v_{(1,t-1)}} [q_t(\text{bal}_t, v_t)] &= \mathbb{E}_{\text{bal}_t^\dagger} [p_t(v_t)], \\ s_\tau(\text{bal}_\tau) &\leq \text{bal}_\tau \leq L_\tau, \\ \implies \text{REV}(M) &= \sum_{\tau=1}^T \mathbb{E}_{\text{bal}_\tau^\dagger, v_\tau} [p_\tau(v_\tau)] \\ &= \text{REV}(B) - \mathbb{E}_{v_{(1,T)}} [\sum_{\tau=1}^T s_\tau(\text{bal}_\tau)] \\ &\geq \text{REV}(B) - \sum_{\tau=1}^T L_\tau > \sum_{\tau=1}^T \rho_\tau(\phi_\tau) \\ \implies \exists \tau, s.t. \mathbb{E}_{\text{bal}_\tau^\dagger, v_\tau} [p_\tau(v_\tau)] &> \rho_\tau(\phi_\tau), \quad (21) \end{aligned}$$

namely, at least in one stage τ , M generates strictly higher revenue than the optimal revenue, $\rho_\tau(\phi_\tau)$.

In particular, by construction M is history independent, and by (6)(8) and $q_t \geq 0$, M is also IC and IR within each single stage. Thus (21) contradicts the optimality of Myerson auction for single-shot case. \square

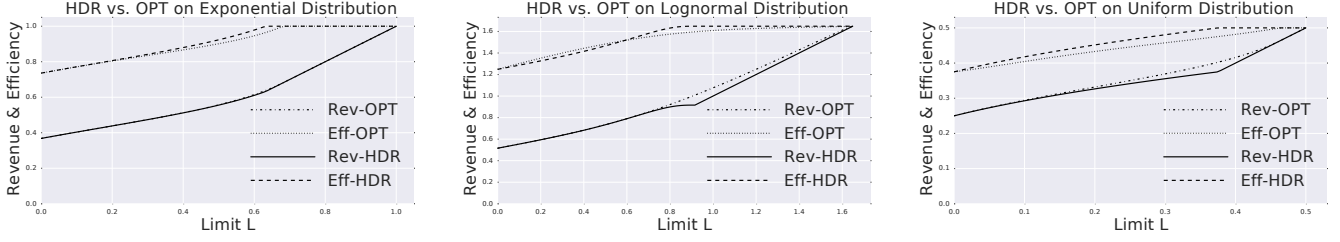


Figure 2: HDR vs. OPT on revenue and efficiency for different distributions.

3.2 Compute Optimal Double-Reserve Auction via Dynamic Programming

As we proved by Theorem 2, the double-reserve auction defined upon a separate posted-price auction has a quantified revenue improvement guarantee. In particular, it is better than the separate Myerson auction when defined on separate Myerson auction.

However, this construction is not guaranteed to be optimal among the double-reserve auctions. In this section, we compute optimal double-reserve auctions.

Theorem 4. *The optimal double-reserve auction subject to given bank account limit can be computed through a dynamic programming algorithm.*

Moreover, for any $\epsilon > 0$, there is an FPTAS³ to achieve an ϵ -approximation of the optimal double-reserve auction.

Proof sketch. Dynamic programming (DP) algorithm.

For any double-reserve auction B , a stage does not depend on latter stages and it can only influence the latter ones through the balance. The balance bal_t is a random variable distributed according to a two-valued distribution, i.e., $\Pr[\text{bal}_t = \bar{L}_t] = \alpha_t$ and $\Pr[\text{bal}_t = 0] = 1 - \alpha_t$.

$$\alpha_{t+1} = 1 - \alpha_t \mathcal{F}_t(r_t(\bar{L}_t)) - (1 - \alpha_t) \mathcal{F}_t(r_t(0)) \quad (22)$$

Moreover, α_{t+1} is uniquely defined by α_t and $r_t(0)$, and so is the expected stage revenue. Hence B is entirely determined by $r_{(1,T)}(0)$, and the optimal double-reserve auction can be computed via DP, where the state update formula is,

$$h_t(\alpha_t) = \max_{r_t(0)} (\kappa_t + \rho_t(r_t(0)) + h_{t+1}(\alpha_{t+1})). \quad (23)$$

$h_t(\alpha_t)$ is the maximum expected revenue from stage t to stage T with specified α_t . Note that κ_t and α_{t+1} are functions of α_t and $r_t(0)$.

FPTAS for the DP: a natural FPTAS to implement the DP algorithm is to discretize h_t over its domain, $[0, 1]$. One critical observation is that h_t is increasing (proof omitted) and bounded (Theorem 3), so can be ϵ -approximated by at most $O(tL/\epsilon)$ many sample points over $[0, 1]$. \square

4 Empirical Evaluations

In this section, we empirically evaluate the expected revenue and efficiency of the optimal double-reserve auction (OPT) and the so-called *heuristic double-reserve auction* (HDR) over the infinite stream of i.i.d. items as functions of the bank

account limit L . The reason that we propose and evaluate HDR is that it is easy to implement, non-sensitive to prior distribution \mathcal{F} , and yields near optimal revenue, compared to the OPT. As a result, HDR can be used as a replacement of OPT in practice.

To simplify notation, we remove all subscript t because of the i.i.d. assumption. Besides, by revenue and efficiency in this section, we refer to per stage revenue and efficiency.

Mechanism 3 (HDR). *HDR is the better one of two double-reserve auctions, B' and B'' .*

- (1) B' : $r'(0) = \phi$; (2) B'' : $r''(L) = 0$.

Intuitively, with given limit L , (1) B' is the double-reserve auction defined upon separate Myerson auction; (2) B'' is the double-reserve auction with the lower reserve $r''(L)$ being 0.

We provide a lower bound analysis for the revenue of HDR as a function of the limit L .

Theorem 5. *The expected revenue per stage π is lower bounded by L . If \mathcal{F} is regular, then π is further lower bounded as follows, where $\mathcal{G} = 1 - \mathcal{F}$, and $\xi = 1 - e^{L\mathcal{F}(\phi)/\rho(\phi)}$.*

$$\pi \geq \frac{\mathcal{G}(\phi)}{\mathcal{G}(\phi)\xi + 1} L + \left(\frac{\mathcal{G}(\phi)^2 \xi}{(\mathcal{G}(\phi)\xi + 1)\mathcal{F}(\phi)} + 1 \right) \rho(\phi)$$

Empirical Evaluations We evaluate both mechanisms on exponential distribution ($\lambda = 1$), lognormal distribution ($\ln \mathcal{N}(0, 1)$), and $[0, 1]$ uniform distribution (Figure 2). For each case, we illustrate the revenue and efficiency of OPT and HDR as functions of the limit L . Note that when the limit is greater than the mean of the valuation distribution, both OPT and HDR converge to fully surplus extraction, therefore the revenue and efficiency will equal to the mean.

In all three cases, the revenue of HDR and OPT are very close, especially when L is less than 50% of the mean. In particular, the revenue improvement (compared with Myerson revenue) can be fitted as a linear function of L with coefficient roughly 0.4 to 0.5 when L is small. As L tends to the mean of each distribution, both HDR and OPT converge to fully surplus extraction. As for efficiency, HDR is better than OPT for most cases, except for the lognormal distribution when $L \leq 0.6$.

In summary, for all the distribution considered above, for reasonably small limit (less than 50% of the mean), the revenue improvements of both mechanism (OPT and HDR) are very close, and grows almost linearly as the limit increases. In contrast, the efficiency of HDR almost dominates that of OPT.

³Fully Polynomial Time Approximation Scheme.

References

- [Athey and Segal, 2013] Susan Athey and Ilya Segal. An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485, 2013.
- [Babaioff *et al.*, 2014] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 21–30. IEEE, 2014.
- [Bergemann and Said, 2011] Dirk Bergemann and Maher Said. Dynamic auctions. *Wiley Encyclopedia of Operations Research and Management Science*, 2011.
- [Bergemann and Välimäki, 2010] Dirk Bergemann and Juuso Välimäki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010.
- [Cai *et al.*, 2012a] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 459–478. ACM, 2012.
- [Cai *et al.*, 2012b] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 130–139. IEEE, 2012.
- [Cai *et al.*, 2016] Yang Cai, Nikhil R Devanur, and S Matthew Weinberg. A duality based unified approach to bayesian mechanism design. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, 2016.
- [Cavallo *et al.*, 2009] Ruggiero Cavallo, David C Parkes, and Satinder Singh. Efficient mechanisms with dynamic populations and dynamic types. 2009.
- [Cavallo, 2008] Ruggiero Cavallo. Efficiency and redistribution in dynamic mechanism design. In *Proceedings of the 9th ACM conference on Electronic commerce*, pages 220–229. ACM, 2008.
- [Daskalakis *et al.*, 2013] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism design via optimal transport. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 269–286. ACM, 2013.
- [Daskalakis *et al.*, 2015] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Strong duality for a multiple-good monopolist. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 449–450. ACM, 2015.
- [Gershkov and Moldovanu, 2009] Alex Gershkov and Benny Moldovanu. Dynamic revenue maximization with heterogeneous objects: A mechanism design approach. *American Economic Journal: Microeconomics*, 1(2):168–198, 2009.
- [Gershkov and Moldovanu, 2010] Alex Gershkov and Benny Moldovanu. Efficient sequential assignment with incomplete information. *Games and Economic Behavior*, 68(1):144–154, 2010.
- [Green and Laffont, 1977] Jerry Green and Jean-Jacques Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica: Journal of the Econometric Society*, pages 427–438, 1977.
- [Hart and Nisan, 2012] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. *arXiv preprint arXiv:1204.1846*, 2012.
- [Kakade *et al.*, 2013] Sham M Kakade, Ilan Lobel, and Hamid Nazerzadeh. Optimal dynamic mechanism design and the virtual-pivot mechanism. *Operations Research*, 61(4):837–854, 2013.
- [Li and Yao, 2013] Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28):11232–11237, 2013.
- [Manelli and Vincent, 2006] Alejandro M Manelli and Daniel R Vincent. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory*, 127(1):1–35, 2006.
- [Manelli and Vincent, 2007] Alejandro M Manelli and Daniel R Vincent. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1):153–185, 2007.
- [Pai and Vohra, 2008] Malleesh Pai and Rakesh V Vohra. Optimal dynamic auctions. Technical report, Discussion paper//Center for Mathematical Studies in Economics and Management Science, 2008.
- [Papadimitriou *et al.*, 2014] Christos Papadimitriou, George Pierrakos, Christos-Alexandros Psomas, and Aviad Rubinfeld. On the complexity of dynamic mechanism design, 2014.
- [Pavan *et al.*, 2008] Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: Revenue equivalence, profit maximization and information disclosure. 2008.
- [Pavan *et al.*, 2014] Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: A myersonian approach. 2014.
- [Tang and Sandholm, 2012] Pingzhong Tang and Tuomas Sandholm. Mixed-bundling auctions with reserve prices. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 729–736. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- [Wang and Tang, 2014] Zihe Wang and Pingzhong Tang. Optimal mechanisms with simple menus. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 227–240. ACM, 2014.
- [Zuo *et al.*, 2016] Song Zuo, Pingzhong Tang, Renato Paes Leme, and Vahab Mirrokni. Optimal dynamic mechanisms with ex-post ir via bank accounts. Technical report, 2016.