

Optimal Dynamic Auctions Are Virtual Welfare Maximizers*

Vahab Mirrokni
Google Research
mirrokni@google.com

Renato Paes Leme
Google Research
renatopl@google.com

Pingzhong Tang
Tsinghua University
kenshinping@gmail.com

Song Zuo
Google Research
szuo@google.com

Abstract

We are interested in the setting where a seller sells sequentially arriving items, one per period, via a *dynamic auction*. At the beginning of each period, each buyer draws a private valuation for the item to be sold in that period and this valuation is independent across buyers and periods. The auction can be dynamic in the sense that the auction at period t can be conditional on the bids in that period and all previous periods, subject to certain appropriately defined *incentive compatible* and *individually rational* conditions. Perhaps not surprisingly, the revenue optimal dynamic auctions are computationally hard to find and existing literatures that aim to approximate the optimal auctions are all based on solving complex dynamic programs. It remains largely open on the structural interpretability of the optimal dynamic auctions.

In this paper, we show that any optimal dynamic auction is a virtual welfare maximizer subject to some monotone allocation constraints. In particular, the explicit definition of the virtual value function above arises naturally from the primal-dual analysis by relaxing the monotone constraints. We further develop an ironing technique that gets rid of the monotone allocation constraints. Quite different from Myerson's ironing approach, our technique is more technically involved due to the interdependence of the virtual value functions across buyers. We nevertheless show that ironing can be done approximately and efficiently, which in turn leads to a Fully Polynomial Time Approximation Scheme of the optimal dynamic auction.

1 Introduction

Recently, the problem of designing multi-period dynamic mechanisms has been shown to be both theoretically challenging and practically important. In particular, when running a sequence of repeated auctions on online advertising platforms, using dynamic auctions optimized across different time periods could potentially bring significant gains both in terms of revenue and social welfare. The power of

dynamic mechanisms has been investigated by a number of recent papers (Bergemann and Välimäki 2002; Parkes and Singh 2004; Cavallo 2008; Athey and Segal 2013; Kakade, Lobel, and Nazerzadeh 2013; Pai and Vohra 2013; Pavan, Segal, and Toikka 2014; Devanur, Peres, and Sivan 2015; Balseiro, Mirrokni, and Paes Leme 2016; Chawla et al. 2016; Bergemann, Castro, and Weintraub 2017; Balseiro et al. 2017; Balseiro, Mirrokni, and Leme 2017; Lobel and Paes Leme 2017; Shen, Wang, and Zuo 2018; Balseiro et al. 2019). We refer to Bergemann and Said (2011) and Bergemann and Välimäki (2017) for comprehensive surveys on the subject.

In particular, we consider a setting where a seller repeatedly interacts with a set of buyers and sells one item per period. The value of each buyer for the item at each period are independently drawn from commonly known prior distributions (no need to be identical) at the beginning of that period. The seller is allowed to sell the item of each period not only depending on the bids submitted in the current period, but also the histories, i.e., all the bids submitted in past periods. In the meanwhile, the seller must guarantee the dynamic auctions to be *ex-post individual rational* — the cumulative utility for each buyer is always positive, and *dynamic incentive compatible* — bidding truthfully (i.e., submitting private values as bids) to the auction is optimal for each buyer by taking into consideration the effect of current bids on future outcomes.

Even though dynamic mechanisms can be much more effective in maximizing revenue and social welfare (Jackson and Sonnenschein 2007; Papadimitriou et al. 2016), they have not been widely adopted in practice. The main issue of implementing dynamic mechanisms in practice is their high complexity. The complexity induced by the exponentially growing design space makes it difficult to solve or even to describe such mechanisms.

A series of recent work has made progress to resolve the complexity issue described above. For example, Ashlagi, Daskalakis, and Haghpahan (2016) and Mirrokni et al. (2018b) show that it is enough for the optimal dynamic auction to depend on scalar summaries of histories instead of the full history. However, the approximately optimal mechanisms described in these papers are solutions of complex dynamic programs that are written into a large table. It is therefore very difficult to understand the structure of these

*The authors thank the anonymous reviewers for their helpful comments. This paper is supported in part by the National Natural Science Foundation of China Grant 61561146398, a China Youth 1000-talent Program, and an Alibaba Innovative Research Program. A full version of this paper is available at <https://ssrn.com/abstract=3282600>. Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

mechanisms. It remains open whether there is an intuitive structural characterization of the optimal dynamic auction.

In this paper, we show that the exact optimal dynamic auction has a very simple structural interpretation: the optimal dynamic auction in each period is a second price auction on a certain appropriately defined virtual value space.¹ More specifically, such virtual values (before ironing^{2,3}) are quite similar to Myerson’s virtual value (Myerson 1981), i.e., they have the form of linear combinations of private values and Myerson’s virtual values (before ironing). However, just like Myerson’s auction, to make the virtual welfare⁴ maximizing allocation rule monotone, one need to first iron the virtual values. Unlike the ironing in Myerson’s auction, the ironing step in our case is interdependent across the values of different buyers. In other words, one’s virtual value after ironing not only depends on his/her own value, but also on other buyers’ values. Although the ironing step here is not as simple as the ironing in Myerson’s auction, its computation is still efficient for constant many buyer cases. Moreover, we provide a Fully Polynomial Time Approximation Scheme to compute the virtual values for any period of the dynamic auction given the histories so far, and such virtual values induce a dynamic auction with revenue arbitrarily close to the optimal.

Techniques There are two main techniques used in our analysis for optimal dynamic auctions. The first is the so-called *bank account mechanisms* which are a subset of dynamic auctions with simple structures and can achieve the optimal revenue of all the dynamic auctions (Mirrokni et al. 2016a; 2016b; 2018a; 2018b). Briefly speaking, a bank account mechanism keeps a state for each buyer as the summary of the history of each buyer, which is a scalar called *balance*. Each period depends on the previous periods through the vector of buyer balances. With the bank account framework, the designer only needs to specify single-period auctions that are single-period incentive compatible together with a valid balance update policy. In other words, the design of an entire dynamic auction breaks into the design of a series of single-period auctions and a balance update policy. The decomposition greatly simplifies the problem and enables clean mathematical programs for each period.

The second is a primal-dual analysis and a sensitivity analysis of the parametric mathematical programs. In fact, primal-dual analysis is commonly used in economic studies (Nisan et al. 2007, Chapter 5; Daskalakis, Deckelbaum, and Tzamos; Cai, Devanur, and Weinberg; Cai and

¹A virtual value function is a map from buyer value space to real numbers. A virtual value is the corresponding real number of some private value of the buyer.

²Informally, the ironing of a virtual value function is an operation that maps the virtual value function to another virtual value function (we call an ironed virtual value function) while preserving the expectation. For the formal definition, see Definition 3.1.

³Some recent works on virtual value and ironing (Elkind 2007; Roughgarden and Schrijvers 2016).

⁴Virtual welfare is the sum of virtual values of the buyers who get the item. For expected virtual welfare, the expectation is taken over the randomness from the allocation rules.

Zhao; Daskalakis, Deckelbaum, and Tzamos 2013; 2016; 2017; 2017). In particular, for our problem, we can prove that these programs are convex and satisfy the Slater’s conditions. Hence the solution is optimal if and only if the KKT conditions are satisfied. From these conditions, we show that the auction in each period maximizes some virtual welfare and we also discover the exact form of the virtual values with partial relaxation on the monotonicity constraints of allocation. Furthermore, we show that adding these relaxed constraints back to the program indeed corresponds to ironing the virtual values. As we mentioned previously, the ironing step here is interdependent across the values of the buyers and hence different from the ironing step in Myerson’s auction. To resolve this difficulty, we show a method to algorithmically accomplish the ironing step.

2 Preliminaries

We study a setting where a seller repeatedly interacts with k buyers selling one item per period over T periods. The value of each buyer $i \in [k]$ for the item in period $t \in [T]$ is $v_i^t \in \mathcal{V}$. If $x_i^t \in [0, 1]$ represents the probability that buyer i is allocated the item in period t , his utility is $v_i^t \cdot x_i^t$. Throughout this paper, (i) we use subscripts as the indices of buyers, bold fonts for vectors of all the buyers (i.e., $\mathbf{a} = (a_1, \dots, a_k)$), and subscript $-i$ for the vector except the i -th element (i.e., $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_k)$); (ii) we use superscripts as the indices of periods and $a^{1..t}$ to denote the sequence a^1, \dots, a^t .

The values v_i^t are assumed to be drawn from independent distributions F_i^t .^{5,6} The distributions F_i^t are assumed to be common knowledge but the realizations of the random variables are initially unknown for both the buyers and the designer. At each period, the following events happen:

1. each buyer i learns his value $v_i^t \sim F_i^t$;
2. each buyer i reports value \hat{v}_i^t to the designer;
3. the designer implements an outcome $\mathbf{x}^t \in [0, 1]^k$ and charges the buyers \mathbf{p}^t ;
4. each buyer accrues utility $u_i^t = v_i^t \cdot x_i^t - p_i^t$.

A dynamic mechanism is then described in terms of a pair of maps for each period, which associate the history of reports $\hat{\mathbf{v}}^{1..t} = (\hat{\mathbf{v}}^1, \hat{\mathbf{v}}^2, \dots, \hat{\mathbf{v}}^t)$ to an outcome \mathbf{x}^t and payment \mathbf{p}^t :

$$\text{Outcome: } \mathbf{x}^t : \mathcal{V}^{kt} \rightarrow [0, 1]^k, \quad \text{Payment: } \mathbf{p}^t : \mathcal{V}^{kt} \rightarrow \mathbb{R}^k.$$

Finally we define: $u_i^t(v_i^t; \hat{\mathbf{v}}^{1..t}) = v_i^t \cdot x_i^t(\hat{\mathbf{v}}^{1..t}) - p_i^t(\hat{\mathbf{v}}^{1..t})$.

2.1 Dynamic Incentive Compatibility

A mechanism is incentive compatible if it provides incentives for buyers to reveal their true types in each iteration. Such conditions for dynamic mechanisms can be easily defined by backward induction: in the last period, regardless

⁵In practice, it is still fair to assume that the values of the buyers (advertisers) are independent conditional on each particular inventory and cookie.

⁶If the distributions are not independent, then the weak version of truthfulness still holds: truthful if all other buyers bid truthfully.

of the history so far and other buyers' reports, it should be incentive compatible for each buyer to report his true value:

$$v_i^T = \operatorname{argmax}_{\hat{v}_i^T} u_i^T(v_i^T; \hat{\mathbf{v}}^{1..T}), \quad \forall i, \hat{\mathbf{v}}^{1..T-1}, \hat{v}_{-i}^T, v_i^T.$$

To simplify notations, from now on we will omit the 'for-all' quantification and assume all expressions are quantified as 'for-all' in its free variables. For the next-to-last-period, it should be incentive compatible for the buyer to report his true value given that he will report his true value in the following period:

$$v_i^{T-1} = \operatorname{argmax}_{\hat{v}_i^{T-1}} u_i^{T-1}(v_i^{T-1}; \hat{\mathbf{v}}^{1..T-1}) + \mathbb{E}_{v_i^T} [u_i^T(v_i^T; \hat{\mathbf{v}}^{1..T-1}, v_i^T, \hat{v}_{-i}^T)].$$

Proceeding by backward induction, we require that:

$$v_i^t = \operatorname{argmax}_{\hat{v}_i^t} u_i^t(v_i^t; \hat{\mathbf{v}}^{1..t}) + U_i^t(\hat{v}_i^{1..t} | \hat{v}_{-i}^{1..T}) \quad (\text{DIC})$$

where the second term is the *continuation utility*, i.e., the expected utility obtained from the subsequent periods of the mechanism assuming the buyer reports truthfully:

$$U_i^t(\hat{v}_i^{1..t} | \hat{v}_{-i}^{1..T}) := \mathbb{E}_{v_i^{t+1..T}} \left[\sum_{\tau=t+1}^T u_i^\tau(v_i^\tau; \hat{\mathbf{v}}^{1..t}, v_i^{t+1..t+\tau}, \hat{v}_{-i}^{t+1..t+\tau}) \right]$$

A well-known fact in dynamic mechanism design is that DIC implies that buyer i 's expected overall utility $U_i^0(\hat{v}_i^{1..T})$ is maximized by reporting truthfully in each period.

2.2 Ex-Post Individual Rationality

Another desirable constraint is ex-post individual rationality which says that a buyer should derive non-negative utility from the mechanism for every realization of the values:

$$\sum_{t=1}^T u_i^t(v_i^t; \mathbf{v}^{1..t}) \geq 0 \quad (\text{eP-IR})$$

We focus on the problem of maximizing revenue subject to DIC, eP-IR, and feasibility constraints:

$$\begin{aligned} \max \quad & \text{REV} = \mathbb{E}[\sum_{t=1}^T \sum_{i=1}^k p_i^t(\mathbf{v}^{1..t})] \\ \text{s.t.} \quad & (\text{DIC}), (\text{eP-IR}), \text{ and feasibility: } \sum_{i=1}^k x_i^t(\mathbf{v}^{1..t}) \leq 1 \end{aligned}$$

2.3 Bank Account Mechanisms (BAM)

The space of mechanisms satisfying DIC and eP-IR is very broad and unstructured. We restrict our attention to a subclass of dynamic mechanisms introduced by Mirrokni et al. (2016a) called *bank account mechanisms*. The mechanisms are simple, dynamic incentive compatible by design and have the following notable features:

Lemma 2.1 ((Mirrokni et al. 2018b)). *Given any dynamic mechanism satisfying DIC and eP-IR, there exists a bank account mechanism with at least the same revenue and welfare.*

In particular, for any given setting, there is an optimal mechanism in the form of a bank account mechanism.

A bank account mechanism keeps a state for each buyer, which is a scalar called *balance*. Each period depends on previous ones through the vector of balances. Another main feature is that in this framework, the designer needs to specify single-period auctions that are single-period incentive

compatible together with a valid balance update policy. Once a valid balance update policy is in place, all the designer needs to worry about are single-period IC constraints.

A bank account mechanism B is defined in terms of the following functions for each period:

- A static single-period auction $\mathbf{x}^{B,t}(\mathbf{v}^t, \mathbf{b}), \mathbf{p}^{B,t}(\mathbf{v}^t, \mathbf{b})$ parameterized by a balance vector $\mathbf{b} \in \mathbb{R}_+^k$ that is (single-period) *incentive-compatible* for each \mathbf{b} , i.e.:

$$\begin{aligned} & v_i^t \cdot x_i^{B,t}(\mathbf{v}^t, \mathbf{b}) - p_i^{B,t}(\mathbf{v}^t, \mathbf{b}) \\ & \geq v_i^t \cdot x_i^{B,t}(\hat{v}_i^t, v_{-i}^t, \mathbf{b}) - p_i^{B,t}(\hat{v}_i^t, v_{-i}^t, \mathbf{b}) \end{aligned} \quad (\text{IC})$$

- Note that we do not require the mechanism to be (single-period) individually rational, while we require the buyer utility to be *balance independent in expectation*, i.e.:

$$\mathbb{E}_{v_i^t \sim F_i^t} \left[v_i^t \cdot x_i^{B,t}(\mathbf{v}^t, \mathbf{b}) - p_i^{B,t}(\mathbf{v}^t, \mathbf{b}) \right] \quad (\text{BI})$$

is a non-negative constant not depending on \mathbf{b}

- A balance update policy $\mathbf{b}^{B,t}(\mathbf{v}^t, \mathbf{b})$ which maps the previous balances and the reports to the current balances, satisfying the following *balance update* conditions:

$$0 \leq b_i^{B,t}(\mathbf{v}^t, \mathbf{b}) \leq b^i + v_i^t \cdot x_i^{B,t}(\mathbf{v}^t, \mathbf{b}) - p_i^{B,t}(\mathbf{v}^t, \mathbf{b}) \quad (\text{BU})$$

Given the balance update functions, we can define $\mathbf{b}^t : \mathcal{V}^t \rightarrow \mathbb{R}_+^k$ recursively as:

$$\begin{aligned} \mathbf{b}^0 &= \mathbf{0} \quad \text{and} \quad \mathbf{b}^1(\mathbf{v}^1) = \mathbf{b}^{B,1}(\mathbf{v}^1, \mathbf{0}) \\ \text{and} \quad \mathbf{b}^t(\mathbf{v}^{1..t}) &= \mathbf{b}^{B,t}(\mathbf{v}^t, \mathbf{b}^{B,t-1}(\mathbf{v}^{1..t-1})) \end{aligned}$$

which allows us to define a dynamic mechanism in the standard sense as:

$$\mathbf{x}^t(\mathbf{v}^{1..t}) = \mathbf{x}^{B,t}(\mathbf{v}^t, \mathbf{b}^{t-1}), \quad \mathbf{p}^t(\mathbf{v}^{1..t}) = \mathbf{p}^{B,t}(\mathbf{v}^t, \mathbf{b}^{t-1}).$$

In what follows we will abuse notations by dropping the superscript B and refer to $\mathbf{x}^t(\mathbf{v}^{1..t})$ and $\mathbf{x}^t(\mathbf{v}^t, \mathbf{b}^{t-1})$ interchangeably. One important theorem from previous studies is that any bank account mechanism satisfies DIC and eP-IR.

Lemma 2.2 ((Mirrokni et al. 2018b)). *Any bank account mechanism satisfying IC, BI, and BU is DIC and eP-IR.*

3 The Structure of Optimal BAMs

In this section, we show our first main result that identifies the underlying structure of the optimal bank account mechanism (hence also the optimal dynamic auctions). Interestingly, the dynamic auctions with the structure can be interpreted as an ironed virtual welfare maximizing auction, where the ironed virtual value is defined as follows.

Definition 3.1 (Ironing). *The ironing operation on the virtual value functions of all buyers transfers these virtual value functions into some other virtual value functions (called ironed virtual value functions) while preserving the following properties:*

- *Monotonicity: The allocation rule maximizes the ironed virtual welfare is monotone: for each buyer i , the allocation probability is weakly increasing in v_i for fixed v_{-i} .*

- Limited transfer: *The expectation of virtual value conditional on each allocation equivalence class is unchanged. An allocation equivalence class of buyer i with any given v_{-i} is a maximal subset of his/her private values where the allocation probability to him/her is constant. The allocation rule here must maximize the ironed virtual welfare.*
- First order dominance: *The ironed virtual value is first order dominated by the virtual value before ironing.*

Informally, we have the following theorem:

Theorem 3.2 (Informal). *Any revenue optimal bank account mechanism is maximizing some virtual value (after ironing) for each period.*

We will restate the formal version after introducing all necessary notations (Theorem 3.8) and the specific form of the virtual value (before ironing) will be provided.

3.1 Formalizing the subprogram for each period

We start with the subproblem of optimizing period t while all other periods are fixed. In particular, the problem can be formalized as a convex program.

Lemma 3.3 (Convex program). *For each period t , any valid balance \mathbf{b}^t , and fixed expected utility for period t , if mechanism is fixed for the remaining periods ($t + 1$ to T), then the optimal auction for period t can be solved via a convex program.*

To proof this statement smoothly, we show how to formalize the program step by step in the rest of this subsection rather than putting everything into a single proof environment. For the ease of presentation, we will hide the superscript t while focusing on a single period.

Let \mathbf{b} be the bank account balance given at the beginning of this period and

$$\xi_i(v_{-i}) := \mathbb{E}_{v_i}[v_i \cdot x(\mathbf{v}, \mathbf{b}) - p(\mathbf{v}, \mathbf{b})] \quad (1)$$

be the expected utility of buyer i in this period. Note that by (BI), it is independent of the bank account balance \mathbf{b} , while it could be different for different bids from other buyers, i.e., v_{-i} . In fact, the ξ are some parameters we will need to determine later based on the distribution of \mathbf{b} induced by the auctions in each period.

Objective The expected revenue acquired within this period can be computed by the expected social welfare minus the expected utility of all the buyers,

$$\text{PERIODREV}(\mathbf{b}, \xi) = \mathbb{E}_{\mathbf{v}}[\sum_i (v_i \cdot x_i(\mathbf{v}, \mathbf{b}) - \xi_i(v_{-i}))].$$

Besides the expected revenue from the current period, we also need to include the expected revenue for future periods in the objective. We then use $g(\mathbf{b}'|M)$ to denote the expected total revenue the seller can collect from all upcoming periods by following a given bank account mechanism M in these periods, where \mathbf{b}' is the vector of the bank account balance at the end of the current period. In fact, $\mathbf{b}' = \mathbf{b} + \Delta\mathbf{b}(\mathbf{v})$ is a function of buyers' bids \mathbf{v} in this period, where $\Delta\mathbf{b}(\mathbf{v})$ are the changes in the bank accounts.

Then the expected revenue since current period is,

$$\text{REV}(\mathbf{b}, \xi) = \mathbb{E}_{\mathbf{v}}[\sum_i (v_i x_i(\mathbf{v}, \mathbf{b}) - \xi_i(v_{-i})) + g(\mathbf{b}'|M)].$$

Constraints Now we proceed to the constraints of the optimization. According to Lemma 2.2, the mechanism must satisfy constraint (IC), (BI), and (BU). Note that (BI) is guaranteed by the definition of $\xi_i(v_{-i})$, we then formalize constraint (IC) and (BU).

For (IC), as we won't introduce the payment variables, it is enough to require monotone allocation rules only, i.e.,

$$(IC) \iff x_i(\mathbf{v}, \mathbf{b}) \text{ is monotone in } v_i \text{ for fixed } v_{-i}. \quad (2)$$

Since by the Envelope theorem (Rochet 1985),

$$\begin{aligned} \partial(v_i \cdot x_i(\mathbf{v}, \mathbf{b}) - p_i(\mathbf{v}, \mathbf{b}))/\partial v_i &= x_i(\mathbf{v}, \mathbf{b}) \\ \implies u_i(\mathbf{v}, \mathbf{b}) &= u_i(0, v_{-i}, \mathbf{b}) + \int_0^{v_i} x_i(s, v_{-i}, \mathbf{b}) ds. \end{aligned}$$

In other words, with $p'_i(\mathbf{v}, \mathbf{b}) := p_i(\mathbf{v}, \mathbf{b}) - p_i(0, v_{-i}, \mathbf{b})$ and $u'_i(\mathbf{v}, \mathbf{b}) := v_i \cdot x_i(\mathbf{v}, \mathbf{b}) - p'_i(\mathbf{v}, \mathbf{b})$, the utility of buyer i in this period is fully determined by the allocation function $x_i(\cdot, \mathbf{b})$ and the minimum payment $p_i(0, v_{-i}, \mathbf{b})$:

$$\begin{aligned} u'_i(\mathbf{v}, \mathbf{b}) &= u_i(\mathbf{v}, \mathbf{b}) + p_i(0, v_{-i}, \mathbf{b}) \\ &= u_i(\mathbf{v}, \mathbf{b}) - u_i(0, v_{-i}, \mathbf{b}) + 0 \cdot x_i(0, v_{-i}, \mathbf{b}) \\ &= \int_0^{v_i} x_i(s, v_{-i}, \mathbf{b}) ds. \end{aligned}$$

Recall that for (BU), we require (i) the increase of balance cannot be more than the utility obtained in the current period and (ii) the updated balance must be nonnegative. Note that by the definition of $\xi_i(v_{-i})$ (1):

$$\begin{aligned} \mathbb{E}_{v_i}[u'_i(\mathbf{v}, \mathbf{b}) - p_i(0, v_{-i}, \mathbf{b})] &= \mathbb{E}_{v_i}[u_i(\mathbf{v}, \mathbf{b})] = \xi_i(v_{-i}) \\ \iff p_i(0, v_{-i}, \mathbf{b}) &= \mathbb{E}_{v_i}[u'_i(\mathbf{v}, \mathbf{b})] - \xi_i(v_{-i}). \end{aligned}$$

Hence on one hand, for (i) and (ii), we can rewrite it as:

$$\begin{aligned} 0 \leq b_i + \Delta b_i(\mathbf{v}) \leq b_i + u_i(\mathbf{v}, \mathbf{b}) &\iff \\ -b_i \leq \Delta b_i(\mathbf{v}) \leq u_i(\mathbf{v}, \mathbf{b}) &= u'_i(\mathbf{v}, \mathbf{b}) - p_i(0, v_{-i}, \mathbf{b}) \\ &= u'_i(\mathbf{v}, \mathbf{b}) - \mathbb{E}_{v_i}[u'_i(\mathbf{v}, \mathbf{b})] + \xi_i(v_{-i}). \end{aligned}$$

On the other hand, to ensure the set of the possible balance increment $\Delta b_i(\mathbf{v})$ satisfying (i) and (ii) is not empty, we need $b_i + u_i(\mathbf{v}, \mathbf{b}) \geq 0$. Note that b_i is given and $u_i(\mathbf{v}, \mathbf{b})$ reaches the minimum when $v_i = 0$, it is equivalent to:

$$\begin{aligned} b_i + u_i(0, v_{-i}, \mathbf{b}) &= b_i - p_i(0, v_{-i}, \mathbf{b}) \geq 0 \\ \iff \mathbb{E}_{v_i}[u'_i(\mathbf{v}, \mathbf{b})] &\leq b_i + \xi_i(v_{-i}). \end{aligned}$$

In summary, (BU) is equivalent to the follows:

$$\begin{cases} \mathbb{E}_{v_i}[u'_i(\mathbf{v}, \mathbf{b})] \leq b_i + \xi_i(v_{-i}) \\ -b_i \leq \Delta b_i(\mathbf{v}) \leq u'_i(\mathbf{v}, \mathbf{b}) - \mathbb{E}_{v_i}[u'_i(\mathbf{v}, \mathbf{b})] + \xi_i(v_{-i}) \end{cases} \quad (3)$$

Finally, we need the feasibility constraint of the allocation rule that the total allocation of each item is no more than 1:

$$\sum_i x_i(\mathbf{v}, \mathbf{b}) \leq 1. \quad (4)$$

Therefore, combining the objective $\text{REV}(\mathbf{b}, \xi)$ and the constraints (2), (3), and (4), for given \mathbf{b} and g , the optimization program for current period can be written as:

$$\max \text{REV}(\mathbf{b}, \xi) \quad (5)$$

subj.t. $x_i(\mathbf{v})$ is monotone in v_i for any fixed v_{-i}

$$\begin{aligned} \mathbb{E}_{v_i}[u'_i(\mathbf{v})] &\leq b_i + \xi_i(v_{-i}), \forall i, v_{-i} \\ -b_i \leq \Delta b_i(\mathbf{v}) &\leq u'_i(\mathbf{v}) - \mathbb{E}_{v_i}[u'_i(\mathbf{v})] + \xi_i(v_{-i}), \forall i, \mathbf{v} \\ \sum_{i \in [n]} x_i(\mathbf{v}) &\leq 1, \forall \mathbf{v} \\ x_i(\mathbf{v}) &\geq 0, \Delta b_i(\mathbf{v}) \geq 0, \forall i, \mathbf{v} \end{aligned}$$

where we hide the \mathbf{b} in the input to $x_i(\mathbf{v}, \mathbf{b})$ and $u'_i(\mathbf{v}, \mathbf{b})$ to emphasize that \mathbf{b} is extraneous.

Simplification We then simplify the optimization program for optimal bank account mechanisms. One key observation is that with fixed ξ (for all periods), the optimization with any finite horizon can be solved via backward induction.

In the program for the last period T , the expected revenue for the future, g^T , is always zero. For any parameter \mathbf{b} of the last period, the optimal value of the program defines a function with respect to the balance vector \mathbf{b} , which, in fact, is the expected revenue for the last period where the mechanism in the last period M_T^* is optimal (for given \mathbf{b} and ξ).

In other words, the optimal solution of the program for period T defines the function $g^{T-1}(\mathbf{b}|M_T^*)$, where the mechanism for period T is fixed to be optimal. Similarly, the optimal mechanism for each period M_t^* and the corresponding $g^t(\mathbf{b}|M_{t+1..T}^*)$ functions in the program for each period can be determined by backward induction. From now on, by omitting the mechanism being conditional on in g^t , we mean the g^t function is conditional on the optimal $M_{t+1..T}^*$, i.e.,

$$g^t(\mathbf{b}) := g^t(\mathbf{b}|M_{t+1..T}^*) = \max_M g^t(\mathbf{b}|M).$$

Note that the maximum always exists because the domain of the mechanisms is compact and the revenue is a continuous function of the mechanism.

Lemma 3.4. $g^t(\mathbf{b})$ is weakly increasing.

Due to limited space, we omit the proof of Lemma 3.4. Hence we can eliminate variable $\Delta\mathbf{b}(\mathbf{v})$ from the original program (5): Since $g^t(\mathbf{b})$ is weakly increasing, it is without loss of generality to enforce

$$\Delta b_i(\mathbf{v}) = u'_i(\mathbf{v}) - \mathbb{E}_{v_i}[u'_i(\mathbf{v})] + \xi_i(v_{-i}), \forall i, \mathbf{v}.$$

Meanwhile, by the following Myerson's lemma, we can further rewrite $\mathbb{E}_{v_i}[u'_i(\mathbf{v})]$.

Lemma 3.5 (Myerson 1981). *For any incentive compatible single item auction, the expected payment of buyer i equals to the expected (Myerson's) virtual welfare contribution of buyer i :*

$$\forall i, v_{-i}, \mathbb{E}_{v_i}[p_i(\mathbf{v})] = \mathbb{E}_{v_i}[(v_i - \vartheta_i(\mathbf{v})) \cdot x_i(\mathbf{v})],$$

where $\vartheta_i(\mathbf{v}) = v_i \cdot (1 - F_i(v_i))/f_i(v_i)$.⁷

Therefore,⁸

$$\begin{aligned} \mathbb{E}_{v_i}[u'_i(\mathbf{v})] &= \mathbb{E}_{v_i}[v_i \cdot x_i(\mathbf{v}) - p'_i(\mathbf{v})] \\ &= \mathbb{E}_{v_i}[v_i \cdot x_i(\mathbf{v}) - (v_i - \vartheta_i(\mathbf{v})) \cdot x_i(\mathbf{v})] = \mathbb{E}_{v_i}[\vartheta_i(\mathbf{v}) \cdot x_i(\mathbf{v})]. \end{aligned} \quad (6)$$

Thus, program (5) can be simplified as follows, where $\Delta b_i(\mathbf{v})$, $u'_i(\mathbf{v})$, and $\bar{u}'_i(v_{-i})$ are notations, not variables.

$$\begin{aligned} \max \quad & \text{REV}(\mathbf{b}, \xi) \\ &= \mathbb{E}_{\mathbf{v}} [\sum_i (v_i \cdot x_i(\mathbf{v}) - \xi_i(v_{-i})) + g(\Delta\mathbf{b}(\mathbf{v}) + \mathbf{b})] \\ \text{subj.t.} \quad & x_i(\mathbf{v}) \text{ is monotone in } v_i \text{ for fixed } v_{-i} \\ & \bar{u}'_i(v_{-i}) := \mathbb{E}_{v_i}[u'_i(\mathbf{v})] = \mathbb{E}_{v_i}[\vartheta_i(\mathbf{v})x_i(\mathbf{v})] \\ & \leq b_i + \xi_i(v_{-i}), \forall i, v_{-i} \\ & \Delta b_i(\mathbf{v}) = u'_i(\mathbf{v}) - \bar{u}'_i(v_{-i}) + \xi_i(v_{-i}), \forall i, \mathbf{v} \\ & \sum_{i \in [n]} x_i(\mathbf{v}) \leq 1, \forall \mathbf{v} \\ & x_i(\mathbf{v}) \geq 0, \forall i, \mathbf{v} \end{aligned} \quad (7)$$

⁷ F_i must be absolutely integrable.

⁸ $\vartheta_i(\mathbf{v})$ is called *virtual value for utility* and equation (6) is from (Hartline and Roughgarden 2008, Lemma 2.6).

Convexity We remains to show that the program is convex. In fact, we have the following lemma.

Lemma 3.6. $g^t(\mathbf{b})$ is and concave.

We omit the proof of Lemma 3.6 here. Since $g^t(\mathbf{b})$ is concave, the objective function (7) is also concave. Meanwhile, all the constraints are linear, so the program is a convex program (note that a standard convex program is minimizing a convex objective function or maximizing a concave objective function).

3.2 Duality, virtual values, and ironing

We first consider the optimal solution to the program with the monotonicity constraint (2) on $x(\mathbf{v})$ relaxed. In fact, in the full version of this paper, we generalize the analysis to the original program. Although the optimal solution to the relaxed program is not a feasible solution, the analysis does capture the most interesting insight to this problem and provides the explicit form of the virtual values (before ironing). In the full version, we show that adding the monotonicity constraint (2) back to the program corresponds to applying the ironing operation on the virtual values.

For the following relaxed program, let $\lambda_i(v_{-i})$ be the Lagrange multiplier of the constraint $\mathbb{E}_{v_i}[\vartheta_i(\mathbf{v})x_i(\mathbf{v})] \leq b_i + \xi_i(v_{-i})$ and $\mu(\mathbf{v})$ be the Lagrange multiplier of the feasibility constraint (4).

$$\begin{aligned} \max \quad & \text{REV}(\mathbf{b}, \xi) \\ \text{subj.t.} \quad & \mathbb{E}_{v_i}[\vartheta_i(\mathbf{v})x_i(\mathbf{v})] \leq b_i + \xi_i(v_{-i}), \forall v_{-i} \\ & \sum_{i \in [n]} x_i(\mathbf{v}) \leq 1, \quad x_i(\mathbf{v}) \geq 0, \forall i, \mathbf{v} \end{aligned} \quad (8)$$

Then the Lagrange $L := L(\mathbf{x}(\mathbf{v}), \lambda_i(v_{-i}), \mu(\mathbf{v}))$ is

$$\begin{aligned} L &= \mathbb{E}_{\mathbf{v}} [\sum_i (v_i \cdot x_i(\mathbf{v}) - \xi_i(v_{-i})) + g(\Delta\mathbf{b}(\mathbf{v}) + \mathbf{b})] \\ & \quad - \sum_{i, v_{-i}} \lambda_i(v_{-i}) (\mathbb{E}_{v_i}[\vartheta_i(\mathbf{v})x_i(\mathbf{v})] - b_i - \xi_i(v_{-i})) \\ & \quad - \sum_{\mathbf{v}} \mu(\mathbf{v}) (\sum_i x_i(\mathbf{v}) - 1). \end{aligned}$$

Note that all constraints in program (8) are linear, therefore the Slater's condition is satisfied and the KKT conditions (Boyd and Vandenberghe 2004, Chapter 5) are necessary and sufficient for any optimal solution to (8). In particular, the explicit form of the virtual values can be derived from the KKT conditions.

Let g_i denote the partial derivative of g with respect to the i -th dimension, i.e., $g_i(\mathbf{b}) = \partial g(\mathbf{b})/\partial b_i$. Let α and β be defined as follows,

$$\begin{aligned} \alpha_i(\mathbf{v}) &:= 1 + g_i(\Delta\mathbf{b}(\mathbf{v}) + \mathbf{b}) \\ \beta_i(v_{-i}) &:= \frac{\lambda_i(v_{-i})}{f(v_{-i})} + \mathbb{E}_{v_i}[g_i(\Delta\mathbf{b}(\mathbf{v}) + \mathbf{b})]. \end{aligned}$$

Lemma 3.7 (Virtual welfare maximizer). *Any optimal solution to (8) must maximize the expected virtual welfare, where the virtual value function $\phi_i(\mathbf{v})$ is given as follows:*

$$\phi_i(\mathbf{v}) = \alpha_i(\mathbf{v}) \cdot v_i - \beta_i(v_{-i}) \cdot \vartheta_i(\mathbf{v}).$$

We refer readers to the full version for the detailed proof.

Ironing As we mentioned at the beginning of this subsection, the optimal solution to the original program (7) must maximize the expected ironed virtual welfare defined according to the ironed virtual values $\tilde{\phi}_i(\mathbf{v})$. Moreover, the ironed virtual value $\tilde{\phi}_i(\mathbf{v})$ is the transformation of $\phi_i(\mathbf{v})$ according to the ironing rules (Definition 3.1). We omit the details of the analysis here and incorporate the conclusion in the formal version of our first main theorem (Theorem 3.8). We refer readers to the full version for the complete proof.

3.3 Sensitivity analysis and the optimality across the periods

So far we have determined the optimal auction for each period when the expected buyer utilities of each period ξ are fixed. Now we show how to optimize the remaining parameters ξ through sensitivity analysis. Let $R(\xi)$ denote the optimal revenue when the auctions in each period are optimal for the given ξ . Then we show how to determine the partial derivatives of $R(\xi)$ via sensitivity analysis and hence enable the gradient descent algorithm to find the optimal ξ^* . In particular, $R(\xi)$ is a concave function.⁹

The standard sensitivity analysis (Boyd and Vandenberghe 2004, Chapter 5.6) indicates how much the optimal objective value of a program will change if some constraints become slightly looser or tighter. Such quantities usually have important physical meanings in economic setups. For example, consider $g^t(\mathbf{b}) = \max \text{REV}^t(\mathbf{b}, \xi)$ (we brought back the superscript t to distinguish the variables for different periods). By sensitivity analysis,¹⁰ we have:

$$g^t(\mathbf{b}) = \mathbb{E}_{\mathbf{v}^t} [g_i^{t+1}(\Delta \mathbf{b}(\mathbf{v}) + \mathbf{b})] + \sum_{v_{-i}^t} \lambda_i^t(v_{-i}^t).$$

Then $g_i^t(\mathbf{b})$, by definition, is the marginal contribution of the balance of buyer i in period t to the expected revenue since the t -th period. Similarly, $\mathbb{E}_{\mathbf{v}^t} [g_i^{t+1}(\Delta \mathbf{b}(\mathbf{v}) + \mathbf{b})]$ is the marginal contribution of b_i to the expected revenue since the $(t+1)$ -th period. Hence their difference, $\sum_{v_{-i}^t} \lambda_i^t(v_{-i}^t)$, is the marginal contribution of b_i to the expected revenue of the t -th period. In particular, $\lambda_i^t(v_{-i}^t)$ is the marginal contribution if the values of other buyers are v_{-i}^t .

Moreover, we can conclude the concrete form of the partial derivatives of $R(\xi)$ by sensitivity analysis:

$$\begin{aligned} \frac{\partial g^t(\mathbf{b})}{\partial \xi_i^t(v_{-i}^t)} &= f^t(v_{-i}^t) \left(-1 + \mathbb{E}_{\mathbf{v}^t} [g_i^{t+1}(\Delta \mathbf{b}(\mathbf{v}) + \mathbf{b})] \right) + \lambda_i^t(v_{-i}^t) \\ &= (\beta_i^t(v_{-i}^t) - 1) f^t(v_{-i}^t) \\ \implies \frac{\partial R(\xi)}{\partial \xi_i^t(v_{-i}^t)} &= \frac{\partial \mathbb{E}_{\mathbf{b}} [g^t(\mathbf{b})]}{\partial \xi_i^t(v_{-i}^t)} = (\mathbb{E}_{\mathbf{b}} [\beta_i^t(v_{-i}^t)] - 1) f^t(v_{-i}^t), \end{aligned}$$

where the expectation taken over \mathbf{b} is computed by simulating the auctions in previous periods. In particular, ξ is selected optimally, if and only if:

$$\mathbb{E}_{\mathbf{b}} [\beta_i^t(v_{-i}^t)] = 1 \quad \text{or} \quad \xi^t(v_{-i}^t) = 0, \quad \mathbb{E}_{\mathbf{b}} [\beta_i^t(v_{-i}^t)] \leq 1.$$

⁹Since by simply taking any convex combination of different ξ and ξ' , the revenue obtained is also the convex combination of the revenue resulted by ξ and ξ' .

¹⁰If g^t is not differentiable at \mathbf{b} , the right-hand-side is a subgradient of g^t .

3.4 Summary

As a summary, we formally restate Theorem 3.2.

Theorem 3.8 (Formal). *A bank account mechanism is optimal if and only if all the following conditions are satisfied:*

- it satisfies all the basic constraints, (IC), (BI), (BU) and the feasibility constraint (4);
- its allocation rule maximizes the ironed virtual welfare, where the virtual value (before ironing) $\phi(\mathbf{v})$ has the following form:

$$\phi(\mathbf{v}) = \alpha_i(\mathbf{v})v_i - \beta_i(v_{-i})\vartheta_i(\mathbf{v}),$$

which can be seen as the combination of the Myerson's virtual value and the private value;

- finally, the expected utility of each period ξ is selected optimally, i.e.,

$$\mathbb{E}_{\mathbf{b}} [\beta_i^t(v_{-i}^t)] = 1 \quad \text{or} \quad \xi^t(v_{-i}^t) = 0, \quad \mathbb{E}_{\mathbf{b}} [\beta_i^t(v_{-i}^t)] \leq 1.$$

4 An Algorithmic Approach to the Structure

So far we showed that the optimal bank account mechanism maximizes the ironed virtual welfare in each period. Although the explicit form of the virtual values (before ironing) is given in Lemma 3.7 and the ironing rule is given in Definition 3.1, how to accomplish the ironing operation is still unknown. One of the major difficulty of the ironing comes from the fact that one's virtual value not only depends on his/her own private value, but also depends on the private values of other buyers through the term $\alpha_i(\mathbf{v})$. In the presence of such interdependence across different buyers on virtual values, monotone virtual value function does not imply monotone allocation rules.

In this section, we algorithmically resolve the difficulty of ironing. In particular, we show a Fully Polynomial Time Approximation Scheme (FPTAS) that can accomplish the ironing step for any constant many buyer cases and hence compute the ironed virtual values. Moreover, the bank account mechanism induced by the ironed virtual values computed is (multiplicatively) $(1 - \epsilon)$ -approximately optimal in terms of revenue.

We also emphasize that the hard core of computing the exact optimal solution is not directly from the ironing step but the step of approximating the concave functions $g^t(\mathbf{b})$. Note that $g^t(\mathbf{b})$ is a continuous function without closed forms and the main effort of this section is to show that we can arbitrarily approximate $g^t(\mathbf{b})$ with piece-wise linear functions and guarantee that (i) the number of pieces is polynomial in the input size and (ii) the final result is approximately optimal.

Recall the program (7), in particular, we bring back the superscripts t to emphasize the periods:

$$\begin{aligned} \max \quad & g^{t-1}(\mathbf{b}^t) = \text{REV}(\mathbf{b}^t, \xi^t) \\ &= \mathbb{E}_{\mathbf{v}^t} [\sum_i (v_i^t x_i^t(\mathbf{v}^t) - \xi_i^t(v_{-i}^t)) + g^t(\Delta \mathbf{b}^t(\mathbf{v}^t) + \mathbf{b}^t)] \\ \text{subj.t.} \quad & x_i^t(v_i^t, v_{-i}^t) - x_i^t(v_i^{t'}, v_{-i}^t) \leq 0, \quad \forall i, v_{-i}^t, v_i^{t'} > v_i^t \\ & \mathbb{E}_{\mathbf{v}_i^t} [u_i^{t'}] = \mathbb{E}_{\mathbf{v}_i^t} [\vartheta_i^t(\mathbf{v}^t) x_i^t(\mathbf{v}^t)] \leq b_i^t + \xi_i^t(v_{-i}^t), \quad \forall v_{-i}^t \\ & \Delta b_i^t(\mathbf{v}^t) = u_i^{t'}(\mathbf{v}^t) - \bar{u}_i^t(v_{-i}^t) + \xi_i^t(v_{-i}^t), \quad \forall i, \mathbf{v}^t \\ & \sum_{i \in [n]} x_i^t(\mathbf{v}^t) \leq 1, \quad x_i^t(\mathbf{v}^t) \geq 0, \quad \forall i, \mathbf{v}^t \end{aligned}$$

In what follows, we formalize the FPTAS via dynamic programming to compute the optimal ironed virtual values for the discrete type case.

Theorem 4.1. *The ironed virtual values of the optimal bank account mechanism can be computed through a dynamic programming based algorithm.*

Moreover, for any $\epsilon > 0$, there is an FPTAS to achieve an ϵ -approximation (multiplicative) of the optimal revenue.

We first outline the main idea of the algorithm: (i) for any fixed ξ , compute the ironed virtual values of the approximately optimal bank account mechanism; (ii) compute the optimal ξ using gradient descent. Since we have shown that the revenue of the bank account mechanism is a concave function with respect to ξ , the second step is standard and can be done with polynomially many queries to the (approximately) optimal revenue as a function of ξ . In what follows, we will focus on the first step.

To compute the ironed virtual value, we need to solve the dual program of (7), which, of course, is equivalent to solve the primal because the strong duality holds. Note that when ξ and \mathbf{b}^t are fixed, (7) is a standard convex program with polynomially many linear constraints. Hence the optimal solution to its dual could be computed efficiently given oracle accesses to the concave function g^t . Then by the definition of g^{t-1} , the value of the optimal solution is $g^{t-1}(\mathbf{b}^t)$. Therefore, the concave function g^t can be evaluated recursively for each t and hence the ironed virtual values of the optimal bank account mechanism can be computed via standard dynamic programming as well.

However, the computation of the entire dynamic programming is not directly efficient: (i) if each g^t is computed recursively upon every query, the depth of the recursion could be up to T and hence the total number of queries required would be exponential in T ; (ii) if each g^t is computed once for all possible \mathbf{b} so that any further queries of g^t can be answered from precalculated values, then the number of input points where g^t need to be computed would be unbounded.

The key step to resolve these issues is to approximate the concave functions $g^t(\mathbf{b})$ by piece-wise linear functions each with at most polynomially many pieces. In addition, each of the piece-wise linear functions can be further expressed as the minimum of a set of affine functions. Hence both the original convex program (7) at each period t and its dual can be approximated by a poly-size linear program.

The following lemma ensures that each g^t can be well approximated by a piece-wise linear function.

Lemma 4.2. *For all $t \in [T]$, g^t can be κ -approximated by two concave piece-wise linear functions, \underline{g}^t and \bar{g}^t :*

$$\forall \mathbf{b}, \underline{g}^t(\mathbf{b}) \leq g^t(\mathbf{b}) \leq \bar{g}^t(\mathbf{b}) \leq \underline{g}^t(\mathbf{b}) + \kappa \max_{\mathbf{b}'} g^t(\mathbf{b}').$$

Moreover, each of \underline{g}^t and \bar{g}^t can be written as the minimum of at most polynomially many pieces.

With such approximations to the concave function g^t , we can approximately solve the convex program (7) by solving the following linear program, where each occurrence of g^t

is replaced by \bar{g}^t :

$$\max \quad \bar{h}^{t-1}(\mathbf{b}^t) := \mathbb{E}_{\mathbf{v}^t} [\sum_i (v_i^t \cdot x_i^t(\mathbf{v}^t) - \xi_i^t(v_{-i}^t)) + \bar{g}^t(\Delta \mathbf{b}^t(\mathbf{v}^t) + \mathbf{b}^t)] \quad (9)$$

$$\begin{aligned} \text{sub.j.t.} \quad & x_i^t(v_i^t, v_{-i}^t) - x_i^t(v_i^{t'}, v_{-i}^t) \leq 0, \forall i, v_{-i}^t, v_i^{t'} > v_i^t \\ & \mathbb{E}_{v_i^t} [u_i^t] = \mathbb{E}_{v_i^t} [\vartheta_i^t(\mathbf{v}^t) x_i^t(\mathbf{v}^t)] \leq b_i^t + \xi_i^t(v_{-i}^t), \forall v_{-i}^t \\ & \Delta b_i^t(\mathbf{v}^t) = u_i^{t'}(\mathbf{v}^t) - \bar{u}_i^t(v_{-i}^t) + \xi_i^t(v_{-i}^t), \forall i, \mathbf{v}^t \\ & \sum_{i \in [n]} x_i^t(\mathbf{v}^t) \leq 1, x_i^t(\mathbf{v}^t) \geq 0, \forall i, \mathbf{v}^t \\ & \bar{g}^t(\Delta \mathbf{b}^t(\mathbf{v}^t) + \mathbf{b}^t) \\ & \leq \alpha_l \cdot (\Delta \mathbf{b}^t(\mathbf{v}^t) + \mathbf{b}^t) + \beta_l, \forall l, \mathbf{v}^t \end{aligned} \quad (10)$$

As we mentioned previously, in the last constraint (10), we assume that the function \bar{g}^t can be expressed by the minimum of a set of affine functions, i.e.,

$$\bar{g}^t(\mathbf{b}) = \min_{l \in L} \alpha_l \cdot \mathbf{b} + \beta_l.$$

Here we slightly abuse the notation of $\bar{g}^t(\Delta \mathbf{b}^t(\mathbf{v}^t) + \mathbf{b}^t)$ as variables in the linear program. Note that having the constraints (10) in the linear program (9) suffices to ensure that the variable values agree with the corresponding function values. Because on one hand, by the constraints, each variable is no more than the corresponding affine functions; on the other hand, since the coefficients of these variables are always positive in the objective, for each of the variables, at least one of the inequalities in (10) must be binding.

Let $\bar{h}^{t-1}(\mathbf{b})$ denote the optimal value of the linear program (9) when $\mathbf{b}^t = \mathbf{b}$. Similarly, we can define \underline{h}^{t-1} by replacing all \bar{g}^t with \underline{g}^t in (9). The following lemma shows that \underline{h}^{t-1} and \bar{h}^{t-1} are lower and upper bounds of g^{t-1} .

Lemma 4.3. *\underline{h}^{t-1} and \bar{h}^{t-1} are concave and for any \mathbf{b} ,*

$$\underline{h}^{t-1}(\mathbf{b}) \leq g^{t-1}(\mathbf{b}) \leq \bar{h}^{t-1}(\mathbf{b})$$

$$\text{and } \bar{h}^{t-1}(\mathbf{b}) \leq \underline{h}^{t-1}(\mathbf{b}) + \max_{\mathbf{b}'} (\bar{g}^t(\mathbf{b}') - \underline{g}^t(\mathbf{b}')).$$

Proof of Lemma 4.3. Let $\underline{\mathbf{x}}^*$, \mathbf{x}^* , and $\bar{\mathbf{x}}^*$ be the optimal solutions for $\underline{h}^{t-1}(\mathbf{b})$, $g^{t-1}(\mathbf{b})$, and $\bar{h}^{t-1}(\mathbf{b})$, respectively. Denote these three programs as \underline{P} , P , and \bar{P} , respectively. Note that $\underline{\mathbf{x}}^*$ is feasible in P , hence by the optimality of \mathbf{x}^* ,

$$P(\underline{\mathbf{x}}^*) \leq P(\mathbf{x}^*),$$

where $P(\mathbf{x})$ is the objective value of program P on \mathbf{x} .

On the other hand, for any \mathbf{x} that is feasible to both \underline{P} and \bar{P} , we have,

$$\underline{P}(\mathbf{x}) \leq P(\mathbf{x}).$$

Because in the objectives, $\underline{g}^{t-1} \leq g^{t-1}$. Hence we conclude:

$$\underline{h}^{t-1}(\mathbf{b}) = \underline{P}(\underline{\mathbf{x}}^*) \leq P(\underline{\mathbf{x}}^*) \leq P(\mathbf{x}^*) = g^{t-1}(\mathbf{b}).$$

Similarly, we can prove that $g^{t-1}(\mathbf{b}) \leq \bar{h}^{t-1}(\mathbf{b})$. For the last inequality, denote $\delta = \max_{\mathbf{b}'} (\bar{g}^t(\mathbf{b}') - \underline{g}^t(\mathbf{b}'))$. Note that for the \bar{g} variables in $\bar{\mathbf{x}}^*$, if we reduce all of them by δ to get $\bar{\mathbf{x}}^{**}$, $\bar{\mathbf{x}}^{**}$ must be feasible to \underline{P} and the objective value is reduced by at most δ , hence

$$\underline{P}(\mathbf{x}^*) \geq \underline{P}(\bar{\mathbf{x}}^{**}) \geq \bar{P}(\bar{\mathbf{x}}^*) - \delta.$$

In other words,

$$\underline{h}^{t-1}(\mathbf{b}) \geq \bar{h}^{t-1}(\mathbf{b}) - \delta.$$

□

Even through both \underline{h}^{t-1} and \bar{h}^{t-1} are in fact piece-wise linear functions, but they cannot be directly used for the computation of period $t - 1$, because they may have exponentially many pieces. However, since they are concave, we can apply Lemma 4.2 to get the lower bound of \underline{h}^{t-1} and the upper bound of \bar{h}^{t-1} :

$$\underline{g}^{t-1}(\mathbf{b}) \leq \underline{h}^{t-1}(\mathbf{b}) \leq g^{t-1}(\mathbf{b}) \leq \bar{h}^{t-1}(\mathbf{b}) \leq \bar{g}^{t-1}(\mathbf{b}).$$

Therefore, we can compute $\underline{g}^1, \dots, \underline{g}^T$ and $\bar{g}^1, \dots, \bar{g}^T$ recursively, and then compute the ironed virtual values of the approximately optimal bank account mechanism for any fixed ξ . Because ξ can then be optimized by using gradient descent, we are done with our algorithm.

References

- Ashlagi, I.; Daskalakis, C.; and Haghpahan, N. 2016. Sequential mechanisms with ex-post participation guarantees. In *EC'16*, 213–214.
- Athey, S., and Segal, I. 2013. An efficient dynamic mechanism. *Econometrica* 81(6):2463–2485.
- Balseiro, S.; Lin, M.; Mirrokni, V.; Paes Leme, R.; and Zuo, S. 2017. Dynamic revenue sharing. In *NIPS'17*, 2678–2686.
- Balseiro, S.; Mirrokni, V.; Paes Leme, R.; and Zuo, S. 2019. Dynamic double auctions: Towards first best. In *SODA'19*.
- Balseiro, S. R.; Mirrokni, V. S.; and Leme, R. P. 2017. Dynamic mechanisms with martingale utilities. forthcoming, *Management Science*.
- Balseiro, S.; Mirrokni, V.; and Paes Leme, R. 2016. Dynamic mechanisms with martingale utilities.
- Bergemann, D., and Said, M. 2011. Dynamic auctions. *Wiley Encyclopedia of Operations Research and Management Science*.
- Bergemann, D., and Välimäki, J. 2002. Information acquisition and efficient mechanism design. *Econometrica* 70(3):1007–1033.
- Bergemann, D., and Välimäki, J. 2017. Dynamic mechanism design: An introduction. Technical report, Cowles Foundation for Research in Economics, Yale University.
- Bergemann, D.; Castro, F.; and Weintraub, G. 2017. The scope of sequential screening with ex post participation constraints. In *EC'17*, 163–164.
- Boyd, S., and Vandenberghe, L. 2004. *Convex optimization*. Cambridge university press.
- Cai, Y., and Zhao, M. 2017. Simple mechanisms for subadditive buyers via duality. In *STOC'17*, 170–183.
- Cai, Y.; Devanur, N. R.; and Weinberg, S. M. 2016. A duality-based unified approach to bayesian mechanism design. *ACM SIGecom Exchanges* 15(1):71–77.
- Cavallo, R. 2008. Efficiency and redistribution in dynamic mechanism design. In *EC'08*, 220–229.
- Chawla, S.; Devanur, N. R.; Karlin, A. R.; and Sivan, B. 2016. Simple pricing schemes for consumers with evolving values. In *SODA'16*, 1476–1490.
- Daskalakis, C.; Deckelbaum, A.; and Tzamos, C. 2013. Mechanism design via optimal transport. In *EC'13*, 269–286.
- Daskalakis, C.; Deckelbaum, A.; and Tzamos, C. 2017. Strong duality for a multiple-good monopolist. *Econometrica* 85(3):735–767.
- Devanur, N. R.; Peres, Y.; and Sivan, B. 2015. Perfect bayesian equilibria in repeated sales. In *SODA'15*, 983–1002.
- Elkind, E. 2007. Designing and learning optimal finite support auctions. In *SODA'07*, 736–745. SIAM.
- Hartline, J. D., and Roughgarden, T. 2008. Optimal mechanism design and money burning. In *STOC'08*, 75–84.
- Jackson, M. O., and Sonnenschein, H. F. 2007. Overcoming incentive constraints by linking decisions. *Econometrica* 75(1):241–257.
- Kakade, S. M.; Lobel, I.; and Nazerzadeh, H. 2013. Optimal dynamic mechanism design and the virtual-pivot mechanism. *Operations Research* 61(4):837–854.
- Lobel, I., and Paes Leme, R. 2017. Dynamic mechanism design under positive commitment.
- Mirroknii, V.; Paes Leme, R.; Tang, P.; and Zuo, S. 2016a. Dynamic auctions with bank accounts. In *IJCAI'16*, 387–393.
- Mirroknii, V.; Paes Leme, R.; Tang, P.; and Zuo, S. 2016b. Optimal dynamic mechanisms with ex-post ir via bank accounts. *arXiv preprint arXiv:1605.08840*.
- Mirroknii, V.; Paes Leme, R.; Ren, R.; and Zuo, S. 2018a. Dynamic mechanism design in the field. In *WWW'18*, 1359–1368.
- Mirroknii, V.; Paes Leme, R.; Tang, P.; and Zuo, S. 2018b. Non-clairvoyant dynamic mechanism design. In *EC'18*, 169–169.
- Myerson, R. B. 1981. Optimal auction design. *Mathematics of operations research* 6(1):58–73.
- Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V. 2007. *Algorithmic game theory*, volume 1. Cambridge University Press Cambridge.
- Pai, M. M., and Vohra, R. 2013. Optimal dynamic auctions and simple index rules. *Mathematics of Operations Research* 38(4):682–697.
- Papadimitriou, C.; Pierrakos, G.; Psomas, C.-A.; and Rubinfeld, A. 2016. On the complexity of dynamic mechanism design. In *SODA'16*, 1458–1475.
- Parkes, D. C., and Singh, S. 2004. An mdp-based approach to online mechanism design.
- Pavan, A.; Segal, I.; and Toikka, J. 2014. Dynamic mechanism design: A myersonian approach. *Econometrica* 82(2):601–653.
- Rochet, J.-C. 1985. The taxation principle and multi-time hamilton-jacobi equations. *Journal of Mathematical Economics* 14(2):113–128.
- Roughgarden, T., and Schrijvers, O. 2016. Ironing in the dark. In *EC'16*, 1–18.
- Shen, W.; Wang, Z.; and Zuo, S. 2018. Ex-post IR dynamic auctions with cost-per-action payments. In *IJCAI'18*, 505–511.