
Robust Pricing in Dynamic Mechanism Design

Yuan Deng¹ Sébastien Lahaie² Vahab Mirrokni²

Abstract

Motivated by the repeated sale of online ads via auctions, optimal pricing in repeated auctions has attracted a large body of research. While dynamic mechanisms offer powerful techniques to improve on both revenue and efficiency by optimizing auctions across different items, their reliance on exact distributional information of buyers' valuations (present and future) limits their use in practice. In this paper, we propose *robust dynamic mechanism design*. We develop a new framework to design dynamic mechanisms that are robust to both estimation errors in value distributions and strategic behavior. We apply the framework in learning environments, leading to the first policy that achieves provably low regret against the optimal *dynamic* mechanism in contextual auctions, where the dynamic benchmark has full and accurate distributional information.

1. Introduction

Motivated by the popularity of selling online ads via auctions, pricing in dynamic auctions has been extensively studied in recent years. Dynamic auctions open up the possibility of linking the auction rules and payments across time to enhance revenue or welfare. Formally, *dynamic mechanism design* considers an environment in which the seller has exact distributional information over the buyers' values for the items, for the current stage and all future stages, and designs revenue-maximizing dynamic mechanisms that adapt the auction rules based on the buyer's historical bids (Thomas & Worrall, 1990; Bergemann & Välimäki, 2010; Ashlagi et al., 2016; Mirrokni et al., 2016a; 2018; Deng et al., 2019b). This line of study provides simple dynamic mechanisms, in terms of descriptive complexity, that compare favorably to the revenue-optimal dynamic benchmark. However, these mechanisms are *clairvoyant* and rely on exact knowledge of

future value distributions to align buyers' incentives (across time) and guarantee desirable outcomes. This strong requirement limits the application of dynamic auctions in practice because the seller may only have access to approximate models of distributions. As one attempt to address this concern, Mirrokni et al. (2018) consider another extreme and propose non-clairvoyant dynamic mechanisms, which do not rely on any information about the future. They show that a non-clairvoyant dynamic mechanism can achieve a constant approximation of the revenue-optimal clairvoyant dynamic mechanism that knows the future in advance.

In this work we take an intermediate stance where the designer can model present and future distributional information subject to an estimation error. Designing revenue-optimal and incentive-compatible auctions in this framework is challenging for the following reason: when the seller's distributional information is not perfectly aligned with the buyer's true value distributions, it is impossible for the seller to offer a prior-dependent dynamic mechanism in which the optimal strategy for the buyer is to report her valuation truthfully at every stage (i.e., dynamic incentive-compatible). Furthermore, in a dynamic mechanism the buyer's misreport can potentially affect auctions for all future items. We overcome these obstacles and provide a *robust dynamic mechanism where the extent of the buyers' misreports and the revenue loss can be related to and bounded by the estimation error of the buyers' distributions*.

We then apply our robust dynamic mechanism to the problem of robust price learning. In particular, we focus on contextual auctions, where a buyer's valuation for an item depends on the context that describes the item, but the relationship between the buyer's valuation and the context is unknown to the seller. The seller's task is to design a policy to adapt the mechanism based on the buyer's historical bids, with the objective of maximizing revenue. Previous results (Amin et al., 2014; Golrezaei et al., 2019) give a no-regret policy against the optimal *static* mechanism in which the auction ignores the history and does not evolve over time. However, Papadimitriou et al. (2016) have shown that the revenue gap between optimal static and dynamic mechanisms can be arbitrarily large. We tailor the structure of our robust dynamic mechanism to a learning environment, leading to a *no-regret policy against the optimal clairvoyant contextual auction that knows the relationship in advance*.

¹Department of Computer Science, Duke University, Durham, NC, USA ²Google Research, New York City, NY, USA. Correspondence to: Yuan Deng <ericdy@cs.duke.edu>.

Related Work

Our work is closely related to the recent work by [Deng et al. \(2019a\)](#), which provides a robust dynamic mechanism design framework for the non-clairvoyant environment. They provide a no-regret policy in contextual auctions against a constant approximation of the optimal clairvoyant mechanism. In contrast, we provide a robust dynamic mechanism design framework for the clairvoyant environment and design a no-regret policy against the optimal clairvoyant contextual auction without any approximation.

Moreover, robust dynamic mechanism design for the clairvoyant environment is more challenging than robust dynamic mechanism design for the non-clairvoyant environment. In a non-clairvoyant environment, there exists a concrete dynamic mechanism from [Mirrokni et al. \(2018\)](#), which is a mixture of the give-for-free auction, the posted-price auction with an entry fee, and the Myerson’s auction. Using such a concrete mechanism as a starting point, one only needs to provide a framework to make this mechanism robust ([Deng et al., 2019a](#)). In contrast, we consider a clairvoyant environment, and the revenue-optimal dynamic mechanism (with perfect prior knowledge) is given by a convex program based dynamic program, i.e., it is computed via a dynamic program in which each transition is computed by a convex program (see Section 3.3). To overcome the difficulty of analyzing a convex program based dynamic program, we both develop new technical tools for analysis and provide structural insights of its optimal solution.

Dynamic Mechanism Design. For a review of the literature, readers are encouraged to refer to ([Bergemann & Välimäki, 2019](#)) for a comprehensive survey. [Bergemann & Välimäki \(2010\)](#) propose a generalized VCG mechanism to the dynamic environment where the buyers receive private information over time, called the dynamic pivot mechanism, which achieves welfare-maximizing outcomes. [Kakade et al. \(2013\)](#) combine the dynamic pivot mechanism and the virtual valuation idea ([Myerson, 1981](#)) to design a virtual-pivot mechanism. [Athey & Segal \(2013\)](#) propose a team mechanism that is efficient and budget-balanced.

The line of research on revenue-maximizing dynamic mechanism design was initiated by [Baron & Besanko \(1984\)](#) and [Courty & Hao \(2000\)](#). [Pavan et al. \(2014\)](#) generalize the Myersonian approach ([Myerson, 1981](#)) to the dynamic setting and provide characterizations of dynamic incentive-compatibility. [Papadimitriou et al. \(2016\)](#) provide an example that demonstrates the revenue gap between the static and dynamic mechanism can be arbitrarily large. Moreover, they show that it is NP-Hard to design the optimal deterministic auctions even in a dynamic environment with a single buyer and two items only. [Ashlagi et al. \(2016\)](#) and [Mirrokni et al. \(2016b\)](#) independently provide fully polynomial-time approximation schemes to compute the

optimal randomized mechanism. Our work is mainly built on top of the framework of *bank account mechanisms* from ([Mirrokni et al., 2018](#); [Deng et al., 2019b](#)), which relies on exact knowledge of valuation distributions. They provide a general framework to design the revenue-maximizing dynamic mechanism, called bank account mechanisms. Inspired by the framework, [Deng & Lahaie \(2019\)](#) and [Deng et al. \(2020\)](#) provide statistical tools to test and measure dynamic incentive compatibility. However, such a framework considers a setting where the seller has a perfect information about the buyer’s distributions. In contrast, our robust dynamic mechanism works in an environment where the seller’s distributional information is not perfect.

Robust Price Learning. Our work is also related to dynamic pricing with learning (see [den Boer \(2015\)](#) for a recent survey). There has been a growing body of literature on price learning with non-strategic buyers ([Cohen et al., 2016](#); [Lobel et al., 2018](#); [Leme & Schneider, 2018](#); [Mao et al., 2018](#)). In their models, the buyers have fixed valuations and are non-strategic, and therefore, the problem can be reduced to a one-shot auction where the buyer acts myopically without considering future. However, [Edelman & Ostrovsky \(2007\)](#) provide empirical evidence that the buyers participating in the online advertising markets do act strategically. The study of robust price learning with strategic buyers was initiated by [Amin et al. \(2013\)](#) and [Medina & Mohri \(2014\)](#). When the valuations are fixed and the buyers are impatient, the revenue regret has been shown to be $\Theta(\log \log T)$ by [Drutsa \(2017; 2018\)](#). For learning in the contextual auctions, [Amin et al. \(2014\)](#) develop a no-regret policy in a setting without market noise. Recently, [Golrezaei et al. \(2019\)](#) enrich the model by incorporating market noise. All of these results are no-regret against optimal *static* mechanisms that ignore the history, while our policy is no-regret against optimal *dynamic* mechanisms.

2. Preliminaries

A dynamic auction model describes an environment where a seller (he) sells a stream of T items that arrive online, based on the reports by strategic buyers. In an online environment, an item must be sold once it arrives. For the sake of clarity, we will focus on the case with a single buyer (she). Our results can be extended to multi-buyer settings by using the techniques from [Cai et al. \(2012\)](#).

In line with the literature ([Deng et al., 2019a](#)), the t -th item arrives at stage t and the buyer’s valuation $v_t \in [0, a_t]$ is drawn independently (but not necessarily identically) according to the cumulative distribution function F_t . We assume that the density function f_t of F_t is upper bounded by c_f/a_t where c_f is a constant. The domain bounds a_t are public and enrich the model to reflect the fact that item valuations may have different scales. We normalize the domain

bound sequence so that $\sum_t a_t = T$. We consider a setting where the seller's distributional information is imperfect: the seller only has access to an *estimated* distribution \hat{F}_t .

After the buyer learns her valuation v_t at the beginning of stage t , she then submits a bid b_t to the seller who then implements an outcome with an allocation probability and a payment. We restrict our attention to the case where the bid b_t is always in the set $V_t = [0, a_t]$. For convenience, let $V^t = \prod_{t'=1}^t V_{t'}$ be the set of all possible sequences of the buyer's bids for the first t stages. Similarly, let ΔV_t be the set of distributions over V_t and let $(\Delta V)^t = \prod_{t'=1}^t (\Delta V_{t'})$ be the set of all possible sequences of distributions for the first t stages. For convenience, we use the notation $a_{(t', t'')}$ to represent a sequence $(a_{t'}, \dots, a_{t''})$ of a between stage t' and stage t'' . In general, a clairvoyant mechanism can be characterized by a sequence of allocation and payment functions: (1) the allocation function maps historical bids and seller's distributional information $\hat{F}_{(1,T)}$ to an allocation probability: $x_t : V^t \times (\Delta V)^T \rightarrow [0, 1]$; (2) the payment function maps historical bids and seller's distributional information $\hat{F}_{(1,T)}$ to a payment: $p_t : V^t \times (\Delta V)^T \rightarrow \mathbb{R}$. Given $b_{(1,t)}$ and $\hat{F}_{(1,T)}$, the utility u_t of the buyer with true valuation v_t is $u_t(v_t; b_{(1,t)}; \hat{F}_{(1,T)}) = v_t \cdot x_t(b_{(1,t)}; \hat{F}_{(1,T)}) - p_t(b_{(1,t)}; \hat{F}_{(1,T)})$. A dynamic mechanism is non-clairvoyant if no prior knowledge about future stages is available, and therefore, the allocation rule and the payment rule of a non-clairvoyant dynamic mechanism at stage t can only depend on $\hat{F}_{(1,t)}$ and $b_{(1,t)}$. We will focus on how to make the revenue-optimal clairvoyant mechanism robust.

Estimated Distributional Information. Following the setup in (Deng et al., 2019a), we relax the standard assumption of exact distributional information (Ashlagi et al., 2016; Mirrokni et al., 2018; Deng et al., 2019b) and consider an environment where the seller's distributional information is estimated with an error Δ .

Assumption 2.1 (Deng et al. (2019a)). *There exists a coupling between a random draw v_t from F_t and a random draw \hat{v}_t from \hat{F}_t such that $v_t = \hat{v}_t + a_t \cdot \epsilon_t$ with $\epsilon_t \in [-\Delta, \Delta]$.*

The assumption states that samples from the estimated distribution have a bounded bias. Looking ahead to our application into contextual auctions, such a bias comes from that the seller does not have perfect information about the relationship between the buyer's valuation and the context.

Utility-maximizing Buyer. We assume the buyer's valuation is additive across items. At stage t , the buyer aims at maximizing her time-discounted cumulative expected utility $\sum_{t'=t}^T \gamma^{t'-t} \cdot \mathbb{E}[u_{t'}]$, where $\gamma \in (0, 1)$ is the discounting factor and the expectation is taken with respect to the true distribution $F_{(1,T)}$. The discounting factor implies that the buyer is less patient than the seller. We note that Amin et al. (2013) showed that it is impossible to obtain a no-regret

policy when the buyer is as patient as the seller.

Incentive Constraints. The buyer's best response in a dynamic mechanism depends on her strategy in future stages. When the seller has perfect distributional information, the classic notion of dynamic incentive-compatibility (DIC) requires that reporting truthfully is always the buyer's optimal strategy, assuming that she plays optimally in the future (Mirrokni et al., 2018). However, exact DIC is no longer possible to achieve in prior-dependent dynamic mechanisms when the seller only has approximate distributional information. We consider $\eta_{(1,T)}$ -approximate DIC (Deng et al., 2019a): assuming the buyer plays optimally in the future (optimally now no longer means truthfully), the buyer's bid should deviate from v_t by at most η_t at stage t . Formally, there exists $\hat{b}_t \in [v_t - \eta_t, v_t + \eta_t]$ that belongs to

$$\arg \max_{b_t} u_t(v_t; b_{(1,t)}; \hat{F}_{(1,T)}) + \gamma \cdot U_t(b_{(1,t)}; F_{(1,T)}; \hat{F}_{(1,T)}) \quad (\eta_{(1,T)}\text{-DIC})$$

for all $v_t, b_{(1,t-1)}$. Here $U_t(b_{(1,t)}; F_{(1,T)}; \hat{F}_{(1,T)})$ is the *continuation utility* that the buyer obtains in the future: for $t < T$, $U_t(b_{(1,t)}; F_{(1,T)}; \hat{F}_{(1,T)})$ is defined recursively as

$$\mathbb{E}_{v_{t+1} \sim F_{t+1}} \left[\max_{b_{t+1}} u_{t+1}(v_{t+1}; b_{(1,t+1)}; \hat{F}_{(1,T)}) + \gamma \cdot U_{t+1}(b_{(1,t+1)}; F_{(1,T)}; \hat{F}_{(1,T)}) \right],$$

while $U_T(b_{(1,T)}; F_{(1,T)}; \hat{F}_{(1,T)}) = 0$.

Participation Constraints. We assume that the buyer weights realized past utilities equally, and therefore, ex-post individual rationality requires that for all $v_{(1,T)}$:

$$\sum_{t=1}^T u_t(v_t; v_{(1,t)}; \hat{F}_{(1,T)}) \geq 0 \quad (\text{ex-post IR})$$

For convenience, we will use the phrasing ‘‘for $F_{(1,T)}$ ’’ to refer to the environment where the true distributions are $F_{(1,T)}$. For example, that a mechanism is $\eta_{(1,T)}$ -DIC for $F_{(1,T)}$ means that the mechanism is $\eta_{(1,T)}$ -DIC when the true distributions are $F_{(1,T)}$.

2.1. Bank Account Mechanism

Even in an environment where the seller has perfect distributional information $\hat{F}_{(1,T)} = F_{(1,T)}$, the first challenge in designing a dynamic mechanism is that the descriptive complexity of the mechanism could be exponentially large. In general, the allocation functions and the payment functions depend on the entire sequence of historical bids. We will focus on a simpler, special class of dynamic mechanisms, called *bank account mechanisms*. For our purposes, this is without loss of generality, because Mirrokni et al. (2018) showed that any dynamic incentive-compatible and ex-post

individual rational mechanism can be converted to a bank account mechanism without loss of revenue.

The salient feature of a bank account mechanism is that it uses a single non-negative real number bal_t , called *bank account balance*, to summarize the history. Henceforth, the allocation and payment function at stage t only depends on bal_t , b_t , and the seller's distributional information.

Definition 2.2 (Bank Account Mechanism (Mirrokni et al., 2018)). *A bank account mechanism $B = \langle x, p, \text{balU} \rangle$ for $\hat{F}_{(1,T)}$ is specified by a tuple constituted by x , p , and balU such that for each stage t :*

(1) *A stage mechanism $x_t(\text{bal}, b_t)$, $p_t(\text{bal}, b_t)$ is parameterized by a balance $\text{bal} \in \mathbb{R}_+$, which is incentive-compatible for the stage for every $\text{bal} \geq 0$: for any v_t and b_t ,*

$$v_t \cdot x_t(\text{bal}, v_t) - p_t(\text{bal}, v_t) \geq v_t \cdot x_t(\text{bal}, b_t) - p_t(\text{bal}, b_t); \quad (\text{stage-IC})$$

(2) *The mechanism is not necessarily individual rational for the stage. However, the expected utility is balance independent if the buyer reports truthfully:*

$$\mathbb{E}_{v_t \sim \hat{F}_t} [v_t \cdot x_t(\text{bal}, v_t) - p_t(\text{bal}, v_t)] = c_t, \quad (\text{BI})$$

where c_t is a constant not dependent on bal ;

(3) *A balance update policy $\text{balU}_t : \mathbb{R}_+ \times V \rightarrow \mathbb{R}_+$ that maps the previous balance and the buyer's bid to a new balance, satisfying $\text{balU}_{t+1}(\text{bal}_t, b_t) \geq 0$ and*

$$\text{balU}_{t+1}(\text{bal}_t, b_t) \leq \text{bal}_t + b_t \cdot x_t(\text{bal}_t, b_t) - p_t(\text{bal}_t, b_t). \quad (\text{BU})$$

bal_{t+1} can be defined recursively as $\text{bal}_1 = 0$ and

$$\text{bal}_{t+1}(b_{(1,t)}) = \text{balU}_{t+1}(\text{bal}_t(b_{(1,t-1)}), b_t).$$

We will use the notation $x_t(b_{(1,t)})$ and $x_t(\text{bal}, b_t)$ where $\text{bal} = \text{bal}_t(b_{(1,t-1)})$ interchangeably since $x_t(b_{(1,t)})$ and $x_t(\text{bal}, b_t)$ are the same; similarly for p_t .

Note that (BI) implies that the buyer's historical reports have no impact on her future expected utilities, assuming she reports truthfully in the future. Therefore, when the seller has perfect distributional information, if the stage mechanism for every stage is *stage-IC*, a backward induction argument can demonstrate that the mechanism is exactly DIC. Moreover, (BU) ensures that the non-negative balance always lower bounds the buyer's utility provided truthful reporting. Thus, the bank mechanism is *ex-post IR*.

3. Core Bank Account Mechanism

It is inconvenient to directly analyze the bank account mechanism since we need to relate the stage mechanisms with different $\text{bal} \geq 0$ to ensure that (BI) is satisfied. A refined

characterization called *core* bank account mechanism (Mirrokni et al., 2016b), provides a more convenient way to ensure (BI). The full proofs of this section are deferred to the full version. After introducing the notion of a core bank account mechanism (Definition 3.1), we develop a novel program to compute the revenue of such a mechanism, even when the seller's distributional information is imperfect (Section 3.1). We then introduce basic operations for editing core bank account mechanisms (Section 3.2), which enable a unification of core bank account mechanisms (Lemma 3.5) as well as a dynamic program for computing the revenue-optimal mechanism when the seller's distributional information is perfect (OPT-BAM). The latter will serve as the base mechanism for our robust mechanisms.

Definition 3.1 (Core Bank Account Mechanism (Mirrokni et al., 2016b)). *A core bank account mechanism $\langle g, y \rangle$ for $\hat{F}_{(1,T)}$ is constituted by a family of functions $g = g_{(1,T)}$ and $y = y_{(1,T)}$. g_t maps a history $b_{(1,t)}$ to a non-negative real number and y_t maps $b_{(1,t)}$ to the stage allocation. Moreover*

(1) *y_t is the sub-gradient of g_t with respect to b_t ;*

(2) *g_t is consistent, symmetric, convex in b_t , and weakly increasing in b_t , where g is consistent if*

$$g_{t-1}(b_{(1,t-1)}) - \mathbb{E}_{v_t \sim \hat{F}_t} [g_t(b_{(1,t-1)}, v_t)] = \chi_t(g),$$

where $\chi_t(g)$ is a number dependent on g but independent of $b_{(1,t-1)}$; g is symmetric if $g_{t-1}(b_{(1,t-1)}) = g_{t-1}(b'_{(1,t-1)})$ implies $g_t(b_{(1,t-1)}, b_t) = g_t(b'_{(1,t-1)}, b_t)$.

A bank account mechanism $B(g, y; \hat{F}_{(1,T)})$ satisfying *stage-IC*, *BI*, and *BU* for $\hat{F}_{(1,T)}$ can be constructed from a core bank account mechanism $\langle g, y \rangle$ as follows:

$$\text{bal}_{t+1}(b_{(1,t)}) = g_t(b_{(1,t)}) - \mu_t(g) \quad (1)$$

$$x_t(\text{bal}_t(b_{(1,t-1)}), b_t) = y_t(b_{(1,t)}) \quad (2)$$

$$\hat{p}_t(\text{bal}_t(b_{(1,t-1)}), b_t) = y_t(b_{(1,t)}) \cdot b_t - \int_0^{b_t} y_t(b_{(1,t-1)}, b) db \quad (3)$$

$$s_t(\text{bal}_t(b_{(1,t-1)})) = \mathbb{E}_{v_t \sim \hat{F}_t} \left[\int_0^{v_t} y_t(b_{(1,t-1)}, v) dv \right] + \chi_t(g) - \mu_{t-1}(g) + \mu_t(g) \quad (4)$$

$$p_t(\text{bal}_t(b_{(1,t-1)}), b_t) = \hat{p}_t(\text{bal}_t(b_{(1,t-1)}), b_t) + s_t(\text{bal}_t(b_{(1,t-1)})) \quad (5)$$

where $\mu_t(g) = \inf_{b_{(1,t)}} g_t(b_{(1,t)})$.

Intuitively, the function g maintains the *state* of the core bank account mechanism, which can be viewed as a variant of the bank account balance from (1). The function y defines the stage allocation rule x_t . By the celebrated Myerson's Lemma (Myerson, 1981), \hat{p}_t is the unique payment rule

derived from x_t such that $\langle x_t, \hat{p}_t \rangle$ constitutes a **stage-IC** and **stage-IR** mechanism. A mechanism $\langle x_t, \hat{p}_t \rangle$ is stage-IR if for any balance $\text{bal} \geq 0$, and for any v_t ,

$$x_t(\text{bal}, v_t) \cdot v_t - \hat{p}_t(\text{bal}, v_t) \geq 0 \quad (\text{stage-IR})$$

We refer to the **stage-IC** and **stage-IR** mechanism $\langle x_t, \hat{p}_t \rangle$ as the *local-stage mechanism*. The payment function p_t is computed as the sum of \hat{p}_t and s_t , where s_t only depends on the historical bids $b_{(1,t-1)}$ and does not depend on b_t . The stage mechanism $\langle x, p \rangle$ can be interpreted as follows:

- (1) The seller first charges a *spend* s_t from the buyer before the buyer learns his valuation v_t ;
- (2) The seller runs the local-stage mechanism $\langle x_t, \hat{p}_t \rangle$.

Note that s_t is independent of the buyer's bid b_t at stage t , and therefore, the stage mechanism $\langle x_t, p_t \rangle$ is **stage-IC**. It is straightforward to verify that the expected utility at stage t is always $-\chi_t(g) + \mu_{t-1}(g) - \mu_t(g)$, independent of $\text{bal}_t(b_{(1,t-1)})$. Moreover, **BU** is satisfied with equality such that $\text{bal}U_{t+1}(\text{bal}_t, b_t) = \text{bal}_t + b_t \cdot x_t(\text{bal}_t, b_t) - p_t(\text{bal}_t, b_t)$. Since both g and y are symmetric, we abuse the notation slightly and let $g_t(g_{t-1}(b_{(1,t-1)}), b_t) = g_t(b_{(1,t)})$ and $y_t(g_{t-1}(b_{(1,t-1)}), b_t) = y_t(b_{(1,t)})$.

3.1. Revenue of Core Bank Account Mechanism

Let $\hat{u}_t(\text{bal}_t(b_{(1,t-1)}), b_t) = x_t(\text{bal}_t(b_{(1,t-1)}), b_t) \cdot b_t - \hat{p}_t(\text{bal}_t(b_{(1,t-1)}), b_t)$ be the buyer's utility from the local-stage mechanism $\langle x_t, \hat{p}_t \rangle$. By taking the equality in **BU**, we have $\text{bal}_{t+1}(b_{(1,t)}) = \text{bal}_t(b_{(1,t-1)}) + \hat{u}_t(\text{bal}_t(b_{(1,t-1)}), b_t) - s_t(\text{bal}_t(b_{(1,t-1)}))$. Plugging in (1) and (4), and noticing that $\int_0^{v_t} y_t(b_{(1,t-1)}, v) dv = \hat{u}_t(\text{bal}_t(b_{(1,t-1)}), v_t)$ we have:

$$g_t(b_{(1,t)}) = g_{t-1}(b_{(1,t-1)}) + \hat{u}_t(\text{bal}_t(b_{(1,t-1)}), b_t) - \chi_t(g) - \mathbb{E}_{v_t \sim \hat{F}_t}[\hat{u}_t(\text{bal}_t(b_{(1,t-1)}), v_t)]. \quad (6)$$

Equation (6) is useful because it connects the transition between $g_{t-1}(b_{(1,t-1)})$ and $g_t(b_{(1,t)})$ to the buyer's utility obtained from the local-stage mechanism at stage t . This connection enables a convenient way to compute the revenue of a core bank account mechanism.

Definition 3.2 (Revenue Tracking Program). *For a core bank account mechanism $\langle g, y \rangle$ for $\hat{F}_{(1,T)}$, we consider a revenue tracking program $\psi_t(\xi; B(g, y; \hat{F}_{(1,T)}); F_{(1,T)})$ to compute the revenue of implementing $B(g, y; \hat{F}_{(1,T)})$ when the buyer's true distribution is $F_{(1,T)}$. We define $\psi_{t-1}(\xi; B(g, y; \hat{F}_{(1,T)}); F_{(1,T)})$ to be $-\xi$ when $t = T$ and*

$$\mathbb{E} \left[y_t(\xi, v'_t) \cdot v'_t + \psi_t(g_t(\xi, v'_t); B(g, y; \hat{F}_{(1,T)}); F_{(1,T)}) \right]$$

when $t < T$, where the expectation is taken over $v_t \sim F_t$ and v'_t is the buyer's bid that maximizes her continuation utility when her true value is v_t .

The revenue tracking program provides a tool to compute the revenue, even when the seller's distributional information is not perfectly aligned with the true distribution. Let $\text{Rev}(B, F_{(1,T)})$ be the revenue of implementing B when the buyer's true valuation is $F_{(1,T)}$.

Lemma 3.3. *$\text{Rev}(B(g, y; \hat{F}_{(1,T)}), F_{(1,T)})$ can be computed as $\psi_0(g_0; B(g, y; \hat{F}_{(1,T)}); F_{(1,T)}) + \mu_T(g)$.*

The proof of Lemma 3.3 is based on the fact that the quantity $\psi_0(g_0; B(g, y; \hat{F}_{(1,T)}); F_{(1,T)})$ can be written as

$$\begin{aligned} & \psi_0(g_0; B(g, y; \hat{F}_{(1,T)}); F_{(1,T)}) \\ &= \mathbb{E}_{v_{(1,T)}} \left[\sum_t y_t(v'_{(1,t)}) \cdot v'_t \right] - \mathbb{E}_{v_{(1,T)}} \left[g_T(v'_{(1,T)}) \right] \end{aligned} \quad (7)$$

where the expectation is taken over $F_{(1,T)}$. Recall that $y_t(v'_{(1,t)})$ defines the allocation rule by (2). Therefore, $\mathbb{E}_{v_{(1,T)}} \left[\sum_t y_t(v'_{(1,t)}) \cdot v'_t \right]$ is exactly the expected reported welfare, i.e., the welfare computed from the buyer's reported bids. Moreover, using (6) that connects g and the buyer's utility, we can show that $g_T(v'_{(1,T)})$ is equal to the buyer's reported utility plus $\mu_T(g)$, i.e.,

$$\begin{aligned} g_T(v'_{(1,T)}) &= \mu_T(g) + \sum_{t=1}^T \left[x_t(\text{bal}_t(v'_{(1,t-1)}), b_t) \cdot b_t \right. \\ & \quad \left. - p_t(\text{bal}_t(v'_{(1,t-1)}), v'_t) \right]. \end{aligned} \quad (8)$$

We can then compute the revenue by taking the difference between reported welfare and reported utility.

3.2. Operations on Core Bank Account Mechanism

Given a core bank account $\langle g, y \rangle$ for $\hat{F}_{(1,T)}$, we can apply modifications on $\langle g, y \rangle$ to obtain a new core bank account mechanism $\langle g', y' \rangle$. We introduce three basic operations that we will use to modify a core bank account mechanism. These operations change the dynamics of the core bank account mechanism, and are useful for us to unify the core bank account mechanism (Lemma 3.5) and make it robust (Definition 4.6).

- (1) A *follow-the-history* operation at stage t is defined as

$$\begin{aligned} g'_t(b_{(1,t)}) &= g'_{t-1}(b_{(1,t-1)}) + g_t(b_{(1,t)}) - g_{t-1}(b_{(1,t-1)}) \\ y'_t(b_{(1,t)}) &= y_t(b_{(1,t)}). \end{aligned}$$

By (2) and (3), $x_t(b_{(1,t)}) = x'_t(b_{(1,t)})$ and $\hat{p}_t(b_{(1,t)}) = \hat{p}'_t(b_{(1,t)})$ for any $b_{(1,t-1)}$. Therefore, $\langle x_t, \hat{p}_t \rangle$ under $\langle g, y \rangle$ is the same as $\langle x'_t, \hat{p}'_t \rangle$ under $\langle g', y' \rangle$ for the same history.

- (2) A *follow-the-state* operation at stage t is defined as

$$\begin{aligned} g'_t(b_{(1,t)}) &= g_t(g'_{t-1}(b_{(1,t-1)}), b_t) \\ y'_t(b_{(1,t)}) &= y_t(g'_{t-1}(b_{(1,t-1)}), b_t). \end{aligned}$$

By (2) and (3), for historical bids $b_{(1,t-1)}$ and $b'_{(1,t-1)}$, if $g_{t-1}(b_{(1,t-1)}) = g'_{t-1}(b'_{(1,t-1)})$, then $x_t(b_{(1,t-1)}, b_t) = x'_t(b'_{(1,t-1)}, b_t)$ and $\hat{p}_t(b_{(1,t-1)}, b_t) = \hat{p}'_t(b'_{(1,t-1)}, b_t)$. Therefore, for two histories mapped to the same state, $\langle x_t, \hat{p}_t \rangle$ under $\langle g, y \rangle$ is the same as $\langle x'_t, \hat{p}'_t \rangle$ under $\langle g', y' \rangle$.

(3) A *state-shift* operation at stage t is defined as,

$$g'_t(b_{(1,t)}) = g_t(b_{(1,t)}) + \delta \quad \text{and} \quad y'_t(b_{(1,t)}) = y_t(b_{(1,t)}),$$

for some δ . Basically, a state shift operation simply follows $\langle g, y \rangle$ so that the local-stage mechanism remains the same. However, there is an additional term δ that is added to the state transition function.

Remark 3.4. *It is worth noting that, although the local-stage mechanisms are maintained for all the operations, the stage mechanism might not be the same since the payment p_t includes an additional term, the spend s_t that depends on $\chi_t(g')$, $\mu_{t-1}(g')$, and $\mu_t(g')$ according to (5).*

3.3. Optimal Core Bank Account Mechanism

The next lemma demonstrates that any core bank account mechanism $\langle g, y \rangle$ can be turned into a core bank account mechanism $\langle g', y' \rangle$ with $\chi_t(g') = 0$ for all t and $\mu_T(g') = 0$ with the same revenue.

Lemma 3.5. *For a core bank account mechanism $\langle g, y \rangle$ for $\hat{F}_{(1,T)}$, we construct a core bank account mechanism $\langle g', y' \rangle$ with $\chi_t(g') = 0$ for all t and $\mu_T(g') = 0$ that shares the same revenue as $\langle g, y \rangle$ as follows:*

- $g'_0 = g_0 - \sum_t \chi_t(g) - \mu_T(g)$;
- For $t > 0$, $g'_t(b_{(1,t)}) = g'_{t-1}(b_{(1,t-1)}) + g_t(b_{(1,t)}) - g_{t-1}(b_{(1,t-1)}) + \chi_t(g)$ and $y'_t(b_{(1,t)}) = y_t(b_{(1,t)})$.

In this construction, we apply a follow-the-history operation and a state-shift operation with $\delta = \chi_t(g)$ for stage $t > 0$, and shift the initial state down by $\sum_t \chi_t(g) + \mu_T(g)$.

Utility Interpretation. Under truthful bidding, the buyer's utility is $\hat{u}_t(\text{bal}_t(b_{(1,t-1)}), v_t) - s_t(\text{bal}_t(b_{(1,t-1)}))$ at stage t and the expected utility is $\mathbb{E}_{v_t \sim \hat{F}_t}[\hat{u}_t(\text{bal}_t(b_{(1,t-1)}), v_t)] - s_t(\text{bal}_t(b_{(1,t-1)}))$. The spend $s_t(\text{bal}_t(b_{(1,t-1)}))$ is cancelled, after taking the difference. When $\chi_t(g) = 0$, the transition function (6) of g at stage t can be interpreted as: first subtract the buyer's expected utility at stage t and then add the buyer's realized utility at stage t . Therefore, for a core bank account mechanism with $\chi_t(g) = 0$ and $\mu_T(g) = 0$, we can interpret the state $g_t(b_{(1,t)})$ as the *promised utility* of the buyer, which is the sum of the realized utility for the first t stages and the expected utility in the future.

Dynamic Programming. We are now ready to design a dynamic programming algorithm (Mirrokni et al., 2016b)

to compute the revenue-optimal core bank account mechanism. Let $\phi_{t-1}(\xi; \hat{F}_{(1,T)})$ be the optimal revenue for the sub-problem consisting of stages from t to T , when the buyer's true distribution is $\hat{F}_{(1,T)}$ and the current state is $\xi \geq 0$. Through backward dynamic programming from stage T to stage 1, we can compute $\phi_{t-1}(\xi; \hat{F}_{(1,T)})$ from the following program (OPT-BAM):

$$\begin{aligned} \max \quad & \mathbb{E}_{v_t \sim \hat{F}_t} \left[z_t(\xi, v_t) \cdot v_t + \phi_t(h_t(\xi, v_t); \hat{F}_{(1,T)}) \right] \\ \text{s.t.} \quad & \langle z_t(\xi, \cdot), \hat{q}_t(\xi, \cdot) \rangle \text{ is a stage-IC and IR mechanism} \\ & \forall v_t, \xi + \hat{u}_t(\xi, v_t; v_t) - \mathbb{E}_{v'_t \sim \hat{F}_t} [\hat{u}_t(\xi, v'_t; v'_t)] \geq 0 \end{aligned} \quad (\text{OPT-BAM})$$

where $\hat{u}_t(\xi, v_t; v_t) = v_t \cdot z_t(\xi, v_t) - \hat{q}_t(\xi, v_t)$. In the above program, the free variables are $z_t(\xi, \cdot)$ and $\hat{q}_t(\xi, \cdot)$ while the state transition function $h_t(\xi, v_t) = \xi + \hat{u}_t(\xi, v_t; v_t) - \mathbb{E}_{v'_t \sim \hat{F}_t} [\hat{u}_t(\xi, v'_t; v'_t)]$ is determined by $z_t(\xi, v_t)$ and $\hat{q}_t(\xi, v_t)$, and ensures consistency. Henceforth, the task of the program is to find a local-stage mechanism $\langle z_t(\xi, \cdot), \hat{q}_t(\xi, \cdot) \rangle$ that is *stage-IC* and *stage-IR* and maximizes the objective.

The optimal initial state is $\xi_0^* = \arg \max_{\xi_0 \geq 0} \phi_0(\xi_0; \hat{F}_{(1,T)})$ and let $B(g|_{\xi_0^*}, y|_{\xi_0^*}; \hat{F}_{(1,T)})$ be the optimal mechanism. A FPTAS can be obtained by approximating $\phi_t(\cdot; \hat{F}_{(1,T)})$ by piece-wise linear functions with polynomial-many pieces (Mirrokni et al., 2016b).

3.4. Mismatch between $F_{(1,T)}$ and $\hat{F}_{(1,T)}$

Before we end this section, we provide the first component of our robust dynamic mechanism that quantifies the revenue loss due to the mismatch in distributional information. The next lemma demonstrates that the gradient of the revenue function in terms of the state is at least -1 , which enables us to relate the revenue loss to the amount of state shift.

Lemma 3.6. *For any stage t , state $\xi \geq 0$, and $\delta > 0$, $\phi_t(\xi + \delta; \hat{F}_{(1,T)}) \geq \phi_t(\xi; \hat{F}_{(1,T)}) - \delta$.*

With Lemma 3.6 at hand, we can show that the revenue of the optimal dynamic mechanism for $\hat{F}_{(1,T)}$ is close to the revenue of the optimal dynamic mechanism for $F_{(1,T)}$.

Lemma 3.7. *Let $\phi_0(\hat{\xi}_0^*; \hat{F}_{(1,T)})$ be the revenue of the optimal dynamic mechanism for $\hat{F}_{(1,T)}$ and let $\phi_0(\xi_0^*; F_{(1,T)})$ be the revenue of the optimal dynamic mechanism for $F_{(1,T)}$. Then, $\phi_0(\hat{\xi}_0^*; \hat{F}_{(1,T)}) \geq \phi_0(\xi_0^*; F_{(1,T)}) - O(\Delta \sum_t a_t)$.*

4. Robust Bank Account Mechanism

In this section we provide the central contribution of this paper: a framework to make the optimal bank account mechanism for $\hat{F}_{(1,T)}$ robust against the estimation error and the buyer's misreport when the true distributions are $F_{(1,T)}$. We

first show that the magnitude of the misreport can be related to the estimation error and the sequence of $a_{(1,T)}$. Next, we demonstrate how to design a revenue robust mechanism in which the revenue loss can be related to the magnitude of the misreport and the estimation error. The full proofs of this section are deferred to the full version.

4.1. Misreport from the Buyer

Since the seller does not have perfect distributional information, there is no way for the seller to compute the buyer's expected future utility exactly. As a result, the seller is not able to design a prior-dependent dynamic mechanism achieving exact dynamic incentive-compatibility.

To this end, we modify a core bank account mechanism $\langle g, y \rangle$ for $\hat{F}_{(1,T)}$ to obtain a dynamic mechanism that is $\eta_{(1,T)}$ -DIC for $F_{(1,T)}$. To begin with, note that for an arbitrary core bank account mechanism $\langle g, y \rangle$ for $\hat{F}_{(1,T)}$, the corresponding bank account mechanism $B(g, y; \hat{F}_{(1,T)})$ is **stage-IC** and **BU** for $\hat{F}_{(1,T)}$. Moreover, both of these properties do not depend on the the buyer's true distributions in a single buyer environment. Hence, $B(g, y; \hat{F}_{(1,T)})$ is **stage-IC** and **BU** for $F_{(1,T)}$. Recall that **BU** ensures **ex-post IR**, and thus, $B(g, y; \hat{F}_{(1,T)})$ is also **ex-post IR** for $F_{(1,T)}$. However, $B(g, y; \hat{F}_{(1,T)})$ is no longer **BI** for $F_{(1,T)}$. To overcome this difficulty, we generalize the definition of **BI**.

Definition 4.1 (Approximate BI). *A bank account mechanism for $F_{(1,T)}$ is $\delta_{(1,T)}$ -BI if, for each t and any $\text{bal} \geq 0$, there exists a constant c_t independent of bal such that*

$$\mathbb{E}_{v_t \sim F_t} [v_t \cdot x_t(\text{bal}, v_t) - p_t(\text{bal}, v_t)] \in c + [-\delta_t/2, \delta_t/2]. \quad (\delta_{(1,T)}\text{-BI})$$

Under Assumption 2.1, for the same stage mechanism, the difference between the expected utility under \hat{F}_t and F_t is at most Δa_t . As a result, $B(g, y; \hat{F}_{(1,T)})$ is $\delta_{(1,T)}$ -BI with $\delta_t = 2\Delta a_t$. Combining these observations, we have the following lemma on $B(g, y; \hat{F}_{(1,T)})$ for $F_{(1,T)}$:

Lemma 4.2. *For a core bank account mechanism $\langle g, y \rangle$ for $\hat{F}_{(1,T)}$, $B(g, y; \hat{F}_{(1,T)})$ is **stage-IC**, $\delta_{(1,T)}$ -BI with $\delta_t = 2\Delta a_t$, **BU** and **ex-post IR** for $F_{(1,T)}$.*

For a bank account mechanism satisfying $\delta_{(1,T)}$ -BI for $F_{(1,T)}$, the range of expected utility is δ_t for the t -th stage. Therefore, no matter how the buyer misreports in the first $(t-1)$ stages, her expected utility in the t -th stage can only fluctuate by at most δ_t under truthful reporting. The fact that the stage mechanisms are **stage-IC** for $F_{(1,T)}$ implies that the buyer's expected utility at a stage is maximized when she reports truthfully. Therefore, we are able to upper bound the continuation utility from a misreport.

Lemma 4.3. *For a dynamic mechanism that is **stage-IC** and $\delta_{(1,T)}$ -BI for $F_{(1,T)}$, for any $b_{(1,t-1)}$ and v_t , the difference between the continuation utility of reporting any*

$b_t \in [0, a_t]$ and the continuation utility of reporting v_t truthfully is bounded by $\sum_{t'=t+1}^T \gamma^{t'-t} \cdot \delta_{t'}$.

As a result, once the dynamic mechanism posts a risk for the buyer to misreport at stage t , we are able to bound the magnitude of the buyer's misreport. To do so, we mix the dynamic mechanism with a random posted-price auction at every stage, with probability λ .

Definition 4.4 ($\eta_{(1,T)}$ -DIC Mechanism). *Given a bank account mechanism $B = \langle x, p, \text{balU} \rangle$ satisfying **stage-IC**, $\delta_{(1,T)}$ -BI and **BU** for $F_{(1,T)}$, we construct a bank account mechanism $\bar{B} = \langle \bar{x}, \bar{p}, \bar{\text{balU}} \rangle$ by mixing $B(x, p; \hat{F}_{(1,T)})$ with a random posted price mechanism with probability λ . In particular, the random posted price mechanism at stage t posts a price uniformly from $[0, a_t]$:*

- (1) $\bar{x}(\text{bal}, b_t) = (1 - \lambda) \cdot x(\text{bal}, b_t) + \lambda \cdot \frac{b_t}{a_t}$;
- (2) $\bar{p}(\text{bal}, b_t) = (1 - \lambda) \cdot p(\text{bal}, b_t) + \lambda \cdot \frac{b_t^2}{2a_t}$;
- (3) $\bar{\text{balU}}(\text{bal}, b_t) = (1 - \lambda) \cdot \text{balU}(\text{bal}, b_t)$.

Note that a random posted price auction is **stage-IC** and **stage-IR**. Moreover, in a random posted price mechanism with a price uniformly drawn from $[0, a_t]$ at stage t , it can be shown that a misreport with magnitude m_t will cause the buyer a utility loss $\frac{m_t^2}{2a_t}$. Since the buyer is a utility-maximizer with discounting factor γ , we have the following lemma on the magnitude of misreport for each stage.

Lemma 4.5. *For $\bar{B} = \langle \bar{x}, \bar{p}, \bar{\text{balU}} \rangle$ constructed according to Definition 4.4 from a $\delta_{(1,T)}$ -BI mechanism, the mechanism \bar{B} is **stage-IC**, $\delta_{(1,T)}$ -BI and **BU** for $F_{(1,T)}$. Moreover, the mechanism is $\eta_{(1,T)}$ -DIC with $\eta_t = \sqrt{\frac{2a_t}{\lambda} \cdot \sum_{t'=t+1}^T \gamma^{t'-t} \cdot \delta_{t'}}$ and **ex-post IR** for $F_{(1,T)}$.*

4.2. Revenue Robust Mechanism

Given the construction in Definition 4.4 and Lemma 4.5, any bank account mechanism B for $\hat{F}_{(1,T)}$ can be turned into a $\eta_{(1,T)}$ -DIC mechanism \bar{B} for $F_{(1,T)}$ by mixing B with a random posted price auction for each stage with probability λ . Excluding the random posted price auction, the remaining mechanism is in fact B with probability $(1 - \lambda)$ such that for stage t , the misreport of the buyer is at most η_t . Given $\eta_{(1,T)}$, we construct a revenue robust mechanism in which the revenue is robust against the buyer's misreport.

Definition 4.6 (Revenue Robust Mechanism). *For a core bank account mechanism $\langle g, y \rangle$ for $\hat{F}_{(1,T)}$, assuming the magnitude of the misreport at stage t is at most η_t , we construct a revenue robust core bank account mechanism $\langle \tilde{g}, \tilde{y} \rangle$ with $\beta_t = \Delta a_t$ for all t such that $\tilde{g}_0 = g_0$ and*

$$\begin{aligned} \tilde{g}_t(b_{(1,t)}) &= g_t(\tilde{g}_{t-1}(b_{(1,t-1)}), b_t) + \beta_t + \eta_t \\ \tilde{y}_t(b_{(1,t)}) &= y_t(\tilde{g}_{t-1}(b_{(1,t-1)}), b_t). \end{aligned}$$

Basically, we apply follow-the-state operations for every stage with an additional state shift $\beta_t + \eta_t$. Therefore, in $B(\tilde{g}, \tilde{y}; \hat{F}_{(1,T)})$, the local-stage mechanism is the same as in $B(g, y; \hat{F}_{(1,T)})$ if the state is the same. Moreover, when we update the state, the next state is shifted up by $\beta_t + \eta_t$. The reason why we shift the state up is that, by Lemma 3.6 we do not have a bound on the revenue loss when the combined effect of the estimation error and the buyer's misreport results in a smaller state in the next stage, but we do have one when the combined effect results in a larger state.

Lemma 4.7. $\langle \tilde{g}, \tilde{y} \rangle$ constructed according to Definition 4.6 is a core bank account mechanism for $\hat{F}_{(1,T)}$. When $\beta_t = \Delta a_t$ and $\theta_t(\xi) = \psi_t(\xi; B(g, y; \hat{F}_{(1,T)}); \hat{F}_{(1,T)})$ satisfies $\theta_t(\xi + \delta) \geq \theta_t(\xi) - \delta$ for all t , $\xi \geq 0$, and $\delta > 0$,

$$\begin{aligned} & \text{Rev}(B(\tilde{g}, \tilde{y}; \hat{F}_{(1,T)}), F_{(1,T)}) \\ & \geq \text{Rev}(B(g, y; \hat{F}_{(1,T)}), \hat{F}_{(1,T)}) - O\left(\sum_t (\Delta a_t + \eta_t)\right). \end{aligned}$$

In particular, the condition in Lemma 4.7 is satisfied by the revenue-optimal mechanism by Lemma 3.6.

4.3. Final Mechanism

We are ready to construct our robust bank account mechanism. We first compute the optimal bank account mechanism $B(g|_{\hat{\xi}_0^*}, y|_{\hat{\xi}_0^*}; \hat{F}_{(1,T)})$ for the estimated distributional information $\hat{F}_{(1,T)}$ (Section 3.3). We then compute the revenue robust mechanism $\langle \tilde{g}, \tilde{y} \rangle$ from $\langle g, y \rangle$ (Definition 4.6), and finally, mix in a random posted price mechanism to obtain $\bar{B}(\tilde{g}|_{\hat{\xi}_0^*}, \tilde{y}|_{\hat{\xi}_0^*}; \hat{F}_{(1,T)})$ (Definition 4.4).

Theorem 4.8. Under Assumption 2.1, $\bar{B}(\tilde{g}|_{\hat{\xi}_0^*}, \tilde{y}|_{\hat{\xi}_0^*}; \hat{F}_{(1,T)})$ is stage-IC, BU and ex-post IR for $F_{(1,T)}$. The mechanism is $\delta_{(1,T)}$ -BI with $\delta_t = 2\Delta a_t$ and $\eta_{(1,T)}$ -DIC and $\eta_t = \sqrt{\frac{2a_t\Delta}{\lambda}} \cdot \sum_{t'=t+1}^T \gamma^{t'-t} a_{t'}$ for $F_{(1,T)}$.

In the worst case, the revenue loss from the random posted price auction is at most the expected welfare $\lambda \sum_t \mathbb{E}_{v_t \sim F_t} [v_t] \leq \lambda \sum_t a_t = \lambda T$. Moreover, we can establish an upper bound on the total magnitude of the buyer's misreport $\sum_t \eta_t = O(\sqrt{\Delta/\lambda} \cdot T)$. Finally, combining Lemma 3.7 and Lemma 4.7, we have:

Theorem 4.9.

$$\begin{aligned} & \text{Rev}(\bar{B}(\tilde{g}|_{\hat{\xi}_0^*}, \tilde{y}|_{\hat{\xi}_0^*}; \hat{F}_{(1,T)}), F_{(1,T)}) \\ & \geq \text{Rev}(B^*(F_{(1,T)}), F_{(1,T)}) - O(\lambda T + \sqrt{\Delta/\lambda} \cdot T) \end{aligned}$$

where $B^*(F_{(1,T)})$ is the optimal clairvoyant mechanism for $F_{(1,T)}$, where Δ is the bias bound in Assumption 2.1. The revenue loss is minimized when $\lambda = \Delta^{\frac{1}{3}}$, which results in revenue loss $O(\Delta^{\frac{1}{3}} T)$.

5. No-Regret Policy in Contextual Auctions

In this section, we apply our robust dynamic mechanism in a learning environment, leading to policies that achieve low regret against the optimal clairvoyant mechanism (which has full and accurate distributional information) in the domain of contextual auctions.

5.1. Contextual Auctions

In a *contextual* auction, the item at stage t is represented by an observable feature vector $\zeta_t \in \mathbb{R}^d$ with $\|\zeta_t\|_2 \leq 1$. In line with the literature, we assume that the feature vectors are drawn independently from a fixed distribution \mathcal{D} with positive-definite covariance matrix (Golrezaei et al., 2019). The buyer's preferences are encoded by a fixed vector $\sigma \in \mathbb{R}^d$ and the buyer's valuation at stage t takes the form $v_t = a_t(\langle \sigma, \zeta_t \rangle + n_t)$, where n_t is a noise term with cumulative distribution M_t . The distribution M_t and the feature vector ζ_t are known to the buyer in advance, but the buyer's preference vector σ remains private. In line with the literature (Deng et al., 2019a), we make the following technical assumption that upper bounds the sequence of domain bounds a_t :

Assumption 5.1 (Deng et al. (2019a)). For all stages t , we assume that $\sum_{t' \leq t} a_{t'} \leq c_a \cdot t$ where c_a is a constant.

Assumption 5.1 limits the portion of welfare and revenue that can arise in the first t stages, for any t . Its purpose is to rule out situations where a large fraction of revenue comes from the initial stages, under which a large revenue loss may be inevitable since it is impossible for the seller to obtain a good estimate of σ from just the first few stages.

Our task is to design a policy π that includes both a learning policy for σ and an associated dynamic mechanism policy to extract revenue. At the beginning of stage t , the learning policy estimates \hat{F}_t using information $(a_t, \zeta_t, M_t, \text{ and } b_{(1,t-1)})$ while the dynamic mechanism policy computes the stage mechanism $\langle x_t, p_t \rangle$ at stage t . Let $\text{Rev}(\pi; F_{(1,T)})$ and $\text{Rev}(B; F_{(1,T)})$ be the revenue of implementing policy π and mechanism B for $F_{(1,T)}$, respectively. Moreover, let $B^*(F_{(1,T)})$ denote the revenue-optimal *clairvoyant* dynamic mechanism that knows $F_{(1,T)}$ in advance. The regret of policy π against the dynamic benchmark is defined as $\text{Regret}^\pi(F_{(1,T)}) = \text{Rev}(B^*(F_{(1,T)}); F_{(1,T)}) - \text{Rev}(\pi; F_{(1,T)})$. Our objective is to design a policy with sublinear regret for both the clairvoyant and the semi-clairvoyant environments.

5.2. Clairvoyant Environment

Due to space limitations, the details of our no-regret policies are deferred to the full version. Our robust dynamic mechanism enables a no-regret policy in the clairvoyant environment.

Theorem 5.2. *There exists a policy such that the T -stage regret of the contextual auction in a clairvoyant environment is $\tilde{O}(T^{\frac{6}{7}})$ against the optimal clairvoyant dynamic mechanism that knows the buyer's preference vector in advance.*

5.3. Semi-clairvoyant Environment

We can generalize our results to a semi-clairvoyant environment. In a semi-clairvoyant environment, the seller does not know the time horizon T and he obtains the estimated distributions in $W > 1$ batches, specified by $(\tau_1 = 1, \tau_2, \dots, \tau_W, \tau_{W+1} = T + 1)$. The j -th batch contains the estimated distributions for items arriving between stage τ_j and stage $(\tau_{j+1} - 1)$. In other words, letting \mathcal{B}_j be the set of stages in batch j , the seller obtains the estimated distributions for batch j at the beginning of stage τ_j and not before, so that the mechanism at stage $t \in \mathcal{B}_j$ can only depend on $\hat{F}_{(1, \tau_{j+1} - 1)}$. Henceforth, in the contextual auction, the seller can learn the information about stage $t \in \mathcal{B}_j$ (i.e., a_t , ζ_t , and M_t) at the beginning of stage τ_j .

However, in the worst-case scenario, a semi-clairvoyant environment will degenerate to a non-clairvoyant environment in which each batch only contains one stage, and [Mirrokni et al. \(2018\)](#) demonstrate that the approximation ratio between the optimal non-clairvoyant mechanism and the optimal clairvoyant mechanism is at most $\frac{1}{2}$. To circumvent this impossibility and obtain a no-regret policy against the optimal clairvoyant mechanism, we introduce a measure to capture the revenue gap between the semi-clairvoyant and the clairvoyant mechanism:

Definition 5.3. *Given $\mathcal{B}_{(1,W)}$ and $a_{(1,T)}$, we define a measure $\mathcal{V}(\mathcal{B}_{(1,W)}, a_{(1,T)}) = \sum_{j=1}^W \sqrt{\sum_{t \in \mathcal{B}_j} a_t^2}$.*

The regret of our policy in the semi-clairvoyant environment depends on $\mathcal{V}(\mathcal{B}_{(1,W)}, a_{(1,T)})$, and the regret is sublinear when $\mathcal{V}(\mathcal{B}_{(1,W)}, a_{(1,T)}) = o(T)$. Observe that, $\sum_{j=1}^W \sqrt{\sum_{t \in \mathcal{B}_j} a_t^2} \leq \sum_{j=1}^W \sum_{t \in \mathcal{B}_j} a_t = T$. Therefore, the difference between $\mathcal{V}(\mathcal{B}_{(1,W)}, a_{(1,T)})$ and T is captured by the sum of the difference between $\sqrt{\sum_{t \in \mathcal{B}_j} a_t^2}$ and $\sum_{t \in \mathcal{B}_j} a_t$. Such a difference is small for batch j when there exists a stage $t' \in \mathcal{B}_j$ such that $a_{t'}$ is close to $\sum_{t \in \mathcal{B}_j} a_t$, which implies that there is no much difference between focusing on stage t' only and the stages in batch \mathcal{B}_j , since the revenue contribution from stages other than t' from \mathcal{B}_j is relatively small. Therefore, when the difference between $\sqrt{\sum_{t \in \mathcal{B}_j} a_t^2}$ and $\sum_{t \in \mathcal{B}_j} a_t$ is small, a semi-clairvoyant environment degenerates to a non-clairvoyant environment.

We obtain a no-regret policy by combining our robust dynamic mechanism with a carefully designed learning policy.

Theorem 5.4. *There exists a policy such that the T -stage regret of the contextual auction in a semi-clairvoyant envi-*

ronment is $\tilde{O}\left(T^{\frac{5}{6}} + \mathcal{V}(\mathcal{B}_{(1,W)}, a_{(1,T)})\right)$ against the optimal clairvoyant dynamic mechanism that knows the buyer's preference vector in advance. In particular, the regret is sublinear when $\mathcal{V}(\mathcal{B}_{(1,W)}, a_{(1,T)}) = o(T)$.

6. Conclusion

In this paper, we provided a new framework for designing dynamic mechanisms that are robust to estimation errors in value distributions as well as to strategic behavior. We applied the framework to design policies for contextual auctions that are no-regret against the revenue-optimal dynamic mechanism that has full information about the buyer's distributions, in both the clairvoyant environment and the semi-clairvoyant environment.

A natural direction to consider in the future is to improve the revenue loss bound of our robust dynamic mechanism as well as our no-regret policies. Is it possible to design a robust dynamic mechanism or no-regret policy with smaller revenue loss? Or could we provide a lower bound for the loss? It would also be interesting to apply our framework to contextual auctions with other kinds of valuation structures, and other dynamic auction environments more generally.

References

- Amin, K., Rostamizadeh, A., and Syed, U. Learning prices for repeated auctions with strategic buyers. In *Advances in Neural Information Processing Systems*, pp. 1169–1177, 2013.
- Amin, K., Rostamizadeh, A., and Syed, U. Repeated contextual auctions with strategic buyers. In *Advances in Neural Information Processing Systems*, pp. 622–630, 2014.
- Ashlagi, I., Daskalakis, C., and Haghpanah, N. Sequential mechanisms with ex-post participation guarantees. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 213–214. ACM, 2016.
- Athey, S. and Segal, I. An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485, 2013.
- Baron, D. P. and Besanko, D. Regulation and information in a continuing relationship. *Information Economics and Policy*, 1(3):267–302, 1984.
- Bergemann, D. and Välimäki, J. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010.
- Bergemann, D. and Välimäki, J. Dynamic mechanism design: An introduction. *Journal of Economic Literature*, 57(2):235–74, 2019.
- Cai, Y., Daskalakis, C., and Weinberg, S. M. An algorithmic characterization of multi-dimensional mechanisms.

- In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, pp. 459–478, 2012.
- Cohen, M. C., Lobel, I., and Paes Leme, R. Feature-based dynamic pricing. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 817–817, 2016.
- Courty, P. and Hao, L. Sequential screening. *The Review of Economic Studies*, 67(4):697–717, 2000.
- den Boer, A. V. Dynamic pricing and learning: Historical origins, current research, and new directions. *Surveys in Operations Research and Management Science*, 20(1): 1–18, 2015.
- Deng, Y. and Lahaie, S. Testing dynamic incentive compatibility in display ad auctions. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1616–1624, 2019.
- Deng, Y., Lahaie, S., and Mirrokni, V. A robust non-clairvoyant dynamic mechanism for contextual auctions. In *Advances in Neural Information Processing Systems*, pp. 8654–8664, 2019a.
- Deng, Y., Mirrokni, V., and Zuo, S. Non-clairvoyant dynamic mechanism design with budget constraints and beyond. *Available at SSRN 3383231*, 2019b.
- Deng, Y., Lahaie, S., Mirrokni, V., and Zuo, S. A data-driven metric of incentive compatibility. In *Proceedings of The Web Conference 2020*, pp. 1796–1806, 2020.
- Drutsa, A. Horizon-independent optimal pricing in repeated auctions with truthful and strategic buyers. In *Proceedings of the 26th International Conference on World Wide Web*, pp. 33–42. International World Wide Web Conferences Steering Committee, 2017.
- Drutsa, A. Weakly consistent optimal pricing algorithms in repeated posted-price auctions with strategic buyer. In *International Conference on Machine Learning*, pp. 1318–1327, 2018.
- Edelman, B. and Ostrovsky, M. Strategic bidder behavior in sponsored search auctions. *Decision support systems*, 43(1):192–198, 2007.
- Golrezaei, N., Javanmard, A., and Mirrokni, V. Dynamic incentive-aware learning: Robust pricing in contextual auctions. In *Advances in Neural Information Processing Systems*, pp. 9756–9766, 2019.
- Kakade, S. M., Lobel, I., and Nazerzadeh, H. Optimal dynamic mechanism design and the virtual-pivot mechanism. *Operations Research*, 61(4):837–854, 2013.
- Leme, R. P. and Schneider, J. Contextual search via intrinsic volumes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 268–282. IEEE, 2018.
- Lobel, I., Leme, R. P., and Vladu, A. Multidimensional binary search for contextual decision-making. *Operations Research*, 66(5):1346–1361, 2018.
- Mao, J., Leme, R., and Schneider, J. Contextual pricing for lipschitz buyers. In *Advances in Neural Information Processing Systems*, pp. 5643–5651, 2018.
- Medina, A. M. and Mohri, M. Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of the 31st International Conference on Machine Learning*, pp. 262–270, 2014.
- Mirroknii, V., Leme, R. P., Tang, P., and Zuo, S. Dynamic auctions with bank accounts. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pp. 387–393, 2016a.
- Mirroknii, V., Leme, R. P., Tang, P., and Zuo, S. Optimal dynamic mechanisms with ex-post IR via bank accounts. *arXiv preprint arXiv:1605.08840*, 2016b.
- Mirroknii, V., Paes Leme, R., Tang, P., and Zuo, S. Non-clairvoyant dynamic mechanism design. In *Proceedings of the 19th ACM Conference on Economics and Computation*, pp. 169–169, 2018.
- Myerson, R. B. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- Papadimitriou, C., Pierrakos, G., Psomas, C.-A., and Rubinstein, A. On the complexity of dynamic mechanism design. In *Proceedings of the 27th annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1458–1475, 2016.
- Pavan, A., Segal, I., and Toikka, J. Dynamic mechanism design: A myersonian approach. *Econometrica*, 82(2): 601–653, 2014.
- Thomas, J. and Worrall, T. Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *Journal of Economic Theory*, 51(2):367–390, 1990.