Robust Pricing in Dynamic Mechanism Design

Yuan Deng [1] **Sébastien Lahaie** [2] **Vahab Mirrokni** [2]

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.*

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

Proceedings of the 37^{th} International Conference on Machine Learning, Online, PMLR 119, 2020.

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 $(a_{t'}, \ldots, a_{t''})$ 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: $x_t : V^t \times (\Delta V)^T \to [0, 1]$; (2) the payment function maps historical bids and seller's distributional information $\hat{F}_{(1,T)}$ to a payment: $p_t : V^t \times (\Delta V)^T \to \mathbb{R}$. Given $b_{(1,t)}$ and $\hat{F}_{(1,T)}$, the utility u_t of the buyer with true valuation v_t is $u_t\big(v_t; b_{(1,t)}; \hat{F}_{(1,T)}\big) = v_t \cdot x_t\big(b_{(1,t)}; \hat{F}_{(1,T)}\big) - p_t\big(b_{(1,t)}; \hat{F}_{(1,T)}\big)$. 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

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

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 $t < T$, $U_t\big(b_{(1,t)}; F_{(1,T)}; \hat{F}_{(1,T)}\big)$ is defined recursively as

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

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

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

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

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 $B = \langle x, p, \mathrm{balU}\rangle$ 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 $x_t(\mathrm{bal}, b_t)$, $p_t(\mathrm{bal}, b_t)$ is parameterized by a balance $\mathrm{bal} \in \mathbb{R}_+$, which is incentive-compatible for the stage for every $\mathrm{bal} \geq 0$: for any v_t and b_t,

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

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

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

where c_t is a constant not dependent on bal*;*

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

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

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

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

We will use the notation $x_t(b_{(1,t)})$ and $x_t(\mathrm{bal}, b_t)$ where $\mathrm{bal} = \mathrm{bal}_t(b_{(1,t-1)})$ interchangeably since $x_t(b_{(1,t)})$ and $x_t(\mathrm{bal}, b_t)$ 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 $\mathrm{bal} \geq 0$ 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

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

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

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:

$$\mathrm{bal}_{t+1}\big(b_{(1,t)}\big) = g_t\big(b_{(1,t)}\big) - \mu_t(g) \tag{1}$$

$$x_t\big(\mathrm{bal}_t(b_{(1,t-1)}), b_t\big) = y_t\big(b_{(1,t)}\big) \tag{2}$$

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

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

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

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,

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

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 $\text{bal}_t\big(b_{(1,t-1)}\big)$. Moreover, BU is satisfied with equality such that $\text{balU}_{t+1}(\text{bal}_t, b_t) = \text{bal}_t + b_t \cdot x_t(\text{bal}_t, b_t) - p_t(\text{bal}_t, b_t)$. Since both g and y are symmetric, we abuse the notation slightly and let $g_t\big(g_{t-1}(b_{(1,t-1)}), b_t\big) = g_t(b_{(1,t)})$ and $y_t\big(g_{t-1}(b_{(1,t-1)}), b_t\big) = y_t(b_{(1,t)})$.

3.1. Revenue of Core Bank Account Mechanism

Let $\hat{u}_t\big(\text{bal}_t(b_{(1,t-1)}), b_t\big) = x_t\big(\text{bal}_t(b_{(1,t-1)}), b_t\big) \cdot b_t - \hat{p}_t\big(\text{bal}_t(b_{(1,t-1)}), b_t\big)$ 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 $\text{bal}_{t+1}(b_{(1,t)}) = \text{bal}_t(b_{(1,t-1)}) + \hat{u}_t\big(\text{bal}_t(b_{(1,t-1)}), b_t\big) - s_t\big(\text{bal}_t(b_{(1,t-1)})\big)$. Plugging in (1) and (4), and noticing that $\int_0^{v_t} y_t\big(b_{(1,t-1)}, v\big) dv = \hat{u}_t\big(\text{bal}_t(b_{(1,t-1)}), v_t)\big)$ we have:

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

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*

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

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 $Rev(B, F_{(1,T)})$ be the revenue of implementing B when the buyer's true valuation is $F_{(1,T)}$.

Lemma 3.3. *$\text{Rev}\big(B(g, y; \hat{F}_{(1,T)}), F_{(1,T)}\big)$ 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

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

where the expectation is taken over $F_{(1,T)}$. Recall that $y_t(v'_{(1,t)})$ defines the allocation rule by (2). Therefore, $\mathbb{E}_{v_{(1,T)}}[\sum_t y_t(v'_{(1,t)}) \cdot v_t']$ 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.,

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

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 $\hat{u}_t\big(\mathrm{bal}_t(b_{(1,t-1)}), v_t\big) - s_t\big(\mathrm{bal}_t(b_{(1,t-1)})\big)$ at stage t and the expected utility is $\mathbb{E}_{v_t \sim \hat{F}_t}[\hat{u}_t\big(\mathrm{bal}_t(b_{(1,t-1)}), v_t\big)] - s_t\big(\mathrm{bal}_t(b_{(1,t-1)})\big)$. The spend $s_t\big(\mathrm{bal}_t(b_{(1,t-1)})\big)$ 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):

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

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

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 $\mathrm{bal} \geq 0$, there exists a constant c_t independent of bal such that*

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

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 $B = \langle x, p, \mathrm{balU}\rangle$ satisfying stage-IC, $\delta_{(1,T)}$-BI and BU for $F_{(1,T)}$, we construct a bank account mechanism $\bar{B} = \langle \bar{x}, \bar{p}, \overline{\mathrm{balU}}\rangle$ 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) $\bar{x}(\mathrm{bal}, b_t) = (1-\lambda) \cdot x(\mathrm{bal}, b_t) + \lambda \cdot \frac{b_t}{a_t}$;

(2) $\bar{p}(\mathrm{bal}, b_t) = (1-\lambda) \cdot p(\mathrm{bal}, b_t) + \lambda \cdot \frac{b_t^2}{2a_t}$;

(3) $\overline{\mathrm{balU}}(\mathrm{bal}, b_t) = (1-\lambda) \cdot \mathrm{balU}(\mathrm{bal}, b_t)$.

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 $\bar{B} = \langle \bar{x}, \bar{p}, \overline{\mathrm{balU}}\rangle$ 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*

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

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 $\delta > 0$,*

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

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.

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

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 $\boldsymbol{\zeta}_t \in \mathbb{R}^d$ with $\|\boldsymbol{\zeta}_t\|_2 \leq 1$. 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 $\boldsymbol{\sigma} \in \mathbb{R}^d$ and the buyer's valuation at stage t takes the form $v_t = a_t(\langle \boldsymbol{\sigma}, \boldsymbol{\zeta}_t\rangle + n_t)$, where n_t is a noise term with cumulative distribution M_t. The distribution M_t and the feature vector $\boldsymbol{\zeta}_t$ are known to the buyer in advance, but the buyer's preference vector $\boldsymbol{\sigma}$ 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 $\boldsymbol{\sigma}$ from just the first few stages.

Our task is to design a policy π that includes both a learning policy for $\boldsymbol{\sigma}$ 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, $\boldsymbol{\zeta}_t$, M_t, and $b_{(1,t-1)}$) while the dynamic mechanism policy computes the stage mechanism $\langle x_t, p_t\rangle$ at stage t. Let $\mathrm{Rev}(\pi; F_{(1,T)})$ and $\mathrm{Rev}(B; F_{(1,T)})$ 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 $\mathsf{Regret}^{\pi}(F_{(1,T)}) = \mathrm{Rev}\big(B^*(F_{(1,T)}); F_{(1,T)}\big) - \mathrm{Rev}(\pi; F_{(1,T)})$. 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 $(\tau_1 = 1, \tau_2, \cdots, \tau_W, \tau_{W+1} = T+1)$. 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, $\boldsymbol{\zeta}_t$, 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 *environment 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.

Koufogiannakis, C. and Young, N. E. A nearly linear-time PTAS for explicit fractional packing and covering linear programs. *Algorithmica*, 70(4):648–674, 2014.

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.

Mirrokni, 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.

Mirrokni, 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.

Mirrokni, 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.

Appendix

A. Omitted Details for No-regret Policies in Contextual Auctions

We provide high-level idea of our no-regret policies in this section while the full proofs are deferred to Appendix E.

A.1. The Learning Policy

Our learning policy is adapted from the learning policy proposed in (Golrezaei et al., 2019). In both environments, our learning policy partitions the entire time horizon into K phases, such that the partition is specified by $(t_1 = 1, t_2, \cdots, t_K, t_{K+1} = T+1)$. The k-th phase spans between the t_k-th stage and the $(t_{k+1}-1)$-th stage with length $\ell_k = t_{k+1} - t_k$. Let E_k be the set of stages in phase k.

At the beginning of the k-th phase, we update the estimation of $\boldsymbol{\sigma}$ using the buyer's bids from the $(k-1)$-th phase, denoted by $\hat{\boldsymbol{\sigma}}_k$. For the estimation, we sample w_t uniformly from $[0,1]$ for $t \in \hat{E}_{k-1}$, where $\hat{E}_{k-1} = \{t \in E_{k-1} \mid t_{k-1} - t > c\log \ell_{k-1}\}$ for some constant c. In other words, we will only use the information from the stages that are at least $c \log \ell_{k-1}$ ahead of the end of phase $(k-1)$. $\hat{\boldsymbol{\sigma}}_k$ is set to be

$$\arg\min_{\|\boldsymbol{\sigma}\|\leq 1} \mathcal{L}_{k-1}(\boldsymbol{\sigma}),$$

where

$$\mathcal{L}_{k-1}(\boldsymbol{\sigma}) = -\sum_{t\in\hat{E}_{k-1}} \Big[\mathbf{1}\{b_t \geq a_t \cdot w_t\} \log\left(1 - M_t(w_t - \langle \boldsymbol{\sigma}, \boldsymbol{\zeta}_t\rangle)\right) + \mathbf{1}\{b_t < a_t \cdot w_t\} \log\left(M_t(w_t - \langle \boldsymbol{\sigma}, \boldsymbol{\zeta}_t\rangle)\right)\Big].$$

Since we run the estimation after the $(k-1)$-th phase, the seller has access to all information required to compute the loss function $\mathcal{L}_{k-1}(\boldsymbol{\sigma})$, i.e., the buyer's bids in the $(k-1)$-th phase.

Note that when the buyer reports truthfully, $\mathcal{L}_{k-1}(\boldsymbol{\sigma})$ is exactly the negative of log-likelihood corresponding to $\boldsymbol{\sigma}$. We do not change our estimation throughout the k-th phase and the next update happens at the beginning of the $(k+1)$-phase. As a result, based on the estimate $\hat{\boldsymbol{\sigma}}_k$, we compute the estimated distribution in phase k as $\hat{F}_t(v_t) = M_t\left(\frac{v_t}{a_t} - \langle \hat{\boldsymbol{\sigma}}_k, \boldsymbol{\zeta}_t\rangle\right)$ for all $t \in E_k$.

We say a *lie* is a misreport from the buyer that results in $\mathbf{1}\{b_t \geq a_t \cdot w_t\} \neq \mathbf{1}\{v_t \geq a_t \cdot w_t\}$. Let $L_{k-1} = \left\{t \in \hat{E}_{k-1} \mid \mathbf{1}\{b_t \geq a_t \cdot w_t\} \neq \mathbf{1}\{v_t \geq a_t \cdot w_t\}\right\}$ be the set of stages in which the buyer lies. For a dynamic mechanism that is $\eta_{(1,T)}$-DIC, we have $v_t - \eta_t \leq b_t \leq v_t + \eta_t$. Hence, if $|a_t \cdot w_t - v_t| > \eta_t$, any misreport from the buyer does not result in a lie. Therefore, a lie occurs only if $v_t \in [a_t \cdot w_t - \eta_t, a_t \cdot w_t + \eta_t]$. By a martingale argument on the sequence of lies, we can obtain that the total number of lies caused by the dynamic mechanism is $O\left(\sum_{t\in\hat{E}_{k-1}} \frac{\eta_t}{a_t}\right)$. Moreover, the buyer has an additional motivation to misreport to change the seller's estimation for the future phases. However, for $t \in \hat{E}_{k-1}$, such a gain is relatively small since the buyer discounts the future.

Let $B(\lambda_{(1,K)})$ be a mechanism by mixing a dynamic mechanism B with the random posted-price auction with probability λ_k for stage $t \in E_k$. In the random posted-price auction at stage t, the price is drawn from $[0, a_t]$ uniformly at random.

Lemma A.1. *In a dynamic mechanism $B(\lambda_{(1,K)})$ constructed from a $\eta_{(1,T)}$-DIC dynamic mechanism B, the additional misreport at stage $t \in \hat{E}_k$ is $O\left(\frac{1}{\sqrt{\lambda_k}\cdot\ell_k^2}\right)$. Moreover, $|L_k| = O\left(\frac{1}{\sqrt{\lambda_k}\cdot\ell_k} + \log \ell_k + \sum_{t\in\hat{E}_k}\frac{\eta_t}{a_t}\right)$ with probability $1 - \frac{1}{\ell_k}$.*

Given this upper bound on $|L_{k-1}|$, the following lemma bounds the estimation error of $\hat{\boldsymbol{\sigma}}_k$.

Lemma A.2 (Proposition 7.1 (Golrezaei et al., 2019))**.** *With probability $1 - \frac{1}{\ell_{k-1}}$, the estimation error for phase k is $\Delta_k \equiv \|\hat{\boldsymbol{\sigma}}_k - \boldsymbol{\sigma}\| = O\left(d \cdot \frac{|L_{k-1}|}{\ell_{k-1}} + \sqrt{\frac{\log(\ell_{k-1}\cdot d)}{\ell_{k-1}}}\right)$.*

A.2. No-regret Policy in the Clairvoyant Environment

In the clairvoyant environment, our learning policy simply partitions the entire time horizon into two phases, such that the first phase spans between stage 1 and stage T^α and the second phase spans between stage $(T^\alpha + 1)$ and T. For stages in the first phase, we will run the random posted-price auction with probability 1. Therefore, by Lemma A.1 and letting $\eta_t = 0$ for $t \in E_1$, we have $|L_1| = O(\alpha \log T)$, which implies that $\Delta_2 = \tilde{O}(T^{-\alpha/2})$ the with probability at least $1 - T^{-\alpha}$ by Lemma A.2.

Applying Theorem 4.9, the revenue loss at phase 2 is at most $\tilde{O}(T^{1-\alpha/6})$ when $\Delta_2 = \tilde{O}(T^{-\alpha/2})$. Under Assumption 5.1, the total revenue loss is at most $\tilde{O}(T^{\alpha} + T^{1-\alpha/6})$. By setting $\alpha = \frac{6}{7}$, we have,

Theorem A.3. *The T-stage regret of a 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.*

A.3. No-regret Policy in the Semi-Clairvoyant Environment

A.3.1. Semi-clairvoyant Mechanism

To begin with, we first design a semi-clairvoyant mechanism when the seller's distributional information is perfect, i.e., $\hat{F}_{(1,T)} = F_{(1,T)}$. Let the optimal bank account mechanism for stages in $\mathcal{B}_j$ be $\langle g^{j*}, y^{j*}\rangle$. We show that the revenue loss of running $\langle g^{j*}, y^{j*}\rangle$ separately for $\mathcal{B}_j$ against the optimal clairvoyant mechanism is $\tilde{O}\left(\mathcal{V}(\mathcal{B}_{(1,W)}, a_{(1,T)}\right)$.

Lemma A.4.

$$\begin{aligned}&\sum_{j=1}^{W} \text{Rev}\left(B(g^{j*}, y^{j*}; \hat{F}_{(\tau_j,\tau_{j+1}-1)}), \hat{F}_{(\tau_j,\tau_{j+1}-1)})\right)\\ &\geq \text{Rev}\left(B^*(\hat{F}_{(1,T)}), \hat{F}_{(1,T)}\right) - \tilde{O}\left(\mathcal{V}(\mathcal{B}_{(1,W)}, a_{(1,T)})\right).\end{aligned}$$

To obtain a no-regret policy, we make an assumption on the measure that $\mathcal{V}(\mathcal{B}_{(1,W)}, a_{(1,T)}) = o(T)$; or otherwise the semi-clairvoyant environment will degenerate to a non-clairovyant environment, in which it is impossible to obtain a no-regret policy.

A.3.2. Learning Policy

In the semi-clairvoyant environment, Our learning policy partitions the entire time horizon into $K = \lceil \log T \rceil$ phases where T is the time horizon, such that the partition is specified by $(t_1 = 1, t_2, \cdots, t_K, t_{K+1} = T+1)$, in which $t_k = 2^{k-1}$. The k-th phase spans between the t_k-th stage and the $(t_{k+1}-1)$-th stage, and therefore, the length of phase k is exactly $\ell_k = t_k$. Note that the partition can be implemented even when T is not known in advance.

To align our learning policy with the arrival of batches in the semi-clairvoyant environment, for each $\mathcal{B}_j$ that has stages in phase k and also stages in future phases, we will split $\mathcal{B}_j$ into two batches such that the first batch is $\mathcal{B}_j \cap E_k$ and the second batch is $\mathcal{B}_j \setminus E_k$. We will continue this process until all batches are contained in some phase E_k. Notice that such an operation will create at most $K = O(\log T)$ more batches, and therefore, the revenue loss by running the optimal dynamic mechanisms for new batches is at most $\tilde{O}(\mathcal{V}(\mathcal{B}_{(1,W)}, a_{(1,T)}))$ by Lemma A.4. Therefore, we will focus on the case in which all batches are contained in some phase E_k from now on.

A.3.3. Hybrid Mechanism

If a static mechanism is offered in each phase, since the estimation is fixed in each phase, the only incentive for the buyer to misreport is to change the estimation in the future. However, in a dynamic mechanism, the buyer has an incentive to misreport in order to change the stage mechanism offered in the future stages within the same phase. As previously mentioned, the buyer's misreport is related to estimation error for each phase, but if the average magnitude of the misreport in a phase is close to the estimation error, or even worse, higher than the estimation error, it is impossible to obtain a much better estimation in the next phase.

Although the static mechanism does not have a good revenue performance, it does remove the buyer's incentive to misreport in order to change the stage mechanism of the same phase. Motivated by this insight, we develop a *hybrid mechanism*, which contains both dynamic stages dependent on the historical bids and static stages independent of the history. In other words, we make a trade-off between the revenue loss and the number of lies. We are able to show that an additional sublinear revenue loss is enough to substantially reduce the number of lies.

Let $\langle g, y\rangle$ be any core bank account mechanism and $\langle g^{hybrid}, y^{hybrid}\rangle$ be the hybrid mechanism that modifies $\langle g, y\rangle$. In

$\langle g^{hybrid}, y^{hybrid}\rangle$, we will modify $\langle g, y\rangle$ for each phase k. The high level idea is that we are going to offer static mechanisms for the stages in which a_t is small and offer dynamic mechanisms for stages in which a_t is large. One can interpret this procedure as an implementation of *online bundling* for the stages in which a_t is small. More formally, we partition the stages in E_k into two sets according to the size of a_t: for a fixed $\omega \in (0, 1)$, let $E_k^\omega = \{t \in E_k | a_t \leq \ell_k^\omega\}$. We are going to offer static mechanisms for stages in E_k^ω and dynamic mechanism for stages in $E_k \setminus E_k^\omega$.

Let t^* be the minimum index of a stage in E_k^ω such that $\sum_{t\leq t^*, t\in E_k^\omega} a_t \geq \iota(k)$, where ω and $\iota(k)$ is an adjustable parameter. We will then offer a give-for-free mechanism for stages in $\{t \in E_k \mid t \leq t^*\}$ such that we always allocate the item to the buyer without charging anything; and for for stages in $\{t \in E_k \mid t > t^*\}$,

- If the state $g_t \geq \mathbb{E}[v_t]$, we will offer a give-for-free mechanism with an extra payment $\mathbb{E}[v_t]$, i.e., we allocate the item to the buyer and charge the buyer $\mathbb{E}[v_t]$, no matter what her bid is;
- otherwise, always allocate nothing to the buyer and charge her nothing;

Due to the martingale nature of the state g, the probability that $g_t < \mathbb{E}_t[v_t]$ is small if we properly choose $\iota(k)$ that is sublinear in terms of ℓ_k. To limit the magnitude of the buyer's misreport, we again mix $B(g^{hybrid}, y^{hybrid}; \hat{F}_{(1,T)})$ with a random posted price mechanism with probability λ_k for stage $t \in E_k$ as constructed in Definition 4.4 to obtain $\bar{B}(g^{hybrid}, y^{hybrid}; \hat{F}_{(1,T)})$.

Lemma A.5. *With properly chosen $\iota(k)$,*

$$\begin{aligned}&\mathrm{Rev}\left(B(g^{hybrid}, y^{hybrid}; \hat{F}_{(1,T)}), \hat{F}_{(1,T)}\right)\\ \geq\ &\mathrm{Rev}\left(B(g, y; \hat{F}_{(1,T)}), \hat{F}_{(1,T)}\right) - \tilde{O}\left(T^{\frac{1}{2}(\omega+1)}\right) - \sum_k \lambda_k \ell_k\end{aligned}$$

and in $\bar{B}(g^{hybrid}, y^{hybrid}; \hat{F}_{(1,T)})$, $\sum_{t\in\hat{E}_k} \frac{\eta_t}{a_t} \leq \tilde{O}\left(\ell_k^{1-\omega}\right)$ with probability at least $1 - \frac{1}{\ell_k}$ for any phase k.

A.3.4. FINAL POLICY

We are now ready to put together our learning policy, the robust bank account mechanism, the semi-clairvoyant mechanism, and the hybrid mechanism to obtain a robust dynamic contextual auction policy for the semi-clairvoyant environment: (1) at the start of phase k, estimate $\hat{\boldsymbol{\sigma}}_k = \arg\min_{\|\boldsymbol{\sigma}\|\leq 1} \mathcal{L}_{k-1}(\boldsymbol{\sigma})$; (2) then construct the semi-clairvoyant hybrid mechanism $\langle g^{semi,hybrid,k}, y^{semi,hybrid,k}\rangle$; and then implement the robust semi-clairvoyant hybrid mechanism $\langle \tilde{g}^{semi,hybrid,k}, \tilde{y}^{semi,hybrid,k}\rangle$.

For the first few phases with length less than $T^{5/6}$, we simply bound the regret by $\sum_k \sum_{t\in E_k} a_t = O(T^{5/6})$. For $\ell_k > T^{5/6}$, from Lemma A.4 and Lemma A.5 the revenue loss from the semi-clairvoyant hybrid mechanism is $\tilde{O}\left(\mathcal{V}(\mathcal{B}_{(1,W)}, a_{(1,T)}) + T^{\frac{1}{2}(\omega+1)} + \sum_k \lambda_k \ell_k\right)$. As for the robust semi-clairvoyant hybrid mechanism, by Lemma 4.7, the revenue loss is $O\left(\sum_k \ell_k \lambda_k + \ell_k \sqrt{\frac{\Delta_k}{\lambda_k}}\right)$. By Lemma A.2, $\Delta_k = \tilde{O}(\ell_k^{-\omega})$ and we can set $\lambda_k = \ell_k^{-\frac{1}{3}\omega}$ and $\omega = \frac{1}{2}$, which yields:

Theorem A.6. *The T-stage regret of the robust dynamic contextual auction policy in the semi-clairvoyant environment is $\tilde{O}\left(T^{\frac{5}{6}} + \mathcal{V}(\mathcal{B}_{(1,W)}, a_{(1,T)})\right)$ against the optimal dynamic mechanism. When $\mathcal{V}(\mathcal{B}_{(1,W)}, a_{(1,T)}) = o(T)$, then our policy is with no-regret.*

B. Helper Lemmas

Lemma B.1. *In a single buyer setting, every stage IC and IR mechanism $\langle x, p\rangle$ can be represented by a mixture of posted-price auctions such that the probability density to offer a posted price r is $f(r) = \frac{dx(r)}{dr}$.*

Proof. We show that such a mixture of posted price auctions preserve the allocation rule and payment rule. By the celebrated Myerson's lemma (Myerson, 1981), a mechanism is IC if and only if the allocation rule is monotonically non-decreasing, i.e., $\frac{dx(r)}{dr} \geq 0$ for all valid r. Therefore, the density function of posted prices $f(r)$ is well-defined. Moreover, for a buyer with bid b, his allocation probability is $\int_0^b f(r)dr = \int_0^b \frac{dx(r)}{dr}dr = x(r)$, which implies that the allocation probability is preserved. Moreover, Myerson's lemma (Myerson, 1981) demonstrated that the payment rule is uniquely determined by the allocation rule: $p(b) = \int_0^b r \cdot \frac{dx(r)}{dr}dr$, which is exactly the payment collected from our mixture of posted price auctions for valuation v. □

Lemma B.2. *For $v \in [\hat{v} - \Delta, \hat{v} + \Delta]$ and any stage IC and IR mechanism $\langle x, p \rangle$, we have $u(v) - \Delta \le u(\hat{v}) \le u(v) + \Delta$ where $u(v) = x(v) \cdot v - p(v)$.*

Proof. Since $\langle x, p \rangle$ is a stage IC and IR mechanism, by Lemma B.1, we can equivalently offer a mixture of posted price auctions such that the probability density to post a price r is $f(r) = \frac{dx(r)}{dr}$. Therefore, we can express the utility of the buyer for valuation v_t as $\int \frac{dx(r)}{dr}(v - r)^+dr$. We first show the first inequality:

$$
\begin{aligned}
x(v) \cdot v - p(v) &= \int \frac{dx(r)}{dr}(v - r)^+dr \\
&\le \int \frac{dx(r)}{dr}(\hat{v} + \Delta - r)^+dr \\
&\le \int \frac{dx(r)}{dr}(\hat{v} - r)^+dr + \Delta \\
&= x(\hat{v}) \cdot \hat{v} - p(\hat{v}) + \Delta
\end{aligned}
$$

where the first equality follows that $\langle x, p \rangle$ is stage-IC and Lemma B.1. By a similar argument, we can prove the second inequality. □

The following is a corollary of Lemma B.2, which demonstrates that the difference of expected utility due to the mismatch of distributional information can be related to the estimation error.

Corollary B.3. *For $F_{(1,T)}$ and $\hat{F}_{(1,T)}$ satisfying Assumption 2.1, and for any stage IC and IR mechanism $\langle x, p \rangle$, we have*

$$\mathbb{E}_{v_t \sim F_t}[u(v_t)] - \Delta a_t \le \mathbb{E}_{v_t \sim \hat{F}_t}[u(v_t)] \le \mathbb{E}_{v_t \sim F_t}[u(v_t)] + \Delta a_t.$$

The following lemma demonstrates that the difference of expected welfare due to the mismatch of distributional information can be related to the estimation error.

Lemma B.4. *For $F_{(1,T)}$ and $\hat{F}_{(1,T)}$ satisfying Assumption 2.1, and for any stage IC and IR mechanism $\langle x, p \rangle$, we have*

$$\mathbb{E}_{v_t \sim F_t}[x_t(v_t) \cdot v_t] - (c_f + 1)\Delta a_t \le \mathbb{E}_{v_t \sim \hat{F}_t}[x_t(v_t) \cdot v_t] \le \mathbb{E}_{v_t \sim F_t}[x_t(v_t) \cdot v_t] + (c_f + 1)\Delta a_t$$

Proof. Since $\langle x, p \rangle$ is an stage IC and IR mechanism, by Lemma B.1, we can equivalently offer a mixture of posted price auctions such that the probability density to post a price r is $f(r) = \frac{dx(r)}{dr}$. Denote the distribution that ϵ_t follows as G_t.

We first show the first inequality:

$$
\begin{aligned}
\mathbb{E}_{v_t \sim F_t}[x_t(v_t) \cdot v_t] &= \mathbb{E}_{v_t \sim \hat{F}_t}[x_t(v_t + \epsilon_t a_t) \cdot (v_t + \epsilon_t a_t)] \\
&= \int \frac{dx_t(r)}{dr} \cdot \mathbb{E}_{v_t \sim \hat{F}_t}[\mathbf{1}\{v_t + \epsilon_t a_t \ge r\} \cdot (v_t + \epsilon_t a_t)]dr \\
&\le \int \frac{dx_t(r)}{dr} \cdot \mathbb{E}_{v_t \sim \hat{F}_t}[\mathbf{1}\{v_t + \Delta a_t \ge r\} \cdot (v_t + \Delta a_t)]dr \\
&\le \int \frac{dx_t(r)}{dr} \cdot \mathbb{E}_{v_t \sim \hat{F}_t}[\mathbf{1}\{v_t + \Delta a_t \ge r\} \cdot (v_t + \Delta a_t)]dr \\
&\le \int \frac{dx_t(r)}{dr} \cdot \mathbb{E}_{v_t \sim \hat{F}_t}[\mathbf{1}\{v_t + \Delta a_t \ge r\} \cdot v_t]dr + \Delta a_t \\
&= \int \frac{dx_t(r)}{dr} \cdot \mathbb{E}_{v_t \sim \hat{F}_t}[\mathbf{1}\{v_t \ge r\} \cdot v_t + \mathbf{1}\{r - \Delta a_t \le v_t \le r\} \cdot v_t]dr + \Delta a_t \\
&= \mathbb{E}_{v_t \sim \hat{F}_t}[x_t(v_t) \cdot v_t] + \int \frac{dx_t(r)}{dr} \cdot \left[\int_{r - \Delta a_t}^{r} v_t \cdot f_t(v_t)dv_t\right] dr + \Delta a_t \\
&\le \mathbb{E}_{v_t \sim \hat{F}_t}[x_t(v_t) \cdot v_t] + (c_f + 1)\Delta a_t
\end{aligned} \tag{9}
$$

where the last inequality follows that $v_t \leq a_t$ and $f_t(v_t) \leq \frac{c_f}{a_t}$. By a similar argument, we can prove the second inequality. □

C. Omitted Proofs of Section 3

C.1. Proof of Lemma 3.3

Proof of Lemma 3.3. Consider the expression in (7) and notice that from Definition 3.1, $y_t(v'_{(1,t)})$ corresponds to the stage allocation rule, and thus, $y_t(v'_{(1,t)}) \cdot v'_t$ adds the reported stage welfare, i.e., the stage welfare computed from the reported bid, to the result. Therefore, all that remains to show is that $g_T(v'_{(1,T)})$ equals to the buyer's reported utility plus $\mu_T(g)$. Notice that that

$$g_T(b_{(1,T)}) = g_0 + \sum_{t=1}^{T} \big(g_t(b_{(1,t)}) - g_{t-1}(b_{(1,t-1)})\big) \tag{10}$$

where by (6), we have

$$g_t(b_{(1,t)}) - g_{t-1}(b_{(1,t-1)}) = \hat{u}_t(\mathrm{bal}_t(b_{(1,t-1)}), b_t) - \mathbb{E}_{v_t \sim \hat{F}_t}[\hat{u}_t(\mathrm{bal}_t(b_{(1,t-1)}), v_t)] - \chi_t(g) \tag{11}$$

Moreover, by (4), the additional payment s_t under g is

$$s_t\big(\mathrm{bal}_t(b_{(1,t-1)})\big) = \mathbb{E}_{v_t \sim \hat{F}_t}[\hat{u}_t(\mathrm{bal}_t(b_{(1,t-1)}), v_t)] + \chi_t(g) - \mu_{t-1}(g) + \mu_t(g) \tag{12}$$

Let $u_t(\mathrm{bal}, \cdot) = \hat{u}_t(\mathrm{bal}, \cdot) - s_t(\mathrm{bal})$. Plugging (12) into (11), we have

$$g_t(b_{(1,t)}) - g_{t-1}(b_{(1,t-1)}) = u_t(\mathrm{bal}_t(b_{(1,t-1)}), b_t) - \mu_{t-1}(g) + \mu_t(g) \tag{13}$$

Plugging (13) into (10), we have

$$\begin{aligned} g_T(b_{(1,T)}) &= g_0 + \sum_{t=1}^{T} \big(u_t\big(\mathrm{bal}_t(b_{(1,t-1)}), b_t\big) - \mu_{t-1}(g) + \mu_t(g)\big) \\ &= \sum_{t=1}^{T} u_t\big(\mathrm{bal}_t(b_{(1,t-1)}), b_t\big) + \mu_T(g) \end{aligned}$$

where the second equality follows the telescoping sum of μ and the fact that $\mu_0 = g_0$. □

C.2. Proof of Lemma 3.5

Proof. First, it is straightforward to verify that y'_t is a sub-gradient of g'_t with respect to b_t and g'_t is symmetric, convex in b_t, and weakly increasing in b_t. For the consistency of g', we will show that $\chi_t(g') = 0$ for all $t \in A$:

$$\begin{aligned} \mathbb{E}_{v_t \sim \hat{F}_t}[g'_t(b_{(1,t-1)}, v_t)] &= \mathbb{E}_{v_t \sim \hat{F}_t}[g'_{t-1}(b_{(1,t-1)}) + g_t(b_{(1,t-1)}, v_t) - g_{t-1}(b_{(1,t-1)}) + \chi_t(g)] \\ &= g'_{t-1}(b_{(1,t-1)}) + \mathbb{E}_{v_t \sim \hat{F}_t}[g_t(b_{(1,t-1)}, v_t)] - g_{t-1}(b_{(1,t-1)}) + \chi_t(g) \\ &= g'_{t-1}(b_{(1,t-1)}) - \chi_t(g) + \chi_t(g) \\ &= g'_{t-1}(b_{(1,t-1)}). \end{aligned}$$

where the third equality follows the fact that g is consistent. Finally, we show that for all t and $b_{(1,t)}$, $g'_t(b_{(1,t)}) \geq 0$.

Claim C.1. *For a core bank account mechanism $\langle g, y \rangle$ with $\chi_t(g) = 0$ for all t, we have $\mu_0 \geq \mu_1 \geq \cdots \geq \mu_T(g)$ where $\mu_t(g) = \inf_{b_{(1,t)}} g_t(b_{(1,t)})$.*

Proof. Let $(b_{(1,t)})_{\min} = \arg\min_{b_{(1,t)}} g_t(b_{(1,t)})$. Notice that we have

$$\begin{aligned} \mu_t &\leq g_t\big((b_{(1,t-1)})_{\min}, 0\big) \\ &= g_t\big((b_{(1,t-1)})_{\min}\big) + \hat{u}_t\big(\mathrm{bal}_t(b_{(1,t-1)})_{\min}, 0\big) - \mathbb{E}_{v_t \sim \hat{F}_t}[\hat{u}_t\big(\mathrm{bal}_t(b_{(1,t-1)})_{\min}, v_t\big)] - \chi_t(g) \\ &= g_t\big((b_{(1,t-1)})_{\min}\big) - \mathbb{E}_{v_t \sim \hat{F}_t}[\hat{u}_t\big(\mathrm{bal}_t(b_{(1,t-1)})_{\min}, v_t\big)] \\ &= \mu_{t-1} - \mathbb{E}_{v_t \sim \hat{F}_t}[\hat{u}_t\big(\mathrm{bal}_t(b_{(1,t-1)})_{\min}, v_t\big)] \\ &\leq \mu_{t-1} \end{aligned}$$

where $\hat{u}_t(b_{(1,t)}) = x_t(b_{(1,t)}) \cdot b_t - \hat{p}_t(b_{(1,t)})$ and x_t, $\hat{p}_t$, and bal_t are derived from $B(g, y; \hat{F}_{(1,T)})$ according to Definition 3.1. Moreover, the first equality follows the definition of μ_t. □

By Claim C.1, it suffices to show that $\mu_T(g') = \inf_{b_{(1,T)}} g'_T(b_{(1,T)}) \geq 0$. However, from our construction, we have for all $b_{(1,t)}$,

$$g'_t(b_{(1,t)}) = g'_0 + \sum_{t'=1}^{t} \big(g_{t'}(b_{(1,t')}) - g_{t'-1}(b_{(1,t'-1)}) + \chi_{t'}(g)\big) = g_t(b_{(1,t)}) - \sum_{t'=t+1}^{T} \chi_{t'}(g) - \mu_T(g). \tag{14}$$

Therefore, for $b_{(1,T)}$, we have $g'_T(b_{(1,T)}) = g_T(b_{(1,T)}) - \mu_T(g)$, which implies that $\mu_T(g') = \mu_T(g) - \mu_T(g) = 0$. Thus, we have finished showing that $\langle g', y' \rangle$ is a valid core bank account mechanism for $\hat{F}_{(1,T)}$.

As for the revenue performance, notice that, (14) implies that

$$\begin{aligned} g'_t\big(g'_{t-1}(b_{(1,t-1)}), b_t\big) &= g_t\big(g_{t-1}(b_{(1,t-1)}), b_t\big) - \sum_{t'=t+1}^{T} \chi_{t'}(g) - \mu_T(g) \\ &= g_t\left(g'_{t-1}(b_{(1,t-1)}) + \sum_{t'=t}^{T} \chi_{t'}(g) + \mu_T(g), b_t\right) - \sum_{t'=t+1}^{T} \chi_{t'}(g) - \mu_T(g). \end{aligned}$$

Thus, for all $\xi \geq 0$, we have

$$g'_t(\xi, b_t) = g_t\left(\xi + \sum_{t'=t}^{T} \chi_{t'}(g) + \mu_T(g), b_t\right) - \sum_{t'=t+1}^{T} \chi_{t'}(g) - \mu_T(g) \tag{15}$$

For convenience, let $\theta_t(\xi) = \psi_t(\xi; B(g', y'; \hat{F}_{(1,T)}); \hat{F}_{(1,T)})$ and $\rho_t(\xi) = \psi_t(\xi; B(g, y; \hat{F}_{(1,T)}); \hat{F}_{(1,T)})$. By Lemma 3.3 and the fact that $\mu_T(g') = 0$, all that remains to show is that $\theta_0(g_0 - \sum_t \chi_t(g) - \mu_T(g)) = \rho_0(g_0) + \mu_T(g)$. We prove by a backward induction from $t = T$ to $t = 0$ to show that for all t and $\xi \geq 0$.

$$\theta_t\left(\xi - \sum_{t'=t+1}^{T} \chi_{t'}(g) - \mu_T(g)\right) = \rho_t(\xi) + \mu_T(g).$$

It is obvious that the base case $t = T$ is true since $\theta_T(\xi - \mu_T(g)) = -\xi + \mu_T(g) = \rho_T(\xi) + \mu_T(g)$ by the boundary condition. Assume it is true for all $t' \geq t$ and for $t - 1$, notice that by the construction of y'_t and (14), we have

$$y'_t\big(g'_{t-1}(b_{(1,t-1)}), b_t\big) = y_t\big(g_{t-1}(b_{(1,t-1)}), b_t\big) = y_t\left(g'_{t-1}(b_{(1,t-1)}) + \sum_{t'=t}^{T} \chi_{t'}(g) + \mu_T(g), b_t\right) \tag{16}$$

Therefore, we have

$$\begin{aligned} &\theta_{t-1}\left(\xi - \sum_{t'=t}^{T} \chi_{t'}(g) - \mu_T(g)\right) \\ &= \mathbb{E}_{v_t \sim \hat{F}_t}\left[y'_t\left(\xi - \sum_{t'=t}^{T} \chi_{t'}(g) - \mu_T(g), v_t\right) \cdot v_t + \theta_t\left(g'_t(\xi - \sum_{t'=t}^{T} \chi_{t'}(g) - \mu_T(g), v_t)\right)\right] \\ &= \mathbb{E}_{v_t \sim \hat{F}_t}\left[y_t(\xi, v_t) \cdot v_t + \rho_t\left(g'_t(\xi - \sum_{t'=t}^{T} \chi_{t'}(g) - \mu_T(g), v_t) + \sum_{t'=t+1}^{T} \chi_{t'}(g) + \mu_T(g)\right) + \mu_T(g)\right] \\ &= \mathbb{E}_{v_t \sim \hat{F}_t}\left[y_t(\xi, v_t) \cdot v_t + \rho_t\big(g_t(\xi, v_t)\big)\right] + \mu_T(g) \\ &= \rho_{t-1}(\xi) + \mu_T(g) \end{aligned}$$

where the second equality follows (16) and the induction hypothesis, while the third equality follows (15). □

C.3. Proof of Lemma 3.6

Proof. We prove by a backward induction from $t = T$ to $t = 0$. It is true for the base case when $t = T$ since $\phi_T(\xi + \delta; \hat{F}_{(1,T)}) = -\xi - \delta = \phi_{T+1}(\xi; \hat{F}_{(1,T)}) - \delta$ by the boundary condition.

Suppose it is true for all $t' \geq t$. For $t-1$ and $\xi \geq 0$, let the local-stage mechanism $\langle z_t(\xi, \cdot), \hat{q}_t(\xi, \cdot)\rangle$ be the optimal solution of the program OPT-BAM for $\phi_{t-1}(\xi; \hat{F}_{(1,T)})$. Notice that since $\xi + \delta > \xi$, $\langle z_t(\xi, \cdot), \hat{q}_t(\xi, \cdot)\rangle$ is also a feasible solution of the program OPT-BAM for $\phi_{t-1}(\xi + \delta; \hat{F}_{(1,T)})$. By using $\langle z_t(\xi, \cdot), \hat{q}_t(\xi, \cdot)\rangle$ as a solution, the corresponding transition function h'_t satisfies

$$h'_t(\xi + \delta, v_t) = \xi + \delta + \hat{u}_t(\xi, v_t; v_t) - \mathbb{E}_{v'_t \sim \hat{F}_t}[\hat{u}_t(\xi, v'_t; v'_t)] = h_t(\xi, v_t) + \delta.$$

Therefore, we have

$$\begin{aligned}
\phi_{t-1}(\xi + \delta; \hat{F}_{(1,T)}) &\geq \mathbb{E}_{v_t \sim \hat{F}_t}[z_t(\xi, v_t) \cdot v_t + \phi_t(h'_t(\xi + \delta, v_t); \hat{F}_{(1,T)})] \\
&= \mathbb{E}_{v_t \sim \hat{F}_t}[z_t(\xi, v_t) \cdot v_t + \phi_t(h_t(\xi, v_t) + \delta; \hat{F}_{(1,T)})] \\
&\geq \mathbb{E}_{v_t \sim \hat{F}_t}[z_t(\xi, v_t) \cdot v_t + \phi_t(h_t(\xi, v_t); \hat{F}_{(1,T)}) - \delta] \\
&= \phi_{t-1}(\xi; \hat{F}_{(1,T)}) - \delta
\end{aligned}$$

where the last inequality follows the induction hypothesis. □

C.4. Proof of Lemma 3.7

Proof. The first inequality directly follows the choice of $\hat{\xi}^*_0$. For the second inequality, we prove by a backward induction from $t = T$ to $t = 0$ to show that for all t and $\xi \geq 0$

$$\phi_t\left(\xi + 2\Delta \sum_{t'=t+1}^{T} a_{t'}; \hat{F}_{(1,T)}\right) \geq \phi_t(\xi; F_{(1,T)}) - (c_f + 5)\Delta \sum_{t'=t+1}^{T} a_{t'}. \tag{17}$$

It is obvious that the base case is true for $t = T$. Suppose it is true for all $t' \geq t$. For $t-1$ and $\xi \geq 0$, let the local-stage mechanism $\langle z_t(\xi, \cdot; F_{(1,T)}), \hat{q}_t(\xi, \cdot; F_{(1,T)})\rangle$ be the optimal solution of the program OPT-BAM for $\phi_{t-1}(\xi; F_{(1,T)})$.

Since $\langle z_t(\xi, \cdot; F_{(1,T)}), \hat{q}_t(\xi, \cdot; F_{(1,T)})\rangle$ is a stage IC and IR mechanism, such a mechanism can be a feasible solution for $\phi_{t-1}(\xi + 2\Delta \sum_{t'=t}^{T} a_{t'}; \hat{F}_{(1,T)})$ if it satisfies the consistency constraints of program OPT-BAM. By Corollary B.3, we have

$$\mathbb{E}_{v_t \sim F_t}[\hat{u}_t(\xi, v_t; v_t)] - \Delta a_t \leq \mathbb{E}_{v_t \sim \hat{F}_t}[\hat{u}_t(\xi, v_t; v_t)] \leq \mathbb{E}_{v_t \sim F_t}[\hat{u}_t(\xi, v_t; v_t)] + \Delta a_t \tag{18}$$

Let $h_t(\xi, \cdot; F_{(1,T)})$ be the transition function for $\phi_{t-1}(\xi; F_{(1,T)})$ and let $h'_t(\xi + \Delta \sum_{t'=t}^{T} a_{t'}, \cdot; \hat{F}_{(1,T)})$ be the transition function when using $\langle z_t(\xi, \cdot; F_{(1,T)}), \hat{p}_t(\xi, \cdot; F_{(1,T)})\rangle$ for $\phi_{t-1}(\xi + \Delta \sum_{t'=t}^{T} a_{t'}; \hat{F}_{(1,T)})$. Therefore, we have

$$h'_t(\xi + 2\Delta \sum_{t'=t}^{T} a_{t'}, v_t; \hat{F}_{(1,T)}) = \xi + 2\Delta \sum_{t'=t}^{T} a_{t'} + \hat{u}_t(\xi, v_t; v_t; F_{(1,T)}) - \mathbb{E}_{v'_t \sim \hat{F}_t}[\hat{u}_t(\xi, v'_t; v'_t; F_{(1,T)})] \tag{19}$$

where $\hat{u}_t(\xi, v_t; v_t; F_{(1,T)}) = v_t \cdot z_t(\xi, v_t; F_{(1,T)}) - \hat{q}_t(\xi, v_t; F_{(1,T)})$. Moreover,

$$h_t(\xi, v_t; F_{(1,T)}) = \xi + \hat{u}_t(\xi, v_t; v_t; F_{(1,T)}) - \mathbb{E}_{v'_t \sim F_t}[\hat{u}_t(\xi, v'_t; v'_t; F_{(1,T)})] \tag{20}$$

Comparing (19) and (20) and combining with (18), we have

$$\begin{aligned}
h_t(\xi, v_t; F_{(1,T)}) + \Delta a_t + 2\Delta \sum_{t'=t+1}^{T} a_{t'} \ &\leq \ h'_t(\xi + 2\Delta \sum_{t'=t}^{T} a_{t'}, v_t; \hat{F}_{(1,T)}) \\
&\leq \ h_t(\xi, v_t; F_{(1,T)}) + 3\Delta a_t + 2\Delta \sum_{t'=t+1}^{T} a_{t'}
\end{aligned} \tag{21}$$

Therefore, $\langle z_t(\xi,\cdot;F_{(1,T)}),\hat{q}_t(\xi,\cdot;F_{(1,T)})\rangle$ is a feasible solution for $\phi_{t-1}(\xi+2\Delta\sum_{t'=t}^{T}a_t;\hat{F}_{(1,T)})$.

By Lemma 3.6 and (21), we have

$$\begin{aligned}
&\phi_t\big(h'_t(\xi+2\Delta\sum_{t'=t}^{T}a_{t'},v_t;\hat{F}_{(1,T)});\hat{F}_{(1,T)}\big)\\
&\geq \phi_t\big(h_t(\xi,v_t;F_{(1,T)})+\Delta a_t+2\Delta\sum_{t'=t+1}^{T}a_{t'};\hat{F}_{(1,T)}\big)-2\Delta a_t\\
&\geq \phi_t\big(h_t(\xi,v_t;F_{(1,T)})+\Delta a_t;F_{(1,T)}\big)-2\Delta a_t-(c_f+5)\cdot\Delta\cdot\sum_{t'=t+1}^{T}a_{t'}
\end{aligned}\tag{22}$$

where the last inequality is by the induction hypothesis. Since $\phi_{t-1}(\cdot|\hat{F})$ optimizes over all possible choice of local-stage mechanisms and $\langle z_t(\xi,\cdot;F_{(1,T)}),\hat{q}_t(\xi,\cdot;F_{(1,T)})\rangle$ is a feasible solution for $\phi_{t-1}(\xi+\Delta\sum_{t'=t}^{T}a_t;\hat{F}_{(1,T)})$, we have

$$\begin{aligned}
&\phi_{t-1}(\xi+2\Delta\sum_{t'=t}^{T}a_{t'};\hat{F}_{(1,T)})\\
&\geq \mathbb{E}_{v_t\sim\hat{F}_t}\left[z_t(\xi,v_t;F_{(1,T)})\cdot v_t+\phi_t\big(h'_t(\xi+2\Delta\sum_{t'=t}^{T}a_{t'},v_t;\hat{F}_{(1,T)});\hat{F}_{(1,T)}\big)\right]\\
&\geq \mathbb{E}_{v_t\sim\hat{F}_t}\left[z_t(\xi,v_t|F)\cdot v_t+\phi_t\big(h_t(\xi,v_t;F_{(1,T)})+\Delta a_t;F_{(1,T)}\big)-2\Delta a_t-(c_f+5)\cdot\Delta\cdot\sum_{t'=t+1}^{T}a_{t'}\right]\\
&=\phi_t(\xi;F_{(1,T)})-\Big(\mathbb{E}_{v_t\sim F_t}[z_t(\xi,v_t;F_{(1,T)})\cdot v_t]-\mathbb{E}_{v_t\sim\hat{F}_t}[z_t(\xi,v_t;F_{(1,T)})\cdot v_t]\Big)\\
&\quad-\Big(\mathbb{E}_{v_t\sim F_t}[\phi_t\big(h_t(\xi,v_t;F_{(1,T)});F_{(1,T)}\big)]-\mathbb{E}_{v_t\sim\hat{F}_t}[\phi_t\big(h_t(\xi,v_t;F_{(1,T)})+\Delta a_t;F_{(1,T)}\big)]\Big)\\
&\quad-(c_f+5)\cdot\Delta\cdot\sum_{t'=t+1}^{T}a_{t'}-2\Delta a_t
\end{aligned}\tag{23}$$

where the second inequality follows (22). All that remains to show is the two terms in the big brackets of (23) can be bounded. Using Lemma B.4 for the first term, we have

$$\mathbb{E}_{v_t\sim F_t}[z_t(\xi,v_t;F_{(1,T)})\cdot v_t]\leq\mathbb{E}_{v_t\sim\hat{F}_t}[z_t(\xi,v_t;F_{(1,T)})\cdot v_t]+(c_f+1)\Delta a_t\tag{24}$$

Moreover, for the second term, we have

$$\begin{aligned}
&\mathbb{E}_{v_t\sim F_t}[\phi_t\big(h_t(\xi,v_t;F_{(1,T)});F_{(1,T)}\big)]\\
&=\mathbb{E}_{v_t\sim F_t}\left[\phi_t\big(\xi+\hat{u}_t(\xi,v_t;v_t;F_{(1,T)})-\mathbb{E}_{v'_t\sim F_t}[\hat{u}_t(\xi,v'_t;v'_t;F_{(1,T)})];F_{(1,T)}\big)\right]\\
&=\mathbb{E}_{v_t\sim\hat{F}_t}\left[\phi_t\big(\xi+\hat{u}_t(\xi,v_t+\epsilon_t a_t;v_t+\epsilon_t a_t;F_{(1,T)})-\mathbb{E}_{v'_t\sim F_t}[\hat{u}_t(\xi,v'_t;v'_t;F_{(1,T)})];F_{(1,T)}\big)\right]
\end{aligned}\tag{25}$$

By Lemma B.2, we have

$$\begin{aligned}
&\xi+\hat{u}_t(\xi,v_t;v_t;F_{(1,T)})-\mathbb{E}_{v'_t\sim F_t}[\hat{u}_t(\xi,v'_t;v'_t;F_{(1,T)})]-\Delta a_t\\
&\leq\xi+\hat{u}_t(\xi,v_t+\epsilon_t a_t;v_t+\epsilon_t a_t;F_{(1,T)})-\mathbb{E}_{v'_t\sim F_t}[\hat{u}_t(\xi,v'_t;v'_t;F_{(1,T)})]\\
&\leq\xi+\hat{u}_t(\xi,v_t;v_t;F_{(1,T)})-\mathbb{E}_{v'_t\sim F_t}[\hat{u}_t(\xi,v'_t;v'_t;F_{(1,T)})]+\Delta a_t
\end{aligned}$$

By Lemma 3.6, we have

$$\begin{aligned}
&\phi_t\left(\xi+\hat{u}_t(\xi,v_t+\epsilon_t a_t;v_t+\epsilon_t a_t;F_{(1,T)})-\mathbb{E}_{v'_t\sim F_t}[\hat{u}_t(\xi,v'_t;v'_t;F_{(1,T)})];F_{(1,T)}\right)\\
&\leq\phi_t\left(\xi+\hat{u}_t(\xi,v_t;v_t;F_{(1,T)})-\mathbb{E}_{v'_t\sim F_t}[\hat{u}_t(\xi,v'_t;v'_t;F_{(1,T)})]+\Delta a_t;F_{(1,T)}\right)+2\Delta a_t
\end{aligned}\tag{26}$$

Combining (25) and (26), we have

$$\begin{aligned}
&\mathbb{E}_{v_t\sim F_t}[\phi_t(h_t(\xi, v_t; F_{(1,T)}); F_{(1,T)})] \\
&= \mathbb{E}_{v_t\sim F_t}\left[\phi_t\big(\xi + \hat{u}_t(\xi, v_t; v_t; F_{(1,T)}) - \mathbb{E}_{v_t'\sim F_t}[\hat{u}_t(\xi, v_t'; v_t'; F_{(1,T)})]; F_{(1,T)}\big)\right] \\
&= \mathbb{E}_{v_t\sim \hat{F}_t}\left[\phi_t\big(\xi + \hat{u}_t(\xi, v_t + \varepsilon_t a_t; v_t + \epsilon_t a_t; F_{(1,T)}) - \mathbb{E}_{v_t'\sim F_t}[\hat{u}_t(\xi, v_t'; v_t'; F_{(1,T)})]; F_{(1,T)}\big)\right] \\
&\le \mathbb{E}_{v_t\sim \hat{F}_t}\left[\phi_t\big(\xi + \hat{u}_t(\xi, v_t; v_t; F_{(1,T)}) - \mathbb{E}_{v_t'\sim F_t}[\hat{u}_t(\xi, v_t'; v_t'; F_{(1,T)})] + \Delta a_t; F_{(1,T)}\big)\right] + 2\Delta a_t \qquad (27)
\end{aligned}$$

We finish the induction by plugging (24) and (27) into (23). □

D. Omitted Proofs of Section 4

D.1. Proof of Lemma 4.2

Proof. Since $B(g, y; \hat{F}_{(1,T)}) = \langle x, p, \text{balU}\rangle$ is stage-IC and BU for $\hat{F}_{(1,T)}$, the mechanism is also stage-IC and BU for $F_{(1,T)}$ by the observation that these two properties do not depend on the underlying distributions. Recall that BU implies ex-post IR. Therefore, $B(g, y; \hat{F}_{(1,T)})$ is also ex-post IR for $F_{(1,T)}$.

As for the balance independence property, since the local-stage mechanism $\langle x_t, \hat{p}_t\rangle$ constructed by $B(g, y; \hat{F}_{(1,T)})$ according to Definition 3.1 is a stage-IC and IR mechanism, by Corollary B.3, for any $\text{bal} \ge 0$, we have

$$\begin{aligned}
\mathbb{E}_{v_t\sim F_t}[v_t \cdot x_t(\text{bal}, v_t) - \hat{p}_t(\text{bal}, v_t)] - \Delta a_t &\le \mathbb{E}_{v_t\sim \hat{F}_t}[v_t \cdot x_t(\text{bal}, v_t) - \hat{p}_t(\text{bal}, v_t)] \\
&\le \mathbb{E}_{v_t\sim F_t}[v_t \cdot x_t(\text{bal}, v_t) - \hat{p}_t(\text{bal}, v_t)] + \Delta a_t.
\end{aligned}$$

Using the fact that $p_t(\text{bal}, v_t) = \hat{p}_t(\text{bal}, v_t) + s_t(\text{bal})$ and $\langle x_t, p_t\rangle$ is BI for $\hat{F}_{(1,T)}$, we have

$$\mathbb{E}_{v_t\sim F_t}[v_t \cdot x_t(\text{bal}, v_t) - p_t(\text{bal}, v_t)] \in [c - \Delta a_t, c + \Delta a_t]$$

where $c = \mathbb{E}_{v_t\sim \hat{F}_t}[v_t \cdot x_t(\text{bal}, v_t) - p_t(\text{bal}, v_t)]$ is a constant independent of bal. Therefore, we prove that $B(g, y; \hat{F}_{(1,T)})$ is $\delta_{(1,T)}$-BI with $\delta_t = 2\Delta a_t$ for $F_{(1,T)}$. □

D.2. Proof of Lemma 4.3

Proof. We consider a fixed combination of $b_{(1,t)}$ and v_t. Let $(X_{t'}, P_{t'})$ be a random variable representing the stage mechanism at stage t'. Let $\left(X_{t'}^{OPT}, P_{t'}^{OPT}\right)_{t'=t+1}^{T}$ be the sequence of stage mechanisms corresponding to the optimal play for stages between t and $(T-1)$ and let $\left(X_{t'}^{Truthful}, P_{t'}^{Truthful}\right)_{t'=t+1}^{T}$ be the sequence of stage mechanisms corresponding to playing truthfully for stages between t and $(T-1)$.

By playing truthfully for all stages between t and T, the buyer's utility is

$$u_t^{Truthful} = \mathbb{E}_{\left(X_{t'}^{Truthful}, P_{t'}^{Truthful}\right)_{t'=t+1}^{T}}\left[\sum_{t'=t+1}^{T} \gamma^{t'-t} \cdot \mathbb{E}_{v_{t'}\sim F_{t'}}\left[v_{t'} \cdot X_{t'}^{Truthful}(v_{t'}) - P_{t'}^{Truthful}(v_{t'})\right]\right].$$

As for the optimal play, the buyer's utility is at most

$$\begin{aligned}
u_t^{OPT} &= \mathbb{E}_{\left(X_{t'}^{OPT}, P_{t'}^{OPT}\right)_{t'=t+1}^{T}}\left[\sum_{t'=t+1}^{T} \gamma^{t'-t} \cdot \mathbb{E}_{v_{t'}\sim F_{t'}}\left[\max_b \left\{v_{t'} \cdot X_{t'}^{OPT}(b) - P_{t'}^{OPT}(b)\right\}\right]\right] \\
&= \mathbb{E}_{\left(X_{t'}^{OPT}, P_{t'}^{OPT}\right)_{t'=t+1}^{T}}\left[\sum_{t'=t+1}^{T} \gamma^{t'-t} \cdot \mathbb{E}_{v_{t'}\sim F_{t'}}\left[v_{t'} \cdot X_{t'}^{OPT}(v_{t'}) - P_{t'}^{OPT}(v_{t'})\right]\right].
\end{aligned}$$

where the second equality is due to the fact that the mechanism is stage-IC for $F_{(1,T)}$. Since the mechanism is $\delta_{(1,T)}$-BI, we have

$$\begin{aligned}U_t(b_{(1,t)};F_{(1,T)};\hat{F}_{(1,T)}) &\le u_t^{OPT}\\ &\le u_t^{Truthful}+\sum_{t'=t+1}^{T}\gamma^{t'-t}\cdot\beta_{t'}\\ &\le U_t(b_{(1,t-1)},v_t;F_{(1,T)};\hat{F}_{(1,T)})+\sum_{t'=t+1}^{T}\gamma^{t'-t}\cdot\beta_{t'}.\end{aligned}$$

□

D.3. Proof of Lemma 4.5

Proof. Since for each stage, both the stage mechanism $\langle x,p\rangle$ and the random posted price mechanism are stage-IC, the mixing of them are also stage-IC. For $\delta_{(1,T)}$-BI, notice that we apply the same random posted price mechanism for all $\mathrm{bal}\ge 0$, which is indeed a mechanism does not depend on bal. Therefore, the random posted price mechanism is BI for any distribution. Mixture of a $\delta_{(1,T)}$-BI mechanism and a BI mechanism result in a $\delta_{(1,T)}$-BI mechanism. Moreover, it is easy to check that the mechanism is BU since the balance update rule is scaled to keep track on $\langle x,p\rangle$ only and the random posted price mechanism generates non-negative utility for the buyer. Recall that BU implies ex-post IR. Therefore, the mechanism is also ex-post IR.

For $\eta_{(1,T)}$-DIC, notice that the buyer's expected utility gain in the t-th round is at most δ_t because the mechanism is $\delta_{(1,T)}$-BI. Therefore, for a buyer who discounts the future with discounting factor γ, the expected gain in the future by misreporting at round t is at most $\sum_{t'=t+1}^{T}\gamma^{t'-t}\delta_{t'}$. However, in the random posted price mechanism at round t, the utility loss of a buyer with true valuation v_t from overbidding in a magnitude of m_t is

$$\int_{v_t}^{v_t+m_t}\frac{b-v_t}{a_t}db=\frac{m_t^2}{2a_t}.$$

By a similar calculation, the utility loss of a buyer with true valuation v_t from underbidding in a magnitude of m_t is also $\frac{m_t^2}{2a_t}$. Thus, we have

$$\lambda\cdot\frac{m_t^2}{2a_t}\le\sum_{t'=t+1}^{T}\gamma^{t'-t}\cdot\delta_{t'}\Rightarrow m_t\le\sqrt{\frac{2a_t}{\lambda}\cdot\sum_{t'=t+1}^{T}\gamma^{t'-t}\cdot\delta_{t'}}.$$

□

D.4. Proof of Claim D.1

Claim D.1. $\sum_t\eta_t=O(\sqrt{\frac{\Delta}{\lambda}}T)$.

Proof. Notice that we have

$$\begin{aligned}\sum_t\eta_t &=\sum_t\sqrt{\frac{4a_t\Delta}{\lambda}\cdot\sum_{t'=t+1}^{T}\gamma^{t'-t}a_{t'}}=\sqrt{\frac{4\Delta}{\lambda}}\sum_t\sqrt{a_t}\sqrt{\sum_{t'=t+1}^{T}\gamma^{t'-t}a_{t'}}\\ &\le\sqrt{\frac{4\Delta}{\lambda}}\sqrt{\sum_t a_t}\sqrt{\sum_t\sum_{t'=t+1}^{T}\gamma^{t'-t}a_{t'}}\le\sqrt{\frac{4\Delta}{\lambda}}\sqrt{\sum_t a_t}\sqrt{\frac{\gamma}{1-\gamma}\sum_t a_t}\\ &\le\sqrt{\frac{4\gamma\Delta}{(1-\gamma)\lambda}}\cdot c_aT.\end{aligned}$$

where the first inequality follows the Cauchy-Schwarz inequality. □

D.5. Revenue Robust Mechanism

Definition D.2 (Revenue Robust Mechanism)**.** *For a core bank account mechanism $\langle g, y\rangle$ for $\hat{F}_{(1,T)}$, assuming the misreport of the buyer at stage t is at most η_t, we construct a revenue robust core bank account mechanism $\langle \tilde{g}, \tilde{y}\rangle$ such that $\tilde{g}_0 = g_0$ and*

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

where $\beta_t \geq \Delta a_t$ for all t.

Lemma D.3. *$\langle \tilde{g}, \tilde{y}\rangle$ constructed according to Definition D.2 is a core bank account mechanism for $\hat{F}_{(1,T)}$.*

Proof. We verify that $\langle \tilde{g}, \tilde{y}\rangle$ is a core bank account mechanism for $\hat{F}_{(1,T)}$ by checking the requirements in Definition 3.1.

First, $\tilde{y}_t(b_{(1,t)})$ is the sub-gradient of $\tilde{g}_t(b_{(1,t)})$ with respect to b_t since

$$\frac{\partial \tilde{g}_t(b_{(1,t)})}{\partial b_t} = \frac{\partial g_t(\tilde{g}_{t-1}(b_{(1,t-1)}), b_t)}{\partial b_t} = y_t(\tilde{g}_{t-1}(b_{(1,t-1)}), b_t) = \tilde{y}_t(b_{(1,t)})$$

where the second equality uses the fact that y_t is a sub-gradient of g_t with respect to b_t.

Next, $\tilde{g}_t(b_{(1,t)})$ is convex in b_t and weakly increasing in b_t since $g_t(\tilde{g}_{t-1}(b_{(1,t-1)}), b_t)$ is convex in b_t and weakly increasing in b_t. Moreover, by the definition of $\tilde{g}$, $\tilde{g}$ is symmetric. Finally, to check the consistency of $\tilde{g}$, we have

$$\begin{aligned}\tilde{g}_{t-1}(b_{(1,t-1)}) - \mathbb{E}_{v_t\sim\hat{F}_t}[g_t(b_{(1,t)})] &= \tilde{g}_{t-1}(b_{(1,t-1)}) - \mathbb{E}_{v_t\sim\hat{F}_t}[g_t(\tilde{g}_{t-1}(b_{(1,t-1)}), b_t) + \beta_t + \eta_t] \\ &= \tilde{g}_{t-1}(b_{(1,t-1)}) - \big(\tilde{g}_{t-1}(b_{(1,t-1)}) + \beta_t + \eta_t\big) \\ &= -\beta_t - \eta_t\end{aligned}$$

where the second inequality uses the fact that g is consistent. Therefore, we have

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

where $\chi_t(\tilde{g}) = -\beta_t - \eta_t$. Moreover, since we shift the state up, $\tilde{g}_t(b_{(1,t)}) \geq 0$ for all $b_{(1,t)}$. $\square$

For convenience, we consider a program $\theta_t(\xi) = \psi_t(\xi; B(g, y; \hat{F}_{(1,T)}); \hat{F}_{(1,T)})$.

Lemma D.4. *When $\beta_t = \Delta a_t$ and θ satisfies $\theta_t(\xi+\delta) \geq \theta_t(\xi) - \delta$ for all t, $\xi \geq 0$, and $\delta \geq 0$, we have*

$$\begin{aligned}&\operatorname{Rev}\big(B(\tilde{g}, \tilde{y}; \hat{F}_{(1,T)}), F_{(1,T)}\big) \\ &\geq \operatorname{Rev}\big(B(g, y; \hat{F}_{(1,T)}), \hat{F}_{(1,T)}\big) - (c_f + 3)\sum_t (\Delta a_t + \eta_t).\end{aligned}$$

Proof. Let $\langle z_t(\xi, \cdot), \hat{q}_t(\xi, \cdot)\rangle$ be the local-stage mechanism in $B(g, y; \hat{F}_{(1,T)})$. Moreover, let $\hat{u}_t(\xi, \cdot)$ be the utility function of $\langle z_t(\xi, \cdot), \hat{q}_t(\xi, \cdot)\rangle$. Notice that by our construction, the local-stage mechanism used by $B(\tilde{g}, \tilde{y}; \hat{F}_{(1,T)})$ is the same as $B(g, y; \hat{F}_{(1,T)})$ when they share the same state.

For convenience, let $\rho_t(\xi) = \psi_t(\xi; B(\tilde{g}, \tilde{y}; \hat{F}_{(1,T)}); F_{(1,T)})$. By Lemma 3.3 and the fact that $\mu_T(\tilde{g}) \geq 0$ from our construction, we have

$$\rho_0(\tilde{g}_0) \leq \operatorname{Rev}\big(B(\tilde{g}, \tilde{y}; \hat{F}_{(1,T)}), F_{(1,T)}\big) \tag{28}$$

$$\theta_0(g_0) = \operatorname{Rev}\big(B(g, y; \hat{F}_{(1,T)}), \hat{F}_{(1,T)}\big) \tag{29}$$

Next, we prove by a backward induction from $t = T$ to $t = 0$ to show that for all t,

$$\rho_t(\xi) \geq \theta_t(\xi) - \sum_{t'=t+1}^{T} \big((c_f + 2)\Delta a_{t'} + \beta_{t'} + (c_f + 3)\eta_t\big) \tag{30}$$

For the base case where $t = T$, it is true since we have $\theta_t(\xi) = \rho_t(\xi) = -\xi$ by the boundary condition. Assume it is true for all $t' \geq t$, for $t-1$, we have

$$\begin{aligned}
\rho_{t-1}(\xi) &= \mathbb{E}_{v_t \sim F_t}\left[z_t(\xi, v_t') \cdot v_t' + \rho_t\big(\tilde{g}_t(\xi, v_t')\big)\right] \\
&= \mathbb{E}_{v_t \sim F_t}\left[z_t(\xi, v_t') \cdot v_t' + \rho_t\big(\xi + \hat{u}_t(\xi, v_t') - \mathbb{E}_{v_t'' \sim \hat{F}_t}[\hat{u}_t(\xi, v_t'')] + \beta_t + \eta_t\big)\right] \\
&= \mathbb{E}_{\hat{v}_t \sim \hat{F}_t}\left[z_t(\xi, v_t') \cdot v_t' + \rho_t\big(\xi + \hat{u}_t(\xi, v_t') - \mathbb{E}_{v_t'' \sim \hat{F}_t}[\hat{u}_t(\xi, v_t'')] + \beta_t + \eta_t\big)\right] \qquad (31)
\end{aligned}$$

In the last equality, we assume ϵ follows G_t and we let $\hat{v}_t$ be the valuation drawn from $\hat{F}_t$ and v_t' be the reported bid given the buyer's true valuation is $v_t + a_t\epsilon_t$. Therefore, we have $v_t' \in [v_t - \Delta a_t - \eta_t, v_t + \Delta a_t + \eta_t]$. By Lemma B.4, we have

$$\mathbb{E}_{\hat{v}_t \sim \hat{F}_t}[z_t(\xi, v_t') \cdot v_t'] \geq \mathbb{E}_{\hat{v}_t \sim \hat{F}_t}[z_t(\xi, \hat{v}_t) \cdot \hat{v}_t] - (c_f + 1)(\Delta a_t + \eta_t) \qquad (32)$$

Moreover, by Lemma B.2, we have

$$\begin{aligned}
&\xi + \hat{u}_t(\xi, \hat{v}_t) - \mathbb{E}_{v_t'' \sim \hat{F}_t}[\hat{u}_t(\xi, v_t'')] - \Delta a_t + \beta_t \\
&\leq \xi + \hat{u}_t(\xi, v_t') - \mathbb{E}_{v_t'' \sim \hat{F}_t}[\hat{u}_t(\xi, v_t'')] + \beta_t + \eta_t \\
&\leq \xi + \hat{u}_t(\xi, \hat{v}_t) - \mathbb{E}_{v_t'' \sim \hat{F}_t}[\hat{u}_t(\xi, v_t'')] + \Delta a_t + \beta_t + 2\eta_t \qquad (33)
\end{aligned}$$

Henceforth, from (33), we have

$$\begin{aligned}
&\rho_t\big(\xi + \hat{u}_t(\xi, v_t') - \mathbb{E}_{v_t'' \sim \hat{F}_t}[\hat{u}_t(\xi, v_t'')] + \beta_t + \eta_t\big) \\
&\geq \theta_t\big(\xi + \hat{u}_t(\xi, v_t') - \mathbb{E}_{v_t'' \sim \hat{F}_t}[\hat{u}_t(\xi, v_t'')] + \beta_t + \eta_t\big) - \sum_{t'=t+1}^{T}\big((c_f+2)\Delta a_{t'} + \beta_{t'} + (c_f+3)\eta_{t'}\big) \\
&\geq \theta_t\big(\xi + \hat{u}_t(\xi, v_t) - \mathbb{E}_{v_t'' \sim \hat{F}_t}[\hat{u}_t(\xi, v_t'')] - \Delta a_t + \beta_t\big) - 2\Delta a_t - 2\eta_t - \sum_{t'=t+1}^{T}\big((c_f+2)\Delta a_{t'} + \beta_{t'} + (c_f+3)\eta_{t'}\big) \\
&\geq \theta_t\big(\xi + \hat{u}_t(\xi, v_t) - \mathbb{E}_{v_t'' \sim \hat{F}_t}[\hat{u}_t(\xi, v_t'')]\big) - \Delta a_t - \beta_t - 2\eta_t - \sum_{t'=t+1}^{T}\big((c_f+2)\Delta a_{t'} + \beta_{t'} + (c_f+3)\eta_{t'}\big) \qquad (34)
\end{aligned}$$

where the first inequality follows the induction hypothesis and the last two inequalities follows the assumption in the lemma statement that $\theta_t(\xi + \delta) \geq \theta_t(\xi) - \delta$ for $\delta > 0$. We finish the induction by plugging (32) and (34) into (31). Finally, we finish the proof by combining (30) with $t = 0$, (28), (29), and the fact that $\tilde{g}_0 = g_0$. □

E. Omitted Proofs of Appendix A

E.1. Proof of Lemma A.1

Proof. First, since the mechanism is $\eta_{(1,T)}$-DIC, the misreport within phase k at stage t is bounded by η_t. We next bound the additional misreport for changing the estimation for the next phase. Note that the utility gain starting from phase k is at most $\sum_{t' \geq \ell_k} \gamma^{t'-t} \cdot a_{t'}$. Under Assumption 5.1, $a_{t'} \leq c_a \cdot t'$. Therefore, we have for $t \in \hat{E}_{k-1}$,

$$\sum_{t' \geq \ell_k} \gamma^{t'-t} \cdot a_{t'} \leq c_a \cdot \frac{\gamma^{\ell_k - t}}{(1-\gamma)^2} \leq \frac{c_a}{(1-\gamma)^2 \cdot \ell_k^5}$$

Recall that at round t, our robust dynamic mechanism is mixed with a random posted price auction with price uniformly drawn from $[0, a_t]$ with probability λ. Therefore, the additional misreport $\bar{m}_t$ for $t \in \hat{E}_{k-1}$ is at most

$$\lambda_k \cdot \frac{\bar{m}_t^2}{2a_t} \leq \frac{c_a}{(1-\gamma)^2 \cdot \ell_k^5} \Rightarrow \bar{m}_t \leq \sqrt{\frac{2c_a \cdot a_t}{\lambda_k \cdot (1-\gamma)^2 \cdot \ell_k^5}} \leq \sqrt{\frac{2}{\lambda}} \cdot \frac{c_a}{(1-\gamma) \cdot \ell_k^2}$$

where the last inequality is due to $a_t \leq c_a \cdot t \leq c_a \cdot \ell_k$.

To bound the number of lies, for $t \in \hat{E}_{k-1}$, Let $L(j)$ be the number of lies for the first j stages in $\hat{E}_{k-1}$ and $\mathsf{EL}(j)$ be the expected number of lies from stage $(\ell_{k-1} + j)$. Recall that since we sample w_t uniformly from $[0, 1]$ and notice that a lie occurs only if

$$v_t - \eta_t - \bar{m}_t \le a_t \cdot w_t \le v_t + \eta_t + \bar{m}_t,$$

which happens with probability at most $2c_f \cdot \frac{\eta_t + \bar{m}_t}{a_t}$. Therefore,

$$\mathsf{EL}(j) \le 2c_f \cdot \left(\frac{\eta_t}{a_t} + \frac{c'}{\sqrt{\lambda}_k \cdot \ell_k^2} \right).$$

with $t = \ell_{k-1} + j$ and $c' = \frac{\sqrt{2}c_a}{1-\gamma}$. Notice that $\mathbb{E}[L(j) - L(j-1) - \mathsf{EL}(j)] = 0$, which implies that $L(j) - \sum_{j'=0}^{j} \mathsf{EL}(j')$ forms a martingale. Henceforth, by multiplicative Azuma's inequality (see Lemma 10 (Koufogiannakis & Young, 2014)) and denoting $\ell = |\hat{E}_{k-1}|$, we have

$$\Pr[L(\ell) \ge 2(1+\delta) \sum_{j'=0}^{\ell-1} \mathsf{EL}(j')] \le \exp\left(-\frac{\delta}{2} \cdot \sum_{j'=0}^{\ell-1} \mathsf{EL}(j') \right)$$

By setting $\delta = 2\log \ell_k / \big(\sum_{j'=0}^{\ell-1} \mathsf{EL}(j') \big)$, with probability at least $1 - \frac{1}{\ell_k}$, we have

$$L(\ell) = O\left(\log \ell_k + \sum_{t \in \hat{E}_k} \left(\frac{\eta_t}{a_t} + \frac{1}{\sqrt{\lambda}_k \cdot \ell_k^2} \right) \right) = O\left(\log \ell_k + \sum_{t \in \hat{E}_k} \frac{\eta_t}{a_t} \right).$$

□

E.2. Omitted Proofs for Semi-clairvoyant Mechanism

Our semi-clairvoyant mechanism is simply to run the optimal dynamic mechanism for each batch separately. To prove Lemma A.4, we first show that:

Lemma E.1. *There exists a dynamic mechanism for $F_{(1,T)}$ such that its revenue performance is $\sum_{t=1}^{T} \mathbb{E}[v_t] - O\left(\log T \cdot \sqrt{\sum_{t=1}^{T} a_t^2} \right)$.*

Proof. Let t^* be the first stage t such that $\sum_{t'=1}^{t} \mathbb{E}[v_{t'}] \ge 3 \log T \cdot \sqrt{\sum_{t'=1}^{T} a_{t'}^2}$. Consider the following simple bank account mechanism.

- For the first t^* stages, run a give for free mechanism, in which the item is always allocated to the buyer without payment while the balance update policy increases the balance by the buyer's reported bid. Formally, for $t \le t^*$,

$$x_t(\mathrm{bal}, b_t) = 1, p_t(\mathrm{bal}, b_t) = 0, \text{ and } \mathrm{balU}_t(\mathrm{bal}, b_t) = \mathrm{bal} + b_t;$$

- For the remaining stages, run a give for free mechanism with entry fee $\mathbb{E}[v_t]$ if the balance is enough. Formally, for $t > t^*$, if $\mathrm{bal} \ge \mathbb{E}[v_t]$,

$$x_t(\mathrm{bal}, b_t) = 1, p_t(\mathrm{bal}, b_t) = \mathbb{E}[v_t], \text{ and } \mathrm{balU}_t(\mathrm{bal}, b_t) = \mathrm{bal} + b_t - \mathbb{E}[v_t];$$

  otherwise, when $\mathrm{bal} < \mathbb{E}[v_t]$, run a null mechanism

$$x_t(\mathrm{bal}, b_t) = 0, p_t(\mathrm{bal}, b_t) = 0, \text{ and } \mathrm{balU}_t(\mathrm{bal}, b_t) = \mathrm{bal}.$$

Let $y_t = \sum_{t'=1}^{t} \mathbb{E}[v_{t'}]$ for $t \le t^*$ and $y_t = \sum_{t'=1}^{t^*} \mathbb{E}[v_{t'}]$. Notice that the balance $\mathrm{bal}_t - y_t$ forms a martingale. Therefore, by Azuma's inequality, we have for $t > t^*$,

$$\begin{aligned}\Pr\left[\mathrm{bal}_t < \mathbb{E}[v_t]\right] &= \Pr\left[\mathrm{bal}_t - y_{t^*} < -(y_{t^*} - \mathbb{E}[v_t])\right] \\ &\le \exp\left(-\frac{(y_{t^*} - \mathbb{E}[v_t])^2}{2\sum_{t'=1}^{t} a_{t'}^2}\right) \\ &\le \exp\left(-\frac{4\sum_{t'=1}^{T} a_{t'}^2 \log T}{2\sum_{t'=1}^{T} a_{t'}^2}\right) \le \frac{1}{T^2}.\end{aligned}$$

where the second inequality is due to $\mathbb{E}[v_t] \le a_t \le \sqrt{\sum_{t'=1}^{T} a_{t'}^2}$. Therefore, by union bound, we have

$$\Pr\left[\exists\, t, \mathrm{bal}_t < \mathbb{E}[v_t]\right] \le \frac{1}{T}.$$

As a result, with probability at least $1 - \frac{1}{T}$, we can obtain revenue

$$\sum_{t'=t^*+1}^{T} \mathbb{E}[v_t] = \sum_{t'=1}^{T} \mathbb{E}[v_t] - \sum_{t'=1}^{t^*} \mathbb{E}[v_t] \ge \sum_{t'=1}^{T} \mathbb{E}[v_t] - 4\log T \cdot \sqrt{\sum_{t'=1}^{T} a_{t'}^2},$$

which concludes the proof. □

Therefore, by running the optimal dynamic mechanism separately for batch j suffers a revenue loss at most $\tilde{O}(\sqrt{\sum_{t\in\mathcal{B}_j} a_t^2})$. Taking the summation over batches concludes the proof of Lemma A.4.

F. Hybrid Mechanism

To reduce the the number of lies, observe that the buyer has no incentive to misreport if the mechanism is static, i.e., the mechanism does not depend on the history. However, offering a static mechanism may cause huge revenue loss (Papadimitriou et al., 2016). We make a trade-off between the dynamic mechanism and the static mechanism to develop a hybrid mechanism, which contains both dynamic stages dependent on the history and static stages independent of the history. The full proofs of this section are deferred to Appendix F.3.

We first partition the stages in E_k into two sets according to the size of a_t. More precisely, for a fixed $\omega \in (0, 1)$, let $E_k^\omega = \{t \in E_k | a_t \le \ell_k^\omega\}$ be a subset of E_k in which the magnitude of a_t is bounded by ℓ_k^ω. We are going to offer static mechanisms for stages in E_k^ω and dynamic mechanism for stages in $E_k \setminus E_k^\omega$.

Moreover, let t^* be the minimum index of a stage in E_k such that $\sum_{t\le t^*, t\in E_k^\omega} a_t \ge 4\ell_k^{\frac{1}{2}(\omega+1)} \log \ell_k$ and let $\tilde{E}_k^\omega = \{t | t \le t^*, a_t \in E_k^\omega\}$ be the set of stages that are in E_k^ω and before t^*. If such t^* does not exist, simply set $t^* = \ell_{k+1} - 1$.

Definition F.1 (Hybrid Mechanism). *For a core bank account mechanism $\langle g, y\rangle$ for $\hat{F}_{(1,\ell_{k+1}-1)}$, we construct a hybrid mechanism $\langle g^{hybrid,k}, y^{hybrid,k}\rangle$ for phase k:*

- $g_0^{hybrid,k} = g_0 + \sum_{t\in\tilde{E}_k^\omega} \mathbb{E}_{v_t\sim\hat{F}_t}[v_t]$;
- *For* $t < \ell_k$,
  - $g_t^{hybrid,k}(b_{(1,t)}) = g_{t-1}^{hybrid,k}(b_{(1,t-1)}) + g_t(b_{(1,t)}) - g_{t-1}(b_{(1,t-1)})$;
  - $y_t^{hybrid,k}(b_{(1,t)}) = y_t(b_{(1,t)})$;

As for $t \in E_k$,

- *For* $t \in E_k \setminus E_k^\omega$, *if* $g_{t-1}^{hybrid,k}(b_{(1,t-1)}) \ge g_{t-1}(b_{(1,t-1)})$

- $g_t^{hybrid,k}(b_{(1,t)}) = g_{t-1}^{hybrid,k}(b_{(1,t-1)}) + g_t(b_{(1,t)}) - g_{t-1}(b_{(1,t-1)})$;
- $y_t^{hybrid,k}(b_{(1,t)}) = y_t(b_{(1,t)})$;

otherwise, $g_t^{hybrid,k}(b_{(1,t)}) = g_{t-1}^{hybrid,k}(b_{(1,t-1)})$ *and* $y_t^{hybrid,k}(b_{(1,t)}) = 0$.

- *and for* $t \in \tilde{E}_k^{\omega}$*:* $g_t^{hybrid,k}(b_{(1,t)}) = g_{t-1}^{hybrid,k}(b_{(1,t-1)}) + b_t - \mathbb{E}_{v_t \sim \hat{F}_t}[v_t]$ *and* $y_t^{hybrid,k}(b_{(1,t)}) = 1$*;*
- *and finally for* $t \in E_k^{\omega} \setminus \tilde{E}_k^{\omega}$*, if* $g_{t-1}^{hybrid,k}(b_{(1,t-1)}) \geq \mathbb{E}_{v_t \sim \hat{F}_t}[v_t]$
  - $g_t^{hybrid,k}(b_{(1,t)}) = g_{t-1}^{hybrid,k}(b_{(1,t-1)}) + b_t - \mathbb{E}_{v_t \sim \hat{F}_t}[v_t]$;
  - $y_t^{hybrid,k}(b_{(1,t)}) = 1$;

  otherwise, $g_t^{hybrid,k}(b_{(1,t)}) = g_{t-1}^{hybrid,k}(b_{(1,t-1)})$ *and* $y_t^{hybrid,k}(b_{(1,t)}) = 0$.

It is easy to verify that $\langle g^{hybrid,k}, y^{hybrid,k}\rangle$ satisfies the requirements in Definition 3.1, except that $g^{hybrid,k}$ might not be symmetric. However, (1)-(5) is still well-defined with respect to the historical bids $b_{(1,t)}$ and Lemma 3.3 still holds. We can apply the technique in (Mirrokni et al., 2016b; 2018) to make it symmetric without any revenue loss while maintaining all the desired properties given in this section.

Intuitively, we first shift the initial state up by $\sum_{t\in \tilde{E}_k^{\omega}} \mathbb{E}_{v_t\sim \hat{F}_t}[v_t]$ and for the first $(k-1)$ phases, we apply follow-the-history operations so that the constructed hybrid mechanism implements the same local-stage mechanism as $\langle g, y\rangle$. For phase k, if the state is smaller than a threshold dependent on t, we offer a null local-stage mechanism in which the allocation probability is always 0. Otherwise, for stage $t \in E_k \setminus E_k^{\omega}$, i.e., the stage in which a_t is large, we apply a follow-the-history operation so that the local-stage mechanism is the same as $\langle g, y\rangle$. As for the stages with small a_t, we change the local-mechanism such that the allocation probability is always 1 no matter what the reported bid is. In particular, for $t \in \tilde{E}_k^{\omega}$, $g_t^{hybrid,k}(b_{(1,t)}) \geq 0$ for all $b_{(1,t)}$ since the initial state is shifted up by $\sum_{t\in \tilde{E}_k^{\omega}} \mathbb{E}_{v_t\sim \hat{F}_t}[v_t]$.

Lemma F.2. *Let* $I_1 = \bigcup_{t\in E_k} I_1^t$ *and* $I_2 = \bigcup_{t\in E_k} I_2^t$ *be the events such that*

$$I_1^t = \{v_{(1,t-1)} \mid t \in E_k \setminus E_k^{\omega}, g_{t-1}^{hybrid,k}(v_{(1,t-1)}) < g_{t-1}(v_{(1,t-1)})\}$$

and

$$I_2^t = \{v_{(1,t-1)} \mid t \in E_k^{\omega} \setminus \tilde{E}_k^{\omega}, g_{t-1}^{hybrid,k}(v_{(1,t-1)}) < \mathbb{E}_{v_t\sim \hat{F}_t}[v_t]\}.$$

We have $\Pr_{v_{(1,\ell_{k+1}-1)}\sim \hat{F}_{(1,\ell_{k+1}-1)}}[I_1 \cup I_2] \leq \frac{1}{\ell_k}$.

Lemma F.2 demonstrate that the probability that a null mechanism is offered is small, assuming the buyer reports truthfully. The argument is based on the martingale nature of $g_t^{hybrid,k}(b_{(1,t)}) - g_t(b_{(1,t)})$. In particular, due to our construction of $g^{hybrid,k}$, we have

$$g_{t-1}^{hybrid,k}(b_{(1,t-1)}) - g_{t-1}(b_{(1,t-1)}) = g_t^{hybrid,k}(b_{(1,t)}) - g_t(b_{(1,t)})$$

for $t < \ell_k$ or $t \in E_k \setminus E_k^{\omega}$ when I_1 does not happen. As a result, the martingale dynamic of $g_t^{hybrid,k}(b_{(1,t)}) - g_t(b_{(1,t)})$ only happens for $t \in E_k^{\omega}$, in which a_t is small. The next lemma demonstrates the revenue loss of our hybrid mechanism is small.

Lemma F.3.

$$\begin{aligned}&\operatorname{Rev}\big(B(g^{hybrid,k}, y^{hybrid,k}; \hat{F}_{(1,\ell_{k+1})}), \hat{F}_{(1,\ell_{k+1})}\big)\\ \geq\ &\operatorname{Rev}\big(B(g, y; \hat{F}_{(1,\ell_{k+1}-1)}), \hat{F}_{(1,\ell_{k+1})}\big) - O(\ell_k^{\frac{1}{2}(\omega+1)} \log \ell_k).\end{aligned}$$

The revenue loss simply depends on the choice of ω. To show the revenue performance, the key observation is that once neither event I_1 nor I_2 happens, the total welfare obtained in $\langle g^{hybrid,k}, y^{hybrid,k}\rangle$ is at least the total welfare obtained in $\langle g, y\rangle$ since $\langle g^{hybrid,k}, y^{hybrid,k}\rangle$ either follow the local-stage mechanism of $\langle g, y\rangle$ or its local-stage mechanism is a mechanism that always allocates the item. Moreover, we exploit the martingale nature of $g^{hybrid,k}$ and g to show that the expected final utility of the buyer in $\langle g^{hybrid,k}, y^{hybrid,k}\rangle$ is at most $O(\ell_k^{\frac{1}{2}(\omega+1)} \log \ell_k)$ higher than the expected final utility in $\langle g, y\rangle$.

F.1. Misreport in Hybrid Mechanism

Let $\bar{B}(g^{hybrid,k}, y^{hybrid,k}; \hat{F}_{(1,\ell_{k+1}-1)})$ be the bank account mechanism mixing $B(g^{hybrid,k}, y^{hybrid,k}; \hat{F}_{(1,\ell_{k+1}-1)})$ with a random posted price mechanism with probability λ_k as constructed in Definition 4.4. As a result, Theorem 4.8 and Claim D.1 applies. We will then shift the initial state in Definition F.1 with an addition amount $O(\sqrt{\frac{\Delta_k}{\lambda_k}}\ell_k)$ so that Lemma F.2 holds even when the buyer misreports while the additional revenue loss is $O(\sqrt{\frac{\Delta_k}{\lambda_k}}\ell_k)$. Let $next(t) = \min\{t' \in E_k \mid t' \in E_k \setminus E_k^\omega, t' > t\}$, which computes the minimum index of the future stage that is not in E_k^ω. The following lemma characterizes the number of lies.

Lemma F.4. *In $\bar{B}(g^{hybrid,k}, y^{hybrid,k}; \hat{F}_{(1,\ell_{k+1}-1)})$, $|L_k| = O\left(\ell_k^{1-\omega}\log \ell_k\right)$ with probability at least $1 - \frac{1}{\ell_k}$.*

Intuitively, the local-stage mechanisms for $t \in \tilde{E}_k^\omega$ are static and moreover, the local-stage mechanisms for $t \in E_k^\omega \setminus \tilde{E}_k^\omega$ are static, because the probability that event I_2 happens is small even when the buyer misreports. Therefore, the gain from the misreport in stage $t \in E_k^\omega$ with $\mathrm{next}(t) - t \geq 3\log \ell_k$ can only be obtained after $3\log \ell_k$ rounds, which is relatively small. Thus, the buyer has less incentive to misreport in these stages, which results in a small number of lies.

F.2. Direct Computation of the Hybrid Mechanism

In order to apply the framework of revenue robust mechanism in Section 4.2, it requires that $\theta_t(\xi) = \psi_t(\xi; B(g^{hybrid}, y^{hybrid}; \hat{F}_{(1,T)}); \hat{F}_{(1,T)})$ satisfies that $\theta_t(\xi + \delta) \geq \theta_t(\xi) - \delta$ for all t, $\xi \geq 0$, and $\delta \geq 0$.

In order to fulfill the requirement, we directly compute the optimal hybrid mechanism. For convenience, let $\phi_{t-1}^{hybrid}(\xi; \hat{F}_{(1,\ell_{k+1}-1)})$ be the optimal revenue for the sub-problem consisting of stages t to $(\ell_{k+1} - 1)$ with distribution $\hat{F}_{(1,\ell_{k+1}-1)}$, when g_{t-1} maps the history to ξ. For the boundary cases, $\phi_{\ell_{k+1}}^{hybrid,k}(\xi; \hat{F}_{(1,\ell_{k+1}-1)}) = -\xi$. We compute $\phi_{t-1}^{hybrid,k}(\xi; \hat{F}_{(1,T)})$ from program OPT-BAM with $\phi_t(\xi; \hat{F}_{(1,T)})$ replaced by $\phi_{t-1}^{hybrid,k}(\xi; \hat{F}_{(1,T)})$:

- For $t \leq \ell_k$ or $t \in E_k \setminus E_k^\omega$, apply program OPT-BAM to compute $\phi_{t-1}^{hybrid,k}(\xi; \hat{F}_{(1,T)})$;
- For $t \in \tilde{E}_k^\omega$,
  - if $\xi \geq \mathbb{E}_{v_t \sim \hat{F}_t}[v_t]$, set $\langle z_t, \hat{q}_t\rangle$ such that $z_t(\xi, v_t) = 1$ and $q_t(\xi, v_t) = 0$ for all v_t;
  - otherwise, set $\phi_{t-1}^{hybrid,k}(\xi; \hat{F}_{(1,T)}) = -\infty$.
- For $t \in E_k^\omega \setminus \tilde{E}_k^\omega$,
  - if $\xi \geq \mathbb{E}_{v_t \sim \hat{F}_t}[v_t]$, set $\langle z_t, \hat{q}_t\rangle$ such that $z_t(\xi, v_t) = 1$ and $q_t(\xi, v_t) = 0$ for all v_t;
  - otherwise, set $\langle z_t, \hat{q}_t\rangle$ such that $z_t(\xi, v_t) = 0$ and $q_t(\xi, v_t) = 0$ for all v_t.

Notice that the mechanism derived from $\phi_t^{hybrid,k}$ is static for $t \in \tilde{E}_k^\omega$ and for $t \in E_k^\omega \setminus \tilde{E}_k^\omega$ is static with high probability, since the probability that event I_2 happens is small. Moreover, for other stages, the local-stage mechanism is computed from program OPT-BAM, in which Lemma 3.6 is applicable. Therefore, with high probability, $\phi_t^{hybrid,k}$ has the desired property for us to apply the framework of revenue robust mechanism in Section 4.2.

Lemma F.5. *For any t, $\xi \geq 0$, and $\delta \geq 0$, with probability $1 - \frac{1}{\ell_k}$, $\phi_t^{hybrid,k}(\xi + \delta; \hat{F}_{(1,T)}) \geq \phi_t^{hybrid,k}(\xi; \hat{F}_{(1,T)}) - \delta$, except for $t \in E_k^\omega \setminus \tilde{E}_k^\omega$ and $\xi < \mathbb{E}_{v_t \sim \hat{F}_t}[v_t]$.*

F.3. Proofs of Hybrid Mechanism

F.3.1. A VALID CORE BANK ACCOUNT MECHANISM

We first verify the hybrid mechanism in Definition F.1 is a valid core bank account mechanism, except that $g^{hybrid,k}$ might not be symmetric.

First, it is straightforward to verify that $y_t^{hybrid,k}$ is a sub-gradient of $g_t^{hybrid,k}$ with respect to b_t and $g_t^{hybrid,k}$ is symmetric, convex in b_t, and weakly increasing in b_t.

For the consistency of $g^{hybrid,k}$, we will show that $\chi_t(g^{hybrid,k}) = 0$ for all t. First, for $t < \ell_k$, $g^{hybrid,k}$ simply follows g, and thus, $\chi_t(g^{hybrid,k}) = \chi_t(g) = 0$. For $t \in E_k \setminus E_k^\omega$, if $g_{t-1}^{hybrid,k}(b_{(1,t-1)}) \geq g_{t-1}(b_{(1,t-1)})$, we have

$$\begin{aligned}
\mathbb{E}_{v_t\sim\hat{F}_t}[g_t^{hybrid,k}(b_{(1,t-1)}, v_t)] &= \mathbb{E}_{v_t\sim\hat{F}_t}[g_{t-1}^{hybrid,k}(b_{(1,t-1)}) + g_t(b_{(1,t-1)}, v_t) - g_{t-1}(b_{(1,t-1)})] \\
&= g_{t-1}^{hybrid,k}(b_{(1,t-1)}) + \mathbb{E}_{v_t\sim\hat{F}_t}[g_t(b_{(1,t-1)}, v_t)] - g_{t-1}(b_{(1,t-1)})] \\
&= g_{t-1}^{hybrid,k}(b_{(1,t-1)}) - \chi_t(g) \\
&= g_{t-1}^{hybrid,k}(b_{(1,t-1)})
\end{aligned}$$

where the third equality follows the fact that g is consistent and the last equality uses the fact that $\chi_t(g) = 0$. If $g_{t-1}^{hybrid,k}(b_{(1,t-1)}) < g_{t-1}(b_{(1,t-1)})$, we have $g_t^{hybrid,k}(b_{(1,t-1)}, v_t) = g_{t-1}^{hybrid,k}(b_{(1,t-1)})$. Therefore, $\mathbb{E}_{v_t\sim\hat{F}_t}[g_t^{hybrid,k}(b_{(1,t-1)}, v_t)] = g_{t-1}^{hybrid,k}(b_{(1,t-1)})$, and thus, $\chi_t(g^{hybrid,k}) = 0$.

For $t \in E_k^\omega$, if $g_{t-1}^{hybrid,k}(b_{(1,t)}) \geq \mathbb{E}_{v_t\sim\hat{F}_t}[v_t]$, then we have

$$\begin{aligned}
\mathbb{E}_{v_t\sim\hat{F}_t}[g_t^{hybrid,k}(b_{(1,t-1)}, v_t)] &= \mathbb{E}_{v_t\sim\hat{F}_t}[g_{t-1}^{hybrid,k}(b_{(1,t-1)}) + v_t - \mathbb{E}_{v_t\sim\hat{F}_t}[v_t]] \\
&= g_{t-1}^{hybrid,k}(b_{(1,t-1)}) + \mathbb{E}_{v_t\sim\hat{F}_t}[v_t - \mathbb{E}_{v_t\sim\hat{F}_t}[v_t]] \\
&= g_{t-1}^{hybrid,k}(b_{(1,t-1)})
\end{aligned}$$

Otherwise, when $g_{t-1}^{hybrid,k}(b_{(1,t)}) < \mathbb{E}_{v_t\sim\hat{F}_t}[v_t]$, we provide a null mechanism and thus $\chi_t(g^{hybrid,k}) = 0$. Finally, we show that for all t and $b_{(1,t)}$, $g_t^{hybrid,k}(b_{(1,t)}) \geq 0$. By Claim C.1, it suffices to show that $\mu_{\ell_{k+1}-1}(g^{hybrid,k}) \geq 0$.

We first show that for all $t \leq t^*$,

$$g_t^{hybrid,k}(b_{(1,t)}) \geq g_t(b_{(1,t)}) + \sum_{t'\in\tilde{E}_k^\omega, t<t'\leq t^*} \mathbb{E}_{v_{t'}\sim\hat{F}_{t'}}[v_{t'}]$$

by an induction from $t = 0$ to $t = t^*$. The base case $t = 0$ is obviously true since $g_0^{hybrid,k} = g_0 + \sum_{t\in\tilde{E}_k^\omega} \mathbb{E}_{v_t\sim\hat{F}_t}[v_t]$. Assume it is true for all $t' < t$ and for t, if $t \in \tilde{E}_k^\omega$, we have

$$\begin{aligned}
g_t^{hybrid,k}(b_{(1,t)}) &= g_{t-1}^{hybrid,k}(b_{(1,t-1)}) + b_t - \mathbb{E}_{v_t\sim\hat{F}_t}[v_t] \\
&\geq g_{t-1}(b_{(1,t-1)}) + b_t + \sum_{t'\in\tilde{E}_k^\omega, t<t'\leq t^*} \mathbb{E}_{v_{t'}\sim\hat{F}_{t'}}[v_{t'}] \\
&\geq g_t(b_{(1,t)}) + \sum_{t'\in\tilde{E}_k^\omega, t<t'\leq t^*} \mathbb{E}_{v_{t'}\sim\hat{F}_{t'}}[v_{t'}]
\end{aligned}$$

where the first inequality follows the induction hypothesis and the second inequality follows (6) such that

$$g_t(b_{(1,t)}) - g_{t-1}(b_{(1,t-1)}) = \hat{u}_t\big(\mathrm{bal}_t(b_{(1,t-1)}), b_t\big) - \mathbb{E}_{v_t\sim\hat{F}_t}[\hat{u}_t\big(\mathrm{bal}_t(b_{(1,t-1)}), v_t\big)] \leq b_t$$

where $\hat{u}_t(b_{(1,t)}) = x_t(b_{(1,t)}) \cdot b_t - \hat{p}_t(b_{(1,t)})$ and x_t and $\hat{p}_t$ is derived from $B(g, y; \hat{F}_{(1,T)})$ according to Definition 3.1. The inequality follows $\hat{u}_t\big(\mathrm{bal}_t(b_{(1,t-1)}), b_t\big) \leq b_t$ and $\mathbb{E}_{v_t\sim\hat{F}_t}[\hat{u}_t\big(\mathrm{bal}_t(b_{(1,t-1)}), v_t\big)] \geq 0$. On the other hand, if $t \notin \tilde{E}_k^\omega$, we have

$$\begin{aligned}
g_t^{hybrid,k}(b_{(1,t)}) &= g_{t-1}^{hybrid,k}(b_{(1,t-1)}) + g_t(b_{(1,t)}) - g_{t-1}(b_{(1,t-1)}) \\
&\geq g_{t-1}(b_{(1,t-1)}) + g_t(b_{(1,t)}) - g_{t-1}(b_{(1,t-1)}) + \sum_{t'\in\tilde{E}_k^\omega, t<t'\leq t^*} \mathbb{E}_{v_{t'}\sim\hat{F}_{t'}}[v_{t'}] \\
&= g_t(b_{(1,t)}) + \sum_{t'\in\tilde{E}_k^\omega, t<t'\leq t^*} \mathbb{E}_{v_{t'}\sim\hat{F}_{t'}}[v_{t'}].
\end{aligned}$$

Therefore, we finish the proof of the induction. As a result, we have $\mu_{t^*} \geq 0$. As for $t > t^*$, by our construction, the update rule is in the form of either $g_t^{hybrid,k}(b_{(1,t)}) = g_{t-1}^{hybrid,k}(b_{(1,t-1)}) + C_a - C_b$ if $g_{t-1}^{hybrid,k}(b_{(1,t-1)}) > C_b$ or $g_t^{hybrid,k}(b_{(1,t)}) = g_{t-1}^{hybrid,k}(b_{(1,t-1)})$. Thus, $g_t^{hybrid,k}(b_{(1,t)}) \geq 0$ for all $b_{(1,t)}$ and $t > t^*$, which implies that $\mu_{\ell_{k+1}-1}(g^{hybrid,k}) \geq 0$.

F.3.2. PROOF OF LEMMA F.2

Proof. Let $d_t(v_{(1,t)}) = g_t^{hybrid,k}(v_{(1,t)}) - g_t(v_{(1,t)}) - \sum_{t\in\tilde{E}_k^\omega} \mathbb{E}_{v_t\sim\hat{F}_t}[v_t]$ and we have

$$d_t(v_{(1,t)}) = d_{t-1}(v_{(1,t-1)}) + \left(g_t^{hybrid,k}(v_{(1,t)}) - g_{t-1}^{hybrid,k}(v_{(1,t-1)})\right) - \left(g_t(v_{(1,t)}) - g_{t-1}(v_{(1,t-1)})\right).$$

For $t \in E_k \setminus E_k^\omega$, by the construction of $g_t^{hybrid,k}$, we have $d_t(v_{(1,t)}) = d_{t-1}(v_{(1,t-1)})$. Moreover, since $\chi_t(g) = \chi_t(g^{hybrid,k}) = 0$, for $t \in E_k^\omega$, we have

$$\begin{aligned}
&\mathbb{E}_{v_t\sim\hat{F}_t}[d_t(v_{(1,t)})] \\
&= d_{t-1}(v_{(1,t-1)}) + \left(\mathbb{E}_{v_t\sim\hat{F}_t}[g_t^{hybrid,k}(v_{(1,t)})] - g_{t-1}^{hybrid,k}(v_{(1,t-1)})\right) - \left(\mathbb{E}_{v_t\sim\hat{F}_t}[g_t(v_{(1,t)})] - g_{t-1}(v_{(1,t-1)})\right) \\
&= d_{t-1}(v_{(1,t-1)}).
\end{aligned}$$

and in addition, we have $|d_t(v_{(1,t)}) - d_{t-1}(v_{(1,t-1)})| \le 2a_t$. Therefore, $d_t(v_{(1,t)})$ forms a martingale with bounded difference $2a_t$ for $t \in E_k^\omega$ and $d_0 = 0$.

Notice that event I_1 occurs when $t \in E_k \setminus E_k^\omega$ and

$$g_t^{hybrid,k}(v_{(1,t)}) - g_t(v_{(1,t)}) = d_t(v_{(1,t)}) + \sum_{t'\in\tilde{E}_k^\omega} \mathbb{E}_{v_{t'}\sim\hat{F}_{t'}}[v_{t'}] < 0.$$

Moreover, event I_2 occurs when $t \in E_k^\omega \setminus \tilde{E}_k^\omega$ and

$$g_t^{hybrid,k}(v_{(1,t)}) = d_t(v_{(1,t)}) + g_t(v_{(1,t)}) + \sum_{t'\in\tilde{E}_k^\omega} \mathbb{E}_{v_{t'}\sim\hat{F}_{t'}}[v_{t'}] < \mathbb{E}_{v_t\sim\hat{F}_t}[v_t] \le a_t \le \ell_k^\omega.$$

Since $g_t(v_{(1,t)}) \ge 0$, the above inequality holds only if

$$d_t(v_{(1,t)}) \le \ell_k^\omega - \sum_{t'\in\tilde{E}_k^\omega} \mathbb{E}_{v_{t'}\sim\hat{F}_{t'}}[v_{t'}] \le -\frac{1}{2}\sum_{t'\in\tilde{E}_k^\omega} \mathbb{E}_{v_{t'}\sim\hat{F}_{t'}}[v_{t'}]$$

where the last inequality is due to $\sum_{t\in\tilde{E}_k^\omega} \mathbb{E}_{v_t\sim\hat{F}_t}[v_t] \ge 4\ell_k^{\frac{1}{2}(\omega+1)} \log \ell_k \gg \ell_k^\omega$. Thus, either event I_1 or event I_2 happens only if

$$d_t(v_{(1,t)}) \le -\frac{1}{2}\sum_{t'\in\tilde{E}_k^\omega} \mathbb{E}_{v_{t'}\sim\hat{F}_{t'}}[v_{t'}]$$

From Azuma's inequality again, we have

$$\Pr\left[d_t(v_{(1,t)}) < -\frac{1}{2}\sum_{t\in\tilde{E}_k^\omega} \mathbb{E}_{v_t\sim\hat{F}_t}[v_t]\right] \le \exp\left(-\frac{4\ell_k^{\omega+1}\log\ell_k}{2\sum_{t'\in E_k^\omega, t'\le t} a_t^2}\right) \le \exp\left(-\frac{2\ell_k^{\omega+1}\log\ell_k}{\sum_{t'\in E_k^\omega} a_t^2}\right) \le \frac{1}{\ell_k^2}.$$

where we use the fact that $\sum_{t\in E_k^\omega} a_t^2 \le \ell_k^{1+\omega}$ due to $\sum_{t\in E_k} a_t = \Theta(\ell_k)$, in which the maximum is obtained when there are $\ell_k^{1-\omega}$ stages with $a_t = \ell_k^\omega$. Finally, applying union bounds on ℓ_k stages, the probability that either event I_1 or event I_2 is at most $O(\frac{1}{\ell_k})$. □

F.3.3. PROOF OF LEMMA F.3

Proof. For simlicity, let $T = \ell_{k+1} - 1$ in this proof. Let

$$\rho_t(\xi) = \psi_t(\xi; B(g^{hybrid,k}, y^{hybrid,k}; \hat{F}_{(1,T)}); \hat{F}_{(1,T)})$$

and

$$\theta_t(\xi) = \psi_t(\xi; B(g, y; \hat{F}_{(1,T)}); \hat{F}_{(1,T)}).$$

From Lemma 3.3, we have

$$\mathrm{Rev}\big(B(g^{hybrid,k}, y^{hybrid,k}; \hat{F}_{(1,T)}), \hat{F}_{(1,T)}\big) = \rho_0(g_0^{hybrid,k}) + \mu_T(g^{hybrid,k}) \geq \rho_0(g_0^{hybrid,k})$$
$$\mathrm{Rev}\big(B(g, y; \hat{F}_{(1,T)}), \hat{F}_{(1,T)}\big) = \theta_0(g_0) + \mu_T(g) = \theta_0(g_0)$$

where we use the fact that $\mu_T(g^{hybrid,k}) \geq 0$ and $\mu_T(g) = 0$. If neither I_1 nor I_2 happens, from our construction of $\langle g^{hybrid,k}, y^{hybrid,k}\rangle$, we have for each $b_{(1,T)}$, if $t \in E_k^{\omega}$, then $y_t^{hybrid,k}(b_{(1,t)}) = 1 \geq y_t(b_{(1,t)})$; otherwise, $y_t^{hybrid,k}(b_{(1,t)}) = y_t(b_{(1,t)})$. Recall that from (7), we have

$$\rho_0(g_0^{hybrid,k}) = \mathbb{E}_{v_{(1,T)}\sim\hat{F}_{(1,T)}}[\sum_t y_t^{hybrid,k}(v_{(1,t)}) \cdot v_t] - \mathbb{E}_{v_{(1,T)}}[g_T^{hybrid,k}(v_{(1,T)})]$$
$$\theta_0(g_0^{hybrid,k}) = \mathbb{E}_{v_{(1,T)}\sim\hat{F}_{(1,T)}}[\sum_t y_t(v_{(1,t)}) \cdot v_t] - \mathbb{E}_{v_{(1,T)}}[g_T(v_{(1,T)})]$$

Therefore, we have

$$\mathbb{E}_{v_{(1,T)}\sim\hat{F}_{(1,T)}}[\sum_t y_t^{hybrid,k}(v_{(1,t)}) \cdot v_t] \geq \mathbb{E}_{v_{(1,T)}\sim\hat{F}_{(1,T)}}[\sum_t y_t(v_{(1,t)}) \cdot v_t]. \tag{35}$$

Next, we compare $\mathbb{E}_{v_{(1,T)}}[g_T^{hybrid,k}(v_{(1,T)})]$ and $\mathbb{E}_{v_{(1,T)}}[g_T(v_{(1,T)})]$. Notice that we have

$$\begin{aligned}
&g_T^{hybrid,k}(v_{(1,T)})\\
&= g_0^{hybrid,k} + \sum_t \Big(g_t^{hybrid,k}(v_{(1,t)}) - g_{t-1}^{hybrid,k}(v_{(1,t-1)})\Big)\\
&= g_0^{hybrid,k} + \sum_{t\in E_k^{\omega}} \Big(g_t^{hybrid,k}(v_{(1,t)}) - g_{t-1}^{hybrid,k}(v_{(1,t-1)})\Big) + \sum_{t\in E_k\setminus E_k^{\omega}} \Big(g_t^{hybrid,k}(v_{(1,t)}) - g_{t-1}^{hybrid,k}(v_{(1,t-1)})\Big)\\
&= g_0 + \sum_{t\in\tilde{E}_k^{\omega}} \mathbb{E}_{v_t\sim\hat{F}_t}[v_t] + \sum_{t\in E_k^{\omega}} \Big(g_t^{hybrid,k}(v_{(1,t)}) - g_{t-1}^{hybrid,k}(v_{(1,t-1)})\Big) + \sum_{t\in E_k\setminus E_k^{\omega}} \big(g_t(v_{(1,t)}) - g_{t-1}(v_{(1,t-1)})\big)
\end{aligned}$$

Let $d_t(v_{(1,t)}) = g_t^{hybrid,k}(v_{(1,t)}) - g_t(v_{(1,t)}) - \sum_{t\in\tilde{E}_k^{\omega}} \mathbb{E}_{v_t\sim\hat{F}_t}[v_t]$, which forms a martingale with $d_0 = 0$, and bounded difference $2a_t$ for $t \in E_k^{\omega}$ and difference 0 for $t \in E_k \setminus E_k^{\omega}$.

Notice that $\sum_{t\in E_k^{\omega}} a_t^2 \leq \ell_k^{1+\omega}$, in which the maximum is obtained when there are $T^{1-\omega}$ stages with $a_t = \ell_k^{\omega}$. By Azuma's inequality, we have

$$\Pr[|d_T(v_{(1,T)})| \geq t] \leq 2\exp\left(\frac{-t^2}{2\sum_{t\in E_k^{\omega}} 4a_t^2}\right) \leq 2\exp\left(\frac{-t^2}{8\ell_k^{1+\omega}}\right).$$

Let $t = 4\ell_k^{\frac{1}{2}(\omega+1)}\log\ell_k$ and we have

$$\Pr\Big[|d_T(v_{(1,T)})| \geq 4\ell_k^{\frac{1}{2}(\omega+1)}\log\ell_k\Big] \leq \frac{2}{T^2}.$$

Finally, since $|d_T(v_{(1,T)})| \leq 2\sum_t a_t = cT$ for some constant c, we have

$$\begin{aligned}
-\mathbb{E}_{v_{(1,T)}}[g_T^{hybrid,k}(v_{(1,T)})] &\geq -\mathbb{E}_{v_{(1,T)}}[g_T(v_{(1,T)})] - \sum_{t\in\tilde{E}_k^{\omega}} \mathbb{E}_{v_t\sim\hat{F}_t}[v_t] - \left(4\ell_k^{\frac{1}{2}(\omega+1)}\log\ell_k + \frac{2}{T^2}\cdot 2cT\right)\\
&\geq -\mathbb{E}_{v_{(1,T)}}[g_T(v_{(1,T)})] - 8\ell_k^{\frac{1}{2}(\omega+1)}\log\ell_k - \frac{2c}{T}
\end{aligned} \tag{36}$$

By Lemma F.2, the probability that either I_1 or I_2 happens is $O(\frac{1}{T})$. Moreover, the revenue loss when either event happens is at most $\sum_t a_t = cT$. Therefore, the revenue loss caused from events I_1 and I_2 is at most c. Combining with (35) and (36), we finish the proof. □

F.3.4. PROOF OF LEMMA F.4

Proof. Since $\langle g^{hybrid,k}, y^{hybrid,k}\rangle$ is a valid core bank account mechanism for $\hat{F}_{(1,T)}$, the bank account mechanism $B(g^{hybrid,k}, y^{hybrid,k}; \hat{F}_{(1,T)})$ is stage-IC, $\delta_{(1,T)}$-BI, BU and ex-post IR for $F_{(1,T)}$ with $\delta_t \le \Delta_k a_t$ for all t by Lemma 4.2. As for $t \in \tilde{E}_k^\omega$, the stage mechanism $\langle x_t, \hat{p}_t\rangle$ is indeed static such that $x_t(\xi, b_t) = 1$ and $\hat{p}_t(\xi, b_t) = 0$ for all $\xi \ge 0$ and b_t. Therefore, $\delta_t = 0$ for $t \in E_k^\omega$. To bound the total magnitude of misreport, we apply Claim D.1 to obtain an upper bound $O(\ell_k\sqrt{\frac{\Delta_k}{\lambda_k}})$.

By Lemma 4.5, we have $\bar{B}(g^{hybrid,k}, y^{hybrid,k}; \hat{F}_{(1,T)})$ is stage-IC, $\delta_{(1,T)}$-BI, BU and ex-post IR for $F_{(1,T)}$. Moreover, $\bar{B}(g^{hybrid,k}, y^{hybrid,k}; \hat{F}_{(1,T)})$ is $\eta_{(1,T)}$-DIC with $\eta_t = \sqrt{\frac{2a_t}{\lambda_k}\cdot\sum_{t'=t+1}^{T}\gamma^{t'-t}\delta_{t'}}$ for $F_{(1,T)}$. Therefore, for $t \in E_k \setminus E_k^\omega$, we have

$$\begin{aligned}\sum_{t\in E_k\setminus E_k^\omega}\frac{\eta_t}{a_t} &= \sqrt{\frac{2\Delta_k}{\lambda_k}}\sum_{t\in E_k\setminus E_k^\omega}\sqrt{\frac{1}{a_t}}\sqrt{\sum_{t'=t+1}^{T}\gamma^{t'-t}a_{t'}}\\ &\le \sqrt{\frac{2\Delta_k}{\lambda_k}}\cdot\sqrt{\sum_{t\in E_k\setminus E_k^\omega}\frac{1}{a_t}}\sqrt{\sum_{t\in E_k\setminus E_k^\omega}\sum_{t'=t+1}^{T}\gamma^{t'-t}a_{t'}}\\ &\le \sqrt{\frac{2\Delta_k}{\lambda_k}}\cdot\sqrt{\ell_k^{1-2\omega}}\cdot\sqrt{c\ell_k} = \ell_k^{1-\omega}\sqrt{\frac{2c\Delta_k}{\lambda_k}}\end{aligned} \tag{37}$$

for some constant c, where the first inequality follows the Cauchy-Schwarz inequality and we use the fact that $\sum_{t\in E_k\setminus E_k^\omega}\frac{1}{a_t} \le \ell_k^{1-2\omega}$ in which the maximum is obtained when $a_t = \ell_k^\omega$ for all $t \in E_k\setminus E_k^\omega$. When $\ell_k^{1-\omega}\sqrt{\frac{\Delta_k}{\lambda_k}} \ll \ell_k^{\frac{1}{2}(1+\omega)}$, the misreport is relatively small compared to $\sum_{t\in\tilde{E}_k^\omega}\mathbb{E}_{v_t\sim\hat{F}_t}[v_t]$. Therefore, Lemma F.2 still holds even the buyer misreports. Therefore, with probability $1 - O(\frac{1}{T})$, the buyer is indeed facing a static mechanism even for stage $t \in E_k^\omega \setminus \tilde{E}_k^\omega$. Finally, for $t \in E_k^\omega$ with $\mathrm{next}(t) - t \ge 3\log_{\frac{1}{\gamma}} T$, we have

$$\begin{aligned}\eta_t &= \sqrt{\frac{2a_t}{\lambda_k}\cdot\sum_{t'=t+1}^{T}\gamma^{t'-t}\delta_{t'}} \le \sqrt{\frac{2a_t}{\lambda_k}\cdot\sum_{t'=t+3\log_{\frac{1}{\gamma}}T}^{T}\gamma^{t'-t}\delta_{t'}}\\ &\le \sqrt{\frac{2a_t\Delta_k}{\lambda_k}\cdot\sum_{t'=t+3\log_{\frac{1}{\gamma}}T}^{T}\frac{a_{t'}}{T^3}} \le \frac{1}{T}\sqrt{\frac{2ca_t\Delta_k}{\lambda_k}} \le \frac{a_t}{T}\sqrt{\frac{2c\Delta_k}{\lambda_k}}.\end{aligned}$$

where the last inequality is due to $a_t \ge 1$. Therefore, we have

$$\sum_{t\in E_k^\omega,\mathrm{next}(t)-t\ge 3\log_{\frac{1}{\gamma}}T}\eta_t \le \sqrt{\frac{2c\Delta_k}{\lambda_k}}.$$

For $t \in E_k^\omega$ but $\mathrm{next}(t) - t < 3\log_{\frac{1}{\gamma}} T$, we simply use the fact that $\eta_t \le a_t$, which implies $\eta_t/a_t \le 1$. Moreover, the number of such t is at most $|E_k \setminus E_k^\omega|\cdot 3\log_{\frac{1}{\gamma}} T = O(\ell_k^{1-\omega}\log\ell_k)$. We finish the proof of bounding the number of lies by combining all three cases and applying Lemma A.1, and thus, we have

$$|L_k| = O\left(\ell_k^{1-\omega}(\sqrt{\frac{\Delta_k}{\lambda_k}} + \log\ell_k) + \log\ell_k\right) = O\left(\ell_k^{1-\omega}\log\ell_k\right)$$

because we assume $\frac{\Delta_k}{\lambda_k} < 1$. □

F.3.5. PROOF OF LEMMA F.5

Proof. We prove by a backward induction from $t = T$ to $t = 0$. It is true for the base case when $t = T$ since $\phi_T(\xi + \delta; \hat{F}_{(1,T)}) = -\xi - \delta = \phi_{T+1}(\xi; \hat{F}_{(1,T)}) - \delta$ by the boundary condition.

Suppose it is true for all $t' \geq t$. For $t-1$ and $\xi \geq 0$, if $t \in E_k \setminus E_k^{\omega}$, we can apply a similar argument as in the proof of Lemma 3.6. For $t \in \tilde{E}_k^{\omega}$, we can again apply a similar argument as in the proof of Lemma 3.6 since the stage mechanisms are the same for all $\xi \geq 0$.

For $t \in E_k^{\omega} \setminus \tilde{E}_k^{\omega}$, Lemma F.2 holds when we shift the initial state up by an additional $O(\sqrt{\frac{\Delta_k}{\lambda_k}} \cdot \ell_k)$. Therefore, with probability $1 - O(\frac{1}{T})$, the buyer is facing a static mechanism, and thus, we can apply a similar argument as in the proof of Lemma 3.6. □