A Robust Non-Clairvoyant Dynamic Mechanism for Contextual Auctions

Anonymous Author(s)
Affiliation
Address
email

Abstract

Dynamic mechanisms offer powerful techniques to improve on both revenue and efficiency by linking sequential auctions using state information, but these techniques rely on exact distributional information of the buyers' valuations (present and future), which limits their use in learning settings. In this paper, we consider the problem of contextual auctions where the seller gradually learns a model of the buyer's valuation as a function of the context (e.g., item features) and seeks a pricing policy that optimizes revenue. Building on the concept of a bank account mechanism—a special class of dynamic mechanisms that is known to be revenue-optimal—we develop a non-clairvoyant dynamic mechanism that is robust to both estimation errors in the buyer's value distribution and strategic behavior on the part of the buyer. We then tailor its structure to achieve a policy with provably low regret against a constant approximation of the optimal dynamic mechanism in contextual auctions. Our result substantially improves on previous results that only provide revenue guarantees against static benchmarks.

1 Introduction

As a fundamental problem in mechanism design, pricing in repeated auctions has been extensively studied in recent years. This is partly motivated by the popularity of selling online ads via auctions, an industry totalling annual revenue of hundreds of billions of dollars. Repeated auctions open up the possibility of linking auctions across time using state information in order to enhance revenue or welfare, but this introduces several challenges. To guarantee optimal outcomes, the process must take into account the bidders' incentives to possibly manipulate each individual auction as well as the auction state across time. In practice, the seller must also rely on approximate models of the buyers' preferences to effectively set auction parameters like reserve prices. These aspects of the problem have so far been explored in two separate strands of the literature on repeated auctions, where items arrive online and the allocation and payment decisions must be made as soon as an item arrives.

One strand, known as *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 state based on the buyer's historical bids [Thomas and Worrall, 1990, Bergemann and Välimäki, 2010, Ashlagi et al., 2016, Mirrokni et al., 2016a,b]. However, this *clairvoyant* framework relies on the seller having an accurate forecast of the buyer's valuation distributions in future auctions. To address this concern, Mirrokni et al. [2018] propose *non-clairvoyant* dynamic mechanisms, which do not rely on any information about the future (but do rely on an accurate forecast of the present). They show that a non-clairvoyant dynamic mechanism can achieve a constant approximation to the revenue of the optimal clairvoyant mechanism. The other strand of literature, known as *robust price learning*, focuses on a setting where the buyer's value distributions across stages are parameterized by some common private factors that are unknown to the seller, and designs robust policies to learn from the

Submitted to 33rd Conference on Neural Information Processing Systems (NeurIPS 2019). Do not distribute.

buyer's bids and set prices with good revenue performance [Amin et al., 2013, 2014, Medina and Mohri, 2014, Golrezaei et al., 2018]. Although these results also take into account strategic buyer behavior, they only provide guarantees against the revenue-optimal *static* benchmark, which does not take advantage of auction state across time and whose revenue can be arbitrarily smaller than the optimal dynamic benchmark [Papadimitriou et al., 2016].

In this work, we consider a scenario in which the designer can only make use of an estimate of the buyer's value distribution in the present auction stage, which connects dynamic mechanism design with the problem of learning. Designing dynamic auctions in this setting is challenging for several reasons. When the seller's estimate of the distribution is not perfectly aligned with the buyer's true distribution, it is impossible for the seller to offer a dynamic mechanism that is exactly incentive-compatible and also makes use of the prior on values. Furthermore, unlike static mechanisms in which the auction for each item is independent of the buyer's past reports, in a dynamic mechanism a buyer's misreport can potentially affect auctions for all future items. We overcome these obstacles and provide a *robust non-clairvoyant* dynamic mechanism such that the extent of the buyers' misreports and the revenue loss can be related to and bounded by the estimation error. We then apply our robust dynamic mechanism to the concrete problem of contextual auctions, where a buyer's valuation depends on the context that describes the item, but the relationship between the buyer's valuation and the context is unknown to the seller and must be estimated across auctions. The seller's task is to design a policy which adapts the auction mechanism based on the buyer's historical bids, with the objective of maximizing revenue. Previous results give no-regret policies against the optimal *static* mechanism [Amin et al., 2014, Golrezaei et al., 2018], but as mentioned it is known that the revenue gap between optimal static and dynamic mechanisms can be arbitrarily large [Papadimitriou et al., 2016]. We tailor the structure of our robust non-clairvoyant dynamic mechanism to a learning environment, leading to a no-regret policy against the strong benchmark of a *constant approximation* of the *optimal clairvoyant dynamic mechanism*.

Related Work

We briefly discuss research in dynamic mechanism design that is closely related to the present work. For a comprehensive review of the literature readers are encouraged to refer to [Bergemann and Said, 2011]. Our work builds upon the framework of bank account mechanisms developed by Mirrokni et al. [2016a,b, 2018]. Based on the bank account mechanism, Mirrokni et al. [2018] design a non-clairvoyant mechanism achieving $1/3$ of the revenue of a clairvoyant mechanism which can make use of present and future information on the distributions of item values. However, their mechanism relies on exact distributional information, which makes it unsuitable in a learning environment where value distributions are estimated. Our robust dynamic mechanism addresses this limitation.

Our work is closely related to dynamic pricing with learning; see [den Boer, 2015] for a recent survey. The study of robust price learning with strategic buyers was initiated by Amin et al. [2013] and Medina and Mohri [2014]. They design no-regret policies in a non-contextual environment where the seller repeatedly interacts with a single buyer through posted price auctions, where the buyer is less patient than the seller. Amin et al. [2013] show that no learning algorithm can achieve sublinear revenue loss if the buyer is as patient as the seller. For learning in contextual auctions, Amin et al. [2014] develop a no-regret policy in a setting without market noise. Recently, Golrezaei et al. [2018] enrich the model by incorporating market noise and design a no-regret policy for cases where the market noise is known exactly or adversarially selected from a set of distributions. All these results are no-regret against the optimal *static* mechanism as a benchmark, whereas our policy is no-regret against a constant-factor approximation of the optimal *dynamic* mechanism which has all distributional information available in advance.

2 Preliminaries

In a *dynamic* auction a seller (he) sells a stream of T items that arrive online, based on bids placed by strategic buyers. An item must be sold when it arrives. For the sake of simplicity we will focus on the case of a single buyer (she) throughout this paper.[1] At the beginning of stage t a new item arrives and the buyer's valuation $v_t \in [0, a_t]$ for the item is drawn independently from a distribution

[1] Our results can be extended to multi-buyer settings by using the techniques from Cai et al. [2012] and Mirrokni et al. [2018].

F_t with density f_t. The distributions are not necessarily identical across stages. We assume that f_t is continuous and upper bounded by c_f/a_t where c_f is a constant. The domain bounds a_t are known to the seller and may vary across stages to reflect the fact that item valuations may have different scales.[2] As a special case of this framework, 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., 2018]. 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 $v_t = a_t(\langle \boldsymbol{\sigma}, \boldsymbol{\zeta}_t \rangle + \varepsilon_t)$, where ε_t is a noise term with cumulative distribution M_t. The distribution M_t and the feature vector $\boldsymbol{\zeta}_t$ are observed by the seller but the buyer's preference vector $\boldsymbol{\sigma}$ remains private. We make the following technical assumption on the sequence of a_t:

Assumption 1. *For all t, $\sum_{t' \leq t} a_{t'} \leq c_a \cdot t$ where c_a is a constant.*

Assumption 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.

Once the buyer learns her valuation v_t at stage t, she then submits a bid $b_t \in [0, a_t]$ to the seller who then decides whether to allocate the item (perhaps stochastically) and what payment to charge. We write V^t to denote the set of all possible sequences $(b_1, \ldots, b_t)$ of buyer bids for the first t stages, and similarly we write $(\Delta V)^t$ to denote the set of all possible independent distributions over the sequence of first t bids. The seller's distributional beliefs over the buyer's values across stages are denoted as $\hat{F}_{(1,T)} = (\hat{F}_1, \hat{F}_2, \ldots, \hat{F}_T)$. Throughout the paper we will use the notation $\hat{F}_{(t',t'')}$ to represent $(\hat{F}_{t'}, \ldots, \hat{F}_{t''})$, and similarly for $F_{(t',t'')}$, $v_{(t',t'')}$, and $b_{(t',t'')}$. A dynamic mechanism is represented by sequences $(x_1, \ldots, x_T)$ and $(p_1, \ldots, p_T)$ where x_t and p_t denote the allocation rule and the payment rule at stage t, respectively. We refer to $\langle x_t, p_t \rangle$ as the *stage mechanism* at stage t.

Non-Clairvoyant Dynamic Mechanism. In a non-clairvoyant environment, the seller obtains an estimated distribution $\hat{F}_t$ only at stage t and not before, so the mechanism at stage t can only depend on $\hat{F}_{(1,t)}$. The allocation function x_t maps the history of bids $b_{(1,t)}$ and distribution $\hat{F}_{(1,t)}$ to an allocation probability, $x_t : V^t \times (\Delta V)^t \to [0,1]$. The payment function p_t maps the history of bids $b_{(1,t)}$ and the distribution $\hat{F}_{(1,t)}$ to a real-valued payment, $p_t : V^t \times (\Delta V)^t \to \mathbb{R}$. In line with the literature, we assume the buyer has a quasi-linear utility such that the buyer's utility from bidding b_t at stage t is $u_t\big(v_t; b_{(1,t)}; \hat{F}_{(1,t)}\big) = v_t \cdot x_t\big(b_{(1,t)}; \hat{F}_{(1,t)}\big) - p_t\big(b_{(1,t)}; \hat{F}_{(1,t)}\big)$. In the contextual auction setting the seller maintains a model $\hat{\boldsymbol{\sigma}}_t$ for the buyer's preference vector estimated from prior bidding behavior, and combines with a_t, $\boldsymbol{\zeta}_t$, and noise model M_t, which can only be observed at the beginning of stage t and not before, to compute $\hat{F}_t$.

Utility-Maximizing Buyer. We assume that the buyer knows the true distributions $F_{(1,T)}$ in advance so that she can reason about how the mechanism will evolve over time and compute a bidding strategy that maximizes her utility. Specifically, we consider a buyer who aims to maximize her time discounted utility $\sum_{t'=t}^{T} \gamma^{t'-t} \cdot \mathbb{E}[u_t]$ at stage t where $\gamma \in [0,1)$ is the discounting factor and the expectation is taken with respect to $F_{(1,T)}$. We note that it is impossible to obtain a no-regret policy when the buyer is as patient as the seller (the case of $\gamma = 1$) [Amin et al., 2013].

Incentive Constraints. In a dynamic environment, the buyer's best response at stage t depends on her strategy in the future stages. When the seller has perfect distributional information, the classic notion of dynamic incentive-compatibility (DIC) requires that the buyer is incentivized to report truthfully assuming that she plays optimally in the future [Mirrokni et al., 2018].[3] When the seller only has approximate distributional information this is no longer possible to achieve, so we introduce the notion of $\eta_{(1,T)}$-approximate DIC, which requires that the buyer's bid deviate from the truth by at most η_t at stage t, assuming the buyer plays optimally in the future (note that optimally now no longer means truthfully). Formally, at each stage t, there exists $\hat{b}_t \in [v_t - \eta_t, v_t + \eta_t]$ such that

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

[2]For instance, in a dynamic auction for display advertising, the value of a video ad may be orders of magnitude larger than the value of a text ad.

[3]Interested readers can refer to [Mirrokni et al., 2018] for discussions on the choice of DIC notions.

for all v_t, $b_{(1,t-1)}$, $F_{(t+1,T)}$, and $\hat{F}_{(t+1,T)}$, where $U_t(b_{(1,t)}; F_{(1,T)}; \hat{F}_{(1,T)})$ is the continuation utility that the buyer obtains in the future: $U_T\big(b_{(1,T)}; F_{(1,T)}; \hat{F}_{(1,T)}\big) = 0$, and for $t < T$ $U_t\big(b_{(1,t)}; F_{(1,T)}; \hat{F}_{(1,T)}\big)$ is defined as

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

Participation Constraints. We assume that the buyer weighs realized past utilities equally. Therefore, ex-post individual rationality requires that for all $\hat{F}_{(1,T)}$ and for all $v_{(1,T)}$,

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

For convenience, we will use the phrase "for $F_{(1,T)}$" to indicate the environment where the buyer's true distribution is $F_{(1,T)}$. For example, when we say that a mechanism is $\eta_{(1,T)}$-DIC for $F_{(1,T)}$ we mean that it is $\eta_{(1,T)}$-DIC when the buyer's true distribution is $F_{(1,T)}$.

No-Regret Policy. 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 $a_{(1,t)}$, $\boldsymbol{\zeta}_{(1,t)}$, $M_{(1,t)}$, and $b_{(1,t-1)}$, while the dynamic mechanism policy computes the stage mechanism $\langle x_t, p_t \rangle$ at stage t using $\hat{F}_{(1,t)}$ and $b_{(1,t-1)}$. Let $\mathsf{Rev}(\pi; F_{(1,T)})$ and $\mathsf{Rev}(B; F_{(1,T)})$ 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 a c-approximation of the dynamic benchmark is defined as $\mathsf{Regret}^{\pi}(F_{(1,T)}) = c \cdot \mathsf{Rev}\big(B^*(F_{(1,T)}); F_{(1,T)}\big) - \mathsf{Rev}(\pi; F_{(1,T)})$. Our objective is to design a policy with sublinear regret.[4]

3 Robust Non-clairvoyant Mechanism

The literature on dynamic mechanism design relies on the strong assumption that the seller has perfect distributional information at each stage, $\hat{F}_{(1,T)} = F_{(1,T)}$ [Ashlagi et al., 2016, Mirrokni et al., 2016b,a, 2018]. However, in a learning setting like that of contextual auctions, the seller can only obtain a sequence of estimated distributions by estimating $\boldsymbol{\sigma}$. In this section, we design a non-clairvoyant mechanism that is robust to misspecifications in the value distribution in the sense that the buyer is incentivized to place a bid within known bounds from its value, which ultimately allows us to relate the mechanism revenue under the estimated and true value distributions. The misspecifications handled by the mechanism are captured by the following assumption.

Assumption 2. *For all t, let $\hat{v}_t$ be the random variable that is drawn from $\hat{F}_t$. We assume that the buyer's true valuation $v_t = \hat{v}_t + a_t \cdot \epsilon_t$ with $\epsilon_t \in [-\Delta, \Delta]$. Here $\hat{v}_t$ and ϵ_t are not necessarily independent and arbitrary correlation between $\hat{v}_t$ and ϵ_t is allowed.*

3.1 The Mechanism

Building on the $\frac{1}{3}$-approximation non-clairvoyant mechanism from Mirrokni et al. [2018], we design our robust non-clairvoyant mechanism by mixing their mechanism with a random posted-price auction. The mechanism is an instance of a *bank account mechanism* where the state information is captured by a single scalar bal_t.

Mechanism 1. *The robust non-clairvoyant mechanism $B(\hat{F}_{(1,T)}, \lambda)$ consists of a mixture of four mechanisms: the give-for-free mechanism, the posted-price auction with extra fee, the Myerson's auction, and the random posted-price auction. The stage mechanism at stage t is parameterized by a non-negative balance bal_t. When the buyer submits a bid b_t:*

[4] Note that sublinear revenue loss is only meaningful if the available revenue to extract is itself at least linear, which is the case when $\sum_{t=1}^{T} a_t = \Omega(T)$ since the revenue obtained by the optimal dynamic mechanism is $\Omega(\sum_t a_t)$ in our setting. In fact, a *static* mechanism can already achieve $\Omega(\sum_t a_t)$ revenue by offering a posted price pa_t with $p = 1/(2c_f)$ at stage t which induces revenue at least $pa_t(1 - pc_f) = a_t/(4c_f)$ from stage t.

Give-for-free Mechanism. *Allocate the item no matter what the buyer's bid is and increase the balance by the buyer's bid:* $x_t^G = 1$, $p_t^G = 0$, *and* $\mathrm{bal}_{t+1}^G = \mathrm{bal}_t + b_t$

Posted-price Auction with Extra Fee. *Let* $\mathsf{fee}_t(\mathrm{bal}_t; \hat{F}_t) = \min(3\mathrm{bal}_t, \mathbb{E}_{v_t \sim \hat{F}_t}[v_t])$ *and* $r_t(\mathrm{bal}_t)$ *be the posted-price such that* $\mathbb{E}_{v_t \sim \hat{F}_t}\left[\left(v_t - r_t(\mathrm{bal}_t)\right)^+\right] = \mathsf{fee}_t(\mathrm{bal}_t; \hat{F}_t)$. *The mechanism charges the buyer* $\mathsf{fee}_t(\mathrm{bal}_t; \hat{F}_t)$ *before the buyer learns her valuation and then runs a posted-price auction with price* $r_t(\mathrm{bal}_t)$*:* $x_t^P = \mathbf{1}\{b_t \geq r_t(\mathrm{bal}_t)\}$ *and* $p_t^P = \mathsf{fee}_t(\mathrm{bal}_t; \hat{F}_t) + r_t(\mathrm{bal}_t) \cdot \mathbf{1}\{b_t \geq r_t(\mathrm{bal}_t)\}$, *and decrease the balance by* $\mathsf{fee}_t(\mathrm{bal}_t; \hat{F}_t)$*:* $\mathrm{bal}_{t+1}^P = \mathrm{bal}_t - \mathsf{fee}_t(\mathrm{bal}_t; \hat{F}_t)$.

Myerson's Auction. *Let* $r_t^*(\hat{F}_t)$ *be Myerson's optimal reserve price, i.e.,* $r_t^*(\hat{F}_t) = \arg\max_r r \cdot \left(1 - \hat{F}_t(r)\right)$ *and run a posted-price auction with price* $r_t^*(\hat{F}_t)$ *without changing the balance:* $x_t^M = \mathbf{1}\{b_t \geq r_t^*(\hat{F}_t)\}$, $p_t^M = r_t^*(\hat{F}_t) \cdot \mathbf{1}\{b_t \geq r_t^*(\hat{F}_t)\}$, *and* $\mathrm{bal}_{t+1}^M = \mathrm{bal}_t$.

Random Posted-price Auction. *Let* $\hat{r}_t$ *be random reserve price drawn from* $[0, a_t]$ *uniformly and run a posted-price auction with price* $\hat{r}_t$ *without changing the balance:* $x_t^R = \mathbf{1}\{b_t \geq \hat{r}_t\}$, $p_t^R = \hat{r}_t \cdot \mathbf{1}\{b_t \geq \hat{r}_t\}$, *and* $\mathrm{bal}_{t+1}^R = \mathrm{bal}_t$.

The robust non-clairvoyant mechanism at stage t is: $x_t = \lambda \cdot x_t^R + \frac{1-\lambda}{3}\left[x_t^G + x_t^P + x_t^M\right]$, $p_t = \lambda \cdot p_t^R + \frac{1-\lambda}{3}\left[p_t^G + p_t^P + p_t^M\right]$, *and* $\mathrm{bal}_t = \lambda \cdot \mathrm{bal}_t^R + \frac{1-\lambda}{3}\left[\mathrm{bal}_t^G + \mathrm{bal}_t^P + \mathrm{bal}_t^M\right]$.

The following central result gives a guarantee on the revenue performance of our robust non-clairvoyant mechanism against a utility-maximizing buyer subject to an estimation error Δ.

Theorem 3.1. $\mathsf{Rev}\left(B(\hat{F}_{(1,T)}, \lambda), F_{(1,T)}\right) \geq \frac{1}{3}\mathsf{Rev}\left(B^*(F_{(1,T)}), F_{(1,T)}\right) - O\left(\lambda T + \sqrt{\frac{\Delta}{\lambda}} T\right)$.

At the optimal choice of $\lambda = \Delta^{\frac{1}{3}}$ the revenue loss is $O\left(\Delta^{\frac{1}{3}} T\right)$. The remainder of this section is devoted to proving Theorem 3.1.

3.2 Analysis

We start by describing the incentive properties that $B(\hat{F}_{(1,T)}, \lambda)$ satisfies for $\hat{F}_{(1,T)}$. First notice that all four base mechanisms are variants of posted-price auctions, and therefore, all of them are stage-IC:

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

In particular, all mechanisms except the posted-price auction with extra fee are stage-IR:

$$\forall v_t,\ v_t \cdot x_t(v_t) - p_t(v_t) \geq 0 \tag{stage-IR}$$

We emphasize that the posted-price auction with extra fee is different from a classic posted-price auction: the posted-price auction with extra fee will charge the buyer an extra payment $\mathsf{fee}_t(\mathrm{bal}_t; \hat{F}_t)$ no matter what the buyer's bid is, and therefore, it is not stage-IR. Moreover, each stage mechanism is balance-independent (BI) with respect to the estimated distribution $\hat{F}_t$: there exists a constant c_t,

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

In particular, the give-for-free mechanism, the Myerson's auction, and the random posted-price auction are static and independent of the balance; as for the posted-price auction with extra fee, it ensures that the buyer's expected utility is always 0 for all $\mathrm{bal}_t \geq 0$ under $\hat{F}_t$.

The combination of stage-IC and BI implies that the mechanism is DIC: since the mechanism promises the buyer that all future stage mechanisms are BI, the buyer can infer that her action at the current stage does not impact her expected utility in the future. Moreover, notice that the non-negative balance bal always lower-bounds the buyer's cumulative utility, and therefore, $B(\hat{F}_{(1,T)}, \lambda)$ is ex-post IR under the estimated distributions $\hat{F}_{(1,T)}$.

Proposition 3.1. $B(\hat{F}_{(1,T)}, \lambda)$ *is stage-IC, BI, DIC, and ex-post IR for* $\hat{F}_{(1,T)}$.

We next turn to the mechanism's properties under the true distributions $F_{(1,T)}$.

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

We first bound the revenue loss due to the mismatch between $\hat{F}_{(1,T)}$ and $F_{(1,T)}$. Observe that one can interpret the estimation error under Assumption 2 as the buyer's misreport: when the buyer reports truthfully under $F_{(1,T)}$ this is equivalent to the case in which the buyer misreports by a magnitude at most $a_t \cdot \Delta$ under $\hat{F}_{(1,T)}$. We develop a program for computing the revenue of our mechanism even when the buyer misreports. For a non-clairvoyant mechanism $B(\hat{F}_{(1,T)}, \lambda)$, we consider a program $\psi_t(\text{bal}, \hat{F}_{(1,T)}; F_{(1,T)})$ to keep track on the revenue of implementing $B(\hat{F}_{(1,T)}, \lambda)$ when the buyer's true distribution is $F_{(1,T)}$. We define $\psi_T(\text{bal}) = 0$ and for $t < T$,

$$\psi_{t-1}(\text{bal}, \hat{F}_{(1,T)}; F_{(1,T)}) = \mathbb{E}_{v_t \sim F_t} \left[\frac{1}{3} \mathsf{fee}_t(\text{bal}; \hat{F}_t) + \frac{1}{3} r_t^*(\hat{F}_t) \cdot \mathbf{1}\{v_t' \geq r_t^*(\hat{F}_t)\} \right. \\ \left. + \psi_t \left(\text{bal} + \frac{1}{3} v_t' - \frac{1}{3} \mathsf{fee}_t(\text{bal}; \hat{F}_t), \hat{F}_{(1,T)}; F_{(1,T)} \right) \right] \quad (1)$$

where v_t' is the buyer's reported bid that maximizes her continuation utility when her true value is v_t. Note that we omit the revenue $r_t(\text{bal}_t)$ obtained from the posted-price auction with extra fee and the revenue from the random posted-price auction.

Proposition 3.2. $\mathsf{Rev}\left(B(\hat{F}_{(1,T)}, \lambda); F_{(1,T)} \right) \geq (1 - \lambda) \cdot \psi_0(0, \hat{F}_{(1,T)}; F_{(1,T)})$.

According to the revenue analysis in [Mirrokni et al., 2018], we can still obtain $\frac{1}{3}$-approximation of the optimal revenue even when the revenue $r_t(\text{bal}_t)$ is omitted.

Lemma 3.1. *[Mirrokni et al., 2018]* $\psi_0(0, F_{(1,T)}; F_{(1,T)}) \geq \frac{1}{3} \cdot \mathsf{Rev}\left(B^*(F_{(1,T)}), F_{(1,T)}\right)$.

The following lemma establishes a connection between the change of the balance and the change of the revenue, when the seller's distributional information is perfect.

Lemma 3.2. *For all* $0 \leq t \leq T$ *and* $\delta \geq 0$,

$$\psi_t(\text{bal} + \delta, F_{(1,T)}; F_{(1,T)}) - \delta \leq \psi_t(\text{bal}, F_{(1,T)}; F_{(1,T)}) \leq \psi_t(\text{bal} + \delta, F_{(1,T)}; F_{(1,T)}).$$

Applying Lemma 3.2, we can bound the revenue loss due to the mismatch between $F_{(1,T)}$ and $\hat{F}_{(1,T)}$.

Lemma 3.3. $\psi_0(0, \hat{F}_{(1,T)}; \hat{F}_{(1,T)}) \geq \psi_0(0, F_{(1,T)}; F_{(1,T)}) - O(\Delta T)$.

3.2.2 The Buyer's Misreport

Note that in a single-buyer environment, the properties stage-IC and ex-post IR do not depend on the underlying distributions, and therefore, $B(\hat{F}_{(1,T)}, \lambda)$ is also stage-IC and ex-post IR for $F_{(1,T)}$. However, $B(\hat{F}_{(1,T)}, \lambda)$ is no longer BI for $F_{(1,T)}$, which is the key property to ensure DIC. To circumvent this difficulty, we generalize the definition of BI to *approximate balance-independence*.

Definition 3.1. *A dynamic mechanism is $\beta_{(1,T)}$-BI for $F_{(1,T)}$ if $\forall t$, there exists a constant c_t:*

$$\forall \text{bal} \geq 0, \mathbb{E}_{v_t \sim F_t}[v_t \cdot x_t(\text{bal}, v_t) - p_t(\text{bal}, v_t)] \in [c_t - \frac{\beta_t}{2}, c_t + \frac{\beta_t}{2}] \qquad (\beta_{(1,T)}\text{-BI})$$

Since with the same stage mechanism, the difference between the expected utility under $\hat{F}_t$ and F_t is at most Δa_t (Corollary A.1), $B(\hat{F}_{(1,T)}, \lambda)$ is $\beta_{(1,T)}$-BI with $\beta_t = 2\Delta a_t$.

Proposition 3.3. *$B(\hat{F}_{(1,T)}, \lambda)$ is stage-IC, $\beta_{(1,T)}$-BI with $\beta_t = 2\Delta a_t$, and ex-post IR for $F_{(1,T)}$.*

For a dynamic mechanism satisfying $\beta_{(1,T)}$-BI for $F_{(1,T)}$, the range of the expected utility is β_t in 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 at most β_t if she reports truthfully at stage t. Combining this with the fact that the stage mechanisms are stage-IC, we have

Lemma 3.4. *For a dynamic mechanism that is stage-IC and $\beta_{(1,T)}$-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 $\sum_{t'=t+1}^{T} \gamma^{t'-t} \cdot \beta_{t'}$.*

As a result, once the mechanism posts a risk for misreporting, we are able to bound the magnitude of the buyer's misreport. This is the purpose of mixing in the random posted-price mechanism at each stage t: it can be shown that a misreport with magnitude m_t will cause the buyer a utility loss $\lambda \cdot \frac{m_t^2}{2a_t}$.

Lemma 3.5. *$B(\hat{F}_{(1,T)}, \lambda)$ is $\eta_{(1,T)}$-DIC with $\eta_t = \sqrt{\frac{2a_t}{\lambda} \cdot \sum_{t'=t+1}^{T} \gamma^{t'-t} \beta_{t'}}$.*

Applying Lemma 3.2, we can show that $B(\hat{F}_{(1,T)}, \lambda)$ is robust against the buyer's misreport.

Lemma 3.6. $\psi_0(0, \hat{F}_{(1,T)}; F_{(1,T)}) \geq \psi_0(0, \hat{F}_{(1,T)}; \hat{F}_{(1,T)}) - O\left(\sqrt{\frac{\Delta}{\lambda}} T\right)$.

Finally, combining Proposition 3.2, Lemma 3.3 and Lemma 3.6, completes the proof of Theorem 3.1.

4 No-Regret Policy in Contextual Auctions

4.1 Learning Policy

Our learning policy is adapted from the contextual robust pricing policy proposed in [Golrezaei et al., 2018]. 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 $(\ell_1 = 1, \ell_2, \cdots, \ell_K, \ell_{K+1} = T + 1)$, in which $\ell_k = 2^{k-1}$. The k-th phase spans between the ℓ_k-th stage and the $(\ell_{k+1} - 1)$-th stage, and therefore, the length of phase k is exactly ℓ_k. Note that the partition can be implemented even when T is not known in advance. We use $E_k = \{\ell_k, \cdots, \ell_{k+1} - 1\}$ to refer to the stages in the k-th phase.

At the beginning of the k-th phase, we update the estimation of the buyer's preference vector $\boldsymbol{\sigma}$ using the buyer's bids from the $(k-1)$-th phase, denoted by $\hat{\boldsymbol{\sigma}}_k$. To estimate $\hat{\boldsymbol{\sigma}}_k$, we sample w_t uniformly from $[0, 1]$ for $t \in \hat{E}_{k-1}$, where $\hat{E}_{k-1} = \{t \in E_{k-1} \mid \ell_k - t > c \log \ell_k\}$ for some constant c. In other words, we will only use the information from the stages that are at least $c \log \ell_k$ 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 $\mathcal{L}_{k-1}(\boldsymbol{\sigma}) = -\sum_{t \in \hat{E}_{k-1}} \Big[\mathbf{1}\{b_t \geq a_t \cdot w_t\} \log\left(1 - M_t(w_t - \langle \boldsymbol{\sigma}, \boldsymbol{\zeta}_t \rangle)\right) + \mathbf{1}\{b_t < a_t \cdot w_t\} \log\left(M_t(w_t - \langle \boldsymbol{\sigma}, \boldsymbol{\zeta}_t \rangle)\right)\Big]$. 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. 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 $B(\hat{F}_{(1,T)}, \lambda_{(1,K)})$ be a mechanism generalized from $B(\hat{F}_{(1,T)}, \lambda)$ such that for $t \in E_k$, $B(\hat{F}_{(1,T)}, \lambda_{(1,K)})$ offers the random posted-price auction with probability λ_k instead of λ.

Lemma 4.1. *In $B(\hat{F}_{(1,T)}, \lambda_{(1,K)})$, the additional misreport at stage $t \in \hat{E}_k$ is $O(\frac{1}{\sqrt{\lambda_k} \cdot \ell_k^2})$. Moreover, $|L_k| = O\left(\log \ell_k + \sum_{t \in \hat{E}_k} \frac{\eta_t}{a_t}\right)$ 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 4.2 (Proposition 7.1 [Golrezaei et al., 2018])**.** *With probability $1 - \frac{1}{\ell_k}$, the estimation error for phase k is $\Delta_k \equiv \|\hat{\boldsymbol{\sigma}}_k - \boldsymbol{\sigma}\| = O\left(d \cdot \frac{|L_{k-1}|}{\ell_{k-1}} + \sqrt{\frac{\log(\ell_{k-1} \cdot d)}{\ell_{k-1}}}\right)$.*

4.2 Dynamic Mechanism Policy

We develop a hybrid non-clairvoyant mechanism to reduce the number of lies by reducing the magnitude of misreports. To do so, observe that the buyer has no incentive to misreport in order to affect future stage mechanisms when the latter are static. However, as previously mentioned, offering a purely static mechanism may forego a large amount of revenue [Papadimitriou et al., 2016]. Motivated by this insight, our hybrid mechanism contains both dynamic stages dependent on the history and static stages independent of the history. We adapt $B(\hat{F}_{(1,T)}, \lambda_{(1,K)})$ to obtain a hybrid non-clairvoyant mechanism $B^{hybrid}(\hat{F}_{(1,T)}, \lambda_{(1,K)}, \omega, \tau)$, which is parameterized by $\omega \in (0,1)$ and a function $\tau : \mathbb{Z}_+ \to \mathbb{R}_+$ that maps the phase number to a real number. The stage mechanism at stage t is parameterized by a_t, two balances bal_t and sbal_t, and an additional parameter sw_t.

We provide a high level description of our mechanism in this section while a detailed description is deferred to Appendix. Let $E_k^\omega = \{t \in E_k \mid a_t < \ell_k^\omega\}$. Intuitively, the hybrid non-clairvoyant mechanism runs different stage mechanisms conditioned on whether $t \in E_k^\omega$ or not: the stage mechanism is dynamic for $t \notin E_k^\omega$ and the stage mechanism is static for $t \in E_k^\omega$ with high probability.

More precisely, for $t \notin E_k^\omega$, the stage mechanisms are exactly the same as $B(\hat{F}_{(1,T)}, \lambda_{(1,K)})$ and in particular, the posted-price auction with extra fee only uses the balance from bal_t. For $t \in E_k^\omega$, the give-for-free mechanism and the Myerson's auction remain the same. We use sw_t to keep track of the summation of expected valuations, i.e., $\mathrm{sw}_t = \frac{1}{3}\sum_{t' \in E_k^\omega, t' < t} \mathbb{E}_{v_{t'} \sim \hat{F}_{t'}}[v_{t'}]$. If $\mathrm{sw}_t < \tau(k)$, we turn the posted-price auction with extra fee into a give-for-free mechanism, but we increase the balance sbal instead of bal; otherwise, we run the posted-price auction with extra fee, except that it only uses the balance from sbal and it will in addition deposit the buyer's utility to sbal.

For $t \in E_k^\omega$ and $\mathrm{sw}_t < \tau(k)$, the stage mechanism is static since it in fact runs a give-for-free mechanism with probability $\frac{2(1-\lambda_k)}{3}$ and a Myerson's auction with probability $\frac{1-\lambda_k}{3}$, both of which are independent of the history. For $t \in E_k^\omega$ and $\mathrm{sw}_t \geq \tau(k)$, by choosing τ properly, we show that with high probability, even if the buyer plays strategically, $3\mathrm{sbal}_t \geq \mathbb{E}_{v_t \sim \hat{F}_t}[v_t]$, which implies that $\min\left(3\mathrm{sbal}_t, \mathbb{E}_{v_t \sim \hat{F}_t}[v_t]\right) = \mathbb{E}_{v_t \sim \hat{F}_t}[v_t]$ so that the posted-price would be 0. Therefore, with high probability, the hybrid posted-price auction with extra fee is a give-for-free mechanism with fee $\mathbb{E}_{v_t \sim \hat{F}_t}[v_t]$, which is static and independent of bal_t and sbal_t. To formally prove these statements, we exploit the fact that the dynamics of sbal_t forms a martingale for stage t with $\mathrm{sw}_t \geq \tau(k)$.

Lemma 4.3. *With $\tau(k) = \Omega\left(\ell_k^{\frac{1}{2}(1+\omega)}\sqrt{\log \ell_k} + \sqrt{\frac{\Delta_k}{\lambda_k}}\ell_k\right)$ for all k, we have*

$$\mathsf{Rev}\big(B^{hybrid}(\hat{F}_{(1,T)}, \lambda_{(1,K)}, \omega, \tau), F_{(1,T)}\big) \geq \frac{1}{3}\mathsf{Rev}\big(B^*(F_{(1,T)}), F_{(1,T)}\big) - \sum_k (\tau(k) + \lambda_k \cdot \ell_k)$$

and with probability at least $1 - \frac{1}{\ell_k}$, $\sum_{t \in \hat{E}_k} \frac{\eta_t}{a_t} \leq \tilde{O}\left(\ell_k^{1-\omega}\right)$.

4.3 The Final Policy

<u>*Learning Policy*</u>: At the start of phase k, estimate $\hat{\boldsymbol{\sigma}}_k = \arg\min_{\|\boldsymbol{\sigma}\| \leq 1} \mathcal{L}_{k-1}(\boldsymbol{\sigma})$.

<u>*Dynamic Mechanism Policy*</u>: $B^{hybrid}(\hat{F}_{(1,T)}, \lambda_{(1,K)}, \frac{1}{2}, \tau)$: at phase k

- $\lambda_k = \ell_k^{-\frac{1}{6}}$ and $\tau(k) = c^* \ell_k^{\frac{5}{6}}$;
- Compute the distributional information $\hat{F}_t$ for $t \in E_k$ according to the estimation $\hat{\boldsymbol{\sigma}}_k$;

Figure 1: Robust Non-clairvoyant Dynamic Contextual Auction Policy

We are now ready to combine our learning policy and dynamic mechanism policy to obtain our no-regret policy for contextual auctions in a non-clairvoyant environment (Figure 1). For our hybrid non-clairvoyant mechanism, we will set $\omega = \frac{1}{2}$, $\lambda_k = \ell_k^{-\frac{1}{6}}$, and $\tau(k) = c^* \ell_k^{\frac{5}{6}}$ with a large enough constant c^*. In particular, the estimation error for $\hat{\boldsymbol{\sigma}}_k$ is $\Delta_k = O(\ell_k^{-\frac{1}{2}})$ under our policy.

Theorem 4.1. *The T-stage regret of the robust non-clairvoyant dynamic contextual auction policy is $\tilde{O}(T^{\frac{5}{6}})$ against $\frac{1}{3}$-approximation of the optimal clairvoyant dynamic mechanism.*

References

Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Learning prices for repeated auctions with strategic buyers. In *Advances in Neural Information Processing Systems*, pages 1169–1177, 2013.

Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Repeated contextual auctions with strategic buyers. In *Advances in Neural Information Processing Systems*, pages 622–630, 2014.

Itai Ashlagi, Constantinos Daskalakis, and Nima Haghpanah. Sequential mechanisms with ex-post participation guarantees. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 213–214. ACM, 2016.

Dirk Bergemann and Maher Said. Dynamic auctions. In *Wiley Encyclopedia of Operations Research and Management Science*, 2011.

Dirk Bergemann and Juuso Välimäki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010.

Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, pages 459–478, 2012.

Arnoud V. den Boer. Dynamic pricing and learning: Historical origins, current research, and new directions. *Surveys in Operations Research and Management Science*, 20(1):1–18, 2015.

Negin Golrezaei, Adel Javanmard, and Vahab Mirrokni. Dynamic incentive-aware learning: Robust pricing in contextual auctions. *Available at SSRN: https://ssrn.com/abstract=3144034 or http://dx.doi.org/10.2139/ssrn.3144034*, 2018.

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

Andres M. Medina and Mehryar Mohri. Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of the 31st International Conference on Machine Learning*, pages 262–270, 2014.

Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. Dynamic auctions with bank accounts. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 387–393, 2016a.

Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. Optimal dynamic mechanisms with ex-post IR via bank accounts. *arXiv preprint arXiv:1605.08840*, 2016b.

Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. Non-clairvoyant dynamic mechanism design. In *Proceedings of the 19th ACM Conference on Economics and Computation*, pages 169–169, 2018.

Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

Christos Papadimitriou, George Pierrakos, Christos-Alexandros Psomas, and Aviad Rubinstein. On the complexity of dynamic mechanism design. In *Proceedings of the 27th annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1458–1475, 2016.

Jonathan Thomas and Tim Worrall. Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *Journal of Economic Theory*, 51(2):367–390, 1990.

Appendix

A Helper Lemmas

Lemma A.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 A.2. *For $v \in [\hat{v} - \Delta, \hat{v} + \Delta]$ and any stage IC and IR mechanism $\langle x, p\rangle$, we have*

$$u(v) - \Delta \leq u(\hat{v}) \leq u(v) + \Delta$$

where $u(v) = x(v) \cdot v - p(v)$.

Proof. Since $\langle x, p\rangle$ is a stage IC and IR mechanism, by Lemma A.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:

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

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

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

Corollary A.1. *For $F_{(1,T)}$ and $\hat{F}_{(1,T)}$ satisfying Assumption 2, and for any stage IC and IR mechanism $\langle x, p\rangle$, we have*

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

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 A.3. *For $F_{(1,T)}$ and $\hat{F}_{(1,T)}$ satisfying Assumption 2, and for any stage IC and IR mechanism $\langle x, p\rangle$, we have*

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

Proof. Since $\langle x, p\rangle$ is an stage IC and IR mechanism, by Lemma A.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:

$$\begin{aligned}
\mathbb{E}_{v_t\sim F_t}[x_t(v_t)\cdot v_t] &= \mathbb{E}_{v_t\sim\hat{F}_t,\epsilon_t\sim G_t}[x_t(v_t+\epsilon_t a_t)\cdot(v_t+\epsilon_t a_t)] \\
&= \int \frac{dx_t(r)}{dr}\cdot \mathbb{E}_{v_t\sim\hat{F}_t,\epsilon_t\sim G_t}[\mathbf{1}\{v_t+\epsilon_t a_t\ge r\}\cdot(v_t+\epsilon_t a_t)]dr \\
&\le \int \frac{dx_t(r)}{dr}\cdot \mathbb{E}_{v_t\sim\hat{F}_t,\epsilon_t\sim G_t}[\mathbf{1}\{v_t+\Delta a_t\ge r\}\cdot(v_t+\Delta a_t)]dr \\
&\le \int \frac{dx_t(r)}{dr}\cdot \mathbb{E}_{v_t\sim\hat{F}_t}[\mathbf{1}\{v_t+\Delta a_t\ge r\}\cdot(v_t+\Delta a_t)]dr \\
&\le \int \frac{dx_t(r)}{dr}\cdot \mathbb{E}_{v_t\sim\hat{F}_t}[\mathbf{1}\{v_t+\Delta a_t\ge r\}\cdot v_t]dr+\Delta a_t \\
&= \int \frac{dx_t(r)}{dr}\cdot \mathbb{E}_{v_t\sim\hat{F}_t}[\mathbf{1}\{v_t\ge r\}\cdot v_t+\mathbf{1}\{r-\Delta a_t\le v_t\le r\}\cdot v_t]dr+\Delta a_t \\
&= \mathbb{E}_{v_t\sim\hat{F}_t}[x_t(v_t)\cdot v_t]+\int \frac{dx_t(r)}{dr}\cdot\left[\int_{r-\Delta a_t}^{r} v_t\cdot f_t(v_t)dv_t\right]dr+\Delta a_t \\
&\le \mathbb{E}_{v_t\sim\hat{F}_t}[x_t(v_t)\cdot v_t]+(c_f+1)\Delta a_t
\end{aligned} \tag{2}$$

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

B Omitted Proofs in Section 3

B.1 Proof of Lemma 3.2

Proof. For ease of presentation, let $\psi_t(\mathrm{bal})=\psi_t(\mathrm{bal},F_{(1,T)},F_{(1,T)})$ and $\mathsf{fee}_t(\mathrm{bal})=\mathsf{fee}_t(\mathrm{bal},F_t)$. We use a backward induction from $t=T$ to $t=0$ to show that for all t and $\mathrm{bal}\ge 0$, the inequalities in the statement hold.

For the base case, it is true since for all $\mathrm{bal}\ge 0$, $\psi_T(\mathrm{bal})=0$. Assume the induction hypothesis is true for all $t'\ge t$. Then for $t-1$, first notice that by the computation of $\mathsf{fee}_t(\mathrm{bal})$, we have $\mathsf{fee}_t(\mathrm{bal}+\delta)=\mathsf{fee}_t(\mathrm{bal})+\delta'$ with $0\le\delta'\le 3\delta$. Therefore, we first have

$$\begin{aligned}
\psi_{t-1}(\mathrm{bal}+\delta) &= \mathbb{E}_{v_t\sim F_t}\left[\frac{1}{3}\mathsf{fee}_t(\mathrm{bal}+\delta)+\frac{1}{3}r_t^*(F_t)\cdot\mathbf{1}\{v_t\ge r_t^*(F_t)\}+\psi_t\big(\mathrm{bal}+\delta+\frac{1}{3}v_t-\frac{1}{3}\mathsf{fee}_t(\mathrm{bal}+\delta)\big)\right] \\
&\ge \mathbb{E}_{v_t\sim F_t}\left[\frac{1}{3}\mathsf{fee}_t(\mathrm{bal})+\frac{1}{3}r_t^*(F_t)\cdot\mathbf{1}\{v_t\ge r_t^*(F_t)\}+\psi_t\big(\mathrm{bal}+\frac{1}{3}v_t-\frac{1}{3}\mathsf{fee}_t(\mathrm{bal})\big)\right] \\
&= \psi_{t-1}(\mathrm{bal})
\end{aligned}$$

where the inequality follows $\delta-\frac{1}{3}\mathsf{fee}_t(\mathrm{bal}+\delta)\ge-\frac{1}{3}\mathsf{fee}_t(\mathrm{bal})$ and the induction hypothesis. We also have

$$\begin{aligned}
\psi_{t-1}(\mathrm{bal}+\delta) &= \mathbb{E}_{v_t\sim F_t}\left[\frac{1}{3}\mathsf{fee}_t(\mathrm{bal}+\delta)+\frac{1}{3}r_t^*(F_t)\cdot\mathbf{1}\{v_t\ge r_t^*(F_t)\}+\psi_t\big(\mathrm{bal}+\delta+\frac{1}{3}v_t-\frac{1}{3}\mathsf{fee}_t(\mathrm{bal}+\delta)\big)\right] \\
&= \mathbb{E}_{v_t\sim F_t}\left[\frac{1}{3}\mathsf{fee}_t(\mathrm{bal})+\frac{1}{3}r_t^*(F_t)\cdot\mathbf{1}\{v_t\ge r_t^*(F_t)\}+\frac{1}{3}\delta'+\psi_t\big(\mathrm{bal}+\delta+\frac{1}{3}v_t-\frac{1}{3}\mathsf{fee}_t(\mathrm{bal})-\frac{1}{3}\delta'\big)\right] \\
&\le \mathbb{E}_{v_t\sim F_t}\left[\frac{1}{3}\mathsf{fee}_t(\mathrm{bal})+\frac{1}{3}r_t^*(F_t)\cdot\mathbf{1}\{v_t\ge r_t^*(F_t)\}+\frac{1}{3}\delta'+\psi_t\big(\mathrm{bal}+\frac{1}{3}v_t-\frac{1}{3}\mathsf{fee}_t(\mathrm{bal})\big)+\delta-\frac{1}{3}\delta'\right] \\
&= \psi_{t-1}(\mathrm{bal})+\frac{1}{3}\delta'+(\delta-\frac{1}{3}\delta') \\
&= \psi_{t-1}(\mathrm{bal})+\delta.
\end{aligned}$$

where the inequality uses the fact that $\delta-\frac{1}{3}\delta'\ge 0$ and follows the induction hypothesis. $\square$

B.2 Proof of Lemma 3.3

Proof. For ease of presentation, let $\phi_t(\mathrm{bal})=\psi_t(\mathrm{bal},\hat{F}_{(1,T)};\hat{F}_{(1,T)})$ and $\theta_t(\mathrm{bal})=\psi_t(\mathrm{bal},F_{(1,T)};F_{(1,T)})$. We proceed by a backward induction from $t=T$ to $t=0$ to show that for all t,

$$\phi_t(\text{bal}) \geq \theta_t(\text{bal}) - (\frac{1}{3}c_f + \frac{5}{3})\Delta \sum_{t'=t+1}^{T} a_{t'}.$$

The base case for $t = T$ is clearly true since $\phi_T(\text{bal}) = \theta_T(\text{bal}) = 0$. Assume it is true for all $t' \geq t$ and we consider for stage $t-1$.

First, if we use $r_t^*(F_t)$ as the reserve price for $\hat{F}_t$, then by Lemma A.3 and Corollary A.1, we have

$$\begin{aligned} \mathbb{E}_{v_t\sim\hat{F}_t}[r_t^*(\hat{F}_t)\cdot \mathbf{1}\{v_t \geq r_t^*(\hat{F}_t)\}] &\geq \mathbb{E}_{v_t\sim\hat{F}_t}[r_t^*(F_t)\cdot \mathbf{1}\{v_t \geq r_t^*(F_t)\}] \\ &\geq \mathbb{E}_{v_t\sim F_t}[r_t^*(F_t)\cdot \mathbf{1}\{v_t \geq r_t^*(F_t)\}] - (c_f+2)\Delta a_t. \end{aligned} \tag{3}$$

where the first inequality follows that $r_t^*(\hat{F}_t)$ is the Myerson's reserve for $\hat{F}_t$. As for the spend, recall that $\mathsf{fee}_t(\text{bal}; F_t) = \min(3\text{bal}, \mathbb{E}_{v_t\sim F_t}[v_t])$, and thus, we have

$$\mathsf{fee}_t(\text{bal}; F_t) - \Delta a_t \leq \mathsf{fee}_t(\text{bal}; \hat{F}_t) \leq \mathsf{fee}_t(\text{bal}; F_t) + \Delta a_t \tag{4}$$

Therefore, combining (4) and Lemma 3.2, we have

$$\theta_t\left(\text{bal} + \frac{1}{3}(v_t + \epsilon_t a_t) - \frac{1}{3}\mathsf{fee}_t(\text{bal}; \hat{F}_t)\right) \geq \theta_t\left(\text{bal} + \frac{1}{3}v_t - \frac{1}{3}\mathsf{fee}_t(\text{bal}; F_t)\right) - \frac{2}{3}\Delta a_t \tag{5}$$

for $|\epsilon_t| \leq \Delta$. Henceforth, we have

$$\begin{aligned} \phi_{t-1}(\text{bal}) &= \mathbb{E}_{v_t\sim\hat{F}_t}\left[\frac{1}{3}\mathsf{fee}_t(\text{bal}; \hat{F}_t) + \frac{1}{3}r_t^*(\hat{F}_t)\cdot \mathbf{1}\{v_t \geq r_t^*(\hat{F}_t)\} + \phi_t\left(\text{bal} + \frac{1}{3}v_t - \frac{1}{3}\mathsf{fee}_t(\text{bal}; \hat{F}_t)\right)\right] \\ &\geq \mathbb{E}_{v_t\sim F_t}\left[\frac{1}{3}\mathsf{fee}_t(\text{bal}; F_t) + \frac{1}{3}r_t^*(F_t)\cdot \mathbf{1}\{v_t \geq r_t^*(F_t)\}\right] - (\frac{1}{3}c_f + 1)\Delta a_t \\ &\quad + \mathbb{E}_{v_t\sim\hat{F}_t}\left[\theta_t\left(\text{bal} + \frac{1}{3}v_t - \frac{1}{3}\mathsf{fee}_t(\text{bal}; \hat{F}_t)\right)\right] - (\frac{1}{3}c_f + \frac{5}{3})\Delta \sum_{t'=t+1}^{T} a_{t'} \\ &\geq \mathbb{E}_{v_t\sim F_t}\left[\frac{1}{3}\mathsf{fee}_t(\text{bal}; F_t) + \frac{1}{3}r_t^*(F_t)\cdot \mathbf{1}\{v_t \geq r_t^*(F_t)\}\right] \\ &\quad + \mathbb{E}_{v_t\sim F_t}\left[\theta_t\left(\text{bal} + \frac{1}{3}v_t - \frac{1}{3}\mathsf{fee}_t(\text{bal}; F_t)\right)\right] - (\frac{1}{3}c_f + \frac{5}{3})\Delta \sum_{t'=t}^{T} a_{t'} \\ &= \theta_{t-1}(\text{bal}) - (\frac{1}{3}c_f + \frac{5}{3})\Delta \sum_{t'=t}^{T} a_{t'} \end{aligned}$$

where the first inequality follows (3), (4), and the induction hypothesis, and the second inequality follows (5). □

B.3 Proof of Lemma 3.4

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

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

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

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

□

B.4 Proof of Lemma 3.5

Proof. By Lemma 3.4, 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 $\sum_{t'=t+1}^{T} \gamma^{t'-t} \beta_{t'}$. 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

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

□

B.5 Proof of Lemma 3.6

Proof. For ease of presentation, let $\phi_t(\mathrm{bal}) = \psi_t(\mathrm{bal}, \hat{F}_{(1,T)}; F_{(1,T)})$ and $\theta_t(\mathrm{bal}) = \psi_t(\mathrm{bal}, \hat{F}_{(1,T)}; \hat{F}_{(1,T)})$. We proceed by a backward induction from $t = T$ to $t = 0$ to show that for all t,

$$\phi_t(\mathrm{bal}) \geq \theta_t(\mathrm{bal}) - (\frac{1}{3}c_f + \frac{5}{3}) \sum_{t'=t+1}^{T} (\Delta a_{t'} + \eta_{t'})$$

The base case for $t = T$ is clearly true since $\phi_T(\mathrm{bal}) = \theta_T(\mathrm{bal}) = 0$. Assume it is true for all $t' \geq t$ and we consider for stage $t - 1$. First, recall that

$$\begin{aligned} \phi_{t-1}(\mathrm{bal}) &= \mathbb{E}_{v_t \sim F_t} \left[\frac{1}{3} \mathsf{fee}_t(\mathrm{bal}; \hat{F}_t) + \frac{1}{3} r_t^*(\hat{F}_t) \cdot \mathbf{1}\{v_t' \geq r_t^*(\hat{F}_t)\} + \phi_t \left(\mathrm{bal} + \frac{1}{3} v_t' - \frac{1}{3} \mathsf{fee}_t(\mathrm{bal}; \hat{F}_t) \right) \right] \\ &= \mathbb{E}_{v_t \sim \hat{F}_t, \epsilon_t} \left[\frac{1}{3} \mathsf{fee}_t(\mathrm{bal}; \hat{F}_t) + \frac{1}{3} r_t^*(\hat{F}_t) \cdot \mathbf{1}\{v_t' \geq r_t^*(\hat{F}_t)\} + \phi_t \left(\mathrm{bal} + \frac{1}{3} v_t' - \frac{1}{3} \mathsf{fee}_t(\mathrm{bal}; \hat{F}_t) \right) \right] \end{aligned}$$

where in the last equality, v_t is the valuation drawn from $\hat{F}_t$ and v_t' is the reported bid given the buyer's true valuation is $v_t + \epsilon_t a_t$ with $|\epsilon_t| \leq \Delta_t$. Therefore, we have $v_t' \in [v_t - \Delta a_t - \eta_t, v_t + \Delta a_t + \eta_t]$. By Lemma A.3 and Corollary A.1,

$$\mathbb{E}_{v_t \sim \hat{F}_t, \epsilon_t}[r_t^*(\hat{F}_t) \cdot \mathbf{1}\{v_t' \geq r_t^*(\hat{F}_t)\}] \geq \mathbb{E}_{v_t \sim \hat{F}_t}[r_t^*(\hat{F}_t) \cdot \mathbf{1}\{v_t \geq r_t^*(\hat{F}_t)\}] - (c_f + 2)(\Delta a_t + \eta_t). \tag{6}$$

As for the spend, by Lemma 3.2 and the fact that $\mathsf{fee}_t(\mathrm{bal}; \hat{F}_t) \leq \mathsf{fee}_t(\mathrm{bal}; F_t) + \Delta a_t$ we first have

$$\theta_t\left(\mathrm{bal} + \frac{1}{3}v_t' - \frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; \hat{F}_t)\right) \geq \theta_t\left(\mathrm{bal} + \frac{1}{3}v_t - \frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; F_t)\right) - \frac{2}{3}\Delta a_t - \frac{1}{3}\eta_t \tag{7}$$

Henceforth, we have

$$\begin{aligned}
\phi_{t-1}(\mathrm{bal}) &= \mathbb{E}_{v_t\sim F_t}\left[\frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; \hat{F}_t) + \frac{1}{3}r_t^*(\hat{F}_t)\cdot \mathbf{1}\{v_t' \geq r_t^*(\hat{F}_t)\} + \phi_t\left(\mathrm{bal} + \frac{1}{3}v_t' - \frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; \hat{F}_t)\right)\right] \\
&\geq \mathbb{E}_{v_t\sim \hat{F}_t}\left[\frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; \hat{F}_t) + \frac{1}{3}r_t^*(\hat{F}_t)\cdot \mathbf{1}\{v_t \geq r_t^*(\hat{F}_t)\}\right] - (\frac{1}{3}c_f + 1)(\Delta a_t + \eta_t) \\
&\quad + \mathbb{E}_{v_t\sim \hat{F}_t}\left[\theta_t\left(\mathrm{bal} + \frac{1}{3}v_t' - \frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; \hat{F}_t)\right)\right] - (\frac{1}{3}c_f + \frac{5}{3})\sum_{t'=t+1}^{T}(\Delta a_{t'} + \eta_{t'}) \\
&\geq \mathbb{E}_{v_t\sim \hat{F}_t}\left[\frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; \hat{F}_t) + \frac{1}{3}r_t^*(\hat{F}_t)\cdot \mathbf{1}\{v_t \geq r_t^*(\hat{F}_t)\}\right] \\
&\quad + \mathbb{E}_{v_t\sim \hat{F}_t}\left[\theta_t\left(\mathrm{bal} + \frac{1}{3}v_t - \frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; F_t)\right)\right] - (\frac{1}{3}c_f + \frac{5}{3})\sum_{t'=t}^{T}(\Delta a_{t'} + \eta_{t'}) \\
&= \theta_{t-1}(\mathrm{bal}) - (\frac{1}{3}c_f + \frac{5}{3})\sum_{t'=t}^{T}(\Delta a_{t'} + \eta_{t'})
\end{aligned}$$

where the first inequality follows (6) and the induction hypothesis, and the second inequality follows (7).

Moreover, notice that we have

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

where the first inequality follows the Cauchy-Schwarz inequality and the last inequality is due to Assumption 1. $\square$

C Proof of Lemma 4.1

Sketch: Recall that in a contextual auction, the buyer's true valuation is $v_t = a_t(\langle \boldsymbol{\sigma}, \boldsymbol{\zeta}_t\rangle + \varepsilon_t)$ where a_t is the intrinsic value of the item, $\boldsymbol{\zeta}_t$ is the contextual vector, and ε_t is a random variable following the market noise distribution M_t. Notice that $M_t(w_t - \langle \boldsymbol{\sigma}, \boldsymbol{\zeta}_t\rangle)$ is the probability of the event that $\langle \boldsymbol{\sigma}, \boldsymbol{\zeta}_t\rangle + \varepsilon_t = \frac{v_t}{a_t} \leq w_t$, which is equivalent to $v_t \leq a_t \cdot w_t$. As a result, assuming the buyer reports truthfully, $\mathcal{L}_{k-1}(\boldsymbol{\sigma})$ is exactly the negative of log-likelihood corresponding to $\boldsymbol{\sigma}$.

Under truthful reporting, we have $\mathbf{1}\{b_t \geq a_t \cdot w_t\} = \mathbf{1}\{v_t \geq a_t \cdot w_t\}$. For a $\eta_{(1,T)}$-DIC robust dynamic mechanism, we have $v_t - \eta_t \leq b_t \leq v_t + \eta_t$. As a result, if $|a_t \cdot w_t - v_t| > \eta_t$, then any misreport from the buyer does not result in a lie. Therefore, a lie occurs only if the true valuation $v_t \in [a_t \cdot w_t - \eta_t, a_t \cdot w_t + \eta_t]$. By a martingale argument on the sequence of lies, we can obtain that the total number of lies caused by the dynamic mechanism within phase k 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 $t \in \hat{E}_{k-1}$, the gain from changing the mechanism for the future phases via changing the seller's estimation is relatively small, since the buyer discounts the future.

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 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

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

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

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

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 and Young, 2014]) and denoting $\ell = |\hat{E}_{k-1}|$, we have

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

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

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

$\square$

D Hybrid Non-clairvoyant Mechanism

We adapt $B(\hat{F}_{(1,T)}, \lambda_{(1,K)})$ to obtain a hybrid non-clairvoyant mechanism $B^{hybrid}(\hat{F}_{(1,T)}, \lambda_{(1,K)}, \omega, \tau)$, which is parameterized by a real number $\omega \in (0,1)$ and a function $\tau : \mathbb{Z}_+ \to \mathbb{R}_+$ that maps the phase number to a real number. The stage mechanism at stage t is parameterized by two non-negative balances bal_t and sbal_t, and an additional parameter sw_t. In particular, sw_t is reset to 0 at the beginning of each phase, i.e., for $t = \ell_k$.

For the give-for-free mechanism, the Myerson's auction, and the random posted-price auction, their allocation rules, payment rules, and the update rule for bal remain the same, while they keep sw and sbal the same, i.e., $\mathrm{sw}_{t+1} = \mathrm{sw}_t$ and $\mathrm{sbal}_{t+1} = \mathrm{sbal}_t$. We replace the posted-price auction with extra fee by a hybrid posted-price auction with extra fee.

Definition D.1 (Hybrid Posted-price Auction with Extra Fee). *For $t \in E_k$, let $E_k^{\omega} = \{t \mid a_t < \ell_k^{\omega}\}$.*

- *If $t \notin E_k^\omega$: let* $\mathsf{fee}_t^b(\mathrm{bal}_t; \hat{F}_t) = \min\left(3\mathrm{bal}_t, \mathbb{E}_{v_t \sim \hat{F}_t}[v_t]\right)$ *and* $r_t(\mathrm{bal}_t)$ *be the posted-price such that*

$$\mathbb{E}_{v_t \sim \hat{F}_t}\left[\left(v_t - r_t(\mathrm{bal}_t)\right)^+\right] = \mathsf{fee}_t^b(\mathrm{bal}_t; \hat{F}_t).$$

The mechanism charges the buyer $\mathsf{fee}_t^b(\mathrm{bal}_t; \hat{F}_t)$ *before the buyer learns her valuation and then run a posted-price auction with price* $r_t(\mathrm{bal}_t)$

$$\begin{aligned} x_t^H &= \mathbf{1}\{b_t \geq r_t(\mathrm{bal}_t)\}, \\ p_t^H &= \mathsf{fee}_t^b(\mathrm{bal}_t; \hat{F}_t) + r_t(\mathrm{bal}_t) \cdot \mathbf{1}\{b_t \geq r_t(\mathrm{bal}_t)\} \end{aligned}$$

and update the balances: $\mathrm{bal}_{t+1}^H = \mathrm{bal}_t - \mathsf{fee}_t^b(\mathrm{bal}_t; \hat{F}_t)$, $\mathrm{sbal}_{t+1}^H = \mathrm{sbal}_t$, *and* $\mathrm{sw}_{t+1}^H = \mathrm{sw}_t$.

- *otherwise, if $t \in E_k^\omega$: we first update* $\mathrm{sw}_{t+1}^H = \mathrm{sw}_t + \mathbb{E}_{v_t \sim \hat{F}_t}[v_t]$;

  - *if* $\mathrm{sw}_t \geq \tau(k)$, *let* $\mathsf{fee}_t^s(\mathrm{sbal}_t; \hat{F}_t) = \min\left(3\mathrm{sbal}_t, \mathbb{E}_{v_t \sim \hat{F}_t}[v_t]\right)$ *and* $r_t(\mathrm{sbal}_t)$ *be the posted-price such that*

  $$\mathbb{E}_{v_t \sim \hat{F}_t}\left[\left(v_t - r_t(\mathrm{sbal}_t)\right)^+\right] = \mathsf{fee}_t^s(\mathrm{sbal}_t; \hat{F}_t)$$

  The mechanism charges the buyer $\mathsf{fee}_t^s(\mathrm{sbal}_t; \hat{F}_t)$ *before the buyer learns her valuation and then run a posted-price auction with price* $r_t(\mathrm{sbal}_t)$

  $$\begin{aligned} x_t^H &= \mathbf{1}\{b_t \geq r_t(\mathrm{sbal}_t)\}, \\ p_t^H &= \mathsf{fee}_t^s(\mathrm{sbal}_t; \hat{F}_t) + r_t(\mathrm{sbal}_t) \cdot \mathbf{1}\{b_t \geq r_t(\mathrm{sbal}_t)\} \end{aligned}$$

  and update the balances: $\mathrm{bal}_{t+1}^H = \mathrm{bal}_t$ *and*

  $$\mathrm{sbal}_{t+1}^H = \mathrm{sbal}_t - \mathsf{fee}_t^s(\mathrm{sbal}_t; \hat{F}_t) + \mathbf{1}\{b_t \geq r_t(\mathrm{sbal}_t)\} \cdot (b_t - r_t(\mathrm{sbal}_t));$$

  - *otherwise: allocate the item no matter what the buyer's bid is. Moreover, increase the balance* sbal_t *by the buyer's bid:*

  $$\begin{aligned} x_t^H &= 1, \qquad & p_t^H &= 0, \\ \mathrm{bal}_{t+1}^H &= \mathrm{bal}_t, \qquad & \mathrm{sbal}_{t+1}^H &= \mathrm{sbal}_t + b_t. \end{aligned}$$

We prove Lemma 4.3 in this section. Lemma 4.3 states that by choosing $\tau(k)$ properly: (1) the revenue loss from running a hybrid non-clairvoyant mechanism against the non-clairvoyant mechanism is small; (2) the number of lies is small. The proof of the first property based on a new revenue tracking program that separates the revenue contribution related to bal (from stages $t \notin E_k^\omega$) and the revenue contribution related to sbal (from stages $t \in E_k^\omega$). The argument for the revenue from stages $t \notin E_k^\omega$ simply follows the argument of $\frac{1}{3}$-approximation of the non-clairvoyant mechanism, while the argument for the revenue from stages $t \in E_k^\omega$ exploits the martingale property of sbal and the fact that $\mathbb{E}_{v_t \sim \hat{F}_t}[v_t]$ is exactly the maximum extra fee we can charge in the posted-price auction with extra fee (Section D.2). We then combine the martingale natural of sbal and techniques in robust non-clairvoyant mechanism to show the number of lies is small (Section D.3).

D.1 Bank Account Property

We generalize the definition of BI to accommodate the introduction of sbal_t:

- The mechanism ensures that the expected utility is balance independent if the buyer reports truthfully:

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

  is a non-negative constant independent of bal_t and sbal_t.

- A balance update rule never uses more than the total balance from bal_t and sbal_t, and never deposits more than the buyer's utility into bal_t and sbal_t in total:

$$\begin{aligned} &\mathrm{bal}_{t+1} \geq 0, \qquad \mathrm{sbal}_{t+1} \geq 0 \\ &\mathrm{bal}_{t+1} + \mathrm{sbal}_{t+1} \leq \mathrm{bal}_t + \mathrm{sbal}_t + b_t \cdot x_t(\mathrm{bal}_t, \mathrm{sbal}_t, \mathrm{sw}_t, b_t) - p_t^B(\mathrm{bal}_t, \mathrm{sbal}_t, \mathrm{sw}_t, b_t) \end{aligned} \tag{sBU}$$

Notice that we allow dependence on sw_t in sBI. This is because sw_t is a global parameter such that it is the same at stage t for all possible historical bids in the past.

Lemma D.1. *The hybrid non-clairvoyant mechanism $B^{hybrid}(\hat{F}_{(1,T)}, \lambda_{(1,K)}, \omega, \tau)$ is stage-IC, sBI and sBU for $\hat{F}_{(1,T)}$. Therefore, $B^{hybrid}(\hat{F}_{(1,T)}, \lambda_{(1,K)}, \omega, \tau)$ is $\eta_{(1,T)}$-DIC with $\eta_t = 0$ and ex-post IR for $\hat{F}_{(1,T)}$.*

Proof. Since all mechanisms are variants of leave-it-or-take-it mechanisms, the mixture of them is clearly stage-IC. For sBU, notice that we only decrease bal and sbal in the posted price auction, and moreover, by the construction of $r_t(\mathrm{bal}_t, \mathrm{sbal}_t; \hat{F}_t)$, it is at most bal_t (sbal_t) when bal_t (sbal_t) is deducted. Furthermore, it is straightforward to verify that the sum of the deposit to bal and sbal is at most the buyer's utility at stage t. Therefore, the mechanism is sBU. To demonstrate the mechanism is sBI, notice that when $a_t \geq \ell_k^\omega$ or $\mathrm{sw}_t \geq \tau(k)$, the buyer's expected utility is exactly $\mathbb{E}_{v_t \sim \hat{F}_t}[v_t] + 0 + \mathbb{E}_{v_t \sim \hat{F}_t}[(v_t - r_t^*(\hat{F}_t)]$ for all bal_t and sbal_t; otherwise, the buyer's expected utility is $2\mathbb{E}_{v_t \sim \hat{F}_t}[v_t] + \mathbb{E}_{v_t \sim \hat{F}_t}[(v_t - r_t^*(\hat{F}_t)]$ for all bal_t and sbal_t. Thus, the mechanism is sBI.

Since the mechanism is sBI, the buyer's historical reports have no impact on her future expected utilities, assuming she reports truthfully in the future. Combining with the fact that the mechanism is stage-IC for every stage, the mechanism is $\eta_{(1,T)}$-DIC with $\eta_{(1,T)} = (0, \ldots, 0)$. Moreover, by the balance update property sBU, the nonnegative $\mathrm{bal}_t + \mathrm{sbal}_t$ always lower bounds the buyer's utility provided truthful reporting. Thus, the mechanism is ex-post IR. $\square$

D.2 Revenue Tracking Program

We develop a program to compute the revenue obtained from the hybrid non-clairvoyant mechanism. For convenience, let $\mathsf{fee}_t^b(\mathrm{bal}; \hat{F}_t) = 0$ for $t \in E_k^\omega$ (in which $\mathsf{fee}_t^b(\mathrm{bal}; \hat{F}_t)$ is not defined). Moreover, for stage t such that $t \in E_k^\omega$ and $\mathrm{sw}_t < \tau(k)$ (in which $\mathsf{fee}_t^s(\mathrm{sbal}; \hat{F}_t)$ and $r_t(\mathrm{sbal})$ are not defined), let $\mathsf{fee}_t^s(\mathrm{sbal}; \hat{F}_t) = r_t(\mathrm{sbal}) = 0$.

Definition D.2. *For a hybrid non-clairvoyant mechanism $B^{hybrid}(\hat{F}_{(1,T)}, \lambda_{(1,K)}, \omega, \tau)$, we consider revenue tracking programs $\psi_t^b(\mathrm{bal}, \hat{F}_{(1,T)}; F_{(1,T)})$ and $\psi_t^s(\mathrm{sbal}, \hat{F}_{(1,T)}; F_{(1,T)})$ to keep track on the revenue of implementing $B^{hybrid}(\hat{F}_{(1,T)}, \lambda_{(1,K)}, \omega, \tau)$ when the buyer's true distribution is $F_{(1,T)}$. We define $\psi_T^b(\mathrm{bal}) = \psi_T^s(\mathrm{sbal}) = 0$ and for $t < T$,*

$$\begin{aligned} \psi_{t-1}^b(\mathrm{bal}, \hat{F}_{(1,T)}; F_{(1,T)}) = \mathbb{E}_{v_t \sim F_t} \; & \Bigg[\frac{1}{3}\mathsf{fee}_t^b(\mathrm{bal}; \hat{F}_t) + \frac{1}{3} r_t^*(\hat{F}_t) \cdot \mathbf{1}\{v_t' \geq r_t^*(\hat{F}_t)\} \\ & + \psi_t^b\left(\mathrm{bal} + \frac{1}{3}v_t' - \frac{1}{3}\mathsf{fee}_t^b(\mathrm{bal}; \hat{F}_t), \hat{F}_{(1,T)}; F_{(1,T)}\right) \Bigg] \end{aligned} \tag{8}$$

where v_t' is the buyer's reported bid that maximizes her continuation utility when the buyer's true valuation is v_t.

Moreover, for $t \notin E_k^\omega$, $\psi_{t-1}^s(\mathrm{sbal}, \hat{F}_{(1,T)}; F_{(1,T)}) = \psi_t^s(\mathrm{sbal}, \hat{F}_{(1,T)}; F_{(1,T)})$; otherwise,

$$\begin{aligned} &\psi_{t-1}^s(\mathrm{sbal}, \hat{F}_{(1,T)}; F_{(1,T)}) \\ &= \mathbb{E}_{v_t \sim F_t}\left[\frac{1}{3}\mathsf{fee}_t^s(\mathrm{sbal}; \hat{F}_t) + \psi_t^s\left(\mathrm{sbal} + \frac{1}{3}\big(v_t' - \mathsf{fee}_t^s(\mathrm{sbal}; \hat{F}_t) - r_t(\mathrm{sbal})\big), \hat{F}_{(1,T)}; F_{(1,T)}\right)\right] \end{aligned} \tag{9}$$

where v_t' is the reported bid when the buyer's true valuation is v_t.

Notice that we separate the revenue tracking for bal and sbal. Moreover, the revenue obtained from the Myerson's auction are counted in ψ_t^b. Similar to the revenue tracking program for the non-clairvoyant mechanism (1), we record the revenue from each stage t while omitting the possible revenue $r_t(\mathrm{bal})$ or $r_t(\mathrm{sbal})$ from the posted-price auction with extra fee.

Proposition D.1. $\mathsf{Rev}\left(B^{hybrid}(\hat{F}_{(1,T)}, \lambda_{(1,K)}, \omega, \tau); F_{(1,T)}\right) \geq \psi_0^b(0, \hat{F}_{(1,T)}; F_{(1,T)}) + \psi_0^s(0, \hat{F}_{(1,T)}; F_{(1,T)}) - O(\lambda T)$.

Revenue Performance with Perfect Distributional Information We first compare the revenue obtained by the hybrid non-clairvoyant mechanism and the non-clairvoyant mechanism, when the seller's distributional information is perfect, i.e., $\hat{F}_{(1,T)} = F_{(1,T)}$. Notice that the definition of ψ_t^b (8) is exactly the same as ψ_t for the non-clairvoyant mechanism (1). However, the difference is that for stage t in which $a_t < \ell_k^\omega$, $\mathsf{fee}_t^b(\mathrm{bal}; \hat{F}_t) = 0$ in the hybrid non-clairvoyant mechanism. Following an argument similar to the proof of Lemma 3.2, we have the following lemma:

Lemma D.2. *For any $F_{(1,T)}$, we have for all $0 \leq t \leq T$,*

$$\psi_t^b(\mathrm{bal} + \delta, F_{(1,T)}; F_{(1,T)}) - \delta \leq \psi_t^b(\mathrm{bal}, F_{(1,T)}; F_{(1,T)}) \leq \psi_t^b(\mathrm{bal} + \delta, F_{(1,T)}; F_{(1,T)}).$$

Therefore, all our results for the robust non-clairvoyant mechanism (Section 3) works for the revenue obtained from ψ_t^b. We then compute the revenue obtain from ψ_t^b when the seller has perfect distributional information.

Lemma D.3. $\psi_0^b(0, F_{(1,T)}; F_{(1,T)}) \geq \psi_0(0, F_{(1,T)}; F_{(1,T)}) - \frac{1}{3}\sum_k \sum_{t \in E_k^\omega} \mathbb{E}_{v_t \sim F_t}[v_t]$.

Proof. For simplicity, let $\phi_t(\mathrm{bal}) = \psi_t^b(\mathrm{bal}, F_{(1,T)}; F_{(1,T)})$ and $\theta_t(\mathrm{bal}) = \psi_t(\mathrm{bal}, F_{(1,T)}; F_{(1,T)})$. We prove by a backward induction from $t = T$ to $t = 0$ to show that for all t and $\mathrm{bal} \geq 0$,

$$\phi_t(\mathrm{bal}) \geq \theta_t(\mathrm{bal}) - \frac{1}{3}\sum_k \sum_{t' \in E_k^\omega, t' > t} \mathbb{E}_{v_{t'} \sim F_{t'}}[v_{t'}].$$

The base case is true for $t = T$ since $\phi_T(\mathrm{bal}) = \theta_T(\mathrm{bal}) = 0$ for all $\mathrm{bal} \geq 0$. Assume the induction hypothesis is true for $t' \geq t$ and we consider stage $t-1$. For $t \in E_k$, if $t \notin E_k^\omega$, we have $\mathsf{fee}_t^b(\mathrm{bal}; F_t) = \mathsf{fee}_t(\mathrm{bal}; F_t)$. Therefore, we have

$$\begin{aligned}
\phi_{t-1}(\mathrm{bal}) &= \mathbb{E}_{v_t \sim F_t}\left[\frac{1}{3}\mathsf{fee}_t^b(\mathrm{bal}; F_t) + \frac{1}{3} r_t^*(F_t) \cdot \mathbf{1}\{v_t \geq r_t^*(F_t)\} + \phi_t\left(\mathrm{bal} + \frac{1}{3}v_t - \frac{1}{3}\mathsf{fee}_t^b(\mathrm{bal}; F_t)\right)\right] \\
&= \mathbb{E}_{v_t \sim F_t}\left[\frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; F_t) + \frac{1}{3} r_t^*(F_t) \cdot \mathbf{1}\{v_t \geq r_t^*(F_t)\} + \phi_t\left(\mathrm{bal} + \frac{1}{3}v_t - \frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; F_t)\right)\right] \\
&= \theta_{t-1}(\mathrm{bal}) + \mathbb{E}_{v_t \sim F_t}\left[\phi_t\left(\mathrm{bal} + \frac{1}{3}v_t - \frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; F_t)\right) - \theta_t\left(\mathrm{bal} + \frac{1}{3}v_t - \frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; F_t)\right)\right] \\
&\geq \theta_{t-1}(\mathrm{bal}) - \frac{1}{3}\sum_k \sum_{t' \in E_k^\omega, t' \geq t} \mathbb{E}_{v_{t'} \sim F_{t'}}[v_{t'}]
\end{aligned}$$

where the inequality follows the induction hypothesis. On the other hand, if $t \in E_k^\omega$, we have $\mathsf{fee}_t^b(\mathrm{bal}; F_t) = 0$. As a result, we have

$$\begin{aligned}
\phi_{t-1}(\mathrm{bal}) &= \mathbb{E}_{v_t \sim F_t}\left[\frac{1}{3}\mathsf{fee}_t^b(\mathrm{bal}; F_t) + \frac{1}{3} r_t^*(F_t) \cdot \mathbf{1}\{v_t \geq r_t^*(F_t)\} + \phi_t\left(\mathrm{bal} + \frac{1}{3}v_t - \frac{1}{3}\mathsf{fee}_t^b(\mathrm{bal}; F_t)\right)\right] \\
&\geq \mathbb{E}_{v_t \sim F_t}\left[\frac{1}{3} r_t^*(F_t) \cdot \mathbf{1}\{v_t \geq r_t^*(F_t)\} + \phi_t\left(\mathrm{bal} + \frac{1}{3}v_t - \mathsf{fee}_t(\mathrm{bal}; F_t)\right)\right] \\
&= \theta_{t-1}(\mathrm{bal}) - \frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; F_t) \\
&\quad + \mathbb{E}_{v_t \sim F_t}\left[\phi_t\left(\mathrm{bal} + \frac{1}{3}v_t - \frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; F_t)\right) - \theta_t\left(\mathrm{bal} + \frac{1}{3}v_t - \frac{1}{3}\mathsf{fee}_t(\mathrm{bal}; F_t)\right)\right] \\
&\geq \theta_{t-1}(\mathrm{bal}) - \frac{1}{3}\sum_k \sum_{t' \in E_k^\omega, t' \geq t} \mathbb{E}_{v_{t'} \sim F_{t'}}[v_{t'}]
\end{aligned}$$

where the first inequality follows Lemma D.2 and the second inequality uses the induction hypothesis and the fact that $\mathsf{fee}_t(\mathrm{bal}; F_t) \le \mathbb{E}_{v_t \sim F_t}[v_t]$. $\square$

Let $\tilde{E}_k^\omega = \{t \in E_k^\omega \mid \mathrm{sw}_t < \tau(k)\}$. Let $t^*(k) = \max \tilde{E}_k^\omega$ and consider a sequence y_t for $t \in E_k^\omega$ such that

$$y_t = \begin{cases} \frac{1}{3}\sum_{t' \in E_k, t' \le t} E_{v_{t'} \sim F_{t'}}[v_{t'}] & t' \in \tilde{E}_k^\omega \\ y_{t^*(k)} & t' \notin \tilde{E}_k^\omega \end{cases}$$

Henceforth, the key observation is that the sequence $\{\mathrm{sbal}_t - y_t\}_{t \in E_k^\omega}$ forms a martingale with bounded difference a_t at stage t: for $t \in \tilde{E}_k^\omega$, we have

$$\mathbb{E}_{v_t \sim F_t}[\mathrm{sbal}_{t+1} - \mathrm{sbal}_t] = \mathbb{E}_{v_t \sim F_t}[\mathrm{sbal}_t + \frac{1}{3}v_t - \mathrm{sbal}_t] = \frac{1}{3}\mathbb{E}_{v_t \sim F_t}[v_t]$$

and for $t \in E_k^\omega \setminus \tilde{E}_k^\omega$, we have

$$\mathbb{E}_{v_t \sim F_t}[\mathrm{sbal}_{t+1} - \mathrm{sbal}_t] = \mathbb{E}_{v_t \sim F_t}\left[\mathrm{sbal}_t + \frac{1}{3}\Big(v_t - \mathsf{fee}_t^s(\mathrm{sbal}_t; F_t) - r_t(\mathrm{sbal}_t)\Big) - \mathrm{sbal}_t\right] = 0$$

where the last equality follows the fact that $\mathbb{E}_{v_t \sim \hat{F}_t}\left[\big(v_t - r_t(\mathrm{sbal}_t)\big)^+\right] = \mathsf{fee}_t^s(\mathrm{sbal}_t; \hat{F}_t)$ from the construction of the hybrid non-clairvoyant mechanism.

Lemma D.4. *If $\tau(k) \ge 4\sqrt{c_a} \cdot \ell_k^{\frac{1}{2}(1+\omega)}\sqrt{\log \ell_k}$, for any $t \in E_k^\omega \setminus \tilde{E}_k^\omega$,* $\Pr\left[\mathrm{sbal}_t < y_{t^*(k)} - \delta\right] \le \exp\left(-\frac{\delta^2}{4c_a \ell_k^{1+\omega}}\right)$.

Proof. Notice that $y_{t^*(k)} \ge \tau(k)$ and by Azuma's inequality, we have for any $t \in E_k^\omega \setminus \tilde{E}_k^\omega$,

$$\Pr\left[\mathrm{sbal}_t < y_{t^*(k)} - \delta\right] \le \exp\left(-\frac{\delta^2}{2\sum_{t \in E_k^\omega} a_t^2}\right) \le \exp\left(-\frac{\delta^2}{4c_a \ell_k^{1+\omega}}\right)$$

where the second inequality follows that

$$\sum_{t \in E_k^\omega} a_t^2 \le (\ell_k^\omega)^2 \cdot \frac{\sum_{t \in E_k^\omega} a_t}{\ell_k^\omega} \le (\ell_k^\omega)^2 \cdot \frac{2c_a \ell_k}{\ell_k^\omega} = 2c_a \cdot \ell_k^{1+\omega}$$

where the second inequality follows Assumption 1. $\square$

Lemma D.5. *If $\tau(k) \ge 4\sqrt{c_a} \cdot \ell_k^{\frac{1}{2}(1+\omega)}\sqrt{\log \ell_k}$, we have*

$$\psi_0^s(0, F_{(1,T)}; F_{(1,T)}) \ge \frac{1}{3}\sum_k \left(\sum_{t \in E_k^\omega} \mathbb{E}_{v_t \sim F_t}[v_t] - \tau(k)\right) - \tilde{O}(T^\omega).$$

Proof. For convenience, let $\phi_t(\mathrm{sbal}) = \psi_t^s(\mathrm{sbal}, F_{(1,T)}; F_{(1,T)})$. Recall that $\phi_{t-1}(\mathrm{sbal}) = \phi_t(\mathrm{sbal})$ if $t \in E_k$ and $t \notin E_k^\omega$. Moreover, recall that sw_t is set to 0 at stage ℓ_k for all k and when $t \in E_k$ and $\mathrm{sw}_t < \tau(k)$, we in fact offer a give-for-free mechanism in the hybrid posted-price auction. Therefore, the mechanism does not accrue any revenue from stages with $t \in \tilde{E}_k^\omega$.

Plugging in $\delta = \tau(k) - \frac{1}{3}\ell_k^\omega$ in Lemma D.4, we have

$$\Pr\left[\mathrm{sbal}_t < \frac{1}{3}\ell_k^\omega\right] \le \Pr\left[\mathrm{sbal}_t < y_{t^*(k)} - \tau(k) + \frac{1}{3}\ell_k^\omega\right] \le \exp\left(-\frac{\left(\tau(k) - \frac{1}{3}\ell_k^\omega\right)^2}{4c_a \ell_k^{1+\omega}}\right) \le \frac{1}{\ell_k^2}$$

where the first inequality is due to $y_{t^*(k)} \ge \tau(k)$. Applying the union bound, we have

$$\Pr\left[\exists t \in E_k^\omega \setminus \tilde{E}_k^\omega, \mathrm{sbal}_t < \frac{1}{3}\ell_k^\omega\right] \le \frac{1}{\ell_k}.$$

Therefore, with probability at least $(1-\frac{1}{\ell_k})$, for all $t \in E_k^\omega \setminus \tilde{E}_k^\omega$, $3\mathrm{sbal}_t \geq \ell_k^\omega \geq a_t \geq \mathbb{E}_{v_t\sim F_t}[v_t]$, which implies that $\mathsf{fee}_t^s(\mathrm{sbal}_t; F_t) = \mathbb{E}_{v_t\sim F_t}[v_t]$. Thus, combining with the fact that $y_{t^*(k)} \leq \tau(k) + \ell_k^\omega$ for all k, we have for the revenue obtained from $\frac{1}{3}\mathsf{fee}_t^s(\mathrm{sbal}_t; F_t)$ is at least

$$(1-\frac{1}{\ell_k})\cdot\frac{1}{3}\left(\sum_{t\in E_k^\omega}\mathbb{E}_{v_t\sim F_t}[v_t] - y_{t^*(k)}\right) = \frac{1}{3}\left(\sum_{t\in E_k^\omega}\mathbb{E}_{v_t\sim F_t}[v_t] - \tau(k)\right) - O(\ell_k^\omega)$$

We conclude the proof of the lemma by taking the summation over all the phases. □

Combining Lemma D.3 and D.5, we can conclude that

Corollary D.1. *By setting $\tau(k) = 4\sqrt{c_a}\cdot\ell_k^{\frac{1}{2}(1+\omega)}\sqrt{\log \ell_k}$, we have*

$$\psi_0^b(0, F_{(1,T)}; F_{(1,T)}) + \psi_0^s(0, F_{(1,T)}; F_{(1,T)}) \geq \psi_0(0, F_{(1,T)}; F_{(1,T)}) - \tilde{O}\left(T^{\frac{1}{2}(1+\omega)}\right).$$

Therefore, the revenue loss of the hybrid non-clairvoyant mechanism against the optimal clairvoyant mechanism is sublinear in T when $\omega \in (0,1)$.

D.3 Analysis on the Misreport

We analyze the buyer's misreport in this section. By the discussion in Section 4.1, we focus on $\hat{E}_k$ instead of E_k. We first provide a naive bound for the property of $\eta_{(1,T)}$-DIC in $\hat{E}_k$ for $B^{hybrid}(\hat{F}_{(1,T)}, \lambda_{(1,K)}, \omega, \tau)$.

Proposition D.2. *In $B^{hybrid}(\hat{F}_{(1,T)}, \lambda_{(1,K)}, \omega, \tau)$, for $t \in \hat{E}_k$, we have*

$$\eta_t \leq 4\sqrt{\frac{a_t\Delta_k}{\lambda_k}\cdot\sum_{t'\in E_k, t'>t}\gamma^{t'-t}a_{t'}}$$

and moreover,

$$\sum_{t\in\hat{E}_k}\eta_t = 4c_a\sqrt{\frac{\Delta_k}{(1-\gamma)\lambda_k}}\cdot\ell_k.$$

Proof. By Lemma 3.5, we have

$$\begin{aligned}
\eta_t &\leq \sqrt{\frac{4a_t\Delta_k}{\lambda_k}\cdot\sum_{t'=t+1}^{T}\gamma^{t'-t}a_{t'}} \\
&= \eta_t \leq \sqrt{\frac{4a_t\Delta_k}{\lambda_k}\cdot\left(\sum_{t'\in E_k, t'>t}\gamma^{t'-t}a_{t'} + \sum_{t'\geq \ell_{k+1}}\gamma^{t'-t}a_{t'}\right)} \\
&\leq \eta_t \leq \sqrt{\frac{4a_t\Delta_k}{\lambda_k}\cdot\left(\sum_{t'\in E_k, t'>t}\gamma^{t'-t}a_{t'} + \frac{c_a}{(1-\gamma)^2\cdot\ell_{k+1}^5}\right)} \\
&\leq \eta_t \leq 4\sqrt{\frac{a_t\Delta_k}{\lambda_k}\cdot\left(\sum_{t'\in E_k, t'>t}\gamma^{t'-t}a_{t'}\right)}
\end{aligned}$$

Combining with the argument similar to the proof of Lemma 3.6, we can conclude that

$$\sum_{t\in\hat{E}_k}\eta_t = 4c_a\sqrt{\frac{\Delta_k}{(1-\gamma)\lambda_k}}\cdot\ell_k.$$

□

For convenience, let $\hat{E}_k^\omega = E_k^\omega \cap \hat{E}_k$ and let

$$A_k^\omega = \{t \in \hat{E}_k^\omega \mid \mathsf{next}(t) > 6\log_{1/\gamma}\ell_k\}$$

where $\mathsf{next}(t) = \min\left(\{t' > t \mid t' \in E_k \setminus E_k^\omega\} \cup \{\ell_{k+1}\}\right) - t'$. Intuitively, $\mathsf{next}(t)$ is the distance between stage t and the first future stage not in E_k^ω. Henceforth, A_k^ω is a set of stages in which the first future stage not in E_k^ω is at least $6\log_{1/\gamma}\ell_k$ far away. By Lemma D.4, with probability at least $(1-\frac{1}{\ell_k})$, the mechanism we offer in E_k^ω is static, which implies that the buyer has little incentive to misreport for stages in A_k^ω since she discounts the future. We formalize this intuition in Lemma D.6. For convenience, let

$$\tau^{\Delta_k,\lambda_k}(k) = 4\sqrt{c_a}\cdot\ell_k^{\frac{1}{2}(1+\omega)}\sqrt{\log\ell_k} + 5c_a\sqrt{\frac{\Delta_k}{(1-\gamma)\lambda_k}}\cdot\ell_k + 6\ell_k^\omega\log_{1/\gamma}\ell_k.$$

Lemma D.6. *In $B^{hybrid}(\hat{F}_{(1,T)},\lambda_{(1,K)},\omega,\tau)$, if $\tau(k)\geq\tau^{\Delta_k,\lambda_k}(k)$ then with probability at least $(1-\frac{1}{\ell_k})$, for all $t\in A_k^\omega$, $\eta_t\leq O\left(\frac{1}{\sqrt{\lambda_k}\cdot\ell_k^2}\right)$ and moreover,*

$$\psi_{\ell_k}^s(0,\hat{F}_{(\ell_k,\ell_{k+1}-1)};F_{(\ell_k,\ell_{k+1}-1)}) \geq \frac{1}{3}\left(\sum_{t\in E_k^\omega}\mathbb{E}_{v_t\sim\hat{F}_t}[v_t]-\tau(k)\right) - O(\ell_k^\omega).$$

Proof. First, plugging $\delta = \tau(k) - 5c_a\sqrt{\frac{\Delta_k}{(1-\gamma)\lambda_k}}\cdot\ell_k - 6\ell_k^\omega\log_{1/\gamma}\ell_k - \frac{1}{3}\ell_k^\omega$ into Lemma D.4, then conditioned on the assumption that the buyer reports truthfully according to $\hat{F}_{(1,T)}$, we have

$$\begin{aligned}
&\Pr\left[\mathrm{sbal}_t < 5c_a\sqrt{\frac{\Delta_k}{(1-\gamma)\lambda_k}}\cdot\ell_k + 6\ell_k^\omega\log_{1/\gamma}\ell_k + \frac{1}{3}\ell_k^\omega\right]\\
\leq\ &\Pr\left[\mathrm{sbal}_t < y_{t^*(k)} - \left(\tau(k) - 5c_a\sqrt{\frac{\Delta_k}{(1-\gamma)\lambda_k}}\cdot\ell_k - 6\ell_k^\omega\log_{1/\gamma}\ell_k - \frac{1}{3}\ell_k^\omega\right)\right]\\
\leq\ &\exp\left(-\frac{\left(5\sqrt{c_a}\cdot\ell_k^{\frac{1}{2}(1+\omega)}\sqrt{\log\ell_k}-\frac{1}{3}\ell_k^\omega\right)^2}{4c_a\ell_k^{1+\omega}}\right)\leq\frac{1}{\ell_k^2}.
\end{aligned}$$

Applying a union bound, we have

$$\Pr\left[\exists t\in A_k^\omega,\ \mathrm{sbal}_t < 5c_a\sqrt{\frac{\Delta_k}{(1-\gamma)\lambda_k}}\cdot\ell_k + 6\ell_k^\omega\log_{1/\gamma}\ell_k + \frac{1}{3}\ell_k^\omega\right]\leq\frac{1}{\ell_k}$$

conditioned on the assumption that the buyer reports truthfully according to $\hat{F}_{(1,T)}$. If the buyer misreports under $F_{(1,T)}$, then by Proposition D.2 and Assumption 1, we have

$$\Pr\left[\exists t\in A_k^\omega,\ \mathrm{sbal}_t < 6\ell_k^\omega\log_{1/\gamma}\ell_k + \frac{1}{3}\ell_k^\omega\right]\leq\frac{1}{\ell_k}.$$

By the definition of A_k^ω, for $t'\in\left[t+1, t+6\log_{1/\gamma}\ell_k\right]$, the mechanism at stage t' is the same if $\mathrm{sbal}_{t'}\geq\frac{1}{3}\ell_k^\omega$, which implies that $3\mathrm{sbal}_{t'}\geq a_{t'}\geq\mathbb{E}_{v_t\sim\hat{F}_t}[v_t]$. Moreover, notice that at stage t', since $a_{t'}\leq\ell_k^\omega$, the decrement of sbal is at most ℓ_k^ω. Therefore, we have

$$\Pr\left[\exists t\in A_k^\omega, \exists t < t'\leq t+6\log_{1/\gamma}\ell_k,\ \mathrm{sbal}_t < \frac{1}{3}\ell_k^\omega\right]\leq\frac{1}{\ell_k}.$$

As a result, for a buyer who misreports at stage $t\in A_k^\omega$, she can only earn benefit from stages at least $6\log_{1/\gamma}\ell_k$ away in the future. Thus, the amount of misreport m_t must satisfy

$$\lambda_k\cdot\frac{m_t^2}{2a_t}\leq\sum_{t'\geq t+6\log_{1/\gamma}\ell_k}\gamma^{t'-t}a_{t'}\leq\frac{c_a}{(1-\gamma)^2\cdot\ell_k^5}\Rightarrow m_t = O\left(\frac{1}{\sqrt{\lambda_k}\cdot\ell_k^2}\right).$$

For the revenue guarantee in phase k, notice that our analysis demonstrates that once $\mathrm{sw}_t \geq \tau(k) \geq \tau^{\Delta_k,\lambda_k}(k)$ for some t, even the buyer misreports, with probability at least $(1-\frac{1}{\ell_k})$, $\mathrm{sbal}_{t'} \geq \frac{1}{3}\ell_k^\omega$ for all $t' \geq t$ and $t' \in E_k^\omega$. Therefore, the mechanism can obtain revenue $\frac{1}{3}\mathbb{E}_{v_{t'}\sim\hat{F}_{t'}}[v_{t'}]$ for these stages. □

We are now ready to bound the estimation error of our learning policy (Section 4.1) in our robust hybrid non-clairvoyant bank account mechanism $B^{hybrid}(\hat{F}_{(1,T)}, \lambda_{(1,K)}, \omega, \tau)$.

Lemma D.7. *If $\tau(k) \geq \tau^{\Delta_k,\lambda_k}(k)$ and $\lambda_k \geq \ell_k^{-2}$, then with probability at least $(1-\frac{1}{\ell_k})$, $B^{hybrid}(\hat{F}_{(1,T)}, \lambda_{(1,K)}, \omega, \tau)$ is $\eta_{(1,T)}$-DIC for $F_{(1,T)}$ such that*

$$\sum_{t\in\hat{E}_k} \frac{\eta_t}{a_t} \leq \tilde{O}\left(\ell_k^{1-\omega}\right).$$

Proof. First of all, for $t \in A_k^\omega$, by Lemma D.6, we have

$$\sum_{t\in A_k^\omega} \frac{\eta_t}{a_t} \leq \sum_{t\in A_k^\omega} \eta_t = O\left(\frac{1}{\sqrt{\lambda_k}\cdot \ell_k}\right)$$

where the inequality is due to $a_t \geq 1$. As for $t \in \hat{E}_k^\omega \setminus A_k^\omega$, we simply apply the bound $\frac{\eta_t}{a_t} \leq 1$ since $\eta_t \leq a_t$. Moreover, by the definition of A_k^ω, for $t \in \hat{E}_k^\omega$, there are at most $|E_k \setminus E_k^\omega| \cdot 6\log_{1/\gamma}\ell_k$ stages not in A_k^ω. In addition, by Assumption 1, we have $|E_k \setminus E_k^\omega| = O(\ell_k^{1-\omega})$. Therefore, we have

$$\sum_{t\in\hat{E}_k^\omega\setminus A_k^\omega} \frac{\eta_t}{a_t} \leq |\hat{E}_k^\omega \setminus A_k^\omega| = \tilde{O}\left(\ell_k^{1-\omega}\right).$$

Finally, for stages in $\hat{E}_k \setminus \hat{E}_k^\omega$, we have $|\hat{E}_k \setminus \hat{E}_k^\omega| = O(\ell_k^{1-\omega})$, and therefore,

$$\sum_{t\in\hat{E}_k\setminus\hat{E}_k^\omega} \frac{\eta_t}{a_t} \leq |\hat{E}_k \setminus \hat{E}_k^\omega| = \tilde{O}\left(\ell_k^{1-\omega}\right).$$

We conclude the proof of the lemma by summing over all three cases. □