
Prior-Free Dynamic Auctions with Low Regret Buyers

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 We study the problem of how to repeatedly sell to a buyer running a no-regret,
2 mean-based algorithm. Previous work [Braverman et al., 2018] shows that it is
3 possible to design effective mechanisms in such a setting that extract almost all
4 of the economic surplus, but these mechanisms require the buyer’s values each
5 round to be drawn independently and identically from a fixed distribution. In this
6 work, we do away with this assumption and consider the *prior-free setting* where
7 the buyer’s value each round is chosen adversarially (possibly adaptively).

8 We show that even in this prior-free setting, it is possible to extract a $(1 - \varepsilon)$ -
9 approximation of the full economic surplus for any $\varepsilon > 0$. The number of options
10 offered to a buyer in any round scales independently of the number of rounds T
11 and polynomially in ε . We show that this is optimal up to a polynomial factor;
12 any mechanism achieving this approximation factor, even when values are drawn
13 stochastically, requires at least $\Omega(1/\varepsilon)$ options. Finally, we examine what is
14 possible when we constrain our mechanism to a natural auction format where
15 overbidding is dominated. Braverman et al. [2018] show that even when values are
16 drawn from a known stochastic distribution supported on $[1/H, 1]$, it is impossible
17 in general to extract more than $O(\log \log H / \log H)$ of the economic surplus. We
18 show how to achieve the same approximation factor in the *prior-independent* setting
19 (where the distribution is unknown to the seller), and an approximation factor of
20 $O(1/\log H)$ in the *prior-free setting* (where the values are chosen adversarially).

21 1 Introduction

22 Revenue optimal auction design in settings where a seller interacts repeatedly with a buyer (like in
23 the sale of Internet ads) is a problem of high commercial relevance. The promise of *dynamic auctions*,
24 that allow the linking of buyers’ decisions across time, is the significantly higher revenue they can
25 achieve over running independent/decoupled auctions across time. The technical challenges that
26 dynamic auctions introduce, along with their practical impact has inspired a lot of recent work in this
27 area [Papadimitriou et al., 2016, Ashlagi et al., 2016, Mirrokni et al., 2018].

28 Traditionally, almost all work in dynamic mechanism design operates in the regime where the
29 players’ types (e.g. bidders’ values) are drawn stochastically from a fixed distribution. In many
30 situations this is far from a realistic assumption – for example, if the values of a buyer are modelled
31 as a distribution, this underlying distribution likely drifts over time and is also subject to shocks
32 determined by uncontrolled exogenous events. But this assumption is also in many ways critical: in a
33 dynamic mechanism in an adversarial setting, a fully rational buyer (who cares about the effect of his
34 current action on his future utility) would be unable to compute his future utility at any point of time
35 in the game and thus unable to meaningfully best-respond.

36 On the other hand, auctions for digital ads have become increasingly more complex over time. The
 37 design space of dynamic auctions, in which a buyer bids on many items over the course of many
 38 rounds, is very rich and has room for exceedingly complex auctions. A bidder may have difficulty
 39 behaving fully rationally in such an auction: the bidder may not have accurate priors for bidders,
 40 the bidder may not completely understand the mechanism, and finding an equilibrium might be
 41 computationally hard. Instead of acting fully rationally, a bidder might instead choose to try to learn
 42 how to bid over time, for example by using a no-regret learning algorithm. Recently, several streams
 43 of work (e.g. Agrawal et al. [2018], Braverman et al. [2018]) have explored the problem of how
 44 to design dynamic auctions for such bidders. In all cases these works assume, as is standard, that
 45 bidders' values are stochastically generated.

46 However, one intriguing feature of modelling a bidder as a learning agent is that it no longer restricts
 47 us to the stochastic setting – the actions taken by a learning algorithm are perfectly well-defined in
 48 (and ostensibly even designed for) the *prior-free* setting where values are drawn adversarially. This
 49 opens a wealth of questions of how to robustly design dynamic mechanisms that perform well in the
 50 worst-case against some class of learning agents. In this paper, we explore this question for one of the
 51 simplest problems in dynamic mechanism design: repeatedly selling a single item to a single buyer
 52 for T rounds.

53 We build off the setting of [Braverman et al., 2018], where they model the buyer as a learner running
 54 a mean-based low-regret algorithm. Intuitively, mean-based algorithms prefer to select actions that
 55 have performed historically well on average (it can be shown that many classic learning algorithms,
 56 like EXP3, Multiplicative Weights, and Follow-the-Perturbed-Leader, are all mean-based low-regret
 57 algorithms). In [Braverman et al., 2018], the authors show that surprisingly, when the buyer's values
 58 $v_t \in [0, 1]$ are drawn from a fixed distribution, it is possible to design a simple mechanism that obtains
 59 almost the full economic surplus (i.e., $\text{Val} = \mathbb{E}[\sum_t v_t]$) as revenue. Their mechanism, however, relies
 60 crucially on the fact that the buyer's values are drawn from the same distribution every round. In
 61 particular, it is straightforward to verify that there exist sequences of values for the buyer that result
 62 in this mechanism receiving asymptotically zero total revenue.

63 In this paper we design mechanisms for this problem in the *prior-free* setting, when the buyer's values
 64 $v_t \in [0, 1]$ are chosen adversarially (possibly adaptively). In the course of doing this, we aim to
 65 minimize the complexity of our mechanisms, measured in terms of the number of distinct options (i.e.
 66 “bids”) the mechanism presents to the bidder in any round. We call this quantity the option-complexity
 67 of the mechanism. Note that in mechanisms with high option-complexity it becomes harder to learn
 68 how to bid. If the option-complexity of the mechanism begins to scale with the number of rounds T ,
 69 this may even nullify any sort of low-regret or mean-based guarantee the learning algorithm has (it
 70 may not even be possible to explore all potential options).

71 **Upper bound in the adversarial setting:** We design a non-adaptive (i.e., does not use historical
 72 bids/allocation/prices) *option-based mechanism* that yields a revenue of $\text{Val} - O(\varepsilon T)$ with
 73 $O\left(\frac{\ln(1/\varepsilon)}{\varepsilon^3}\right)$ options, where the instance (v_1, \dots, v_T) is chosen by a (possibly adaptive) adversary
 74 and Val is the total economic surplus defined by $\text{Val} = \sum_{t=1}^T v_t$.

75 **Lower bound in the stochastic (and hence adversarial) setting:** We show that even if values are
 76 drawn from an unknown stochastic distribution (i.e. in every round the buyer's value was drawn
 77 independently from some distribution \mathcal{D}), any non-adaptive option-based mechanism needs to offer
 78 at least $\Omega(1/\varepsilon)$ options to attain a $\text{Val} - O(\varepsilon T)$ revenue. This implies the option-complexity of our
 79 algorithm is tight up to a polynomial factor in $1/\varepsilon$.

80 **Upper bound in the stochastic setting with unknown distribution via critical mechanisms:** Fi-
 81 nally, although our mechanisms have relatively low option-complexity, they can still appear unnatural
 82 and complex. We examine what is possible by further restricting our mechanisms to *critical mecha-*
 83 *nisms* [Braverman et al., 2018], by imposing the desiderata of individual rationality, monotonicity of
 84 price and allocation in bid, and overbidding being dominated (see Section 2). Braverman et al. [2018]
 85 show that the seller can use a critical mechanism to extract a good revenue but not all of surplus,
 86 in particular showing the seller can always guarantee revenue equal to an $O\left(\frac{\log \log H}{\log H}\right)$ fraction of
 87 total economic surplus when buyer values lie in the interval $[\frac{1}{H}, 1]$, and that this competitive ratio
 88 is tight. This critical mechanism requires full knowledge of the value distribution \mathcal{D} . We design a

89 critical mechanism that achieves this same approximation factor, but in a *prior-independent setting*
90 *where the distribution \mathcal{D} is unknown*. In addition, we show that it is possible to achieve a slightly
91 worse competitive ratio of $O(\frac{1}{\log H})$ in the *prior-free* (adversarial values) setting by adapting existing
92 prior-free mechanisms for the single-shot instance of this problem.

93 We emphasize that all the mechanisms we present are non-adaptive (i.e. allocation and payment rules
94 at all times are fixed starting at the beginning of the protocol, and are not functions of the historical
95 bids/allocations/payments) as in [Braverman et al., 2018].

96 1.1 Related Work

97 Our work is closely related to the dynamic mechanism design literature, such as [Balseiro et al.,
98 2017, Liu and Psomas, 2017, Agrawal et al., 2018, Mirrokni et al., 2018, Balseiro et al., 2019],
99 which studies how to sell items online to a fixed set of buyers, whose valuations are drawn from
100 some distributions. However, the buyers are fully strategic such that their bidding strategies aim to
101 maximize their accumulative utility throughout the auction. In contrast to these works, we model the
102 buyers are running no-regret algorithms.

103 No-regret algorithms were first introduced in the context of the multi-armed bandit problem and have
104 been widely studied (see Bubeck et al. [2012] for a survey). Applications of low-regret learning
105 to algorithmic game theory are widespread (e.g. [Roughgarden, 2012, Syrgkanis and Tardos,
106 2013, Nekipelov et al., 2015, Daskalakis and Syrgkanis, 2016]). Most applications to dynamic
107 auction design are from the perspective of seller attempting to learn the optimal auction [Cole and
108 Roughgarden, 2014, Devanur et al., 2016, Morgenstern and Roughgarden, 2015, 2016, Gonczarowski
109 and Nisan, 2017, Cai and Daskalakis, 2017, Dudík et al., 2017] but some recent papers have studied
110 the problem of applying learning algorithms to the problem of learning how to bid Feng et al. [2018],
111 Balseiro et al. [2018].

112 As pointed out in a seminal empirical work [Nekipelov et al., 2015], bidders' behavior on Bing is
113 largely consistent with a no-regret learning algorithm, which motivates a question of designing a
114 dynamic mechanism against such a no-regret learning behavior. Braverman et al. [2018] initiated the
115 study of mechanism design against a no-regret buyer when the buyer's valuations are drawn from a
116 fixed and known distribution. In contrast to their works, we design mechanisms against a no-regret
117 buyer in a prior-free / prior-independent setting.

118 2 Model and Preliminaries

119 Our setting is similar to the setting considered in [Braverman et al., 2018]: we consider a multiple
120 round auction where every round a seller attempts to sell an item to a buyer running a low-regret (in
121 fact, mean-based) algorithm to learn how to bid.

122 Specifically, we consider a T -round auction with one buyer and one seller. In each round t , there is
123 one item for sale. At the beginning of this round, the buyer learns his private valuation $v_t \in \mathcal{V} \subseteq [0, 1]$
124 for this item. These valuations v_t can be generated in one of two ways: (1) *Adversarial*, where v_t is
125 chosen arbitrarily by a (possibly adaptive) adversary; and (2) *Stochastic*, where v_t is independently
126 drawn from some distribution \mathcal{D} . This distribution \mathcal{D} may either be known to the seller or not and we
127 will mostly consider the case where \mathcal{D} is unknown to the seller (i.e., the prior-independent setting).

128 For simplicity, we assume the values v_t belong to a finite set \mathcal{V} . This is solely for the purpose of
129 providing a finite number of different contexts to the buyer's learning algorithm and otherwise does
130 not affect our mechanism at all.

131 To measure the performance of our mechanisms, we compare the revenue extracted from the mecha-
132 nism to the *welfare*, the total value the buyer assigns to all the items.

133 **Definition 1.** The *welfare* $\text{Val}(v_1, \dots, v_T)$ is equal to $\sum_{t=1}^T v_t$.

134 The welfare clearly provides an upper bound on the revenue of our mechanisms. In cases where v_t is
135 drawn from some distribution \mathcal{D} , we will write $\text{Val}(\mathcal{D}) = \mathbb{E}_{x \sim \mathcal{D}}[x] \cdot T$ to denote the expected welfare
136 under this distribution.

137 **2.1 Mechanism format**

138 Since the buyer is running a learning algorithm, it is especially important to specify the manner of
 139 interaction between the buyer and the seller. We consider two classes of mechanisms for the seller:
 140 *option-based mechanisms*, and *critical mechanisms*.

141 In an option-based mechanism, the seller offers the buyer K options (labeled 1 through K) each round.
 142 If the buyer selects choice i at time t , the buyer receives the item with probability $a_{i,t}$ and pays a price
 143 $p_{i,t}$. A natural measure of complexity for such mechanisms is the number of options K presented to
 144 the buyer, which we refer to as the *option-complexity* of the mechanism. Limiting this complexity is
 145 especially important when interacting with learning agents, as they require some time to explore each
 146 option (indeed, as K approaches T , the low regret guarantees of the learning algorithms we consider
 147 become vacuous).

148 Critical mechanisms [Braverman et al., 2018] are a subset of option-based mechanisms that are
 149 *reasonable*. In a critical mechanism, the buyer interacts with the mechanism each round by submitting
 150 a bid b . The buyer then receives the item with probability $a_t(b)$ and pays a price $p_t(b)$. These
 151 allocation/payment rules should satisfy the following properties:

- 152 • *Individual rationality*: $p_t(b)$ satisfies $p_t(b) \leq b \cdot a_t(b)$, i.e. a bidder should never be charged
 153 more than their bid in expectation.
- 154 • *Monotonicity*: $p_t(b)$ and $a_t(b)$ are weakly increasing in b , i.e., submitting a higher bid should
 155 never decrease the winning probability or the payment.
- 156 • *Overbidding is dominated*: If the bidder’s value is v , it should never be in their interest to
 157 submit a bid $b > v$, i.e. if $b > v$ then $v \cdot a_t(v) - p_t(v) > v \cdot a_t(b) - p_t(b)$ for all t .

158 In both option-based mechanisms and critical mechanisms, we assume that the seller is completely
 159 non-adaptive and sets the allocation / payment functions at the beginning of the protocol.

160 **2.2 No-regret learner**

161 In contrast to a utility-maximizing buyer, we consider a buyer who follows some no-regret strategy
 162 for the multi-armed bandit problem. In a classic multi-armed bandit problem with T rounds, the
 163 learner (in our setting, the buyer) selects one of K options (‘arms’) on round t and receives a reward
 164 $r_{i,t} \in [0, 1]$ if he selects option i . The rewards can be chosen adversarially and the learner’s objective
 165 is to maximize his total reward.

166 Let i_t be the arm pulled by the learner at round t . The *regret* for a (possibly randomized) strategy
 167 \mathcal{A} is defined as the difference between performance of the strategy \mathcal{A} and the best arm: $\text{Reg}(\mathcal{A}) =$
 168 $\max_i \sum_{t=1}^T r_{i,t} - r_{i_t,t}$. A strategy \mathcal{A} for the multi-armed bandit problem is *no-regret* if the expected
 169 regret is sub-linear in T , i.e., $\mathbb{E}[\text{Reg}(\mathcal{A})] = o(T)$. In addition to the *bandits* setting in which the
 170 learner only learns the reward of the arm he pulls, our results also apply to the *experts* setting in
 171 which the learner can learn the rewards of all arms for every round. In our setting, the buyer learns
 172 $a_{i,t}$ and $p_{i,t}$, allowing him to compute the reward as $r_{i,t} = a_{i,t} \cdot v_t - p_{i,t}$. Moreover, the buyer has
 173 the additional information of her value v_t , and thus is in fact facing a *contextual bandit* problem.

174 **Contextual Bandits** In a contextual bandit problem, the learner is additionally provided a *context*
 175 c_t from a finite set \mathcal{C} . The reward of pulling arm i under context c on round t is now given by $r_{i,t}(c)$.
 176 In the experts setting, the learner can obtain the values of $r_{i,t}(c_t)$ for all arms i under context c_t after
 177 round t , while the learner only learns $r_{i_t,t}(c_t)$ for the arm i he pulls in the bandits setting.

178 The notion of regret for a strategy \mathcal{M} can be easily extended to the contextual bandit problem by
 179 considering the best context-specific policy π : $\text{Reg}(\mathcal{M}) = \max_{\pi: \mathcal{C} \rightarrow [K]} \sum_{t=1}^T r_{\pi(c_t),t}(c_t) - r_{i_t,t}(c_t)$.
 180 As before, a strategy \mathcal{M} is no-regret if $\mathbb{E}[\text{Reg}(\mathcal{M})] = o(T)$. When the size of the context \mathcal{C} is a
 181 constant with respect to T , a no-regret strategy \mathcal{M} for the contextual bandits can be simply constructed
 182 from a no-regret strategy \mathcal{A} for the classic bandit problem: maintain a separate instance of \mathcal{A} for
 183 every context $c \in \mathcal{C}$ [Bubeck et al., 2012].

184 Among no-regret strategies, we are interested in a special class of *mean-based* strategies:

185 **Definition 2** (Mean-based Strategy). Let $\sigma_{i,t}(c) = \sum_{s=1}^t r_{i,s}(c)$ be the cumulative rewards for
 186 pulling arm i under context c for the first t rounds. A strategy is γ -mean-based if whenever $\sigma_{i,t}(c_t) <$

187 $\sigma_{j,t}(c_t) - \gamma T$, the probability for the strategy to pull arm i on round t is at most γ . A strategy is
 188 mean-based if it is γ -mean-based with $\gamma = o(1)$.

189 Intuitively, mean-based strategies are strategies that will pick the arm that historically performs
 190 the best. Braverman et al. [2018] shows that many no-regret algorithms are mean-based, including
 191 commonly used variants of EXP3 (for the bandits setting), the Multiplicative Weights algorithm (for
 192 the experts setting) and the Follow-the-Perturbed-Leader algorithm (for the experts setting).

193 3 Option-based Mechanisms

194 In this section, we demonstrate a mechanism that can extract full welfare from a mean-based no-regret
 195 learner even when the values are chosen adversarially.

196 3.1 Warm-up: Extracting Full Welfare for $\mathcal{V} = \{1, 2\}$

197 Consider an additive approximation target $\varepsilon > 0$. It is without loss of generality to consider the
 198 case with $2(1 - \varepsilon) > 1$: when $2(1 - \varepsilon) \leq 1$, the seller can simply implement a scheme with only
 199 one option that always allocates the item and charges a payment $2(1 - \varepsilon)$. We design a menu-based
 200 mechanism with $K = \lceil \frac{\log \varepsilon}{\log(1-\varepsilon)} \rceil + 1$ choices in addition to the null choice in which the buyer
 201 receives and pays nothing for the entire time horizon. For the 0-th option, the buyer receives the item
 202 with probability $a_{0,t} = 1$ and pays a price $p_{0,t} = 2(1 - \varepsilon)$ for all t . As for the remaining $K - 1$
 203 options, let $\kappa_i = \frac{\varepsilon}{(1-\varepsilon)^{i-1}} T$. We will divide the timeline of the i -th option with $1 \leq i \leq K$ into five
 204 sessions (see Table 1 for details).

205 For convenience, let $S_i = (\kappa_i, \kappa_{i+1}]$. Intuitively, the i -th option is *active* when $t \in S_i$, which spans
 206 $L_i = \kappa_{i+1} - \kappa_i = \frac{\varepsilon^2}{(1-\varepsilon)^i} T$ rounds. Among these L_i rounds, the item is always allocated to the buyer
 207 with probability 1 while the payment changes in a way such that: the payment for the first εL_i rounds
 208 is 0, the payment for the last εL_i rounds is 2, and the payment for the remaining rounds is 1.

Session	Start Time	End Time	Allocation Prob.	Payment
\emptyset_1	0	κ_i	0	0
0	κ_i	$\left(\kappa_i + \frac{\varepsilon^3}{(1-\varepsilon)^i}\right) T$	1	0
1	$\left(\kappa_i + \frac{\varepsilon^3}{(1-\varepsilon)^i}\right) T$	$\left(\kappa_{i+1} - \frac{\varepsilon^3}{(1-\varepsilon)^i}\right) T$	1	1
2	$\left(\kappa_{i+1} - \frac{\varepsilon^3}{(1-\varepsilon)^i}\right) T$	κ_{i+1}	1	2
\emptyset_2	κ_{i+1}	T	0	0

Table 1: Construction of the i -th option

209 Assume the buyer is running a γ -mean-based algorithm. To analyze the revenue guarantee of our
 210 mechanism, we consider an arbitrary sequence of valuations (v_1, \dots, v_T) and $\text{Val} = \sum_t v_t$. The
 211 high level idea behind this construction is that for the high valuations, i.e. $v_t = 2$, the utility $\sigma_{i,t}(2)$
 212 does keep increasing as t increases for *the high option* ($i = 0$) while for *the low options* ($i > 0$),
 213 it only increases within the active period S_i . Therefore, with sufficiently large t , one could expect
 214 that $\sigma_{0,t}(2) > \sigma_{i,t}(2)$ for all $i > 0$. As for $v_t = 1$, the buyer does not play the high option since its
 215 payment is too high and we argue that the buyer will play the i -th option if $t \in S_i$.

216 **High valuation** Assume that $v_t = 2$. First notice that the cumulative utility for playing the 0-th
 217 option is $\sigma_{0,t}(2) = \varepsilon t \cdot 2$. Suppose $t \in S_{i^*}$ for some i^* . For $i < i^*$, the active period of the i -th
 218 option with $i < i^*$ is already past and the cumulative utility for playing the i -th option is at most

$$\sigma_{i,t}(2) \leq L_i \cdot 2 = \frac{\varepsilon^2}{(1-\varepsilon)^i} T \cdot 2 \leq \frac{\varepsilon^2}{(1-\varepsilon)^{i^*-1}} T \cdot 2 = \varepsilon \cdot \kappa_{i^*} \cdot 2 = \sigma_{0,t}(2) - \varepsilon \cdot (t - \kappa_{i^*}) \cdot 2$$

219 As for the i^* -th option, we have $\sigma_{i^*,t}(2) \leq (t - \kappa_{i^*}) \cdot 2 = \sigma_{0,t}(2) - (\kappa_{i^*} - (1-\varepsilon)t) \cdot 2$.

220 Moreover, for any i -th option with $i > i^*$, we simply have $\sigma_{i,t}(2) = 0$. Therefore, the buyer
 221 with valuation $v_t = 2$ for $t \in S_{i^*}$ will play the 0-th option with probability at least $1 - K\gamma$

222 when $\varepsilon t \cdot 2 > \gamma T$, $\varepsilon \cdot (t - \kappa_{i^*}) \cdot 2 > \gamma T$, and $(\kappa_{i^*} - (1 - \varepsilon)t) \cdot 2 > \gamma T$, which implies that
 223 $\kappa_{i^*} + \frac{\gamma}{2\varepsilon} \cdot T < t < \kappa_{i^*+1} - \frac{\gamma}{2(1-\varepsilon)} \cdot T$. Therefore, for each time period S_i with $1 \leq i \leq K$, there
 224 are at least $L_i - \left(\frac{\gamma}{2\varepsilon} + \frac{\gamma}{2(1-\varepsilon)}\right)T$ rounds where the buyer has probability at least $1 - K\gamma$ to play
 225 the 0-th option, which contributes $2(1 - \varepsilon)$ revenue per round. Therefore, the expected revenue loss
 226 from time period S_i is at most

$$2\varepsilon \cdot L_i + \left(\frac{\gamma}{2\varepsilon} + \frac{\gamma}{2(1-\varepsilon)}\right)T \cdot 2 + K\gamma \cdot L_i \cdot 2$$

227 where $2\varepsilon \cdot L_i$ is the revenue loss of charging $2(1 - \varepsilon)$ and $K\gamma \cdot L_i \cdot 2$ is the expected revenue loss
 228 from playing an option other than the 0-th option. Thus, the total expected revenue loss from the
 229 rounds when $v_t = 2$ is at most

$$(\varepsilon T) \cdot 2 + \sum_i \left[2\varepsilon \cdot L_i + \left(\frac{\gamma}{2\varepsilon} + \frac{\gamma}{2(1-\varepsilon)}\right)T \cdot 2 + K\gamma \cdot L_i \cdot 2 \right] = O(\varepsilon T)$$

230 where $(\varepsilon T) \cdot 2$ is the revenue loss from the first εT rounds.

231 **Low valuation** Assume that $v_t = 1$. First notice that after the first εT rounds, the cumulative utility
 232 for playing the 0-th option is $\sigma_{0,t}(1) = (1 - 2(1 - \varepsilon))t = -\Omega(T)$. Since there is a null arm that
 233 provides cumulative utility 0, the buyer's probability of playing the 0-th option is at most γ .

234 Suppose $t \in S_{i^*}$ for some i^* . From our construction of the i -th option for any $i \neq i^*$, the buyer's
 235 cumulative utility of playing the i -th option is exactly 0: the buyer's utility gain is 0 in from session
 236 \emptyset_1 , \emptyset_2 , and 1, while her utility gain from session 0 is exactly cancelled out with his utility loss from
 237 session 2. As for the i^* -th option, we have

$$\sigma_{i^*,t}(1) = \begin{cases} t - \kappa_i & \text{for } t \text{ in session 0} \\ \frac{\varepsilon^3}{(1-\varepsilon)^{i^*}}T & \text{for } t \text{ in session 1} \\ \kappa_{i+1} - t & \text{for } t \text{ in session 2} \end{cases}$$

238 Therefore, once $\kappa_i + \gamma T < t < \kappa_{i+1} - \gamma T$, the buyer with $v_t = 1$ will play the i^* -th option with
 239 probability $1 - K\gamma$. Therefore, the expected revenue loss within the time period S_i is $2\gamma T + K\gamma \cdot L_i$,
 240 where $K\gamma \cdot L_i$ is the expected revenue loss from playing an option other than the i^* -th option. Thus,
 241 the total revenue loss from the rounds with $v_t = 1$ is at most $\varepsilon T + \sum_{i=1}^K K\gamma \cdot L_i = O(\varepsilon T)$ where
 242 εT is the revenue loss from the first εT rounds.

243 3.2 Extracting Full Welfare for $\mathcal{V} = \{1, \dots, H\}$

244 We provide an option-based mechanism with $K = H \cdot \lceil \frac{3H^2}{\varepsilon} \rceil$ options that achieves an additive revenue
 245 loss $O(\ln H \cdot \varepsilon T)$ for $\mathcal{V} = \{1, \dots, H\}$. As usual, we assume that there is always a null choice in which
 246 the buyer receives and pays nothing for the entire time horizon. For convenience, let $G_i = \sum_{\tau=1}^i \frac{1}{\tau}$ be
 247 the sum of the harmonic series up to i and $\alpha = 1 - \frac{1}{3H}$. Moreover, $\kappa_{i,j} = (G_H + 2\alpha) \cdot \frac{\varepsilon T}{H} + (j-1) \cdot \frac{\varepsilon T}{3H^2}$
 248 where $i \in \mathcal{V}$ and $1 \leq j \leq \lceil \frac{3H^2}{\varepsilon} \rceil$. Although $\kappa_{i,j}$ only depends on j , we still use the notation $\kappa_{i,j}$ for
 249 clarity. We will divide the timeline of the (i, j) -th option into five sessions (see Table 2).

Session	Start Time	End Time	Allocation Prob.	Payment
<i>init</i>	0	$\alpha \cdot \frac{\varepsilon T}{H}$	0	i
0	$\alpha \cdot \frac{\varepsilon T}{H}$	$\kappa_{i,j} - (G_i + \alpha) \cdot \frac{\varepsilon T}{H}$	0	0
<i>ready</i>	$\kappa_{i,j} - (G_i + \alpha) \cdot \frac{\varepsilon T}{H}$	$\kappa_{i,j}$	1	0
1	$\kappa_{i,j}$	$\kappa_{i,j+1}$	1	i
\emptyset	$\kappa_{i,j+1}$	T	0	H

Table 2: Construction of the (i, j) -th option

250 Intuitively, the (i, j) -th option starts with a *init* session in which it does not allocate the item but
 251 charges a payment i , followed by a 0 session in which the option allocates and charges nothing.
 252 Therefore, the buyer will not play the (i, j) -th option before its *ready* session. In the *ready* session,

253 the option allocates the item for free while in the 1 session, the option allocates the item with a
 254 payment i . The objective of our construction is to try to ensure that if $v_t = i$ for $t \in (\kappa_{i,j}, \kappa_{i,j+1}]$,
 255 then the buyer will play the option (i, j) , which generates revenue i .

256 Assume the buyer is running a γ -mean-based algorithm. To analyze the revenue guarantee of our
 257 mechanisms, we consider an arbitrary sequence of valuations (v_1, \dots, v_T) and $\text{Val} = \sum_t v_t$.

258 **Lemma 3.** *If $t \in (\kappa_{i,j} + \gamma T, \kappa_{i,j+1} - \gamma T]$, then for any option (i', j') with $i' \neq i$ or $j' \neq j$,
 259 $\sigma_{(i,j),t}(i) - \sigma_{(i',j'),t}(i) > \gamma T$.*

260 Therefore, for $v_t = i$ with $t \in (\kappa_{i,j} + \gamma T, \kappa_{i,j+1} - \gamma T]$, the buyer will play option (i, j) with
 261 probability at least $1 - K\gamma$, which generates revenue i per round. Thus, the revenue loss is at most

$$H \cdot (G_H + 2\alpha) \cdot \frac{\varepsilon T}{H} + H \cdot 2\gamma T \cdot K + K\gamma \cdot H \cdot T = O(\ln H \cdot \varepsilon T)$$

262 where $H \cdot (G_H + 2\alpha) \cdot \frac{\varepsilon T}{H}$ is the revenue loss for the first $\max_i \kappa_{i,1} = (G_H + 2\alpha) \cdot \frac{\varepsilon T}{H}$ rounds,
 263 $H \cdot 2\gamma T \cdot K$ is the revenue loss for $t \in (\kappa_{i,j}, \kappa_{i,j} + \gamma T]$ or $t \in (\kappa_{i,j+1} - \gamma T, \kappa_{i,j+1}]$, and $K\gamma \cdot H \cdot T$
 264 is the revenue loss from playing an undesired option.

265 **Theorem 4.** *If the buyer with $\mathcal{V} = \{1, 2, \dots, H\}$ is running a mean-based algorithm, for any
 266 constant $\varepsilon > 0$, there exists a non-adaptive option-based mechanism with $O(\frac{H^3 \ln H}{\varepsilon})$ options for the
 267 seller which obtains revenue at least $\text{Val} - O(\varepsilon T)$.*

268 3.3 Extracting Full Welfare for $\mathcal{V} \subseteq [0, 1]$

269 Let ε be parameter for the target additive revenue loss $O(\varepsilon T)$. For ease of presentation, we will
 270 rescale \mathcal{V} to $[0, H]$ such that $H = 1/\varepsilon$, and thus, it suffices to show that we can obtain $O(T)$ loss in
 271 the scaled version. First notice that it suffices to consider $\mathcal{V} \subseteq [1, H]$ since for all valuations less than
 272 1, we will suffer revenue loss at most 1 from each of them.

273 **Lemma 5.** *Consider v_t such that $i < v_t < i + 1$ and $t \in (\kappa_{i,j}, \kappa_{i,j+1}]$. Then, for any option (i', j')
 274 with $i' \notin \{i, i + 1\}$ or $j' > j$, $\max\{\sigma_{(i,j),t}(v_t), \sigma_{(i+1,j),t}(v_t)\} - \sigma_{(i',j'),t}(v_t) > \gamma T$.*

275 Therefore, with probability at least $1 - K\gamma$, the buyer satisfying the requirement of Lemma 5 will
 276 play either option (i, j') or option $(i + 1, j')$ with $j' \leq j$. Recall that it is in fact that $\kappa_{i,j} = \kappa_{i+1,j}$
 277 for all i . Therefore, if the buyer plays option $(i + 1, j)$, it will generate revenue $i + 1$ since option
 278 $(i + 1, j)$ is also in its 1 session. Moreover, if the buyer plays option (i, j') or $(i + 1, j')$ with $j' < j$,
 279 then the option is already in its \emptyset session and the buyer needs to pay H .

280 Thus, the revenue loss from v_t is at most 1. Applying a similar argument as in Section 3.2, we can
 281 conclude that the expected revenue loss is $O(T)$. Rescale it back to $\mathcal{V} = [0, 1]$, we have

282 **Theorem 6.** *If the buyer with $\mathcal{V} \subseteq [0, 1]$ is running a mean-based algorithm, for any constant $\varepsilon > 0$,
 283 there exists a non-adaptive option-based mechanism with $O(\frac{\ln 1/\varepsilon}{\varepsilon^3})$ options for the seller which
 284 obtains revenue at least $\text{Val} - O(\varepsilon T)$.*

285 Meanwhile, we provide a lower-bound on the option-complexity, which implies the option-complexity
 286 of our algorithm is tight up to a polynomial factor in $\frac{1}{\varepsilon}$.

287 **Theorem 7.** *If the buyer with $\mathcal{V} \subseteq [0, 1]$ is running a mean-based algorithm, an option-based
 288 mechanism, which obtains expected revenue at least $\text{Val} - O(\varepsilon T)$, must have $\Omega(\frac{1}{\varepsilon})$ options.*

289 4 Critical mechanisms

290 In this section we examine what the seller can accomplish when restricted to a critical mechanism.

291 With option-based mechanisms, we have shown in the previous section that it is possible to extract
 292 arbitrarily close to the full welfare even when the buyer's values are chosen adversarially. In contrast
 293 to this, Braverman et al. [2018] show that with a critical mechanism, it is impossible to achieve even
 294 a constant-factor approximation to the buyer's welfare, even when the buyer's values are drawn from
 295 a distribution known to the seller.

296 **Theorem 8** (Corollary C.13 of [Braverman et al., 2018]). *Let $R(\mathcal{D})$ be the maximum possible revenue
 297 a seller using a non-adaptive critical mechanism can achieve when the buyer's values are drawn*

298 *independently each round from distribution \mathcal{D} . Then the ratio $R(\mathcal{D})/\text{Val}(\mathcal{D})$ can grow arbitrarily*
299 *small. If \mathcal{D} is supported on an interval $[1, H]$, then this ratio can be as small as $O(\log \log H / \log H)$.*

300 In [Braverman et al. \[2018\]](#), the authors also demonstrate how to construct a simple mechanism which
301 achieves this maximum possible revenue (and hence this $O(\log \log H / \log H)$ competitive ratio to
302 the welfare), but their construction requires detailed knowledge of the distribution \mathcal{D} .

303 4.1 Values from an unknown distribution

304 We show that it is possible achieve this same competitive ratio to the welfare in the prior-independent
305 setting, where the seller does not know the distribution \mathcal{D} but only a range $[1, H]$ it is supported on. In
306 our mechanism, at each time t the seller specifies a reserve price $f(t)$, where f is a decreasing function
307 with range $[1, H]$ such that $f(t) = \max\left(\frac{1}{C} \cdot (1 - \eta - \frac{t}{T}), 1\right)$, where $\eta = (1 + \log H)^{-\varepsilon}$
308 and $C = \frac{1-\eta}{1+\log H}$ for $\varepsilon \in (0, 1)$. In each round, if the buyer bids above $f(t)$ they receive the item
309 and pay b ; otherwise, they do not receive the item and pay nothing. More formally, the allocation
310 and payment rules $(a_t(b), p_t(b))$ are defined as follows: if $b \geq f(t)$, then $p_t(b) = b$, and $a_t(b) = 1$;
311 otherwise, $p_t(b) = a_t(b) = 0$.

312 **Theorem 9.** *There is a non-adaptive critical mechanism for the seller which obtains expected revenue*
313 *at least $O(\log \log H / \log H)\text{Val}(\mathcal{D})$ from any buyer running a mean-based algorithm whose values*
314 *are drawn independently each round from some distribution \mathcal{D} supported on $[1, H]$. This mechanism*
315 *depends only on H and not on \mathcal{D} .*

316 4.2 Adversarial values

317 Additionally, when the buyer's values are drawn adversarially, we show that it is possible to achieve a
318 slightly worse competitive ratio of $O(1/\log H)$.

319 This follows naturally from the known fact that it is possible to achieve the same approximation
320 guarantee against a *strategic* buyer with value in $[1, H]$ playing a single-round version of this game
321 (see e.g. Chapter 6 of [Hartline \[2013\]](#)) – we simply show that if we run this mechanism every round,
322 mean-based buyers will learn to bid in the same manner as strategic buyers.

323 Our mechanism (equivalent to the mechanism presented in Theorem 6.5 of [Hartline \[2013\]](#)) is as
324 follows: for each b and for all t , we set the allocation probability $a_t(b) = (1 + \log b)/(1 + \log H)$
325 and the expected price charged to $p_t(b) = b/(1 + \log H)$. Note that this mechanism can also be
326 interpreted as a second-price auction where the seller draws a random reserve from the distribution
327 with a cumulative density function $F(r) = \frac{1+\log r}{1+\log H}$. It can be seen that for any $v \in [1, H]$, the
328 expected utility $U(v, b) = v \cdot a_t(b) - p_t(b)$ of bidding b with value v , is maximized when $b = v$. A
329 strategic buyer therefore will always bid their value, and pay $1/(1 + \log H)$ of their value in total.

330 Intuitively, the mean-based guarantee ensures that a mean-based buyer will (most of the time) choose
331 a bid close to v , and thus it contributes a similar amount of revenue as a strategic buyer.

332 **Theorem 10.** *There is a non-adaptive critical mechanism for the seller which obtains expected*
333 *revenue at least $O(1/\log H)\text{Val}$ from any buyer running a mean-based algorithm whose values are*
334 *adversarially set but lie in the interval $[1, H]$. This mechanism depends only on H and not on \mathcal{D} .*

335 5 Conclusion

336 In this work, we design mechanisms against a no-regret, mean-based buyer in prior-independent and
337 prior-free setting. We show that using option-based mechanism can extract almost full welfare in a
338 prior-free setting. For critical mechanisms, our mechanism in the prior-independent setting matches
339 the best-known guarantee for the prior-dependent setting in the literature, and we obtain a slightly
340 worse guarantee for the prior-free setting.

341 A nature direction for future work is to understand what can be achieved in an environment with
342 multiple learning buyers. Moreover, while both our works and [\[Braverman et al., 2018\]](#) focus on the
343 revenue guarantee of the seller against a no-regret buyer, it is interesting to understand what kinds of
344 the buyer's learning strategy can lead to a good utility performance. Furthermore, what combinations
345 of the buyer's learning strategy and the seller's mechanism can achieve a socially desirable outcome?

References

- 346
- 347 Shipra Agrawal, Constantinos Daskalakis, Vahab S. Mirrokni, and Balasubramanian Sivan. Robust
348 repeated auctions under heterogeneous buyer behavior. In *Proceedings of the 2018 ACM Conference*
349 *on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, page 171, 2018.
- 350 Itai Ashlagi, Constantinos Daskalakis, and Nima Haghpanah. Sequential mechanisms with ex-
351 post participation guarantees. In *Proceedings of the 2016 ACM Conference on Economics and*
352 *Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 213–214, 2016.
- 353 Santiago Balseiro, Negin Golrezaei, Mohammad Mahdian, Vahab Mirrokni, and Jon Schneider.
354 Contextual bandits with cross-learning. *arXiv preprint arXiv:1809.09582*, 2018.
- 355 Santiago R. Balseiro, Vahab S. Mirrokni, and Renato Paes Leme. Dynamic mechanisms with
356 martingale utilities. In *Proceedings of the 2017 ACM Conference on Economics and Computation,*
357 *EC '17, Cambridge, MA, USA, June 26-30, 2017*, page 165, 2017.
- 358 Santiago R. Balseiro, Vahab S. Mirrokni, Renato Paes Leme, and Song Zuo. Dynamic double
359 auctions: Towards first best. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on*
360 *Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 157–172,
361 2019.
- 362 Mark Braverman, Jieming Mao, Jon Schneider, and Matt Weinberg. Selling to a no-regret buyer. In
363 *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 523–538. ACM,
364 2018.
- 365 Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic
366 multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012.
- 367 Yang Cai and Constantinos Daskalakis. Learning multi-item auctions with (or without) samples. In
368 *FOCS*, 2017.
- 369 Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings*
370 *of the Forty-sixth Annual ACM Symposium on Theory of Computing, STOC '14*, pages 243–252,
371 New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2710-7. doi: 10.1145/2591796.2591867.
372 URL <http://doi.acm.org/10.1145/2591796.2591867>.
- 373 Constantinos Daskalakis and Vasilis Syrgkanis. Learning in auctions: Regret is hard, envy is
374 easy. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11*
375 *October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 219–228, 2016. doi:
376 10.1109/FOCS.2016.31. URL <https://doi.org/10.1109/FOCS.2016.31>.
- 377 Nikhil R. Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. The sample complexity of
378 auctions with side information. In *Proceedings of the Forty-eighth Annual ACM Symposium on*
379 *Theory of Computing, STOC '16*, pages 426–439, New York, NY, USA, 2016. ACM. ISBN
380 978-1-4503-4132-5. doi: 10.1145/2897518.2897553. URL [http://doi.acm.org/10.1145/](http://doi.acm.org/10.1145/2897518.2897553)
381 [2897518.2897553](http://doi.acm.org/10.1145/2897518.2897553).
- 382 Miroslav Dudík, Nika Haghtalab, Haipeng Luo, Robert E. Schapire, Vasilis Syrgkanis, and Jen-
383 nifer Wortman Vaughan. Oracle-efficient learning and auction design. In *FOCS*, 2017.
- 384 Zhe Feng, Chara Podimata, and Vasilis Syrgkanis. Learning to bid without knowing your value. In
385 *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 505–522. ACM,
386 2018.
- 387 Yannai A. Gonczarowski and Noam Nisan. Efficient empirical revenue maximization in single-
388 parameter auction environments. In *Proceedings of the 49th Annual ACM SIGACT Symposium*
389 *on Theory of Computing, STOC 2017*, pages 856–868, New York, NY, USA, 2017. ACM. ISBN
390 978-1-4503-4528-6. doi: 10.1145/3055399.3055427. URL [http://doi.acm.org/10.1145/](http://doi.acm.org/10.1145/3055399.3055427)
391 [3055399.3055427](http://doi.acm.org/10.1145/3055399.3055427).
- 392 Jason D Hartline. Mechanism design and approximation. *Book draft. October*, 122, 2013.

- 393 Siqi Liu and Christos-Alexandros Psomas. On the competition complexity of dynamic mechanism
394 design. *CoRR*, abs/1709.07955, 2017. URL <http://arxiv.org/abs/1709.07955>.
- 395 Vahab S. Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. Non-clairvoyant dynamic
396 mechanism design. In *Proceedings of the 2018 ACM Conference on Economics and Computation*,
397 *Ithaca, NY, USA, June 18-22, 2018*, page 169, 2018.
- 398 Jamie Morgenstern and Tim Roughgarden. The pseudo-dimension of near-optimal auctions. In
399 *Proceedings of the 28th International Conference on Neural Information Processing Systems*
400 *- Volume 1*, NIPS'15, pages 136–144, Cambridge, MA, USA, 2015. MIT Press. URL <http://dl.acm.org/citation.cfm?id=2969239.2969255>.
- 402 Jamie Morgenstern and Tim Roughgarden. Learning simple auctions. In Vitaly Feldman, Alexander
403 Rakhlin, and Ohad Shamir, editors, *29th Annual Conference on Learning Theory*, volume 49
404 of *Proceedings of Machine Learning Research*, pages 1298–1318, Columbia University, New
405 York, New York, USA, 23–26 Jun 2016. PMLR. URL <http://proceedings.mlr.press/v49/morgenstern16.html>.
- 407 Denis Nekipelov, Vasilis Syrgkanis, and Eva Tardos. Econometrics for learning agents. In *Proceedings*
408 *of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 1–18, New
409 York, NY, USA, 2015. ACM. ISBN 978-1-4503-3410-5. doi: 10.1145/2764468.2764522. URL
410 <http://doi.acm.org/10.1145/2764468.2764522>.
- 411 Christos Papadimitriou, George Pierrakos, Christos-Alexandros Psomas, and Aviad Rubinfeld. On
412 the complexity of dynamic mechanism design. In *Proceedings of the Twenty-seventh Annual*
413 *ACM-SIAM Symposium on Discrete Algorithms*, SODA '16, pages 1458–1475, Philadelphia, PA,
414 USA, 2016. Society for Industrial and Applied Mathematics. ISBN 978-1-611974-33-1. URL
415 <http://dl.acm.org/citation.cfm?id=2884435.2884535>.
- 416 Tim Roughgarden. The price of anarchy in games of incomplete information. In *Proceedings*
417 *of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 862–879, New York,
418 NY, USA, 2012. ACM. ISBN 978-1-4503-1415-2. doi: 10.1145/2229012.2229078. URL
419 <http://doi.acm.org/10.1145/2229012.2229078>.
- 420 Vasilis Syrgkanis and Eva Tardos. Composable and efficient mechanisms. In *Proceedings of the*
421 *Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 211–220, New
422 York, NY, USA, 2013. ACM. ISBN 978-1-4503-2029-0. doi: 10.1145/2488608.2488635. URL
423 <http://doi.acm.org/10.1145/2488608.2488635>.

424 **Appendix**

425 **A Appendix for Option-based Mechanism**

426 **A.1 Proof of Lemma 3**

Session	$\sigma_{(i,j),t}(v)$
<i>init</i>	$-t \cdot i$
0	$-\alpha \cdot \frac{\varepsilon T}{H} \cdot i$
<i>ready</i>	$-\alpha \cdot \frac{\varepsilon T}{H} \cdot i + (t - \kappa_{i,j} + (G_i + \alpha) \cdot \frac{\varepsilon T}{H}) \cdot v$
1	$-\alpha \cdot \frac{\varepsilon T}{H} \cdot i + (G_i + \alpha) \cdot \frac{\varepsilon T}{H} \cdot v + (t - \kappa_{i,j}) \cdot (v - i)$
\emptyset	$-\alpha \cdot \frac{\varepsilon T}{H} \cdot i + (G_i + \alpha) \cdot \frac{\varepsilon T}{H} \cdot v + \frac{\varepsilon T}{3H^2} \cdot (v - i) - (t - \kappa_{i,j+1})H$

Table 3: Cumulative utility of the (i, j) -th option

427 *Proof.* Notice that for a valuation $v_t = i$ with $t \in (\kappa_{i,j}, \kappa_{i,j+1}]$, its cumulative utility for playing
 428 option (i, j) is

$$\sigma_{(i,j),t}(i) = -\alpha \cdot \frac{\varepsilon T}{H} \cdot i + (G_i + \alpha) \cdot \frac{\varepsilon T}{H} \cdot i + (t - \kappa_{i,j}) \cdot (i - i) = G_i \cdot \frac{\varepsilon T}{H} \cdot i. \quad (1)$$

429 v_t 's cumulative utility for playing option (i', j') with $i' < i$ is at most

$$\begin{aligned} \sigma_{(i',j'),t}(i) &\leq -\alpha \cdot \frac{\varepsilon T}{H} \cdot i' + (G_{i'} + \alpha) \cdot \frac{\varepsilon T}{H} \cdot i + (t - \kappa_{i',j'}) \cdot (i - i') \\ &\leq -\alpha \cdot \frac{\varepsilon T}{H} \cdot i' + (G_{i'} + \alpha) \cdot \frac{\varepsilon T}{H} \cdot i + \frac{\varepsilon T}{3H^2} \cdot (i - i') \\ &= G_{i'} \cdot \frac{\varepsilon T}{H} \cdot i + \left(\alpha + \frac{1}{3H}\right) \cdot \frac{\varepsilon T}{H} \cdot (i - i') \\ &= G_{i'} \cdot \frac{\varepsilon T}{H} \cdot i + \frac{\varepsilon T}{H} \cdot (i - i') \end{aligned}$$

430 Taking the difference between $\sigma_{(i',j'),t}(i)$ and $\sigma_{(i,j),t}(i)$, we have

$$\begin{aligned} \sigma_{(i,j),t}(i) - \sigma_{(i',j'),t}(i) &\geq (G_i - G_{i'}) \cdot \frac{\varepsilon T}{H} \cdot i - \frac{\varepsilon T}{H} \cdot (i - i') \\ &= \left(\sum_{\tau=i'+1}^i \frac{i}{\tau} \right) \cdot \frac{\varepsilon T}{H} - \frac{\varepsilon T}{H} \cdot (i - i') \\ &= \left(i - i' + \sum_{\tau=i'+1}^i \left(\frac{i}{\tau} - 1 \right) \right) \cdot \frac{\varepsilon T}{H} - \frac{\varepsilon T}{H} \cdot (i - i') \\ &= \left(\sum_{\tau=i'+1}^i \left(\frac{i}{\tau} - 1 \right) \right) \cdot \frac{\varepsilon T}{H} > \gamma T \end{aligned}$$

431 In addition, v_t 's cumulative utility for playing option (i', j') with $i' > i$ is at most

$$\sigma_{(i',j'),t}(i) \leq \max \left\{ -\alpha \cdot \frac{\varepsilon T}{H} \cdot i' + (G_{i'} + \alpha) \cdot \frac{\varepsilon T}{H} \cdot i + (t - \kappa_{i',j'}) \cdot (i - i'), 0 \right\}$$

432 If the maximum is taken by 0, then it is clear that $\sigma_{(i,j),t}(i) - \sigma_{(i',j'),t}(i) > \gamma T$. On the other hand,

$$\begin{aligned} \sigma_{(i',j'),t}(i) &\leq -\alpha \cdot \frac{\varepsilon T}{H} \cdot i' + (G_{i'} + \alpha) \cdot \frac{\varepsilon T}{H} \cdot i + (t - \kappa_{i',j'}) \cdot (i - i') \\ &\leq -\alpha \cdot \frac{\varepsilon T}{H} \cdot i' + (G_{i'} + \alpha) \cdot \frac{\varepsilon T}{H} \cdot i \\ &= G_{i'} \cdot \frac{\varepsilon T}{H} \cdot i + \alpha \cdot \frac{\varepsilon T}{H} \cdot (i - i') \end{aligned}$$

433 Taking the difference between $\sigma_{(i',j'),t}(i)$ and $\sigma_{(i,j),t}(i)$, we have

$$\begin{aligned}
\sigma_{(i,j),t}(i) - \sigma_{(i',j'),t}(i) &\geq (G_i - G_{i'}) \cdot \frac{\varepsilon T}{H} \cdot i + \alpha \cdot \frac{\varepsilon T}{H} \cdot (i' - i) \\
&= \left(-\sum_{\tau=i+1}^{i'} \frac{i}{\tau} \right) \cdot \frac{\varepsilon T}{H} + \alpha \cdot \frac{\varepsilon T}{H} \cdot (i' - i) \\
&= \left(-(i' - i) + \sum_{\tau=i+1}^{i'} \frac{\tau - i}{\tau} \right) \cdot \frac{\varepsilon T}{H} + \alpha \cdot \frac{\varepsilon T}{H} \cdot (i' - i) \\
&\geq \left(-(i' - i) + \frac{i' - i}{H} \right) \cdot \frac{\varepsilon T}{H} + \alpha \cdot \frac{\varepsilon T}{H} \cdot (i' - i) \\
&\geq \frac{2}{3H} \cdot \frac{\varepsilon T}{H} \cdot (i' - i) > \gamma T
\end{aligned}$$

434 Moreover, notice that $(\kappa_{i,j}, \kappa_{i,j+1}]$ are disjoint intervals for a fix i . In addition, for $\kappa_{i,j} + \gamma T < t <$
435 $\kappa_{i,j+1} - \gamma T$, we have for any $j' < j$,

$$\sigma_{(i,j'),t}(i) = \sigma_{(i,j'),\kappa_{i,j'+1}}(i) - (t - \kappa_{i,j'+1})H \leq G_i \cdot \frac{\varepsilon T}{H} \cdot i - H \cdot \gamma T$$

436 where we use the fact that $\sigma_{(i,j'),\kappa_{i,j'+1}}(i) = G_i \cdot \frac{\varepsilon T}{H} \cdot i$ for all j' due to (1). Similarly, for $j' > j$,

$$\sigma_{(i,j'),t}(i) = \sigma_{(i,j'),\kappa_{i,j'}}(i) - (\kappa_{i,j'} - t)i \leq G_i \cdot \frac{\varepsilon T}{H} \cdot i - i \cdot \gamma T$$

437 where we use the fact that $\sigma_{(i,j'),\kappa_{i,j'}}(i) = G_i \cdot \frac{\varepsilon T}{H} \cdot i$ for all j' due to (1). □

438 A.2 Analysis for Extracting Full Welfare for $\mathcal{V} \subseteq [0, 1]$

439 Let ε' be parameter for the target additive revenue loss $O(\varepsilon' T)$. For ease of presentation, we will
440 rescale \mathcal{V} to $[0, H]$ such that $H = 1/\varepsilon'$, and thus, it suffices to show that we can obtain $O(T)$ loss in
441 the scaled version. First notice that it suffices to consider $\mathcal{V} = [1, H]$ since for all valuations less than
442 1, we will suffer revenue loss at most 1 from each of them.

443 To extend our result to $\mathcal{V} = [1, H]$, consider $i < v_t < i + 1$. Its cumulative utility of playing the
444 option (i, j) with $\kappa_{i,j} < t \leq \kappa_{i,j+1}$ is

$$\begin{aligned}
\sigma_{(i,j),t}(v_t) &= -\alpha \cdot \frac{\varepsilon T}{H} \cdot i + (G_i + \alpha) \cdot \frac{\varepsilon T}{H} \cdot v_t + (t - \kappa_{i,j}) \cdot (v_t - i) \\
&\geq G_i \cdot \frac{\varepsilon T}{H} \cdot v_t + \alpha \cdot \frac{\varepsilon T}{H} \cdot (v_t - i) \\
&= \sigma_{(i,j),t}(i) + (G_i + \alpha) \cdot \frac{\varepsilon T}{H} \cdot (v_t - i)
\end{aligned}$$

445 The next lemma demonstrates that the buyer at round t with valuation $i < v_t < i + 1$ and $t \in$
446 $(\kappa_{i,j}, \kappa_{i,j+1}]$, prefers the option (i, j) over any option (i', j') with $i' \notin \{i, i + 1\}$.

447 **Lemma 11.** For any option (i', j') with $i' \neq i$, $\sigma_{(i,j),t}(v_t) - \sigma_{(i',j'),t}(v_t) > \gamma T$.

448 *Proof.* Notice that its cumulative utility of playing the option (i', j') with $i' < i$ is at most

$$\begin{aligned}
\sigma_{(i',j'),t}(v_t) &= -\alpha \cdot \frac{\varepsilon T}{H} \cdot i' + (G_{i'} + \alpha) \cdot \frac{\varepsilon T}{H} \cdot v_t + (t - \kappa_{i',j'}) \cdot (v_t - i') \\
&\leq G_{i'} \cdot \frac{\varepsilon T}{H} \cdot v_t + \alpha \cdot \frac{\varepsilon T}{H} \cdot (v_t - i') + \frac{\varepsilon T}{3H^2} \cdot (v_t - i') \\
&= G_{i'} \cdot \frac{\varepsilon T}{H} \cdot v_t + \frac{\varepsilon T}{H} \cdot (v_t - i') \\
&= \max_t \{ \sigma_{(i',j'),t}(i) \} + (G_{i'} + 1) \cdot \frac{\varepsilon T}{H} \cdot (v_t - i)
\end{aligned}$$

449 Taking the difference between $\sigma_{(i,j),t}(v_t)$ and $\sigma_{(i',j'),t}(v_t)$ when $i' < i$, we have

$$\begin{aligned}\sigma_{(i,j),t}(v_t) - \sigma_{(i',j'),t}(v_t) &\geq \left(\sum_{\tau=i'+1}^i \left(\frac{i}{\tau} - 1 \right) \right) \cdot \frac{\varepsilon T}{H} + \left(-\frac{1}{3H} + \sum_{\tau=i'+1}^i \frac{1}{\tau} \right) \cdot \frac{\varepsilon T}{H} \cdot (v_t - i) \\ &\geq 0 + \left(-\frac{1}{3H} + \frac{1}{i} \right) \cdot \frac{\varepsilon T}{H} \cdot (v_t - i) > \gamma T\end{aligned}$$

450 where the first inequality partly follows the difference between $\sigma_{(i,j),t}(i)$ and $\max_t \{ \sigma_{(i',j'),t}(i) \}$.

451 Next, notice that its cumulative utility of playing the option (i', j') with $i' > i$ is at most

$$\begin{aligned}\sigma_{(i',j'),t}(v_t) &= -\alpha \cdot \frac{\varepsilon T}{H} \cdot i' + (G_{i'} + \alpha) \cdot \frac{\varepsilon T}{H} \cdot v_t + (t - \kappa_{i',j'}) \cdot (v_t - i') \\ &\leq G_{i'} \cdot \frac{\varepsilon T}{H} \cdot v_t + \alpha \cdot \frac{\varepsilon T}{H} \cdot (v_t - i')\end{aligned}$$

452 Taking the difference between $\sigma_{(i,j),t}(v_t)$ and $\sigma_{(i',j'),t}(v_t)$ when $i' > i + 1$, we have

$$\begin{aligned}\sigma_{(i,j),t}(v_t) - \sigma_{(i',j'),t}(v_t) &\geq (G_i - G_{i'}) \cdot \frac{\varepsilon T}{H} \cdot v_t + \alpha \cdot \frac{\varepsilon T}{H} \cdot (i' - i) \\ &\geq (G_i - G_{i'}) \cdot \frac{\varepsilon T}{H} \cdot (i + 1) + \alpha \cdot \frac{\varepsilon T}{H} \cdot (i' - i) \\ &\geq \left(-\sum_{\tau=i+1}^{i'} \frac{i+1}{\tau} \right) \cdot \frac{\varepsilon T}{H} + \alpha \cdot \frac{\varepsilon T}{H} \cdot (i' - i) \\ &= \left(-(i' - i) + \sum_{\tau=i+1}^{i'} \frac{\tau - i - 1}{\tau} \right) \cdot \frac{\varepsilon T}{H} + \alpha \cdot \frac{\varepsilon T}{H} \cdot (i' - i) \\ &\geq \left(-(i' - i) + \frac{i' - i - 1}{i'} \right) \cdot \frac{\varepsilon T}{H} + \alpha \cdot \frac{\varepsilon T}{H} \cdot (i' - i) \\ &= \left(-\frac{i' - i}{3H} - \frac{i + 1}{i'} + 1 \right) \cdot \frac{\varepsilon T}{H}\end{aligned}$$

453 The minimum is obtained when $i' = i + 2$ or $i' = H$. When $i' = i + 2$, we have

$$-\frac{i' - i}{3H} - \frac{i + 1}{i'} + 1 = -\frac{2}{3H} - \frac{i + 1}{i + 2} + 1 = \frac{1}{i + 2} - \frac{2}{3H} \geq \frac{1}{H} - \frac{2}{3H} > 0$$

454 while when $i' = H$, we have $i \leq H - 2$ and

$$-\frac{H - i}{3H} - \frac{i + 1}{H} + 1 = -\frac{H + 2i + 3}{3H} + 1 \geq -\frac{H + 2(H - 2) + 3}{3H} + 1 = \frac{1}{3H} > 0.$$

455 □

456 Therefore, for $i < v_t < i + 1$ and $t \in (\kappa_{i,j}, \kappa_{i,j+1}]$, the buyer will play option (i', j') with
457 $i' \notin \{i, i + 1\}$ with probability at most γ .

458 Moreover, recall that it is indeed that $\kappa_{i,j} = \kappa_{i+1,j}$ for all i . Therefore, if the buyer plays option
459 $(i + 1, j)$, it will generate revenue $i + 1$. Moreover, if the buyer plays option (i', j') with $j' < j$ and
460 $i' \in \{i, i + 1\}$, then the option is already in its \emptyset session and the buyer is going to pay H . Finally, for
461 any option (i', j') with $j' > j$ and $i' \in \{i, i + 1\}$, the utility of playing such an option is at most

$$\sigma_{(i',j'),t}(v_t) = -\alpha \cdot \frac{\varepsilon T}{H} \cdot i' + \left(t - \kappa_{i',j'} + (G_{i'} + \alpha) \cdot \frac{\varepsilon T}{H} \right) \cdot v_t$$

462 while the utility of playing the option (i', j) is

$$\sigma_{(i',j),t}(v_t) = -\alpha \cdot \frac{\varepsilon T}{H} \cdot i' + (G_i + \alpha) \cdot \frac{\varepsilon T}{H} \cdot v_t + (t - \kappa_{i',j}) \cdot (v_t - i')$$

463 Taking the difference, we have

$$\sigma_{(i',j),t}(v_t) - \sigma_{(i',j'),t}(v_t) = (\kappa_{i',j'} - \kappa_{i',j}) \cdot v_t + (t - \kappa_{i',j}) \cdot (v_t - i')$$

464 Therefore, when $i' = i$, then $\sigma_{(i',j),t}(v_t) - \sigma_{(i',j'),t}(v_t)$ is clear positive. Moreover, when $i' = i + 1$,
465 we have

$$\begin{aligned} \sigma_{(i+1,j),t}(v_t) - \sigma_{(i+1,j'),t}(v_t) &= (\kappa_{i+1,j'} - \kappa_{i+1,j}) \cdot v_t + (t - \kappa_{i+1,j}) \cdot (v_t - i - 1) \\ &\geq (\kappa_{i+1,j+1} - \kappa_{i+1,j}) \cdot v_t + (\kappa_{i+1,j+1} - \kappa_{i+1,j}) \cdot (v_t - i - 1) \\ &= (\kappa_{i+1,j+1} - \kappa_{i+1,j}) \cdot (2v_t - i - 1) \\ &> (\kappa_{i+1,j+1} - \kappa_{i+1,j}) \cdot (2i - i - 1) > \gamma T. \end{aligned}$$

466 Therefore, we finish showing that for $i < v_t < i + 1$ and $t \in (\kappa_{i,j}, \kappa_{i,j+1}]$, the buyer will play option
467 (i', j') with $i' \in \{i, i + 1\}$ and $j' \leq j$ with probability at least $1 - K\gamma$. Thus, the revenue loss from
468 v_t is at most 1. Applying a similar argument as in Section 3.2, we can conclude that the expected
469 revenue loss is $O(T)$.

470 A.3 Proof of Theorem 7

471 We first prove a lower-bound for $\mathcal{V} = \{1, 2, \dots, H\}$.

472 **Lemma 12.** *If the buyer with $\mathcal{V} = \{1, 2, \dots, H\}$ is running a mean-based algorithm, a non-adaptive
473 menu-based mechanism, which obtains expected revenue at least $\text{Val} - O(\varepsilon T)$, must have $\Omega(\frac{H}{\varepsilon})$
474 options, when there exists a null option that always allocates and charges nothing. $\Omega(\frac{H}{\varepsilon})$ options are
475 necessary even when the values of the buyer are drawn from an unknown distribution.*

476 *Proof.* Let $I_{i,t}(c)$ be a binary variable indicating whether the buyer with value $v_t = c$ plays the i -th
477 option. Suppose there are K options in total and let

$$P_i(c) = \sum_{t=1}^T \Pr[I_{i,t}(c) = 1] \cdot p_{i,t}$$

478 be the expected total revenue obtained from the i -th option when the buyer's valuations are $v_t = c$
479 for all t . Since the expected total revenue is at least $\text{Val} - O(\varepsilon T)$, when the buyer's valuations are
480 $v_t = 1$ for all t in which the total expected revenue is at least $T - \mu\varepsilon T$ for some constant μ , there
481 must exist an option i^* such that

$$P_{i^*}(1) \geq \frac{(1 - \mu\varepsilon)T}{K}.$$

482 Moreover, let $t^* = \sup\{t \mid \sigma_{i^*,t}(1) \geq -\gamma T\}$. t^* is well-defined since $\sigma_{i^*,0}(1) = 0$. Notice that for
483 all $t > t^*$, since the buyer is running a mean-based algorithm, we have $\Pr[I_{i^*,t}(1)] \leq \gamma$ due to the
484 presence of the null option. Therefore, we have

$$\sum_{t \leq t^*} p_{i^*,t} + \sum_{t > t^*} \gamma \cdot p_{i^*,t} \geq P_{i^*}(1) \geq \frac{(1 - \mu\varepsilon)T}{K} \Rightarrow \sum_{t \leq t^*} p_{i^*,t} \geq \frac{(1 - \mu\varepsilon)T}{K} - \gamma HT.$$

485 where we use the fact that $0 \leq p_{i^*,t} \leq H$. Notice that the cumulative utility $\sigma_{i^*,t^*}(H)$ is

$$\begin{aligned} \sigma_{i^*,t^*}(H) &= \sum_{t \leq t^*} H \cdot a_{i^*,t} - p_{i^*,t} \\ &= H \cdot \sigma_{i^*,t^*}(1) + (H - 1) \sum_{t \leq t^*} p_{i^*,t} \\ &\geq \frac{(H - 1)(1 - \mu\varepsilon)T}{K} - \gamma H^2 T \end{aligned}$$

486 Consider an environment when the buyer's valuations are $v_t = H$ for all t . Since the buyer is running
487 a no-regret algorithm, her cumulative utility for the first t^* rounds is at least $\sigma_{i^*,t^*}(H) - o(T)$. This
488 is true because although the standard no-regret guarantee only applies to the final round T , the regret
489 for the first t rounds must also be $o(T)$, for any $t < T$. For the sake of contradiction, assume that
490 the regret for the first t rounds is $\Omega(T)$. Notice that the no-regret algorithm does not depend on the

491 future. Therefore, consider an environment where the rewards for all options after round t are set to
 492 be 0, which results in a $\Omega(T)$ regret for the final round T . A contradiction.

493 In addition, notice that the revenue loss from the first t^* rounds is at least the buyer's cumulative
 494 utility, and thus, the revenue loss is at least

$$\sigma_{i^*, t^*}(H) - o(T) = \frac{(H-1)T}{K} - O(\varepsilon T).$$

495 Finally, since the total revenue loss for T rounds is at least the total revenue loss for the first t^* rounds,
 496 in order to achieve $O(\varepsilon T)$ revenue loss, we must have $K = \Omega(\frac{H}{\varepsilon})$. \square

497 Moreover, observe that our proof only uses two sequences of valuations: a sequence with all 1 and
 498 a sequence of all H . Thus, our lower bound also applies to the stochastic settings with unknown
 499 distributions. Finally, Theorem 7 is a simple corollary of Lemma 12.

500 B Critical mechanisms

501 B.1 Values from an unknown distribution

502 We will show that it is possible to achieve the approximation guarantees in Theorem 9 via a *first-price*
 503 *auction with decreasing reserve prices*. This is the same type of mechanism used in Braverman et al.
 504 [2018] to extract the full surplus from buyers drawn with a known distribution. Their construction
 505 requires complete knowledge of \mathcal{D} in order to set these reserve prices over time (specifically, by
 506 solving a linear program whose coefficients depend on \mathcal{D}) – we show we can design a single
 507 construction that gets the same approximation factor to the total economic surplus without knowledge
 508 of \mathcal{D} .

509 In such a mechanism, at each time t the seller specifies a reserve price $f(t)$, where f is a decreasing
 510 function with range $[1, H]$. In each round, if the buyer bids above $f(t)$ they receive the item and pay
 511 b ; otherwise, they do not receive the item and pay nothing. More formally, the allocation and payment
 512 rules $(a_t(b), p_t(b))$ are defined as follows: if $b \geq f(t)$, then $p_t(b) = b$, and $a_t(b) = 1$; otherwise,
 513 $p_t(b) = q_t(b) = 0$.

514 We will often find it easier to work with the function $x(v) : [1, H] \rightarrow [0, T]$ where $x(v) =$
 515 $1 - \frac{1}{T} \cdot \min_{f(t) \leq v} t$. Note that $x(v)$ equals the number of rounds where a bidder with value v
 516 has value higher than the reserve price $f(t)$ (in particular, $x(v)$ is an increasing function of v).

517 It can be shown (Braverman et al. [2018], Section C) that if the buyer is mean-based, the revenue
 518 obtained by the seller by using such an auction is given by

$$R(\mathcal{D}) = T \cdot \mathbb{E}_{v \sim \mathcal{D}} \left[vx(v) - \max_w (v-w)x(w) \right] - o(T).$$

519 Since this is an expectation over v , we have the following lemma:

520 **Lemma 13.** *If the seller is using a first-price auction with decreasing reserve price, then*
 521 *$R(\mathcal{D})/\text{Val}(\mathcal{D})$ is maximized when \mathcal{D} is a singleton distribution.*

522 We now proceed to prove Theorem 9.

523 *Proof of Theorem 9.* Fix any constant $0 < \varepsilon < 1$ (e.g. $\varepsilon = 1/2$). The seller will use a first price
 524 auction with descending reserve price given by

$$f(t) = \max \left(\exp \left(\frac{1}{C} \cdot \left(1 - \eta - \frac{t}{T} \right) \right), 1 \right),$$

525 where $\eta = (1 + \log H)^{-\varepsilon}$ and $C = \frac{1-\eta}{1+\log H}$. Note that for this choice of f , $x(v) = \eta + C \log v$. By
 526 Lemma 13, $R(\mathcal{D})/\text{Val}(\mathcal{D})$ is maximized when \mathcal{D} is a singleton distribution. We therefore have that:

$$\begin{aligned}
\frac{R(\mathcal{D})}{\text{Val}(\mathcal{D})T} &\geq \min_v \frac{vx(v) - \max_w (v-w)x(w)}{v} \\
&= \min_{v, w < v} \left(x(v) - \left(1 - \frac{w}{v}\right) x(w) \right) \\
&= \min_{v, w < v} \left((\eta + C \log v) - \left(1 - \frac{w}{v}\right) (\eta + C \log w) \right) \\
&= \min_{v, w < v} \left(C \log \frac{v}{w} + \frac{w}{v} (\eta + C \log w) \right).
\end{aligned}$$

527 For a fixed w , this is minimized when $v = w \left(\frac{\eta}{C} + \log w \right)$. It follows that

$$\begin{aligned}
\frac{R(\mathcal{D})}{\text{Val}(\mathcal{D})T} &\geq \min_w \left(C \log \left(\frac{\eta}{C} + \log w \right) + 1 \right) \\
&\geq C \log \left(\frac{\eta}{C} \right) \\
&\geq \frac{(1 - (1 + \log H)^{-\epsilon}) \log((1 + \log H)^{1-\epsilon})}{1 + \log H} \\
&= \Theta \left(\frac{\log \log H}{\log H} \right).
\end{aligned}$$

528

□

529 B.2 Adversarial values

530 In this section we will show that it is possible to achieve a $(1/\log H)$ -approximation to the buyer's
531 welfare if they are playing a mean-based algorithm with adversarial values supported on $[1, H]$.
532 This will follow naturally from the known fact that it is possible to achieve the same approximation
533 guarantee against a *strategic* buyer with value in $[1, H]$ playing a single-round version of this game
534 (see e.g. Chapter 6 of [Hartline \[2013\]](#)) – we simply show that if we run this mechanism every round,
535 mean-based buyers will learn to bid in the same manner as strategic buyers.

536 Our mechanism (equivalent to the mechanism presented in Theorem 6.5 of [Hartline \[2013\]](#)) is as
537 follows: for each b and for all t , we set the allocation probability $a_t(b) = (1 + \log b)/(1 + \log H)$ and
538 the expected price charged to $p_t(b) = b/(1 + \log H)$. Note that this mechanism can also be interpreted
539 as a second-price auction where the buyer draws a random reserve from the distribution with CDF
540 $F(r) = \frac{1 + \log r}{1 + \log H}$. It can be seen that for any $v \in [1, H]$, the expected utility $U(v, b) = v \cdot a_t(b) - p_t(b)$
541 of bidding b with value v , is maximized when $b = v$. A strategic buyer therefore will always bid their
542 value, and pay $1/(1 + \log H)$ of their value in total.

543 We will show that the mean-based guarantee ensures that a mean-based buyer will (most of the time)
544 choose a bid close to v , and thus contribute a similar amount of revenue as a strategic buyer.

545 **Lemma 14.** *Let*

$$U(v, b) = v \cdot a_t(b) - p_t(b) = \frac{v \cdot (1 + \log b) - b}{1 + \log H}.$$

546 *If $U(v, v) - U(v, b) \leq \delta$, then $|v - b| \leq H \sqrt{2(1 + \log H)\delta}$.*

547 *Proof.* Let $f(v, b) = v(1 + \log b) - b$. Note that

$$\frac{\partial}{\partial b} f(v, b) = \frac{v}{b} - 1$$

548 and

$$\frac{\partial^2}{\partial b^2} f(v, b) = -\frac{v}{b^2} \leq -\frac{1}{H^2}.$$

549 This implies f is $(1/2H^2)$ -strongly concave and maximized when $b = v$, so

$$f(v, v) - f(v, b) \geq \frac{1}{2H^2}(v - b)^2.$$

550 Since $U(v, b) = f(v, b)/(1 + \log H)$, this implies that if $U(v, v) - U(v, b) \leq \delta$, then $|v - b| \leq$
 551 $H\sqrt{2(1 + \log H)\delta}$. \square

552 *Proof of Theorem 10.* Consider the critical mechanism defined by $a_t(b) = (1 + \log b)/(1 + \log H)$
 553 and $p_t(b) = b/(1 + \log H)$. We claim that this mechanism obtains expected revenue at least
 554 $\frac{1}{1 + \log(H)} \text{Val} - o(T)$ against any mean-based bidder with values adversarially chosen from $[1, H]$.

555 Note that by the mean-based guarantee, with probability at least $1 - \gamma$ a mean-based algorithm will
 556 pick a bid b_t satisfying

$$\sigma_{v_t, t}(b^*) - \sigma_{v_t, t}(b_t) \leq \gamma T, \quad (2)$$

557 where $b^* = \arg\max_b \sigma_{v_t, t}(b)$. Now, note that $\sigma_{v_t, t}(b) = t \cdot U(v, b)$; since $U(v, b)$ is maximized when
 558 $b = v$, $b^* = v_t$. It then follows from Lemma 14, that if (2) holds, then with probability at least $(1 - \gamma)$

$$|v_t - b_t| \leq H\sqrt{2(1 + \log H)\frac{\gamma T}{t}}. \quad (3)$$

559 Since $v_t - b_t \leq H$ is always true (since $v_t \leq H$), it follows that in expectation,

$$\mathbb{E}[b_t] \geq v_t - \gamma H - H\sqrt{2(1 + \log H)\frac{\gamma T}{t}}$$

560 and therefore we have that

$$\begin{aligned} \mathbb{E}[\text{Rev}] &= \mathbb{E}\left[\sum_{t=1}^T p_t(b_t)\right] \\ &= \frac{1}{1 + \log H} \sum_{t=1}^T \mathbb{E}[b_t] \\ &\geq \frac{1}{1 + \log H} \sum_{t=1}^T (v_t - \gamma H - H\sqrt{2(1 + \log H)\frac{\gamma T}{t}}) \\ &= \frac{\text{Val}}{1 + \log H} - \frac{H}{1 + \log H} \gamma T - H\sqrt{\frac{2}{1 + \log H}} \gamma T \sum_{t=1}^T \frac{1}{t} \\ &\geq \frac{\text{Val}}{1 + \log H} - \frac{H}{1 + \log H} \gamma T - H\sqrt{\frac{2}{1 + \log H}} \gamma T \\ &= \frac{\text{Val}}{1 + \log H} - o(T), \end{aligned}$$

561 where the last inequality holds since γ and $\sqrt{\gamma}$ are both $o(1)$.

562 \square