

# Clock Auctions Augmented with Unreliable Advice

Vasilis Gkatzelis\*

Daniel Schoepflin†

Xizhi Tan‡

## Abstract

We provide the first analysis of (deferred acceptance) clock auctions in the learning-augmented framework. These auctions satisfy a unique list of very appealing properties, including obvious strategyproofness, transparency, and unconditional winner privacy, making them particularly well-suited for real-world applications. However, early work that evaluated their performance from a worst-case analysis perspective concluded that no deterministic clock auction with  $n$  bidders can achieve a  $O(\log^{1-\epsilon} n)$  approximation of the optimal social welfare for a constant  $\epsilon > 0$ , even in very simple settings. This overly pessimistic impossibility result heavily depends on the assumption that the designer has no information regarding the bidders' values. Leveraging the learning-augmented framework, we instead consider a designer equipped with some (machine-learned) advice regarding the optimal solution; this advice can provide useful guidance if accurate, but it may be unreliable.

Our main results are learning-augmented clock auctions that use this advice to achieve much stronger performance guarantees whenever the advice is accurate (known as *consistency*), while maintaining worst-case guarantees even if this advice is arbitrarily inaccurate (known as *robustness*). Our first clock auction achieves the best of both worlds:  $(1 + \epsilon)$ -consistency for any desired constant  $\epsilon > 0$  and  $O(\log n)$  robustness; we also extend this auction to achieve error tolerance. We then consider a much stronger notion of consistency, which we refer to as *consistency* $^\infty$ , and provide an auction that achieves a near-optimal trade-off between *consistency* $^\infty$  and robustness. Finally, using our impossibility results regarding this trade-off, we prove lower bounds on the “cost of smoothness,” i.e., on the robustness that is achievable if we also require that the performance of the auction degrades smoothly as a function of the prediction error.

## 1 Introduction

Our focus in this paper is on a classic family of mechanism design problems involving a set  $N$  of  $n$  buyers competing for some type of resource or service, and a seller (the service provider) who needs to decide which subset of buyers to serve. Each buyer  $i \in N$  has a value  $v_i$  for receiving the service (i.e., the price they are willing to pay for it), and there is a feasibility constraint  $\mathcal{F} \subseteq 2^N$  that restricts the sets of agents that can be served simultaneously. The goal of the service provider is to achieve an efficient outcome, i.e., to serve a feasible set of buyers  $F \in \mathcal{F}$  that maximizes social welfare:  $\sum_{i \in F} v_i$ . Depending on the nature of the feasibility constraint  $\mathcal{F}$ , this captures many well-studied optimization problems, like the knapsack problem, the maximum weight independent set problem, and many more. However, apart from the computational challenges of the underlying optimization problem, the seller also faces the additional obstacle that the values of the buyers are private (not known to the seller), making it even harder to maximize social welfare.

To address the fact that buyers' values are private, a common solution in the mechanism design literature is to use *direct revelation mechanisms*: these mechanisms ask each buyer to report their value to the mechanism, which can then use this information to decide which buyers should receive the service and how much they should pay for it. However, unless these mechanisms are carefully designed, they are very likely to introduce incentives for buyers to misreport their true values (e.g., to report a lower value to reduce their payment for the service). A central goal in mechanism design is to design *strategyproof* mechanisms, which ensure that buyers' optimal strategy is to always report their true values. A classic example of a strategyproof direct revelation mechanism

\*Drexel University. Partially supported by NSF CAREER award CCF-2047907 and NSF grant CCF-2210502, CCF-2008280 and CCF-1755955.

†Rutgers University – DIMACS. Partially supported by NSF grants CCF-2008280 and CCF-1755955 and by a grant to DIMACS from the Simons Foundation (820931). A portion of this work was supported by the NSF under Grant No. DMS-1928930 and by the Alfred P. Sloan Foundation under grant G-2021-16778, while this author was in residence at the Simons Laufer Mathematical Sciences Institute (formerly MSRI) in Berkeley, California, during the Fall 2023 semester.

‡Drexel University. Partially supported by NSF CAREER award CCF-2047907 and NSF grant CCF-2210502.

that maximizes social welfare is the Vickrey-Clarke-Groves (VCG) mechanism. However, despite its elegance and great theoretical appeal, the VCG mechanism is rarely used in practice due to a variety of shortcomings [4]. For example, its strategyproofness heavily depends on its ability to solve the underlying optimization problem optimally, even though it may be NP-hard. In fact, direct revelation mechanisms in general face several issues that make them impractical for many applications. For example, (i) even if they are actually strategyproof, it may be non-trivial for buyers to verify it [40], (ii) they require that buyers put a lot of trust in the designer [2], and (iii) buyers need to directly reveal all of their private information, which can be sensitive and could potentially be used against them in the future (e.g., to set higher reserve prices) [54].

To overcome the issues that direct revelation mechanisms face, Milgrom and Segal [50, 51] proposed *deferred-acceptance clock auctions* as a much more practical alternative for high-stakes auctions. In a clock auction, buyers are not asked to report their private values. Instead, the auction takes place over a sequence of rounds, and each buyer is offered a personalized price that weakly increases over time. Buyers can then remain active as long as they are willing to pay the price offered to them, and they can drop out whenever the price offered exceeds their value. These auctions have a wide variety of benefits that make them very practical: they guarantee (i) transparency (there is no way the auctioneer can mishandle the information provided by the buyers behind the scenes), (ii) unconditional winner privacy (the winners of the auction never need to reveal their true value), and (iii) simplicity (buyers do not need to understand the inner workings of the auction; all they need to know is the price offered to them in each round). These are properties that most strategyproof auctions do not satisfy.

Motivated by this unique list of very appealing properties, subsequent work focused on analyzing the worst-case approximation guarantees that clock auctions can achieve for a variety of different feasibility constraints, often setting aside computational constraints and focusing on information-theoretic constraints. The first set of results was rather pessimistic, showing that no deterministic clock auction can guarantee an approximation much better than  $O(\log n)$  even in seemingly simple settings [24]. On the positive side, subsequent work showed that there exists a natural clock auction, the water-filling clock auction (WFCA), that can guarantee an approximation of  $O(\log n)$  for any type of feasibility constraint [21]. Although the water-filling auction achieves the optimal worst-case approximation guarantee, a logarithmic approximation may not be quite as appealing from a practical standpoint. This is a common issue when analyzing algorithms and mechanisms from a worst-case perspective, leading to results that may be overly pessimistic.

To overcome the overly pessimistic nature of worst-case analysis, a surge of recent work has used the *learning-augmented framework*, which assumes that the designer is provided with some useful, though unreliable, prediction or advice regarding the instance at hand. Using this advice, the goal is to provide stronger and much more practical performance guarantees whenever this advice is accurate (the *consistency* guarantee), while simultaneously maintaining non-trivial worst-case guarantees even if the advice is arbitrarily inaccurate (the *robustness* guarantee). As a result, this framework provides a natural way to leverage machine-learned predictions to guide the design of mechanisms while maintaining the important robustness that comes with worst-case analysis.

**1.1 Our Results.** We provide the first analysis of clock auctions within the learning-augmented framework. Specifically, rather than assuming that the auction has absolutely no information regarding the values of the bidders, we assume it is provided with a prediction,  $\hat{\text{OPT}} \in \mathcal{F}$ , regarding which feasible set in  $\mathcal{F}$  has the optimal social welfare. Crucially, this prediction can be arbitrarily inaccurate, so the auction must use it carefully to maintain any bounded robustness.

In stark contrast to prior work, our first main result is highly optimistic: we propose a learning-augmented clock auction that simultaneously achieves the "best of both worlds." Specifically, our auction achieves a constant approximation of the optimal social welfare whenever the prediction is accurate, i.e.,  $\hat{\text{OPT}}$  is indeed optimal, while maintaining the worst-case approximation guarantee of  $O(\log n)$ , regardless of how inaccurate the prediction may be. In fact, we show that it achieves a consistency of  $1 + \epsilon$  for any desired, arbitrarily small constant  $\epsilon > 0$ , while maintaining a robustness of  $O(\frac{1}{\epsilon} \log n)$ , and we prove that this is asymptotically optimal even with respect to its dependence on  $\epsilon$ . We also extend this auction to achieve error tolerance, i.e., an approximation guarantee that degrades gracefully as a function of the error, up to some predetermined error tolerance threshold.

For our second main result, we move beyond the standard notion of consistency and consider a significantly stronger notion, which we refer to as consistency $^\infty$ . In contrast to the standard consistency constraint that binds only if the prediction is fully accurate, consistency $^\infty$  binds on every input and requires that the social welfare

achieved by the auction always approximates the social welfare of the predicted set,  $\hat{\mathcal{O}}\hat{\mathcal{T}}$ . Note that, whenever  $\hat{\mathcal{O}}\hat{\mathcal{T}}$  is indeed optimal, approximating its social welfare reduces to approximating the optimal social welfare, so consistency $^\infty$  is strictly more demanding than the standard notion of consistency.

This notion of consistency dates back to earlier work by Mahdian et al. [48] and has also been studied explicitly or implicitly in some subsequent work in the learning-augmented literature. Our main result for this notion is a learning-augmented clock auction that achieves a near-optimal trade-off between robustness and consistency $^\infty$ . Specifically, we show that the optimal robustness achievable by any learning-augmented clock auction that satisfies  $\alpha$ -consistency $^\infty$  is  $\Omega(\frac{1}{\alpha}n^{1/(\alpha-1)}\log n)$  and our auction combines  $\alpha$ -consistency $^\infty$  with  $O(n^{1/(\alpha-1)}\log n)$  robustness. Note that for any constant  $\alpha$  consistency $^\infty$ , the robustness of our auction is asymptotically optimal.

Finally, using our robustness lower bound for constant consistency $^\infty$ , we also prove a lower bound on the “cost of smoothness” for clock auctions, i.e., a lower bound on the best achievable robustness if we also require that the performance of the auction degrades smoothly as a function of the prediction error. This result exhibits an interesting connection between consistency $^\infty$  and smoothness and is in contrast with almost all prior work in the learning-augmented framework, which does not provide impossibility results regarding smoothness guarantees.

**1.2 Related Work.** Our paper lies at the intersection of two lines of literature: the literature on (deferred acceptance) clock auctions and the literature on learning-augmented algorithm and mechanism design. Below, we provide an overview of the most relevant papers in each of these two areas.

**Clock auctions.** Since their formal introduction in [50, 51], clock auctions have attracted attention from both the economics and computer science communities, partly due to their many practical advantages over the sealed-bid format. These important practical properties include: (i) obvious strategyproofness, a strong notion of incentive compatibility originally defined by Li [43]; (ii) unconditional winner privacy, defined by Milgrom and Segal [51], which guarantees that clock auctions reveal the minimum possible information about the winners’ values; and (iii) transparency, including the fact that ascending clock auctions are “credible” auctions, as defined in [2]. Notably, clock auctions are the *unique* class of auctions satisfying these properties, making them particularly well-suited to practical implementation. We refer the reader to [51] and [28, Section 1.2] for a more complete discussion of the properties of clock auctions.

The papers technically closest to our work include [6, 21, 24]. Before [50, 51] explicitly defined clock auctions, Babaioff et al. [6] proposed a deterministic, prior-free clock auction that achieved an  $O(\log v_{\max})$ -approximation, where  $v_{\max}$  is the highest value of any single bidder. However, their auction notably achieved only a  $\Theta(n)$ -approximation. In terms of lower bounds, Dütting et al. [24] demonstrated that no deterministic, prior-free clock auction could achieve a  $\log^\tau n$ -approximation for any constant  $\tau < 1$ . Over a decade after [6], Christodoulou et al. [21] showed that the bound in [24] was essentially tight by providing a deterministic, prior-free clock auction that achieved an  $O(\log n)$ -approximation (and simultaneously an  $O(\log v_{\max})$ -approximation).

In single-parameter, forward auction settings, Christodoulou et al. [21] were the first to propose using *randomization* to circumvent the impossibility results of [24]. They presented an auction that guarantees *expected* welfare, achieving an  $O(\sqrt{\log k})$ -approximation to the optimal social welfare, where  $k$  is the number of maximal feasible sets in the set system  $\mathcal{F}$ . This was later improved to an  $O(\log \log k)$ -approximation guarantee by [28].

Clock auctions have also been studied beyond the single-parameter, binary service level, forward auction context. For example, [32] demonstrated how to extend clock auctions to settings with multiple levels of service and analyzed several special cases of this general framework. Many works have also examined *reverse* (i.e., procurement) clock auctions, both with [38, 9, 27, 16, 34] and without [42, 18] an auctioneer budget constraint. The works of [45, 26] studied double clock auctions involving both strategic sellers and buyers. The work of [33] studied clock auctions in the economically significant setting of interdependent values. Finally, the work of [30] examined clock auctions aimed at maximizing consumer utility.

**Learning-augmented framework.** In recent years, the learning-augmented framework has gained widespread recognition as a valuable paradigm for designing and analyzing algorithms. We refer to [52] for a survey of early contributions and [44] for a frequently updated list of papers in this field. This approach aims to overcome the limitations of overly pessimistic worst-case bounds. In the past five years alone, over 200 papers have revisited classic algorithmic problems using this framework, with prominent examples including online paging [47], scheduling [53], optimization problems involving covering [13] and knapsack constraints [35], as well as Nash social welfare maximization [14], the secretary problem [3, 25, 29], and a variety of graph-related challenges [5]. As we discuss in our results section, there is another, stronger notion of consistency that requires a good

approximation of the *predicted solution* irrespective of the quality of predictions. This notion was first proposed in [48], and subsequent work, including [13, 41], also defines it as consistency. Additionally, we note that while some works do not explicitly state it, their results actually hold for this stronger notion of consistency [55, 39, 1].

More closely related to our work, the line of research on learning-augmented mechanisms interacting with strategic agents is recent and was initiated by Agrawal et al. [1] and Xu and Lu [56]. This line of work encompasses strategic facility location [1, 56, 36, 15, 20, 10], strategic scheduling [56, 11], auctions [49, 56, 46, 19, 12], bicriteria mechanism design (optimizing both social welfare and revenue) [7], graph problems with private input [23], metric distortion [17], and equilibrium analysis [31, 37]. Recently, [22] revisited mechanism design problems with predictions on the outcome space instead of the input. The prediction discussed in our work is one such example. For further discussion regarding this line of work, we refer the reader to [8].

## 2 Preliminaries

We consider a canonical single-parameter auction setting where an auctioneer aims to allocate a service among a set  $N$  of  $n$  bidders. There exists a public family of sets  $\mathcal{F} \subseteq 2^N$  indicating subsets of bidders that can feasibly be served by the auctioneer. We assume  $\mathcal{F}$  is *downward-closed*, meaning if  $S \in \mathcal{F}$ , then  $T \subseteq S$  implies  $T \in \mathcal{F}$ . The auctioneer's goal is to select a feasible subset  $S \in \mathcal{F}$  of bidders to maximize social welfare, i.e.,  $v(S) = \sum_{i \in S} v_i$ . Since bidders have private values, the auctioneer must design a mechanism to elicit these values. However, bidders are selfish and may strategically misreport their values to achieve their preferred outcomes (i.e., receiving the service at a lower price). Therefore, the auctioneer must carefully design payment rules to incentivize accurate reporting.

We design clock auctions for this problem. A clock auction is a dynamic mechanism proceeding over several rounds. In each round  $t$ , each bidder  $i$  is offered a personalized “clock” price  $p_{i,t}$ , which represents the current amount they would need to pay if the auction were to terminate at that moment. At the outset of the auction, the clock price  $p_{i,0}$  for each bidder  $i$  is initialized to some arbitrarily small positive value  $v_{\min}$ , and each bidder is placed in an “active” bidder set  $A$ . The clock prices are non-decreasing throughout the auction, i.e.,  $p_{i,t} \geq p_{i,t-1}$  for all rounds  $t$ .<sup>1</sup> In each round  $t$ , the auctioneer announces a price  $p_{i,t}$  for each bidder  $i \in A$ , and bidder  $i$  can then choose to remain in the auction (staying “active”) or permanently exit (becoming “inactive” and being removed from  $A$ ). If  $i$  exits at any point in the auction, they will not receive the service and pay nothing. The auction then terminates when  $A \in \mathcal{F}$ , i.e., all active bidders may be feasibly served.

Note that in each round, the clock price for each bidder can only depend on public information, i.e., the feasibility constraint, the history of prices, and the points at which bidders exited the auction. We use  $\text{rev}(S, \mathbf{p}_t) = \sum_{i \in S \cap A} p_{i,t}$  to denote the *revenue* of a set  $S$  at the current price vector  $\mathbf{p}_t$ . In other words, the revenue of a set  $S$  in round  $t$  of the auction is the sum of the clock prices in round  $t$  of the active bidders in  $S$ .

To analyze the performance of a clock auction  $\mathcal{M}$ , we consider its worst-case approximation guarantee. On an instance  $I$ , let  $I$  denote the welfare obtained by the clock auction on  $I$  and  $\text{OPT}(I)$  denote the maximum social welfare among all feasible sets in  $I$ . Then,  $\mathcal{M}$  obtains an  $\alpha$ -approximation of the optimal social welfare for a class of instances  $\mathcal{I}$  if

$$\max_{I \in \mathcal{I}} \frac{v(\text{OPT}(I))}{v(\mathcal{M}(I))} \leq \alpha.$$

**The Water-Filling Clock Auction.** The central result of [21] is a deterministic clock auction called the “Water-Filling Clock Auction” (WFCA), which achieves the (near) optimal approximation guarantee of  $O(\log n)$  for any family of instances where  $\mathcal{F}$  is downward-closed. The WFCA computes the revenue of each set in each round and marks the set of bidders  $W$  with the highest revenue as the “conditional winners” and increments the price of the lowest-priced bidders outside  $W$ . This, intuitively, makes sense since the revenue of a set is a lower bound on its remaining possible welfare. As the WFCA is an important subroutine in our auctions, we include pseudocode for it here as Subroutine 1 for completion.

We summarize two critical features of the WFCA shown in [21] that we use in our analysis as Lemma 2.1 below. We note that the authors of [21] show that the WFCA achieves a  $4 \log n$ -approximation [21, Theorem 2].

<sup>1</sup>We assume that each bidder with positive value has a value above a publicly known minimum value  $v_{\min}$ , which is consistent with [6] and [21]. Thus, rather than continuous prices, we may increment prices by an arbitrarily small positive value  $\delta$ , say,  $v_{\min}/n^2$ . As argued in [21], discretizing prices to this range causes negligible loss in approximation since the optimal social welfare is lower-bounded by  $v_{\min}$ .

---

**SUBROUTINE 1:** The Water-Filling Clock Auction (WFCA) of [21]

---

```

Input: An initial price vector  $\mathbf{p}$  and set of bidders  $\tilde{N}$ 
1  $t \leftarrow 0$ ,  $A \leftarrow \tilde{N}$ , and  $p_{i,t} \leftarrow \tilde{p}_i$  for all  $i \in \tilde{N}$ 
2 while  $A \notin \mathcal{F}$  do
3    $t \leftarrow t + 1$ 
4    $W \leftarrow \arg \max_{S \in \mathcal{F}: S \subseteq A} \left\{ \sum_{i \in S} p_{i,t-1} \right\}$  be the set of highest current revenue
5    $\ell \leftarrow \min_{i \in A \setminus W} \{p_{i,t-1}\}$  be the lowest price among active bidders not in  $W$ 
6   foreach bidder  $i \in A \setminus W$  with  $p_{i,t-1} = \ell$  do
7      $p_{i,t} \leftarrow p_{i,t-1} + \delta$ ; // Increment prices of lowest-priced ‘‘losers’’
8     if  $i$  rejects updated price then
9        $A \leftarrow A \setminus \{i\}$ 
10    foreach bidder  $i \in W$  and each bidder with  $p_{i,t-1} > \ell$  do
11       $p_{i,t} \leftarrow p_{i,t-1}$ ; // Keep all other prices the same
12 return  $A, \mathbf{p}$ 

```

---

The proof of this theorem demonstrates a slightly stronger guarantee that is more useful for our bounds, namely that the welfare achieved is a  $2H_n$ -approximation, where  $H_n = \sum_{i=1}^n 1/i$  is the  $n$ -th harmonic number. It is known that  $H_n$  is asymptotically  $O(\log n)$ .

LEMMA 2.1. ([21]) *The Water-Filling Clock Auction is revenue monotone, i.e., the sum of prices of bidders in the set with the highest sum of prices is monotonically increasing throughout the WFCA. In addition, the WFCA obtains a  $2H_n$ -approximation to the optimal social welfare in any downward-closed set system.*

**The Learning-Augmented Framework.** In this work, we adopt the learning-augmented framework and study clock auctions that are also equipped with a (potentially erroneous) prediction  $\hat{\text{OPT}}$  regarding the feasible set of highest welfare. Given the predictions and an instance, we denote the welfare achieved by a mechanism  $\mathcal{M}$  as  $\mathcal{M}(I, \hat{\text{OPT}})$ , and we evaluate the performance of  $\mathcal{M}$  using its *robustness* and *consistency*.<sup>2</sup>

The robustness of a mechanism refers to the worst-case approximation ratio of the mechanism given an adversarially chosen, possibly erroneous, prediction. Mathematically,

$$\text{robustness}(\mathcal{M}) = \max_{I, \hat{\text{OPT}}} \frac{v(\text{OPT}(I))}{\mathcal{M}(I, \hat{\text{OPT}})}.$$

The consistency of a mechanism refers to the worst-case approximation ratio of the mechanism when the prediction it is provided with is accurate, i.e.,  $\hat{\text{OPT}} = \text{OPT}$ . Mathematically,

$$(2.1) \quad \text{consistency}(\mathcal{M}) = \max_I \frac{v(\text{OPT}(I))}{\mathcal{M}(I, \text{OPT}(I))}.$$

Additionally, we consider a stronger notion of consistency first proposed by [48], which we denote as  $\text{consistency}^\infty$ . It is defined to be the worst-case approximation ratio with respect to the *value of the predicted set*, regardless of the correctness of the prediction. Mathematically,

$$(2.2) \quad \text{consistency}^\infty(\mathcal{M}) = \max_{I, \hat{\text{OPT}}} \frac{v(\hat{\text{OPT}})}{\mathcal{M}(I, \hat{\text{OPT}})}.$$

We emphasize that an  $\alpha$ -consistency $^\infty$  guarantee implies an  $\alpha$ -consistency guarantee for any  $\alpha$  by considering a subset of instances where  $\hat{\text{OPT}} = \text{OPT}(I)$ .

---

<sup>2</sup>We note that the predictions used by mechanisms are *public*, i.e., the mechanism designer and bidders observe the predictions. Although our clock auctions utilize the predictions to guide the price-increase process, the bidders still face only monotonically increasing clock prices and, thus, they still have a simple interface with an obviously dominant strategy of exiting the auction when the clock price becomes undesirable (higher than their private value).

### 3 Clock Auction Achieving the Best of Both Worlds

In this section, we present a deterministic clock auction augmented with a prediction  $\hat{OPT} \in \mathcal{F}$  that suggests the feasible set which maximizes social welfare. Our auction leverages this prediction to achieve a favorable trade-off between consistency (see Equation (2.1)) and robustness. Recall that without predictions, no mechanism can achieve an  $O(\log^\tau n)$ -approximation for any constant  $\tau < 1$ . The aim is to bypass this logarithmic barrier whenever the prediction is accurate, without significantly sacrificing robustness.

One naive approach is to simply serve the predicted set, which would yield 1-consistency; however, this would suffer from unbounded robustness. Conversely, ignoring the prediction and simply running the WFCA would provide  $O(\log n)$  consistency and  $O(\log n)$  robustness. A desirable outcome is small constant consistency with  $O(\log n)$  robustness, often referred to as a “best-of-both-worlds” guarantee. This is precisely what our clock auction, formally defined in Mechanism 1, achieves. In fact, we can set the consistency arbitrarily close to  $1 + \epsilon$  for any small constant  $\epsilon > 0$ , while maintaining  $O(\frac{1}{\epsilon} \log n)$ -robustness. We complement this positive result with a lower bound, showing that this robustness bound is tight with respect to its dependence on both  $n$  and  $\epsilon$ .

**THEOREM 3.1.** *For any downward-closed set system, Mechanism 1 with parameter  $\gamma = \frac{10(1+\epsilon)}{9\epsilon}$  is  $(1+\epsilon)$ -consistent and  $O(\frac{1}{\epsilon} \log n)$ -robust for any constant  $\epsilon > 0$ .*

The mechanism, called FOLLOWTHEUNPREDICTEDLEADER, is a clock auction that alternates between raising the price for bidders in the predicted and unpredicted sets, while maintaining a balanced ratio between the target revenue from the predicted set and the rejected welfare of the unpredicted set. At a high level, in each iteration, the mechanism first sets a new, higher revenue “target”  $R_t$  and pushes the unpredicted bidders only until the robustness guarantee based on current revenue is nearly at risk. It then further pushes the unpredicted bidders, aiming to ensure the robustness guarantee is satisfied solely by the unpredicted set. If all unpredicted bidders reject the offers during this process, the mechanism outputs the remaining active bidders in the predicted set. Otherwise, it switches to pushing the predicted set to match the current revenue obtained. If the target revenue is met, the process repeats; if all predicted bidders reject, the WFCA is run on the remaining unpredicted bidders. We assume that bidders choose the obviously dominant strategy of exiting the auction whenever the clock price exceeds their value, and we learn each bidder’s value upon their exit.

For a formal description, please refer to Mechanism 1.

---

#### MECHANISM 1: FOLLOWTHEUNPREDICTEDLEADER

---

**Input:** A system  $\mathcal{F}$  of feasible sets of bidders, a predicted optimal set  $\hat{OPT}$ , a parameter  $\gamma$

- 1  $F \leftarrow F \setminus \{F \cap \hat{OPT}\}$  for all  $F \neq \hat{OPT} \in \mathcal{F}$  // Make  $\hat{OPT}$  disjoint from all unpredicted sets
- 2  $p_i \leftarrow v_{\min}$  for all  $i \in N$  // Raise price of each bidder to the initial small amount
- 3  $R_0 \leftarrow \text{rev}(\hat{OPT}, \mathbf{p})$
- 4  $t \leftarrow 1$  // Initialize the counter
- 5 **while** *TRUE* **do**
- 6      $R_t \leftarrow 10 \cdot R_{t-1}$
- 7     Run  $\text{UNIFORMPRICE}(N \setminus \hat{OPT}, \mathbf{p})$  until  $\max_F v(F \setminus N) \in [R_t \gamma H_n, 2R_t \gamma H_n]$  // Reject “safe” amount of welfare to determine minimum total value in optimal set<sup>3</sup>
- 8     **if**  $N = \hat{OPT}$  **then return**  $\hat{OPT}, \mathbf{p}$
- 9     Run  $\text{UNIFORMPRICE}(N \setminus \hat{OPT}, \mathbf{p})$  until  $\max_F \text{rev}(F \cap N, \mathbf{p}) \geq R_t$  // Ensure unpredicted sets “cover” their own lost welfare
- 10    **if**  $N = \hat{OPT}$  **then return**  $\hat{OPT}, \mathbf{p}$
- 11    Run  $\text{UNIFORMPRICE}(\hat{OPT}, \mathbf{p})$  until  $\text{rev}(\hat{OPT}, \mathbf{p}) \geq R_t$  // Ensure predicted set “covers” welfare lost from unpredicted sets
- 12    **if**  $\hat{OPT} = \emptyset$  **then**
- 13      **return**  $\text{WFCA}(N, \mathbf{p})$
- 14     $t \leftarrow t + 1$

---

<sup>3</sup>We can make sure the total welfare in some unpredicted set is at least  $R_t \gamma H_n$  but the rejected welfare is no more than  $2R_t \gamma H_n$  by stopping when current price of active unpredicted bidders reaches  $R_t \gamma H_n$ .

---

**SUBROUTINE 2: UNIFORMPRICE**


---

**Input:** A set  $S$  and the current price  $p_i$  for all  $i \in S$ , the current price vector  $\mathbf{p}$

- 1 **while** do
- 2      $\ell \leftarrow \min_{i \in S} p_i$
- 3     **for**  $i$  with  $p_i = \ell$  **do**
- 4          $p_i \leftarrow p_i + \epsilon$
- 5         **if**  $i$  reject the updated price  $i$  **then**
- 6              $S \leftarrow S \setminus \{i\}$

---

We first show that Line 1 affects neither consistency nor robustness by too much.

**OBSERVATION 3.1.** *The social welfare of the optimal set in  $\mathcal{F}$  before making the transformation in line 1 of Mechanism 1 is at most 2 times greater than the social welfare of the optimal set after the transformation.*

*Proof.* We first note that the social welfare of the predicted set is the same before and after completing the transformation in line 1 of Mechanism 1. Then if the predicted set was the optimal set before the transformation, then the value of the optimal set does not decrease. So suppose some non-predicted set  $O^*$  is optimal and denote its total welfare  $v(O^*)$ . If at least half of the welfare contained in  $O^*$  comes from bidders contained in the predicted set then after the transformation the predicted set has welfare at least  $v(O^*)/2$  and, hence, the optimal set has welfare at least  $v(O^*)/2$  after the transformation. On the other hand, if more than half of the welfare contained in  $O^*$  comes from bidders not in the predicted set then after the transformation  $O^*$  still has social welfare at least  $v(O^*)/2$ .  $\square$

For ease of exposition, for the remainder of this section, we assume that  $\mathcal{F}$  is such that the predicted set is completely disjoint from all maximal unpredicted sets. According to Observation 3.1, the *consistency* guarantees we obtain in this special case are exactly the same as those we obtain for general set systems. Moreover, by Observation 3.1, the *robustness* guarantees we obtain in this special case are within a factor of 2 of the robustness guarantees we obtain for general set systems. Therefore, it suffices to show that we obtain  $1 + \epsilon$ -consistent and  $O(\log n)$ -robust algorithms for the special case where the predicted set is completely disjoint from all unpredicted sets. We devote the rest of this section to demonstrating just that. For ease of presentation, without loss of generality, we normalize  $R_0$  to 1. All the missing proofs can be found in Appendix A.

**Robustness Argument.** We first establish the robustness of Mechanism 1. To begin, we demonstrate the robustness guarantee if a subset of the predicted set is output (Lemma 3.2). For cases where an unpredicted set is output, we argue that the revenue obtained at the time point where  $\mathbf{OPT} = \emptyset$  provides a good approximation of the best welfare achievable from any rejected subset (Lemma 3.3). We then utilize the revenue monotonicity of WFCA to achieve the desired robustness bound (Lemma 3.4).

**LEMMA 3.1.** *The total welfare rejected from any unpredicted set in line 9 in any iteration  $t$  of while loop of Mechanism 1 is at most  $R_t \cdot H_n$ .*

**LEMMA 3.2.** *For any  $\gamma > 10/9$ , if Mechanism 1 outputs a subset of the predicted set, then it is  $\frac{9}{100(2\gamma+1)H_n}$ -robust.*

*Proof.* Consider the iteration  $\hat{t}$  in which the while loop of Mechanism 1 terminates. By definition, the mechanism terminates in either line 8 or line 10. By Line 7 and Lemma 3.1 we have that the total welfare rejected from any unpredicted set in any round  $t'$  is at most  $2R_{t'}\gamma H_n + R_{t'}H_n = R_{t'} \cdot (2\gamma + 1)H_n$ . But then, summing over all rounds we have that the total welfare contained in any unpredicted set is at most

$$\sum_{t'=1}^{\hat{t}} R_{t'} \cdot (2\gamma + 1)H_n = (2\gamma + 1)H_n \cdot \sum_{t'=1}^{\hat{t}} 10^{t'} = (2\gamma + 1)H_n \cdot \frac{10^{\hat{t}+1} - 1}{9} \leq R_{\hat{t}} \cdot \frac{(2\gamma + 1)10H_n}{9}.$$

On the other hand, since we continued to iteration  $\hat{t}$  we have that the revenue of the bidders in the predicted set reached  $R_{\hat{t}-1} = \frac{R_{\hat{t}}}{10}$ . Hence, Mechanism 1 obtains at least a  $\frac{R_{\hat{t}}/10}{R_{\hat{t}} \cdot \frac{(2\gamma+1)10H_n}{9}} = \frac{9}{100(2\gamma+1)H_n}$ -fraction of the optimal social welfare.  $\square$

LEMMA 3.3. *For any  $\gamma > 10/9$ , if Mechanism 1 reaches line 13, then the revenue of active bidders in the highest revenue set before running WFCA is at least a  $\frac{9}{10(2\gamma+1)H_n}$ -approximation of the maximum social welfare obtainable by rejected unpredicted bidders, i.e.,*

$$\left( \frac{10(2\gamma+1)H_n}{9} \right) \max_{F \in \mathcal{F}} \text{rev}(F \cap N, \mathbf{p}) \geq \max_{F \in \mathcal{F}} v(F \setminus N).$$

*Proof.* As before we consider the iteration  $\hat{t}$  in which the while loop of Mechanism 1 hit Line 13. By assumption, our while loop terminates in line 13. Following the same reasoning as Lemma 3.2, the total welfare rejected from any unpredicted set over all rounds from 1 to  $\hat{t}$  is at most

$$\sum_{i'=1}^{\hat{t}} R_{i'} \cdot (2\gamma+1)H_n \leq R_{\hat{t}} \cdot \frac{(2\gamma+1)10H_n}{9}.$$

On the other hand, because we arrive at line 13 we know that there exists some set unpredicted set which obtains revenue  $R_{\hat{t}}$  in round  $\hat{t}$ . Thus, the revenue among active bidders in the highest revenue set when we reach line 13 is a  $\frac{9}{10(2\gamma+1)H_n}$ -fraction of the maximum social welfare obtained from any feasible set of rejected unpredicted bidders.  $\square$

LEMMA 3.4. *For any  $\gamma > 10/9$ , if Mechanism 1 outputs a subset of some unpredicted sets, then it is  $\frac{9}{10(2\gamma+1)H_n} + \frac{1}{2H_n}$ -robust.*

*Proof.* In case Mechanism 1 outputs a subset of unpredicted bidders it must arrive at line 13. As such, we can decompose the social welfare from the optimal (unpredicted) set into two components, the portion of the welfare coming from bidders rejected before Mechanism 1 runs the Water-filling Clock Auction and the portion of the welfare coming from bidders who participate in the WFCA. From Lemma 3.3 we know that the *revenue* reached by Mechanism 1 before running the WFCA is within a  $\frac{9}{10(2\gamma+1)H_n}$ -factor of the welfare rejected from the optimal set before running the WFCA. Moreover, from Lemma 2.1 we have that the revenue reached by Mechanism 1 after running the WFCA is weakly higher than the revenue reached before running the WFCA. As such, since the welfare obtained by serving a set of bidders is always weakly higher than the revenue collected from these bidders we have that the social welfare obtained by Mechanism 1 is within a  $\frac{9}{10(2\gamma+1)H_n}$ -factor of the welfare rejected from the optimal set before running the WFCA. Finally, we have from Lemma 2.1 that the social welfare achieved by running the WFCA only on bidders who are active when line 13 is reached is within a  $1/(2H_n)$ -factor of the optimal social welfare achievable from these bidders. Combining these two guarantees, we have that the social welfare obtained by Mechanism 1 is within a  $\frac{9}{10(2\gamma+1)H_n} + \frac{1}{2H_n}$ -factor of the optimal social welfare.  $\square$

**Consistency Argument.** We now establish the consistency of Mechanism 1. The high-level idea is to upper bound the total rejected welfare when the mechanism terminates and argue that: (i) If the predicted set is indeed optimal, the mechanism will always output a subset of the predicted set; and (ii) the remaining welfare in the predicted set (when it is output) is high as a result of us “favoring” the predicted set in Mechanism 1.

LEMMA 3.5. *The total welfare rejected from the predicted set up to the end of the  $t$ -th iteration of the while loop is at most  $R_t \cdot 10H_n/9$*

LEMMA 3.6. *For any  $\gamma > 10/9$ , Mechanism 1 is  $\left(1 - \frac{10}{9\gamma}\right)$ -consistent.*

*Proof.* Fix an instance  $I$  where the optimal set (also the predicted set by definition of consistency) has social welfare  $V^*$ . Consider the minimum value  $k$  such that  $10^k \cdot \gamma H_n$  exceeds  $V^*$ . Let  $\hat{t}$  denote the value of  $t$  at the beginning of the final iteration of the while loop of Mechanism 1 when run on instance  $I$ . We first note that the Mechanism 1 terminates in either Line 8 or Line 10, otherwise there exist a unpredicted set with value more than  $V^*$ . Since we only continue to line 9 after the preceding if statement when there are still active unpredicted bidders it must be that  $\hat{t} \leq k$  (as, otherwise, some unpredicted set would have total welfare greater than  $V^*$ ) and  $\hat{t} = k$  if we terminate in Line 8 (i.e., before we cause any additional predicted bidders to exit the auction in the iteration  $k$ ).

Now consider the total amount of welfare lost from the predicted set from rounds 1 to  $\hat{t} - 1$ . By Lemma 3.5 we have that this total welfare is at most  $R_{\hat{t}-1} \cdot 10H_n/9 = 10^{\hat{t}-1} \cdot \frac{10H_n}{9} \leq 10^{k-1} \cdot \frac{10H_n}{9}$ . On the other hand, by our definition of  $k$ , we know that the total social welfare in the predicted set is at least  $10^{k-1} \cdot \gamma H_n$ . But then, at the end of the  $\hat{t} - 1$  iteration there must exist active predicted bidders and, moreover, the fraction of the predicted set's social welfare that we retain is at least

$$\frac{10^{k-1} \cdot \gamma H_n - 10^{k-1} \cdot \frac{10}{9} H_n}{10^{k-1} \cdot \gamma H_n} = \frac{\gamma - \frac{10}{9}}{\gamma},$$

which completes the proof.  $\square$

We are now ready to prove the main theorem of the section.

*Proof. (Proof for Theorem 3.1)* First observe that  $\epsilon > 0$  ensures that  $\gamma > 10/9$ . In the case that the predicted set, indeed, optimal, Lemma 3.6 gives that some subset of the predicted set is served. Moreover, we obtain that the total social welfare of served bidders is within a factor

$$\left(1 - \frac{10}{9\gamma}\right) = \left(1 - \frac{10}{9\frac{10(1+\epsilon)}{9\epsilon}}\right) = \frac{1}{1+\epsilon}$$

of the total social welfare in the predicted set, thereby guaranteeing  $(1 + \epsilon)$ -consistency. On the other hand, if some subset of the predicted set is served, Lemma 3.2 guarantees that the total social welfare obtained by Mechanism 1 is within a factor  $\frac{9}{100(2\gamma+1)H_n}$  of the optimal social welfare. Since  $H_n = O(\log n)$  and  $\gamma$  is a constant for any choice of constant  $\epsilon$  we obtain  $O(\log n)$ -robustness. Finally, if some subset of an unpredicted set is served, Lemma 3.4 guarantees that Mechanism 1 obtains welfare within a factor  $\frac{9}{10(2\gamma+1)H_n} + \frac{1}{2H_n}$  of the optimal welfare. Again, since  $H_n = O(\log n)$  and  $\gamma$  is constant for any constant  $\epsilon$  we obtain  $O(\log n)$ -robustness in this case.  $\square$

While the focus of our work is on leveraging advice to circumvent the informational limitations faced by clock auctions (recall that the lower bound from [24] precludes clock auctions from achieving a  $\log^\tau(n)$ -approximation for any constant  $\tau < 1$ ), FOLLOWTHEUNPREDICTEDLEADER could still require exponential time for certain feasibility constraints (e.g., when  $\mathcal{F}$  corresponds to the independent sets of a graph). Fortunately, a key advantage of FOLLOWTHEUNPREDICTEDLEADER is that it admits the use of a black-box approximation algorithm for the welfare and revenue maximization steps in lines 7 and 9. Specifically, if there exists an algorithm that outputs a set of rejected bidders which gives a  $\rho$ -approximation to the welfare of the best-rejected set in line 7 and an active set which gives  $\rho$ -approximation to the revenue in line 9, then our algorithm achieves robustness that scales linearly with  $\rho$  (i.e., robustness becomes  $O(\rho \cdot \frac{1}{\epsilon} \log n)$ ). This result is formalized in Corollary 3.1. For completeness, a full proof is provided in Appendix A.

**COROLLARY 3.1.** *For any downward-closed set system, Mechanism 1 with parameter  $\gamma = \frac{10(1+\epsilon)}{9\epsilon}$  equipped with a  $\rho$ -approximate algorithm for the underlying problem of maximizing revenue/welfare (at some fixed clock prices) is  $(1 + \epsilon)$ -consistent and  $O(\rho \cdot \frac{1}{\epsilon} \cdot \log n)$ -robust for any constant  $\epsilon > 0$ .*

**3.1 Lower Bound.** We now show that apart from achieving the “best of both worlds” in terms of asymptotic robustness and consistency guarantees, our learning-augmented clock auction also achieves a tight asymptotic dependence on the  $\epsilon$  parameter. Specifically, we show that there exists a family of instances for which any  $(1 + \epsilon)$ -consistent auction has robustness at least  $\Omega(\frac{1}{\epsilon} \log n)$  (the proof is deferred to Appendix B).

**THEOREM 3.2.** *The robustness of any deterministic clock auction that is augmented with a prediction  $\hat{\text{OPT}}$  and satisfies  $(1 + \epsilon)$ -consistency for some constant  $\epsilon > 0$  is  $\Omega(\frac{1}{\epsilon} \log n)$ .*

**3.2 The Error-Tolerant Auction.** We now demonstrate that Mechanism 1 can be extended to achieve a constant approximation not only when the prediction is optimal, but also when the predicted set deviates from the optimal set by a constant factor in terms of value. To this end, we define the *prediction error* of a prediction  $\hat{\text{OPT}}$  in a given instance  $I$  as the multiplicative difference between the value of the predicted set and the value of the actual optimal set:

$$(3.3) \quad \eta(\hat{\text{OPT}}, I) = \frac{v(\text{OPT}(I))}{v(\hat{\text{OPT}})}.$$

This type of error is called the *quality of recommendation* by [22]. Next, we introduce the ERRORTOLERANT auction, which is an extension of Mechanism 1. This auction takes as input a parameter  $\bar{\eta} \geq 1$ , referred to as the error-tolerance parameter, which is chosen by the mechanism designer. The only modification from Mechanism 1 to ERRORTOLERANT is that in Line 11, the targeted revenue is adjusted from  $R_t$  to  $R_t/\bar{\eta}$ . When the prediction error  $\eta$  is at most  $\bar{\eta}$ , ERRORTOLERANT achieves an approximation of  $(1+\epsilon)\eta$  to the optimal welfare. Additionally, it maintains an approximation of  $\left(\frac{2000\bar{\eta}(1+\epsilon)}{9\epsilon} + 1\right)H_n$  to the optimal welfare at all times. Note that if  $\bar{\eta} = 1$ , the auction reduces to Mechanism 1, preserving the corresponding consistency and robustness guarantees. The formal description of ERRORTOLERANT and the proof of the following theorem are deferred to Appendix C.

**THEOREM 3.3.** *For any downward-closed set system and predicted optimal set, ERRORTOLERANT with parameter  $\gamma = \frac{10(1+\epsilon)}{9\epsilon}$  and  $\bar{\eta} \geq 1$  achieves an approximation of  $(1+\epsilon)\eta$  if  $\eta \leq \bar{\eta}$  and  $\left(\frac{2000\bar{\eta}(1+\epsilon)}{9\epsilon} + 1\right)H_n$  otherwise, where  $\eta$  is the prediction error defined in (3.3).*

#### 4 Clock Auction Guaranteeing a Stronger Notion of Consistency

In this section, we turn to the more demanding notion of consistency $^\infty$ , which requires a good approximation of the predicted solution irrespective of the prediction quality. We show that achieving good consistency $^\infty$  is much harder than achieving good consistency. In particular, there is a separation in terms of the robustness that one can achieve for each of these two notions of consistency. Specifically, we show that if one wants  $\alpha$ -consistency $^\infty$  for some constant  $\alpha$ , not only is it impossible to achieve  $O(\log n)$  robustness, as we achieved for the standard notion of consistency in the previous section, but one must actually suffer  $\Omega\left(n^{\frac{1}{\alpha-1}} \log n\right)$  robustness. Before proving this separation, we provide a learning-augmented clock auction that matches this bound for any constant  $\alpha$ .

**THEOREM 4.1.** *For any downward-closed set system, Mechanism 2 with parameters  $\alpha \in [1+\epsilon, H_n]$  and*

$$\beta \geq \frac{8H_n \Gamma\left(n + \frac{\alpha}{\alpha-1}\right)}{\Gamma\left(1 + \frac{\alpha}{\alpha-1}\right) n!} - \frac{4H_n(\alpha-1)}{\alpha} = \Theta\left(n^{\frac{1}{\alpha-1}} \log n\right),$$

where  $\Gamma(x) = \int_0^\infty t^{x-1} e^{-t} dt$ , is  $\alpha$ -consistent $^\infty$  and  $\beta$ -robust.

The clock auction defined in Mechanism 2, named FOLLOWTHEBINDINGBENCHMARK, also alternates between uniformly raising the bids of the unpredicted and predicted bidders while maintaining a careful balance between the revenue from the predicted set and the welfare of the unpredicted set. However, to achieve consistency $^\infty$ , it requires that in every round the revenue collected from the predicted set be at least  $\frac{1}{\alpha}$  of the rejected welfare of the unpredicted set plus the current revenue, i.e., the lower bound of the welfare of the predicted set. Unless this constraint is satisfied, the clock auction will continue raising the bids for the predicted set. See Figure 1 for a demonstration of the asymptotic trade-off the mechanism obtains for different sizes of instances.

As before, we assume that  $\mathcal{F}$  is a set system where the predicted set is completely disjoint from all maximal unpredicted sets; by Observation 3.1, this is without loss of generality. We first prove the robustness guarantee. All the missing proofs can be found in Appendix D.

**Robustness Argument.** To demonstrate that Mechanism 2 achieves the desired robustness we consider two cases depending on whether or not the auction terminates by serving a subset of the predicted set or a subset of some unpredicted set. Since we assume that  $\alpha < \beta$  (and we will eventually show that our mechanism achieves social welfare within a  $\alpha$ -factor of the welfare in the predicted set), we show that our mechanism obtains welfare within a  $\beta$ -factor of the best unpredicted set. Building toward this guarantee, we first prove a useful lemma regarding the total welfare lost from an unpredicted set.

**LEMMA 4.1.** *The total welfare rejected in any unpredicted set  $S$  in some iteration  $t$  is at most  $\frac{\beta}{2} \cdot R_{t-1}^P$ .*

As a corollary of this lemma, we can bound the total welfare lost from any unpredicted set of bidders by observing that the revenue and welfare targets increase by at least a factor 2 in each round.

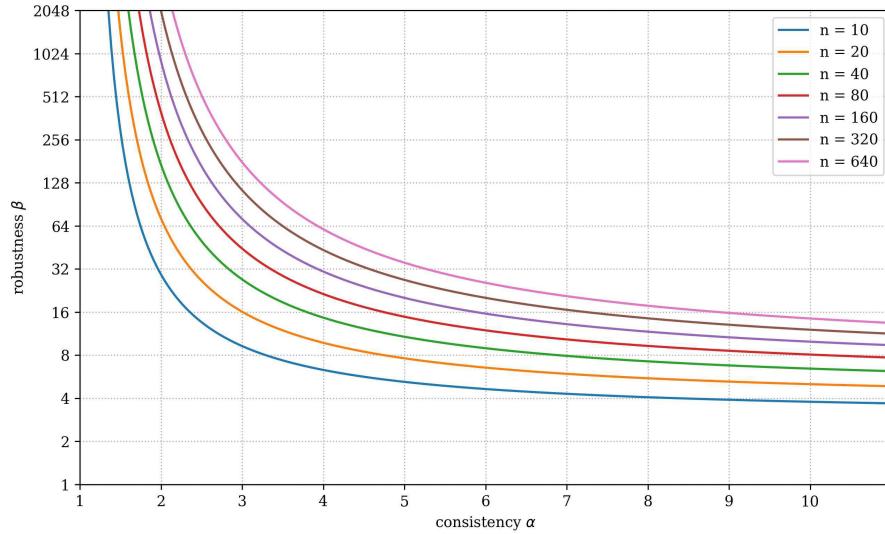


Figure 1: The trade-off between robustness and consistency $^\infty$  when  $\beta = n^{1/(\alpha-1)} H_n$  for different values of  $n$

**COROLLARY 4.1.** *The total welfare rejected in any unpredicted set  $S$  throughout the first  $t'$  iterations is at most  $\beta R_{t'-1}^P$ .*

We are now ready to prove the robustness guarantees of Mechanism 2. Similarly as in Section 3 we will analyze the two cases which depend on whether or not a subset of  $\hat{OPT}$  is output separately. We first consider the case where a subset of  $\hat{OPT}$  is output. The intuition is that we only serve a subset of predicted bidders if we do not satisfy either condition in Line 7 but we secured good revenue from the predicted set in the previous iteration of the while loop. We formalize this intuition in the following lemma.

**LEMMA 4.2.** *If Mechanism 2 outputs a subset of the predicted set, then Mechanism 2 is  $\beta$ -robust.*

*Proof.* Let  $\hat{t}$  denote the value of  $t$  when the auction terminates. By assumption that we serve a subset of the predicted set and definition of the mechanism, this was a result of not satisfying either condition in Line 7. Note that we proceeded to round  $\hat{t}$  because we satisfied the conditions of Line 11 in round  $\hat{t} - 1$ . As such, the welfare remaining in the the predicted set is at least  $R_{\hat{t}-1}^P$  (the revenue of the predicted set at the end of the previous round). On the other hand, we know that the total welfare in any unpredicted set (since all unpredicted bidders were rejected) is at most  $\beta R_{\hat{t}-1}^P$  by Corollary 4.1. Therefore we obtain welfare within a  $\beta$  fraction of the optimal unpredicted set.  $\square$

We now turn toward the case where a subset of an unpredicted set is output. As in Section 3, we consider the welfare contributions from bidders who exit the auction before the WFCA and those active during the WFCA separately.

**LEMMA 4.3.** *If Mechanism 2 outputs a subset of some unpredicted set, then Mechanism 2 is  $\beta$ -robust.*

*Proof.* Let  $\hat{t}$  denote the value of  $t$  when the auction terminates. By assumption that we serve a subset of the unpredicted set and definition of the mechanism, the auction terminates in Line 12. We may divide the welfare of the best unpredicted set into two components – the portion of the welfare coming from bidders rejected during before the WFCA is run and the portion of the welfare coming from bidders which are active when the WFCA is run in Line 12. By Corollary 4.1 we have that the total welfare rejected in any unpredicted set through the first  $\hat{t}$  iterations is  $\beta R_{\hat{t}-1}^P$ . On the other hand, we have that the revenue of the best unpredicted set before we run Line

---

**MECHANISM 2: FOLLOWTHEBINDINGBENCHMARK**


---

**Input:** A system  $\mathcal{F}$  of feasible sets of bidders, a predicted optimal set  $\hat{OPT}$ , a consistency parameter  $\alpha$ , a robustness parameter  $\beta$

```

1  $F \leftarrow F \setminus \{F \cap \hat{OPT}\}$  for all  $F \neq \hat{OPT} \in \mathcal{F}$            // Make  $\hat{OPT}$  disjoint from all unpredicted sets
2  $p_i \leftarrow v_{\min}$  for all  $i \in N$            // Raise price of each bidder to the initial small amount
3  $R_0^P \leftarrow \text{rev}(\hat{OPT}, \mathbf{p})$ 
4  $S \leftarrow \hat{OPT}$ 
5  $t \leftarrow 1$ 
6 while TRUE do
7   Run UNIFORMPRICE( $N \setminus \hat{OPT}$ ,  $\mathbf{p}$ ) until: (i)  $\max_{F \in \mathcal{F}} \text{rev}(F \cap N, \mathbf{p}) \geq \frac{\beta}{4} \cdot R_{t-1}^P$ ; OR (ii)
         $\max_{F \in \mathcal{F}} v(F \setminus N) \geq \frac{\beta}{4} \cdot R_{t-1}^P$  and  $\max_{F \in \mathcal{F}} \text{rev}(F \cap N, \mathbf{p}) \geq \frac{\beta}{4H_n} \cdot R_{t-1}^P$ 
8   if  $N = \hat{OPT}$  then return  $\hat{OPT}, \mathbf{p}$            // All unpredicted bidders reject case
9
10   $\tilde{R}_t^P \leftarrow R_{t-1}^P \cdot 2$ 
11  Run UNIFORMPRICE( $\hat{OPT}$ ,  $\mathbf{p}$ ) until: (i)  $\text{rev}(\hat{OPT}, \mathbf{p}) \geq \tilde{R}_t^P$  AND (ii)  $(\alpha - 1) \cdot \text{rev}(\hat{OPT}, \mathbf{p}) \geq v(S \setminus \hat{OPT})$ 
12  if  $\hat{OPT} = \emptyset$  then return WFCA( $N, \mathbf{p}$ )           // All predicted bidders reject case
13
14   $R_t^P \leftarrow \text{rev}(\hat{OPT}, \mathbf{p})$ 
15   $t \leftarrow t + 1$ 

```

---

12 is at least  $\frac{\beta}{4H_n} \cdot R_{t-1}^P$ . In addition, by Lemma 2.1 we know that the revenue after completing Line 12 must then be at least  $\frac{\beta}{4H_n} \cdot R_{t-1}^P$ . This gives that the revenue collected by Mechanism 2 is within a  $4H_n$ -factor of the welfare in the best unpredicted set contributed from bidders rejected before Line 12 is run. Moreover, we have by Lemma 2.1 that the welfare achieved by our auction is within a  $2H_n$ -factor of the welfare in the best unpredicted set contributed from bidders who are active when Line 12 is run. Combining these we have that if a subset of the unpredicted bidders is served then Mechanism 2 obtains welfare within a factor  $4H_n + 2H_n \leq \beta$  of the best unpredicted set.  $\square$

**Consistency Argument.** We now turn toward the more technically challenging consistency $^\infty$  guarantee. Recall again that a mechanism achieves  $\alpha$ -consistency $^\infty$  if it achieves an  $\alpha$ -approximation to the value of the *predicted* set, irrespective of the prediction quality. To this end, we first resolve the simpler case where we serve some subset of predicted set. In this case the guarantee is achieved by design, but we include a formal proof of this for completeness in Appendix D.

LEMMA 4.4. *If Mechanism 2 outputs a subset of the predicted set, then Mechanism 2 is  $\alpha$ -consistent $^\infty$ .*

The more difficult case involves ensuring  $\alpha$ -consistency $^\infty$  given that we output a subset of some unpredicted set. Intuitively we want to argue that by defining a large enough  $\beta$ , the revenue we obtained from the best unpredicted set is enough to “cover” the maximum potential welfare the predicted set can contain. To this end, we begin with some individual bounds of bidders’ values in the predicted set based on when they rejected the price offer.

LEMMA 4.5. *Consider any iteration  $t$ . At the start of running Line 11, condition (ii) is true. In addition, while (ii) remains true, the current price offered to the bidders is less than  $\tilde{R}_t^P/k$  where  $k$  is the number of active bidders in the predicted set.*

*Proof.* We begin by demonstrating the first part of the lemma statement. First observe that, by design, we satisfied both conditions (i) and (ii) of Line 11 at the end of iteration  $t - 1$ . Furthermore, since we do not raise any prices for predicted bidders in iteration  $t$  before reaching Line 11 again, it must be that in iteration  $t$  we satisfy condition (ii) at the start of Line 11. For the second part of the lemma statement, suppose that condition (ii) of Line 11 is satisfied and  $k$  bidders are active. If we continue to uniformly raise prices for these  $k$  bidders,

then it must be we do not satisfy condition (i). As such, it must be that these bidders have not yet accepted a price of  $\tilde{R}_t^P/k$ .  $\square$

CLAIM 4.1. *Fix a set of bidders  $S$  and index bidders in  $S$  in non-increasing value order. Two conditions hold:*

1. *Suppose that  $(\alpha - 1)kv_k \leq \sum_{i=k+1}^{|S|} v_i$ . Then, if  $(\alpha - 1)(k - 1)v_{k-1} \geq \sum_{i=k}^{|S|} v_i$  we have that*

$$v_{k-1} \geq \frac{(\alpha - 1)k + 1}{(\alpha - 1)(k - 1)} v_k.$$

2. *Suppose that  $(\alpha - 1)kv_k \geq \sum_{i=k+1}^{|S|} v_i$ . Then, if  $(\alpha - 1)(k - 1)v_{k-1} \leq \sum_{i=k}^{|S|} v_i$  we have that*

$$v_{k-1} \leq \frac{(\alpha - 1)k + 1}{(\alpha - 1)(k - 1)} v_k.$$

*Proof.* We show part 1 of the claim as part 2 follows, symmetrically, by flipping the inequalities. By assumption, we have that  $\sum_{i=k+1}^{|S|} v_i \geq (\alpha - 1)kv_k$  so substituting gives

$$(\alpha - 1)(k - 1)v_{k-1} \geq \sum_{i=k}^{|S|} v_i \geq v_k + (\alpha - 1)kv_k,$$

or, equivalently,

$$v_{k-1} \geq \frac{(\alpha - 1)k + 1}{(\alpha - 1)(k - 1)} v_k,$$

as desired.  $\square$

LEMMA 4.6. *Consider any point when we are running UNIFORMPRICE on the predicted set (i.e., we are running Line 11) where condition (ii) is satisfied and let  $k$  denote the number of active bidders in the predicted set. If the smallest valued predicted bidder with value  $p_k$  rejecting causes condition (ii) to be violated, then if condition (ii) is not satisfied before the  $k - 1$ -th largest bidder exits the auction her value  $p_{k-1}$  is upper bounded by*

$$(4.4) \quad p_{k-1} < \frac{(\alpha - 1)k + 1}{(\alpha - 1)(k - 1)} \cdot p_k.$$

*Proof.* Fix some point in the auction when the revenue of the predicted set satisfied  $\alpha$ -consistency $^\infty$ , i.e.,  $(\alpha - 1)$  times the current predicted revenue is greater than or equal to the rejected welfare in the predicted set, and let  $k$  denote the number of remaining active predicted bidders and  $\bar{W}$  denote the total rejected welfare at this point. Let  $p_k$  denote the price offered to these  $k$  active bidders just before one of them exits the auction. We then have  $(\alpha - 1) \cdot kp_k \geq \bar{W}$ . If the lowest valued bidder among these  $k$  exiting the auction causes  $\alpha$ -consistency $^\infty$  to be violated, i.e.,  $\alpha \cdot (k - 1)p_k < p_k + \bar{W}$  we may obtain an upper bound on the price  $p_{k-1}$  corresponding to the point at which  $\alpha$ -consistency $^\infty$  would be satisfied again if  $k - 1$  bidders accept  $p_{k-1}$ . By applying part 2 of Claim 4.1, we obtain that

$$p_{k-1} < \frac{(\alpha - 1)k + 1}{(\alpha - 1)(k - 1)} \cdot p_k,$$

as desired.  $\square$

OBSERVATION 4.1. *For any  $k > 2$  and  $\alpha > 1$  we have that*

$$\frac{1}{k-1} \leq \frac{1}{k} \cdot \frac{(\alpha - 1)k + 1}{(\alpha - 1)(k - 1)} = \frac{1}{k-1} + \frac{1}{k(\alpha - 1)(k - 1)}.$$

With Lemma 4.6 and Observation 4.1 in hand we can now move to show that the auction achieves  $\alpha$ -consistency $^\infty$  when a subset of the unpredicted set is output provided we have  $\beta$  large enough. A crucial function in both the upper and lower bounds we demonstrate is  $\Gamma(x)$  where  $\Gamma(x) = \int_0^\infty t^{x-1} e^{-t} dt$  is the extension of the factorial function to the complex numbers, i.e.,  $\Gamma(k) = (k - 1)!$  for  $k \in \mathbb{N}$ .

LEMMA 4.7. If Mechanism 2 with parameters  $\alpha \in [1 + \epsilon, H_n]$  and

$$\beta \geq \frac{8H_n \Gamma\left(n + \frac{\alpha}{\alpha-1}\right)}{\Gamma\left(1 + \frac{\alpha}{\alpha-1}\right) n!} - \frac{4H_n(\alpha-1)}{\alpha} = \Theta\left(n^{\frac{1}{\alpha-1}} \log n\right)$$

outputs a subset of some unpredicted set, then it is  $\alpha$ -consistent $^\infty$ .

*Proof.* Consider the round  $\hat{t}$  in which the auction terminates and assume that the auction proceeds to Line 12, i.e., a subset of the unpredicted bidders is served. By Lemma 4.5, note that the revenue  $R_{\hat{t}-1}^P$  reached by the predicted set in iteration  $\hat{t}-1$  was such that  $(\alpha-1) \cdot R_{\hat{t}-1}^P$  was more than the total rejected welfare in the predicted set throughout the first  $\hat{t}-1$  iterations of the while loop, i.e., the iteration  $\hat{t}-1$  concluded with the predicted set satisfying  $\alpha$ -consistency $^\infty$ . On the other hand, since we proceeded past Line 7 we have that the revenue in the unpredicted set reached at least  $\frac{\beta}{4H_n} \cdot R_{\hat{t}-1}^P$ . We just then want to bound the total welfare lost from the predicted set in the final iteration  $\hat{t}$ . By Lemma 4.5 we have that the first rejected bidder has value at most  $2R_{\hat{t}-1}^P/k$  where  $k$  is the number of active predicted bidders at the start of iteration  $\hat{t}$ . Applying Lemma 4.6 and Observation 4.1 to upper bound the value  $v_{k-1}$  of the next rejecting bidder, we obtain that

$$v_{k-1} \leq \frac{2R_{\hat{t}-1}^P}{k} \cdot \frac{(\alpha-1)k+1}{(\alpha-1)(k-1)}.$$

Iteratively applying Lemma 4.6 and Observation 4.1 to find an upper bound to each rejected bidder's value, we have that the total welfare rejected in this round is at most

$$\begin{aligned} \frac{2R_{\hat{t}-1}^P}{k} + \sum_{i=2}^k \frac{2R_{\hat{t}-1}^P}{k} \cdot \prod_{j=i}^k \frac{(\alpha-1)j+1}{(\alpha-1)(j-1)} &\leq \frac{2R_{\hat{t}-1}^P}{n} + \sum_{i=2}^n \frac{2R_{\hat{t}-1}^P}{n} \cdot \prod_{j=i}^n \frac{(\alpha-1)j+1}{(\alpha-1)(j-1)} \\ &= \frac{2R_{\hat{t}-1}^P}{n} + \frac{2R_{\hat{t}-1}^P}{n} \cdot \sum_{i=2}^n \prod_{j=i}^n \frac{j+1/(\alpha-1)}{j-1} \\ &= \frac{2R_{\hat{t}-1}^P}{n} + \frac{2R_{\hat{t}-1}^P}{n} \cdot \sum_{i=2}^n \frac{\Gamma(i-1)\Gamma\left(n+1+\frac{1}{\alpha-1}\right)}{\Gamma\left(\frac{1}{\alpha-1}+i\right)\Gamma(n)} \\ &= \frac{2R_{\hat{t}-1}^P}{n} + \frac{2R_{\hat{t}-1}^P}{n} \cdot \left( -\alpha n + \frac{\alpha\Gamma\left(\frac{n\alpha+\alpha-n}{\alpha-1}\right)}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)\Gamma(n)} + n-1 \right) \\ &= \frac{2R_{\hat{t}-1}^P}{n} \cdot \left( -\alpha n + \frac{\alpha\Gamma\left(\frac{n\alpha+\alpha-n}{\alpha-1}\right)}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)(n-1)!} + n \right), \end{aligned}$$

where the inequality comes from the fact that for any  $\alpha > 1$  the partial product from  $j = k+1$  to  $n$  times  $2R_{\hat{t}-1}^P/n$  yields a value at least  $2R_{\hat{t}-1}^P/k$ , the second equality comes from the fact that  $z \cdot \Gamma(z) = \Gamma(z+1)$  for all  $z \in \mathbb{R}$  and, hence,  $\Gamma(x+y)/\Gamma(x) = x \cdot (x+1) \cdot \dots \cdot (x+y-1)$  for all  $y \in \mathbb{N}$  and all  $x \in \mathbb{R}$ , and the third equality is due to Lemma E.1 which we defer to Appendix E.

Combining the welfare rejected from the first  $\hat{t}-1$  rounds and the welfare from the last round we obtain that the total welfare rejected from the predicted set is at most

$$(\alpha-1) \cdot R_{\hat{t}-1}^P + \frac{2R_{\hat{t}-1}^P}{n} \cdot \left( -\alpha n + \frac{\alpha\Gamma\left(\frac{n\alpha+\alpha-n}{\alpha-1}\right)}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)(n-1)!} + n \right).$$

But then, since our auction reached revenue at least  $\frac{\beta}{4H_n} \cdot R_{\hat{t}-1}^P$  before running the WFCA and the WFCA is

revenue monotone by Lemma 2.1 it suffices that

$$\alpha \frac{\beta}{4H_n} \geq (\alpha - 1) + \frac{2}{n} \cdot \left( -\alpha n + \frac{\alpha \Gamma\left(\frac{n\alpha+\alpha-n}{\alpha-1}\right)}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)(n-1)!} + n \right)$$

or, equivalently,

$$\beta \geq \frac{4H_n}{\alpha} \cdot \left( (\alpha - 1) + \frac{2}{n} \cdot \left( -\alpha n + \frac{\alpha \Gamma\left(\frac{n\alpha+\alpha-n}{\alpha-1}\right)}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)(n-1)!} + n \right) \right).$$

Therefore, for the proof to go through it suffices for  $\beta$  to satisfy

$$\beta \geq \frac{4H_n}{\alpha} (\alpha - 1 - 2\alpha + 2) + \frac{8H_n \Gamma\left(n + \frac{\alpha}{\alpha-1}\right)}{\Gamma\left(1 + \frac{\alpha}{\alpha-1}\right) n!} = \frac{8H_n \Gamma\left(n + \frac{\alpha}{\alpha-1}\right)}{\Gamma\left(1 + \frac{\alpha}{\alpha-1}\right) n!} - \frac{4H_n(\alpha - 1)}{\alpha}.$$

As we care about the asymptotic behavior of this bound and  $n + \frac{\alpha}{\alpha-1}$  is a positive real number, we can apply Stirling's approximation, i.e.,

$$\Gamma(z) = \sqrt{\frac{2\pi}{z}} \left(\frac{z}{e}\right)^z \left(1 + O\left(\frac{1}{z}\right)\right),$$

to obtain that the proof goes through with

$$\begin{aligned} \beta &= \frac{8H_n}{\Gamma\left(1 + \frac{\alpha}{\alpha-1}\right)} \cdot \frac{\sqrt{\frac{2\pi}{n+1+1/(\alpha-1)}} \left(\frac{n+1+1/(\alpha-1)}{e}\right)^{n+1+1/(\alpha-1)} \left(1 + O\left(\frac{1}{n+1+1/(\alpha-1)}\right)\right)}{\sqrt{\frac{2\pi}{n+1}} \left(\frac{n+1}{e}\right)^{n+1} \left(1 + O\left(\frac{1}{n+1}\right)\right)} \\ &= \frac{8H_n}{\Gamma\left(1 + \frac{\alpha}{\alpha-1}\right)} \cdot O(1) \cdot \left(\frac{n+1+1/(\alpha-1)}{e}\right)^{n+1+1/(\alpha-1)} \cdot \left(\frac{e}{n+1}\right)^{n+1} \\ &= \frac{8H_n}{\Gamma\left(1 + \frac{\alpha}{\alpha-1}\right)} \cdot O(1) \cdot \frac{(n+1+1/(\alpha-1))^{n+1+1/(\alpha-1)}}{(n+1)^{n+1}} \cdot \frac{1}{e^{1/(\alpha-1)}} \\ &= \Theta\left(n^{\frac{1}{\alpha-1}} \log n\right) \end{aligned}$$

where the last line is since  $\Gamma(c)$  is a constant for any constant real number  $c$  and  $H_n \geq \alpha > 1 + \epsilon$  for constant  $\epsilon > 0$  so  $1/(\alpha-1) = O(1)$  and  $\Gamma(1 + \alpha/(\alpha-1)) = O(1)$ .  $\square$

**4.1 Lower Bound.** We now provide a lower bound for all deterministic clock auctions in terms of the trade-off between consistency $^\infty$  and robustness. This lower bound shows that the trade-off achieved by Mechanism 2 is near-optimal; in particular for any constant  $\alpha$ , the trade-off achieved by the mechanism is asymptotically optimal.

**THEOREM 4.2.** *The robustness of any deterministic clock auction that is augmented with a prediction  $\hat{OPT}$  and satisfies  $\alpha$ -consistency $^\infty$  for  $\alpha \in [1 + \epsilon, H_n]$  is  $\Omega(\frac{1}{\alpha} n^{1/(\alpha-1)} \log n)$ .*

Note that for  $\alpha = 2$ , we have  $\beta = \Theta(n \log n)$ . As the consistency $^\infty$  guarantee  $\alpha$  approaches 1, the robustness grows exponentially (e.g., for  $\alpha = 4/3$ ,  $\beta = \Theta(n^3 \log n)$ ). This is in sharp contrast with the standard consistency notion, for which we achieve  $(1 + \epsilon)$ -consistency with  $O(\log n)$ -robustness for any arbitrarily small  $\epsilon$ .

**Implications of the impossibility result on the feasible smoothness guarantees.** Before presenting the proof of this impossibility result, we would like to first point out its implications regarding the “price of smoothness,” i.e., the best achievable robustness if we require that the performance of the auction degrades smoothly as a function of the prediction error. Although a lot of prior work in the learning-augmented framework aims to design mechanisms with smoothness guarantees, to the best of our knowledge, (almost) none of this work provides explicit impossibility results regarding the price of smoothness. In contrast to this paucity of

impossibility results, the robustness lower bounds that we prove for  $\alpha$ -consistent $^\infty$  auctions imply that any clock auction whose performance degrades linearly as a function of the error,  $\eta$ , needs to be  $\Omega(\text{poly}(n))$ -robust.

This bound on the price of smoothness is implied by the following general observation (which holds beyond our setting, as long as the error is measured as in (3.3), termed *quality of recommendation* by [22]: any mechanism whose performance degrades linearly as a function of the error  $\eta$ , e.g., guaranteeing an approximation factor of  $\alpha\eta$  if the prediction error is  $\eta$ , also satisfies  $\alpha$ -consistency $^\infty$ . To verify this, note that the value of a predicted set  $\hat{OPT}$  with error  $\eta$  is  $OPT/\eta$  (by definition). A mechanism with approximation factor  $\alpha\eta$  guarantees welfare  $\frac{OPT}{\alpha\eta} = \frac{\hat{OPT}}{\alpha}$ , which implies a consistency $^\infty$  of  $\alpha$ . This leads to the following Corollary regarding the price of smoothness.

**COROLLARY 4.2.** *If a deterministic clock auction augmented with a prediction  $\hat{OPT}$  guarantees an approximation of  $\alpha\eta$  for every  $\alpha \in [1 + \epsilon, H_n]$ , where  $\eta$  is the prediction error (as defined in (3.3)), then its robustness cannot be better than  $\Omega(\frac{1}{\alpha}n^{1/(\alpha-1)} \log n)$ .*

We now provide the proof of Theorem 4.2.

*Proof. (Proof of Theorem 4.2)* Let  $\alpha$  be the targeted consistency any mechanism is aiming for. Consider a feasibility constraint  $\mathcal{F}$  defined by two disjoint maximal feasible sets  $F_1$  and  $F_2$ , i.e.,  $N = F_1 \cup F_2$ , and a set  $F \subseteq N$  is feasible if and only if  $F \subseteq F_1$  or  $F \subseteq F_2$ . Let the sizes of these sets be  $|F_1| = k_1$  and  $|F_2| = \alpha k_2$ , and let  $F_2$  be the set predicted to be optimal, i.e.,  $\hat{OPT} = F_2$ .

Using this feasibility constraint, we now define a family of instances  $\mathcal{I}$  by defining the set of values that can be assigned to the bidders in  $F_1$  and the bidders in  $F_2$ ; the class of instances in  $\mathcal{I}$  corresponds to all possible assignments of these values to the bidders in the corresponding sets. Let

$$V_1 = \left\{ 1, \frac{1}{2}, \frac{1}{3}, \dots, \frac{1}{k_1} \right\}$$

be a set of  $k_1$  values that can be matched to the  $k_1$  bidders in  $F_1$  and let

$$V_2 = \{v_1, v_2, \dots, v_{k_2}\}$$

be a set of  $k_2$  values that can be assigned to the  $\alpha k_2$  bidders in  $F_2$ . In particular, each value  $v_i$  for  $i \in [1, k_2 - 1]$  is assigned to a single bidder, while  $v_{k_2}$  is assigned to the remaining  $(\alpha - 1)k_2 + 1$  bidders. The values in  $V_2$  are in decreasing in sequence such that the highest value among them is  $v_1 = 1$  for every  $i \in \{2, 3, \dots, k_2\}$ :

$$(4.5) \quad v_i = \frac{(\alpha - 1)(i - 1) + \delta}{(\alpha - 1)i + 1} \cdot v_{i-1},$$

where  $\delta > 0$  is some arbitrarily small constant. Note that as  $\delta \rightarrow 0$  and  $\alpha \rightarrow \infty$ , this sequence of values converges to the harmonic sequence  $\left\{1, \frac{1}{2}, \frac{1}{3}, \dots, \frac{1}{k_2}\right\}$ . However, for smaller values of  $\alpha$  the values in  $V_2$  drop much faster than that. Specifically, the lowest value in  $V_2$  is

$$(4.6) \quad v_{k_2} = \prod_{i=2}^{k_2} \frac{(\alpha - 1)(i - 1) + \delta}{(\alpha - 1)i + 1} = \frac{\Gamma\left(1 + \frac{\alpha}{\alpha-1}\right)(k_2 - 1)!}{\Gamma\left(k_2 + \frac{\alpha}{\alpha-1}\right)} + \delta',$$

where  $\delta' \rightarrow 0$  as  $\delta \rightarrow 0$ .

By Equation (4.5), we have that  $(\alpha - 1) \cdot (k_2 - 1) \cdot v_{k_2-1} < ((\alpha - 1)k_2 + 1)v_{k_2}$ . Note that the right-hand side is the sum of the values of the bidders in  $F_2$  with value  $v_{k_2}$ . Again by Equation (4.5) and part 1 of Claim 4.1 we have, for every  $i \in \{1, 2, \dots, k_2 - 1\}$  it satisfies:

$$(4.7) \quad (\alpha - 1) \cdot i \cdot v_i < ((\alpha - 1)k_2 + 1)v_{k_2} + \sum_{j=i+1}^{k_2-1} v_j.$$

To prove this theorem, we consider any learning-augmented clock auction  $\mathcal{M}$  and simulate it on an instance from  $\mathcal{I}$ , chosen adversarially to maximize the number of bidders that drop out. Note that, since  $\mathcal{M}$  is deterministic,

the adversary can essentially determine who to assign each value to after observing the price trajectory that  $\mathcal{M}$  will follow. Specifically, whenever  $\mathcal{M}$  raises the price of an active bidder  $i \in F_1$  to an amount that exceeds some value  $v \in V_1$  that remains active (i.e., that was not already assigned to a bidder that has dropped out), then the adversary can assign the value  $v$  to  $i$  and, as a result,  $i$  drops out. This implies that, if at any point during the execution of  $\mathcal{M}$  the maximum price offered to an active bidder in  $F_1$  is  $p_1$ , then the value of every active bidder in  $F_1$  is at least  $p_1$  (otherwise the adversary could have assigned that value to the bidder facing price  $p_1$ , causing him to drop out). Using the same argument for  $F_2$ , we can conclude that if at any point during the execution of  $\mathcal{M}$  the maximum price offered to an active bidder in  $F_2$  is  $p_2$ , then the values of all active bidders in  $F_2$  are at least  $p_2$ .

We now partition the set of all clock auctions into three cases, based on their outcome facing the adversarially chosen instance from  $\mathcal{I}$ .

**Case 1 (the auction  $\mathcal{M}$  accepts a subset of  $F_1$ ):** For this case, we show that there exists an instance (not in  $\mathcal{I}$ ) for which the auction fails to achieve  $\alpha$ -consistency $^\infty$ . To verify this, let  $x$  be the number of bidders in  $F_1$  that did not drop out and note that, given the harmonic structure of the values in  $V_1$ , the smallest value among them will be at most  $1/x$ . As we showed above, this means that the maximum price offered to any one of the winning bidders in  $F_1$  was  $p_1 \leq 1/x$ . Now, consider an alternative instance that is identical in terms of the values of all the bidders that dropped out, but has value  $1/x$  for all of the bidders that won. Note that the outcome of the auction would be the same for this alternative instance, since the only difference is regarding values of bidders that are higher than the price offered to them, so the welfare achieved by  $\mathcal{M}$  in this alternative instance would be  $x \cdot \frac{1}{x} = 1$ . However, the social welfare of  $\hat{\mathcal{OPT}} = F_2$  is more than  $\alpha$ , which is implied by Inequality (4.7) for  $i = 1$  since the left-hand side of  $(\alpha - 1)v_1 = \alpha - 1$  is less than the sum of all the remaining values in  $F_2$ , which violates the  $\alpha$ -consistency $^\infty$  constraint.

**Case 2 (the auction  $\mathcal{M}$  accepts a subset of  $\hat{\mathcal{OPT}}$  and raises at least  $(\alpha - 1)k_2 + 1$  bidders' clocks above  $v_{k_2}$ ):** By the fact that we may assign values adversarially to ensure that the first  $(\alpha - 1)k_2 + 1$  bidders whose clocks are raised above  $v_{k_2}$  have value  $v_{k_2}$ , we may assume that all bidders with value  $v_{k_2}$  exit the auction. For this case, we once again show that there exists an instance (not in  $\mathcal{I}$ ) for which the auction fails to achieve  $\alpha$ -consistency $^\infty$ . To verify this, let  $x \leq k_2 - 1$  be the number of bidders in  $\hat{\mathcal{OPT}}$  that did not drop out and note that, by the definition of the values in  $V_2$ , the smallest value among them is  $v_x$ . As we showed above, this means that the maximum price offered to any one of the winning bidders in  $\hat{\mathcal{OPT}}$  was  $p_2 \leq v_x$ . We, once again, consider the alternative instance that is identical in terms of the values of all the bidders that dropped out, but has value  $v_x$  for all the bidders in  $\hat{\mathcal{OPT}}$  that won. The outcome of the auction would not be affected by this change, so the resulting welfare would be  $x \cdot v_x$  but, using Inequality (4.7) for  $i = x$  for any  $x \leq k_2 - 1$ , this violates the  $\alpha$ -consistency $^\infty$  constraint since the left-hand side of Inequality (4.7) is  $(\alpha - 1)$  times the welfare the auction obtains and the right-hand side is the rejected welfare from the predicted set.

**Case 3 (the auction  $\mathcal{M}$  accepts a subset of  $\hat{\mathcal{OPT}}$  and raises fewer than  $(\alpha - 1)k_2 + 1$  bidders' clocks above  $v_{k_2}$ ):** For this case, note that by the adversarial assignment of values, the auction maximizes its welfare if no price offered to a bidder in  $\hat{\mathcal{OPT}}$  exceeded  $v_{k_2}$ , otherwise the adversarial assignment would ensure that the corresponding bidder would have dropped out (and  $\mathcal{M}$  then would not accept all bidders of value  $v_{k_2}$ ). Now consider an alternative instance that is identical with respect to the values of the bidders in  $F_1$ , but the values of everyone in  $F_2$  are equal to  $v_{k_2}$ . Note that the outcome of the auction would be the same in this alternative instance, since the only difference is regarding values of bidders that are higher than the price offered to them. Therefore, using Equation (4.6), we conclude that the welfare achieved by the auction in this instance would be

$$\alpha \cdot k_2 \cdot v_{k_2} = \alpha \cdot k_2 \cdot \frac{\Gamma\left(1 + \frac{\alpha}{\alpha-1}\right) (k_2 - 1)!}{\Gamma\left(k_2 + \frac{\alpha}{\alpha-1}\right)} + \delta' = \frac{\alpha \Gamma\left(1 + \frac{\alpha}{\alpha-1}\right) k_2!}{\Gamma\left(k_2 + \frac{\alpha}{\alpha-1}\right)} + \delta',$$

whereas the welfare of  $F_1$  is equal to  $H_{k_1} = 1 + \frac{1}{2} + \dots + \frac{1}{k_1}$ , leading to a robustness of

$$\beta = \Omega\left(\frac{H_{k_1} \Gamma\left(k_2 + \frac{\alpha}{\alpha-1}\right)}{\alpha \Gamma\left(1 + \frac{\alpha}{\alpha-1}\right) k_2!}\right) = \Omega\left(\frac{1}{\alpha} k_2^{1/(\alpha-1)} \log k_1\right),$$

where the last equation is inferred by following the same steps as we did at the end of the proof of Lemma 4.7. If we let  $k_1 = n^c$  and  $k_2 = \frac{n-n^c}{\alpha}$ , this yields the claimed robustness lower bound.  $\square$

## 5 Conclusion and Further Directions

In this work, we provide the first analysis of clock auctions through the learning-augmented framework. Apart from the standard notion of consistency, we also consider a stronger notion of consistency, which we refer to as consistency $^\infty$ . We prove upper and lower bounds on the trade-off between robustness and each of these notions of consistency. Additionally, we observe that the impossibility results regarding the trade-off between consistency $^\infty$  and robustness also provide insights into the “cost of smoothness,” i.e., the robustness that is achievable if we require the performance to degrade smoothly as a function of the prediction error.

In our setting, the achievable trade-off between consistency and robustness and the achievable trade-off between consistency $^\infty$  and robustness are drastically different. This is in contrast to some other settings (e.g., [1, 17, 13]), where there is no separation between the two. An interesting open question is to identify which kinds of problems exhibit such separation and which do not. Another type of result currently lacking in the literature is the “cost of smoothness”, i.e., the best-possible trade-off between good error-dependence and robustness.

To address the limitations of deterministic clock auctions without prior information, we utilize unreliable predictions. Prior work has also explored Bayesian clock auctions using distributional information on bidders’ values [28] and randomized clock auctions to eliminate the need for prior information altogether [21]. Despite the improvements that these alternative approaches have achieved, there remain large gaps. Determining the tight bounds for Bayesian and randomized clock auctions is one of the most compelling open questions in this line of work. More broadly, the compelling properties of clock auctions make them highly suited to practical applications. Thus, further studying their performance in other auction settings (e.g., procurement auctions and double auction settings), potentially augmented with predictions, is an interesting future direction with important implications for “real-world” auction implementation.

## References

- [1] P. Agrawal, E. Balkanski, V. Gkatzelis, T. Ou, and X. Tan. Learning-augmented mechanism design: Leveraging predictions for facility location. In *EC ’22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*, pages 497–528. ACM, 2022.
- [2] M. Akbarpour and S. Li. Credible auctions: A trilemma. *Econometrica*, 88(2):425–467, 2020.
- [3] A. Antoniadis, T. Gouleakis, P. Kleer, and P. Kolev. Secretary and online matching problems with machine learned advice. *Discret. Optim.*, 48(Part 2):100778, 2023.
- [4] L. Ausubel and P. Milgrom. The lovely but lonely vickrey auction. *Comb. Auct.*, 17, 01 2006.
- [5] Y. Azar, D. Panigrahi, and N. Touitou. Online graph algorithms with predictions. *Proceedings of the Thirty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, 2022.
- [6] M. Babaioff, R. Lavi, and E. Pavlov. Single-value combinatorial auctions and algorithmic implementation in undominated strategies. *Journal of the ACM (JACM)*, 56(1):4, 2009.
- [7] M.-F. Balcan, S. Prasad, and T. Sandholm. Bicriteria multidimensional mechanism design with side information. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS ’23, Red Hook, NY, USA, 2024. Curran Associates Inc.
- [8] E. Balkanski, V. Gkatzelis, and X. Tan. Mechanism design with predictions: An annotated reading list. *SIGecom Exchanges*, 21(1):54–57, 2023.
- [9] E. Balkanski, P. Garimidi, V. Gkatzelis, D. Schoepflin, and X. Tan. Deterministic budget-feasible clock auctions. In *Proceedings of the Thirty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA 2022, USA, 2022. Society for Industrial and Applied Mathematics.
- [10] E. Balkanski, V. Gkatzelis, and G. Shahkarami. Randomized strategic facility location with predictions. *arXiv preprint arXiv:2409.07142*, 2024.
- [11] E. Balkanski, V. Gkatzelis, and X. Tan. Strategyproof scheduling with predictions. In Y. T. Kalai, editor, *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA*, volume 251 of *LIPICS*, pages 11:1–11:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- [12] E. Balkanski, V. Gkatzelis, X. Tan, and C. Zhu. Online mechanism design with predictions. In *Proceedings of the 25th ACM Conference on Economics and Computation (to appear)*, 2024.
- [13] É. Bamas, A. Maggiori, and O. Svensson. The primal-dual method for learning augmented algorithms. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [14] S. Banerjee, V. Gkatzelis, A. Gorokh, and B. Jin. Online nash social welfare maximization with predictions. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms*, SODA 2022. SIAM, 2022.

- [15] Z. Barak, A. Gupta, and I. Talgam-Cohen. MAC advice for facility location mechanism design. *CoRR*, abs/2403.12181, 2024.
- [16] E. Batziou and M. Bichler. Budget-feasible market design for biodiversity conservation: Considering incentives and spatial coordination. *Wirtschaftsinformatik 2023 Proceedings.*, 44., 2023.
- [17] B. Berger, M. Feldman, V. Gkatzelis, and X. Tan. Optimal metric distortion with predictions. In *Proceedings of the 25th ACM Conference on Economics and Computation (to appear)*, 2024.
- [18] M. Bichler, Z. Hao, R. Littmann, and S. Waldherr. Strategyproof auction mechanisms for network procurement. *OR Spectrum*, 42(4):965–994, 2020.
- [19] I. Caragiannis and G. Kalantzi. Randomized learning-augmented auctions with revenue guarantees. In K. Larson, editor, *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, pages 2687–2694. International Joint Conferences on Artificial Intelligence Organization, 8 2024. Main Track.
- [20] Q. Chen, N. Gravin, and S. Im. Strategic facility location via predictions. *arXiv preprint arXiv:2410.07497*, 2024.
- [21] G. Christodoulou, V. Gkatzelis, and D. Schoepflin. Optimal deterministic clock auctions and beyond. In M. Braverman, editor, *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA*, volume 215 of *LIPICS*, pages 49:1–49:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [22] G. Christodoulou, A. Sgouritsa, and I. Vlachos. Mechanism design augmented with output advice. *arXiv preprint arXiv:2406.14165*, 2024.
- [23] R. Colini-Baldeschi, S. Klumper, G. Schäfer, and A. Tsikiridis. To trust or not to trust: Assignment mechanisms with predictions in the private graph model. *CoRR*, abs/2403.03725, 2024.
- [24] P. Dütting, V. Gkatzelis, and T. Roughgarden. The performance of deferred-acceptance auctions. *Math. Oper. Res.*, 42(4):897–914, 2017.
- [25] P. Dütting, S. Lattanzi, R. P. Leme, and S. Vassilvitskii. Secretaries with advice. In P. Biró, S. Chawla, and F. Echenique, editors, *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, pages 409–429. ACM, 2021.
- [26] P. Dütting, I. Talgam-Cohen, and T. Roughgarden. Modularity and greed in double auctions. *Games and Economic Behavior*, 105:59–83, 2017.
- [27] L. Ensthaler and T. Giebe. A dynamic auction for multi-object procurement under a hard budget constraint. *Research Policy*, 43(1):179–189, 2014.
- [28] M. Feldman, V. Gkatzelis, N. Gravin, and D. Schoepflin. Bayesian and randomized clock auctions. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 820–845, 2022.
- [29] K. Fujii and Y. Yoshida. The secretary problem with predictions. *CoRR*, abs/2306.08340, 2023.
- [30] A. Ganesh and J. Hartline. Combinatorial pen testing (or consumer surplus of deferred-acceptance auctions). *arXiv preprint arXiv:2301.12462*, 2023.
- [31] V. Gkatzelis, K. Kollias, A. Sgouritsa, and X. Tan. Improved price of anarchy via predictions. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 529–557, 2022.
- [32] V. Gkatzelis, E. Markakis, and T. Roughgarden. Deferred-acceptance auctions for multiple levels of service. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 21–38, 2017.
- [33] V. Gkatzelis, R. Patel, E. Pountourakis, and D. Schoepflin. Prior-free clock auctions for bidders with interdependent values. In I. Caragiannis and K. A. Hansen, editors, *Algorithmic Game Theory - 14th International Symposium, SAGT 2021, Aarhus, Denmark, September 21-24, 2021, Proceedings*, volume 12885 of *Lecture Notes in Computer Science*, pages 64–78. Springer, 2021.
- [34] H. Huang, K. Han, S. Cui, and J. Tang. Randomized pricing with deferred acceptance for revenue maximization with submodular objectives. In *Proceedings of the ACM Web Conference 2023*, pages 3530–3540, 2023.
- [35] S. Im, R. Kumar, M. M. Qaem, and M. Purohit. Online knapsack with frequency predictions. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 2733–2743, 2021.
- [36] G. Istrate and C. Bonchis. Mechanism design with predictions for obnoxious facility location. *CoRR*, abs/2212.09521, 2022.
- [37] G. Istrate, C. Bonchis, and V. Bogdan. Equilibria in multiagent online problems with predictions. *CoRR*, abs/2405.11873, 2024.
- [38] F. Jarman and V. Meisner. Ex-post optimal knapsack procurement. *Journal of Economic Theory*, 171:35–63, 2017.
- [39] S. H. Jiang, E. Liu, Y. Lyu, Z. G. Tang, and Y. Zhang. Online facility location with predictions. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [40] J. H. Kagel, R. M. Harstad, and D. Levin. Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica: Journal of the Econometric Society*, pages 1275–1304, 1987.
- [41] E. Kevi and K. T. Nguyen. Primal-dual algorithms with predictions for online bounded allocation and ad-auctions

- problems. In S. Agrawal and F. Orabona, editors, *International Conference on Algorithmic Learning Theory, February 20-23, 2023, Singapore*, volume 201 of *Proceedings of Machine Learning Research*, pages 891–908. PMLR, 2023.
- [42] A. Kim. Welfare maximization with deferred acceptance auctions in reallocation problems. In *Algorithms-ESA 2015*, pages 804–815. Springer, 2015.
- [43] S. Li. Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–87, 2017.
- [44] A. Lindermayr and N. Megow. algorithms-with-predictions. <https://algorithms-with-predictions.github.io/>.
- [45] S. Loertscher and L. M. Marx. Asymptotically optimal prior-free clock auctions. *Journal of Economic Theory*, page 105030, 2020.
- [46] P. Lu, Z. Wan, and J. Zhang. Competitive auctions with imperfect predictions. In *Proceedings of the 25th ACM Conference on Economics and Computation (to appear)*, 2024.
- [47] T. Lykouris and S. Vassilvitskii. Competitive caching with machine learned advice. In *International Conference on Machine Learning*, pages 3296–3305. PMLR, 2018.
- [48] M. Mahdian, H. Nazerzadeh, and A. Saberi. Allocating online advertisement space with unreliable estimates. In J. K. MacKie-Mason, D. C. Parkes, and P. Resnick, editors, *Proceedings 8th ACM Conference on Electronic Commerce (EC-2007), San Diego, California, USA, June 11-15, 2007*, pages 288–294. ACM, 2007.
- [49] A. M. Medina and S. Vassilvitskii. Revenue optimization with approximate bid predictions. In I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 1858–1866, 2017.
- [50] P. Milgrom and I. Segal. Deferred-acceptance auctions and radio spectrum reallocation. In M. Babaioff, V. Conitzer, and D. A. Easley, editors, *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*, pages 185–186. ACM, 2014.
- [51] P. Milgrom and I. Segal. Clock auctions and radio spectrum reallocation. *Journal of Political Economy*, 2019.
- [52] M. Mitzenmacher and S. Vassilvitskii. Algorithms with predictions. *Commun. ACM*, 65(7):33–35, 2022.
- [53] M. Purohit, Z. Svitkina, and R. Kumar. Improving online algorithms via ML predictions. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 9684–9693, 2018.
- [54] M. H. Rothkopf, T. J. Teisberg, and E. P. Kahn. Why are vickrey auctions rare? *Journal of Political Economy*, 98(1):94–109, 1990.
- [55] A. Wei and F. Zhang. Optimal robustness-consistency trade-offs for learning-augmented online algorithms. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [56] C. Xu and P. Lu. Mechanism design with predictions. In L. D. Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 571–577. ijcai.org, 2022.

## A Missing Proofs from Section 3

### A.1 Proof of Lemma 3.1

*Proof.* Fix some set  $S$  and consider the process of raising the price offered to bidders in this set as in line 9. If at least  $k$  bidders in  $S$  accept a price of  $R_t/k$  for any  $k$  then we terminate this process (since the total revenue in  $S$  would reach the revenue target of  $R_t$ ). This would happen if the  $k$ -th highest value bidder in  $S$  had value more than  $R_t/k$ . As such, rejecting the  $k'$ -th highest value bidder in iteration  $t$  of the while loop implies that this bidder has value at most  $R_t/k'$ . On the other hand, since there are at most  $n$  bidders contained in set  $S$  the total value of the bidders rejected in iteration  $i$  of the while loop is at most

$$\sum_{k'=1}^n \frac{R_t}{k'} = R_t \cdot H_n,$$

completing the proof.  $\square$

### A.2 Proof of Lemma 3.5

*Proof.* Following the same argument of Lemma 3.1 we first observe that the total welfare rejected from the predicted set in line 11 in any iteration  $t'$  of the while loop of FOLLOWTHEUNPREDICTEDLEADER is at most

$R_{t'} \cdot H_n$ . Taking a sum over all rounds from 1 to  $t$  we then obtain that the total welfare rejected is at most

$$\sum_{t'=1}^t R_{t'} \cdot H_n = H_n \sum_{t'=1}^t 10^{t'} = H_n \cdot \frac{10^{t+1} - 1}{9} \leq R_t \cdot \frac{10H_n}{9},$$

as desired.  $\square$

### A.3 Proof of Corollary 3.1

*Proof.* Suppose we have an algorithm which gives a  $\rho$ -approximation to the underlying algorithmic problem of maximizing welfare/revenue on a fixed set of values. We first argue that the consistency of FOLLOWTHEUNPREDICTEDLEADER when using this approximation algorithm (in place of an exact algorithm) remains  $(1 + \epsilon)$ . To see this, observe that we can exactly compute the revenue within the predicted set in every round in polynomial time since we can simply sum the prices of the bidders in the predicted set. As such, the arguments in Lemma 3.6 still hold when we use an approximation algorithm in Lines 7 and 9 of FOLLOWTHEUNPREDICTEDLEADER on the *unpredicted* sets.

It remains to verify that the robustness of FOLLOWTHEUNPREDICTEDLEADER is  $O(\rho \cdot \frac{1}{\epsilon} \log n)$ . First observe that, by definition of the fact that we are using a  $\rho$ -approximate algorithm in Line 7, the maximum welfare rejected from any unpredicted set in iteration  $t$  of the while loop of Mechanism 1 at the end of line 7 is  $\rho \cdot 2R_t \gamma H_n$ . Similarly, by the same reasoning as in the proof of Lemma 3.1, the total welfare rejected from any unpredicted set in line 9 of Mechanism 1 in iteration  $t$  is at most  $\rho R_t \cdot H_n$  (since  $k$  bidders in any set accepting a price  $\rho R_t / k$  would mean that the best revenue is weakly higher than  $\rho R_t$  and a  $\rho$ -approximate algorithm would find a set with revenue at least  $R_t$ ). Combining these two gives that the total welfare rejected from any unpredicted set in any iteration  $t$  is at most  $R_t \cdot (2\gamma + 1)\rho H_n$ . As in Lemma 3.2, summing over all rounds then gives that the total welfare rejected from any unpredicted set from rounds 1 to  $t$  is at most  $R_t \cdot \frac{(2\gamma + 1)10\rho H_n}{9}$ .

Following similar reasoning to Lemmas 3.2 and 3.4, we consider cases depending on whether a predicted or unpredicted set wins. Let  $\hat{t}$  denote the final iteration of the while loop of Mechanism 1. If Mechanism 1 outputs a subset of the predicted set, observe that the revenue of the predicted bidders is at least  $R_{\hat{t}-1} = \frac{R_{\hat{t}}}{10}$ . But then, since the total welfare in any unpredicted set is at most  $R_{\hat{t}} \cdot \frac{(2\gamma + 1)10\rho H_n}{9}$  (since we rejected all the bidders in some round between 1 and  $\hat{t}$ ), we obtain that Mechanism 1 outputs a set with at least a  $\frac{9}{\rho \cdot (2\gamma + 1)100H_n}$ -fraction of the optimal welfare. Suppose, instead, that Mechanism 1 outputs a subset of some unpredicted set (i.e., we reach line 13) in the final iteration of the while loop  $\hat{t}$ . By the same arguments in Lemma 3.3, we know that the revenue among active bidders in the highest revenue set is  $R_{\hat{t}}$  when we reach line 13. As above, the total welfare rejected from any unpredicted set by this point is at most  $R_{\hat{t}} \cdot \frac{(2\gamma + 1)10\rho H_n}{9}$ . But then, we have that the *revenue* among active bidders in the highest revenue set when we reach line 13 is a  $\frac{9}{\rho \cdot 10(2\gamma + 1)H_n}$ -fraction of the maximum social welfare from any feasible set of rejected unpredicted bidders. We can then apply similar reasoning to the proof of Lemma 3.4 to complete the proof. As noted in [21], the WFCA can also readily admit the use of a  $\rho$ -approximation for maximizing the revenue at a fixed set of clock prices. In this case, the approximation guarantee given the WFCA is  $\rho \cdot 2H_n$  [21, Theorem 5]. Moreover, the WFCA remains revenue monotone even when using an approximation algorithm. Combining these observations with our above reasoning, we have that the revenue from the set output by Mechanism 1 is a  $\frac{9}{\rho \cdot 10(2\gamma + 1)H_n}$ -fraction of the welfare rejected before running the WFCA and the welfare from the set output by Mechanism 1 is a  $\frac{1}{\rho \cdot 2H_n}$ . Putting this all together, we have that Mechanism 1 with any  $\gamma > 10/9$  and a  $\rho$ -approximate black-box achieves at least a  $\frac{9}{\rho \cdot 10(2\gamma + 1)H_n} + \frac{1}{\rho \cdot 2H_n}$ -fraction of the optimal welfare when the prediction is incorrect and Mechanism 1 outputs a subset of some unpredicted set. Combining our two cases gives that Mechanism 1 is  $O(\rho \cdot \frac{1}{\epsilon} \log n)$ -robust, as desired.  $\square$

### B Missing Proof from Section 3.1

*Proof.* [Proof of Theorem 3.2] Consider any deterministic clock auction  $\mathcal{M}$  that is augmented with a prediction  $\hat{\text{OPT}}$  and achieves  $(1 + \epsilon)$ -consistency and robustness better than  $\Omega(\frac{1}{\epsilon} \log n)$ . Also, consider a family of instances involving an arbitrarily large number of bidders  $n$  and a feasibility constraint  $\mathcal{F}$  defined by two disjoint maximal feasible sets  $F_1$  and  $F_2$ , i.e.,  $N = F_1 \cup F_2$ , and a set  $F \subseteq N$  is feasible if and only if  $F \subseteq F_1$  or  $F \subseteq F_2$ .  $F_1$  contains only a single bidder and  $|F_2| = n - 1$ . Also, let  $F_2$  be the set predicted to be optimal, i.e.,  $\hat{\text{OPT}} = F_2$ .

Using this feasibility constraint, we now define a family of instances  $\mathcal{I}$  where the value of the single bidder in  $F_1$  is 0.99 and the values of the bidders in  $F_2$  correspond to the following set of values:

$$V_2 = \left\{ 0.99, \frac{\epsilon}{2(H_{n-1} - 1)}, \frac{\epsilon}{3(H_{n-1} - 1)}, \dots, \frac{\epsilon}{(n-1)(H_{n-1} - 1)} \right\}.$$

We now consider what  $\mathcal{M}$  does facing an instance in  $\mathcal{I}$  that is chosen by adversarially matching the values in  $V_2$  to bidders in  $F_2$ , depending on the deterministic price trajectory that  $\mathcal{M}$  will follow. Specifically, at each time  $t$  during the execution of  $\mathcal{M}$ , let  $V_2(t) \subseteq V_2$  be the set of values from  $V_2$  assigned to bidders that have not dropped out yet. Then, the adversarial assignment of values to bidders in  $F_2$  ensure that if at any time  $t$  the price offered to an active bidder  $i \in F_2$  exceeds a value  $v \in V_2(t)$ , we can assume that the adversary assigns  $v$  to  $i$ , so  $i$  would drop out. As a result, if number of active bidders at any time  $t$  is  $x \geq 2$ , then the highest value offered to any one of them is at most  $\frac{\epsilon}{x(H_{n-1} - 1)}$ , so the revenue never exceeds  $x \cdot \frac{\epsilon}{x(H_{n-1} - 1)} = \frac{\epsilon}{H_{n-1} - 1}$ .

Using this fact, with the assumption that the auction achieves better than  $\Omega(\frac{1}{\epsilon} \log n)$ , we can infer that the auction cannot terminate with at least two active bidder in  $F_2$ . Specifically, if  $\mathcal{M}$  terminated with  $x \geq 2$  active bidders in  $F_2$  and, as we argued above, their final price was at most  $\frac{\epsilon}{x(H_{n-1} - 1)}$ , then  $\mathcal{M}$  would fail to satisfy the claimed robustness guarantee on an alternative instance (not in  $\mathcal{I}$ ). This alternative instance is identical in terms of the values of all the bidders that dropped out and the only difference is in the value of the  $x$  winning bidders, whose values are changed to  $\frac{\epsilon}{x(H_{n-1} - 1)}$ . Note that the outcome of  $\mathcal{M}$  facing this alternative instance would be the same, since the prices offered to the winning bidders would still be below their values. However, in this alternative instance the welfare of  $\mathcal{M}$  would be  $\frac{\epsilon}{H_{n-1} - 1}$  while the optimal social welfare is 0.99, leading to a robustness of  $\Omega(\frac{1}{\epsilon} \log n)$ .

If, on the other hand, the auction accepted at most 1 bidder from  $F_2$ , then the resulting welfare would be 0.99. Consider the total value rejected so far, we would have:  $\sum_{i=2}^{n-1} \frac{\epsilon}{i \cdot (H_{n-1} - 1)} = \epsilon$ , since  $\sum_{i=2}^{n-1} = H_{n-1} - 1$ . Therefore the consistency would be

$$\frac{0.99 + \epsilon}{0.99} > 1 + \epsilon,$$

violating the consistency guarantee.  $\square$

## C Missing Proof from Section 3.2

For completeness we include the full description of `ERRORTOLERANT` below:

**LEMMA C.1.** *Suppose that the predicted set has total welfare at least  $1/\bar{\eta}$  times the optimal social welfare (i.e., the prediction is within the error tolerance). Then for a fixed parameter  $\gamma > 10/9$  `ERRORTOLERANT` terminates in either Line 8 or Line 10 and obtains at least a  $\left(1 - \frac{10}{9\gamma}\right)$ -fraction of the social welfare of the predicted set.*

*Proof.* Fix an instance  $I$  where the optimal set has social welfare  $W^*$  and where the predicted set has welfare at least  $W^*/\bar{\eta}$ . Consider the minimum value  $k$  such that  $10^k \cdot \gamma H_n$  exceeds  $W^*$ . Let  $i'$  denote the value of  $i$  at the beginning of the final iteration of the while loop of `ERRORTOLERANT` when run on instance  $I$ . Since we only continue to line 9 after the preceding if statement when there are still active unpredicted bidders it must be that  $i' \leq k$  (as, otherwise, some unpredicted set would have total welfare greater than  $W^*$ ) and if  $i' = k$  then we terminate before line 9 (i.e., before we cause any additional predicted bidders to exit the auction in the iteration  $k$ ).

Now consider the total amount of welfare lost from the predicted set from rounds 1 to  $i' - 1$ . By Lemma 3.5 we have that this total welfare is at most  $\frac{R_{i'-1}}{\bar{\eta}} \cdot 10H_n/9 = \frac{10^{i'-1}}{\bar{\eta}} \cdot \frac{10H_n}{9} \leq \frac{10^{i^*-1}}{\bar{\eta}} \cdot \frac{10H_n}{9}$ . On the other hand, by our definition of  $k$ , we know that the total social welfare in the predicted set is at least  $\frac{10^{k-1}}{\bar{\eta}} \cdot \gamma H_n$ . But then, at the end of the  $i' - 1$  iteration there must be active predicted bidders and moreover, we retain at least a

$$\frac{\frac{10^{k-1}}{\bar{\eta}} \cdot \gamma H_n - \frac{10^{i^*-1}}{\bar{\eta}} \cdot \frac{10}{9} H_n}{\frac{10^{k-1}}{\bar{\eta}} \cdot \gamma H_n} = \frac{\gamma - \frac{10}{9}}{\gamma}$$

fraction of the social welfare of the predicted set.  $\square$

---

**MECHANISM 3: ERRORTOLERANT**


---

**Input:** A system  $\mathcal{F}$  of feasible sets of bidders, a predicted optimal set  $\hat{\mathcal{O}}\mathcal{P}\mathcal{T}$ , a parameter  $\gamma$

```

1  $F \leftarrow F \setminus \{F \cap \hat{\mathcal{O}}\mathcal{P}\mathcal{T}\}$  for all  $F \neq \hat{\mathcal{O}}\mathcal{P}\mathcal{T} \in \mathcal{F}$            // Make  $\hat{\mathcal{O}}\mathcal{P}\mathcal{T}$  disjoint from all unpredicted sets
2  $p_i \leftarrow v_{\min}$  for all  $i \in N$            // Raise price of each bidder to the initial small amount
3  $R_0 \leftarrow \text{rev}(\hat{\mathcal{O}}\mathcal{P}\mathcal{T}, \mathbf{p})$ 
4  $t \leftarrow 1$            // Initialize the counter
5 while TRUE do
6    $R_t \leftarrow 10^t R_{t-1}$ 
7   Run  $\text{UNIFORMPRICE}(N \setminus \hat{\mathcal{O}}\mathcal{P}\mathcal{T}, \mathbf{p})$  until  $\max_F v(F \setminus N) \in [R_t \gamma H_n, 2R_t \gamma H_n]^4$  // Reject ‘‘safe’’
      amount of welfare to determine minimum total value in optimal set
8   if  $N = \hat{\mathcal{O}}\mathcal{P}\mathcal{T}$  then return  $\hat{\mathcal{O}}\mathcal{P}\mathcal{T}, \mathbf{p}$ 
9   Run  $\text{UNIFORMPRICE}(N \setminus \hat{\mathcal{O}}\mathcal{P}\mathcal{T}, \mathbf{p})$  until  $\max_F \text{rev}(F \cap N, \mathbf{p}) \geq R_t$  // Ensure unpredicted sets
      ‘‘cover’’ their own lost welfare
10  if  $N = \hat{\mathcal{O}}\mathcal{P}\mathcal{T}$  then return  $\hat{\mathcal{O}}\mathcal{P}\mathcal{T}, \mathbf{p}$ 
11  Run  $\text{UNIFORMPRICE}(\hat{\mathcal{O}}\mathcal{P}\mathcal{T}, \mathbf{p})$  until  $\text{rev}(\hat{\mathcal{O}}\mathcal{P}\mathcal{T}, \mathbf{p}) \geq R_t / \bar{\eta}$  // Ensure predicted set ‘‘covers’’
      welfare lost from unpredicted sets
12  if  $\hat{\mathcal{O}}\mathcal{P}\mathcal{T} = \emptyset$  then
13    return  $\text{WFCA}(N, \mathbf{p})$ 
14   $t \leftarrow t + 1$ 

```

---

**LEMMA C.2.** *Consider some fixed constant parameter  $\gamma > 10/9$  **ERRORTOLERANT** outputs a subset of the predicted set. Then **ERRORTOLERANT** obtains social welfare within a  $\frac{9}{\bar{\eta} \cdot 100(2\gamma+1)H_n}$ -factor of the optimal social welfare.*

*Proof.* Consider the iteration  $\hat{t}$  in which the while loop of **ERRORTOLERANT** terminates. By definition, the mechanism terminates in either line 8 or line 10. By Line 7 and Lemma 3.1 we have that the total welfare rejected from any unpredicted set in any round  $\hat{t}$  is at most  $2W_{\hat{t}} + R_{\hat{t}}H_n = R_{\hat{t}} \cdot (2\gamma + 1)H_n$ . But then, summing over all rounds we have that the total welfare contained in any unpredicted set is at most

$$\sum_{\hat{t}=1}^{\hat{t}} R_{\hat{t}} \cdot (2\gamma + 1)H_n = (2\gamma + 1)H_n \cdot \sum_{\hat{t}=1}^{\hat{t}} 10^{\hat{t}} = (2\gamma + 1)H_n \cdot \frac{10^{\hat{t}+1} - 1}{9} \leq R_{\hat{t}} \cdot \frac{(2\gamma + 1)10H_n}{9}.$$

On the other hand, since we continued to iteration  $\hat{t}$  we have that the revenue of the bidders in the predicted set reached  $\frac{R_{\hat{t}-1}}{\bar{\eta}} = \frac{R_{\hat{t}}}{10\bar{\eta}}$ . Moreover, this means that the welfare of the active predicted bidders is at least  $\frac{R_{\hat{t}}}{10\bar{\eta}}$  and we serve all these bidders. Hence, **ERRORTOLERANT** obtains at least a  $\frac{R_{\hat{t}}/(10\bar{\eta})}{R_{\hat{t}} \cdot \frac{(2\gamma+1)10H_n}{9}} = \frac{9}{\bar{\eta} \cdot 100(2\gamma+1)H_n}$ -fraction of the optimal social welfare.  $\square$

**LEMMA C.3.** *Consider some parameter  $\gamma > 10/9$  **ERRORTOLERANT** reaches line 13. Then the revenue of active bidders in the highest revenue set before running the Water-filling Clock Auction is at least a  $\frac{9}{10(2\gamma+1)H_n}$ -approximation of the maximum social welfare obtainable by rejected unpredicted bidders.*

*Proof.* As before we consider the iteration  $\hat{t}$  in which the while loop of **ERRORTOLERANT** terminates. By assumption, our while loop terminates in line 13. Following the same reasoning as Lemma C.2, the total welfare rejected from any unpredicted set over all rounds from 1 to  $\hat{t}$  is at most

$$\sum_{i'=1}^{\hat{t}} R_{i'} \cdot (2\gamma + 1)H_n \leq R_{\hat{t}} \cdot \frac{(2\gamma + 1)10H_n}{9}.$$

On the other hand, because we arrive at line 13 we know that there exists some set unpredicted set which obtains revenue  $R_{\hat{t}}$  in round  $\hat{t}$ . Thus, the revenue among active bidders in the highest revenue set when we reach line

13 is a  $\frac{9}{10(2\gamma+1)H_n}$ -fraction of the maximum social welfare obtained from any feasible set of rejected unpredicted bidders.  $\square$

LEMMA C.4. *Consider some parameter  $\gamma > 10/9$  ERRORTOLERANT outputs a subset of unpredicted bidders. Then ERRORTOLERANT obtains social welfare within a  $\frac{9}{10(2\gamma+1)H_n} + \frac{1}{4\log n}$ -approximation to the optimal social welfare.*

*Proof.* In the case that ERRORTOLERANT outputs a subset of unpredicted bidders it must arrive at line 13. As such, we can decompose the social welfare from the optimal (unpredicted) set into two components, the portion of the welfare coming from bidders rejected before ERRORTOLERANT runs the Water-filling Clock Auction and the portion of the welfare coming from bidders who participate in the WFCA. From Lemma C.3 we know that the *revenue* reached by ERRORTOLERANT before running the WFCA is within a  $\frac{9}{10(2\gamma+1)H_n}$ -factor of the welfare rejected from the optimal set before running the WFCA. Moreover, from Lemma 2.1 we have that the revenue reached by ERRORTOLERANT after running the WFCA is weakly higher than the revenue reached before running the WFCA. As such, since the welfare obtained by serving a set of bidders is always weakly higher than the revenue collected from these bidders we have that the social welfare obtained by ERRORTOLERANT is within a  $\frac{9}{10(2\gamma+1)H_n}$ -factor of the welfare rejected from the optimal set before running the WFCA. Finally, we have from Lemma 2.1 that the social welfare achieved by running the WFCA only on bidders who are active when line 13 is reached is within a  $1/(4\log n)$ -factor of the optimal social welfare achievable from these bidders. Combining these two guarantees, we have that the social welfare obtained by ERRORTOLERANT is within a  $\frac{9}{10(2\gamma+1)H_n} + \frac{1}{4\log n}$ -factor of the optimal social welfare.  $\square$

*Proof.* [Proof of Theorem 3.3] First observe that  $\epsilon > 0$  ensures that  $\gamma > 10/9$ . In the case that the predicted set, indeed, has social welfare at least  $1/\bar{\eta}$  times the optimal social welfare then Lemma C.1 gives that some subset of the predicted set is served. Moreover, we obtain that the total social welfare of served bidders is within a factor

$$\left(1 - \frac{10}{9\gamma}\right) = \left(1 - \frac{10}{9\frac{10(1+\epsilon)}{9\epsilon}}\right) = \frac{1}{1+\epsilon}$$

of the total social welfare in the predicted set, thereby guaranteeing  $(1+\epsilon)\eta$  approximation with error tolerance  $\bar{\eta}$ . On the other hand, if the predicted set has welfare less than  $1/\bar{\eta}$  times the optimal (i.e.,  $\eta \leq \bar{\eta}$ ) then we can consider two cases depending on whether or not a subset of the predicted set is ultimately served. If some subset of the predicted set is served, Lemma C.2 guarantees that the total social welfare obtained by ERRORTOLERANT is within a factor  $\frac{9}{\bar{\eta} \cdot 100(2\gamma+1)H_n}$  of the optimal social welfare. Since  $H_n = \Theta(\log n)$  and  $\gamma$  is a constant for any choice of constant  $\epsilon$  we obtain  $\Theta(\log n)$ -robustness in this case when  $\bar{\eta} = \Theta(1)$ . Finally, if some subset of an unpredicted set is served, Lemma C.4 guarantees that ERRORTOLERANT obtains welfare within a factor  $\frac{9}{10(2\gamma+1)H_n} + \frac{1}{4\log n}$  of the optimal social welfare. Again, since  $H_n = \Theta(\log n)$  and  $\gamma$  is constant for any constant  $\epsilon$  we obtain  $\Theta(\log n)$ -robustness in this case, completing the proof.  $\square$

## D Missing Proofs from Section 4

### D.1 Proof of Lemma 4.1

*Proof.* Fix a round  $t$  in our auction and consider the total welfare lost by some unpredicted set  $S$  in Line 7 during  $t$ . First observe that if we move to Line 11 as a result of satisfying the second condition of the price increase for the unpredicted sets then it must be that every unpredicted set has rejected welfare less than  $\frac{\beta}{4} \cdot R_{t-1}^P$  (since we would have otherwise have stopped raising prices sooner, having met the lower revenue target of  $\frac{\beta}{4H_n} \cdot R_{t-1}^P$  before reaching  $\frac{\beta}{2} \cdot R_{t-1}^P$ ). On the other hand, suppose we move to Line 11 as a result of satisfying the first condition. Once some unpredicted set has rejected welfare  $\frac{\beta}{4} \cdot R_{t-1}^P$ , this condition will be satisfied if there exists any set with  $k$  remaining bidders of value at least  $\frac{\beta}{4H_n} \cdot \frac{R_{t-1}^P}{k}$ . As such, if we reject the  $k$ -th highest value bidder from any unpredicted set in round  $t$  after having rejected after having rejected  $\frac{\beta}{4} \cdot R_{t-1}^P$  from some unpredicted set it must be that her value was at most  $\frac{\beta}{4H_n} \cdot \frac{R_{t-1}^P}{k}$ . However, since each set has size at most  $n$  the total value rejected from

any unpredicted set  $S$  in round  $t$  after having rejected welfare at least  $\frac{\beta}{4} \cdot R_{t-1}^P$  from some unpredicted set is at most

$$\sum_{k=1}^n \frac{\beta}{4H_n} \cdot \frac{R_{t-1}^P}{k} = \frac{\beta}{4} \cdot R_{t-1}^P.$$

In total, we cannot reject total value from any set more than  $\frac{\beta}{2} \cdot R_{t-1}^P$ . Similarly, if neither condition is met (i.e., we reject all unpredicted bidders before reaching either (i) or (ii) in Line 7) we cannot reject total value from any set more than  $\frac{\beta}{4} \cdot R_{t-1}^P + \frac{\beta}{4} \cdot R_{t-1}^P = \frac{\beta}{2} \cdot R_{t-1}^P$ .  $\square$

## D.2 Proof of Corollary 4.1

*Proof.* Consider that in each round  $\tilde{R}_t^P$  is set to double of  $R_{t-1}^P$  in Line 10. And, by the conditions of Line 11 we have that the value of  $R_t^P$  at the end of iteration  $t$  in the while loop is at least twice  $R_{t-1}^P$ . Applying Lemma 4.1 to each round  $t < t'$  we then have that the total welfare rejected through the first  $t'$  iterations from any unpredicted set is at most

$$\sum_{t=1}^{t'} \frac{\beta}{2} R_{t-1}^P = \sum_{t=1}^{t'-1} \frac{\beta}{2} R_{t-1}^P + \frac{\beta}{2} R_{t'-1}^P \leq \beta R_{t'-1}^P,$$

as desired.  $\square$

## D.3 Proof of Lemma 4.4

*Proof.* Observe that when we terminate, by assumption that we output a subset of the predicted set, we did so because we did not satisfy either condition in Line 7. But in the previous iteration of the while loop we have that the revenue of the predicted set was  $(\alpha - 1)$  times the total welfare rejected from the predicted set. The social welfare remaining in the predicted set is then at least a  $\alpha$ -factor of the total welfare of the predicted set.  $\square$

## E Auxiliary Lemmas

LEMMA E.1. *For any choice of  $\alpha > 1$  and all  $n > 2$ , we have that*

$$\sum_{i=2}^n \frac{\Gamma(i-1)\Gamma\left(n+1+\frac{1}{\alpha-1}\right)}{\Gamma\left(\frac{1}{\alpha-1}+i\right)\Gamma(n)} = -\alpha n + \frac{\alpha\Gamma\left(\frac{n\alpha+\alpha-n}{\alpha-1}\right)}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)\Gamma(n)} + n - 1.$$

*Proof.* We demonstrate first a simpler summation in the form of Equation (E.1) below which we argue holds whenever  $\alpha > 1$

$$(E.1) \quad \sum_{i=2}^n \frac{\Gamma(i-1)}{\Gamma\left(i+\frac{1}{\alpha-1}\right)} = \frac{(-\alpha n + n - 1)\Gamma(n)}{\Gamma\left(\frac{n\alpha+\alpha-n}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)}.$$

Our proof that Equation E.1 holds proceeds via induction on  $n$ . Consider the base case of  $n = 2$ . The left-hand side of Eq. (E.1) is then

$$\frac{\Gamma(1)}{\Gamma\left(2+\frac{1}{\alpha-1}\right)} = \frac{1}{\Gamma\left(2+\frac{1}{\alpha-1}\right)}.$$

On the other hand we have that the right-hand side is

$$\begin{aligned} \frac{(-2\alpha+1)\Gamma(2)}{\Gamma\left(\frac{2\alpha+\alpha-2}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)} &= \frac{(-2\alpha+1)}{\Gamma\left(\frac{3\alpha-2}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)} \\ &= \frac{1-2\alpha}{\Gamma\left(3+\frac{1}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)} \\ &= \frac{1-2\alpha}{\left(2+\frac{1}{\alpha-1}\right)\Gamma\left(2+\frac{1}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)} \end{aligned}$$

$$\begin{aligned}
&= \frac{1-\alpha}{\Gamma\left(2 + \frac{1}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(2 + \frac{1}{\alpha-1}\right)} \\
&= \frac{1}{\Gamma\left(2 + \frac{1}{\alpha-1}\right)},
\end{aligned}$$

where the third equality is due to the fact that  $\Gamma(z) = (z-1)\Gamma(z-1)$  and the fourth equality is due to the fact that  $\alpha > 1$  means that  $\alpha \neq 1/2$  (and hence  $(1-2\alpha)/(2+1/(\alpha-1)) = 1-2\alpha$ ). As such, we have that the base case of our induction holds.

We now proceed to the inductive step. Assume that

$$\sum_{i=2}^n \frac{\Gamma(i-1)}{\Gamma\left(i + \frac{1}{\alpha-1}\right)} = \frac{(-\alpha n + n-1)\Gamma(n)}{\Gamma\left(\frac{n\alpha+\alpha-n}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)}$$

for all  $n < n' + 1$ . We thus want to show the equation remains true for  $n = n' + 1$ , i.e., we want to show

$$(E.2) \quad \sum_{i=2}^{n'+1} \frac{\Gamma(i-1)}{\Gamma\left(i + \frac{1}{\alpha-1}\right)} = \frac{(-\alpha(n'+1) + (n'+1)-1)\Gamma(n'+1)}{\Gamma\left(\frac{(n'+1)\alpha+\alpha-(n'+1)}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)}.$$

We rewrite the left hand side of Equation (E.2) as

$$\begin{aligned}
\frac{\Gamma(n')}{\Gamma\left(n'+1 + \frac{1}{\alpha-1}\right)} + \sum_{i=2}^{n'} \frac{\Gamma(i-1)}{\Gamma\left(i + \frac{1}{\alpha-1}\right)} &= \frac{\Gamma(n')}{\Gamma\left(n'+1 + \frac{1}{\alpha-1}\right)} + \frac{(-\alpha n' + n'-1)\Gamma(n')}{\Gamma\left(\frac{n'\alpha+\alpha-n'}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)} \\
&= \frac{\Gamma(n')}{\Gamma\left(\frac{n'(\alpha-1)+\alpha}{\alpha-1}\right)} + \frac{(-\alpha n' + n'-1)\Gamma(n')}{\Gamma\left(\frac{n'(\alpha-1)+\alpha}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)} \\
&= \frac{(-\alpha n' + n')\Gamma(n')}{\Gamma\left(\frac{n'(\alpha-1)+\alpha}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)} \\
&= \frac{n'(1-\alpha)\Gamma(n')}{\Gamma\left(\frac{n'(\alpha-1)+\alpha}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)} \\
&= \frac{(1-\alpha)\Gamma(n'+1)}{\Gamma\left(\frac{n'(\alpha-1)+\alpha}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)}
\end{aligned}$$

where the first equality applies the inductive hypothesis and the last equality applies the identity  $\Gamma(z) = (z-1)\Gamma(z-1)$ .

We can similarly rewrite the right hand side of Equation (E.2) as

$$\begin{aligned}
\frac{(-\alpha(n'+1) + (n'+1)-1)\Gamma(n'+1)}{\Gamma\left(\frac{(n'+1)\alpha+\alpha-(n'+1)}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)} &= \frac{((1-\alpha)(n'+1)-1)\Gamma(n'+1)}{\Gamma\left(\frac{(\alpha-1)(n'+1)+\alpha}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)} \\
&= \frac{((1-\alpha)(n'+1)-1)\Gamma(n'+1)}{\Gamma\left(1 + \frac{n'(\alpha-1)+\alpha}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)} \\
&= \frac{((1-\alpha)(n'+1)-1)\Gamma(n'+1)}{\frac{n'(\alpha-1)+\alpha}{\alpha-1} \cdot \Gamma\left(\frac{n'(\alpha-1)+\alpha}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)} \\
&= \frac{((1-\alpha)n' - \alpha)\Gamma(n'+1)}{\frac{n'(\alpha-1)+\alpha}{\alpha-1} \cdot \Gamma\left(\frac{n'(\alpha-1)+\alpha}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)} \\
&= \frac{(1-\alpha)\Gamma(n'+1)}{\Gamma\left(\frac{n'(\alpha-1)+\alpha}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)},
\end{aligned}$$

where the third equality applies the identity  $\Gamma(z) = (z - 1)\Gamma(z - 1)$ , completing the inductive step as desired.

Our main summation then follows, essentially, as a direct consequence of Equation (E.1). Factoring out the common term  $\frac{\Gamma(n+1+\frac{1}{\alpha-1})}{\Gamma(n)}$  and applying Equation (E.1) we obtain

$$\begin{aligned}
\sum_{i=2}^n \frac{\Gamma(i-1)\Gamma\left(n+1+\frac{1}{\alpha-1}\right)}{\Gamma\left(\frac{1}{\alpha-1}+i\right)\Gamma(n)} &= \frac{\Gamma\left(n+1+\frac{1}{\alpha-1}\right)}{\Gamma(n)} \sum_{i=2}^n \frac{\Gamma(i-1)}{\Gamma\left(i+\frac{1}{\alpha-1}\right)} \\
&= \frac{\Gamma\left(n+1+\frac{1}{\alpha-1}\right)}{\Gamma(n)} \cdot \left( \frac{(-\alpha n + n - 1)\Gamma(n)}{\Gamma\left(\frac{n\alpha+\alpha-n}{\alpha-1}\right)} + \frac{\alpha}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)} \right) \\
&= -\alpha n + \frac{\alpha\Gamma\left(\frac{n\alpha+\alpha-n}{\alpha-1}\right)}{\Gamma\left(\frac{2\alpha-1}{\alpha-1}\right)\Gamma(n)} + n - 1,
\end{aligned}$$

as desired.  $\square$