

Randomized Learning-Augmented Auctions with Revenue Guarantees

Ioannis Caragiannis and Georgios Kalantzis

Department of Computer Science, Aarhus University

iannis@cs.au.dk, 202202869@post.au.dk

Abstract

We consider the fundamental problem of designing a truthful single-item auction with the challenging objective of extracting a large fraction of the highest agent valuation as revenue. Following a recent trend in algorithm design, we assume that the agent valuations belong to a known interval, and a (possibly erroneous) prediction for the highest valuation is available. Then, auction design aims for high consistency and robustness, meaning that, for appropriate pairs of values γ and ρ , the extracted revenue should be at least a γ - or ρ -fraction of the highest valuation when the prediction is correct for the input instance or not. We characterize all pairs of parameters γ and ρ so that a randomized γ -consistent and ρ -robust auction exists. Furthermore, for the setting in which robustness can be a function of the prediction error, we give sufficient and necessary conditions for the existence of robust auctions and present randomized auctions that extract a revenue that is only a polylogarithmic (in terms of the prediction error) factor away from the highest agent valuation.

1 Introduction

We revisit one of the most central topics in microeconomic theory, namely the design of *single-item auctions*, under the lens of the recent trend of *learning-augmented algorithms* [Mitzenmacher and Vassilvitskii, 2020]. The basic auction setting consists of a seller who has an item for sale and potential buyers (or agents), each having private valuations for the item. An auction is then used to sell the item. The agents submit their bids to the seller, who then decides who should get the item and at which price.

Agents are strategic; they may consider bidding non-truthfully (i.e., submit a bid that is different than their valuation for the item) if this can increase their utility. Since non-truthful bids may make reasoning about the agents' behaviour and the auction's outcome difficult, a large part of auction design theory has focused on truthful auctions.

For single-item auctions (and, single-parameter environments, more generally), the auction designer has essentially to define two components: an allocation rule and a payment

rule. Both are functions taking the agent bids as input. The allocation rule returns the agent who should get the item (if any), and the payment rule returns the payment the agents should give to the seller.

The design of truthful auctions requires careful selection of both the allocation and the payment rules. For example, the celebrated second-price auction [Vickrey, 1961] allocates the item to the agent with the highest bid, who pays the second-highest bid to the seller. The second-price auction is truthful. Even though proving this formally is rather easy and usually serves as warming-up material in introductory textbooks to mechanism design (e.g., see Roughgarden [2016]), doing so for more sophisticated auctions can be tricky. Fortunately, the beautiful theory of Myerson [1981] provides us with a powerful toolset for single-item auction design (and, more generally, for mechanism design in single-parameter environments). The infamous Myerson's Lemma restricts the allocation rules that can be used in truthful auctions to those having a monotonicity property. Furthermore, the allocation rule identifies the payment function in an almost unique way.

The second-price auction has additional nice economic properties, e.g., it maximizes social welfare. Unfortunately, it may result in poor revenue. To get revenue guarantees, knowledge of statistical information (e.g., in the form of the probability distributions of the agent valuations) can be very useful to the auction designer. Then, a usual trick is to design auctions that use reserve prices, which can secure a high payment by the highest bidder (or the auction winner, in general), even when the competition is low.

In this paper, we assume that no statistical information for the potential buyers is available. Instead, we assume that their valuations can have any value in a known interval $[1, H]$. An ambitious goal would be to design truthful auctions that always extract a revenue that is close to (i.e., a multiplicative approximation of) the highest valuation among the agents. For the second-price auction, this is an unrealistic goal; just imagine two bidders with valuations 1 and H . Even worse, this is unrealistic for any deterministic auction, as the next argument shows. To get non-zero revenue from two bidders, both having a valuation of 1, one should get the item at a price of at most 1. Then, the monotonicity of truthful allocation rules implies that the same agent should get the item at the same price if her valuation was H instead (and the other bidder still had a valuation of 1).

To solve this issue keeping the ambitious goal of extracting a revenue that approximates the highest valuation, we resort to randomization. Randomized auctions can use allocation rules, which, for each bid profile, define the probability that each agent will get the item. They are still bound to Myerson’s monotonicity, which is not restrictive anymore but rather reveals a much richer design space. For example, randomized auctions can exploit reserve prices, which would yield no revenue if used by deterministic auctions.

Furthermore, following the practice in the literature of learning-augmented algorithms and the model of Xu and Lu [2022] in particular, we assume that a *prediction* \hat{u} of the highest valuation is available. This means that both the range $[1, H]$ of valuations and the value \hat{u} can be hardwired in the definition of the auction, which should now satisfy two different revenue guarantees, known as *consistency* and *robustness* in the literature. First, for valuation profiles for which the prediction is correct (and the highest valuation is indeed equal to \hat{u}), the auction should extract a γ -approximation of the highest valuation as revenue. For the remaining inputs, in which the prediction is incorrect, a typically worse ρ -approximation is sought. Such an auction will be said to have consistency and robustness of γ and ρ , respectively.

Alternatively, we can allow the revenue guarantee to depend on the prediction error η , denoting how far (multiplicatively) the highest valuation in a given valuation profile is from the prediction \hat{u} . Then, robustness is expressed as a function ρ , with $\rho(\eta)$ indicating the required lower bound on the revenue-over-highest-valuation ratio extracted by the auction in all valuation profiles with prediction error η . Hence, in the terminology of the setting of the previous paragraph, the value $\rho(1)$ is essentially a consistency guarantee.

1.1 Our Contribution

We study a class of randomized auctions with very appealing characteristics. We refer to them as *intuitive auctions*. An intuitive auction is anonymous (i.e., the outcome does not depend on the identifiers of the agents in any way) and, furthermore, has the property that only the highest bidder(s) get the item with positive probability. These properties make intuitive auctions particularly simple.

For the first setting, we give an optimal trade-off between the consistency and robustness of truthful (more precisely, of dominant-strategy incentive-compatible) intuitive auctions. As a corollary, we present auctions that have constant consistency and robustness $\Omega(\ln^{-1} H)$. For the second setting, we present a sufficient and necessary condition for intuitive auctions with a given function ρ of prediction error as robustness guarantee. As corollaries, we obtain robustness guarantees so that $1/\rho(\eta)$ is polylogarithmic in η (even for unbounded valuations) or (sub)logarithmic, with a small dependency on H as well. In our proofs, we use extensively a convenient new variation of Myerson’s Lemma.

1.2 Related Work

Learning-augmented algorithms have emerged as a very hot topic nowadays, with numerous related contributions in recent years. Many classical problems have been reconsidered, and new algorithms, enhanced with (possibly erro-

neous) machine-learned predictions about their input, are designed and analyzed with respect to the consistency and robustness guarantees they can achieve. Representative problem domains include data structures, online and approximation algorithms for combinatorial optimization, streaming and sublinear algorithms, and many more. See the early survey by Mitzenmacher and Vassilvitskii [2020] as well as the online repository algorithms-with-predictions.github.io.

In algorithmic game theory and computational social choice, the concept of prediction has been considered for problems related to mechanism design [Agrawal *et al.*, 2022; Balcan *et al.*, 2023; Balkanski *et al.*, 2023a; Balkanski *et al.*, 2023b; Istrate and Bonchis, 2022; Xu and Lu, 2022], price of anarchy of cost sharing [Gkatzelis *et al.*, 2022], and distortion of voting [Berger *et al.*, 2023]. The paper by Xu and Lu [2022] is closest to ours. Among other problems, the authors study single-item auctions with the revenue-over-highest valuation objective and present consistency and robustness bounds for deterministic auctions. In particular, they present a truthful auction which is proved to have consistency γ and per-instance robustness that is a function of γ , the prediction error, and the upper bound on the valuation of agents. We remark that the use of randomization in the current paper allows us to obtain considerably better consistency vs. robustness tradeoffs compared to the results of Xu and Lu [2022].

Our assumptions of agents with worst-case valuations are similar in spirit to early work on competitive auctions, initiated with the work of Goldberg *et al.* [2006]. Even though revenue has been the main concern in that line of research, weaker benchmarks than the highest valuation among all agents have been mainly considered. See the related discussion in the very recent paper by Lu *et al.* [2023], who study competitive auctions with predictions. The more distant but also extremely important field of Bayesian mechanism design, which uses extensively statistical information about the agent valuations, is surveyed by Hartline [2013].

2 Preliminaries

We begin with some background on auctions and mechanism design. The interested reader may find more information in standard textbooks, e.g., see [Roughgarden, 2016, Chapter 3]. In single-item auctions, a set of n agents (or *bidders*) compete for an item. Each agent has a private *valuation* v_i for the item. All valuations belong to the interval $[1, H]$ for $H > 1$. An *auction mechanism* (or, simply, auction) receives *bids* from the agents as input (to be thought of as reportings of their valuations) and decides the agent who will get the item (or that no agent should get the item) and the payment that will be received from each agent. Auctions are, in general, randomized. Formally, the auction consists of an *allocation rule* $\mathbf{x} = (x_1, x_2, \dots, x_n)$ and a *payment rule* $\mathbf{p} = (p_1, p_2, \dots, p_n)$. Both receive the *bid vector* $\mathbf{b} = (b_1, b_2, \dots, b_n)$ as input, consisting of one bid per agent. The allocation function $x_i(\mathbf{b})$ denotes the probability that agent i gets the item while the payment $p_i(\mathbf{b})$ denotes the payment by agent i . Thus, $x_i(\mathbf{b}) \in [0, 1]$ and $p_i(\mathbf{b}) \geq 0$. An allocation rule \mathbf{x} is *feasible* if $\sum_{i \in [n]} x_i(\mathbf{b}) \leq 1$ for every bid vector \mathbf{b} .

Agents are (expected) utility maximizers. Agent i ’s util-

ity, from the outcome of an auction that uses the allocation rule \mathbf{x} and the payment rule \mathbf{p} when the agents submit the bid vector \mathbf{b} , is defined as $u_i(\mathbf{b}) = v_i \cdot x_i(\mathbf{b}) - p_i(\mathbf{b})$. Ideally, we would like to use auctions that motivate agents to report their private valuations *truthfully* as bids. Truthfulness requires that each agent i maximizes her utility by reporting her valuation as bid. Formally, this requirement can be written as $u_i(v_i, \mathbf{b}_{-i}) \geq u_i(z, \mathbf{b}_{-i})$, for every bid vector \mathbf{b}_{-i} submitted by the agents different than i and for every possible bid $z \in [1, H]$ agent i may consider to submit. In addition, the property of *individual rationality* requires that every truthful agent i has non-negative utility $u_i(v_i, \mathbf{b}_{-i})$, for any possible bids \mathbf{b}_{-i} by other agents. An auction that is both truthful and individually rational is called *dominant-strategy incentive-compatible* (DSIC).

One of the most fundamental results in auction theory is the characterization of DSIC auctions by Myerson [1981], which connects the property of truthfulness to the monotonicity of the allocation rule. An allocation rule \mathbf{x} is *monotone* if, for every agent i and any possible bids \mathbf{b}_{-i} by the other agents, the function $x_i(z, \mathbf{b}_{-i})$ is non-decreasing in terms of z . We say that an allocation rule \mathbf{x} is *implementable* if there exist a payment rule \mathbf{p} so that the auction consisting of these two rules is DSIC.

Lemma 1 (Myerson’s Lemma). *The allocation rule \mathbf{x} is implementable if and only if it is monotone. If \mathbf{x} is monotone, it is implementable through the payment rule \mathbf{p} , which, for every agent i and any bids \mathbf{b}_{-i} by the other agents, it holds $x_i(1, \mathbf{b}_{-i}) \geq p_i(1, \mathbf{b}_{-i})$ and*

$$\begin{aligned} & p_i(t, \mathbf{b}_{-i}) - p_i(s, \mathbf{b}_{-i}) \\ &= t \cdot x_i(t, \mathbf{b}_{-i}) - s \cdot x_i(s, \mathbf{b}_{-i}) - \int_s^t x_i(z, \mathbf{b}_{-i}) \, dz \end{aligned}$$

for every $s, t \in [1, H]$ with $s \leq t$.

We now give an alternative to Myerson’s Lemma, which will be more convenient for our proofs. Notice that, while Lemma 1 identifies the payment rule necessary and sufficient to implement a monotone allocation rule, Lemma 2 does the opposite: Given a monotone payment rule, it gives us the allocation curve that is implemented with this payment rule.

Lemma 2. *A monotone allocation rule \mathbf{x} is implementable through a payment rule \mathbf{p} if and only if, for every agent i and any bids \mathbf{b}_{-i} by the other agents, it holds $x_i(1, \mathbf{b}_{-i}) \geq p_i(1, \mathbf{b}_{-i})$ and*

$$\begin{aligned} x_i(t, \mathbf{b}_{-i}) &= \frac{p_i(t, \mathbf{b}_{-i})}{t} + \int_s^t \frac{p_i(z, \mathbf{b}_{-i})}{z^2} \, dz \\ &+ x_i(s, \mathbf{b}_{-i}) - \frac{p_i(s, \mathbf{b}_{-i})}{s} \end{aligned}$$

for every $s, t \in [1, H]$ with $s \leq t$.

Before proving Lemma 2, we remark that the formulas connecting the allocation and payment functions in Lemmas 1 and 2 characterize truthfulness while the inequality $x_i(1, \mathbf{b}_{-i}) \geq p_i(1, \mathbf{b}_{-i})$ yields individual rationality. Notice that neither the particular version of Myerson’s Lemma that we use here nor our Lemma 2 make any differentiability assumptions for the allocation or payment functions.

Proof. It suffices to show that the two different formulas in the statements of Lemmas 1 and 2 are equivalent.

For valuations $s, t \in [1, H]$ with $s \leq t$, let

$$\sigma_s(t) = \frac{\int_s^t x_i(z, \mathbf{b}_{-i}) \, dz}{t}.$$

By applying Myerson’s Lemma, we get

$$\begin{aligned} & p_i(t, \mathbf{b}_{-i}) + s \cdot x_i(s, \mathbf{b}_{-i}) - p(s, \mathbf{b}_{-i}) \\ &= t \cdot x_i(t, \mathbf{b}_{-i}) - \int_s^t x_i(z, \mathbf{b}_{-i}) \, dz = t^2 \cdot \sigma'_s(t) \end{aligned}$$

and, equivalently,

$$\sigma'_s(t) = \frac{p_i(t, \mathbf{b}_{-i})}{t^2} + \frac{s \cdot x_i(s, \mathbf{b}_{-i}) - p_i(s, \mathbf{b}_{-i})}{t^2}.$$

Thus,

$$\begin{aligned} \sigma_s(t) &= \sigma_s(s) + \int_s^t \frac{p_i(z, \mathbf{b}_{-i})}{z^2} \, dz \\ &+ \left(\frac{1}{s} - \frac{1}{t} \right) \cdot (s \cdot x_i(s, \mathbf{b}_{-i}) - p_i(s, \mathbf{b}_{-i})), \end{aligned}$$

implying, by the definition of $\sigma_s(t)$, that

$$\begin{aligned} \int_s^t x_i(z, \mathbf{b}_{-i}) \, dz &= t \int_s^t \frac{p_i(z, \mathbf{b}_{-i})}{z^2} \, dz \\ &+ \left(\frac{t}{s} - 1 \right) \cdot (s \cdot x_i(s, \mathbf{b}_{-i}) - p_i(s, \mathbf{b}_{-i})). \end{aligned}$$

The lemma now follows after differentiating both sides with respect to t . \square

We consider *anonymous* auctions, in which the allocation function $x_i(z, \mathbf{b}_{-i})$ and the payment function $p_i(z, \mathbf{b}_{-i})$ do not depend on i . Furthermore, we consider auctions which give the item with positive probability to the highest bidder(s) only, and this probability depends on the highest and second highest bid. Due to anonymity, ties are resolved uniformly at random among the tied agents. We use the term *intuitive* to refer to such auctions. We also simplify notation for bid vectors, and allocation and payment functions as follows. We describe a bid vector \mathbf{b} as the triplet (t, b, ν) , where t is the bid of a particular agent i , b denotes the highest bid among the other agents, and ν denotes the number of agents different than i with a bid of b . Thus, anonymity allows us to use $x(t, b, \nu)$ and $p(t, b, \nu)$ to refer to allocations and payments in general. Of course, for auctions with a single agent, which will be important in our study, the bid of the agent is the only parameter required to describe the bid vector; we use $x(t)$ and $p(t)$ to refer to allocation and payments in this case.

We use the term *revenue* to refer to the total payment received by all agents and aim to design auctions that extract a high revenue. Ideally, an auction applied on agents with a highest valuation of t should extract a revenue of as close to t as possible. We consider the *revenue-over-highest-valuation ratio* as our main objective. We assume that in addition to parameter H denoting the upper bound on valuations, we are

given a prediction \hat{u} for the highest valuation. Then, for parameters γ and ρ , we aim to design auctions with revenue-over-highest valuation guarantees of γ and ρ when the prediction is correct and incorrect, respectively. Then, such auctions will be said to have *consistency* of γ and *robustness* of ρ . Alternatively, we can allow the robustness requirement to be a function of the prediction error $\max\{t/\hat{u}, \hat{u}/t\}$ indicating how far the highest valuation t is from the prediction \hat{u} . The robustness requirement can then be described by a non-increasing function $\rho : [1, H] \rightarrow [0, 1]$, asking for a revenue that is at least $\rho (\max\{t/\hat{u}, \hat{u}/t\}) \cdot t$.

3 A Consistency vs. Robustness Trade-off

We devote this section to proving the following statement, which gives sufficient and necessary conditions for the existence of consistent and robust auctions.

Theorem 3. *Let $0 \leq \rho \leq \gamma \leq 1$. There exists a γ -consistent and ρ -robust DSIC intuitive auction for bidders with valuations from $[1, H]$ and a prediction for the highest bid $\hat{u} \in [1, H]$ if and only if $\gamma + \rho \cdot \ln \max\left\{\hat{u}, \frac{H \cdot \rho}{\gamma}\right\} \leq 1$.*

The proof of Theorem 3 is given in Sections 3.1 and 3.2 below. As a corollary of Theorem 3, we obtain that auctions with constant consistency and robustness $\Omega(\ln^{-1} H)$ do exist. E.g., notice that the parameters $\gamma = 1/2$ and $\rho = \frac{1}{2(1+\ln H)}$ satisfy the condition of Theorem 3. Even though Theorem 3 just claims the existence of the desired auction, its proof is constructive. Once we have parameters γ and ρ satisfying the condition in the theorem, the definition of a corresponding auction is given explicitly in Section 3.1 below.

3.1 Proving the ‘‘if’’ Part of Theorem 3

We prove the ‘‘if’’ part of Theorem 3 by constructing a DSIC intuitive auction which uses the allocation function \bar{x} and the payment function \bar{p} defined below. Consider an agent i with valuation t . Let b be the highest bid among the remaining agents, and let ν denote the number of bidders different than i with valuation b .

The allocation fraction $\bar{x}(t, b, \nu)$ is defined as

$$\bar{x}(t, b, \nu) = \begin{cases} 0, & t \in [1, b) \\ \frac{\rho}{\nu+1}, & t = b \\ \rho + \rho \cdot \ln \frac{t}{b}, & t \in (b, H] \end{cases}$$

if $1 \leq \hat{u} < b \leq H$, as

$$\bar{x}(t, b, \nu) = \begin{cases} 0, & t \in [1, \hat{u}) \\ \frac{\gamma}{\nu+1}, & t = \hat{u} \\ \gamma, & t \in (\hat{u}, \min\left\{\frac{\gamma \cdot \hat{u}}{\rho}, H\right\}] \\ \gamma + \rho \cdot \ln \frac{t \cdot \rho}{b \cdot \gamma}, & t \in \left(\min\left\{\frac{\gamma \cdot \hat{u}}{\rho}, H\right\}, H\right] \end{cases}$$

if $1 \leq \hat{u} = b \leq H$, and as

$$\bar{x}(t, b, \nu) = \begin{cases} 0, & t \in [1, b) \\ \frac{\rho}{\nu+1}, & t = b \\ \rho + \rho \cdot \ln \frac{t}{b}, & t \in (b, \hat{u}) \\ \gamma + \rho \cdot \ln \frac{\hat{u}}{b}, & t \in \left[\hat{u}, \min\left\{\frac{\gamma \cdot \hat{u}}{\rho}, H\right\}\right] \\ \gamma + \rho \cdot \ln \frac{t \cdot \rho}{b \cdot \gamma}, & t \in \left(\min\left\{\frac{\gamma \cdot \hat{u}}{\rho}, H\right\}, H\right] \end{cases}$$




Figure 1: The most general form of the allocation function \bar{x} for an agent in terms of her valuation t , assuming $1 < b < \hat{u} < \frac{\gamma \cdot \hat{u}}{\rho} < H$ and $\nu = 1$. Notice that $\bar{x}(t, b, \nu)$ consists of five parts: the left-most part in which the item is not allocated to the agent, the point corresponding to a tie for the highest bid, and three more parts in which the allocation function has logarithmic, constant, and again logarithmic form. The remaining cases for the relative values of b , \hat{u} , $\frac{\gamma \cdot \hat{u}}{\rho}$ and H do not include some of the three rightmost parts. The black and white dots are used at points in which the allocation function ‘‘jumps’’; the black dot represents the allocation value at these points.

if $1 \leq b < \hat{u} \leq H$. See Figure 1.

Notice that, in all cases, \bar{x} is non-negative and non-decreasing in t . Whenever agent i is tied as highest bidder with a value of $t = b$ (together with ν more agents), the total fraction allocated is γ when $b = \hat{u}$ and ρ otherwise, i.e., at most 1. When agent i is the unique highest bidder, she is the only agent who gets a positive fraction of the item, which is maximized to either $\rho + \rho \cdot \ln \frac{H}{b}$, or $\gamma + \rho \cdot \ln \frac{\hat{u}}{b}$, or $\gamma + \rho \cdot \ln \frac{H \cdot \rho}{b \cdot \gamma}$. We have

$$\begin{aligned} \rho + \rho \cdot \ln \frac{H}{b} &\leq \rho + \rho \cdot \ln H \\ &\leq \rho + \rho \cdot \ln H + \rho \cdot \left(\frac{\gamma}{\rho} - 1 - \ln \frac{\gamma}{\rho} \right) \\ &\leq \gamma + \rho \cdot \ln \max\left\{\hat{u}, \frac{H \cdot \rho}{\gamma}\right\} \leq 1. \end{aligned}$$

The first two inequalities follow since $b \geq 1$ and using the inequality $z - 1 - \ln z \geq 0$ for $z > 0$. Furthermore,

$$\begin{aligned} &\max\left\{\gamma + \rho \cdot \ln \frac{\hat{u}}{b}, \gamma + \rho \cdot \ln \frac{H \cdot \rho}{b \cdot \gamma}\right\} \\ &\leq \gamma + \rho \cdot \ln \max\left\{\hat{u}, \frac{H \cdot \rho}{\gamma}\right\} \leq 1, \end{aligned}$$

by the assumption of Theorem 3. Thus, the allocation function \bar{x} is feasible and monotone.

The corresponding payment function \bar{p} is defined as

$$\bar{p}(t, b, \nu) = \begin{cases} 0, & t \in [1, b) \\ \frac{\rho \cdot b}{\nu+1}, & t = b \\ \rho \cdot t, & t \in (b, H] \end{cases}$$

if $1 \leq \hat{u} < b \leq H$, as

$$\bar{p}(t, b, \nu) = \begin{cases} 0, & t \in [1, \hat{u}) \\ \frac{\gamma \cdot \hat{u}}{\nu+1}, & t = \hat{u} \\ \gamma \cdot \hat{u}, & t \in \left(\hat{u}, \min \left\{ \frac{\gamma \cdot \hat{u}}{\rho}, H \right\} \right] \\ \rho \cdot t, & t \in \left(\min \left\{ \frac{\gamma \cdot \hat{u}}{\rho}, H \right\}, H \right] \end{cases}$$

if $1 \leq \hat{u} = b \leq H$, and as

$$\bar{p}(t, b, \nu) = \begin{cases} 0, & t \in [1, b) \\ \frac{\rho \cdot b}{\nu+1}, & t = b \\ \rho \cdot b, & t \in (b, \hat{u}) \\ \gamma \cdot \hat{u}, & t \in \left[\hat{u}, \min \left\{ \frac{\gamma \cdot \hat{u}}{\rho}, H \right\} \right] \\ \rho \cdot t, & t \in \left(\min \left\{ \frac{\gamma \cdot \hat{u}}{\rho}, H \right\}, H \right] \end{cases}$$

if $1 \leq b < \hat{u} \leq H$.

It is straightforward to verify that, indeed, the payment function \bar{p} implements the allocation function \bar{x} , i.e., that \bar{x} and \bar{p} are consistent with Lemma 2. Furthermore, we can verify that the auction's revenue yields the required consistency and robustness. When the prediction is correct, and the highest bid is equal to \hat{u} , the revenue is exactly $\gamma \cdot \hat{u}$, which implies a consistency of γ . When the highest bid t is different than \hat{u} , the revenue is either equal to $\rho \cdot t$ or equal to $\gamma \cdot \hat{u}$ for $t \in \left(\hat{u}, \min \left\{ \frac{\gamma \cdot \hat{u}}{\rho}, H \right\} \right]$, i.e., for $t \leq \frac{\gamma \cdot \hat{u}}{\rho}$. Notice that $\gamma \cdot \hat{u} \geq \rho \cdot t$ in this case, implying the bound of ρ for robustness. The completes the proof of the “if” part of Theorem 3.

Before we proceed to the “only if” part of the proof, let us attempt a “reverse engineering” of the proof above. Our main goal has been to decide the allocation rule \bar{x} . For intuitive auctions, this means that we need to define the monotone allocation curve $\bar{x}(t, b, \nu)$ for values $t \geq b$ (by definition, $\bar{x}(t, b, \nu) = 0$ for $t \in [1, b)$). The (unique) payment rule and, consequently, the revenue for every highest valuation t is then given by Myerson's Lemma. Pictorially (e.g., see Figure 1), this revenue is equal to the area between the vertical lines at points 0 and t that is above the allocation curve and below the horizontal line at point $\bar{x}(t, b, \nu)$. Thus, to get the guarantees of γ and ρ for consistency and robustness, the constraints are that this area is $\gamma \cdot t$ for $t = \hat{u}$ if $\hat{u} \geq b$ (the case $b > \hat{u}$ suggests that the prediction \hat{u} is incorrect and we can ignore it) and at least $\rho \cdot t$ for any other value of $t > b$. Our construction of the allocation rule \bar{x} essentially optimizes the revenue under these two constraints.

As a comparison to the work of Xu and Lu [2022], we remark that their auction is deterministic. Hence, as Myerson's Lemma and monotonicity suggest, they can only follow a “critical bid” definition. In particular, they have to use a threshold (possibly depending on the non-winning bids), so that the winning bidder gets the item (with certainty) when the value is above the threshold. This significantly constrains the shape of their allocation rule compared to the freedom we have with randomized auctions.

3.2 Proving the “Only If” Part of Theorem 3

We will now prove the “only if” part of Theorem 3 by considering the case of a single bidder of valuation t . We will

use the simplest univariate form of allocation and payment functions. For the sake of contradiction, let us assume that

$$\gamma + \rho \cdot \ln \max \left\{ \hat{u}, \frac{H \cdot \rho}{\gamma} \right\} > 1 \quad (1)$$

and, furthermore, that there is a feasible allocation rule x which is implementable through a payment function p that satisfies $p(t) \geq \rho \cdot t$ for $t \in [1, H]$ and $p(\hat{u}) \geq \gamma \cdot \hat{u}$. I.e., the auction defined by the pair (x, p) is γ -consistent and ρ -robust.

By the robustness requirement, we have $p(t) \geq \rho \cdot t$ for $t \in [1, \hat{u}]$ and, hence,

$$\int_1^{\hat{u}} \frac{p(z)}{z^2} dz \geq \rho \int_1^{\hat{u}} \frac{dz}{z} = \rho \cdot \ln \hat{u}.$$

By applying Lemma 2, we have the individual rationality condition $x(1) - p(1) \geq 0$ and, furthermore, for $s = 1$ and $t = \hat{u}$,

$$\begin{aligned} x(\hat{u}) &= \frac{p(\hat{u})}{\hat{u}} + \int_1^{\hat{u}} \frac{p(z)}{z^2} dz + x(1) - p(1) \\ &\geq \frac{p(\hat{u})}{\hat{u}} + \rho \cdot \ln \hat{u}. \end{aligned} \quad (2)$$

We now distinguish between two cases. First, if $\hat{u} \geq \frac{H \cdot \rho}{\gamma}$, inequality (2) yields (using the consistency requirement $p(\hat{u}) \geq \gamma \cdot \hat{u}$)

$$x(\hat{u}) \geq \gamma + \rho \cdot \ln \hat{u} = \gamma + \rho \cdot \ln \max \left\{ \hat{u}, \frac{H \cdot \rho}{\gamma} \right\} > 1.$$

The strict inequality follows by assumption (1) and contradicts the feasibility of the allocation function x .

In the following, we consider the second case, where $\hat{u} < \frac{H \cdot \rho}{\gamma}$. Define $v = \frac{\gamma \cdot \hat{u}}{\rho}$ and observe that $v \in [\hat{u}, H)$.

By the consistency requirement and payment monotonicity, we have $p(t) \geq \gamma \cdot \hat{u}$ for $t \in [\hat{u}, v)$ and, hence,

$$\int_{\hat{u}}^v \frac{p(z)}{z^2} dz \geq \gamma \cdot \hat{u} \int_{\hat{u}}^v \frac{dz}{z^2} = \gamma - \frac{\gamma \cdot \hat{u}}{v} = \gamma - \rho.$$

Thus, by applying Lemma 2 for $s = \hat{u}$ and $t = v$, we get

$$\begin{aligned} x(v) &= \frac{p(v)}{v} + \int_{\hat{u}}^v \frac{p(z)}{z^2} dz + x(\hat{u}) - \frac{p(\hat{u})}{\hat{u}} \\ &\geq \frac{p(v)}{v} + \gamma - \rho + x(\hat{u}) - \frac{p(\hat{u})}{\hat{u}}. \end{aligned} \quad (3)$$

Also, by the robustness requirement, we have $p(t) \geq \rho \cdot t$ for $t \in [v, H]$ and, hence (using the definition of v),

$$\int_v^H \frac{p(z)}{z^2} dz \geq \rho \int_v^H \frac{dz}{z} = \rho \cdot \ln \frac{H}{v} = \rho \cdot \ln \frac{H \cdot \rho}{\hat{u} \cdot \gamma}.$$

Thus, by applying Lemma 2 for $s = v$ and $t = H$ and using the robustness requirement $p(H) \geq \rho \cdot H$, we get

$$\begin{aligned} x(H) &= \frac{p(H)}{H} + \int_v^H \frac{p(z)}{z^2} dz + x(v) - \frac{p(v)}{v} \\ &\geq \rho + \rho \cdot \ln \frac{H \cdot \rho}{\hat{u} \cdot \gamma} + x(v) - \frac{p(v)}{v}. \end{aligned} \quad (4)$$

Now, by summing equations (2), (3), and (4), and using the assumption $\hat{u} < \frac{H \cdot \rho}{\gamma}$ in the second case and assumption (1), we obtain

$$x(H) \geq \gamma + \ln \frac{H \cdot \rho}{\gamma} = \gamma + \ln \max \left\{ \hat{u}, \frac{H \cdot \rho}{\gamma} \right\} > 1,$$

again contradicting the feasibility of the allocation function. This completes the proof of the “only if” part of Theorem 3.

4 Robustness as Function of Prediction Error

In this section, we give the formal statement and proof of our second result.

Theorem 4. *Let $\rho : [1, H] \rightarrow [0, 1]$ be a differentiable function so that $\rho(\eta)$ is non-increasing and $\eta \cdot \rho(\eta)$ is non-decreasing in η . There exists a ρ -robust DSIC intuitive auction for bidders with valuations from $[1, H]$ and a prediction for the highest bid $\hat{u} \in [1, H]$ if and only if*

$$\rho(H/\hat{u}) + \int_1^{\hat{u}} \frac{\rho(z)}{z} dz + \int_1^{H/\hat{u}} \frac{\rho(z)}{z} dz \leq 1.$$

The proof follows in Sections 4.1 and 4.2. Then, we present specific robustness guarantees (i.e., functions of prediction error) as applications in Section 4.3.

4.1 Proving the “If” Part of Theorem 4

We prove the “if” part of Theorem 4 by constructing a DSIC intuitive auction, which uses the allocation function \tilde{x} and the payment function \tilde{p} defined below. Of course, the definition of both \tilde{x} and \tilde{p} uses the function ρ .

Again, consider an agent i with valuation t . Let b be the highest bid among the remaining agents, and let ν denote the number of bidders different than i with valuation b . The allocation fraction $\tilde{x}(t, b, \nu)$ is defined as

$$\tilde{x}(t, b, \nu) = \begin{cases} 0, & t \in [1, b) \\ \frac{\rho(b/\hat{u})}{\nu+1}, & t = b \\ \rho(t/\hat{u}) + \int_{b/\hat{u}}^{t/\hat{u}} \frac{\rho(z)}{z} dz, & t \in (b, H] \end{cases}$$

if $1 \leq \hat{u} \leq b \leq H$, and as

$$\tilde{x}(t, b, \nu) = \begin{cases} 0, & t \in [1, b) \\ \frac{\rho(\hat{u}/b)}{\nu+1}, & t = b \\ \rho(\hat{u}/t) + \int_{\hat{u}/t}^{\hat{u}/b} \frac{\rho(z)}{z} dz, & t \in (b, \hat{u}) \\ \rho(t/\hat{u}) + \int_1^{\hat{u}/b} \frac{\rho(z)}{z} dz \\ \quad + \int_1^{t/\hat{u}} \frac{\rho(z)}{z} dz, & t \in [\hat{u}, H] \end{cases}$$

if $1 \leq b < \hat{u} \leq H$.

The quantity $\tilde{x}(t, b, \nu)$ is clearly non-negative. We will show that it is also non-decreasing. Let $\sigma_1(\eta) = \eta \cdot \rho(\eta)$ and $\sigma_2(\eta) = \rho(\eta)/\eta$. By the definition of the allocation function for t such that $1 \leq \hat{u} \leq b < t \leq H$ and $1 \leq b < \hat{u} \leq t \leq H$, its derivative with respect to t is

$$\begin{aligned} \tilde{x}'(t, b, \nu) &= \frac{1}{t} \cdot \left(\frac{t}{\hat{u}} \cdot \rho'(t/\hat{u}) + \rho(t/\hat{u}) \right) \\ &= \frac{1}{t} \cdot \sigma_1'(t/\hat{u}). \end{aligned} \quad (5)$$

Also, for t such that $1 \leq b < t < \hat{u} \leq H$, the derivative of the allocation function with respect to t is

$$\begin{aligned} \tilde{x}'(t, b, \nu) &= -\frac{1}{t} \cdot \left(\frac{\hat{u}}{t} \cdot \rho'(\hat{u}/t) - \rho(\hat{u}/t) \right) \\ &= -\frac{u^2}{t^3} \cdot \sigma_2'(\hat{u}/t). \end{aligned} \quad (6)$$

Furthermore, notice that, in any case, function $\tilde{x}(t, b, \nu)$ is continuous at $t = \hat{u}$. Thus, equations (5) and (6) imply that $\tilde{x}(t, b, \nu)$ is indeed non-decreasing (recall that $\sigma_1(\eta)$ is non-decreasing and $\sigma_2(\eta)$ is non-increasing in η and their derivatives are non-negative and non-positive, respectively). To show feasibility, it suffices to show that the highest fraction of the item that is allocated does not exceed 1. Indeed, if $1 \leq \hat{u} \leq b \leq H$, by the definition of the allocation function and the condition of Theorem 4, we have

$$\begin{aligned} \tilde{x}(H, b, \nu) &= \rho(H/\hat{u}) + \int_{b/\hat{u}}^{H/\hat{u}} \frac{\rho(z)}{z} dz \\ &\leq \rho(H/\hat{u}) + \int_1^{H/\hat{u}} \frac{\rho(z)}{z} dz \leq 1, \end{aligned}$$

while, if $1 \leq b < \hat{u} \leq H$, we have

$$\begin{aligned} \tilde{x}(H, b, \nu) &= \rho(H/\hat{u}) + \int_1^{\hat{u}/b} \frac{\rho(z)}{z} dz + \int_1^{H/\hat{u}} \frac{\rho(z)}{z} dz \\ &\leq \rho(H/\hat{u}) + \int_1^{\hat{u}} \frac{\rho(z)}{z} dz + \int_1^{H/\hat{u}} \frac{\rho(z)}{z} dz \leq 1, \end{aligned}$$

as desired. Thus, the allocation function \tilde{x} is feasible and monotone.

The corresponding payment function \tilde{p} is defined as

$$\tilde{p}(t, b, \nu) = \begin{cases} 0, & t \in [1, b) \\ \frac{\rho(b/\hat{u}) \cdot b}{\nu+1}, & t = b \\ \rho(t/\hat{u}) \cdot t, & t \in (b, H] \end{cases}$$

if $1 \leq \hat{u} \leq b \leq H$, and as

$$\tilde{p}(t, b, \nu) = \begin{cases} 0, & t \in [1, b) \\ \frac{\rho(\hat{u}/b) \cdot b}{\nu+1}, & t = b \\ \rho(\hat{u}/t) \cdot t, & t \in (b, \hat{u}) \\ \rho(t/\hat{u}) \cdot t, & t \in [\hat{u}, H] \end{cases}$$

if $1 \leq b < \hat{u} \leq H$.

We will show that \tilde{p} indeed implements the allocation rule \tilde{x} . i.e., \tilde{x} and \tilde{p} are consistent with Lemma 2. For t such that $1 \leq \hat{u} \leq b \leq t \leq H$, the RHS of the expression in Lemma 2 for t and $s = b$ becomes

$$\begin{aligned} &\frac{p(t)}{t} + \int_b^t \frac{p(z)}{z^2} dz + x(b) - \frac{p(b)}{b} \\ &= \rho(t/\hat{u}) + \int_b^t \frac{\rho(z/\hat{u})}{z} dz = \rho(t/\hat{u}) + \int_{t/\hat{u}}^{b/\hat{u}} \frac{\rho(y)}{y} dy \\ &= x(t), \end{aligned}$$

as desired. The second equality follows using the substitution $y = z/\hat{u}$. For t such that $1 \leq b < t < \hat{u} \leq H$, the RHS of the expression in Lemma 2 for t and $s = b$ becomes

$$\rho(\hat{u}/t) + \int_b^t \frac{\rho(\hat{u}/z)}{z} dz = \rho(\hat{u}/t) + \int_{\hat{u}/t}^{\hat{u}/b} \frac{\rho(y)}{y} dy = x(t),$$

as desired. The first equality follows using the substitution $y = \hat{u}/z$. Finally, for t such that $1 \leq b < \hat{u} \leq t \leq H$, the RHS of the expression in Lemma 2 for t and $s = \hat{u}$ becomes

$$\begin{aligned} \frac{p(t)}{t} + \int_{\hat{u}}^t \frac{p(z)}{z^2} dz + x(\hat{u}) - \frac{p(\hat{u})}{\hat{u}} \\ = \rho(t/\hat{u}) + \int_{\hat{u}}^t \frac{\rho(z/\hat{u})}{z} dz + \int_1^{\hat{u}/b} \frac{\rho(z)}{z} dz \\ = \rho(t/\hat{u}) + \int_1^{t/\hat{u}} \frac{\rho(y)}{y} dy + \int_1^{\hat{u}/b} \frac{\rho(z)}{z} dz = x(t), \end{aligned}$$

as desired, again. The second equality follows using the substitution $y = z/\hat{u}$.

Notice that, clearly, the auction's revenue yields the required robustness. In particular, assuming, without loss of generality, that agent i has the highest valuation (i.e., $t \geq b$), the revenue is exactly $\rho(\hat{u}/t) \cdot t$ for $t \in [b, \hat{u}]$ and exactly $\rho(t/\hat{u}) \cdot t$ for $t \in [\hat{u}, H]$. The proof of the ‘‘if’’ part of Theorem 4 is now complete.

4.2 Proving the ‘‘Only If’’ Part of Theorem 4

Like in the corresponding proof for Theorem 3, we will prove the ‘‘only if’’ part of Theorem 4 by considering the case of a single bidder. For the sake of contradiction, let us assume that

$$\rho(H/\hat{u}) + \int_1^{\hat{u}} \frac{\rho(z)}{z} dz + \int_1^{H/\hat{u}} \frac{\rho(z)}{z} dz > 1 \quad (7)$$

and, furthermore, that there exists a feasible and monotone allocation function x which is implementable through a payment function p that satisfies $p(t) \geq \rho(\hat{u}/t) \cdot t$ for $t \in [1, \hat{u}]$ and $p(t) \geq \rho(t/\hat{u}) \cdot t$ for $t \in [\hat{u}, H]$.

By applying Lemma 2, we have the individual rationality condition $x(1) - p(1) \geq 0$ and, for $s = 1$ and $t = \hat{u}$,

$$\begin{aligned} x(\hat{u}) &= \frac{p(\hat{u})}{\hat{u}} + \int_1^{\hat{u}} \frac{p(z)}{z^2} dz + x(1) - p(1) \\ &\geq \frac{p(\hat{u})}{\hat{u}} + \int_1^{\hat{u}} \frac{\rho(\hat{u}/z)}{z} dz. \end{aligned} \quad (8)$$

By applying Lemma 2 for $s = \hat{u}$ and $t = H$ and using our assumptions for the payment function p , we get

$$\begin{aligned} x(H) &= \frac{p(H)}{H} + \int_{\hat{u}}^H \frac{p(z)}{z^2} dz + x(\hat{u}) - \frac{p(\hat{u})}{\hat{u}} \\ &\geq \rho(H/\hat{u}) + \int_{\hat{u}}^H \frac{\rho(z/\hat{u})}{z} dz + x(\hat{u}) - \frac{p(\hat{u})}{\hat{u}}. \end{aligned} \quad (9)$$

By inequalities (8) and (9), we have

$$x(H) \geq \rho(H/\hat{u}) + \int_1^{\hat{u}} \frac{\rho(\hat{u}/z)}{z} dz + \int_{\hat{u}}^H \frac{\rho(z/\hat{u})}{z} dz$$

$$\begin{aligned} &= \rho(H/\hat{u}) + \int_1^{\hat{u}} \frac{\rho(y)}{y} dy + \int_1^{H/\hat{u}} \frac{\rho(y)}{y} dy \\ &> 1, \end{aligned}$$

contradicting the feasibility of the allocation function x . The equality follows by the substitution $y = \hat{u}/z$ and $y = z/\hat{u}$ in the two integrals and the second inequality uses our assumption (7). The proof of the ‘‘only if’’ part of Theorem 4 is complete.

4.3 Applications of Theorem 4

We now present applications of Theorem 4. In particular, in Corollary 5, we present robustness functions that satisfy the condition of Theorem 4. These are just indicative of what Theorem 4 can give us, and have the characteristic that the quantity $1/\rho(\eta)$ depends polylogarithmically, logarithmically, and sublogarithmically on the prediction error, respectively. The explicit allocation and payment functions of the corresponding auctions then follow by applying the machinery of Section 4.1. We remark that for the first ρ -function in Corollary 5, which does not depend on H , the corresponding auction satisfies the claimed robustness requirement even when it is applied to settings with agent valuations from the interval $[1, \infty)$.

Corollary 5. *Let $H > 1$. For the functions*

$$\begin{aligned} \rho(\eta) &:= \frac{1}{(\pi/\varepsilon + 1) \cdot (1 + \ln^{1+\varepsilon}(\eta))}, \text{ with } \varepsilon \in (0, 1], \\ \rho(\eta) &:= \frac{1}{1 + 2 \ln(1 + \ln H)} \cdot \frac{1}{1 + \ln \eta}, \text{ and} \\ \rho(\eta) &:= \frac{1 - \varepsilon}{2(1 + \ln H)^{1-\varepsilon}} \frac{1}{(1 + \ln \eta)^\varepsilon}, \text{ with } \varepsilon \in (0, 1), \end{aligned}$$

there exist ρ -robust DSIC intuitive auctions for bidders with valuations from $[1, H]$ and a prediction for the highest bid.

We omit the proof due to lack of space.

5 Conclusion

We have presented an optimal trade-off for the consistency and robustness of DSIC intuitive single-item auctions that use predictions. This result gives us auctions with constant consistency and robustness $\Omega(\ln^{-1} H)$, where H upper bounds the agent valuations. We have also given a sufficient and necessary condition for DSIC intuitive auctions with a robustness guarantee that is a function of the prediction error. As a corollary, we have obtained auctions with a robustness that depends only on the prediction error η (e.g., as $\Omega(\ln^{-2} \eta)$ and even better) and not in H , or auctions that have a small (logarithmic or sublogarithmic) dependence on H and better dependence on the prediction error. An obvious extension of our work would be to explore trade-offs between consistency and robustness in more general single-parameter mechanism design environments. In addition to the revenue extracted, these can also involve the social welfare. We expect that our variant of Myerson's Lemma will find more applications there.

References

- [Agrawal *et al.*, 2022] Priyank Agrawal, Eric Balkanski, Vasilis Gkatzelis, Tingting Ou, and Xizhi Tan. Learning-augmented mechanism design: Leveraging predictions for facility location. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*, pages 497–528, 2022.
- [Balcan *et al.*, 2023] Maria-Florina Balcan, Siddharth Prasad, and Tuomas Sandholm. Bicriteria multidimensional mechanism design with side information. In *Proceedings of the 37th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2023.
- [Balkanski *et al.*, 2023a] Eric Balkanski, Vasilis Gkatzelis, and Xizhi Tan. Strategyproof scheduling with predictions. In *Proceedings of the 14th Innovations in Theoretical Computer Science Conference (ITCS)*, pages 11:1–11:22, 2023.
- [Balkanski *et al.*, 2023b] Eric Balkanski, Vasilis Gkatzelis, Xizhi Tan, and Cherlin Zhu. Online mechanism design with predictions. *CoRR*, abs/2310.02879, 2023.
- [Berger *et al.*, 2023] Ben Berger, Michal Feldman, Vasilis Gkatzelis, and Xizhi Tan. Optimal metric distortion with predictions. *CoRR*, abs/2307.07495, 2023.
- [Gkatzelis *et al.*, 2022] Vasilis Gkatzelis, Kostas Kollias, Alkmini Sgouritsa, and Xizhi Tan. Improved price of anarchy via predictions. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*, pages 529–557, 2022.
- [Goldberg *et al.*, 2006] Andrew V. Goldberg, Jason D. Hartline, Anna R. Karlin, Michael Saks, and Andrew Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.
- [Hartline, 2013] Jason D. Hartline. Bayesian mechanism design. *Foundations and Trends® in Theoretical Computer Science*, 8(3):143–263, 2013.
- [Istrate and Bonchis, 2022] Gabriel Istrate and Cosmin Bonchis. Mechanism design with predictions for obnoxious facility location. *CoRR*, abs/2212.09521, 2022.
- [Lu *et al.*, 2023] Pinyan Lu, Zongqi Wan, and Jialin Zhang. Competitive auctions with imperfect predictions. *CoRR*, abs/2309.15414, 2023.
- [Mitzenmacher and Vassilvitskii, 2020] Michael Mitzenmacher and Sergei Vassilvitskii. Algorithms with predictions. In Tim Roughgarden, editor, *Beyond the Worst-Case Analysis of Algorithms*, pages 646–662. Cambridge University Press, 2020.
- [Myerson, 1981] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [Roughgarden, 2016] Tim Roughgarden. *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press, 2016.
- [Vickrey, 1961] William Vickrey. Counterspeculation, auctions, and competitive tenders. *The Journal of Finance*, 16(1):8–37, 1961.
- [Xu and Lu, 2022] Chenyang Xu and Pinyan Lu. Mechanism design with predictions. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 571–577, 2022.