

Envy, Regret, and Social Welfare Loss

Riccardo Colini-Baldeschi
rickuz@fb.com
Facebook, Core Data Science
London, UK

Okke Schrijvers

okke@fb.com
Facebook, Core Data Science
Menlo Park, US

Stefano Leonardi*

leonardi@diag.uniroma1.it
Sapienza University of Rome
Rome, Italy
Facebook, Core Data Science
London, UK

Eric Sodomka

sodomka@fb.com
Facebook, Core Data Science
London, UK

ABSTRACT

Incentive compatibility (IC) is a desirable property for any auction mechanism, including those used in online advertising. However, in real world applications practical constraints and complex environments often result in mechanisms that lack incentive compatibility. Recently, several papers investigated the problem of deploying black-box statistical tests to determine if an auction mechanism is incentive compatible by using the notion of IC-Regret that measures the regret of a truthful bidder. Unfortunately, most of those methods are computationally intensive, since they require the execution of many counterfactual experiments.

In this work, we show that similar results can be obtained using the notion of IC-Envy. The advantage of IC-Envy is its efficiency: it can be computed using only the auction’s outcome. In particular, we focus on position auctions. For position auctions, we show that for a large class of pricing schemes (which includes e.g. VCG and GSP), $IC\text{-Envy} \geq IC\text{-Regret}$ (and $IC\text{-Envy} = IC\text{-Regret}$ under mild supplementary conditions). Our theoretical results are completed showing that, in the position auction environment, IC-Envy can be used to bound the loss in social welfare due to the advertiser untruthful behavior.

Finally, we show experimentally that IC-Envy can be used as a feature to predict IC-Regret in settings not covered by the theoretical results. In particular, using IC-Envy yields better results than training models using only price and value features.

KEYWORDS

Position Auctions, Envy-freeness, Incentive-Compatibility Measurement, Social Welfare Loss

*Partially supported by ERC Advanced Grant 788893 AMDROMA “Algorithmic and Mechanism Design Research in Online Markets” and MIUR PRIN project ALGADIMAR “Algorithms, Games, and Digital Markets”

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW ’20, April 20–24, 2020, Taipei, Taiwan

© 2020 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-7023-3/20/04.

<https://doi.org/10.1145/3366423.3380057>

ACM Reference Format:

Riccardo Colini-Baldeschi, Stefano Leonardi, Okke Schrijvers, and Eric Sodomka. 2020. Envy, Regret, and Social Welfare Loss. In *Proceedings of The Web Conference 2020 (WWW ’20)*, April 20–24, 2020, Taipei, Taiwan. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3366423.3380057>

1 INTRODUCTION

Over the past decades, online advertising has grown into a huge industry, with many different online publishers offering impression opportunities. Auction theory has played a major role in shaping this ecosystem, and many ad auctions strive to be *Incentive Compatible* (IC), which means that an advertiser achieves the best outcome by truthfully reporting their willingness-to-pay. Despite the role that auction theory has played, the resulting systems may not be IC. For example, intermediaries (called Demand Side Platforms or DSPs) first run an auction to determine the best ad among their clients, and then pass this along to a publisher who runs their own auction including bids from other sources. Even when both auctions in isolation are IC, their composition is not. Furthermore, some publishers use past bids to set a reserve price (or minimum bid), and others are moving to a pay-your-bid model entirely (which has strong incentives to misreport willingness-to-pay) [35, 36].

Recently, there have been several works addressing the problem of determining whether an auction is IC based on statistical tests using only inputs and outputs of an (unobserved) auction mechanism [17, 26]. This gives advertisers the power to test whether an auction is IC without having access to the code. Feng et al. [17] use regret [21], similar to Lubin and Parkes [28], as a way to measure “how far” an auction is from being IC:

$$IC\text{-Regret}_i(v_i) = \max_{b_i} \mathbb{E}_{b_{-i}} [u_i(b_i, b_{-i}) - u_i(v_i, b_{-i})], \quad (1)$$

where v_i is the true value of advertiser i , b_i the bid of i , b_{-i} the bids of other advertisers, and $u_i(\cdot)$ the (expected) utility of i . IC-Regret captures the difference in utility between bidding truthfully, and the maximum utility achievable. By definition, IC mechanisms have IC-Regret 0, while higher IC-Regret indicates a stronger incentive to misreport.

While measuring IC is most naturally a concern for advertisers (who cannot observe the auction mechanism directly), it is also both important and non-trivial for the auctioneer. The auctioneer cares about IC auctions because they admit simple optimal bidding

strategies (namely truthful reporting), and simple bidding strategies may in turn lead to lower churn of advertisers.¹ Additionally, when auctions are not IC, bidders don't truthfully report their value, which may harm social welfare² and thus the quality of the service provided to the advertisers. While important, it may not be straightforward for publishers to guarantee that their auction is IC for many different reasons: complex interaction between different layers in the advertising ecosystem, running-time constraints, bugs in the auction code, and so forth.

Feng et al. [17] proposed a method to determine IC-Regret for publishers (by taking a worst-case perspective over the advertiser value v_i). A downside of their method is that it requires many counterfactual evaluations of the auction's outcomes for alternative bids. This means that the auction code needs to be run many times over. While this may be the best thing one can do with only black-box access to the auction mechanism, could we do better if we're using intermediate data from the auction mechanism?

To overcome the practical difficulties to measure IC-Regret, we propose to use Envy [20] as a proxy for IC-Regret, by identifying relevant classes of auction mechanisms where Envy and IC-Regret coincide or where Envy is a good proxy for IC-Regret. So what is Envy? Instead of comparing the advertiser's utility against her utility for alternative bids, Envy takes a single outcome, and measures to what extent advertisers are happy with the outcome. Ad auctions typically simultaneously sell multiple ad slots with varying click-through rates [14, 40]. Let X be an expected allocation vector (where allocation corresponds to the ad being clicked) and p be an expected pricing vector. Envy of bidder i given the outcome (X, p) is:

$$\text{Envy}_i(v_i, X, p) = \max_j (X_j \cdot v_i - p_j) - (X_i \cdot v_i - p_i) \quad (2)$$

Envy is defined with respect to some outcome. However, in an auction, changing one's bid may change the outcome in the auction. Therefore, we define IC-Envy as the Envy experienced in the outcome *when a bidder bids truthfully*. In the following, let $X(v_i, b_{-i})$ is the expected allocation of all bidders and $p(v_i, b_{-i})$ be the expected payment vector of all bidders, when bidder i bids truthfully and for bids b_{-i} of the remaining bidders. IC-Envy is then:

$$\text{IC-Envy}_i(v_i) = \max_j (X_j(v_i, b_{-i}) \cdot v_i - p_j(v_i, b_{-i})) \quad (3)$$

$$- (X_i(v_i, b_{-i}) \cdot v_i - p_i(v_i, b_{-i})) \quad (4)$$

In arbitrary auction environments, IC-Envy and IC-Regret don't necessarily coincide. There are natural auction environments with envy-free outcomes, that still have positive IC-Regret and vice versa.

Example 1.1 (IC-Envy = 0, IC-Regret is positive). Consider a single-item, first-price auction, with two bidders with values $v_1 = \$10$ and $v_2 = \$8$ and assume bidder 1 bids truthfully. The IC-Regret for bidder 1 is $\$2 - \epsilon$ for arbitrary small ϵ , as the best alternative for

bidder 1 is $\$8 + \epsilon$. However, IC-Envy is $\$0$ as the only alternative allocation for bidder 1 is to not receive the item.

Example 1.2 (IC-Regret = 0, IC-Envy is positive). Consider the IC auction for a single item with 2 bidders who face different reserve prices³ $r_1 = \$1$ and $r_2 = \$5$. With bids $b_1 = b_2 = \$3$, bidder 1 received the good at their reserve price of $\$1$. Bidder 2 has IC-Envy of $\$2$, but there is no counterfactual bid that will give her positive utility (hence IC-Regret = 0).

While in general IC-Regret and IC-Envy can be quite different, in this paper we show $\text{IC-Regret} \leq \text{IC-Envy}$ for a large class of position auction environments which includes VCG and GSP [14, 40], and under mild conditions we prove that $\text{IC-Regret} = \text{IC-Envy}$. Since computing IC-Envy requires no counterfactual evaluation of the algorithm, it can serve as an efficient certificate that IC-Regret is low.⁴ In Section 4, we also use IC-Envy to bound the loss in social welfare due to misreports. Finally, we show empirically that IC-Envy can be used as a feature in an estimator for IC-Regret in auction environments beyond those for which we have theoretical results (Section 5).

1.1 Related Work

We propose to connect IC-Envy and IC-Regret directly by defining a large class of auction mechanisms for which $\text{IC-Envy} = \text{IC-Regret}$ (and a larger class where $\text{IC-Envy} \geq \text{IC-Regret}$). The line of work that's closest in spirit aims to identify classes of auction mechanisms that are simultaneously envy-free and IC (in our notation: classes for which $\text{IC-Regret} = \text{IC-Envy} = 0$). Feldman et al. [16] and Goldberg et al. [22] studied the conditions that are required in order to have mechanisms that are efficient, truthful and envy-free and that VCG satisfy these properties for capacitated valuation functions. For homogeneous capacities there's a class of mechanisms that achieve this, while for heterogeneous capacities there is no mechanisms that simultaneously achieved all 3 conditions. Cohen et al. [7] provided a characterization based on cycle-monotonicity of the allocation functions that are incentive-compatible and envy-free without considering the efficiency of the algorithms.

The notion of envy-freeness was initially introduced by Varian [39] and Foley [20]. The key property of an envy-free allocation is that buyers prefer the bundle of goods they receive over any other allocated bundles (given bundle prices). The notion is particularly appealing due to its connection to markets: in an envy-free allocation, given the prices for goods, all buyers prefer to buy the bundle that's assigned to them. More recently, the notion of envy-freeness has been deeply studied with a different perspective that involves item-pricing [10, 23] and bundle-pricing [8, 15, 18]. In our setting there is no difference between those two models. A similar line of work focused in studying envy-free algorithms, both in the item-pricing and the bundle pricing models, when the bidders have budget constraints [2, 8, 10, 15, 37].

In much of the other related work, envy-freeness is taken as an alternative solution concept to IC (e.g. [13, 15, 18, 24]) and in

¹Moreover, additional advertisers are better for revenue than being clever about devising a revenue-optimal auction mechanism, see [3] (and follow-ups).

²While in some cases there are symmetric equilibria in which social welfare is not harmed [40], these need not always exist and bidders may not reach equilibrium [1, 12, 34].

³While non-anonymous reserve prices may seem contrived, they naturally occur e.g. for the revenue-optimal auction on non-i.i.d. bidders [31].

⁴For an auctioneer who cares about incentive compatibility, false positives (i.e. high IC-Envy but low IC-Regret) are acceptable while false negatives are not. Therefore the inequality goes in the right direction.

contexts outside of the auction domain (e.g. [5, 6, 19, 30]). Of particular note: Daskalakis and Syrgkanis [13] address the relation between envy-freeness and incentive-compatibility in the context of algorithmic learning. In particular, the authors discussed the computational complexity of no-regret learning algorithms and no-envy algorithms in simultaneous second price auctions. Hartline and Yan [24] studied the relation between envy-freeness and incentive compatibility in revenue-maximizing prior-free mechanisms. Lipton et al. [27] investigated envy-free mechanisms in the context of indivisible items with focus on the computational complexity of finding allocations with minimum envy. Moreover, they proved that it is possible to obtain truthful mechanisms with bounded envy. Those results have been simplified and extended by Caragiannis et al. [4]. While this line of work is interesting, it does not quantitatively address the relationship of envy and IC regret.

1.2 Our Contributions

This paper has four main contributions:

- (1) First, in Section 3, we define a class of auction mechanisms— which includes VCG, GSP, and GFP for position auctions— where IC-Envy is tightly related to IC-Regret. We show that almost all payment rules in this class have $\text{IC-Envy} \geq \text{IC-Regret}$ and we give supplementary conditions under which they are exactly equal.
- (2) Secondly, in Section 4, we upperbound the social welfare loss in terms of IC-Envy for the same sets of mechanisms introduced in Sections 3. We show that in equilibrium, the social welfare loss is *at most* $4 \cdot \text{IC-Envy}$ (under a technical condition we introduce in the section).
- (3) Finally, in Section 5, we use bidding data from a major online publisher to show that IC-Envy can be used as a feature to learn an estimator for IC-Regret. The estimator has low mean-squared error, and performs better than comparable estimators that are trained using other features from the auction like values and prices for different slots.

All missing proofs are given in the full version of the paper [9]. The extended version [9] also considers the more general Ad Types auction environment [11] in which different ads have different discount curves. It is shown that for a suitable generalization of GSP it still holds that $\text{IC-Envy} \geq \text{IC-Regret}$.

2 PRELIMINARIES

There are n bidders and m slots. Let I be the set of bidders and J be an (ordered) set of slots. Each bidder $i \in I$ has a valuation vector $\mathbf{v}_i = \langle v_{i,1}, v_{i,2}, \dots, v_{i,m} \rangle$ that is the willingness to pay of bidder i for each slot j , with $v_{i,1} \geq v_{i,2} \geq \dots \geq v_{i,m}$. Bidders are unit demand. In the standard position auction environment, slots have common quality factor $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_m$ such that for each bidder i and slot j we have $v_{i,j} = v_i \cdot \alpha_j$ for private value $v_i \in \mathbb{R}$ of the bidder.

The slots are allocated to the bidders by a (direct-revelation) mechanism \mathcal{M} . The mechanism \mathcal{M} is defined by an allocation function $\mathcal{A} : \mathbb{R}^n \rightarrow \mathbb{N}^n$ and a payment function $\mathcal{P} : \mathbb{R}^n \rightarrow \mathbb{R}^n$. Since bidders' values v_i are private, the mechanism solicits bids \mathbf{b}_i to represent the values, though reports may not be truthful. Let \mathbf{v} be the vector representing the valuations of all the bidders and \mathbf{b}

the vector representing the bids. After receiving \mathbf{b} , the mechanism $\mathcal{M} = \langle \mathcal{A}, \mathcal{P} \rangle$ computes an outcome $\langle X, \mathbf{p} \rangle$, i.e., $\mathcal{A}(\mathbf{b}) = X$ and $\mathcal{P}(\mathbf{b}) = \mathbf{p}$.

X describes the allocation of the slots to the bidders and \mathbf{p} describes how much each bidder is charged for the obtained slot. In particular, $X = \langle X_1, X_2, \dots, X_n \rangle$ where $X_i = j$, if the bidder i obtains the slot j and 0 if she does not receive any slot. And $\mathbf{p} = \langle p_1, p_2, \dots, p_n \rangle$ where $p_i \in \mathbb{R}_{\geq 0}$ is the price that the bidder i pays for slot X_i .

For an allocation X and a valuation vector \mathbf{v} , the social welfare of the allocation X is defined as $SW(\mathbf{v}, X) = \sum_{i \in I} v_i \alpha_{X_i}$. In turn, the optimal social welfare is defined as $SW^{OPT}(\mathbf{v}) = \max_X \sum_{i \in I} v_i \alpha_{X_i}$. Finally, the *Social Welfare Loss* (SWL) is defined as $SWL(\mathbf{v}, X) = SW^{OPT}(\mathbf{v}) - SW(\mathbf{v}, X)$. When the valuation vector \mathbf{v} is clear from the context, we will use $SW(X)$, SW^{OPT} , and $SWL(X)$.

Given an outcome $\langle X, \mathbf{p} \rangle$, the utility of a bidder i is $u_i(v_i, X, \mathbf{p}) = v_i \alpha_{X_i} - p_i$. Since the outcome of a mechanism \mathcal{M} is a function of the bids, and the auctions we consider are not necessarily IC, bidders may be incentivized to report a type \mathbf{b} different from \mathbf{v} in order to produce an outcome with higher utility.

IC-Regret. IC-Regret describes the outcome for bidding truthfully, compared to the optimal alternative bid (given constant competition \mathbf{b}_{-i}). Formally, the regret of a bidder i for bidding truthfully compared to a *specific* alternative bid b_i is:

$$r_i(b_i, \mathbf{b}_{-i}, v_i) = u_i(v_i, \mathcal{A}(b_i, \mathbf{b}_{-i}), \mathcal{P}(b_i, \mathbf{b}_{-i})) - u_i(v_i, \mathcal{A}(v_i, \mathbf{b}_{-i}), \mathcal{P}(v_i, \mathbf{b}_{-i})),$$

which is used in the formal definition for IC-Regret.

Definition 2.1 (IC-Regret). The IC-Regret that bidder i experiences is⁵

$$\text{IC-Regret}_i(v_i, \mathbf{b}_{-i}) = \max_{b_i \in \mathbb{R}^+} \{r_i(b_i, \mathbf{b}_{-i}, v_i)\}.$$

IC-Regret can be directly connected to *incentive-compatibility (IC)*. Indeed, a mechanism \mathcal{M} is IC iff for all v_i, \mathbf{b}_{-i} , and $i \in I$, we have $r_i(b_i, \mathbf{b}_{-i}, v_i) = 0$ for all $b_i \in \mathbb{R}^+$.

IC-Envy. Given an allocation X and payments \mathbf{p} , Envy describes how much a bidder prefers the allocation and price of another buyer, compared to what she received. Since different bids may lead to a different auction outcome, we define IC-Envy_i as Envy with respect to the allocation X and payments \mathbf{p} when bidder i bids *truthfully*. IC-Envy is some notion of fairness of the produced outcome whereas IC-Regret measures how much the underlying mechanism incentives misreported types.

Formally, for given an allocation X and payments \mathbf{p} , the envy that bidder i experiences compared to bidder j is

$$e_i^j(v_i, X, \mathbf{p}) = (X_j(v_i, \mathbf{b}_{-i}) \cdot v_i - p_j(v_i, \mathbf{b}_{-i})) - (X_i(v_i, \mathbf{b}_{-i}) \cdot v_i - p_i(v_i, \mathbf{b}_{-i})), \quad (5)$$

and the envy of bidder i in the outcome $\langle X, \mathbf{p} \rangle$ is which is used in the formal definition for IC-Envy in the outcome $\langle X, \mathbf{p} \rangle$:

⁵Equation (1) in the introduction takes an expectation over competition \mathbf{b}_{-i} since the work of Feng et al. [17] considers the auction mechanism as a black box. In our setting (from the perspective of the auctioneer) the alternative bids are known, and we define IC-Regret on an auction-by-auction basis.

Definition 2.2 (IC-Envy). The IC-Envy of bidder i in the outcome $(\mathcal{A}(v_i, \mathbf{b}_{-i}), \mathcal{P}(v_i, \mathbf{b}_{-i}))$:

$$\text{IC-Envy}_i(v_i, \mathbf{b}_{-i}) = \max_{j \in I} \{e_i^j(v_i, \mathcal{A}(v_i, \mathbf{b}_{-i}), \mathcal{P}(v_i, \mathbf{b}_{-i}))\}.$$

Note that computing IC-Envy requires one single execution of the auction whereas the computation of IC-Regret requires the execution of the auction for multiple bid values of each bidder.

Along the paper we often use $\text{IC-Regret} \leq \text{IC-Envy}$ with the meaning that $\text{IC-Regret}_i(v_i, \mathbf{b}_{-i}) \leq \text{IC-Envy}_i(v_i, \mathbf{b}_{-i})$ holds for all bidder $i \in I$.

3 POSITION AUCTION ENVIRONMENTS

As stated before, IC-Envy and IC-Regret measure different things: IC-Envy provides some measure of fairness of the outcome, whereas IC-Regret measures the incentive-compatibility of the mechanism. In this section, we study the relationship between IC-Envy and IC-Regret for position auctions that are widely used in search and feed advertising.⁶ In order to characterize the class of position auctions that have $\text{IC-Envy} \geq \text{IC-Regret}$, we introduce the definition of *regular mechanism*.

Definition 3.1 (Regular Mechanisms for Position Auctions). A mechanism \mathcal{M} is called a *regular mechanism for position auctions* if there exist a matrix $A^{n \times m} \in \mathbb{R}_{\geq 0}^{n \times m}$ depending on the quality vector $(\alpha_1, \dots, \alpha_m)$ such that for every bid vector \mathbf{b} the following conditions hold:

- (1) Slots are assigned in order of non-increasing α_j to bidders ordered by non-increasing bid value b_i . Ties are broken lexicographically.
- (2) The payment of bidder i is defined as $p_i = \sum_{j=1}^n a_{i,j} \cdot b_j$ with $a_{i,j} \in A$.
- (3) The mechanism is individually rational (IR), i.e., $p_i \leq \alpha_i b_i$.

This definition includes several widely used auction mechanisms, as for example:

- VCG: $a_{i,j} = 0$ for $j = 1, \dots, i$, and $a_{i,j} = \alpha_{j-1} - \alpha_j$ for $k = i+1, \dots, n$.
- GSP: $a_{i,j} = 0$ for $j = 1, \dots, i$, $a_{i,i+1} = \alpha_{i+1}$, and $a_{i,j} = 0$ for $j = i+2, \dots, n$.
- GFP: $a_{i,j} = 0$ for $j = 1, \dots, i-1$, $a_{i,i} = \alpha_i$, and $a_{i,j} = 0$ for $j = i+1, \dots, n$.

First of all, note the individual rationality of a regular mechanisms for position auctions implies that

LEMMA 3.2. *Regular mechanisms for position auctions must have $a_{i,j} = 0$ for all $j < i$.*

PROOF. We prove by contraposition: If $a_{i,j} > 0$ for some $j < i$, then consider bids $b_1 = \dots = b_j = 1$ and $b_{j+1} = \dots = b_n = 0$ and $\alpha_1 = \dots = \alpha_j = 1$ and $\alpha_{j+1} = \dots = \alpha_n = 0$. We have $p_i = a_{i,j} b_j > 0$, while $\alpha_i b_i = 0$, hence $p_i > \alpha_i b_i$ which violates IR. \square

We next provide conditions for a regular mechanism in order to have $\text{IC-Envy} \leq \text{IC-Regret}$.

⁶In display advertising it is more common to sell ad slots one-by-one, which is a special case of position auctions, though one which is arguably mathematically less interesting.

THEOREM 3.3. *For a regular mechanism for position auctions, for each bidder i with true value v_i and all possible bid profiles \mathbf{b}_{-i} , if*

B1 $a_{i,i} = 0$, for each bidder i ,

then it must hold that

A1 $\text{IC-Envy}_i(v_i, \mathbf{b}_{-i}) \geq \text{IC-Regret}_i(v_i, \mathbf{b}_{-i})$.

PROOF. We first prove that B1 and regularity imply $\text{IC-Envy}_i(v_i, \mathbf{b}_{-i}) \geq \text{IC-Regret}_i(v_i, \mathbf{b}_{-i})$ if $b_i > v_i$. If bidder i is still assigned to slot i with bid $b_i > v_i$, condition B1 ($a_{i,i} = 0$) implies that price p_i is not increased. Assume otherwise bidder i wins a higher slot $j < i$ with bid b_i . Price $p_j = \sum_{k=j+1}^n a_{j,k} \cdot b_k$ charged to bidder i is computed on the basis of a monotonically higher set of bids, since coefficients $a_{j,k}$ are non-negative and $a_{j,j} = 0$. This implies that the counterfactual price must be higher or equal than the price computed for slot j with truthful bid $b_j = v_j$. Indeed, the corresponding price p_j is computed on a bid vector $(b_j, \dots, b_{i-1}, b_{i+1}, \dots, b_n)$ that weakly dominates term by term the bid vector $(b_{j+1}, \dots, v_i, b_{i+1}, \dots, b_n)$.

Let us now consider the case $b_i < v_i$. From condition B1 we immediately derive that p_i is not changed if bidder i is still allocated to the same slot i , and therefore there is no positive regret. Assume bid b_i is such that bidder i is allocated slot $j > i$. We observe that $e_i^j(v_i, \mathbf{X}, \mathbf{p}) = r_i(b_i, \mathbf{b}_{-i}, v_i)$ since the computation of the price for slot j on both (v_i, \mathbf{b}_{-i}) and (b_i, \mathbf{b}_{-i}) is based on bid vector (b_{j+1}, \dots, b_n) . \square

Ideally, we would like to state condition A1 of Theorem 3.3 with equality in order to provide with IC-Envy the same exact information of IC-Regret. The two following obstacles may prevent the possibility of proving this stronger claim: i) Bidder i bidding a value $b_i > v_i$ could obtain a slot $j < i$ with a payment $p_j(b_i, \mathbf{b}_{-i}) > p_j(v_i, \mathbf{b}_{-i})$. ii) Bidder i may not be able to obtain a slots $j > i$ by bidding some value $b_i < v_i$ if there are at two agents with identical bids since ties are broken lexicographically. In both cases we may have $\text{IC-Envy} > \text{IC-Regret}$. However, in the next theorem we state additional realistic conditions that imply $\text{IC-Envy}_i(v_i, \mathbf{b}_{-i}) = \text{IC-Regret}_i(v_i, \mathbf{b}_{-i})$.

THEOREM 3.4. *For a regular mechanism for position auctions, for each bidder i with true value v_i and all possible bid profiles \mathbf{b}_{-i} , if*

B1 $a_{i,i} = 0$, for each bidder i , and

B2 $p_j - p_{j+1} \geq (\alpha_j - \alpha_{j+1}) b_{j+1}$, for $j = 1, \dots, n-1$, and

B3 all bids b_1, \dots, b_n are distinct

then it must hold that

A2 $\text{IC-Envy}_i(v_i, \mathbf{b}_{-i}) = \text{IC-Regret}_i(v_i, \mathbf{b}_{-i})$.

The two additional conditions B2 and B3 allow to remove the two obstacles mentioned above. The reason of condition B2 is to enforce a difference in payment between any two slots i and $j < i$ in order to avoid any envy from bidder i towards bidder j . The reason of condition B3 is to allow bidder i to obtain all slots $j > i$ by posting a suitable bid $b_i < v_i$.

PROOF. We first prove that B2 implies A2 if bidder i increases her bid and obtain a higher slot $j < i$. Indeed, for all slots $j < i$ that are obtained with some value $b_i < v_i$, since the increase in payment is at least as large as the increase in utility, we have that

the following expression for the envy of bidder i on slot j :

$$\begin{aligned} e_i^j(v_i, \mathbf{X}, \mathbf{p}) &= (\alpha_j \cdot v_i - p_j) - (\alpha_i \cdot v_i - p_i) \\ &= (\alpha_j - \alpha_i) v_i - (p_j - p_i) \leq 0, \end{aligned}$$

with the last inequality implied by

$$\begin{aligned} p_j - p_i &= \sum_{k=j}^{i-1} (p_k - p_{k+1}) \geq \sum_{k=j}^{i-1} (\alpha_k - \alpha_{k+1}) b_{k+1} \\ &\geq \sum_{k=j}^{i-1} (\alpha_k - \alpha_{k+1}) v_i = (\alpha_j - \alpha_i) v_i \end{aligned}$$

We conclude the first part of the proof by showing that regret is upper bounded by envy for any bid $b_i \geq v_i$ that makes bidder i to win a higher slot $j < i$. Indeed, the price $p_j = \sum_{k=j+1}^n a_{j,k} \cdot b_k$ charged to bidder i is based on a monotonically higher set of bids, since bid vector $(b_j, \dots, b_{i-1}, b_{i+1}, \dots, b_n)$ dominates term by term the bid vector (b_{j+1}, \dots, b_n) . We have therefore proved $r_i(b_i, \mathbf{b}_{-i}, v_i) \leq e_i^j(v_i, \mathbf{X}, \mathbf{p}) \leq 0$ for any $b_i > v_i$ that gives bidder i a slot $j \leq i$.

We have established so far that envy for bidder i can only be positive for a slot $j > i$ and regret can only be positive if bid b_i is decreased. Assume bid b_i is decreased in a way that bidder i is allocated to a slot $j > i$. We have already shown in the proof of Theorem 3.3 that $e_i^j(v_i, \mathbf{X}, \mathbf{p}) = r_i(b_i, \mathbf{b}_{-i}, v_i)$. Moreover, given condition B3, all slots $j > i$ can be obtained from bidder i by decreasing her bid and therefore $\text{IC-Envy}(v_i, \mathbf{b}_{-i}) = r_i(b_i, \mathbf{b}_{-i}, v_i)$ for all values $b_i < v_i$. \square

The conditions of Theorem 3.3 and Theorem 3.4 hold for mechanisms like VCG, GSP, and any combination of the two. For Theorem 3.4 we additionally require that all bids are different and this can be obtained through a small perturbation of the bids. Condition B1 is clearly true for the two mechanisms. Condition B2 is true with $p_i - p_{i+1} = (\alpha_i - \alpha_{i+1})b_{i+1}$ for VCG and with $p_i - p_{i+1} = \alpha_i b_i - \alpha_{i+1} b_{i+1} \geq (\alpha_i - \alpha_{i+1})b_{i+1}$ for GSP. On the contrary, even the first condition of the theorems is violated for GFP.

4 MEASURING SOCIAL WELFARE LOSS WITH IC-ENVY

We proved in Section 3 that IC-Envy is a good proxy for IC-Regret for a wide class of mechanisms for position auctions including GSP and VCG. Mechanisms like GSP are non truthful and therefore suffer a loss in efficiency that has been quantified in previous works [29, 38] by studying the price of anarchy of the equilibria induced by the GSP mechanism. A natural question we address is whether efficiency loss can also be quantified by measuring IC-Envy instead of analyzing the equilibria of the mechanism. In this section we positively answer this question by providing an estimation of the Social Welfare Loss based on the analysis of IC-Envy for the mechanisms characterized in Section 3.

In the following we show a direct connection for GSP between *Social Welfare Loss* and the total IC-Envy of the bidders. This connection will be established under the assumption that bidding b_i is

not worse for bidder i than bidding truthfully. We remark that this is a very mild and realistic assumption since bidding truthfully is always the simplest alternative for the advertiser. This implies that we do not require the bidders to play a min-regret strategy, neither we assume the bid vector \mathbf{b} to be a Nash Equilibrium. Finding a min-regret strategy is indeed computationally expensive and may require a large number of counterfactual experiments.

We prove the connection between total envy and social welfare loss by resorting to the notion of *semi-smoothness*, that has been introduced [29] with the goal of studying the efficiency of a position auctions off equilibrium.

Let $\text{Regret}_i(b_i, \mathbf{b}_{-i})$ be the maximum regret that the bidder i can experience with respect to the bid b_i , i.e.,

$$\text{Regret}_i(b_i, \mathbf{b}_{-i}) = \max_{b'_i \in \mathbb{R}_{\geq 0}} \{r_i(b'_i, \mathbf{b}_{-i}, b_i)\}.$$

We prove the following claim:

CLAIM 4.1. *Given a valuation profile \mathbf{v} and a bid profile \mathbf{b} , then it holds: $\sum_{i \in I} \text{Regret}_i(b_i, \mathbf{b}_{-i}) \geq \frac{1}{2} \text{SW}^{OPT} - 2\text{SW}(\mathbf{b})$.*

The next claim show that for bidder i that the regret at v_i is not smaller than the regret at b_i .

CLAIM 4.2. *$\text{Regret}_i(b_i, \mathbf{b}_{-i}) \leq \text{IC-Regret}_i(v_i, \mathbf{b}_{-i})$*

We therefore conclude:

THEOREM 4.1. *If $\text{SW}^{OPT}(\mathbf{v}) \geq 8\text{SW}(\mathbf{b})$ then*

$$\sum_i \text{IC-Envy}_{i \in I}(v_i, \mathbf{b}_{-i}) \geq \frac{1}{4} \text{SWL}(\mathbf{b})$$

PROOF. The proof follows from Claim 4.1 and Claim 4.2:

$$\begin{aligned} \sum_{i \in I} \text{IC-Envy}_i(v_i, \mathbf{b}_{-i}) &\geq \sum_{i \in I} R_i(b_i, \mathbf{b}_{-i}) \\ &\geq \frac{1}{2} \text{SW}^{OPT}(\mathbf{v}) - 2\text{SW}(\mathbf{b}) \geq \frac{1}{4} \text{SW}^{OPT}(\mathbf{v}) \\ &\geq \frac{1}{4} \text{SWL}(\mathbf{b}), \end{aligned}$$

with the second to last inequality following from $\text{SW}^{OPT}(\mathbf{v}) \geq 8\text{SW}(\mathbf{b})$. \square

The conclusion that we draw from the previous theorem is that either the social welfare loss is relatively small or the social welfare loss is away from the total envy for at most a constant multiplicative factor. We remark that the total envy is measured with one single execution of the mechanism.

5 USING ENVY AS A FEATURE

In Sections 3 we showed that for reasonably large classes of mechanisms, IC-Regret can be expressed in terms of IC-Envy and the two quantities are equal when all bids are different. In this section we move beyond linear relationships between envy and regret, and show that using envy as a feature can lead to better ML algorithms. In particular, we'll show that we can predict regret with reasonable accuracy in auction environments that are far more general than the ones discussed on Section 3.

5.1 Using Envy to Predict Regret

5.1.1 Experimental Setup. Auction Environment. We look at auctions with 5 slots, where different bidders have different monotonically decreasing discount curves over the slots (cf. the Ad Types model in [9, 11]). In this setting, not all ads can target all slots (as a consequence of the different discount curves, not by assumption), and the greedy allocation algorithm is no longer optimal. The auction mechanism that we consider use the greedy allocation (for each slot from highest to lowest, assign the slot to the unassigned ad with the highest discounted value). The goal is to show that IC-Envy is useful even outside the setting covered by the theory. We consider 2 pricing rules:

- Generalized Second Price (GSP). The discounted value of the next highest bidder, i.e. during the greedy algorithm, the next-highest value ad.
- Externality pricing.⁷ The social welfare loss of other buyers due to the presence of buyer i .

Datasets. We generate the datasets by drawing bids from a log-normal distribution⁸ and using 3 classes of bidders with geometric discount curves with parameters $\alpha_1 = 0.9$, $\alpha_2 = 0.7$, $\alpha_3 = 0.5$. For each bidder in an auction, a datapoint corresponds to the envy profile (meaning for each of the 5 slots, the unclamped, possibly negative, envy) and the label is the regret.

Baseline. We compare the performance of the ML models trained on envy, with models that were trained using the (value, price) profile (meaning for each slot, what is the discounted value, and what is the current slot price).

Implementation. We use scikit-learn [33] to train the different models. In particular we use support vector regression (SVR) with the RBF kernel; gradient-boosted regression trees (GBRT) with least-squares loss function, learning rate of 0.1, and 100 trees; and neural nets (NN) with 2 hidden layers (of 100, and 20 nodes each) and Adam solver [25].

5.1.2 Results. Figure 1 shows the training and cross-validation mean-squared error (MSE) as a function of the number of training samples for the GBDT. The MSE quickly decreases to about 0.02 after 30K iterations and remains relatively stable after that.

So using envy, we can construct a model that accurately predicts regret. To show that envy uniquely does this compared to reasonable benchmarks, we compare it against models that were trained using price and discounted value for each slot as features; the results are in Table 1. The models here are trained using 100K datapoints, the point being not to train as accurate as possible of a model, but rather to compare the performance of models trained on different features given an equal amount of data. Across all 3 models, the regressor trained on the envy feature alone does better than one that is trained on both the values and prices for slots. This remains qualitatively true for smaller training data sets as well. None of the regressors in the table are necessarily great, but the goal here is not to tweak a regressor to perform well; rather it is to




Figure 1: The training and validation MSE of the GBDT model on GSP data as a function of the number data points used to train the model.

	R^2 Price and Value	R^2 Envy
SVR	55.9%	78.4%
GBRT	46.8%	84.4%
NN	77.1%	86.2%

Table 1: Comparing using price and value as features vs. using envy as features across a range of models trained on 100K datapoints.

show that using envy as a feature gives better results across a wide variety of models without tuning the model for this particular case.

6 CONCLUSIONS

In this paper we proposed to use IC-Envy to give insight in an auction in four ways. First, we defined a class of auction mechanisms for position auctions—which includes VCG, GSP, and GFP—where IC-Envy and IC-Regret are tightly related. For this class we gave necessary and sufficient conditions for IC-Envy to upperbound IC-Regret and mild supplementary conditions under which they are exactly equal. Secondly, we upperbounded the social welfare loss in terms of IC-Envy for the same set of mechanisms. We show that the social welfare loss is *at most* $4 \cdot \text{IC-Envy}$ (under a technical condition we introduce in the section). Finally, we used bidding data from a major online publisher to show that IC-Envy can be used as a feature to learn an estimator for IC-Regret. The estimator has low MSE, and performs better than comparable estimators that are trained using other features from the auction like values and prices for different slots. In the extended version of the paper [9] we also consider the Ad Types setting, with multiple discount curves, and show that a suitable generalization of GSP (as well as VCG) continue to have $\text{IC-Envy} \geq \text{IC-Regret}$. For future work, we plan to extend our study in case of bidders with different ad types to monotonic anonymous prices such that $\text{IC-Envy} \geq \text{IC-Regret}$.

⁷If the allocation algorithm optimized social welfare, then externality pricing would be VCG pricing, and the resulting auction would have 0 envy and 0 regret. Since greedy isn't optimal, generally both envy and regret are positive.

⁸Real-world bids in online auctions typically follow a log-normal distribution, see e.g. [32].

REFERENCES

- [1] Yakov Babichenko and Aviad Rubinfeld. 2017. Communication complexity of approximate Nash equilibria. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 878–889.
- [2] Simina Brânzei, Aris Filos-Ratsikas, Peter Bro Miltersen, and Yulong Zeng. 2017. Walrasian Pricing in Multi-Unit Auctions. In *42nd International Symposium on Mathematical Foundations of Computer Science, MFCS 2017, August 21–25, 2017 - Aalborg, Denmark (LIPIcs)*, Kim G. Larsen, Hans L. Bodlaender, and Jean-François Raskin (Eds.), Vol. 83. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 80:1–80:14. <https://doi.org/10.4230/LIPIcs.MFCS.2017.80>
- [3] Jeremy Bulow and Paul Klemperer. 1994. *Auctions vs. negotiations*. Technical Report. National Bureau of Economic Research.
- [4] Ioannis Caragiannis, Christos Kaklamani, Panagiotis Kanellopoulos, and Maria Kyropoulou. 2009. On low-envy truthful allocations. In *International Conference on Algorithmic Decision Theory*. Springer, 111–119.
- [5] George Christodoulou and Annamária Kovács. 2011. A global characterization of envy-free truthful scheduling of two tasks. In *International Workshop on Internet and Network Economics*. Springer, 84–96.
- [6] Edith Cohen, Michal Feldman, Amos Fiat, Haim Kaplan, and Svetlana Olonetsky. 2010. Envy-free makespan approximation. In *Proceedings of the 11th ACM conference on Electronic commerce*. ACM, 159–166.
- [7] Edith Cohen, Michal Feldman, Amos Fiat, Haim Kaplan, and Svetlana Olonetsky. 2010. On the interplay between incentive compatibility and envy freeness. *arXiv preprint arXiv:1003.5328* (2010).
- [8] Riccardo Colini-Baldeschi, Stefano Leonardi, Piotr Sankowski, and Qiang Zhang. 2014. Revenue Maximizing Envy-Free Fixed-Price Auctions with Budgets. In *Web and Internet Economics - 10th International Conference, WINE 2014, Beijing, China, December 14–17, 2014. Proceedings (Lecture Notes in Computer Science)*, Tie-Yan Liu, Qi Qi, and Yinyu Ye (Eds.), Vol. 8877. Springer, 233–246. https://doi.org/10.1007/978-3-319-13129-0_18
- [9] Riccardo Colini-Baldeschi, Stefano Leonardi, Okke Schrijvers, and Eric Sodomka. 2019. Envy, Regret, and Social Welfare Loss. *arXiv, preprint arXiv:1907.07721* (2019).
- [10] Riccardo Colini-Baldeschi, Stefano Leonardi, and Qiang Zhang. 2016. Revenue Maximizing Envy-Free Pricing in Matching Markets with Budgets. In *Web and Internet Economics - 12th International Conference, WINE 2016, Montreal, Canada, December 11–14, 2016. Proceedings (Lecture Notes in Computer Science)*, Yang Cai and Adrian Vetta (Eds.), Vol. 10123. Springer, 207–220. https://doi.org/10.1007/978-3-662-54110-4_15
- [11] Riccardo Colini-Baldeschi, Julian Mestre, Okke Schrijvers, and Christopher A. Wilkens. 2019. The Ad Types Problem. *arXiv, preprint arXiv:1907.04400* (2019).
- [12] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. 2009. The complexity of computing a Nash equilibrium. *SIAM J. Comput.* 39, 1 (2009), 195–259.
- [13] Constantinos Daskalakis and Vasilis Syrgkanis. 2016. Learning in auctions: Regret is hard, envy is easy. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*. IEEE, 219–228.
- [14] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review* 97, 1 (2007), 242–259.
- [15] Michal Feldman, Amos Fiat, Stefano Leonardi, and Piotr Sankowski. 2012. Revenue maximizing envy-free multi-unit auctions with budgets. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 532–549.
- [16] Michal Feldman and John Lai. 2012. Mechanisms and impossibilities for truthful, envy-free allocations. In *Algorithmic Game Theory*. Springer, 120–131.
- [17] Zhe Feng, Okke Schrijvers, and Eric Sodomka. 2019. Online Learning for Measuring Incentive Compatibility in Ad Auctions. In *Proceedings of the 2019 World Wide Web Conference*.
- [18] Amos Fiat and Amiram Wingarten. 2009. Envy, multi envy, and revenue maximization. In *International Workshop on Internet and Network Economics*. Springer, 498–504.
- [19] Lisa Fleischer and Zhenghui Wang. 2011. Lower bound for envy-free and truthful makespan approximation on related machines. In *International Symposium on Algorithmic Game Theory*. Springer, 166–177.
- [20] D Foley. 1967. Resource allocation and the public sector. *Yale Economic Essays* 7 (1967), 45–98.
- [21] Yoav Freund and Robert E Schapire. 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences* 55, 1 (1997), 119–139.
- [22] Andrew V Goldberg and Jason D Hartline. 2003. Envy-free auctions for digital goods. In *Proceedings of the 4th ACM conference on Electronic commerce*. ACM, 29–35.
- [23] Venkatesan Guruswami, Jason D Hartline, Anna R Karlin, David Kempe, Claire Kenyon, and Frank McSherry. 2005. On profit-maximizing envy-free pricing. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 1164–1173.
- [24] Jason Hartline and Qiqi Yan. 2011. Envy, truth, and profit. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, 243–252.
- [25] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [26] Sébastien Lahaie, Andrés Muñoz Medina, Balasubramanian Sivan, and Sergei Vassilvitskii. 2018. Testing Incentive Compatibility in Display Ad Auctions. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 1419–1428.
- [27] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM conference on Electronic commerce*. ACM, 125–131.
- [28] Benjamin Lubin and David C. Parkes. 2009. Quantifying the Strategyproofness of Mechanisms via Metrics on Payoff Distributions. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (Montreal, Quebec, Canada) (UAI '09)*. AUAI Press, Arlington, Virginia, United States, 349–358. <http://dl.acm.org/citation.cfm?id=1795114.1795155>
- [29] Brendan Lucier and Renato Paes Leme. 2011. GSP auctions with correlated types. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, 71–80.
- [30] Ahuva Mu'alem and Michael Schapira. 2018. Setting lower bounds on truthfulness. *Games and Economic Behavior* 110 (2018), 174–193.
- [31] Roger Myerson. 1981. Optimal Auctions. (1981).
- [32] Michael Ostrovsky and Michael Schwarz. 2011. Reserve prices in internet advertising auctions: a field experiment. *EC* 11 (2011), 59–60.
- [33] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [34] Aviad Rubinfeld. 2018. Inapproximability of Nash equilibrium. *SIAM J. Comput.* 47, 3 (2018), 917–959.
- [35] Sarah Sluis. 2017. Big Changes Coming To Auctions, As Exchanges Roll The Dice On First-Price. <https://adexchanger.com/platforms/big-changes-coming-auctions-exchanges-roll-dice-first-price/>. Last accessed: 2019-05-23.
- [36] Sarah Sluis. 2019. Google Switches To First-Price Auction. <https://adexchanger.com/online-advertising/google-switches-to-first-price-auction/>. Last accessed: 2019-05-23.
- [37] Bo Tang and Jinshan Zhang. 2015. Envy-Free Sponsored Search Auctions with Budgets. In *IJCAI*. 653–659.
- [38] Éva Tardos and Renato Paes Leme. 2010. Pure and Bayes-Nash Price of Anarchy for Generalized Second Price Auction. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*. IEEE, 735–744.
- [39] Hal R Varian. 1973. Equity, envy, and efficiency. (1973).
- [40] Hal R Varian. 2007. Position auctions. *international Journal of industrial Organization* 25, 6 (2007), 1163–1178.