Mechanism Design via Consensus Estimates, Cross Checking, and Profit Extraction

BACH Q. HA and JASON D. HARTLINE, Northwestern University

There is only one technique for *prior-free optimal mechanism design* that generalizes beyond the structurally benevolent setting of digital goods. This technique uses random sampling to estimate the distribution of agent values and then employs the *Bayesian optimal mechanism* for this estimated distribution on the remaining players. Though quite general, even for digital goods, this random sampling auction has a complicated analysis and is known to be suboptimal. To overcome these issues we generalize the consensus and profit-extraction techniques from Goldberg and Hartline [2003] to structurally rich environments that include, e.g., single-minded combinatorial auctions.

Categories and Subject Descriptors: F.2.0 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity—*General*

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: consensus, envy-freedom, prior-free mechanism design, profit extraction

ACM Reference Format:
ACM 1, 1, Article 1 (December 2012), 14 pages.
DOI = 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

1. INTRODUCTION

The classical economic theory of mechanism design is Bayesian: it is assumed that the preferences of the agents are drawn at random from a known probability distribution and the designer aims to optimize her objective in expectation over this randomization. This leads to mechanisms that are tailored to the distributional setting. In contrast, prior-free mechanism design looks at mechanisms that perform well without knowledge or assumptions on agent preferences. While these mechanisms do not perform as well as ones tailored to the distribution, in many environments they provide good approximations and are more robust.

Prior-free mechanisms are evaluated in terms of the approximation factor they obtain relative to a given performance benchmark. For instance, prior-free auctions for digital goods, i.e., where there are n agents and n identical units of an item, have been shown to obtain a constant approximation to the revenue of the best posted price, ex post. Meaning: there is a constant β such that the expected revenue of the auction is at least a β-fraction of the revenue from posting the best, in hindsight, price. Of course, the Bayesian optimal auction for an i.i.d. distribution on agent preferences is a posted price, so approximation relative to the best posted price provides a very strong robustness guarantee.

In large markets (under reasonable assumptions) it is possible to design prior-free mechanisms that obtain arbitrarily close approximations to the Bayesian opti-

This work is supported in part by the National Science Foundation CAREER Award CCF-0846113.

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

mal mechanism. This follows by sampling theory and the law of large numbers; explicit constructions have been given by Segal [2003], Baliga and Vohra [2003], and Balcan et al. [2008]. Furthermore, Jackson and Kremer [2007] show that agents become "price takers" in the limit and, therefore, mechanism design is easy. Our focus, therefore, is on small, a.k.a., *thin*, markets where we look for constant approximations (arbitrarily close approximations are impossible).

In thin markets the Bayesian prior assumption, that the designer has modeled the agents' preferences as draws from a continuous distribution, breaks down. Market analysis cannot hope to give such an accurate prior; and historical data must only be used with care due to Coase-conjecture-type phenomena [Coase 1972]. Furthermore, many such thin markets are ones where it is not possible to reoptimize the mechanism for each scenario. In these markets a prior-free mechanism, i.e., one that works pretty well under any market conditions, may be preferred.

The simplest environment in which to explore mechanism design is that of selling a *digital good*, i.e., where the seller has no constraint over the subsets of agents that can be served simultaneously. Recent contributions to the literature on prior-free mechanism design have focused on extending results for digital good environments to ones that are more structurally rich. From least-general to most-general, these include *multi-unit environments*, where there is a given number k of units available for sale (i.e., any subset of the agents of size at most k can be served); *matching environments*, where feasible sets correspond to one side of a bipartite matching; and *downward-closed environments*, where the only constraint on feasible sets is that any subset of a feasible set is feasible.

The only prior-free mechanisms known to give good approximations for general downward-closed environments are variants of the random sampling auction. This auction first gathers distributional information from a random sample of the agents and then runs the Bayesian optimal auction for the empirical distribution on the remaining agents. Tight analysis of the random sampling auction is difficult, upper bounds on its approximation factor are 4.68, 25, 50, and 2560 for digital good [Alaei et al. 2009], multi-unit [Devanur and Hartline 2009], matching [Hartline and Yan 2011], and downward-closed environments [Hartline and Yan 2011], respectively. Other mechanism design techniques give 3.25 and 6.5-approximations for digital-good [Hartline and McGrew 2005] and multi-unit environments [Hartline and Yan 2011], respectively, and notably the 3.25 approximation for digital goods surpasses the lower-bound of 4 which is known for the random sampling auction. These limitations suggest the need to consider other techniques for obtaining good approximations for general downward-closed environments.

To obtain good approximation mechanisms for general downward-closed environments, we generalize the digital-good auction technique of consensus estimates from Goldberg and Hartline [2003]. The two main ingredients of this approach are a *profit extraction* mechanism and a *consensus* function. Given a target profit, the profit extraction mechanism should obtain the target if the target is less than the optimal revenue possible. If we had a good estimate of the revenue, we could then obtain a good revenue with the profit extractor. The consensus function is used to get an estimate of the revenue from the reports of the agents in a way that is non-manipulable. In particular, for each agent we can calculate the optimal profit from the other agents, plug this profit into the consensus function, and with high probability the estimated profit produced will be the same for all agents. We can then simulate the profit extraction mechanism for each agent with their consensus estimate. If the estimates agree, the result of this simulation is the agreed-upon profit, otherwise, it is at least zero.

There are two main challenges to extending this approach for general downward-closed environments. The first challenge is in designing a profit-extraction mechanism for these environments. Our profit extraction mechanism will be parameterized by a *revenue curve*, the revenue as a function of number of winners (without taking into account any feasibility constraints). Given a target revenue curve that is below the actual revenue curve, our mechanism obtains revenue comparable to that which would be obtained by the optimal mechanism on the input that corresponds to the target revenue curve. The second challenge is in ensuring infeasible outcomes are not produced in the case that the estimates do not have a consensus. Note that for digital good environments, there is no feasibility constraint that could be violated when the estimates do not reach consensus. The same is not so for general downward-closed environments. A parameterized mechanism (such as a profit extractor) is of course required to always produced feasible outcomes. However, if we determine the outcome for each agent by simulating the parameterized mechanism with different parameters for different agents, the combined outcome may not be feasible. To address this potential inconsistency we give a *cross checking* approach for identifying a subset of agents for which consensus is achieved.

Our mechanism is a 30.4-approximation in general downward-closed environments.

1.1. Related Work

The most important prior-free mechanism is the second-price, a.k.a. Vickrey, auction and its generalization to complex environments which is known as the Vickrey-Clarke-Groves (VCG) mechanism [Vickrey 1961; Clarke 1971; Groves 1973]. This mechanism maximizes the social surplus, it is dominant strategy incentive compatible, and it relies on no distributional assumptions on agent preferences. Unfortunately, the objective of social surplus is singular in this respect; for other objectives such as revenue the optimal mechanism for any given prior distribution depends on the distribution. Myerson [1981] solved for this *Bayesian optimal* mechanism in single-item environments with independently distributed agent preferences; Bulow and Roberts [1989] provided a natural economic interpretation of Myerson's approach; and the approach generalizes beyond single-item auctions to any abstract environment where the agents have single-dimensional preferences.

The last decade has seen a resurgence of interest in optimal mechanism design; however, with the objective of describing good mechanisms for revenue that are more like VCG in assumptions: they are dominant strategy incentive compatible and have good revenue regardless of the prior distribution or even adversarially. This direction was initiated by Goldberg et al. [2001; 2006] who studied a digital good environment. In this digital good environment there are n agents and the seller has n identical units of an item. Goldberg et al. gave several digital good auctions which guaranteed good (i.e., constant) approximation to the revenue of the optimal posted pricing (in hindsight). The best approximation factor for digital good auctions of 3.25 was obtained by an auction due to Hartline and McGrew [2005].

An important issue in the design of prior-free mechanism design for objectives like revenue where the Bayesian optimal mechanism is not prior-free, is how mechanisms should be judged. For instance, the digital good auctions above were judged relative to optimal posted pricing. Hartline and Roughgarden [2008] observed that optimal posted pricing is a good benchmark for comparison of digital good auctions because the Bayesian optimal mechanism for digital goods under an i.i.d. distribution is a posted price. Therefore, by approximating the optimal posted pricing revenue, a digital good auction would simultaneously approximate the Bayesian optimal auction for any i.i.d. distribution. Hartline and Yan [2011] showed that a good generalization of this benchmark to structurally rich environments is the optimal envy-free revenue (See Section 2). Their justification of the envy-free benchmark is reinforced by prior work of

Jackson and Kremer [2007] who observe that envy freedom and incentive compatibility are identical in the limit as the market grows.

There are two relaxations of the above approach prior-free mechanism design that are noteworthy. One can relax the assumption that the agent preferences could be adversarial and instead assume they are drawn from a distribution. Dhangwatnotai, Roughgarden, and Yan [2010] considered the design of dominant strategy incentive compatible mechanisms in such an environment. They consider abstract service provision in downward-closed environments with the assumption that the values of the agents are distributed according to a unknown distribution that satisfies a standard *monotone hazard rate* condition. Under this assumption, they give (essentially) a 4-approximation mechanism. Caillaud and Robert [2005] further relax the design considerations by allowing for Bayes-Nash implementation (for single-item auctions). Their auction relies on the agents' knowledge of the distribution and obtains the optimal revenue. Furthermore, standard approaches in non-parametric Bayes-Nash implementation theory suggest that the Caillaud-Robert result can approximately extended to general environments, see e.g., Jackson [2001]. The mechanisms suggested by this last approach requires agents to report their knowledge about the distribution of other agents' preferences.

Relative to the above discussion, the present paper considers the original problem of prior-free mechanism design: a single dominant strategy incentive compatible mechanism is sought to approximate the envy-free benchmark of Hartline and Yan [2011]. This work can be viewed as generalizing that of Dhangwatnotai et al. [2010] by relaxing the monotone hazard rate assumption or improving on the mechanisms of Hartline and Yan [2011]. The mechanism design techniques we develop are extensions of profit-extraction and consensus techniques for digital good environments that come from Goldberg and Hartline [2003; 2005].

1.2. Organization

In Section 2 we will formally describe our auction environment, design and analysis framework, and review the consensus technique. In Section 3 we describe our cross-checking approach as it applies to obtaining a consistent consensus estimate. In Section 4 we describe a mechanism for extracting the profit suggested by a given target revenue curve. In Section 5 we describe an approach for obtaining a consensus estimate on revenue curves. Finally, in Section 6 we combine the three parts to give a good mechanism and we analyze its performance.

2. PRELIMINARIES

Here we describe the abstract setting in which we consider mechanism design and the structural tools that we will be using to design and analyze mechanisms.

2.1. Incentives

Let $[n] = \{1, \ldots, n\}$ be a set of $n \geq 2$ bidders. Each bidder $i \in [n]$ has a private *valuation* v_i for receiving some abstract service. A bidder i, upon reporting his valuation, will be served with a probability x_i and charged a payment of p_i. We denote the *valuation profile*, *allocation vector*, and *payment vector* by $\mathbf{v} = (v_1, \ldots, v_n)$, $\mathbf{x} = (x_1, \ldots, x_n)$, and $\mathbf{p} = (p_1, \ldots, p_n)$ respectively.

We assume the standard *risk-neutral quasi-linear* utility model, i.e., an agent i wishes to maximize his expected utility which is given by $u_i = v_i x_i - p_i$. We will focus solely on dominant strategy incentive compatible (IC) mechanisms; which means for any agent, reporting his true valuation would be a dominant strategy; and we assume that agents follow this dominant strategy. We view a mechanism as a function from reports to allocation and payments and denote these functions by $\mathbf{x}(\cdot)$ and $\mathbf{p}(\cdot)$.

Incentive compatibility is defined as, for all i, $\mathbf{v}$, and z:

$$v_i x_i(\mathbf{v}) - p_i(\mathbf{v}) \geq v_i x_i(z, \mathbf{v}_{-i}) - p_i(z, \mathbf{v}_{-i}) \tag{1}$$

where $(z, \mathbf{v}_{-i})$ is the valuation profile $\mathbf{v}$ with v_i replaced with z. A mechanism is incentive compatible if and only if for all i [Myerson 1981]:

(1) $x_i(\mathbf{v})$ is monotone non-decreasing in v_i, and
(2) $p_i(\mathbf{v})$ satisfies the payment identity:

$$p_i(\mathbf{v}) = v_i x_i(\mathbf{v}) - \int_0^{v_i} x_i(z, \mathbf{v}_{-i})\, dz. \tag{2}$$

When we give a mechanism we will describe only the *allocation rule* and infer the *payment rule* from the payment identity.[1] For allocation rule $\mathbf{x}(\cdot)$ we will denote the incentive compatible payment for agent i by $\mathrm{IC}_i^{\mathbf{x}}(\mathbf{v})$ and the total revenue by $\mathrm{IC}^{\mathbf{x}}(\mathbf{v})$.

2.2. Feasibility

The designer faces a *feasibility constraint* which describes the subsets of the agents that can be served simultaneously. We assume that this feasibility constraint is downward closed, i.e., any subset of a feasible set is feasible. As described in the introduction, many common environments for mechanism design are downward closed. We allow the feasibility constraint to be probabilistic, i.e., given by a convex combination of downward-closed set systems.[2] For the purpose of calculating revenue the allocation $\mathbf{x}$ and payments $\mathbf{p}$ are taken in expectation over the randomization in the mechanism and the set system.

Symmetry will play an important role in our performance analysis. Given an asymmetric set system we can always make it symmetric by randomly permuting the identities of the agents. This assumption is akin to standard assumptions in the performance analysis of the secretary problem and in settings where one might consider the agents to be a priori indistinguishable. For instance, symmetry is without loss in Bayesian optimal mechanism design when values drawn from an i.i.d. distribution. The resulting feasibility constraint we refer to as a *downward-closed permutation environment*.

2.3. Algorithms

Our mechanisms will be based on an algorithm for weighted optimization.[3] Given weights $\mathbf{w} = (w_1, \ldots, w_n)$, for each agent, this algorithm selects a feasible set to optimize the sum of the selected weights (for the realized set system) with ties broken randomly. As suggested above, x_i will denote the probability (over randomization in the set system and random tie-breaking) that agent i is selected. Clearly $\mathbf{x}$ maximizes $\sum_i w_i x_i$ subject to feasibility and so we will refer to $\mathbf{x}$ as the *maximizer* for weights $\mathbf{w}$.

In downward-closed permutation environments it is without loss to index the agents in decreasing order of value, i.e., $v_i \geq v_{i+1}$. For weights $\mathbf{w}$ sorted in decreasing order, the maximizer $\mathbf{x}$ for weights $\mathbf{w}$ is monotone, i.e., $x_i \geq x_{i+1}$ for all i. Of course, the maximizer $\mathbf{x}$ for $\mathbf{w}$ is not generally the maximizer for $\mathbf{v}$.

2.4. Performance Benchmarks

We adopt the framework from Hartline and Yan [2011] wherein the revenue of the designed prior-free mechanism is compared to the *envy-free revenue* benchmark. We will

[1] Of course such a payment can be easily calculated, e.g., with techniques from Archer et al. [2003].
[2] For instance, the position auction environment which models advertising on Internet search engines is structurally equivalent to a convex combination of multi-unit auctions. See, e.g., Hartline and Yan [2011].
[3] We make no comment on this algorithm's computational tractability.

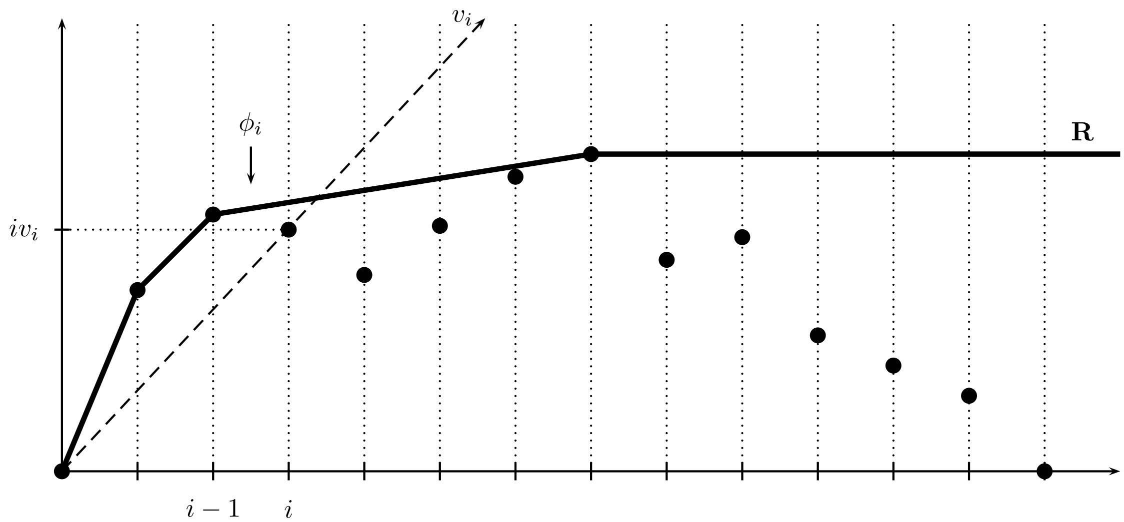

Fig. 1. Agents are indexed in decreasing order by value. The point set $\{(i, iv_i) : i \in [n]\}$ is depicted. The revenue curve is the smallest non-decreasing concave function that upper-bounds this point set and the origin. The value of agent i can be represented by a line from the origin with slope v_i; the virtual value of agent i is the left-slope of the revenue curve at i.

denote the maximum envy-free revenue by $\mathrm{EFO}(\mathbf{v})$. For technical reasons the envy-free benchmark is defined to be $\mathrm{EFO}^{(2)}(\mathbf{v}) = \mathrm{EFO}(v_2, v_2, v_3, \ldots, v_n)$.[4] In our downward-closed permutation environment, the goal of such a design and analysis framework is then to give a mechanism that, in expectation over the random permutation of agent identities, obtains a revenue that is a good approximation, in worst case over agent valuation profiles $\mathbf{v}$, to the benchmark $\mathrm{EFO}^{(2)}(\mathbf{v})$. Hartline and Yan [2011] give formal Bayesian justification for this benchmark by showing that a prior-free mechanism that approximates it simultaneously approximates the Bayesian optimal mechanisms for most i.i.d. distributions.

We will only consider envy freedom in permutation environments were $\mathbf{x}$ is the allocation in expectation over the permutation. An outcome $(\mathbf{x}, \mathbf{p})$ is envy free if no agent wants to swap outcome with another agent, i.e., for all $i, j \in [n]$,

$$v_i x_i - p_i \geq v_i x_j - p_j. \quad (3)$$

An outcome is envy free if and only if [Hartline and Yan 2011]:

(1) $\mathbf{x}$ is monotone non-decreasing (i.e., $x_i \geq x_{i+1}$), and
(2) $\mathbf{p}$ (for maximum payments given $\mathbf{x}$) satisfies the payment identity (for all i):

$$p_i = \sum\nolimits_{j \geq i}^{n} v_j \cdot (x_j - x_{j+1}). \quad (4)$$

For allocation $\mathbf{x}$, we denote the envy-free payments from the payment identity as $\mathrm{EF}_i^{\mathbf{x}}(\mathbf{v})$ and the total envy-free revenue by $\mathrm{EF}^{\mathbf{x}}(\mathbf{v})$. Notice that unlike incentive compatibility, which is defined on allocation and payment rules, envy freedom is defined point wise on allocations and payments.

The envy-free revenue can be understood structurally in terms of the *revenue curve* and *virtual values*. The revenue curve $\mathbf{R}(\mathbf{v})$ for $\mathbf{v}$ is a vector that describes the optimal

[4]No prior-free mechanism can approximate $\mathrm{EFO}(\mathbf{v})$ in the case where there is one agent with an extremely large value, see Goldberg et al. [2006] for discussion.

revenue (when feasibility constraints are ignored) indexed by the number of agents served. The ith coordinate of the revenue curve, $R_i(\mathbf{v})$, can be calculated by evaluating at i the smallest concave non-decreasing function that contains the point set $\{(i, iv_i) : i \in [n]\}$ and the origin. The virtual value at i is the left-slope of this function, i.e., $\phi_i(\mathbf{v}) = R_i(\mathbf{v}) - R_{i-1}(\mathbf{v})$. Where there is no ambiguity with respect to the valuation profile we will denote $\mathbf{R}(\mathbf{v})$ and $\boldsymbol{\phi}(\mathbf{v})$ as $\mathbf{R}$ and $\boldsymbol{\phi}$, respectively. See Figure 1.

LEMMA 2.1. *[Hartline and Yan 2011] The envy-free revenue of monotone allocation* $\mathbf{x}$ *satisfies*

$$\mathrm{EF}^{\mathbf{x}}(\mathbf{v}) \leq \sum\nolimits_{i=1}^{n} \phi_i(\mathbf{v}) \cdot x_i = \sum\nolimits_{i=1}^{n} R_i(\mathbf{v}) \cdot (x_i - x_{i+1}) \tag{5}$$

with equality when $x_i = x_{i+1}$ *whenever* $\phi_i(\mathbf{v}) = \phi_{i+1}(\mathbf{v})$.

The optimal envy-free revenue, $\mathrm{EFO}(\mathbf{v})$, can be found from Lemma 2.1, in particular, from the maximizer for *virtual surplus*, i.e., $\sum_i \phi_i(\mathbf{v}) \cdot x_i$, with random tie breaking. Random tie breaking results in an allocation $\mathbf{x}$ that satisfies $x_i = x_{i+1}$ whenever $\phi_i(\mathbf{v}) = \phi_{i+1}(\mathbf{v})$.

Notice that for the same allocation rule $\mathbf{x}(\cdot)$, the envy-free payments (4) and incentive compatible payments (2) are distinct, i.e., $\mathrm{IC}_i^{\mathbf{x}}(\mathbf{v}) \neq \mathrm{EF}_i^{\mathbf{x}}(\mathbf{v})$.

2.5. Consensus estimates.

A central ingredient in our approach is the technique of *consensus estimates* that was introduced by Goldberg and Hartline [2003]. A *consensus function* maps shared randomness and a statistic to an estimate of the statistic. The objective of such a consensus function is that, when applied individually to each of a set of statistics that are within some bounded range, with high probability (in the shared randomness) the estimates will coincide, i.e., there will be a *consensus* among the *estimates*.

Definition 2.2. For implicit parameter $c > 1$ and shared randomness $\sigma \sim U[0,1]$, the *consensus function* on statistic s is,

$$\mathrm{Consens}(\sigma, s) = \lfloor s \rfloor_{\{c^{\sigma+d} \,:\, d \in \mathbb{Z}\}},$$

where $\lfloor s \rfloor_S$ denotes s rounded down to the nearest element of S.

LEMMA 2.3. *[Goldberg and Hartline 2003] For* $c \geq \beta$*, The probability (over randomization of* σ*) that the consensus function is constant on interval* $[s/\beta, s]$ *is* $1 - \log_c \beta$.

3. CROSS-CHECKING

Given some statistic on valuation profiles $s(\cdot)$, we will be using the consensus function (Definition 2.2) to get an estimate of this statistic, e.g., by calculating $\mathrm{Consens}(\sigma, s(\mathbf{v}_{-i}))$ for each i where $\mathbf{v}_{-i}$ is the valuation profile $\mathbf{v}$ without coordinate i. Notice that if we had some mechanism $\mathcal{M}_s$ that was parameterized by statistic s then, if there is consensus, simulations of $\mathcal{M}_{\mathrm{Consens}(\sigma, s(\mathbf{v}_{-i}))}$ to determine the allocation x_i and payment p_i for all agents i are internally consistent. I.e., the outcome produced by $\mathcal{M}_s$ is feasible for any s; therefore, so is the combined outcome. Unfortunately, when consensus is not achieved then these simulations may not be consistent.

In this section we give a method of cross-checking to ensure that consistent estimates of the statistic for some subset of the agents. For environments with downward closed feasibility constraints, such a method can be used in mechanism design as agents outside this consistent subset can be rejected.

Definition 3.1. For shared randomness σ, statistic s, consensus function $\mathrm{Consens}$, and valuation profile $\mathbf{v}$ calculate the following:

(1) For all pairs $i \neq j \in [n]$, calculate the estimate $\tilde{s}_{i,j} = \mathrm{Consens}(\sigma, s(\mathbf{v}_{-i,j}))$ where $\mathbf{v}_{-i,j}$ is the valuation profile $\mathbf{v}$ without coordinates i and j.
(2) I is the set of agents i that have consensus on $\tilde{s}_{i,j}$ for all j; or $\emptyset$ if no such i exists.
(3) $\tilde{s}$ is the estimate $\tilde{s}_{i,j}$ of any $i \in I$ and any j (they are all the same).

The *cross-checked consensus function* is defined as

$$\mathrm{CrossConsens}(\sigma, s, \mathbf{v}) = (\tilde{s}, I).$$

Cross-checked consensus estimates are non-manipulable in a strong sense. Whether or not an agent i is in I is not a function of that agents value. Furthermore, the final estimate is not a function of the report of any agent $i \in I$. This implies that mechanisms which take the following form are incentive compatible.

Definition 3.2 (*cross-checked consensus estimate composition*). Given an incentive compatible mechanism $\mathcal{M}_s$ that is parameterized by some statistic s and a consensus function Consens for the statistic, compose them as follows:

(1) Calculate cross-checked consensus estimate $(\tilde{s}, I) = \mathrm{CrossConsens}(\sigma, s, \mathbf{v})$.
(2) Simulate incentive compatible mechanism $\mathcal{M}_{\tilde{s}}$ on $\mathbf{v}$.
(3) For agents $i \in I$ output result of simulation, reject all others.

THEOREM 3.3. *Mechanisms produced by the cross-checked consensus estimate composition are incentive compatible.*

PROOF. Let $(\tilde{s}, I) = \mathrm{CrossConsens}(\sigma, s, \mathbf{v})$. Agent $i \notin I$ has no report he can make to win; therefore, he has no incentive not to report truthfully. Agent $i \in I$ has no report he can make to change the value of $\tilde{s}$; therefore, the incentive compatibility of $\mathcal{M}_s$ for any fixed value of s implies that he has no incentive not to report truthfully. □

4. PROFIT EXTRACTION MECHANISM

In this section we will show how to design a mechanism with good revenue that is parameterized by an approximation of the revenue curve. Such a mechanism is termed a *profit extractor*. Given a target revenue curve that is upper-bounded by our actual revenue curve, this mechanism will obtain at least the optimal envy-free revenue for the target revenue curve. The target revenue curve will be provided to the mechanism in the form of the valuation profile $\tilde{\mathbf{v}}$ that generates it. We will denote by $\tilde{\mathbf{R}}$ and $\tilde{\boldsymbol{\phi}}$ the revenue curve and virtual values for $\tilde{\mathbf{v}}$.

Definition 4.1 (*profit extractor,* $\mathrm{PE}_{\tilde{\mathbf{v}}}$). Parameterized by non-increasing valuation vector $\tilde{\mathbf{v}}$:

(1) Sort the bids in a non-increasing order, break ties arbitrary. If $\tilde{v}_i > v_i$ for some i, reject everyone.
(2) Assign weights $\tilde{\boldsymbol{\phi}}$ to agents in the same order as their values.
(3) Serve the set of agents to maximize the sum of their assigned weights.

We will show that the IC revenue obtained from the profit extractor for $\tilde{\mathbf{v}}$ on $\mathbf{v}$ is higher than the optimal envy-free revenue for $\tilde{\mathbf{v}}$. Furthermore, for appropriately chosen $\tilde{\mathbf{v}}$, this revenue approximates the optimal envy-free revenue for $\mathbf{v}$.

THEOREM 4.2. *For any downward-closed environment and any* $\tilde{\mathbf{v}}$*, the profit extractor for* $\tilde{\mathbf{v}}$ *is dominant strategy incentive compatible.*

PROOF. Fix the randomization in the set system, fix the values $\mathbf{v}_{-i}$ of all agents except agent i, and suppose i is served with value v_i. We will show that agent i continues to be served when he reports a higher value z between the jth and $j-1$st highest values (for $j < i$); consequently the allocation rule is a step function which is monotone and sufficient for incentive compatibility.

First, since i was served with value v_i, Step 1 did not bind and will continue to not bind when i increases his bid to z.

Second, consider what happens in the weighted maximization when i out bids j. He assumes the weight w_j of j (positive change of weight: $d_i = w_j - w_i$), all agents $i' \in \{j, \ldots, i-1\}$ assume the weight of the agent below them (negative change of weight: $d_{i'} = w_{i'+1} - w_{i'}$), and all other agents $i' \notin \{j, \ldots, i\}$ keep their old weights (zero change of weight: $d_{i'} = 0$). Of course the total weight change is conserved and $\sum_{i' \in [n]} d_{i'} = 0$. Therefore, the total weight change of any subset of agents that contains i is non-negative and the total weight change of any subset of agents that does not contain i is non-positive. Since agent i was served with value v_i, he was in the weight-maximizing set with report v_i, when i out bids j he continues to be in the weight-maximizing set (though this set may change). Therefore, agent i continues to be served when bidding $z > v_i$. □

THEOREM 4.3. *For any $\tilde{\mathbf{v}} \leq \mathbf{v}$ and any downward-closed permutation environment, the revenue of the profit extractor for $\tilde{\mathbf{v}}$ on $\mathbf{v}$ is at least the envy-free optimal revenue for $\tilde{\mathbf{v}}$. Moreover, the inequality holds on each agent's payment, i.e.,* $\mathrm{IC}_i^{\mathrm{PE}_{\tilde{\mathbf{v}}}}(\mathbf{v}) \geq \mathrm{EFO}_i(\tilde{\mathbf{v}})$.

PROOF. We show the second condition of the theorem for any agent i (the first condition follows). Let $\tilde{\mathbf{x}}$ be the allocation for $\mathrm{EFO}(\tilde{\mathbf{v}})$ (in expectation over the permutation). This is the same allocation as used by $\mathrm{PE}_{\tilde{\mathbf{v}}}$ unless $\mathbf{v} \geq \tilde{\mathbf{v}}$ fails to hold.

First, notice that the IC payments, from equation (2), and EF payments, from equation (4), correspond to the area in the region bounded by $x \leq \tilde{x}_i$ (above), $0 \leq v$ (left), and the "allocation rule" (bottom right). For IC payments, this allocation rule is the probability that the agent is served for any possible misreport z. For EF payments, this "allocation rule" is the smallest monotone function that upper-bounds the point set $\{(\tilde{v}_i, \tilde{x}_i) : i \in [n]\}$. To prove the lemma we need only show that the IC allocation rule gives a weaker bound than the EF "allocation rule."

For any $j > i$, the EF allocation rule drops from $\tilde{x}_j$ to $\tilde{x}_{j+1}$ at $\tilde{v}_j$. We claim that the IC allocation rule makes the same drop but at a value that is at least $\tilde{v}_j$. To see this claim, consider the minimum bid z that agent i can make to secure allocation probability at least $\tilde{x}_j$. If $v_{j+1} \geq \tilde{v}_j$ then obtaining slot j requires bidding z at least $v_{j+1} \geq \tilde{v}_j$. Otherwise, $\tilde{v}_j > v_{j+1}$ and preventing Step 1 from rejecting everyone requires bidding z at least $\tilde{v}_j$. In either case, i must bid at least $\tilde{v}_j$ to get allocation probability at least $\tilde{x}_j$. □

LEMMA 4.4. *For any $\tilde{\mathbf{v}}$ and $\mathbf{v}$ with $\tilde{\mathbf{R}} \geq \frac{1}{\beta}\mathbf{R}$, the envy-free optimal revenue for $\tilde{\mathbf{v}}$ is a β-approximation to that from $\mathbf{v}$, i.e.,* $\mathrm{EFO}(\tilde{\mathbf{v}}) \geq \frac{1}{\beta} \cdot \mathrm{EFO}(\mathbf{v})$.

PROOF. Let $\mathbf{x}$ and $\tilde{\mathbf{x}}$ be the allocations $\mathrm{EFO}(\mathbf{v})$ and $\mathrm{EFO}(\tilde{\mathbf{v}})$, respectively.

$$\mathrm{EFO}(\tilde{\mathbf{v}}) = \sum\nolimits_i \tilde{R}_i \cdot (\tilde{x}_i - \tilde{x}_{i+1}) \geq \sum\nolimits_i \tilde{R}_i \cdot (x_i - x_{i+1}) \geq \tfrac{1}{\beta} \sum\nolimits_i R_i \cdot (x_i - x_{i+1}) = \tfrac{1}{\beta}\,\mathrm{EFO}(\mathbf{v}).$$

The first inequality follows from the optimality of $\tilde{\mathbf{x}}$ for $\tilde{\mathbf{v}}$ and Lemma 2.1. The second inequality follows from monotonicity of $\mathbf{x}$ and the assumption that $\forall i,\ \tilde{R}_i \geq \frac{1}{\beta} \cdot R_i$. □

Combining these lemmas, we see that with the right $\tilde{\mathbf{v}}$, $\mathrm{PE}_{\tilde{\mathbf{v}}}(\mathbf{v})$ can approximate the optimal envy free revenue on $\mathbf{v}$.

THEOREM 4.5. *For any $\tilde{\mathbf{v}} \leq \mathbf{v}$ with $\tilde{\mathbf{R}} \geq \frac{1}{\beta}\mathbf{R}$ and any downward-closed permutation environment, the profit extractor for $\tilde{\mathbf{v}}$ on $\mathbf{v}$ is a β-approximation to the optimal envy-free revenue for $\mathbf{v}$, i.e., $\mathrm{IC}^{\mathrm{PE}_{\tilde{\mathbf{v}}}}(\mathbf{v}) \geq \frac{1}{\beta} \cdot \mathrm{EFO}(\mathbf{v})$.*

5. CONSENSUS ESTIMATES OF REVENUE CURVES

Our objective now is to get a consensus estimate of the revenue curve. We will express the estimated revenue curve, $\tilde{\mathbf{R}}$, in terms of the estimated valuation profile, $\tilde{\mathbf{v}}$, that generates it.

In addition to the implicit parameter $c > 1$ in the definition of consensus (Definition 2.2) we will also use implicit parameter $\alpha > 1$ and a minimum required support $m \in \mathbb{Z}_+$. A statistic we will be interested in getting consensus on is the number of agents with value at least α^j for any given j. We will use $n_j(\mathbf{v})$ to denote this statistic. As per our notation in the previous section, we will denote $\tilde{n}_j(\sigma, \mathbf{v}) = \lceil \mathrm{Consens}(\sigma, n_j(\mathbf{v})) \rceil$ where we round the estimate up to the nearest integer because it is an integer statistic. Estimates that do not have the minimum required support of $\tilde{n}_j(\sigma, \mathbf{v}) \geq m$ will be discarded. We will use the remaining estimates to construct an estimate of the valuation profile $\tilde{\mathbf{v}}(\sigma, \mathbf{v})$ and revenue curve $\tilde{\mathbf{R}}(\sigma, \mathbf{v})$ as follows.

Definition 5.1 (*estimated revenue curve and valuation profile*). For any j for which the estimate $\tilde{n}_j$ of the number of agents with values α^j is at least the minimum required support m, define point $Q_j = (\tilde{n}_j, \alpha^j \tilde{n}_j)$. The *estimated revenue curve*, $\tilde{\mathbf{R}}$, is the minimum non-decreasing concave function that upper-bounds the point set $\{Q_j\}_{j \in \mathbb{Z}}$ and the origin. Let j_k denote the k^{th} largest index such that point Q_{j_k} is on $\tilde{\mathbf{R}}$. For each k, the *estimated valuation profile*, $\tilde{\mathbf{v}}$, has $\tilde{n}_{j_k} - \tilde{n}_{j_{k-1}}$ values equal to α^{j_k}. Pad the remainder of $\tilde{\mathbf{v}}$ with zeros to get an n-vector. (See Figure 2.)

In the above construction $\tilde{\mathbf{v}}$ is the smallest (point-wise) valuation profile that has revenue curve $\tilde{\mathbf{R}}$, and furthermore, $\tilde{\mathbf{v}} \leq \mathbf{v}$. For statistical estimates $\tilde{n}_j(\sigma, \mathbf{v})$ the estimated revenue curve and valuation profile will be denoted $\tilde{\mathbf{v}}(\sigma, \mathbf{v})$ and $\tilde{\mathbf{R}}(\sigma, \mathbf{v})$. Our goal now is to show that with high probability the estimated revenue curve (and valuation profile) has consensus when a few agents, S, are omitted, i.e., $\tilde{\mathbf{R}}(\sigma, \mathbf{v}) = \tilde{\mathbf{R}}(\sigma, \mathbf{v}_{-S})$. To do this we define a notion of relevance for statistics j and show that with high probability there is simultaneous consensus for all relevant statistics.

Definition 5.2 (t*-consensus on* $\mathbf{v}$). Given a fixed valuation vector $\mathbf{v}$, a positive integer constant t and a fixed choice of σ, the jth statistic has a t*-consensus on* $\mathbf{v}$ if for every set $S \subset [n]$ of no more than t elements,

$$\tilde{n}_j(\sigma, \mathbf{v}) = \tilde{n}_j(\sigma, \mathbf{v}_{-S}).$$

Definition 5.3 (*relavant statistic*). For a given valuation vector $\mathbf{v}$ and a positive integer t, the jth statistic n_j is *relevant* if there exists $\sigma \in [0, 1]$, and a set $S \subseteq [n]$ of no more than t elements such that the point $Q_j(\sigma, \mathbf{v}_{-S})$ is on $\tilde{\mathbf{R}}(\sigma, \mathbf{v}_{-S})$.

Notice that when t-consensus happens for a relevant statistic j, then points $Q_j(\sigma, \mathbf{v})$ and $Q_j(\sigma, \mathbf{v}_{-S})$ in the construction of the estimated revenue curve are identical.

We now argue that the probability that any statistic j does not have consensus is roughly proportional to $1/n_j(\mathbf{v})$, that for relevant statistics the $n_j(\mathbf{v})$ values are geometrically increasing, and thus the union bound implies that all estimates of relevant statistics, and thus the estimated revenue curve, have consensus with high probability. This approach is adapted from Goldberg and Hartline [2005].

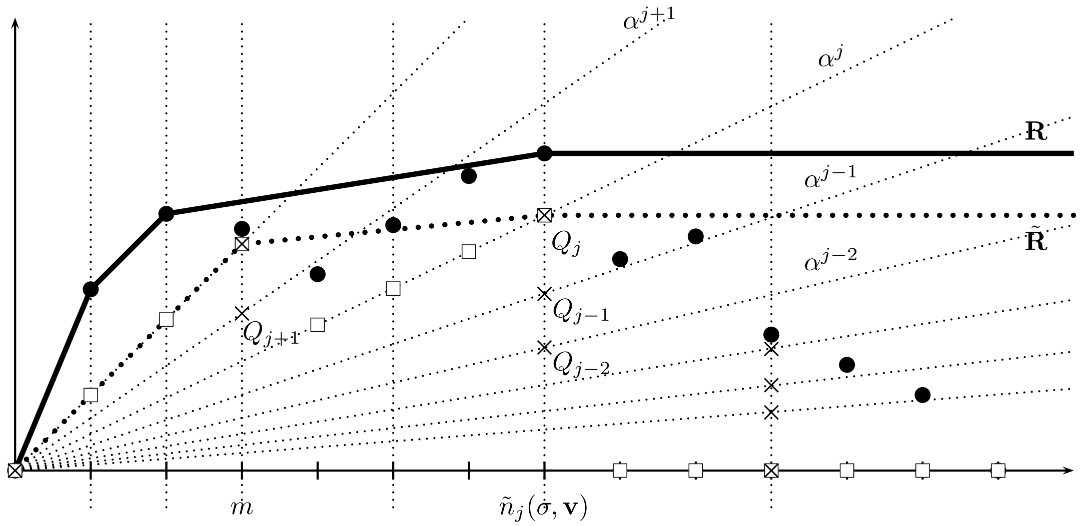

Fig. 2. The dotted vertical lines are $\lceil c^{d+\sigma} \rceil$ for $d \in \mathbb{Z}$; while the diagonal lines from the origin represent value α^j (as their slopes) for $j \in \mathbb{Z}$. Each value v_i is represented with "•" at point (i, iv_i). The line upper-bounding the black dots represents the original revenue curve, $\mathbf{R}$. For each diagonal line representing α^j, $n_j(\mathbf{v})$ is the number of values that lie above it; the dotted vertical line immediately to the left of the right-most among these values represents $\tilde{n}_j(\sigma, \mathbf{v})$; finally the intersection between this line and α^j is the point Q_j, represented with "$\times$". $\tilde{\mathbf{R}}$ and $\tilde{\mathbf{v}}$ are constructed from estimates $\tilde{n}_j(\sigma, \mathbf{v}) \geq m$ and their corresponding Q_js. The thick dotted function upper-bounding these Q_js is $\tilde{\mathbf{R}}$, and each $\tilde{v}_i$ is represented by the point $(i, i\tilde{v}_i)$ with a "□".

LEMMA 5.4. *For any* $\mathbf{v}$ *and* $\sigma \sim U[0, 1]$*, the probability that the* j*th statistic has a* t*-consensus on* $\mathbf{v}$ *is at least* $1 + \log_c \left(1 - \frac{t}{n_j(\mathbf{v})}\right)$.

PROOF. Observe that for every set S of no more than t elements, $n_j(\mathbf{v}) - t \leq n_j(\mathbf{v}_{-S}) \leq n_j(\mathbf{v})$. These inequalities hold since when some bids are removed, the number of bids above any α^j decrease, but only by at most the size of the removed set. Thus the probability that $\text{Consens}(\sigma, n_j(\mathbf{v})) = \text{Consens}(\sigma, n_j(\mathbf{v}_{-S}))$ is at least $1 - \log_c \frac{n_j(\mathbf{v})}{n_j(\mathbf{v}) - t}$ as suggested by Lemma 2.3. The lemma follows from the power rule for logarithm. □

LEMMA 5.5. *The values above successive relevant statistics are bounded by a geometrically increasing function: for any relevant statistic* j*,* $n_j(\mathbf{v}) \geq m\alpha^{r-j}$ *where* r *is the largest index of any relevant statistic.*

PROOF. First note that the largest index of a relevant statistic r is well defined. For any j that is relevant, it must be that $n_j(\mathbf{v}) \geq m$; otherwise, j would be discarded by the estimated revenue curve construction. Thus, the largest index that may not be discarded is r that satisfies $\alpha^{r+1} > v_m \geq \alpha^r$.

From the definition of $\tilde{n}_j(\sigma, \mathbf{v}_{-S})$, we have $\alpha^j \tilde{n}_j(\sigma, \mathbf{v}_{-S}) \leq \alpha^j n_j(\mathbf{v}_{-S}) \leq \alpha^j n_j(\mathbf{v})$ for any σ and S. Since j is relevant, there exists a σ and S and with size at most t such that the corresponding point $Q_j(\sigma, \mathbf{v}_{-S})$ is higher than $Q_r(\sigma, \mathbf{v}_{-S})$; therefore, $\alpha^j \tilde{n}_j(\sigma, \mathbf{v}_{-S}) \geq \alpha^r \tilde{n}_r(\sigma, \mathbf{v}_{-S}) \geq m\alpha^r$. Combining this with the previous inequality, we have the desired claim. □

LEMMA 5.6. *The probability of t-consensus at all relevant values is at least*

$$1 + \log_c \left[1 - \tfrac{t\alpha}{m(\alpha-1)}\right].$$

PROOF. We will first bound the probability of consensus at one relevant value, then use the union bound to find a lower bound on the probability of consensus at all relevant values. For any relevant statistic j, let $\mathcal{E}_j$ denote the event that n_j has a t-consensus on $\mathbf{v}$.

$$\Pr[\mathcal{E}_j] \geq 1 + \log_c \left(1 - \tfrac{t}{n_j(\mathbf{v})}\right) \geq 1 + \log_c \left(1 - \tfrac{t}{m}\alpha^{-(r-j)}\right).$$

The first inequality is from Lemma 5.4, while the second inequality is from Lemma 5.5. Let $J = \{j \,:\, n_j \text{ is relevant}\}$. The probability that all relevant statistics have t-consensus on $\mathbf{v}$, using the union bound, is

$$\begin{aligned}\Pr[t\text{-consensus}] &\geq 1 - \sum\nolimits_{j\in J} \Pr[\neg\mathcal{E}_j] \geq 1 + \sum\nolimits_{j\in J} \log_c \left(1 - \tfrac{t}{m}\alpha^{-(r-j)}\right) \\ &\geq 1 + \sum\nolimits_{i\geq 0} \log_c \left(1 - \tfrac{t}{m}\alpha^{-i}\right) = 1 + \log_c \left[\prod\nolimits_{i\geq 0} \left(1 - \tfrac{t}{m}\alpha^{-i}\right)\right] \\ &\geq 1 + \log_c \left[1 - \sum\nolimits_{i\geq 0} \tfrac{t}{m}\alpha^{-i}\right] = 1 + \log_c \left[1 - \tfrac{t}{m}\tfrac{\alpha}{\alpha-1}\right]. \quad \square\end{aligned}$$

The last thing that we need for our estimated revenue curves is for them to be good estimates. This follows directly from their definition.

LEMMA 5.7. *For any σ and $\mathbf{v}$, the consensus revenue curve $\tilde{\mathbf{R}}(\sigma, \mathbf{v})$ is a $c\alpha$-approximation of the revenue curve $\mathbf{R}^{(m')}$ for $m' = \lfloor mc \rfloor$ and truncated valuation profile $\mathbf{v}^{(m')} = (v_{m'}, \ldots, v_{m'}, v_{m'+1}, v_{m'+2}, \ldots, v_n)$, i.e., $\tilde{\mathbf{R}}(\sigma, \mathbf{v}) \geq \frac{1}{c\alpha}\mathbf{R}^{(m')}$.*

PROOF. It is sufficient to show that $\tilde{R}_i \geq \frac{1}{c\alpha} i v_i$ for $i \geq m'$ as concavity of revenue curves would then imply the lemma. Consider then any index $i \geq m'$ and let j be the index of the statistic that satisfies $\alpha^j \leq v_i < \alpha^{j+1}$. Since $n_j(\mathbf{v}) \geq i$, by the definition of $\tilde{n}_j$ and m', respectively, $\tilde{n}_j \geq \lceil i/c \rceil \geq m$; therefore, statistic j is not discarded in the first step of the construction of $\tilde{\mathbf{R}}$. Furthermore, the point $Q_j = (\tilde{n}_j, \alpha^j \tilde{n}_j)$ is above and to the left of $(i, \frac{1}{c\alpha} i v_i)$ because $\tilde{n}_j \leq i$, $\alpha^j \geq v_i/\alpha$, and $\tilde{n}_j \geq i/c$. Monotonicity of $\tilde{\mathbf{R}}$, then, implies the desired $\tilde{R}_i \geq \frac{1}{c\alpha} i v_i$. $\square$

6. DESIGNED MECHANISM

We now proceed to define a mechanism that is a good approximation to $\mathrm{EFO}^{(2)}(\mathbf{v})$. This mechanism will be a convex combination of a primitive cross-checked consensus-estimate profit-extraction mechanism and an extension of the Vickrey [1961] auction to downward-closed permutation environments.

Definition 6.1. The *primitive cross-checked consensus-estimate profit-extraction* mechanism, CCEPE' is the profit extraction mechanism $\mathrm{PE}_{\tilde{\mathbf{v}}}$ (Definition 4.1) composed (Definition 3.2) with the valuation profile estimate (Definition 5.1). CCEPE' is parameterized implicitly by α, c, and m.

Definition 6.2. The pseudo-Vickrey auction, PV, serves the highest valued agent (and charges her the second highest agent's value) if doing so is feasible with respect to the set system; otherwise, it rejects everyone.

Definition 6.3. The *cross-checked consensus-estimate profit-extraction* mechanism, CCEPE, is a convex combination of the pseudo-Vickrey auction (with probability p) and CCEPE' (with probability $1 - p$). CCEPE is parameterized implicitly by α, c, m, and p.

The pseudo-Vickrey mechanism is intended to obtain good revenue from the highest-valued agents where as CCEPE' is intended to obtain good revenue from the lower-valued agents. The convex combination obtains good revenue over all. This analysis is given by the following lemmas and theorem.

LEMMA 6.4. *For any* $\mathbf{v}$*,* $m' = \lfloor mc \rfloor$*, and downward-closed permutation environment; the expected revenue of* CCEPE' *is a* β'*-approximation of* $\mathrm{EFO}^{(m')}$ *with*

$$\beta' \leq c\alpha \left[1 + \log_c \left(1 - \tfrac{2\alpha}{m(\alpha-1)}\right)\right]^{-1}.$$

PROOF. Let $(\tilde{\mathbf{v}}, I)$ denote the outcome of the cross-checked consensus estimate of the valuation profile. By Lemma 5.6 (with $t = 2$) with probability at least $1 + \log_c \left[1 - \frac{2\alpha}{m(\alpha-1)}\right]$ all agents are cross-checked and $I = [n]$. In this case, $\tilde{\mathbf{v}} \leq \mathbf{v}$ (from Definition 5.1) and $\tilde{\mathbf{R}} \geq \frac{1}{c\alpha}\mathbf{R}^{(m')}$ (Lemma 5.7); therefore, the profit-extraction mechanism for $\tilde{\mathbf{v}}$, $\mathrm{PE}_{\tilde{\mathbf{v}}}$, on $\mathbf{v}$ obtains revenue at least $\frac{1}{c\alpha}\,\mathrm{EFO}^{(m')}(\mathbf{v})$ (Theorem 4.5). □

LEMMA 6.5. *For any* $\mathbf{v}$ *and downward-closed permutation environment, the pseudo-Vickrey auction revenue is at least the envy-free optimal payment of the highest-valued agent, i.e.,* $\mathrm{IC}^{\mathrm{PV}}(\mathbf{v}) \geq \mathrm{EFO}_1^{(2)}(\mathbf{v})$.

PROOF. Assume that $v_1 = v_2$, this is without loss for this lemma because both pseudo-Vickrey's revenue and $\mathrm{EFO}^{(2)}$'s revenue is the same on $\mathbf{v}$ and $\mathbf{v}^{(2)} = (v_2, v_2, v_3, \ldots, v_n)$. The payment for agent 1 upon winning in pseudo-Vickrey is $v_2 = v_1$; the payment upon winning in $\mathrm{EFO}^{(2)}$ is at most $v_1 = v_2$. The probability that agent 1 wins in pseudo-Vickrey is the highest of any feasible allocation (because agent 1 wins whenever serving agent 1 is feasible); in particular it is as high as that of $\mathrm{EFO}^{(2)}$. Therefore, the revenue from agent 1 in pseudo-Vickrey is at least that of $\mathrm{EFO}^{(2)}$. □

THEOREM 6.6. *For any* $\mathbf{v}$ *and downward-closed permutation environment,* CCEPE *is a* β*-approximation to* $\mathrm{EFO}^{(2)}(\mathbf{v})$ *where* β *satisfies*

$$\beta \leq \max\left\{\tfrac{\lfloor mc \rfloor}{p}, \tfrac{c\alpha}{1-p}\left[1 + \log_c\left(1 - \tfrac{2\alpha}{m(\alpha-1)}\right)\right]^{-1}\right\}.$$

PROOF. We will separate our revenue into two parts; the first part is obtained from the top $m' = \lfloor mc \rfloor$ agents, denoted $H = \{1, \ldots, m'\}$; and the second part is obtained from the remaining $n - m'$, denoted $L = \{m' + 1, \ldots, n\}$.

The contribution to the envy-free optimal revenue by the top agents satisfies $\mathrm{EFO}_H^{(2)}(\mathbf{v}) \leq m'\,\mathrm{IC}^{\mathrm{PV}}(\mathbf{v})$. This bound follows from Lemma 6.5 and the observation that envy-free payments are monotonically non-increasing in agent values.

The contribution to the envy-free optimal revenue by the bottom agents satisfies $\mathrm{EFO}_L^{(2)}(\mathbf{v}) \leq \mathrm{EFO}^{(m')}(\mathbf{v})$. This bound follows as $\mathrm{EFO}^{(m')}$ could try to simulate the outcome of $\mathrm{EFO}^{(2)}$ and would then receive the same contribution to revenue from agents L as $\mathrm{EFO}^{(2)}$; of course, its revenue from all agents must only be higher.

In conclusion, $\mathrm{EFO}^{(2)}(\mathbf{v}) \leq m'\,\mathrm{IC}^{\mathrm{PV}}(\mathbf{v}) + \mathrm{EFO}^{(m')}(\mathbf{v})$.

The revenue of our mechanism, CCEPE, the sum of a $\beta_1 = m'/p$ approximation to $m'\,\mathrm{IC}^{\mathrm{PV}}(\mathbf{v})$ and a $\beta_2 = \beta'/(1-p)$ approximation to $\mathrm{EFO}^{(m')}(\mathbf{v})$, with β' as defined in Lemma 6.4. Therefore, it is a $\beta = \max(\beta_1, \beta_2)$ approximation to $\mathrm{EFO}^{(2)}(\mathbf{v})$. □

We can optimize the parameters of CCEPE to obtain the following corollary.

COROLLARY 6.7. *For any* $\mathbf{v}$ *and downward-closed permutation environment,* CCEPE *with* $p = 0.627$*,* $c = 1.666$*,* $\alpha = 2.734$ *and* $m = 12$ *is a* 30.4*-approximation to* $\text{EFO}^{(2)}(\mathbf{v})$.

REFERENCES

ALAEI, S., MALEKIAN, A., AND SRINIVASAN, A. 2009. On random sampling auctions for digital goods. In *Proceedings of the 10th ACM Conference on Electronic Commerce*. Association for Computing Machinery, New York, NY, USA, 187–196.

ARCHER, A., PAPADIMITRIOU, C., TALWAR, K., AND TARDOS, E. 2003. An approximate truthful mechanism for combinatorial auctions with single parameter agents. In *Proceedings of the fourteenth annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 205–214.

BALCAN, M.-F., BLUM, A., HARTLINE, J., AND MANSOUR, Y. 2008. Reducing mechanism design to algorithm design via machine learning. *J. Comput. Syst. Sci. 74,* 8, 1245–1270.

BALIGA, S. AND VOHRA, R. 2003. Market research and market design. *Advances in Theoretical Economics 3,* 1.

BULOW, J. AND ROBERTS, J. 1989. The simple economics of optimal auctions. *The Journal of Political Economy 97,* 5, 1060–90.

CAILLAUD, B. AND ROBERT, J. 2005. Implementation of the revenue-maximizing auction by an ignorant seller. *Review of Economic Design 9*, 127–143.

CLARKE, E. H. 1971. Multipart pricing of public goods. *Public Choice 11*, 17–33.

COASE, R. 1972. Durability and monopoly. *Journal of Law and Economics 15,* 1, 143–149.

DEVANUR, N. AND HARTLINE, J. 2009. Limited and online supply and the bayesian foundations of prior-free mechanism design. In *Proceedings of the 10th ACM Conference on Electronic Commerce*. Association for Computing Machinery, New York, NY, USA, 41–50.

DHANGWATNOTAI, P., ROUGHGARDEN, T., AND YAN, Q. 2010. Revenue maximization with a single sample. In *Proceedings of the 11th ACM Conference on Electronic Commerce*. Association for Computing Machinery, New York, NY, USA, 129–138.

GOLDBERG, A. AND HARTLINE, J. 2003. Competitiveness via consensus. In *Proceedings of the fourteenth annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 215–222.

GOLDBERG, A., HARTLINE, J., KARLIN, A., SAKS, M., AND WRIGHT, A. 2006. Competitive auctions. *Games and Economic Behavior 55,* 2, 242–269. Mini Special Issue: Electronic Market Design.

GOLDBERG, A., HARTLINE, J., AND WRIGHT, A. 2001. Competitive auctions and digital goods. In *Proceedings of the 12th annual ACM-SIAM Symposium on Discrete Algorithms*. 735–744.

GROVES, T. 1973. Incentives in teams. *Econometrica 41*, 617–631.

HARTLINE, J. AND GOLDBERG, A. 2005. Collusion-resistant mechanisms for single-parameter agents. In *Proceedings of the 16th Annual ACM-SIAM Sympsium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 620–629.

HARTLINE, J. AND MCGREW, R. 2005. From optimal limited to unlimited supply auctions. In *Proceedings of the 6th ACM Conference on Electronic Commerce*. Association for Computing Machinery, New York, NY, USA, 175–182.

HARTLINE, J. AND ROUGHGARDEN, T. 2008. Optimal mechanism design and money burning. In *Proceedings of the 41st ACM Symposium on Theory of Computing*. Association for Computing Machinery, New York, NY, USA, 75–84.

HARTLINE, J. AND YAN, Q. 2011. Envy, truth, and profit. In *Proceedings of the 12th ACM Conference on Electronic Commerce*. Association for Computing Machinery, New York, NY, USA, 243–252.

JACKSON, M. 2001. A crash course in implementation theory. *Social Choice and Welfare 18,* 4, 655–708.

JACKSON, M. AND KREMER, I. 2007. Envy-freeness and implementation in large economies. *Review of Economic Design 11*, 185–198.

MYERSON, R. B. 1981. Optimal auction design. *Mathematics of Operations Research 6,* 1, 58–73.

SEGAL, I. 2003. Optimal pricing mechanisms with unknown demand. *American Economic Review 93,* 3, 509–529.

VICKREY, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance 16*, 8–37.