

Combinatorial Pen Testing (Or Consumer Surplus of Deferred-Acceptance Auctions)

Aadityan Ganesh  

Princeton University, NJ, USA

Jason Hartline  

Northwestern University, Evanston, IL, USA

Abstract

Pen testing is the problem of selecting high-capacity resources when the only way to measure the capacity of a resource expends its capacity. We have a set of n pens with unknown amounts of ink and our goal is to select a feasible subset of pens maximizing the total ink in them. We are allowed to learn about the ink levels by writing with them, but this uses up ink that was previously in the pens.

We identify optimal and near optimal pen testing algorithms by drawing analogues to auction theoretic frameworks of deferred-acceptance auctions and virtual values. Our framework allows the conversion of any near optimal deferred-acceptance mechanism into a near optimal pen testing algorithm. Moreover, these algorithms guarantee an additional overhead of at most $(1 + o(1)) \ln n$ in the approximation factor to the omniscient algorithm that has access to the ink levels in the pens. We use this framework to give pen testing algorithms for various combinatorial constraints like matroid, knapsack, and general downward-closed constraints, and also for online environments.

2012 ACM Subject Classification Theory of computation → Theory and algorithms for application domains; Theory of computation → Approximation algorithms analysis; Theory of computation → Algorithm design techniques

Keywords and phrases Pen testing, consumer surplus, money-burning, deferred-acceptance auctions

Digital Object Identifier 10.4230/LIPIcs.ITCS.2025.52

Related Version *Full Version:* <https://arxiv.org/abs/2301.12462>

Funding *Jason Hartline:* This work is supported in part by the NSF award CCF-2229162.

1 Introduction

Pen testing is the problem of selecting high capacity resources when the only way to measure the capacity of the resource expends its capacity [23]. We show that any ascending-price auction can be converted into an equivalent pen testing algorithm. We apply the auction theoretic framework of virtual values to identify optimal pen testing algorithms. This connection allows many existing results from auction theory to be applied to pen testing and gives optimal and near optimal pen testing algorithms in combinatorial and online environments.

The pen testing problem of Qiao and Valiant [23] is the following. We have a set of pens with varying amounts of remaining ink and we want to choose a pen with the largest amount of ink left. We have access to the distribution of ink levels in these pens, but we are only allowed to learn about the ink levels by writing with the pens to test them. While writing gives information about whether there is still ink left, it uses up ink that was previously in the pen. Pens that are expended due to testing can be discarded without incurring any penalty. The combinatorial pen testing problem generalizes the pen testing problem to selecting a feasible subset of pens, e.g., according to a matroid or knapsack constraint, to maximize the total remaining ink in the chosen subset.



© Aadityan Ganesh and Jason Hartline;

licensed under Creative Commons License CC-BY 4.0

16th Innovations in Theoretical Computer Science Conference (ITCS 2025).

Editor: Raghu Meka; Article No. 52; pp. 52:1–52:22

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Deferred-acceptance auctions [21] formalize the notion of ascending-price mechanisms for allocating goods to strategic agents. These are mechanisms that greedily reject the least promising agent by increasing the price for getting allocated at each stage. For instance, to auction off one good, the deferred-acceptance mechanism keeps increasing the price, rejecting agents whenever their value for the good falls below the price until exactly one agent remains in the auction. The *surplus* of the mechanism is the sum of the values of all the allocated agents and the *consumer surplus* is the surplus minus the payments made by the agents.

We provide a black box reduction to convert any surplus optimal (or near optimal) deferred-acceptance auction into an optimal (or near optimal) pen testing algorithm with equivalent performance guarantees. In this reduction, pens and their ink levels are analogous to agents and their values. Payments correspond to the ink used up via testing, and as a consequence, the consumer surplus of the deferred-acceptance auction corresponds to the residual ink in the pen testing environment. Therefore, optimizing residual ink reduces to designing consumer-surplus-optimal deferred-acceptance mechanisms.

Qiao and Valiant [23] propose comparing the performance of a given pen testing algorithm against the *omniscient benchmark*, the performance of an omniscient algorithm that has access to the ink levels in the pens. The omniscient benchmark can thereby choose the optimal feasible subset of pens without having to burn any ink (i.e, payments in the auction environment), and is thus, analogous to the surplus optimal (sealed-bid) mechanism in the corresponding auction environment. Thus, the *omniscient approximation*, the ratio between the omniscient benchmark to the residual ink in the pen testing algorithm, is analogous to the ratio between the optimal surplus and the consumer surplus of the corresponding deferred-acceptance auction.

In the literature, near optimal deferred-acceptance mechanisms are evaluated against optimal (sealed bid) mechanisms, an upper bound to the optimal deferred-acceptance mechanism. Define the *standard benchmark* as the optimal (sealed-bid) mechanism for consumer surplus and let the *standard approximation* $\gamma(n)$ be the ratio between the standard benchmark and the consumer surplus of a given deferred-acceptance mechanism. Note that the standard benchmark does not have a natural pen testing analog, but is a much tighter upper bound on the optimal pen testing algorithm than the omniscient benchmark. Thus, the reduction to deferred-acceptance mechanism design enables a much tighter approximation bound for the optimal pen testing algorithm.

The reduction framework decomposes into two steps for designing near optimal deferred-acceptance mechanisms. One of the two steps involves proving upper bounds on the standard approximation $\gamma(n)$. In the other step, for any given feasibility constraint, we consider the *optimal omniscient approximation*, the ratio between the omniscient benchmark (here, the optimal surplus) and the standard benchmark (here, the optimal consumer surplus). As we will argue in Theorem 1, the omniscient approximation $\pi(n)$ is at most the product $\gamma(n)\zeta(n)$ between the two approximation ratios. At a high level, $\gamma(n)$ captures the loss in surplus due to running a deferred-acceptance auction and $\zeta(n)$ captures the loss in surplus due to burning all payments. Beyond consumer surplus optimization, the reduction framework applies more generally to optimization problems with two natural benchmarks, a weaker standard benchmark and much more powerful omniscient benchmark.

► **Theorem 1.** *Consider a combinatorial pen testing environment with n pens, the analogous auction environment with an optimal omniscient approximation $\zeta(n)$, and a deferred-acceptance mechanism with a standard approximation equal to $\gamma(n)$. Then, there is a pen testing algorithm that is a $\gamma(n)$ -approximation to the standard benchmark and $\pi(n) = \gamma(n)\zeta(n)$ -approximation to the omniscient benchmark.*

In the first step of our framework, we bound the optimal omniscient approximation $\zeta(n)$ under various feasibility constraints. Generally, remember that the optimal omniscient approximation $\zeta(n)$ is the ratio between the omniscient benchmark and the standard benchmark, and for our specific application, it corresponds to the loss in surplus due to payments getting burnt instead of getting transferred to the auctioneer. We show that $\zeta(n) \leq (1 + o(1)) \ln n$ in any combinatorial environment. Our result tightens the best known upper bound on $\zeta(n)$ for most settings of interest. Fotakis et al. [13] show that the optimal omniscient approximation is at most $\frac{2}{\ln 2} (1 + o(1)) \ln m$, where m is the number of feasible outcomes. Our bound improves their bound by removing the $\frac{2}{\ln 2}$ constant and by replacing m by n , which is usually smaller.¹ Hartline and Roughgarden [18] show that $\zeta(n) \leq \frac{2}{\ln 2} (1 + o(1)) \ln \frac{n}{k}$ when any k out of the n agents can be allocated the good simultaneously, i.e., any k pens can be chosen in the corresponding pen testing environment. We improve their bound by a multiplicative factor of at least $1.36(1 - o(1))$. We also prove that $\zeta(n) \geq (0.577 + \ln \frac{n}{k})$ for the k -identical goods setting, which we show is tight up to $(1 + o(1))$ terms when $k = o(n)$.

Our upper bound for $\zeta(n)$ is via a reduction from multi-agent environments to a single-agent environment, similar in flavour to Feng et al. [12]. In a single-agent environment where the agent is allocated with an ex-ante allocation probability q , we show that the ratio of the optimal surplus to the optimal consumer surplus is at most $1 - \ln q$. Our reduction allows us to convert the $(1 - \ln q)$ -approximation in the single-agent setting to a $(1 + o(1)) \ln n$ -approximation for any multi-agent environment.

We reduce designing multi-agent approximation mechanisms to single-agent approximations with ex-ante constraints q bounded away from 0. The approach is inspired by Feng et al. [12], who prove a similar reduction when the single-agent approximation is bounded above by a constant for all $q \in [0, 1]$. Their approach cannot be directly applied to the optimal omniscient approximation of consumer surplus because the single-agent approximation $(1 - \ln q)$ approaches ∞ as $q \rightarrow 0$. We adapt ideas from Hartline and Taggart [17] to prove a reduction to the single-agent setting with an ex-ante constraint $q \in [1/n, 1]$. In this range, the single-agent approximation ratio is at most $(1 + \ln n)$. The resulting construction establishes an optimal omniscient approximation of $(1 + o(1)) \ln n$ for multi-agent environments with arbitrary feasibility constraints.

The second step of the reduction framework from pen testing to deferred-acceptance auctions upper bounds the standard approximation $\gamma(n)$, the ratio between the optimal consumer surplus of a sealed-bid mechanism and the consumer surplus of a given deferred-acceptance auction. Through Myerson's virtual value framework [22], we argue that the above is equivalent to finding an upper bound for the optimal expected surplus divided by the expected surplus of some deferred-acceptance auction. Surplus optimal deferred-acceptance auctions are known for k -identical goods and more general matroid feasibility constraints [21, 4]. We give a 2-approximate deferred-acceptance mechanism for knapsack constraints. Beyond matroids and knapsacks, Feldman et al. [11] give a logarithmic approximate deferred-acceptance auction for any downward-closed feasibility constraint. We observe an inherent downward-closure in the pen testing problem that allows the auction of Feldman et al. [11] to be extended to all combinatorial constraints. The set of pens chosen by any pen testing algorithm can be padded by adding (possibly expended) pens to always output a maximal feasible set.

¹ In some non-downward-closed settings like public projects, $m < n$. For example, $m = 2$ for building a bridge that serves everyone. Either the bridge is built or not. However, $m \geq n$ in all downward-closed environments. Indeed, for each bidder, we assume there exists an outcome where they are included; otherwise, the bidder can be dropped from consideration altogether.

■ **Table 1** Prior bounds (if any, references indexed by superscripts) and our upper bounds. $n =$ number of agents, $m =$ number of feasible outcomes. Note that $m \geq n$ for all the environments listed in the table. The bounds for optimal omniscient approximation ζ and omniscient approximation π are normalized by a $(1 + o(1))$ factor. Prior work: ^{*}Hartline and Roughgarden [18]; ^{**}Milgrom and Segal [21], Bikhchandani et al. [4]; [§]Fotakis et al. [13]; [†]Feldman et al. [11]; ^{††}Samuel-Cahn [25], Chawla et al. [6]; ^{§§}Kleinberg and Weinberg [19]; [‡]Qiao and Valiant [23]; ^{‡‡}Chawla et al. [6], Yan [26]. All bounds in the table also hold for the underlying auction environment. Pen testing algorithms with the same guarantees as downward-closed constraints can be achieved for arbitrary combinatorial constraints.

Environment	$\frac{\zeta(n)}{(1+o(1))}$	$\gamma(n)$	$\frac{\pi(n)}{(1+o(1))}$
Select $k = o(n)$	$\frac{2}{\ln 2} \ln \frac{n}{k}^*$	$\ln \frac{n}{k}$	$\ln \frac{n}{k}^{**}$
Select $k = \Theta(n)$	$\frac{2}{\ln 2} \ln \frac{n}{k}^*$	$2.27 (0.577 + \ln \frac{n}{k})$	$2.27 (0.577 + \ln \frac{n}{k})^{**}$
Matroids	$\frac{2}{\ln 2} \ln m^{\S}$	$\ln n$	$\ln n^{**}$
Knapsack	$\frac{2}{\ln 2} \ln m^{\S}$	$\ln n$	$2 \cdot \ln n$
General Downward-Closed	$\frac{2}{\ln 2} \ln m^{\S}$	$\ln n$	$O(\log n \log \log m)^{\dagger}$
Select 1 Online Oblivious	$\ln n$	$2^{\dagger\dagger}$	$O(\log n)^{\ddagger} \cdot 2 \cdot \ln n$
Select 1 Online Sequential	$\ln n$	$\frac{e}{e-1}^{\ddagger\ddagger}$	$O(\log n)^{\ddagger} \cdot \frac{e}{e-1} \cdot \ln n$
Select 1 Online IID	$\ln n$	$\frac{e}{e-1}^{\ddagger\ddagger}$	$e \ln n^{\ddagger} \cdot \ln n$
Matroid Online Oblivious	$\ln n$	$2^{\S\S}$	$2 \cdot \ln n$
Matroid Online Sequential	$\ln n$	$\frac{e}{e-1}^{\ddagger\ddagger}$	$\frac{e}{e-1} \cdot \ln n$

Online posted-price mechanisms are a special case of deferred-acceptance auctions. Online environments are ones where an irrevocable decision regarding choosing a pen needs to be taken before testing the next one. We consider the settings where the order of testing pens is adversarially chosen (the oblivious version, *cf.* prophet inequalities, [6]) and where the algorithm can choose the order (the sequential version, *cf.* correlation gap, [26]). With our analysis of the optimal omniscient approximation ζ and known bounds on the standard approximation γ of online pricing mechanisms, the reduction framework yields better performance guarantees than the $O(\log n)$ bounds by Qiao and Valiant [23] by constant factors for these environments.

The results detailed above are summarized in Table 1.

Our reduction from pen testing to deferred-acceptance auction design improves bounds from previous literature for all the environments that we discuss; however, it does not always yield a tight bound. For the online IID environment (the online environment where the ink levels of all pens are drawn IID), we identify an algorithm with an omniscient approximation $\pi(n)$ that is a constant factor less than $\gamma(n)\zeta(n)$.

Our results and techniques are of interest beyond pen testing. Our upper bound of $(1 + o(1))\ln n$ on the optimal omniscient approximation $\zeta(n)$ is the ratio between the optimal surplus and the optimal consumer surplus, and is a mechanism design problem of independent interest [18, 13, 15]. Recall that the surplus created in an auction equals the sum of the consumer surplus and the revenue generated by the auctioneer. However, when payments made by the winners are burnt instead of being transferred to the auctioneer, the total surplus created is just the consumer surplus. The need to burn payments can be diverse. For instance:

1. Making monetary transactions is not desirable or repugnant, like in setting prices for welfare schemes for vulnerable populations. Ordeal-based mechanisms are employed in such scenarios, where agents make non-transferable non-monetary payments, like spending time waiting in a line to avail the provisions provided in the welfare scheme [20, 3, 5].
2. In computer systems like data packet routing networks, where technological infeasibility to collect payments necessitates charging non-monetary payments [10, 9, 18].

3. The auctioneer cannot be trusted to truthfully implement the auction, like in a blockchain, where a miner runs an auction to decide which set of transaction goes into the next block it proposes. To stop the miner from colluding with agents or from benefiting from injecting fake bids, it becomes essential to burn a substantial portion of the transaction fees collected from the agents [24, 7, 14].

The technique to obtain the $(1 - \ln q)$ -approximation in the single-agent environment can have further applications. We express consumer surplus optimization as a linear program and explicitly compute the worst-case distribution. This provides a proof technique which automatically gives tight approximation guarantees. For example, we reuse the technique to compute the worst-case distribution for the omniscient approximation for the online IID environment. We improve the bound obtained from the reduction framework to a multiplicative approximation ratio within an additive 1 of the H_n lower bound of Qiao and Valiant [23] (where H_n is the n^{th} harmonic number).

Paper Organization. We make the reduction from the pen-testing problem to designing deferred-acceptance auctions formal in Section 2. We execute the first step of the reduction framework – upper bounding the optimal omniscient approximation $\zeta(n)$ – in Section 3. We bound the approximation ratio between the surplus and consumer surplus in a single-agent environment in Section 3.1 and discuss the reduction from multi-agent to single-agent settings in Section 3.2. The second step of the reduction from pen testing algorithms to deferred-acceptance mechanisms is discussed in Section 4, where we review deferred-acceptance mechanisms for various environments and discuss the corresponding bounds we get for pen testing. Finally, we revisit the online IID single item environment and achieve a tight (up to $o(1)$ terms) bound in Section 5.

2 The Pen Testing Problem and Ascending-Price Auctions

In this section, we formally define the combinatorial pen testing problem, generalized from Qiao and Valiant [23], and explore its connections to ascending-price mechanisms.

► **Definition 2.** *A pen testing instance is described by a set $N = \{1, 2, \dots, n\}$ of pens with unknown ink levels v_1, \dots, v_n , each v_i drawn independently from the distribution F_i , and a subset \mathcal{P} of the powerset of the pens denoting feasible subsets of pens. The residual ink level u_i of pen i is initiated to v_i . We can test the pens before making a decision. Test (i, θ_i) is done by writing with pen i for time θ_i . We receive a binary signal at the end of the test:*

1. *If $u_i \geq \theta_i$: The test succeeds. The pen now has an (unobservable) ink level $u_i \leftarrow u_i - \theta_i$.*
2. *If $u_i < \theta_i$: The test fails. The pen now has no ink left, i.e., $u_i \leftarrow 0$.*

We can choose to run multiple tests on the pens. We need to output some feasible subset $P \in \mathcal{P}$ of pens, maximizing

$$\sum_{i \in P} u_i.$$

The goal is to optimize the total amount of ink remaining in the set of chosen pens.

Compare a pen testing instance with n pens and a feasibility constraint \mathcal{P} with an n -agent auction under the same feasibility constraint. The original levels of ink correspond to the value each agent gets upon being allocated. The ink spent through testing is analogous to the prices in the auction. Thus, maximizing total residual ink is equivalent to optimizing consumer surplus, i.e., the sum of the values of the winning agents minus the sum of all payments. Note that auctions generally give the auctioneer the added advantage of the

ability to solicit bids from agents. In pen testing algorithms, we only get to learn whether the ink left in the pen is more than some threshold and in doing so, we irrevocably expend ink up to the threshold. This is similar to ascending-price auctions, where the auctioneer irreversibly increases the price faced by each agent, and in doing so, only learns whether the agent has a value at least the price.

Milgrom and Segal [21] describe the wide class of ascending-price auctions called deferred-acceptance mechanisms.

► **Definition 3** (Milgrom and Segal [21]; Deferred-Acceptance Auctions). *A deferred-acceptance auction is held across multiple stages. For each stage t , the auction maintains a set of active bidders A_t satisfying $A_1 \supseteq A_2 \supseteq \dots \supseteq A_t$, where the initial active set A_1 is the set of all bidders. The auction is characterized by a (possibly randomized) pricing rule \vec{p} , mapping the history of the auction at stage t to (discriminatory) prices for each agent satisfying $p_i(H_t) \geq p_i(H_{t-1})$ for all agents i and for all histories H_t at stage t arising out of stage $t-1$ history H_{t-1} , i.e., the prices are monotonously non-decreasing for all agents over time. Agents can opt to drop out once the prices are updated in stage t . A_{t+1} is the set of agents in A_t that did not drop out in stage t . The mechanism terminates at stage t when A_t becomes feasible. The agents in A_t are charged according to $\vec{p}(H_t)$.*

The definition from Milgrom and Segal [21] restricts deferred-acceptance mechanisms to be deterministic, which we relax for the purpose of designing pen testing algorithms. From the discussion above, note that any deferred-acceptance mechanism can be converted into a pen testing algorithm, by setting thresholds equal to the prices recommended by the mechanism. Thus, the optimal (near optimal) pen testing algorithm corresponds to the consumer surplus optimal (near optimal) deferred-acceptance auction.

We compare the performance of our deferred-acceptance mechanisms (and the corresponding pen testing algorithms) to the following benchmarks:

- The *omniscient benchmark* [23]: The omniscient benchmark corresponds to the omniscient pen testing algorithm that exactly knows the ink levels in the pens, and thus, does not have to waste ink by writing with them. In the reduction to the auction environment, the omniscient benchmark corresponds to a mechanism that neither loses out on the payments nor has to stick to an increasing trajectory of prices. Thus, the omniscient benchmark corresponds to the surplus optimal sealed-bid mechanism. The *omniscient approximation* $\pi(n)$ corresponds to the ratio between the omniscient benchmark and the performance of a given pen testing algorithm in the pen testing setting, and the omniscient benchmark and the consumer surplus of a given deferred-acceptance auction in the auction environment.
- The *standard benchmark*: The standard benchmark corresponds to the consumer-surplus-optimal sealed-bid mechanism. We define the *standard approximation* $\gamma(n)$ to be the ratio between the standard benchmark and the consumer surplus of a candidate deferred-acceptance mechanism. While the standard benchmark does not have a natural pen testing analogy, observe that the standard benchmark is a tighter upper bound to the optimal pen testing algorithm than the omniscient benchmark. Thus, the standard approximation gives a better upper bound on the ratio between the optimal pen testing algorithm and the residual ink gathered by a given pen testing algorithm.

Now that we have established the consumer surplus is a quantity of interest for pen testing, we review consumer surplus optimization through virtual values in the next section.

2.1 Consumer Surplus Optimization and Virtual Valuations

Let us begin in a single-agent environment with one good. Let the agent's value be drawn from the distribution F . The quantile q of an agent with value $v \sim F$ is the measure with respect to F of stronger values, i.e., $q = 1 - F(v)$. For $q \in [0, 1]$, let $v(q)$ be the inverse demand function of F satisfying $F(v(q)) = 1 - q$. In other words, $\Pr_{\hat{v} \sim F}(\hat{v} > v(q)) = q$. Throughout the paper, we assume $v(1) = 0$ and $\int_0^1 v(t)dt = 0$. It can be shown that these assumptions can be made without loss of generality.

Let $V(q) = \int_0^q v(t)dt$ be the price-posting surplus curve. Notice that $V(q)$ is the expected surplus from posting a price $v(q)$ to the agent, thereby allocating the good with probability q . By our assumption on F , $V(0) = 0$. Let $U(q) = V(q) - qv(q)$ be the price-posting consumer surplus curve. Similar to the price-posting surplus curve, $U(q)$ is the consumer-surplus from posting a price $v(q)$. Note that $0 \leq U(0) \leq V(0) = 0$, and thus, $U(0) = 0$.

Let $u(q) = U'(q) = -qv'(q)$ be the marginal price-posting consumer surplus curve. Note that v is monotonously non-increasing, and hence its derivative v' is non-positive. Thus, $u(q) \geq 0$ and U is monotone non-decreasing. A quick note on derivatives: whenever $v'(q)$ is well defined, $u(q)$ is well defined. At other points, $u(q)$ can be calculated using any sub-derivative of v . At points q of discontinuity of v , $v'(q) = -\infty$ and thus $u(q) = \infty$.

► **Theorem 4** (Myerson [22]). *In a Bayesian incentive-compatible mechanism with allocation rule y and payment rule p (over quantiles), the expected consumer surplus of an agent satisfies*

$$E_{q \sim U[0,1]}[v(q)y(q) - p(q)] = E_{q \sim U[0,1]}[u(q)y(q)] = E_{q \sim U[0,1]}[-U(q)y'(q)] + U(0)y(0)$$

Thus, the expected consumer surplus equals the expected virtual surplus with marginal consumer surplus as virtual values. Myerson [22] states that an allocation rule can be implemented as a truthful auction if and only if the allocation rule is monotone non-increasing in quantile space. Thus, the consumer surplus optimal mechanism optimizes for virtual surplus subject to the allocation rule being monotonously non-increasing.

When the marginal price-posting consumer surplus curve is non-increasing (and hence U is concave), optimizing for virtual surplus automatically ensures monotonicity of the allocation rule. We will call these distributions with non-increasing marginal price-posting consumer surplus curves *consumer surplus regular*.

However, u is not monotone non-increasing for common distributions like the uniform and normal distribution. In fact, u is monotone non-decreasing for these distributions. In such a case, the allocation that pointwise optimizes for virtual surplus might not satisfy monotonicity. Myerson [22] prescribes an ironing procedure to optimize for virtual surplus subject to monotonicity.

1. Construct the concave hull \bar{U} of the price-posting consumer surplus curve U
2. Define $\bar{u}(q) = \bar{U}'(q)$
3. Find the virtual surplus optimal mechanism using \bar{u} as virtual surplus

Since \bar{U} is concave, \bar{u} is non-increasing, and hence optimizing for the ironed virtual surplus is easier. Throughout the paper, we will follow the convention of attaching a bar on top of a curve to describe the ironed curve.

In a multi-agent environment, the interim allocation rule for an agent is the single-agent allocation rule that arises in expectation over the quantiles of all other agents. This captures the perspective of the agent after knowing its value, but before learning the values of the other agents.

► **Theorem 5.** *Let y be the interim allocation rule for some agent, satisfying $\frac{d}{dq}y(q) = 0$ for all q such that $U(q) \neq \bar{U}(q)$. Then*

$$E_q[u(q)y(q)] = E_q[\bar{u}(q)y(q)]$$

In other words, if the allocation rule does not differentiate between agents in an ironed region, then the ironed virtual surplus can be used to compute the consumer surplus instead of the actual virtual surplus.

► **Theorem 6** (Alaei et al. [1, 2]). *Let y be some allocation rule with interim allocation rule y_i monotonously non-increasing for each agent i . Then, there exists a mechanism \bar{y} with interim allocation rule \bar{y}_i for agent i such that the expected consumer surplus for agent i equals*

$$E_{q \sim U[0,1]}[u_i(q) \bar{y}_i(q)] = E_{q \sim U[0,1]}[\bar{u}_i(q) y_i(q)]$$

In other words, the expected consumer surplus of \bar{y} is the consumer surplus of the original mechanism if the virtual values were given by \bar{u}_i for agent i instead of u_i

2.2 Near-Optimal Deferred-Acceptance Mechanisms for Consumer Surplus

In this section, we prescribe a recipe similar to Feng et al. [12] to convert any approximately optimal deferred-acceptance mechanism for surplus into an approximately optimal deferred-acceptance mechanism for consumer surplus.

► **Definition 7** (Virtual-Pricing Transformation of a Deferred-Acceptance Mechanism). *Let DA^V be a deferred-acceptance mechanism designed to optimize surplus. Let v_i be the inverse demand function and u_i be the virtual value function (for consumer surplus) for agent i . The virtual-pricing transformation on DA^V , denoted by DA^U , implements DA^V in ironed virtual price space, i.e., whenever DA^V posts a price \hat{v}_i to agent i , DA^U posts a price $v_i(\theta_i)$ satisfying*

$$\theta_i = \sup\{\theta : \bar{u}_i(\theta) \geq \hat{v}_i\}$$

We will make some preliminary observations about the transformation:

1. Consider two runs of the transformed mechanism DA^U with some agent having two different values, both corresponding to the same ironed virtual value (which can happen if both these values are within the same ironed interval). From the definition of the transform, the agent faces the smallest price needed to differentiate itself from virtual values smaller than the threshold set by the mechanism. Consequently, the mechanism does not post prices from the middle of an ironed interval. Thus, in both these runs of the mechanism, the agent behaves identically. Since the mechanism does not discriminate between values with the same ironed virtual value, the consumer surplus of the mechanism equals the ironed virtual surplus.
2. If the original mechanism DA^V is a γ -approximation to the optimal surplus, in expectation over all product distributions, then the transformed mechanism DA^U is a γ -approximation to the optimal ironed virtual surplus, and hence, to the optimal consumer surplus.
3. Finally, it is straightforward to see that the transformed mechanism DA^U posts non-decreasing prices to each agent. Hence, DA^U is also a deferred-acceptance mechanism.

Recall that the standard approximation $\gamma(n)$ is the ratio between the consumer surplus of the optimal deferred-acceptance mechanism and our proposed mechanism while the optimal omniscient approximation $\zeta(n)$ is the ratio between the optimal surplus and the optimal consumer surplus. The proof of Theorem 1 is now straightforward.

Proof of Theorem 1. Let OPT^U denote the expected optimal consumer surplus, OPT^V denote the expected optimal surplus, and DA^U denote the expected consumer surplus of the deferred acceptance mechanism in Definition 7. Then,

$$\text{DA}^U \geq \frac{\text{OPT}^U}{\gamma(n)} \geq \frac{\text{OPT}^V}{\gamma(n)\zeta(n)}$$

where the first inequality follows from the definition of standard approximation $\gamma(n)$ and bullet 2 in the discussion above, while the second follows from the definition of optimal omniscient approximation $\zeta(n)$. The equivalent pen testing algorithm yields the necessary approximation ratios against the two benchmarks. ◀

In the next section, we analyze the optimal omniscient approximation $\zeta(n)$ for general combinatorial environments.

3 Consumer Surplus versus Surplus

Our approximation procedure follows two steps. In Section 3.1, we show that the optimal surplus is at most $1 - \ln q$ times the consumer surplus in a single-agent environment where the ex-ante probability of allocation is constrained to be at most q . In Section 3.2, we move from the single-agent environment to multi-agent environments combining approaches from Feng et al. [12] and Hartline and Taggart [17] to show a $(1 + o(1)) \ln n$ approximation between the consumer surplus and optimal surplus in n -agent environments with general feasibility constraints.

3.1 The Single-Agent Problem

Recall that V, U and \bar{U} are the price-posting surplus curve, price-posting consumer surplus curve and the ironed consumer surplus curve respectively. In this section, we establish a “closeness property” like in Feng et al. [12].

► **Theorem 8.** *For an ex-ante allocation probability $q \in [0, 1]$, the ratio of the optimal surplus $V(q)$ to the optimal consumer surplus $\bar{U}(q)$ is at most $1 - \ln q$.*

We will begin by proving Theorem 8 for consumer surplus regular distributions, and then extend the result to all distributions.

► **Lemma 9.** *For any consumer surplus regular distribution and an ex-ante allocation probability $q \in [0, 1]$, the ratio of the optimal surplus $V(q)$ to the optimal consumer surplus $U(q)$ is at most $1 - \ln q$.*

Proof. Conditional on generating a surplus $V(q)$, we want to compute the distribution that minimizes consumer surplus $U(q)$.

$$\begin{aligned} U(q) &= \int_0^q u(t) dt. \\ V(q) &= U(q) + qv(q) \\ &= \int_0^q u(t) dt - q \int_q^1 v'(t) dt \\ &= \int_0^q u(t) dt + q \int_q^1 \frac{u(t)}{t} dt. \end{aligned} \tag{1}$$

We substitute $u(t) = -tv'(t)$ in the last equation.

Simultaneously, we enforce $u(t) \geq 0$, and $u(t)$ is non-increasing (since the distribution is consumer surplus regular), $v(1) = 0$ and $V(0) = 0$.

52:10 Combinatorial Pen Testing (Or Consumer Surplus of Deferred-Acceptance Auctions)

Finding the minimum consumer surplus distribution reduces to the following program.

$$\min \int_0^q u(t) dt$$

subject to

$$\int_0^q u(t) dt + q \int_q^1 \frac{u(t)}{t} dt = V(q)$$

$$u(t) \geq 0, u(t) \text{ is monotonously non-increasing}$$

$$v(1) = 0, V(0) = 0$$

The constraints in the last line can be enforced after solving for u .

Consumer surplus regularity dictates $u(t) \geq u(q)$ for $t \leq q$. Hence, minimizing $\int_0^q u(t) dt$ would correspond to setting $u(t) = u(q)$ for $t \leq q$. Similarly, $u(t) \leq u(q)$ for $t \geq q$. Minimizing $\int_0^q u(t) dt$ would mean minimizing $V(q) - q \int_q^1 \frac{u(t)}{t} dt$ which is achieved by setting $u(t) = u(q)$ for $t \geq q$. Let $u(0) = u(t) = u(1) = u$. For a constant marginal price-posting consumer surplus curve,

$$U(q) = \int_0^q u(t) dt = q u$$

$$V(q) = \int_0^q u(t) dt + q \int_q^1 \frac{u(t)}{t} dt = [q - q \ln q] u$$

Thus, for any consumer surplus regular distribution, $\frac{V(q)}{U(q)} \leq [1 - \ln q]$. The constant marginal consumer surplus curve that achieves this ratio corresponds to the exponential distribution. ◀

► **Lemma 10.** *Let the consumer surplus curves U, \hat{U} satisfy $\hat{U}(0) = 0$ and $\hat{U}(t) \geq U(t)$ for all $t \in [0, 1]$. Then, for the corresponding surplus curves, $\hat{V}(q) \geq V(q)$.*

Proof. We rewrite equation (1) as follow.

$$\begin{aligned} V(q) &= \int_0^q u(t) dt + q \int_q^1 \frac{u(t)}{t} dt \\ &= U(q) + q \left[\frac{U(t)}{t} \right]_{t=q}^{t=1} + q \int_q^1 \frac{U(t)}{t^2} dt \\ &= qU(1) + q \int_q^1 \frac{U(t)}{t^2} dt \end{aligned} \tag{2}$$

The second equality is obtained through integrating by parts. Increasing U pointwise clearly results in an increase in V . ◀

Proof of Theorem 8. The optimal consumer surplus for an ex-ante allocation probability q equals $\bar{U}(q)$.

Consider the distribution with price-posting consumer surplus curve \bar{U} . Let \bar{V} be the price-posting surplus curve corresponding to \bar{U} . From Lemma 10 and Lemma 9,

$$\frac{V(q)}{\bar{U}(q)} \leq \frac{\bar{V}(q)}{\bar{U}(q)} \leq 1 - \ln q$$

The second inequality holds from Lemma 9, since \bar{V} is the price-posting surplus curve of a consumer surplus regular distribution (we made \bar{U} concave through ironing). ◀

See the full online version for a brief discussion on the surplus generated by the single-agent optimal consumer surplus auction.

3.2 Multi-Agent Environments

In this section, we convert the single-agent bound from Theorem 8 into an upper bound on the ratio between the optimal surplus and optimal consumer surplus (aka the optimal omniscient approximation) for multi-agent environments. We outline our approach before delving into the proof. Consider the interim allocation rule y_i for agent i in the surplus optimal mechanism OPT^V . Without loss of generality, assume $y_i(1) = 0$. The expected surplus for agent i equals

$$E_{q \in U[0,1]}[y_i(q) v_i(q)] = [y_i(q) V_i(q)]_{q=0}^{q=1} + E_{q \in U[0,1]}[-y_i'(q) V(q)] = E_{q \in U[0,1]}[-y_i'(q) V(q)]$$

This equality comes from integrating by parts ($y_i(q)V_i(q)$ vanishes at both, $q = 0$ and 1 , since $V_i(0) = 0$ and $y_i(1) = 0$). Suppose there exists a constant α such that $V_i(q) \leq \alpha \bar{U}_i(q)$ for all $q \in [0, 1]$, then,

$$E_{q \in U[0,1]}[-y_i'(q) V(q)] \leq E_{q \in U[0,1]}[-y_i'(q) \alpha \bar{U}(q)] = \alpha \cdot E_{q \in U[0,1]}[y_i(q) \bar{u}_i(q)]$$

The above inequality connects the expected surplus to the expected ironed consumer surplus of the interim allocation rule y_i . Since $\frac{d}{dq} y_i$ need not be equal to 0 whenever $U(q) \neq \bar{U}(q)$, $E_{q \in U[0,1]}[y_i(q) \bar{u}_i(q)] \neq E_{q \in U[0,1]}[y_i(q) u_i(q)]$. But, Theorem 6 shows the existence of a mechanism \bar{y} with expected consumer surplus of each agent i equal to $E_{q \in U[0,1]}[y_i(q) \bar{u}_i(q)]$. Thus, the expected consumer surplus of \bar{y} is an α -approximation to the optimal surplus. The consumer surplus optimal mechanism OPT^U will only do better.

Unfortunately, as $q \rightarrow 0$, the bound from Theorem 8 of $\frac{V_i(q)}{\bar{U}_i(q)} \leq (1 - \ln q)$ diverges to ∞ . Thus, we need to show that the loss from ignoring the small (strong) quantiles is not much. Suppose we find a mechanism \hat{y} with interim allocation \hat{y}_i for agent i , near optimal surplus, and $\hat{y}_i'(q) = 0$ for $q \in [0, \epsilon]$. Then, $y_i'(q) \times V(q) = \hat{y}_i'(q) \times U(q) = 0$ in this region. For $q \in [\epsilon, 1]$, we know $V(q) \leq (1 - \ln \epsilon) \times \bar{U}(q)$. Thus,

$$(1 - \ln \epsilon) \times \text{OPT}^U \geq \text{surplus}(\hat{y})$$

Hartline and Taggart [17] give the construction of such a mechanism \hat{y} , originally intended to design good mechanisms with only sample access to the distributions of bidders' bids.

► **Definition 11** (Hartline and Taggart [17]; ϵ -buffering rule). *Given an allocation rule $\vec{y} = (y_1, y_2, \dots, y_n)$ and a quantile $\epsilon \in [0, 1]$, the ϵ -buffering rule for \vec{y} simulates y with quantiles transformed on each agent as follows:*

- *Top inflate: for any $q_i \in [0, \epsilon]$, return 0*
- *For any $q_i \in [\epsilon, 1 - \epsilon]$, return $\frac{q_i - \epsilon}{1 - 2\epsilon}$*
- *Bottom deflate: for any $q_i \in [1 - \epsilon, 1]$, return 1*

In essence, the top inflate branch makes the ϵ -buffering rule treat all agents with a small quantile as if they had quantile 0 and thus, $y_i'(q) = 0$ for $q \in [0, \epsilon]$ (the bottom deflate branch is unnecessary for the analysis of monotone payoff curves like surplus V and consumer surplus U ; see Remark 15).

► **Theorem 12** (Hartline and Taggart [17]). *Let \hat{y} be the ϵ -buffering rule of the surplus optimal mechanism OPT^V . Then, $\text{OPT}^V \leq \frac{1}{(1 - \frac{\epsilon}{1-\epsilon})(1-\epsilon)(1-2n\epsilon)} \times \text{surplus}(\hat{y})$.*

► **Theorem 13.** *Consider an n -agent environment with an arbitrary feasibility constraint. The optimal omniscient approximation is at most $\frac{1 - \ln \epsilon}{(1 - \frac{\epsilon}{1-\epsilon})(1-\epsilon)(1-2n\epsilon)}$.*

52:12 Combinatorial Pen Testing (Or Consumer Surplus of Deferred-Acceptance Auctions)

Proof. The proof follows by combining Theorem 12 with the discussion above. Let \hat{y} be the ϵ -buffering rule of the surplus optimal mechanism OPT^V . Then,

$$\text{OPT}^V \leq \frac{1}{(1 - \frac{\epsilon}{1-\epsilon})(1 - \epsilon)(1 - 2n\epsilon)} \times \text{surplus}(\hat{y}) \leq \frac{1 - \ln \epsilon}{(1 - \frac{\epsilon}{1-\epsilon})(1 - \epsilon)(1 - 2n\epsilon)} \times \text{OPT}^U$$

◀

► **Corollary 14.** *In n -agent environments with an arbitrary feasibility constraint, the optimal omniscient approximation $\zeta(n) \leq (1 + o(1)) \ln n$.*

Proof. Setting $\epsilon = \frac{1}{n \ln n}$ in Theorem 13, we get

$$\begin{aligned} \zeta(n) &\leq \frac{1 - \ln \frac{1}{n \ln n}}{(1 - \frac{\frac{1}{n \ln n}}{1 - \frac{1}{n \ln n}})(1 - \frac{1}{n \ln n})(1 - 2n \cdot \frac{1}{n \ln n})} \\ &= \frac{1 + \ln n + \ln \ln n}{(1 - \frac{1}{n \ln n - 1})(1 - \frac{1}{n \ln n})(1 - \frac{2}{\ln n})} \\ &= \frac{n \ln n - 1}{n \ln n - 2} \cdot \frac{n \ln n}{n \ln n - 1} \cdot \frac{\ln n - 2}{\ln n} \cdot \frac{1 + \ln n + \ln \ln n}{\ln n} \times \ln n \\ &= (1 + o(1)) \ln n \end{aligned}$$

◀

► **Remark 15.** The ϵ -buffering rule of Hartline and Taggart [17] can be implemented without a bottom deflate branch. This would strengthen the approximation guarantee to $\frac{1}{(1-\epsilon)(1-n\epsilon)}$. However, substituting $\epsilon = \frac{1}{n \ln n}$ will not result in a bound tighter than $(1 + o(1)) \ln n$.

► **Remark 16.** Fotakis et al. [13] show that the optimal omniscient approximation as a function of the number of outcomes m in the feasibility constraint is at most $\frac{2}{\ln 2}(1 + o(1)) \ln m$. Corollary 14 improves their bound whenever the feasibility constraint has cardinality at least \sqrt{n} . Note that all downward-closed constraints have at least n feasible outcomes. For such feasibility constraints, Corollary 14 yields an exponential improvement. As the number of feasible outcomes become exponential in n , their bound declines to $O(n)$.

► **Remark 17.** The upper bound from Corollary 14 matches the lower bound of Qiao and Valiant [23]. In the single-item environment with values of agents drawn IID from the exponential distribution, they show that the optimal surplus is $H_n = (1 + o(1)) \ln n$ times the optimal consumer surplus, where H_n is the n^{th} harmonic number.

3.3 The k -Identical Goods Environment

Theorem 13 and Corollary 14 give an over-arching bound for the optimal omniscient approximation $\zeta(n)$ that holds in any environment. However, specific environments admit better approximations. For instance, we will give a tighter bound for $\zeta(n)$ for the k -identical goods environment, where any set of at most k -agents can be allocated simultaneously.

► **Theorem 18.** *In the n -agent k -identical goods environment, the optimal omniscient approximation satisfies*

1. $\zeta(n) \leq (1 + o(1)) \ln \frac{n}{k}$ when $k = o(n)$, and
2. $\zeta(n) \leq 2.27(1 + o(1))(0.577 + \ln \frac{n}{k})$, when $k = \Theta(n)$.

To prove Theorem 18, we make use of the following lemma, whose proof we defer to the end of the section.

► **Lemma 19.** *Consider the special case with n agents and k identical goods. The optimal omniscient approximation $\zeta(n) \leq \frac{1}{(1 - \frac{1}{\sqrt{2\pi k}})} \cdot (1 - \ln \epsilon)(1 + [\frac{n}{k} - 1]\epsilon)$ for any $\epsilon > 0$.*

Lemma 19 is particularly useful when $k = \Theta(n)$. For a constant ϵ , $\frac{1}{(1 - \frac{1}{\sqrt{2\pi k}})}$ is $(1 + o(1))$ while $(1 - \ln \epsilon)(1 + [\frac{n}{k} - 1] \epsilon)$ is a constant.

Proof of Theorem 18. We show the theorem separately for the three cases, when $k = \Theta(1)$, $k = \omega(1)$ but $o(n)$, and $k = \Theta(n)$.

- $k = \Theta(1)$: The proof is a direct consequence of Corollary 14, where we show $\zeta(n) \leq (1 + o(1)) \ln n = (1 + o(1)) \ln \frac{n}{k}$. The equality follows since k is a constant.
- $k = \omega(1)$, but $o(n)$: By setting $\epsilon = \frac{1}{\frac{n}{k}(1 + \ln \frac{n}{k})}$ in Lemma 19, we get

$$\zeta(n) \leq \left(\frac{1}{1 - \frac{1}{\sqrt{2\pi k}}}\right) \times \left(1 + \frac{1}{1 + \ln \frac{n}{k}} \left[1 - \frac{k}{n}\right]\right) \times \left(1 + \ln \frac{n}{k} + \ln\left(1 + \ln \frac{n}{k}\right)\right)$$

The first term is $(1 + o(1))$ since $k = \omega(1)$. The second and the third terms are $(1 + o(1))$ and $(1 + o(1)) \ln \frac{n}{k}$ respectively, since $k = o(n)$. Summarizing, $\zeta(n) \leq (1 + o(1)) \ln \frac{n}{k}$. The above two cases conclude the proof of the first part of the theorem. The next case clinches the proof of the second part.

- $k = \Theta(n)$: We set $\epsilon = \frac{1}{\frac{n}{k}(1 + \ln \frac{n}{k})}$ as in the previous case. Once again $\left(\frac{1}{1 - \frac{1}{\sqrt{2\pi k}}}\right) = (1 + o(1))$ since $k = o(1)$. We plot the ratio between our upper bound and the lower bound $(0.577 + \ln \frac{n}{k})$ from Lemma 20 as a function of $\frac{k}{n}$ in Figure 1a to observe that the ratio is at most 2.27, exactly as needed. We skip a more detailed numerical analysis substantiating this observation. ◀

We modify the lower bound from Qiao and Valiant [23] sketched in Remark 17 to get an analogous result for the k -identical goods environment.


► **Lemma 20.** *Consider the n -agent k -identical goods environment with values drawn IID from the exponential distribution with mean 1. Assume, for convenience, that $\frac{n}{k}$ is an integer. The optimal omniscient approximation $\zeta(n) \geq H_{\frac{n}{k}} \geq (0.577 + \ln \frac{n}{k})$.*

Proof. The exponential distribution with mean 1 has a constant marginal consumer surplus curve $u(q) = 1$, and thus, the optimal consumer surplus mechanism allocates to some set of k agents to get an expected consumer surplus equal to k .

Now, consider the following mechanism. Group agents into k batches of $\frac{n}{k}$ agents each, and award the agent with the highest value in each group. Each group is now identical to a single item environment with $\frac{n}{k}$ agents, and thus, generates an expected surplus of $H_{\frac{n}{k}}$ (see Remark 17 on the example sketched in [23]). The k batches together produce an expected surplus equal to $k \cdot H_{\frac{n}{k}}$. The optimal mechanism creates a larger surplus and hence, $\zeta(n)$ is at least $H_{\frac{n}{k}} \geq (0.577 + \ln \frac{n}{k})$. ◀

Thus, the upper bound from Theorem 18 is tight up to $(1 + o(1))$ terms when $k = o(n)$ and up to a multiplicative $2.27(1 + o(1))$ when $k = \Theta(n)$. The best known bound for the optimal omniscient approximation is due to Hartline and Roughgarden [18], where they show $\zeta(n) \leq \frac{2}{\ln 2} (\ln 2 + \ln \frac{n}{k}) \approx 2.88(\ln 2 + \ln \frac{n}{k})$. Apart from closing the $\frac{2}{\ln 2}$ gap between their upper bound and the lower bound when $k = o(n)$, our bound is an improvement by a factor of at least $1.36(1 - o(1))$ when $k = \Theta(n)$ (Figure 1b).

We conclude the section with the proof of Lemma 19. We will adopt an approach similar to the proof of Theorem 13. Increasing the allocation to the agents with a small quantile so that $y_i(q_i) = y_i(0)$ for $q_i \in [0, \epsilon]$ vastly increases the consumer surplus of i without noticeably decreasing the surplus of the mechanism. We consider online posted-price mechanisms with the constraint $y'_i(q_i) = 0$ whenever $q_i \leq \epsilon$ for all agents i . Posted-price mechanisms are known to achieve a near optimal surplus and we bound the loss in surplus due to the additional constraint.



■ **Figure 1** We compare our upper bound against the lower bound and the best known upper bound when $k = \Theta(n)$. Figure (a) plots the ratio between our upper bound from Theorem 18, normalized by $(1 + o(1))$ and our lower bound from Lemma 20, $(0.577 + \ln \frac{n}{k})$. The maximum value of the curve is less than 2.27, and hence, our upper bound can potentially be improved by a factor of at most $2.27(1 + o(1))$. Figure (b) plots the ratio between the current best known bound $\frac{2}{\ln 2} (\ln 2 + \ln \frac{n}{k})$ and our bound from Theorem 18, again normalized by $(1 + o(1))$. The minimum value taken by the ratio is greater than 1.36, and hence, our bound improves the best known bound by a factor of at least $1.36(1 - o(1))$.

Consistent with prior literature on online sequential price-posting [6, 26] we will approximate the performance of the mechanism against the optimal ex-ante relaxed mechanism EAR. In the ex-ante relaxation, the feasibility constraint is relaxed to bind ex-ante rather than ex-post. In other words, the expected number of agents served by the mechanism must be at most k . The optimal ex-ante mechanism obtains a larger surplus than OPT^V , since OPT^V is also a feasible ex-ante mechanism. Note that, if agent i is being allocated with probability q_i by EAR, it would allocate to the q_i smallest quantiles (quantiles corresponding to the largest values) without violating the ex-ante constraint, thereby getting an expected surplus $V_i(q_i)$ from agent i .

Let Δ be the set of all ex-ante feasible allocations $\vec{q} = (q_1, \dots, q_n)$ for selling up to k items such that $q_1, \dots, q_n \in [0, 1]$ and $\sum_{i=1}^n q_i \leq k$. Let $\text{EAR}(\vec{q})$ be the ex-ante relaxed mechanism that allocates to the q_i smallest quantiles of agent i . Note that $\text{EAR}(\vec{q})$ is well-defined even if $\vec{q} \notin \Delta$ as long as $q_1, \dots, q_n \in [0, 1]$. We will abuse notation to denote the surplus from $\text{EAR}(\vec{q})$ by $\text{EAR}(\vec{q})$. Thus, $\text{EAR} = \max_{\vec{q} \in \Delta} \text{EAR}(\vec{q})$.

Yan [26] showed that the following sequential posted-price mechanism $\text{SP}(\vec{q})$ has a surplus comparable to $\text{EAR}(\vec{q})$ for all $\vec{q} \in \Delta$.

1. Order agents in decreasing order of $\frac{V_i(q_i)}{q_i}$, the expected surplus conditioned on selling to agent i .
2. In the order above, offer agent i a price $v_i(q_i)$ until all n agents interact with the mechanism or until all k goods are sold.

As usual, we will abuse notation to denote the surplus of $\text{SP}(\vec{q})$ by $\text{SP}(\vec{q})$.

► **Theorem 21** (Yan [26]). For $\vec{q} \in \Delta$, $\frac{1}{(1 - \frac{1}{\sqrt{2\pi k}})} \times \text{SP}(\vec{q}) \geq \text{EAR}(\vec{q})$.

We propose a sequential-posted-price-mechanism with near optimal surplus and interim allocation rules satisfying the additional properties outlined above. For a vector $\vec{q} \in \Delta$, let the ϵ -inflated sequential posted-price mechanism $\text{SP}(\vec{q}, \epsilon)$ denote $\text{SP}(\vec{p})$, where $p_i = \max\{q_i, \epsilon\}$.

Observe that if $\text{SP}(\vec{q}, \epsilon)$ has an unsold good while interacting with agent i , it always sells the good if the agent has a quantile $q_i \leq \epsilon$. Thus, the interim allocation rule y_i is a constant function in $[0, \epsilon]$ and $y'_i(q) = 0$ in this range. Next, we argue that $\text{SP}(\vec{q}, \epsilon)$ approximates the surplus of $\text{EAR}(\vec{q})$.

► **Lemma 22.** For all $\vec{q} \in \Delta$, $\text{EAR}(\vec{q}) \leq \frac{(1 + [\frac{n}{k} - 1]\epsilon)}{(1 - \frac{1}{\sqrt{2\pi k}})} \text{SP}(\vec{q}, \epsilon)$

Proof. The proof proceeds in two steps. First, for \vec{p} with $p_i = \max\{q_i, \epsilon\}$, we show that \vec{p} is not too far away from the set of ex-ante feasible allocations Δ . We then show that the projection \vec{t} of \vec{p} back onto Δ satisfies $\text{EAR}(\vec{q}) \leq \text{EAR}(\vec{p}) \leq \frac{k + (n-k)\epsilon}{k} \times \text{EAR}(\vec{t})$ and $\text{SP}(\vec{t}) \leq \text{SP}(\vec{p}) = \text{SP}(\vec{q}, \epsilon)$.

Step 1. $\sum_{i=1}^n p_i \leq k + (n-k)\epsilon$.

Without loss of generality, we assume $\sum_{i=1}^n q_i = k$. Increasing q_i so that their sum equals k would not decrease the value of $\sum_{i=1}^n p_i$. The claim reduces to showing the maximum value of $\sum_{i=1}^n (p_i - q_i)$ is at most $(n-k)\epsilon$.

Note that $p_i - q_i \leq \epsilon$, equality holding exactly when $q_i = 0$. Thus, whenever $q_i < \epsilon$, decreasing them to zero increases the value of $\sum_{i=1}^n (p_i - q_i)$. However, this would contradict $\sum_{i=1}^n q_i = k$. For $q_i \geq \epsilon$, $p_i - q_i = 0$. Increasing q_i to 1 whenever $q_i \geq \epsilon$ does not change the value of $\sum_{i=1}^n (p_i - q_i)$, but allows other values q_j to be set to zero without violating $\sum_{i=1}^n q_i = k$. A simple greedy argument would suggest that $\sum_{i=1}^n (p_i - q_i)$ is maximized when $q_i = 1$ for $1 \leq i \leq k$ and $q_i = 0$ when $i > k$, in which case, $\sum_{i=1}^n (p_i - q_i) < (n-k)\epsilon$.

Step 2. Order (and re-index) agents in decreasing order of $V_i(p_i)$. Consider the ex-ante allocation $\frac{k}{k+(n-k)\epsilon} \cdot \text{EAR}(\vec{p})$ that offers agent i with $\frac{k}{k+(n-k)\epsilon}$ fraction of a good whenever $q_i \leq p_i$. This has a surplus $\frac{k}{k+(n-k)\epsilon} \cdot \text{EAR}(\vec{p}) = \sum_{i=1}^n \frac{k}{k+(n-k)\epsilon} \cdot V_i(p_i)$ and allocates up to k units, since $\sum_{i=1}^n p_i \leq k + (n-k)\epsilon$. Let \vec{t} be an alternate ex-ante feasible allocation that allocates $t_i = p_i$ for agents $1, \dots, m$ for some $m \leq n$ until all k items are allocated. \vec{t} redistributes goods from agents with a smaller surplus $V_i(p_i)$ allocated by $\frac{k}{k+(n-k)\epsilon} \cdot \text{EAR}(\vec{p})$ to agents with a larger $V_i(p_i)$ to get a surplus

$$\text{EAR}(\vec{t}) = \sum_{i=1}^m V_i(p_i) \geq \sum_{i=1}^n \frac{k}{k+(n-k)\epsilon} \cdot V_i(p_i) = \frac{k}{k+(n-k)\epsilon} \cdot \text{EAR}(\vec{p})$$

Step 3. $\text{SP}(\vec{t}) \leq \text{SP}(\vec{p})$

$\text{SP}(\vec{t})$ allocates the first m agents with probability p_i and allocates the others with probability 0. $\text{SP}(\vec{p})$ allocates the first m agents with probability p_i and further allocates any unsold good to the others. Thus, $\text{SP}(\vec{t}) \leq \text{SP}(\vec{p})$.

Summarizing,

$$\begin{aligned} \text{EAR}(\vec{q}) &\leq \text{EAR}(\vec{p}) \\ &\leq \frac{k+(n-k)\epsilon}{k} \times \text{EAR}(\vec{t}) = (1 + [\frac{n}{k} - 1]\epsilon) \times \text{EAR}(\vec{t}) \\ &\leq \frac{(1 + [\frac{n}{k} - 1]\epsilon)}{(1 - \frac{1}{\sqrt{2\pi k}})} \times \text{SP}(\vec{t}) \\ &\leq \frac{(1 + [\frac{n}{k} - 1]\epsilon)}{(1 - \frac{1}{\sqrt{2\pi k}})} \times \text{SP}(\vec{p}) = \frac{(1 + [\frac{n}{k} - 1]\epsilon)}{(1 - \frac{1}{\sqrt{2\pi k}})} \times \text{SP}(\vec{q}, \epsilon) \end{aligned}$$

The third inequality follows from Theorem 21 ◀

Proof of Lemma 19. Let $\vec{q} \in \Delta$ be such that $\text{EAR} = \text{EAR}(\vec{q})$. From the review on sequential-price-posting and Lemma 22,

$$\text{OPT}^V \leq \text{EAR} = \text{EAR}(\vec{q}) \leq \frac{(1 + [\frac{n}{k} - 1] \epsilon)}{(1 - \frac{1}{\sqrt{2\pi k}})} \times \text{SP}(\vec{q}, \epsilon)$$

Further, $\text{SP}(\vec{q}, \epsilon)$ has a constant interim allocation rule for agent i when its quantile q_i is in the range $[0, \epsilon]$. Arguments identical to the proof of Theorem 13 can be used to conclude

$$\zeta(n) = \frac{\text{OPT}^V}{\text{OPT}^U} \leq \frac{(1 + [\frac{n}{k} - 1] \epsilon)}{(1 - \frac{1}{\sqrt{2\pi k}})} \cdot \frac{\text{SP}(\vec{q}, \epsilon)}{\text{OPT}^U} \leq \frac{1}{(1 - \frac{1}{\sqrt{2\pi k}})} \cdot (1 - \ln \epsilon) (1 + [\frac{n}{k} - 1] \epsilon) \quad \blacktriangleleft$$

4 Pen Testing Corollaries from Deferred-Acceptance Mechanisms

There are many environments in which deferred-acceptance mechanisms are known to be good. By Theorem 1, these imply good pen testing algorithms (up to an additional $\zeta(n)$ factor in the omniscient approximation, i.e, the ratio of the performances of the omniscient algorithm that knows the amount of ink in each pen and our proposed pen testing algorithm). In this section, we discuss a few notable examples. A summary of these results was given previously in Table 1.

Deferred-acceptance mechanisms that achieve the ex-post surplus optimal outcome are known for various feasibility constraints. While the simple auction that uniformly increases the price until exactly k bidders remain active is surplus optimal in the k -identical goods environment, Milgrom and Segal [21] and Bikhchandani et al. [4] generalize the mechanism for matroid feasibility constraints. We get the following corollaries from the above auctions.

► **Corollary 23.** *For a pen testing environment with a matroid feasibility constraint, there exists a pen testing algorithm with an omniscient approximation ratio $(1 + o(1)) \ln n$.*

► **Corollary 24.** *When choosing any k of the n pens is feasible, there exists a pen testing algorithm with an omniscient approximation ratio $(1 + o(1)) \ln \frac{n}{k}$ when $k = o(n)$ and $2.27(1 + o(1))(0.577 + \ln \frac{n}{k})$ when $k = \Theta(n)$.*

Dütting et al. [8] initiated the study of prior-free deferred-acceptance approximation mechanisms in environments where deferred-acceptance mechanisms are known to be suboptimal. Their bounds can be improved in settings like ours where prior distributions of the agents' values are known. For general downward-closed constraints with a prior distribution on values, Feldman et al. [11] give a $O(\log \log m)$ approximation to the optimal surplus, where m is the number of maximal feasible sets. Note that $m \leq 2^n$ (every subset of agents might be feasible) and hence, $\log \log m$ is at most $\log n$. Thus, we have a poly-logarithmic approximate pen testing algorithm for any downward-closed feasibility environment.

Note, however, that the pen testing model is inherently downward-closed. Specifically, suppose an algorithm wanted to select a set $P \subseteq \bar{P}$, but P is not feasible while \bar{P} is. Since all pens have non-negative residual ink, even the ones that failed their tests, there is no loss in selecting \bar{P} instead of the set P . Thus, near optimal pen testing algorithms for downward-closed constraints can be extended to give near optimal algorithms for general combinatorial constraints, giving the same performance guarantee as the downward-closed environment.

► **Corollary 25.** *For any combinatorial pen testing environment, there exists a pen testing algorithm with an omniscient approximation ratio $O(\log n \log \log m) = O(\log^2 n)$.*

► **Remark 26.** Note that this reduction cannot be used to design deferred-acceptance mechanisms for general feasibility constraints. Allocating to agents that have dropped out of the auction can make the mechanism non-truthful (i.e., staying active till the price reaches the agent's value and dropping out subsequently might not be in the agent's best interests).

The bound for downward-closed and general combinatorial environments can be improved upon in special cases; for example, in the full online version, we give a natural deferred-acceptance mechanism for knapsack constraints that is a 2-approximation. For this knapsack problem, the agents have sizes along with values and feasible subsets are precisely those with a total size at most the knapsack capacity.

► **Corollary 27.** *For a pen testing environment with a knapsack feasibility constraint, there exists a pen testing algorithm with an omniscient approximation ratio $2(1 + o(1)) \ln n$.*

Online pricing mechanisms are a special case of deferred-acceptance mechanisms; thus, Theorem 1 converts online pricing mechanisms with good surplus into good online pen testing algorithms. For online pen testing problems, every pen can be tested exactly once and must either be selected or discarded immediately after testing and before testing another pen. The algorithm might have the ability to determine the order of testing pens (the sequential version) or the order might be adversarially determined (the oblivious version).

For the sequential pricing problem, Chawla et al. [6] gave an $\frac{e}{e-1}$ -approximation mechanism for the single-item environment. Yan [26] connected the sequential pricing problem to the correlation gap and generalized the $\frac{e}{e-1}$ -approximation to matroid environments. Both of these papers considered optimizing revenue, but the results can be easily adapted to optimize surplus.

► **Corollary 28.** *In the online sequential pen testing problem with a matroid feasibility constraint, there exists a pen testing algorithm that achieves an omniscient approximation ratio $\frac{e}{e-1} \cdot (1 + o(1)) \ln n$.*

Hajiaghayi et al. [16] and Chawla et al. [6] adapt the prophet inequality of Samuel-Cahn [25] to give a non-adaptive 2-approximation for the oblivious price posting problem with 1 good. Kleinberg and Weinberg [19] extend this 2-approximation to matroid environments and any arrival order of agents but with adaptive prices.

► **Corollary 29.** *In the online oblivious pen testing problem to select a single pen, there exists a non-adaptive pen testing algorithm that achieves an omniscient approximation ratio $2(1 + o(1)) \ln n$.*

► **Corollary 30.** *In the online oblivious pen testing problem with a matroid feasibility constraint, there exists a pen testing algorithm that achieves an omniscient approximation ratio $2(1 + o(1)) \ln n$.*

The bounds of Corollary 28 and Corollary 29 improve on the online pen testing bounds of Qiao and Valiant [23].

5 The Online IID Environment

Consider the special case of the online pen testing problem where the ink levels are drawn independently from the same distribution. We move away from the reduction framework of Theorem 1 to obtain a better bound for this environment.

► **Theorem 31.** *In online single-item environments with IID agents, there exists a price-posting strategy with an omniscient approximation at most $(1 + o(1)) \ln n$. Equivalently, in the online pen testing problem with IID ink levels, there exists an algorithm that achieves $\pi(n) \leq (1 + o(1)) \ln n$.*

First, note that the above theorem is tight. In Remark 17, we saw an instance (originally from [23]) with n IID agents whose values are drawn from the exponential distribution with consumer surplus only an H_n approximation to the optimal surplus. The optimal online consumer surplus cannot do better. Hence, the $(1 + o(1)) \ln n$ approximation ratio in Theorem 31 is tight.

Before proving Theorem 31, we briefly discuss the guarantee obtained through the reduction framework of Theorem 1. For the online environment with IID agents, it is known that $\gamma(n) = \frac{e}{e-1}$ [6]. Thus, the framework would give us an $\frac{e}{e-1} \cdot (1 + o(1)) \ln n$ bound. The direct analysis below gives a tight $H_n + 1 = (1 + o(1)) \ln n$ bound up to an additive factor 1.

Proof of Theorem 31. Similar to the proof of Theorem 8, we will phrase our program as an optimization problem and compute the optimal solution.

The surplus optimal auction always allocates the good to the agent with the highest value (smallest quantile). Let v_{max} be the random variable denoting the value of the winner.

$$Pr(v_{max} \leq v(t)) = (1 - t)^n$$

Thus, the expected surplus of the surplus optimal mechanism equals

$$\begin{aligned} \int_0^1 v(t) \times n(1-t)^{n-1} dt &= \int_0^1 \int_t^1 \frac{u(r)}{r} dr \times n(1-t)^{n-1} dt \\ &= \int_0^1 \frac{u(t)}{t} dt - \int_0^1 \frac{u(t)}{t} (1-t)^n dt \\ &= \int_0^1 \frac{1 - (1-t)^n}{t} u(t) dt \\ &= U(1) + \int_0^1 \frac{1 - (1-t)^n - nt(1-t)^{n-1}}{t^2} U(t) dt \end{aligned}$$

The first equality follows by integrating from $u(t) = -tv'(t)$ and $v(1) = 0$. We integrate by parts for the second and fourth equation.

Next, we look at the consumer surplus through price-posting. For a price $v(q)$, we allocate the good if at least one agent has a quantile in $[0, q]$. Conditioned on selling the good, we get an expected consumer surplus of $\frac{U(q)}{q}$. Thus, the expected consumer surplus equals

$$\frac{1 - (1-q)^n}{q} \times U(q)$$

As before, we first show Theorem 31 for consumer surplus regular distributions. Without loss of generality, assume that the maximum achievable consumer surplus through anonymous price posting is at most 1. We want to find the consumer surplus regular distribution with the largest surplus.

$$\max U(1) + \int_0^1 \frac{1 - (1-t)^n - nt(1-t)^{n-1}}{t^2} U(t) dt$$

subject to

$$\frac{1 - (1-q)^n}{q} \times U(q) \leq 1 \text{ for all } q \in [0, 1]$$

U is concave

First, observe that $C(q) = \frac{q}{1-(1-q)^n}$ is convex (the proof for convexity is deferred to the full online version). We want to fit a concave curve U under the convex curve C . Let \overline{C}_q be the tangent to C at q . \overline{C}_q is a feasible solution for U for all $\hat{q} \in [0, 1]$. Further, given that U touches the curve C at \hat{q} , $U(t) \leq \overline{C}_{\hat{q}}(t)$ for all $t \in [0, 1]$ (from the concavity of U). Also, it is clear that the objective increases with a pointwise increase in U . Thus, we conclude that the optimal solution to the program is achieved at $U = \overline{C}_{\hat{q}}$ at some $\hat{q} \in [0, 1]$. In particular, the worst-case ratio occurs when U is a straight-line.

Without loss of generality, we can normalize the slope to 1. Let $U(q) = q + a$. The surplus equals

$$1 + a + \int_0^1 \frac{1 - (1-t)^n - nt(1-t)^{n-1}}{t} dt + a \times \int_0^1 \frac{1 - (1-t)^n - nt(1-t)^{n-1}}{t^2} dt = H_n + an$$

See the full online version for the proof of the above equality.

By posting a price $v(q)$, we generate a consumer surplus

$$\frac{1 - (1-q)^n}{q} \times U(q) = (1 + (1-q) + \dots + (1-q)^{n-1}) \times (q + a)$$

If $a \geq \frac{1}{n}$, the ratio between the surplus and the consumer surplus is at most $H_n + 1$ by setting $q = 0$. If $a < \frac{1}{n}$, the surplus is less than $H_n + 1$, and by setting $q = 1$, the consumer surplus is at least 1. Thus, the worst-case ratio between surplus and consumer surplus is at most $H_n + 1 = (1 + o(1)) \ln n$.

We now show the result for all distributions. We rewrite the optimal surplus as follows.

$$\int_0^1 v(t) \times n(1-t)^{n-1} dt = \left[V(t) \times n(1-t)^{n-1} \right]_{t=0}^{t=1} + \int_0^1 V(t) \times n(n-1)(1-t)^{n-2} dt$$

The equality follows through integration by parts. From the above expression, it can be concluded that the optimal surplus increases with a pointwise increase in V .

Let \overline{V} be the price-posting surplus curve of the distribution with a price-posting consumer surplus curve \overline{U} . \overline{U} is pointwise larger than U , and hence, from Lemma 10, \overline{V} is pointwise larger than V . Hence, the optimal surplus is larger in the distribution with price-posting consumer surplus curve \overline{U} . However, from Theorem 5, the consumer surplus from posting prices is identical to both the distributions. Thus, the consumer surplus irregular distribution has a better ratio between surplus and consumer surplus than the consumer surplus regular distribution. Thus, the ratio between optimal surplus to optimal consumer surplus is at most $(1 + o(1)) \ln n$ for all distributions. ◀

6 Conclusion

In this section, we discuss the various advantages and disadvantages of using the reduction framework of Theorem 1. Connecting the pen testing problem to auction theory and building an easy-to-use reduction framework enable us to design near optimal pen testing algorithms for any general feasibility constraint. We are able to construct algorithms that are only a constant factor away from the lower bound described in Qiao and Valiant [23] for various feasibility constraints like matroids ($\gamma(n) = 1$) and knapsack ($\gamma(n) = 2$). However, there is a scope for improvement on the following two fronts:

1. obtaining tight approximation guarantees, and
2. obtaining approximation guarantees under other models of access to the prior distributions.

Tight Approximation Guarantees. The approximation ratio of any pen testing algorithm is at least the optimal omniscient approximation $\zeta(n)$. This follows since the optimal consumer surplus in the underlying auction environment is an upper bound on the performance of the optimal pen testing algorithm. The framework matches this bound up to a multiplicative factor of the approximation $\gamma(n)$ of deferred-acceptance mechanisms. Reducing this gap for various feasibility environments stands as an open problem.

For instance, consider the online environments discussed in the paper, where $\zeta(n) = (1 + o(1)) \ln n$. Our framework yielded an omniscient approximation $2(1 + o(1)) \ln n$ in the oblivious version ($\gamma(n) = 2$) and $\frac{e}{e-1} \cdot (1 + o(1)) \ln n$ in the sequential version ($\gamma(n) = \frac{e}{e-1}$). We conjecture that the approximation guarantees obtained by applying the reduction framework are not tight, and can be strengthened to $(1 + o(1)) \ln n$ in both these environments, like in the IID setting (Section 5).

Other Models of Access to Distributions. Our framework crucially makes use of virtual values, which in turn need complete knowledge of the distribution of ink levels. Thus, tailoring this approach to models where the algorithm gets partial access to the priors is challenging.

Qiao and Valiant [23] study two such models. They study the online oblivious problem where the algorithm gets access to just one sample from the distribution of each pen. They give a $O(\log n)$ -approximation to the omniscient benchmark in this setting. They also look at the online secretary setting, where the pens arrive according to a permutation chosen uniformly at random. The quantity of ink in the pens are adversarially determined. The algorithm is told the quantity of ink in the pen with the largest ink level. They give a $O(\log n)$ -approximation in this environment. Further, when the algorithm is given access to the ink levels in all the n pens (with random arrival order), they tighten their bound to $O(\frac{\log n}{\log \log n})$. It is an open question to identify reduction frameworks for these different models of access to the distributions.

References

- 1 Saeed Alaei, Hu Fu, Nima Haghpanah, Jason D Hartline, and Azarakhsh Malekian. Bayesian optimal auctions via multi-to single-agent reduction. In *13th ACM Conference on Electronic Commerce, EC'12*, page 17, 2012. doi:10.1145/2229012.2229017.
- 2 Saeed Alaei, Hu Fu, Nima Haghpanah, Jason D Hartline, and Azarakhsh Malekian. Efficient computation of optimal auctions via reduced forms. *Mathematics of Operations Research*, 44(3):1058–1086, 2019. doi:10.1287/MOOR.2018.0958.
- 3 Vivi Alatas, Abhijit Banerjee, Rema Hanna, Benjamin A Olken, Ririn Purnamasari, and Matthew Wai-Poi. Ordeal mechanisms in targeting: Theory and evidence from a field experiment in indonesia. Technical report, National Bureau of Economic Research, 2013.
- 4 Sushil Bikhchandani, Sven De Vries, James Schummer, and Rakesh V Vohra. An ascending vickrey auction for selling bases of a matroid. *Operations research*, 59(2):400–413, 2011. doi:10.1287/OPRE.1100.0888.
- 5 Mark Braverman, Jing Chen, and Sampath Kannan. Optimal provision-after-wait in healthcare. *Mathematics of Operations Research*, 41(1):352–376, 2016. doi:10.1287/MOOR.2015.0731.
- 6 Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320, 2010. doi:10.1145/1806689.1806733.
- 7 Hao Chung and Elaine Shi. Foundations of transaction fee mechanism design. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3856–3899. SIAM, 2023. doi:10.1137/1.9781611977554.CH150.

- 8 Paul Dütting, Vasilis Gkatzelis, and Tim Roughgarden. The performance of deferred-acceptance auctions. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 187–204, 2014. doi:10.1145/2600057.2602861.
- 9 Cynthia Dwork, Andrew Goldberg, and Moni Naor. On memory-bound functions for fighting spam. In *Advances in Cryptology-CRYPTO 2003: 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003. Proceedings 23*, pages 426–444. Springer, 2003. doi:10.1007/978-3-540-45146-4_25.
- 10 Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In *Annual international cryptology conference*, pages 139–147. Springer, 1992. doi:10.1007/3-540-48071-4_10.
- 11 Michal Feldman, Vasilis Gkatzelis, Nick Gravin, and Daniel Schoepflin. Bayesian and randomized clock auctions. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 820–845, 2022. doi:10.1145/3490486.3538247.
- 12 Yiding Feng, Jason D Hartline, and Yingkai Li. Simple mechanisms for non-linear agents. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3802–3816. SIAM, 2023. doi:10.1137/1.9781611977554.CH148.
- 13 Dimitris Fotakis, Dimitris Tsipras, Christos Tzamos, and Emmanouil Zampetakis. Efficient money burning in general domains. *Theory of Computing Systems*, 59:619–640, 2016. doi:10.1007/S00224-016-9720-2.
- 14 Aadityan Ganesh, Clayton Thomas, and S. Matthew Weinberg. Revisiting the primitives of transaction fee mechanism design. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC '24), July 8–11, 2024, New Haven, CT, USA. ACM, New York, NY, USA, 1 page.*, 2024. doi:10.1145/123456.654321.
- 15 Kira Goldner and Taylor Lundy. Simple mechanisms for utility maximization: Approximating welfare in the iid unit-demand setting. *arXiv preprint arXiv:2402.12340*, 2024. doi:10.48550/arXiv.2402.12340.
- 16 Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI*, volume 7, pages 58–65, 2007. URL: <http://www.aaai.org/Library/AAAI/2007/aaai07-009.php>.
- 17 Jason Hartline and Samuel Taggart. Sample complexity for non-truthful mechanisms. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 399–416, 2019. doi:10.1145/3328526.3329632.
- 18 Jason D Hartline and Tim Roughgarden. Optimal mechanism design and money burning. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 75–84, 2008. doi:10.1145/1374376.1374390.
- 19 Robert Kleinberg and Seth Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 123–136, 2012.
- 20 Henrik Jacobsen Kleven and Wojciech Kopczuk. Transfer program complexity and the take-up of social benefits. *American Economic Journal: Economic Policy*, 3(1):54–90, 2011.
- 21 Paul Milgrom and Ilya Segal. Deferred-acceptance auctions and radio spectrum reallocation. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 185–186, 2014. doi:10.1145/2600057.2602834.
- 22 Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981. doi:10.1287/MOOR.6.1.58.
- 23 Mingda Qiao and Gregory Valiant. Online pen testing. *14th Innovations in Theoretical Computer Science Conference, (ITCS) 2023, January 10-13, 2023, MIT, Cambridge, USA, 2023*. doi:10.48550/arXiv.2210.00655.
- 24 Tim Roughgarden. Transaction fee mechanism design. In *Proceedings of the 22nd ACM Conference on Economics and Computation, EC '21*, page 792, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3465456.3467591.

52:22 Combinatorial Pen Testing (Or Consumer Surplus of Deferred-Acceptance Auctions)

- 25 Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Ann. Probab.*, 12(4):1213–1216, 1984. URL: <http://dml.mathdoc.fr/item/1176993150>.
- 26 Qiqi Yan. Mechanism design via correlation gap. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 710–719. SIAM, 2011. doi:10.1137/1.9781611973082.56.