

Optimal Platform Design*

Jason D. Hartline[†]

Tim Roughgarden[‡]

Abstract

An auction house cannot generally provide the optimal auction technology to every client. Instead it provides one or several auction technologies, and clients select the most appropriate one. For example, eBay provides ascending auctions and “buy-it-now” pricing. For each client the offered technology may not be optimal, but it would be too costly for clients to create their own. We call these mechanisms, which emphasize generality rather than optimality, platform mechanisms. A platform mechanism will be adopted by a client if its performance exceeds that of the client’s outside option, e.g., hiring (at a cost) a consultant to design the optimal mechanism. We ask two related questions. First, for what costs of the outside option will the platform be universally adopted? Second, what is the structure of good platform mechanisms? We answer these questions using a novel prior-free analysis framework in which we seek mechanisms that are approximately optimal for every prior.

1 Introduction

Auction houses, like Sotheby’s, Christie’s, and eBay, exemplify the commodification of economic mechanisms, like auctions, and warrant an accompanying theory of design. The field of *mechanism design* suggests how special-purpose mechanisms might be optimally designed; however, in commodity industries there is a trade-off between special-purpose and general-purpose products. While for any particular setting an optimal special-purpose product is better, a general-purpose product may be favored, for instance, because of its cheaper cost or greater versatility. We develop a theory for the design of general-purpose mechanisms, henceforth, *platform design*.

Consider the following simple model for platform design. The platform *provider* offers a *platform* mechanism to potential customers (*principals*), who each wish to employ the mechanism in their particular *setting*. For example, the provider is eBay, the platform is the eBay auction, the principals are sellers, and the settings are the distinct markets of the sellers, which comprise of a set of buyers (*agents*) with preferences drawn according to a distribution. Each principal has the option to not adopt the platform and instead to employ a consultant to design the optimal auction for his specific setting. We assume that this outside option comes at a greater cost than the platform, and thus the platform provider has a *competitive advantage*.

*There is some overlap between this paper and the paper “Optimal Mechanism Design and Money Burning,” which appeared in the STOC 2008 conference. However, the focus of this paper is different, with some of our earlier results omitted and several new results included.

[†]Electrical Engineering and Computer Science, Northwestern University, Evanston, IL 60208. Email: hartline@eecs.northwestern.edu. This work was done while author was at Microsoft Research, Silicon Valley.

[‡]Department of Computer Science, Stanford University, Stanford, CA 94305. Email: tim@cs.stanford.edu. Supported in part by NSF CAREER Award CCF-0448664, an ONR Young Investigator Award, and an Alfred P. Sloan Fellowship.

We impose two restrictions to focus on the differences between the special-purpose optimal mechanism design and the general-purpose optimal platform design. First, we restrict the platform to be a single, unparameterized mechanism (unlike eBay where sellers can set their own reserve prices).¹ Second, we require that the platform is universally adopted. Without this assumption, we would need to model in detail the relative value of adoption in each setting, and this would likely give less general results. We ask: *What must the competitive advantage of the platform be to guarantee universal adoption by all principals? What is the platform designer’s mechanism that guarantees universal adoption?*

There are two important points of contact between this theory of platform design and the existing literature. First, the problem of optimal platform design provides a formal setting in which to explore the Wilson (1987) doctrine, which critiques mechanisms that are overly dependent on the details of the setting but does not quantify the cost of this dependence. A universally adopted platform, by definition, performs well in all settings and hence is not dependent on the details of setting. Second, the optimal platform design problem is closely related to *prior-free optimal mechanism design*. Indeed, our study of platform design formally connects the prior-free and Bayesian theories of optimal mechanism design. We make a rigorous comparison between the two settings and quantify the Bayesian designer’s relative advantage over the prior-free designer.

Platform Design. In classical Bayesian optimal mechanism design, a principal designs a mechanism for a set of self-interested agents that have private preferences over the outcomes of the mechanism. These private preferences are drawn from a known probability distribution. The optimal mechanism is the one that maximizes the expected value of the principal’s objective function when the agents’ strategies are in Bayes-Nash equilibrium.

For a given distribution and objective function, the *approximation factor* of a candidate mechanism is the ratio between the expected performance of an optimal mechanism and that of the candidate mechanism. A good mechanism is one with a small approximation factor (close to 1); a bad one has a large approximation factor.

We assume that the cost of designing the optimal mechanism is higher than the cost of adopting the platform. For this reason, a principal might choose to adopt the sub-optimal platform mechanism. We assume this *competitive advantage* of the platform is multiplicative. This assumption is consistent with commission structures in marketing and, from a technical point of view, frees the model from artifacts of scale. The platform’s competitive advantage gives an upper bound on the approximation factor that the platform mechanism needs to induce a principal to adopt the platform instead of hiring a consultant to design the optimal mechanism. Each principal’s decision to adopt is based on the platform mechanism’s performance in the principal’s setting. Therefore, universal adoption demands that the platform mechanism’s approximation factor on every distribution is at most its competitive advantage. Of particular interest is the minimum competitive advantage for which there is a platform that is universally adopted, and also the platform that attains this minimum approximation factor. This *optimal platform* is the mechanism that minimizes (over mechanisms) the maximum (over distributions) approximation factor. Optimal platform design is therefore inherently a min-max design criterion.

The basic formal question of platform design is: *What is the minimum competitive advantage β and optimal platform mechanism \mathcal{M} such that for all distributions F the expected performance of*

¹In a separate study, we consider the technically orthogonal topic of reserve-price based platforms (Hartline and Roughgarden, 2009).

\mathcal{M} when values are drawn i.i.d. from F is at least $\frac{1}{\beta}$ times the expected performance of the optimal mechanism for F ?

Directly answering the platform design questions above is difficult as it requires simultaneous consideration of all distributions. This difficulty motivates a more stringent version of the basic question which has the following economic interpretation. Suppose that instead of requiring the principal to choose ex ante between the optimal mechanism and the platform, we allow him to choose ex post? Clearly, this makes the platform designer’s task even more challenging, in that the minimum achievable β is only higher.

The formal question of platform design now becomes: *What is the minimum competitive advantage β and optimal platform mechanism \mathcal{M} such that for all valuation profiles $\mathbf{v} = (v_1, \dots, v_n)$ the performance of \mathcal{M} on \mathbf{v} is at least $\frac{1}{\beta}$ times the supremum over symmetric² Bayesian optimal mechanisms’ performance on \mathbf{v} ?*

This question motivates the definition of a *performance benchmark* that is defined point-wise on valuation profiles, specifically as the supremum over optimal symmetric mechanisms’ performance on the given valuation profile. Notice that this benchmark is prior-free. The analysis of a platform mechanism is then a comparison of the performance of a prior-free platform mechanism and a prior-free performance benchmark.

Results. Our contributions are two-fold. First, we propose a conceptual framework for the design and analysis of general-purpose platforms. Second, we instantiate this framework to derive novel platform mechanisms for specific problems and, in some cases, prove their optimality.

In more detail, we consider the problem of optimal platform design in general symmetric settings of multi-unit unit-demand allocation problems and for general linear (in agents’ payments and values) objectives of the principal. For much of the paper, we focus on the canonical objective of *residual surplus*, which is the difference between the winning agents’ values and payments. Residual surplus is interesting in its own right (e.g., McAfee and McMillan, 1992; Condorelli, 2012; Chakravarty and Kaplan, 2013) and is, in a sense, technically more general than the objectives of surplus and profit.³ Intuitively, maximizing the residual surplus involves compromising between the competing goals of identifying high-valuation agents and of minimizing payments. For example, with a single item, the Vickrey auction performs well when there is only one high-valuation agent, while giving the item away for free is good when all agents have comparable valuations.

Our approach comprises four steps.

1. We characterize Bayesian optimal mechanisms for multi-unit unit-demand allocation problems and general linear objectives by a straightforward generalization of the literature on optimal mechanism design.
2. We characterize the prior-free performance benchmark, i.e., the supremum over optimal symmetric mechanisms’ performance on a given valuation profile, as an ex post optimal two-level lottery.

²Our study focuses solely on settings where the agents are a priori indistinguishable. This focus motivates our restriction to i.i.d. distributions and symmetric optimal mechanisms. Distinguishable agents are considered by Balcan et al. (2008) and Bhattacharya et al. (2013).

³For surplus maximization, the Vickrey auction is optimal for every distribution. For profit maximization, reserve-price-based auctions are optimal for standard distributions assumptions (Myerson, 1981). For residual surplus, reserve-price-based auctions are not optimal even for standard distributions.

3. We give a general platform design and a finite upper bound on the competitive advantage necessary for universal adoption.
4. We give a lower bound on the competitive advantage for which there exists a platform that achieves universal adoption.

Importantly, the platform mechanisms that we identify as being universally adopted with finite competitive advantage are not standard mechanisms from the literature on Bayesian optimal mechanisms. Indeed, we prove that no standard mechanism is universally adopted with any finite competitive advantage. Instead, general purpose mechanisms for platforms require novel features, which we identify in Step 3.

Example. Our main results are interesting to interpret in the special case of allocating a single item to one of two agents to maximize the residual surplus. Denote the high agent value by $v_{(1)}$ and the low agent value by $v_{(2)}$. We characterize the performance benchmark as $\max(\frac{v_{(1)}+v_{(2)}}{2}, v_{(1)} - \frac{v_{(2)}}{2})$. As the supremum of Bayesian optimal mechanisms, the first term in this benchmark arises from a lottery and the second term from the two-level lottery that serves a random agent with value strictly above price $v_{(2)}$ if one exists and otherwise serves a random agent (at price zero). The optimal platform mechanism randomizes between a lottery and weighted Vickrey auctions. Precisely, it sets $w_1 = 1$, draws w_2 uniformly from $\{0, 1/2, 2, \infty\}$, and serves the agent $i \in \{1, 2\}$ that maximizes $w_i v_i$. It is universally adopted with competitive advantage $\frac{4}{3}$ and no other mechanism is better.

While all possible prior distributions are considered when deriving the performance benchmark above, the actual benchmark for a particular valuation profile is given by a simple formula with no distributional dependence. Consequently, our analysis that shows the $\frac{4}{3}$ competitive advantage is a simple comparison between a (prior-free) platform mechanism and a (prior-free) performance benchmark in the worst case over valuation profiles.

Related Work. Our description of Bayesian optimal mechanisms for general linear objectives follows from the work on optimal mechanism design (see Myerson, 1981, and Riley and Samuelson, 1981). Within this theory, the residual surplus objective coincides with that of the grand coalition in a *weak cartel*, where agents wish to maximize the cartel’s total utility without side payments amongst themselves, so that payments to the auctioneer are effectively “burnt”. Our characterizations are thus related to those in the literature on collusion in multi-unit auctions, e.g., by McAfee and McMillan (1992) and Condorelli (2012). Recently, Chakravarty and Kaplan (2013) also specifically studied Bayesian optimal auctions for residual surplus.

There is a growing literature on “redistribution mechanisms” where, similar to the objective of residual surplus, payments are bad, e.g., see Moulin (2009) and Guo and Conitzer (2009). These mechanisms transfer some of the winners’ payments back to the losers so that the residual payment left over is as small as possible. The mechanisms considered are prior-free.

Finally, as already mentioned, there is a large related literature on prior-free optimal mechanism design. Goldberg et al. (2001), Segal (2003), Baliga and Vohra (2003), and Balcan et al. (2008) consider asymptotic approximation of the Bayesian optimal mechanisms by a single (prior-free) mechanism. This is quite different from our question of platform design as it says nothing about whether or not a principal in a small or moderate-sized market would adopt the platform. The line of research initiated by Goldberg et al. (2006) on prior-free profit maximization can be reinterpreted in the context of platform design; Section 6 describes this connection in detail.

2 Warm-up: Monopoly Pricing

Consider the following monopoly pricing problem. A monopolist seller (principal) of a single item faces a single buyer (agent). The seller has no value for the item and wishes to maximize his revenue, i.e., the payment of the buyer. The buyer's value for the item is $v \in [1, h]$ and she wishes to maximize her utility which is her value less her payment. The seller may post a price p and the buyer may take it or leave it. The buyer will clearly take any price $p \leq v$.

The seller's optimal mechanism, when the buyer's value comes from the distribution F (where $F(z) = \Pr[v \leq z]$), is to post the price p that maximizes $p(1 - F(p))$, a.k.a., the *monopoly price*. The performance benchmark $\mathcal{G}(v)$, i.e., the revenue of the best of the Bayesian optimal mechanism when the buyer's value is v , is then $\mathcal{G}(v) = v$. The platform designer must give a single mechanism with revenue that approximates v for every value v in the support $[1, h]$. The optimal platform and its competitive advantage for universal adoption are given by the theorem below.

Theorem 2.1. *The optimal platform mechanism offers a price drawn from distribution P with cumulative distribution function $P(z) = (1 + \ln z)/(1 + \ln h)$ on $[1, h]$, and a point mass of $1/(1 + \ln h)$ at 1 , and is universally adopted with competitive advantage $1 + \ln h$.*

Proof. An easy calculation verifies that, for every $v \in [1, h]$, the expected revenue from such a random price from P is $v/(1 + \ln h)$. Thus, the competitive advantage for universal adoption is $1 + \ln h$ as claimed.

To show that this is the optimal platform, we can similarly find a distribution F over values v such that the expected revenue of every platform mechanism is 1. The *equal revenue* distribution has distribution function $F(z) = 1 - 1/z$, a point mass of $1/h$ at h , and any price p is accepted by the agent with probability $1/p$ for an expected revenue of 1. The expected value of the benchmark for the equal-revenue distribution can be calculated as $\mathbf{E}[\mathcal{G}(v)] = \mathbf{E}[v] = 1 + \ln h$. Thus, the ratio of these expectations is $1 + \ln h$, and for any platform mechanism there must be some $v \in [1, h]$ that achieves the ratio. We conclude that no platform is universally adopted with competitive advantage less than $1 + \ln h$. \square

This analysis can be viewed as a zero-sum game between the platform designer and Nature where the solution is a mixed strategy on the part of both players, every action in the game achieves equal payoff, and the value of the game is the optimal competitive advantage.

To conclude, we considered a simple monopoly pricing setting and derived for it the optimal platform. While a logarithmic competitive advantage may seem impractical, except when the maximum variation h in values is small, the ideas from this design and analysis play an important role in the developments of this paper. The platform mechanisms we derive subsequently, however, will be universally adopted with a competitive advantage that is an absolute constant, independent of the number of agents, the number of units, and the range of agent values.

3 Review of Bayesian Optimal Mechanism Design

In this section we review Bayesian optimal mechanism design for single-dimensional agents, i.e., with utility given by the value for receiving a good or service less the required payment, and develop the notation employed in the remainder of the paper. Characterizing Bayesian optimal mechanisms is the first step in our approach to platform design.

We consider mechanisms for allocating k units of an indivisible item to n unit-demand agents. The outcome of such a mechanism is an *allocation vector*, $\mathbf{x} = (x_1, \dots, x_n)$, where x_i is 1 if agent i receives a unit and 0 otherwise, and a non-negative *payment vector*, $\mathbf{p} = (p_1, \dots, p_n)$. The allocation vector \mathbf{x} is required to be feasible, i.e., $\sum_i x_i \leq k$, and we denote this set of feasible allocation vectors by \mathcal{X} .

We assume that each agent i is risk-neutral, has a privately known valuation v_i for receiving a unit, and aims to maximize her (quasi-linear) utility, defined as $u_i = v_i x_i - p_i$. Each agent's value is drawn independently and identically from a continuous distribution F , where $F(z)$ and $f(z)$ denote the cumulative distribution and density functions, respectively. We denote the *valuation profile* by $\mathbf{v} = (v_1, \dots, v_n)$.

We consider general symmetric, linear objectives of the mechanism designer. For valuation coefficient γ_v and payment coefficient γ_p , the objective for maximization is:

$$\sum_{i=1}^n \gamma_v v_i x_i + \gamma_p p_i. \quad (1)$$

We single out three such objectives: *surplus* with $\gamma_v = 1$ and $\gamma_p = 0$, *profit* with $\gamma_v = 0$ and $\gamma_p = 1$, and *residual surplus* with $\gamma_v = 1$ and $\gamma_p = -1$. We will not discuss surplus maximization in this paper as the optimal mechanism for this objective is simply the prior-free k -unit Vickrey auction; therefore, we assume that $\gamma_p \neq 0$.

We assume that agents play in Bayes-Nash equilibrium and moreover if truthtelling is a Bayes-Nash equilibrium then agents truthtell. When searching for Bayesian optimal mechanisms, the revelation principle (Myerson, 1981) allows us to restrict attention to Bayesian incentive compatible mechanisms, i.e., ones with a truthtelling Bayes-Nash equilibrium.

Characterization of incentive compatibility. The *allocation rule*, $\mathbf{x}(\mathbf{v})$, is the mapping (in equilibrium) from agent valuations to the outcome of the mechanism. Similarly the *payment rule*, $\mathbf{p}(\mathbf{v})$, is the mapping from valuations to payments. Given an allocation rule $\mathbf{x}(\mathbf{v})$, let $x_i(v_i)$ be the interim probability with which agent i is allocated when her valuation is v_i (over the probability distribution on the other agents' valuations): $x_i(v_i) = \mathbf{E}_{\mathbf{v}_{-i}}[x_i(v_i, \mathbf{v}_{-i})]$. Similarly define $p_i(v_i)$. We require interim individual rationality, i.e., that non-participation in the mechanism is an allowable agent strategy. The following lemma provides the standard characterization of allocation rules that are implementable by Bayesian incentive compatible mechanisms and the accompanying payment rule (which is unique up to additive shifts, and usually fixed by setting $p_i(0) = 0$).

Lemma 3.1. (Myerson, 1981) *Every Bayesian incentive compatible mechanism satisfies, for all i and $v_i \geq v'_i$:*

(a) *Allocation monotonicity:* $x_i(v_i) \geq x_i(v'_i)$.

(b) *Payment identity:* $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z) dz + p_i(0)$.

Virtual valuations. Myerson (1981) defined *virtual valuations* and showed that the virtual surplus of an agent is equal to her expected payment. For $v \sim F$, this *virtual valuation for payment* is:

$$\varphi(v_i) = v_i - \frac{1-F(v_i)}{f(v_i)}. \quad (2)$$

Lemma 3.2. (Myerson, 1981) *In a Bayesian incentive-compatible mechanism with allocation rule $\mathbf{x}(\cdot)$, the expected payment of an agent equals her expected virtual surplus: $\mathbf{E}_{\mathbf{v}}[p_i(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}[\vartheta(v_i) x_i(\mathbf{v})]$.*

The notion of virtual valuations applies generally to linear objectives. By substituting virtual values for payments into the objective (1) we arrive at a formula for *general virtual values*: $\vartheta(v_i) = (\gamma_v + \gamma_p)v_i - \gamma_p \frac{1-F_i(v_i)}{f_i(v_i)}$. For the objective of residual surplus, i.e., the sum of the agent utilities, *virtual values for utility* are given by:

$$\vartheta(v_i) = \frac{1-F_i(v_i)}{f_i(v_i)}. \quad (3)$$

The revenue-optimal mechanism for a given distribution is the one that maximizes the virtual surplus for payment subject to feasibility and monotonicity of the allocation rule. Analogously, optimal mechanisms for general linear objectives are precisely those that maximize the expected (general) virtual surplus subject to feasibility and monotonicity of the allocation rule. Unfortunately, choosing \mathbf{x} to maximize $\sum_i \vartheta(v_i)x_i$ for each valuation profile \mathbf{v} does not generally result in a monotone allocation rule. When $\vartheta(\cdot)$ is not monotone increasing, an increase in an agent’s value may decrease her virtual value and cause her to be allocated less frequently. Notice that under the standard “monotone hazard rate” assumption the virtual value function for utility $\vartheta(v) = \frac{1-F_i(v)}{f_i(v)}$ is monotone *in the wrong direction*.

Ironing. We next generalize the “ironing” procedure of Myerson (1981) that transforms a possibly non-monotone virtual valuation function into an *ironed virtual valuation* function that is monotone; optimizing ironed virtual surplus results in a monotone allocation rule. Furthermore, the ironing procedure preserves the target objective, so that an optimal allocation rule for the ironed virtual valuations is equal to the optimal monotone allocation rule for the original virtual valuations.

Given a distribution function $F(\cdot)$ with virtual valuation function $\vartheta(\cdot)$, the *ironed virtual valuation function*, $\bar{\vartheta}(\cdot)$, is constructed as follows:

1. For $q \in [0, 1]$, define $h(q) = \vartheta(F^{-1}(q))$.
2. Define $H(q) = \int_0^q h(r)dr$.
3. Define G as the convex hull of H — the largest convex function bounded above by H for all $q \in [0, 1]$.
4. Define $g(q)$ as the derivative of $G(q)$, where defined, extended to all of $[0, 1]$ by right-continuity.
5. Finally, define $\bar{\vartheta}(z) = g(F(z))$.

Convexity of G implies that Step 4 of the ironing procedure is well defined and that g , and hence $\bar{\vartheta}$, is a monotone non-decreasing function.

From the main theorem of Myerson (1981), maximizing the expectation of a general linear objective subject to incentive compatibility is equivalent to maximizing the expected ironed virtual surplus. Different tie-breaking rules, however, can yield different optimal mechanisms. In our symmetric settings, with i.i.d. agents and the symmetric feasibility constraint \mathcal{X} of k -unit auctions, it is natural to consider symmetric optimal mechanisms.

Theorem 3.3. *For every general linear objective and distribution F , the k -unit auction that allocates the units to the agents with the highest non-negative ironed virtual values, breaking ties randomly and discarding all leftover units, maximizes the expected value of the objective.*




Figure 1: Ironed virtual value functions in the three distributional cases. For the objective of residual surplus the cases correspond to (a) MHR distributions, (b) anti-MHR distributions, and (c) non-MHR distributions.

Interpretation for residual surplus maximization. Consider the residual surplus objective, where $\vartheta(v) = \frac{1-F(v)}{f(v)}$, and the following three types of distributions (Figure 1). Monotone hazard rate (MHR) distributions; e.g., uniform, normal, and exponential; have monotone non-increasing $\vartheta(v)$. In this case, ironing $\vartheta(\cdot)$ to be non-decreasing results in $\bar{\vartheta}(\cdot) = \mathbf{E}[v]$, a constant function. The optimal (symmetric) mechanism is therefore a lottery that awards the k units to k agents uniformly at random. For distributions with a hazard rate monotone in the opposite direction, henceforth *anti-MHR* distributions, $\vartheta(\cdot)$ is non-negative and monotone non-decreasing. Power-law distributions, such as $F(z) = 1 - 1/z^c$ with $c > 0$ on $[1, \infty)$, are canonical examples. In this case, the optimal mechanism awards the k units to the k highest valued agents, i.e., it is the k -Vickrey auction. Thus, as also observed by McAfee and McMillan (1992), Chakravarty and Kaplan (2006), and Condorelli (2007), the optimal mechanism depends on whether or not the distribution is heavy-tailed.

The final case occurs when the distribution is neither MHR nor anti-MHR, henceforth *non-MHR*. Here, the ironed virtual valuation function $\bar{\vartheta}(\cdot)$ is constant on some intervals and monotone increasing on other intervals. The optimal mechanism can be described, for instance, as an indirect Vickrey auction where agents are not allowed to bid on intervals where the ironed virtual value is constant. For example, consider the two-point distribution with probability mass $\frac{1}{2}$ on 1 and $\frac{1}{2}$ on $h > 1$. Provided h is sufficiently large, the residual-surplus-maximizing mechanism allocates to a random high-value agent or, if there are no high-value agents, to a random (low-value) agent. This final case is the most general, in that it subsumes both the MHR and anti-MHR cases. Our general theory of platform design necessitates understanding this non-MHR case in detail.

4 The Performance Benchmark

In this section we leverage the characterization of Bayesian optimal mechanisms from the preceding section to identify and characterize a simple prior-free performance benchmark. This constitutes the second step of our approach to platform design.

The performance benchmark is derived as follows. As discussed in Section 3, Bayesian optimal mechanisms are ironed virtual surplus optimizers. For k -unit environments, these mechanisms simply select the k agents with the highest non-negative ironed virtual values. Among these optimal

mechanisms, the symmetric one breaks ties randomly. Denote the symmetric optimal mechanism for distribution F by Opt_F . Denote by $\text{Opt}_F(\mathbf{v})$ the expected performance (over the choice of random allocation) obtained by the mechanism Opt_F on the valuation profile \mathbf{v} .

Definition 4.1. *The performance benchmark is the supremum of Bayesian optimal mechanisms, $\mathcal{G}(\mathbf{v}) = \sup_F \text{Opt}_F(\mathbf{v})$.*

For one interpretation of the definition of \mathcal{G} , observe that

$$\mathbf{E}_{\mathbf{v}}[\mathcal{G}(\mathbf{v})] \geq \mathbf{E}_{\mathbf{v}}[\text{Opt}_F(\mathbf{v})] \quad (4)$$

for valuation profiles drawn i.i.d. from an arbitrary distribution F . Thus, the approximation of the performance benchmark \mathcal{G} implies the simultaneous approximation of all symmetric Bayesian optimal mechanisms.

We now give a simple characterization of the performance benchmark for general linear objectives by considering ex post outcomes of symmetric Bayesian optimal mechanisms. When k units are available, a symmetric Bayesian optimal mechanism serves these units to the k agents with the highest non-negative ironed virtual values. Ties, which occur in ironed virtual surplus maximization when two (or more) agents' values are mapped to same ironed virtual value, are broken randomly. Ex post, we can classify the agents into at most three groups: those that win with certainty (winners), those that lose with certainty (losers), and those that win with a common probability strictly between 0 and 1 (partial winners).

Definition 4.2. *A two-level (p, q) -lottery, denoted $\text{Lot}_{p,q}$, first serves agents with values strictly more than p , then serves agents with values strictly more than q , while supplies last (breaking ties randomly, as needed). All agents with values at most q are rejected.*

It will be useful to calculate explicitly, using Lemma 3.1, the payments of a two-level lottery. Let S and T denote the sets of agents with value in the ranges (p, ∞) and $(q, p]$, respectively. Let $s = |S|$ and $t = |T|$. For simplicity, assume that $s \leq k < s + t$, where k is the number of units available. The payments are as follows.

1. Agents $i \in S$ are each allocated a unit and charged

$$p_i = p - (p - q) \frac{k-s+1}{t+1}. \quad (5)$$

2. The remaining $k - s$ units are allocated uniformly at random to the $k - s$ agents $i \in T$, i.e., by lottery; each such winner pays $p_i = q$.

We characterize the performance benchmark for platform design for general linear objectives in terms of two-level lotteries.

Theorem 4.3. $\mathcal{G}(\mathbf{v}) = \sup_F \text{Opt}_F(\mathbf{v}) = \sup_{p,q} \text{Lot}_{p,q}(\mathbf{v})$.

Proof. The outcome of ironed virtual surplus maximization is equivalent to a k -unit (p, q) -lottery. To see this, consider an ironed virtual valuation function $\bar{\vartheta}$ and a valuation profile \mathbf{v} . Set p to be the infimum bid that the highest-valued agent can make and be a winner (possibly larger than the agent's value), and q to be the infimum bid that a partial winner can make and remain a partial winner (or p if there are no partial winners). The two mechanisms have the same outcome on

profile \mathbf{v} . Conversely, every (p, q) -lottery arises in ironed virtual surplus maximization with respect to some i.i.d. distribution, for example with $\bar{v}(v) = 2$ for $v \in (p, \infty)$, $\bar{v}(v) = 1$ for $v \in (q, p]$, and $\bar{v}(v) = -1$ for $v \leq q$.⁴ \square

We conclude with a simple but useful observation: The values of p and q that attain the supremum in Theorem 4.3 must each either be zero, infinity, or an agent’s value. Observe that the objective $\sum_i \gamma_v v_i x_i + \gamma_p p_i$ is linear in payments. If q or p is not in the valuation profile, then it can either be increased or decreased without decreasing the objective. For example, lowering p or q without changing the allocation increases residual surplus.

5 Residual Surplus

In this section we consider platform design for the objective of residual surplus. We consider separately the $n = 2$ agent case and the general $n > 2$ agent case. For $n = 2$ agents (and a single item) we completely execute our template for platform design by reinterpreting the benchmark, giving a platform mechanism that is universally adopted with competitive advantage $4/3$, and proving that no platform mechanism is universally adopted with a smaller competitive advantage. The platform mechanism that achieves this bound is neither a standard auction nor a mixture over standard auctions, where by “standard” we mean a symmetric Bayesian-optimal mechanism with respect to some valuation distribution.

For every number $n > 2$ of agents and $k \geq 1$ of items, we give a heuristic platform that guarantees universal adoption with a constant competitive advantage (independent of k , n , and the support of the valuations). This platform is not a mixture of standard auctions, and we show in Appendix B that no such mixture is universally adopted with any finite competitive advantage (as $n \rightarrow \infty$). This heuristic mechanism identifies properties of good platforms and is a proof-of-concept that good platforms exist.

5.1 Single-unit Two-agent Platforms

We now execute the framework for platform design for two agents, a single unit, and the objective of residual surplus. Bayesian optimal mechanisms and our benchmark are characterized in Sections 3 and 4, respectively; for two agents and a single item, the benchmark takes a simple form.

There are only two relevant (p, q) -lotteries for the performance benchmark, the degenerate $p = q = 0$ lottery, and the $p = v_{(2)}$ and $q = 0$ lottery; here $v_{(1)}$ and $v_{(2)}$ denote the highest and second-highest agent values, respectively. From equation (5), the residual surpluses of these two-level lotteries are $\frac{v_1 + v_2}{2}$ (i.e., the average value) and $v_{(1)} - \frac{v_{(2)}}{2}$, respectively. Thus,

$$\mathcal{G}(\mathbf{v}) = \max\left\{\frac{v_1 + v_2}{2}, v_{(1)} - \frac{v_{(2)}}{2}\right\}. \quad (6)$$

This benchmark is depicted in Figure 2(a).

We now turn to the problem of designing a platform mechanism that is universally adopted with a minimal competitive advantage. As mentioned above, the lottery is adopted with a competitive advantage of 2. A natural approach to platform design is to randomly mix over two platforms that are good in different settings. For example, the Vickrey auction is good on the valuation profile

⁴For objectives like residual surplus where the virtual values are always non-negative, set $\bar{v}(v) = 1/2$ instead of -1 for $v \leq q$. See the construction in Appendix A for details.

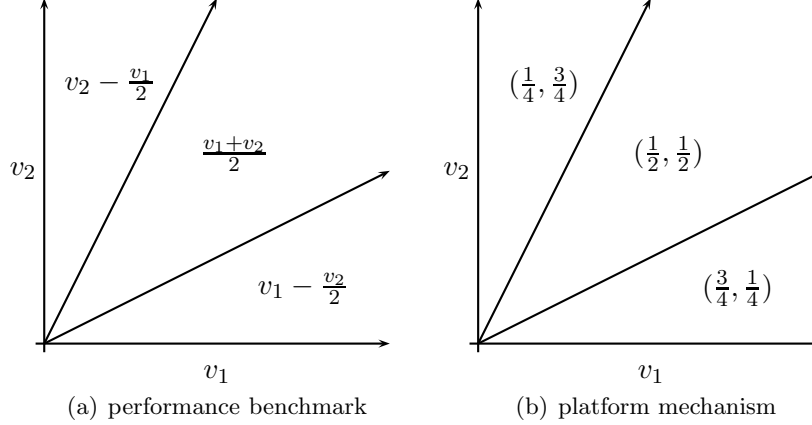


Figure 2: The performance benchmark (6) and optimal platform mechanism for the single-item, two-agent, residual-surplus-maximization problem. The positive quadrant is partitioned by the lines $v_1 = 2v_2$ and $2v_1 = v_2$. The allocation rule of the platform mechanism is given as (x_1, x_2) .

$\mathbf{v} = (1, 0)$, whereas the lottery is good on the valuation profile $\mathbf{v} = (1, 1)$. Considering only these two valuation profiles (where $\mathcal{G}(\mathbf{v}) = 1$), choosing the Vickrey auction with probability $1/3$ and the lottery with probability $2/3$ balances the competitive advantage necessary for adoption of the platform for each profile at $3/2$. In fact, a routine calculation shows that this mixture is universally adopted with competitive advantage $3/2$. This platform mechanism is, however, not optimal.

One approach to solving for the optimal platform mechanism is to look for a mechanism that achieves the same approximation factor to the benchmark for every valuation profile.⁵ Inspecting the benchmark (Figure 2(a)), we conclude that an auction with identical approximation factor on all inputs must have a discontinuity in behavior only where the ratio between the high and low value is 2. Importantly, there should be no discontinuity in behavior when the values are equal, that is, the optimal platform should never mix over the Vickrey auction. These observations suggest the following parameterized class of auctions.

Definition 5.1. *The two-agent single-item ratio auction with ratio $\alpha \geq 1$ and bias $\chi \in [1/2, 1]$ allocates the good according to a fair coin if the agent values are within a factor α of each other and, otherwise, according to a biased coin with probability χ in favor of the high-value agent.⁶*

The Vickrey auction and the lottery are special cases of the ratio auction. With bias $1/2$ the ratio auction is a lottery (for every ratio); with bias $\chi = 1$ and ratio $\alpha = 1$ it is the Vickrey auction. We next show that the optimal two-agent single-item platform for residual surplus is the ratio auction with ratio $\alpha = 2$ and bias $\chi = 3/4$. The allocation probabilities of this auction are depicted in Figure 2(b). It is adopted with competitive advantage $4/3$.

Lemma 5.2. *The ratio auction with ratio $\alpha = 2$ and bias $\chi = 3/4$ is universally adopted with competitive advantage $4/3$.*

⁵Our optimal platform for monopoly pricing in Section 2 also exhibits this property.

⁶Appropriate payments can be derived by reinterpreting the ratio auction as a distribution over weighted Vickrey auctions; see also the proof of Lemma 5.2.

Proof. The ratio auction (with ratio α) can always be expressed as a distribution over weighted Vickrey auctions, where $w_1 = 1$, w_2 is selected randomly from some distribution over the set $\{0, 1/\alpha, \alpha, \infty\}$, and the agent i that maximizes $w_i v_i$ winning the item. With bias $\chi = 3/4$, the distribution over the set is uniform. We calculate the auction's approximation of the benchmark via simple case analysis. the expected residual surplus from the four choices of w_2 averages to $\frac{1}{4} [v_1 + (v_1 - \frac{v_2}{2}) + (v_2 - \frac{v_1}{2}) + v_2] = \frac{3}{4} \frac{v_1 + v_2}{2}$ when $v_1 \in [v_2/2, 2v_2]$ and to $\frac{1}{4} [v_1 + (v_1 - \frac{v_2}{2}) + (v_1 - 2v_2) + v_2] = \frac{3}{4} (v_1 - \frac{v_2}{2})$ when $v_1 > 2v_2$. The case where $v_1 < v_2/2$ is symmetric. In each case, the expected residual surplus is exactly $\frac{3}{4}\mathcal{G}(\mathbf{v})$. \square

We now show that the ratio auction with ratio $\alpha = 2$ and bias $\chi = 3/4$ is an optimal platform; meaning, no platform is universally adopted with competitive advantage less than $4/3$. We first note that, for every distribution F , the expected residual surplus of the ratio auction with ratio $\alpha = 2$ and bias $\chi = 3/4$ is exactly $3/4$ times the expected value of the benchmark \mathcal{G} . Of course, the Bayesian optimal auction for F is no worse.

Corollary 5.3. *For every distribution F and $n = 2$ agents and $k = 1$ item, the expected benchmark is at most $4/3$ times the expected residual surplus of the optimal auction, that is, $\mathbf{E}[\mathcal{G}(\mathbf{v})] \leq \frac{4}{3}\mathbf{E}[\text{Opt}_F(\mathbf{v})]$.*

The following technical lemma exhibits a distribution F for which the inequality in Corollary 5.3 is tight. Intuitively, this distribution is the one with constant virtual value for utility.

Lemma 5.4. *For the exponential distribution $F(z) = 1 - e^{-z}$, $n = 2$ agents, $k = 1$ unit, the expected value of the benchmark is $4/3$ times the expected residual surplus of the optimal auction, that is, $\mathbf{E}[\mathcal{G}(\mathbf{v})] = \frac{4}{3}\mathbf{E}[\text{Opt}_F(\mathbf{v})]$.*

Proof. Since the exponential distribution has a monotone hazard rate, a lottery maximizes the expected residual surplus (Section 3). The expected value of an exponential random variable is 1 so $\mathbf{E}[\text{Opt}_F(\mathbf{v})] = \mathbf{E}[v] = 1$.

We now calculate the expected value of the benchmark $\mathcal{G}(\mathbf{v})$ defined in equation (6). Write the smaller value as $v = v_{(2)}$ and the higher value as $x + v = v_{(1)}$ for $x \geq 0$. In terms of v and x the benchmark is $v + \frac{x}{2}$ when $x \leq v$ and $\frac{v}{2} + x$ when $x \geq v$. Therefore, the expectation of \mathcal{G} conditioned on v is

$$\begin{aligned} \mathbf{E}[\mathcal{G}(x + v, v) \mid v] &= \int_0^v (v + \frac{x}{2}) e^{-x} dx + \int_v^\infty (\frac{v}{2} + x) e^{-x} dx \\ &= v(1 - e^{-v}) + \frac{1}{2}(1 - (v + 1)e^{-v}) + \frac{v}{2}e^{-v} + (v + 1)e^{-v} \\ &= v + \frac{1}{2}(1 + e^{-v}). \end{aligned}$$

The smaller value $v_{(2)} = v$ is distributed according to an exponential distribution with rate 2. Integrating out yields

$$\begin{aligned} \mathbf{E}[\mathcal{G}(x + v, v)] &= \int_0^\infty (v + \frac{1}{2} + \frac{1}{2}e^{-v}) 2e^{-2v} dv \\ &= \frac{1}{2} + \frac{1}{2} + \int_0^\infty e^{-3v} dv = \frac{4}{3}. \end{aligned} \quad \square$$

For the setting of Lemma 5.4, the optimal mechanism has expected residual surplus $\frac{3}{4}\mathbf{E}[\mathcal{G}(\mathbf{v})]$. Any platform mechanism is only worse and, by the definition of expectation, there must be a valuation profile \mathbf{v} where this platform mechanism has residual surplus at most $\frac{3}{4}\mathcal{G}(\mathbf{v})$.

Corollary 5.5. *For $n \geq 2$ agents, $k = 1$ item, and the residual surplus objective, no platform mechanism is universally adopted with competitive advantage less than $4/3$.*

We conclude that the ratio auction with ratio $\alpha = 2$ and bias $\chi = 3/4$ is an optimal platform for two-agent, single-item residual surplus maximization.

5.2 Multi-unit, Multi-agent Platforms

We now turn to markets with $n > 2$ agents and $k \geq 1$ units. We show that the minimum competitive advantage for universal adoption is a finite constant, independent of the number of units, the number of bidders, and the support size of the valuations.

In contrast to the $n = 2$ case, neither the Vickrey auction, the lottery, nor a convex combination thereof obtains a constant-factor approximation of the benchmark \mathcal{G} (Definition 4.1). For instance, with one object and valuation profile $\mathbf{v} = (1, 1, 0, \dots, 0)$, the Vickrey auction has zero residual surplus and the lottery has expected residual surplus $2/n$, while the benchmark residual surplus is $\mathcal{G}(\mathbf{v}) = 1$. In fact, no Bayesian optimal auction (a.k.a., standard auction) or mixture over standard auctions is universally adopted with a competitive advantage that is an absolute constant. This result is stated as Theorem 5.6, below, and proved in Appendix B. We conclude that the derivation of a platform mechanism that is universally adopted with a constant competitive advantage requires non-standard auction techniques.

Theorem 5.6. *For every $\rho > 1$ there is a sufficiently large n such that, for an n -agent, 1-unit setting, no mixture over standard auctions is universally adopted with competitive advantage ρ .*

Due to the complexity of the problem, we relax the goal of determining the optimal platform mechanism and instead look for a heuristic platform that is universally adopted with a constant competitive advantage. We believe this heuristic pinpoints properties of good platforms, while the optimal platform is complex and perhaps difficult to interpret.

This heuristic follows the random sampling paradigm of Goldberg et al. (2001). Half of the agents (henceforth: sample) are used for a market analysis to determine a good mechanism to run on the other half of the agents (henceforth: market). We do not attempt to estimate the distribution of the sample, as distributions are complex objects. Instead, we use the sample to determine a good two-level lottery and then simply run that two-level lottery on the market. Two-level lotteries are described by two numbers and are therefore, statistically, far simpler objects than distributions.

To make this task even simpler, we first argue that two-level lotteries can be approximated by one-level lotteries.

Definition 5.7. *The one-level r -lottery, denoted Lot_r , serves agents with values strictly more than r , while supplies last (breaking ties randomly). Winners are charged r and agents with values below r are rejected.*

Lemma 5.8. *For every valuation profile \mathbf{v} and parameters k, p , and q , there is an r such that the k -unit r -lottery obtains at least half of the expected residual surplus of the k -unit (p, q) -lottery.*

Proof. We prove the lemma by showing that $\text{Lot}_{p,q}(\mathbf{v}) \leq \text{Lot}_p(\mathbf{v}) + \text{Lot}_q(\mathbf{v})$. We argue the stronger statement that each agent enjoys at least as large a combined expected utility in $\text{Lot}_p(\mathbf{v})$ and $\text{Lot}_q(\mathbf{v})$ as in $\text{Lot}_{p,q}(\mathbf{v})$.

Let S and T denote the agents with values in the ranges (p, ∞) and $(q, p]$, respectively. Let $s = |S|$ and $t = |T|$. Assume that $0 < s \leq k < s + t$ as otherwise the k -unit (p, q) lottery is equivalent to a one-level lottery. Each agent in T participates in a k -unit q -lottery in Lot_q and only a $(k - s)$ -unit q -lottery in $\text{Lot}_{p,q}$; her expected utility can only be smaller in the second case. Now consider $i \in S$. Writing $\rho = (k - s + 1)/(t + 1)$ in equation (5) we can upper bound the utility of agent i in $\text{Lot}_{p,q}$ by

$$v_i - p + \rho(p - q) = (1 - \rho)(v_i - p) + \rho(v_i - q) \leq (v_i - p) + \frac{k}{s+t} \cdot (v_i - q),$$

which is the combined expected utility that the agent obtains from participating in both a k -unit p -lottery (with $s \leq k$) and a k -unit q -lottery. \square

Corollary 5.9. *For every valuation profile \mathbf{v} , the benchmark \mathcal{G} is at most twice the expected residual surplus of the best one-level lottery: $\mathcal{G}(\mathbf{v}) \leq 2 \cdot \sup_r \text{Lot}_r(\mathbf{v})$.*

The following auction does market analysis on the fly to identify and run a good one-level lottery. We have deliberately avoided optimizing the parameters of this mechanism in order to keep its description and analysis as simple as possible.

Definition 5.10. *The k -unit Random Sampling Optimal Lottery (RSOL) mechanism works as follows.*

1. *Partition the agents uniformly at random into a market M and a sample S , i.e., each agent is in S or M independently with probability $1/2$ each.*
2. *Calculate the optimal k -unit lottery price r_S for the sample: $r_S = \text{argmax}_r \text{Lot}_r(\mathbf{v}_S)$.*
3. *Run the k -unit r_S -lottery on the market M ; reject the agents in the sample S .*

We show that this RSOL mechanism gives a good approximation to the residual surplus of the optimal one-level lottery unless a majority of its residual surplus is derived from the highest-valued agent. If a majority of its residual surplus is derived from the highest-valued agent, then the k -unit Vickrey auction is a good approximation of the benchmark. Therefore, mixing between the two auctions gives a platform that is universally adopted with constant competitive advantage (independent of k and n).

Theorem 5.11. *For every $n, k \geq 1$, there is an n -agent k -unit platform mechanism that is universally adopted with constant competitive advantage.*

A key fact that enables the analysis of RSOL is that, with constant probability, the relevant statistical properties of the full valuation profile are preserved in the market and the sample. These statistical properties can be summarized in terms of a “balance” condition. Define a partition of the agents $\{1, 2, 3, \dots, n\}$ into a market M and a sample S to be *balanced* if $1 \in M$, $2 \in S$, and for all $i \in \{3, \dots, n\}$, between $i/4$ and $3i/4$ of the i highest-valued agents are in S (and similarly M). In the proof of Theorem 5.11, we use the following adaptation of the “Balanced Sampling Lemma” of Feige et al. (2005) to bound from below the probability that RSOL selects a balanced partitioning.

Lemma 5.12. *When each agent is assigned to the market M or sample S independently according to a fair coin, the resulting partitioning is balanced with probability at least 0.169.*

For completeness, we include a proof of Lemma 5.12 in Appendix C. We now turn to Theorem 5.11.

Proof of Theorem 5.11. We outline the high-level argument and then fill in the details. We focus on the expected residual surplus of RSOL, where the expectation is over the random partition of agents, relative to that of an optimal one-level lottery, on the “truncated” valuation profile $\mathbf{v}^{(2)} = (v_{(2)}, v_{(2)}, v_{(3)}, \dots, v_{(n)})$. We only track the contributions to RSOL’s expected residual surplus when the partitioning of the agents is balanced. In such cases, RSOL’s residual surplus on the truncated valuation profile can only be less than on the original one.

Step 1 of the analysis proves that, conditioned on the partitioning of the agents being balanced, the expected residual surplus of the optimal one-level lottery for the sample is at least $1/2$ times that of the optimal one-level lottery for the full truncated valuation profile. Step 2 of the analysis proves that, conditioned on an arbitrary balanced partition, the residual surplus of every one-level lottery on the market is at least $1/9$ times its residual surplus on the sample. In particular, this inequality holds for the optimal one-level lottery for the sample. Combining these two steps with Lemma 5.12 implies that the expected residual surplus of RSOL is at least $0.169 \times \frac{1}{2} \times \frac{1}{9} \geq 1/107$ times that of the optimal one-level lottery on the truncated valuation profile $\mathbf{v}^{(2)}$. The additional residual surplus achieved by an optimal one-level lottery on the original valuation profile \mathbf{v} over the truncated one is at most $v_{(1)} - v_{(2)}$. The residual surplus of the $(k + 1)$ th-price auction, where k is the number of units for sale, is at least this amount. The platform mechanism that mixes between RSOL with probability $107/108$ and the $(k + 1)$ th-price auction with probability $1/108$ has expected residual surplus at least $1/108$ times that of the optimal one-level lottery on \mathbf{v} , and (by Corollary 5.9) at least $1/216$ times the benchmark \mathcal{G} . Below, we elaborate on the two steps described above.

Step 1: *Conditioned on a balanced partitioning, the expected residual surplus of the optimal one-level lottery for the sample S is at least $1/2$ times that of the optimal one-level lottery for the full truncated valuation profile.* Let r be the price of the optimal one-level lottery for $\mathbf{v}^{(2)}$. Conditioned on a balanced partition, exactly one of the top two (equal-valued) bidders of $\mathbf{v}^{(2)}$ lies in S . By symmetry, each other bidder has probability $1/2$ of lying in S . The winning probability of bidders in S with value at least r is only higher than that when all agents are present. Summing over the bidders’ contributions to the residual surplus and using the linearity of expectation, $\mathbf{E}_S[\text{Lot}_r(S) \mid \text{balanced partition}] \geq \text{Lot}_r(\mathbf{v}^{(2)})/2$. Of course, the optimal one-level lottery for the sample is only better.

Step 2: *Conditioned on an arbitrary balanced partition, for the truncated valuation profile $\mathbf{v}^{(2)}$, the residual surplus of every one-level lottery on the market is at least $1/9$ times its residual surplus on the sample.* Fix a balanced partition into S and M and a one-level lottery at price r . The expected contribution of a bidder j to a r -lottery is $(v_j - r)$ times its winning probability (if $v_j > r$) or 0 (otherwise). The balance condition ensures that, for every $i \geq 2$, the number of the i highest-valued bidders that belong to the market is between $1/3$ and 3 times that of the sample. In particular, the winning probability of bidders with value at least r in M is at least $1/3$ of that of such bidders in S . Moreover, the balance condition implies that

$$\sum_{j \in M} \max\{v_j - r, 0\} \geq \frac{1}{3} \sum_{j \in S} \max\{v_j - r, 0\}$$

for the truncated valuation profile $\mathbf{v}^{(2)}$; the claim follows. \square

It is certainly possible to optimize better the parameters of the platform mechanism defined in the proof of Theorem 5.11. Furthermore, since for simplicity we only keep track of RSOL’s performance when the partition is balanced, the mechanism’s performance is better than the proved bound.

6 Platform Design and Prior-Free Profit Maximization

While the objective of profit maximization is not central to this paper, there have been a number of studies of prior-free mechanisms for profit maximization that are relevant to platform design. This section discusses digital good settings (Section 6.1), multi-unit settings (Section 6.2), and more general settings (Section 6.3). We describe these results using the terminology of platform design. An important goal of our discussion is to compare our performance benchmark, which is justified by Bayesian foundations, with the prior-free benchmarks employed in this literature.

6.1 Digital Good Settings

The simplest setting for platform design is that of a *digital good*, i.e., a multi-unit setting with the same number $k = n$ of units as (unit-demand) agents. This environment admits a trivial optimal mechanism for surplus and residual surplus (serve all agents for free); but for profit maximization, designing a good platform mechanism is a challenging problem.

The Bayesian optimal mechanism for a digital good when values are drawn i.i.d. from the distribution F simply posts the monopoly price for F , i.e., an r that maximizes $r(1 - F(r))$. In the language of the preceding sections, this optimal mechanism can be viewed as an r -lottery. The performance benchmark described in Section 4 simplifies to

$$\mathcal{G}(\mathbf{v}) = \max_i iv_{(i)}. \tag{7}$$

For $n = 1$ agent, the benchmark (7) equals the surplus and, as we concluded in Section 2, it cannot be well approximated by any platform mechanism. Because of this technicality, the benchmark $\mathcal{G}^{(2)}$ to which prior-free digital good auctions have been compared (e.g., Goldberg et al., 2006) explicitly excludes the possibility of deriving all its profit from one agent:

$$\mathcal{G}^{(2)}(\mathbf{v}) = \max_{i \geq 2} iv_{(i)}. \tag{8}$$

Therefore, up to the technical difference between benchmarks (7) and (8), the prior-free literature for digital goods is compatible with our framework for platform design. Some notable results in this literature are as follows. For reasons we explain shortly, we refer to the approximation of $\mathcal{G}^{(2)}$ as giving near-universal adoption. Optimal platform mechanisms are given in Goldberg et al. (2006) and Hartline and McGrew (2005) for two and three-player digital goods settings, where the competitive advantages for near-universal adoption are precisely 2 and 13/6, respectively. As the number n of agents tends to infinity, Goldberg et al. (2006) show that there is no platform mechanism that is near-universally adopted with competitive advantage less than 2.42; and Chen et al. (2014) show that there exists a mechanism that matches this bound. This optimal platform mechanism is fairly complex; Hartline and McGrew (2005) had previously given a simple mechanism that is near-universally adopted with competitive advantage 3.25.

The benchmark $\mathcal{G}^{(2)}$ does not satisfy our most basic requirement for benchmarks: there exist distributions for which the expected Bayesian optimal revenue exceeds the expected value of the benchmark.⁷ Therefore, mechanisms that approximate the benchmark may not be universally adopted. This problem is not an artifact of the $\mathcal{G}^{(2)}$ benchmark and is inherent to profit maximization: pathological distributions show that there is *no* benchmark \mathcal{G}' and constant β that satisfy $\beta \mathbf{E}_{\mathbf{v}}[\text{Opt}_F(\mathbf{v})] \geq \mathbf{E}_{\mathbf{v}}[\mathcal{G}'(\mathbf{v})] \geq \mathbf{E}_{\mathbf{v}}[\text{Opt}_F(\mathbf{v})]$ for every distribution F .

For the profit objective, the requirement of universal adoption can be relaxed to adoption for every distribution in a large permissive class of distributions. Approximation of the benchmark $\mathcal{G}^{(2)}$ implies such a near-universal adoption, in the following sense.

Proposition 6.1. *If mechanism \mathcal{M} is a β approximation to $\mathcal{G}^{(2)}$ on all valuation profiles, i.e., $\mathcal{M}(\mathbf{v}) \geq \mathcal{G}^{(2)}(\mathbf{v})/\beta$ for every \mathbf{v} , then \mathcal{M} is adopted with competitive advantage β on distributions F with $\mathbf{E}_{\mathbf{v}}[\mathcal{G}^{(2)}(\mathbf{v})] \geq \mathbf{E}_{\mathbf{v}}[\text{Opt}_F(\mathbf{v})]$.*

Proposition 6.1 has bite in that it is satisfied by most relevant distributions. The following lemma, which we prove in Appendix D, gives a sufficient condition for the distribution. Intuitively, this condition states that the revenue from posting a price does not drop too quickly as that price is lowered, and it is a strict generalization of the regularity condition of Myerson (1981). This condition is not satisfied in the bad example above, as most of the optimal revenue is derived from one high-valued agent.

Lemma 6.2. *For digital good settings and every distribution F with $v(1 - F(v))/F(v)$ non-increasing, $\mathbf{E}_{\mathbf{v}}[\mathcal{G}^{(2)}(\mathbf{v})] \geq \mathbf{E}_{\mathbf{v}}[\text{Opt}_F(\mathbf{v})]$.*

6.2 Multi-unit Settings

We next consider maximizing profit in a k -unit auction with unit-demand bidders. We assume throughout that $k \geq 2$. We next define a variant of the performance benchmark $\mathcal{G}^{(2)}$ for platform design and compare it to the benchmark $\mathcal{F}^{(2)}(\mathbf{v}) = \max_{2 \leq i \leq k} iv_{(i)}$ that has been employed, without formal justification, in previous work on prior-free multi-unit auctions.

The benchmark \mathcal{G} defined as the supremum of Bayesian optimal mechanisms is, by Theorem 4.3, equivalent to the supremum over two-level lotteries (which need not sell all units). Two-level lotteries are not useful for profit-maximization in digital goods settings, where a (p, q) -lottery is equivalent to a q -lottery which is equivalent to a q price posting. They are useful in limited supply settings, however. The performance benchmark $\mathcal{G}^{(2)}$ is defined by $\mathcal{G}^{(2)}(\mathbf{v}) = \mathcal{G}(v_{(2)}, v_{(2)}, v_{(3)}, \dots, v_{(n)})$. For every valuation profile, the benchmark $\mathcal{G}^{(2)}$ is at most twice the value of $\mathcal{F}^{(2)}$ (cf. Lemma 5.8). Thus, every multi-unit auction that β -approximates the benchmark $\mathcal{F}^{(2)}$ also 2β -approximates the benchmark $\mathcal{G}^{(2)}$. As in Section 6.1, approximation of the benchmark $\mathcal{G}^{(2)}$ implies approximation of the optimal expected revenue in every Bayesian setting with a non-pathological distribution.

The above discussion provides Bayesian foundations for the benchmark $\mathcal{F}^{(2)}$, and translates the known results for approximating that benchmark into good platform designs. Specifically, every digital good auction that β -approximates the $\mathcal{F}^{(2)}$ benchmark — equivalently for a digital good, the $\mathcal{G}^{(2)}$ benchmark — can be easily converted into a multi-unit auction that β -approximates the $\mathcal{F}^{(2)}$ benchmark (Goldberg et al., 2006) and hence 2β -approximates the $\mathcal{G}^{(2)}$ benchmark. The optimal

⁷For example, consider n agents, each having value n^2 with probability $1/n^2$ and 0 otherwise. The expected revenue of a Bayesian optimal mechanism is n . The expected value of the benchmark $\mathcal{G}^{(2)}$ is bounded above by a constant, independent of n .

platform mechanism for digital goods from Chen et al. (2014) can be thus converted into a platform mechanism for limited supply that is near-universally adopted with competitive advantage 4.84, i.e., it is a 4.84-approximation to $\mathcal{G}^{(2)}$.

6.3 General Environments

The approach to Bayesian optimal mechanism design discussed in Section 3 characterizes optimal mechanisms beyond just multi-unit settings. For every single-parameter setting, where agents want service and there is a feasibility constraint over the set of agents that can be simultaneously served, the optimal mechanism is the ironed virtual surplus maximizer.

Our benchmark \mathcal{G} is difficult to analyze beyond multi-unit settings. In follow-up work to this paper, Hartline and Yan (2011) gave a refinement of our benchmark using the notion of envy-freedom. For instance, when the set system that constrains feasible outcomes satisfies a substitutes condition (formally: the set of feasible outcomes are the independent sets of a matroid set system), their benchmark is at least as large as ours. They give a mechanism that is similar to the RSOL (Definition 5.10) that, for these set systems, is near-universally adopted with constant competitive advantage.

7 Discussion

We defined an analysis framework for platform design based on relative approximation of a performance benchmark. Auctions that approximate this benchmark are simultaneously near-optimal in every Bayesian setting with i.i.d. bidder valuations. Optimizing within this analysis framework suggests novel multi-unit auction formats, different from those suggested by Bayesian analysis. The framework is flexible and permits several extensions and modifications, discussed next.

We focused on platform design for the objectives of residual surplus (Section 5) and profit maximization (Section 6), but our platform design approach extends beyond these objectives. As an example, imagine the k -unit auction in an i.i.d. Bayesian setting where the optimal solution is characterized by optimizing the ironed virtual value corresponding to “a 8% government sales tax.” The objective is then the value of the agents and mechanism less the tax deducted by government, and the corresponding virtual value function has the form $\varphi(v) = 0.92v - 0.08\frac{1-F(v)}{f(v)}$. The optimal k -unit (p, q) -priority lottery remains the appropriate benchmark for this and every other linear objective. For every linear objective with $\gamma_p < 0$, see equation (1), there is always an optimal (p, q) -priority lottery with p and q at most the second highest value $v_{(2)}$, and the mechanism RSOL of Section 5.2 can be mixed with a Vickrey auction to approximate the benchmark.

Optimal mechanisms are ironed virtual surplus maximizers in every single-parameter Bayesian setting with independent private values (Myerson, 1981), not just in the multi-unit auction settings studied here. Examples of more general single-parameter settings include constrained matching markets, single-minded combinatorial auctions, and public projects. The performance benchmark (Definition 4.1) can again be defined pointwise as the supremum over the performance of ironed virtual surplus maximizers on a given valuation profile. As discussed in Section 6.3, this performance benchmark seems hard to characterize beyond multi-unit settings (cf., Theorem 4.3). The follow-up work of Hartline and Yan (2011) recently proposed an alternative benchmark based on envy freedom. This benchmark has structure similar to that of Bayesian optimal auctions and, for this reason, is analytically tractable. Hartline and Yan (2011), for the profit objective and the envy-free

benchmark, give platform mechanisms that are near-universally adopted with constant competitive advantage in general settings.

We defined the performance benchmark as the supremum performance of (symmetric) mechanisms that are optimal in some Bayesian setting with i.i.d. valuations. Clearly, one can define such a benchmark with respect to *any* class of mechanisms. As an example alternative, consider residual surplus maximization and the class of symmetric mechanisms that are Bayesian optimal for some i.i.d. distribution that is either MHR or anti-MHR — that is, the class consisting solely of the Vickrey auction and the (zero-price) lottery. This class arises naturally when domain knowledge suggests that only MHR and anti-MHR distributions are relevant, or if outside consultants are only equipped to design optimal mechanisms for these cases. Specializing to the two-bidder single-item case studied in Section 5.1, the platform design benchmark decreases from $\max\{\frac{v_1+v_2}{2}, v_{(1)} - \frac{v_{(2)}}{2}\}$ to $\max\{\frac{v_1+v_2}{2}, v_{(1)} - v_{(2)}\}$. Reworking the analysis of that section for this new benchmark shows that the optimal mechanism remains a ratio auction, just with a different setting of the parameters (namely, $\alpha = 3$ and $\chi = 4/5$, for an approximation ratio of $5/4$). Notably, the format of the optimal platform is robust to this particular change in the benchmark.

The mechanism in Section 5.2 demonstrates the existence of platform mechanisms for residual surplus maximization that are universally adopted with constant competitive advantage, independent of the number of units and bidders. No standard auction, meaning an ironed virtual surplus maximizer, enjoys such a guarantee (Appendix B). Developing our understanding of n -player platform design further is an interesting research direction. For starters, there should be a much tighter analysis of our mixture of RSOL and Vickrey auctions. One avenue for improvement is to track contributions to the residual surplus on the 83% of the probability space for which the chosen partition is not balanced. The “average balance” approach of Alaei et al. (2009), used previously to improve over the balanced partition technique of Feige et al. (2005) in a profit-maximization context, can be used to give such an improved bound. A second idea is to compare the mechanism’s residual surplus directly to the benchmark in Theorem 4.3, rather than to the simpler “approximate benchmark” in Corollary 5.9. Analogously, avoiding the approximate benchmark $\mathcal{F}^{(2)}$ could lead to better profit-maximizing platform designs (see Section 6.2).

There are surely platforms that are universally adopted with smaller competitive advantage than is required by that in Section 5.2. A challenging problem is to characterize optimal platforms for the residual surplus objective. Even the case of three-bidder single-item settings appears challenging. We conjecture that, for residual surplus maximization with n bidders and a single item, the minimum competitive advantage required by an optimal platform for universal adoption is precisely the expected value of the performance benchmark when bidders’ valuations are drawn i.i.d. from the exponential distribution (as in Theorem 4.3).

While solving for optimal platforms is interesting theoretically, we suspect that optimal platforms will suffer from some drawbacks. First, when the number of agents is large, the optimal platform is a complex object, perhaps a distribution over a very large number of different auctions. This complexity is characteristic of exact optimization in any auction analysis framework; other well-known examples include profit-maximizing auctions in Bayesian single-item settings when bidders’ valuations are not identical or are i.i.d. from an irregular distribution (Myerson, 1981). The complexity of optimal auctions motivates the design and analysis of platforms that are relatively simple while requiring a competitive advantage that is almost as small as the minimum possible, e.g., in the spirit of Bulow and Klemperer (1996) and Hartline and Roughgarden (2009). Second, in optimizing a min-max criterion, the optimal platform will equalize the approximation factor

of the benchmark across all valuation profiles (cf., the proof of Lemma 5.2). In practice there might be agreed-upon “common inputs” and “rare inputs” — without there necessarily being a full prior — with auction performance on common inputs being the most important. For example, the random-sampling-based auctions of Balcan et al. (2008) out perform the optimal platform mechanism (Goldberg et al., 2006) on a family of common inputs.

An auction that approximates the performance benchmark is simultaneously near-optimal in every Bayesian setting with i.i.d. bidder valuations. The converse need not hold, and an interesting research direction is to better understand the relationship between these two conditions. Sometimes, as with the monopoly pricing problem studied in Section 2, simultaneous Bayesian near-optimality is as hard as approximation of the performance benchmark.⁸ This is not always the case, however. For example, Dhangwatnotai et al. (2014) showed that the digital good auction that partitions the agents into pairs and runs a Vickrey auction to serve one agent in each pair obtains a 2-approximation to the revenue of the Bayesian optimal mechanism whenever the distribution is *regular*, meaning that virtual values are increasing (cf. Section 3 and Appendix D). The proof of this 2-approximation is a simple consequence of the $n = 1$ special case of the main theorem of Bulow and Klemperer (1996), i.e., that the 2-agent Vickrey auction obtains more revenue than monopoly pricing a single agent. As mentioned in Section 6.1, no auction for a digital good achieves a 2-approximation of the benchmark $\mathcal{G}^{(2)}$ (Goldberg et al., 2006). On the other hand, all work thus far on simultaneous Bayesian near-optimality that avoids the pointwise benchmark approach — termed “prior-independent guarantees” by Dhangwatnotai et al. (2014) — are confined to regular distributions (Dhangwatnotai et al., 2014; Devanur et al., 2011; Roughgarden et al., 2012). By contrast, our benchmark approximations directly imply prior-independent guarantees for most distributions (for profit maximization) and for all distributions (for residual surplus maximization).

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*, pages 187–196.
- Balcan, M.-F., Blum, A., Hartline, J., and Mansour, Y. (2008). Reducing mechanism design to algorithm design via machine learning. *Journal of Computer and System Sciences*, 74(8):1245–1270.
- Baliga, S. and Vohra, R. (2003). Market research and market design. *Advances in Theoretical Economics*, 3. Article 5.
- Bhattacharya, S., Koutsoupias, E., Kulkarni, J., Leonardi, S., Roughgarden, T., and Xu, X. (2013). Near-optimal multi-unit auctions with ordered bidders. In *Proceedings of the 14th ACM conference on Electronic Commerce*, pages 91–102.
- Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86(1):180–194.

⁸For a value of $h \geq 1$, consider the set of distributions that are concentrated at a single point in $[1, h]$. For each such distribution, the corresponding optimal auction extracts full surplus. As in Section 2, no single auction can extract more than a $1/\ln h$ fraction of the surplus for every such distribution.

- Bulow, J. and Roberts, J. (1989). The simple economics of optimal auctions. *The Journal of Political Economy*, 97:1060–1090.
- Chakravarty, S. and Kaplan, T. R. (2006). Manna from heaven or forty years in the desert: Optimal allocation without transfer payments. Working paper.
- Chakravarty, S. and Kaplan, T. R. (2013). Optimal allocation without transfer payments. *Games and Economic Behavior*, 77(1):1–20.
- Chen, N., Gravin, N., and Lu, P. (2014). Optimal competitive auctions. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 253–262. ACM.
- Condorelli, D. (2007). Weak cartels at standard auctions. Working paper.
- Condorelli, D. (2012). What money can’t buy: Efficient mechanism design with costly signals. *Games and Economic Behavior*, 75(2):613–624.
- Devanur, N., Hartline, J., Karlin, A., and Nguyen, T. (2011). Prior-independent multi-parameter mechanism design. In *Proceedings of the 7th Workshop on Internet and Network Economics*, pages 122–133.
- Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2014). Revenue maximization with a single sample. *Games and Economic Behavior*. To appear.
- Feige, U., Flaxman, A., Hartline, J., and Kleinberg, R. (2005). On the Competitive Ratio of the Random Sampling Auction. In *Proceedings of the 1st Workshop on Internet and Network Economics*, pages 878–886.
- Goldberg, A. V., Hartline, J. D., Karlin, A., Saks, M., and Wright, A. (2006). Competitive auctions. *Games and Economic Behavior*, 55(2):242–269.
- Goldberg, A. V., Hartline, J. D., and Wright, A. (2001). Competitive auctions and digital goods. In *Proc. 12th ACM Symp. on Discrete Algorithms*, pages 735–744.
- Guo, M. and Conitzer, V. (2009). Worst-case optimal redistribution of VCG payments in multi-unit auctions. *Games and Economic Behavior*, 67(1):69–98.
- Hartline, J. and McGrew, R. (2005). From optimal limited to unlimited supply auctions. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 175–182.
- Hartline, J. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, pages 225–234.
- Hartline, J. and Yan, Q. (2011). Envy, truth, and profit. In *Proceedings of the 12th ACM conference on Electronic Commerce*, pages 243–252.
- McAfee, P. and McMillan, J. (1992). Bidding rings. *American Economic Review*, 82(3):579–599.
- Moulin, H. (2009). Almost budget balanced VCG mechanisms to assign multiple objects. *Journal of Economic Theory*, 144(1):96–119.
- Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.




Figure 3: Construction of the cumulative distribution function F with virtual value function (for utility) that is piecewise constant on $[0, 1)$, $[1, 2)$, and $[2, \infty)$, with virtual values 1, 2, and 1, respectively.

Riley, J. and Samuelson, W. (1981). Optimal auctions. *American Economic Review*, 71(3):381–392.

Roughgarden, T., Talgam-Cohen, I., and Yan, Q. (2012). Supply-limiting mechanisms. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 844–861.

Segal, I. (2003). Optimal pricing mechanisms with unknown demand. *American Economic Review*, 93(3):509–529.

Wilson, R. (1987). Game theoretic analysis of trading processes. In Bewley, T., editor, *Advances in Economic Theory*. Cambridge University Press.

A Distribution Construction from Virtual Values

Section 3 gives a formula for calculating an agent’s virtual value function from the distribution from which her value is drawn. For the residual surplus objective, this formula is $\vartheta(v) = \frac{1-F(v)}{f(v)}$. This section reverses the calculation and gives a constructive proof that every non-negative piecewise constant function arises as the virtual value (for utility) function for some distribution. This fact is alluded to in Section 4 and is used explicitly in the proof of Theorem 5.6 in Appendix B.

First observe that the exponential distribution has constant virtual value equal to its mean. That is, the exponential distribution with mean μ has rate $1/\mu$, cumulative distribution $F_\mu(z) = 1 - e^{-z/\mu}$, and virtual value $\vartheta_\mu(z) = \mu$.

Now consider a non-negative piecewise constant function $\vartheta : [0, \infty) \rightarrow \mathbb{R}_+$, where the boundaries of each interval are given by $a_0 = 0, a_1, \dots$, and where the value of the function on the interval $[a_j, a_{j+1})$ is ϑ_j . We construct the distribution F with virtual value function $\vartheta(\cdot)$ inductively. The starting interval is given by the distribution function $F(z) = F_{\vartheta_0}(z)$ for $z \in [a_0, a_1]$. With the first $j - 1$ intervals defined by $F(z)$ for $z \in [a_0, a_j]$, we define the distribution function for the j th interval by $F(a_j + z) = F_{\vartheta_j}(F_{\vartheta_j}^{-1}(F(a_j)) + z)$ for $a_j + z \in [a_j, a_{j+1}]$. Intuitively, this construction does a horizontal shift of the distribution function of the exponential distribution with mean μ_j so that its height matches the height of the constructed (so far) distribution function at a_j . This construction is illustrated in Figure 3.

B Inadequacy of Standard Auctions

This section establishes the limitations of *standard auctions*, i.e., ironed virtual value maximizers (including lotteries and the Vickrey auction), as platform mechanisms. Section 5.1 shows that with $n = 2$ agents, the optimal platform mechanism is not a mixture of standard auctions. Here we show that, even with $k = 1$ unit, there is no finite competitive advantage for which a mixture of standard auctions is universally adopted for all n . This result contrasts with Theorem 5.11, which gives a (non-standard) platform mechanism that is universally adopted with a constant competitive advantage, independent of n and k .

Our argument uses the distributions in the construction in Appendix A; we next note some of their salient properties. These distributions are piece-wise exponential distributions, with piece-wise constant virtual values for utility and piece-wise constant hazard rates. Recall that the virtual value for an exponential distribution equals its expected value which equals the reciprocal of its hazard rate. Also, exponential distributions are memoryless: given that the value v from an exponential distribution is at least z , the conditional distribution of v is identical to that of $z + w$ where w is exponential with the same rate. Of particular relevance, the probability that an exponential random variable with mean one exceeds β is $e^{-\beta}$, and the probability that an exponential random variable with mean β exceeds β is $1/e$. For piece-wise exponential distributions, these properties hold within each piece. From the analysis of Section 3, the expected residual surplus of any mechanism \mathcal{M} on distribution F is equal to its expected virtual surplus.

Theorem 5.6. *For every $\rho > 1$ there is a sufficiently large n such that, for an n -agent, 1-unit setting, no mixture over standard auctions is universally adopted with competitive advantage ρ .*

Proof. Define β to be an integer greater than or equal to $\max\{24 \cdot \rho, 2\}$. For $\kappa \in \{0, 1, 2, \dots, \beta - 1\}$, let $F_{\kappa, \beta}$ denote the piece-wise exponential distribution (as in Appendix A) with virtual value for utility equal to 1 everywhere except on the interval $[\kappa\beta, \kappa\beta + \beta)$, where it is equal to β . Such a distribution thus has a “special interval” where the hazard rate is relatively low (and virtual value is relatively high). Let \mathcal{C}_β denote the set of distributions $\{F_{\kappa, \beta}\}_{\kappa=0}^{\beta-1}$. Consider a setting with a single item and $n = e^{\beta^2}$ agents (rounded up to the nearest integer).

We claim the following:

- (a) For every $F_{\kappa, \beta} \in \mathcal{C}_\beta$, there is an auction with expected residual surplus at least $\beta/4$.
- (b) For every standard auction \mathcal{A} , if $F_{\kappa, \beta} \in \mathcal{C}_\beta$ is chosen uniformly at random, then the expected residual surplus of \mathcal{A} , over the choice of $F_{\kappa, \beta}$ and valuations $v_1, \dots, v_n \sim F_{\kappa, \beta}$, is at most 6.

Claims (a) and (b) imply the theorem. To see this, the residual surplus of a convex combination of mechanisms \mathcal{M} is the convex combination of their residual surpluses. As the inequality of property (b) holds for each auction in the support of such a convex combination, it also holds for the combination. Taking expectation over F uniform from \mathcal{C}_β in inequality (4) from Section 4 we have,

$$\mathbf{E}_{\mathbf{v}, F}[\mathcal{G}(\mathbf{v})] \geq \mathbf{E}_{\mathbf{v}, F}[\text{Opt}_F(\mathbf{v})] \geq \frac{\beta}{24} \mathbf{E}_{\mathbf{v}, F}[\mathcal{M}(\mathbf{v})].$$

By the definition of expectation, there must exist a valuation profile \mathbf{v} that achieves this separation, i.e., with $\mathcal{G}(\mathbf{v}) \geq \frac{\beta}{24} \mathcal{M}(\mathbf{v})$. Thus, the competitive advantage needed for universal adoption is at least $\beta/24 \geq \rho$.

We now proceed to the proofs of (a) and (b). Call an agent *high-valued* if her value is at least β^2 . The probability that there is no high-valued agent is at most $1/\beta$.⁹

To prove (a), fix a choice of $\kappa \in \{0, 1, \dots, \beta - 1\}$ and consider the $\kappa\beta$ -lottery. The probability that all agents' values are below $\kappa\beta$ is less than the probability that all agents' values are below β^2 ; thus the probability of a winner in this lottery is at least $1 - 1/\beta \geq 1/2$. Otherwise, the winner is a random agent with value at least $\kappa\beta$. By the memoryless property of exponential distributions, the probability that the winner, with value at least $\kappa\beta$, has value less than $\kappa\beta + \beta$ is $1 - 1/e \geq 1/2$; such a winner has virtual value β . Virtual values for residual surplus are non-negative, so the expected virtual surplus of the $\kappa\beta$ -lottery is at least $\frac{1}{2} \cdot \frac{1}{2} \cdot \beta$, as claimed.

To prove (b), we first warm up by considering the case where \mathcal{A} is an r -lottery. Choose $F_{\kappa, \beta} \in \mathcal{C}_\beta$ uniformly at random. The intuition is that this random choice effectively “hides” the location of the large virtual values.

If $r \geq \kappa\beta + \beta$, then the winner of \mathcal{A} (if any) has virtual value 1. If $r \leq \kappa\beta - \beta$, then by the memoryless property of exponential distributions, the value of a winner is less than $\kappa\beta$ with probability at least $1 - e^{-\beta}$. Thus, the expected virtual value of a winner in this case (if any) is at most $1 + \beta e^{-\beta} \leq 2$. Finally, if $r \in (\kappa\beta - \beta, \kappa\beta + \beta)$, then the virtual value of a winner (if any) is at most β . As the third case occurs with probability at most $2/\beta$ (over the random choice of κ), the expected virtual value of \mathcal{A} is at most $\frac{2}{\beta} \cdot \beta + 1 \cdot 2 = 4$.

We conclude by extending the argument of the preceding paragraph for one-level lotteries to an arbitrary ironed virtual surplus maximizer \mathcal{A} . In our single-item symmetric setting, the auction \mathcal{A} specifies ironed intervals where ties are broken randomly but otherwise awards the item to the agent with the highest value. Let $[r, r']$ denote the ironed interval of \mathcal{A} that contains the value β^2 . (If value β^2 is not ironed, then set $r = r' = \beta^2$.) Valuation profiles without a high-valued bidder occur with probability at most $1/\beta$ (by the above analysis) and give virtual surplus at most β ; thus their contribution to the expected virtual surplus of \mathcal{A} is at most 1. Valuation profiles with a high-valued bidder and highest value in $[r, r']$ contribute the same expected virtual surplus to \mathcal{A} as to a r -lottery; by the previous paragraph, this contribution is at most 4. In every valuation profile with highest value greater than r' , the item is awarded to a bidder with virtual value 1; these profiles contribute at most 1 to the expected virtual surplus of \mathcal{A} . \square

This proof, in fact, gives a lower bound on the competitive advantage for universal adoption of any standard auction that grows proportionally to $\sqrt{\log n}$ with the number n of agents.

C The Balanced Sampling Lemma

Recall that a partitioning of the agents $\{1, 2, 3, \dots, n\}$ into a market M and sample S is *balanced* if $1 \in M$, $2 \in S$, and for all $i \geq 3$, between $i/4$ and $3i/4$ of the i th highest-valued agents are in M (and similarly S). We restate and prove the balanced sampling lemma below.

Lemma 5.12. *When each agent is assigned to the market M or sample S independently according to fair coin, the resulting partitioning is balanced with probability at least 0.169.*

⁹The analysis is elementary. Using the memoryless property of exponential distributions, the probability that a given agent is high-valued is $\eta = (e^{-\beta})^{\beta-1} \cdot e^{-1} = e^{-\beta^2 + (\beta-1)}$. With $n = e^{\beta^2} = e^{\beta-1}/\eta$ agents, the probability of no high-valued agents is $(1 - \eta)^n \leq e^{-\eta n} \leq e^{-e^{\beta-1}} \leq 1/\beta$. The last inequality can be verified by checking the lower endpoint of $\beta = 2$ and comparing the derivatives of $e^{e^{\beta-1}}$ and β .

Proof. Call a subset of the agents *imbalanced* if, for some $i \geq 3$, it contains fewer than $i/4$ of the i highest-valued agents. After conditioning on the events that $1 \in M$ and $2 \in S$, the probability that S is imbalanced can be calculated as at most 0.161 by a simple probability of ruin analysis proposed in Feige et al. (2005) (details given below). By symmetry, the same bound holds for M . By the union bound, the partition is balanced with probability at least 0.678. Agent 1 is in M and 2 is in S with probability $1/4$, so the unconditional probability that the partition is balanced is at least 0.169.

The following analysis from Feige et al. (2005) shows that the conditional probability that S is imbalanced is at most 0.161. Consider the random variable $Z_i = 4|S \cap \{1, \dots, i\}| - i$; the balanced condition is equivalent to $Z_i \geq 0$ for all $i \geq 3$. By the conditioning, $S \cap \{1, 2\} = \{2\}$ and so $Z_2 = 2$. View Z_i as the positions of a random walk on the integers that starts from position two and takes three steps forward (at step i with $i \in S$) or one step back (at step i with $i \notin S$), each with probability $1/2$. The set S is imbalanced if and only if this random walk visits position -1 . The probability r of ever (with $n \rightarrow \infty$) visiting the preceding position in such a random walk can be calculated as the root of $r^4 - 2r + 1$ on the interval $(0, 1)$, which is approximately 0.544. The probability of imbalance, which requires eventually moving backward three steps from position 2, is at most $r^3 \leq 0.161$, as claimed. \square

D Profit Maximization with Near-universal Adoption

Recall from Section 6 the benchmark $\mathcal{G}^{(2)} = \max_{i \geq 2} iv_{(i)}$, which effectively excludes selling to the highest-valued agent at her value. We will show that a mechanism \mathcal{M} that achieves a β -approximation of this benchmark on every valuation profile is near-universally adopted with competitive advantage β , meaning that for every distribution F in a large class, the expected profit of \mathcal{M} is at least a β fraction of that of the Bayesian optimal auction for F . By Proposition 6.1, it suffices to give sufficient condition on F that guarantees that $\mathbf{E}_{\mathbf{v}}[\mathcal{G}^{(2)}(\mathbf{v})] \geq \mathbf{E}_{\mathbf{v}}[\text{Opt}_F(\mathbf{v})]$.

From Bulow and Roberts (1989), virtual values for revenue are given by the *marginal revenue* of the *revenue curve* that plots the revenue $p(1 - F(p))$ against the probability $1 - F(p)$ that the agent buys (i.e., her expected demand). Virtual values are given by the slope of the revenue curve, thus monotonicity of virtual values, as required by the regularity condition of Myerson (1981), is equivalent to the concavity of the revenue curve.

Our sufficient condition, the “inscribed triangle property,” states that for every point $(1 - F(p), p(1 - F(p)))$ on the revenue curve, the triangle formed with the points $(0, 0)$ and $(1, 0)$ lies underneath the revenue curve. This condition is clearly satisfied whenever the revenue curve is concave — equivalently, whenever the distribution is regular — and is also satisfied by a large family of multi-modal distributions that are not regular.

To understand this condition better, observe that, for every distribution F and price p , the line from $(0, 0)$ to $(1 - F(p), p(1 - F(p)))$ lies beneath the revenue curve. The reason is that, for every $\alpha \in [0, 1]$, the price p' with selling probability $\alpha \cdot (1 - F(p))$ is at least p and hence obtains revenue $p'(1 - F(p')) \geq \alpha \cdot p(1 - F(p))$. The inscribed triangle property is therefore equivalent to requiring that, for every price p , the line between $(1 - F(p), p(1 - F(p)))$ and $(1, 0)$ lies beneath the revenue curve. For an economic interpretation of this condition, consider the measure of types that are not served at a price p , i.e., $F(p)$. Viewing the revenue curve as a function of $F(p)$, the condition says that as the price is dropped, the revenue per unit of types that are not served is non-decreasing. In other words, the condition requires $p(1 - F(p))/F(p)$ to be non-increasing.

The inscribed triangle property immediately implies the following lemma, which is reminiscent of the main theorem of Bulow and Klemperer (1996).

Lemma D.1. *For distribution F with non-increasing $v(1 - F(v))/F(v)$, the two-agent Vickrey auction revenue exceeds the single-agent optimal revenue.*

Proof. In the two-agent Vickrey auction, each agent faces a take-it-or-leave-it offer equal to the other agent's bid. Thus, each agent faces a random price p distributed such that the probability of sale to this agent is uniform on $[0, 1]$. Since the revenue of every such price is given by the revenue curve, and the distribution of $1 - F(p)$ is uniform, the expected revenue obtained from the agent equals the area under the revenue curve. Invoking the inscribed triangle property at the point $(1 - F(p^*), p^*(1 - F(p^*)))$ for the monopoly price p^* , we conclude that the expected revenue obtained from one agent is at least $\frac{1}{2} \cdot 1 \cdot p^*(1 - F(p^*))$, half the expected revenue of the monopoly price. Since there are two agents, the total expected Vickrey revenue is at least the optimal single-agent revenue. \square

Lemma D.2. *For distribution F with the non-increasing $v(1 - F(v))/F(v)$ property, conditioning to exceed a price p preserves the property.*

Proof. The original condition is saying that the virtual value is not more negative than the slope of the line that connects that point on the revenue curve to $(1, 0)$. After conditioning on being at least p , the condition can be viewed on the original revenue curve as the virtual value not being more negative than the slope of line that connects the point on the revenue curve to $(1 - F(p), 0)$. As this slope is steeper than the slope of the line through $(1, 0)$, the property is preserved by such conditioning. \square

We now combine the above lemmas to prove Lemma 6.2, restated below. A key intuition in this proof is that Lemma D.1 implies that $\mathbf{E}_{\mathbf{v}}[2v_{(2)}] \geq \mathbf{E}_{\mathbf{v}}[\text{Opt}_F(\mathbf{v})]$ for $k = 2$ items and $n = 2$ agents — the left-hand side is double the revenue of the Vickrey auction (with 1 item) and the right-hand side is double the revenue of the single-agent optimal mechanism.

Lemma 6.2. *For digital good settings and every distribution F with $v(1 - F(v))/F(v)$ non-increasing, $\mathbf{E}_{\mathbf{v}}[\mathcal{G}^{(2)}(\mathbf{v})] \geq \mathbf{E}_{\mathbf{v}}[\text{Opt}_F(\mathbf{v})]$.*

Proof. Let $p^* = \arg\max_p p(1 - F(p))$ be the monopoly price for the distribution. The analysis proceeds by conditioning on $v_{(3)} = z$ and considering the cases where $z \leq p^*$ and $z > p^*$. In the first case, we have

$$\begin{aligned} \mathbf{E}_{\mathbf{v}}[\mathcal{G}^{(2)}(\mathbf{v}) \mid v_{(3)} = z \leq p^*] &\geq \mathbf{E}_{\mathbf{v}}[2v_{(2)} \mid v_{(3)} = z \leq p^*] \\ &\geq \mathbf{E}_{\mathbf{v}}[\text{Opt}_F(\mathbf{v}) \mid v_{(3)} = z \leq p^*]. \end{aligned}$$

The last inequality follows by Lemmas D.1 and D.2 and the fact that, given $v_{(3)} = z \leq p^*$, Opt_F is an auction that sells to at most agents 1 and 2 and therefore has revenue that is at most the optimal auction that sells to these agents for the conditional distribution. In the second case, let $k^* \geq 3$ be a random variable for the number of units sold by Opt_F ; we have

$$\begin{aligned} \mathbf{E}_{\mathbf{v}}[\mathcal{G}^{(2)}(\mathbf{v}) \mid v_{(3)} = z \geq p^*] &\geq \mathbf{E}_{\mathbf{v}}[k^*v_{(k^*)} \mid v_{(3)} = z \geq p^*] \\ &\geq \mathbf{E}_{\mathbf{v}}[k^*p^* \mid v_{(3)} = z \geq p^*] \\ &= \mathbf{E}_{\mathbf{v}}[\text{Opt}_F(\mathbf{v}) \mid v_{(3)} = z \geq p^*]. \end{aligned}$$

Combining the two cases proves the lemma.

□