Optimal Mechanism Design and Money Burning

Jason D. Hartline*
Electrical Engineering and Computer Science
Northwestern University
2133 Sheridan Rd., Evanston, IL 60208
hartline@eecs.northwestern.edu

Tim Roughgarden†
Dept. of Computer Science
Stanford University
353 Serra Mall, Stanford, CA 94305
tim@cs.stanford.edu

ABSTRACT

Mechanism design is now a standard tool in computer science for aligning the incentives of self-interested agents with the objectives of a system designer. There is, however, a fundamental disconnect between the traditional application domains of mechanism design (such as auctions) and those arising in computer science (such as networks): while monetary *transfers* (i.e., payments) are essential for most of the known positive results in mechanism design, they are undesirable or even technologically infeasible in many computer systems. Classical impossibility results imply that the reach of mechanisms without transfers is severely limited.

Computer systems typically do have the ability to reduce service quality—routing systems can drop or delay traffic, scheduling protocols can delay the release of jobs, and computational payment schemes can require computational payments from users (e.g., in spam-fighting systems). Service degradation is tantamount to requiring that users *burn money*, and such "payments" can be used to influence the preferences of the agents at a cost of degrading the social surplus.

We develop a framework for the design and analysis of *money-burning mechanisms* to maximize the residual surplus—the total value of the chosen outcome minus the payments required. Our primary contributions are the following.

- We define a general template for prior-free optimal mechanism design that explicitly connects Bayesian optimal mechanism design, the dominant paradigm in economics, with worst-case analysis. In particular, we establish a general and principled way to identify appropriate performance benchmarks in prior-free mechanism design.
- For general single-parameter agent settings, we characterize the Bayesian optimal money-burning mechanism.
- For multi-unit auctions, we design a near-optimal prior-free money-burning mechanism: for every valuation profile, its expected residual surplus is within a constant factor of our benchmark, the residual surplus of the best Bayesian optimal mechanism for this profile.
- For multi-unit auctions, we quantify the benefit of general transfers over money-burning: optimal money-burning mechanisms always obtain a logarithmic fraction of the full social surplus, and this bound is tight.

Categories and Subject Descriptors

F.0 [**Theory of Computation**]: General

General Terms

Algorithms, Economics, Theory

*This work was done while the author was at Microsoft Research, Silicon Valley.

†Supported in part by NSF CAREER Award CCF-0448664, an ONR Young Investigator Award, and an Alfred P. Sloan Fellowship.

STOC'08, May 17–20, 2008, Victoria, British Columbia, Canada.

1. INTRODUCTION

Mechanism design is now a standard tool in computer science for designing resource allocation protocols (a.k.a. mechanisms) in computer systems used by agents with diverse and selfish interests. The goal of mechanism design is to achieve non-trivial optimization even when the underlying data—the preferences of participants—are unknown a priori. Fundamental for most positive results in mechanism design are monetary transfers (i.e., payments) between participants. For example, in the surplus-maximizing VCG mechanism [35, 7, 20], such transfers enable the mechanism designer to align fully the incentives of the agents with the system's objective.

Most computer systems differ from classical environments for mechanism design, such as traditional markets and auctions, in that monetary transfers are unpopular, undesirable, or technologically infeasible. It is sometimes possible to design mechanisms that eschew transfers completely; see [33] for classical results in economics and [15, 24] for recent applications in interdomain routing. Unfortunately, negative results derived from Arrow's Theorem [3, 17, 32] imply that the reach of mechanisms without transfers is severely limited.

The following observation motivates our work: *computer systems typically have the ability to arbitrarily reduce service quality.* For example, routing systems can drop or delay traffic (e.g. [8]), scheduling protocols can delay the release of jobs (e.g. [6]), and computational payment schemes allow a mechanism to demand computational payments from agents (e.g., in spam-fighting systems [11, 10, 25]).[1] Such service degradation can be used to align the preferences of the agents with the social objective, at a cost: *these "payments" also degrade the social surplus.*

We develop a framework for the design and analysis of *money-burning mechanisms*—mechanisms that can employ arbitrary payments and seek to maximize the *residual surplus*, defined as the total value to the participants of the chosen outcome minus the sum of the ("burnt") payments.[2] Such mechanisms must trade off the social cost of imposing payments with the ability to elicit private information from participants and thereby enable accurate surplus-maximization. For example, suppose we intend to award one of two participants access to a network. Assume that the two agents have valuations (i.e., maximum willingness to pay) v_1 and $v_2 \leq v_1$ for acquiring access, and that these valuations are private (i.e., unknown to the mechanism designer). The *Vickrey* or *second-price auction* [35] would award access to agent 1, charge a payment of v_2, and thereby obtain residual surplus $v_1 - v_2$. A *lottery* would award access to an agent chosen at random, charge nothing, and achieve a (residual) surplus of $(v_1 + v_2)/2$, a better result if and only if $v_1 < 3v_2$. Even in this trivial scenario, it is not clear how to define (let alone design) an optimal money-burning mechanism.[3]

Our goal is to rigorously answer the following two questions:

1. What is the optimal money-burning mechanism?
2. How much more powerful are mechanisms with monetary transfers than money-burning mechanisms?

Our Results.

Our first contribution is to identify a general template for prior-free (i.e., worst-case) optimal mechanism design. The basic idea is to characterize the set of mechanisms that are Bayesian optimal for some i.i.d. distribution on valuations, and then define a prior-free performance benchmark that corresponds to competing simultaneously with all of these on a fixed (worst-case) valuation profile. The template, which we detail below, is general and we expect it to apply in many mechanism design settings beyond money-burning mechanisms.

Second, we characterize Bayesian optimal money-burning mechanisms—the incentive-compatible mechanisms with maximum-possible expected residual surplus. Our characterization applies to general single-parameter agents, meaning that the preferences of each agent is naturally summarized by a single real-valued valuation, with independent but not necessarily identically distributed valuations. The characterization unifies results in the economics literature [5, 26] and also extends them in two important directions. First, the results in [5, 26] concern only multi-unit auctions, where k identical units of an item can be allocated to agents who each desire at most one unit. Our characterization applies to the general, possibly asymmetric, setting of single-parameter agents; for example, agents could be seeking disjoint paths in a multicommodity network.[4] In addition, for multi-unit auctions, we give a simple description of the optimal mechanism even when the "hazard rate" of the valuation distribution is not monotone in either direction. This important case is the most technically interesting and challenging one, and it has not been considered in detail in the literature.

Third, for multi-unit auctions, we design a mechanism that is approximately optimal in the worst case. We derive our benchmark using our characterization of Bayesian optimal mechanisms restricted to i.i.d. valuations and symmetric mechanisms. We prove that such mechanisms are always well approximated by a *k-unit p-lottery*, defined as follows: order the agents randomly, sequentially make each agent a take-it-or-leave-it offer of p, and stop after either k items have been allocated or all agents have been considered. This result reduces the design of a constant-approximation prior-free money-burning mechanism to the problem of approximating the residual surplus achieved by the optimal k-unit p-lottery. Our prior-free mechanism obtains a constant approximation of this benchmark using random sampling to select a good value of p. Surprisingly, we accomplish this even when k is very small (e.g., $k = 1$). Our benchmark definition ensures that such a guarantee is strong: for example, if valuations are drawn from some unknown i.i.d. distribution $\mathbf{F}$, our mechanism obtains a constant fraction of the expected residual surplus of an optimal mechanism tailored specifically for $\mathbf{F}$.

Finally, for multi-unit auctions, we provide a price-of-anarchy-type analysis that measures the social cost of burnt payments. Recall that the full surplus is achievable with monetary transfers using the Vickrey-Clarke-Groves (VCG) mechanism. We prove that the largest-possible relative loss in surplus due to money-burning is precisely logarithmic in the number of participants, in both the Bayesian and worst-case settings. Indeed, our near-optimal money-burning mechanism always obtains residual surplus within a logarithmic factor of the full surplus. This result suggests that the cost of implementing money-burning (e.g., computational payments) rather than general transfers (e.g., micropayments) in a system is relatively modest. Further, our positive result contrasts with the linear lower bound that we prove on the fraction of the full surplus obtainable by mechanisms without any kind of payments.

A Template for Prior-Free Auction Design.

The following template forges an explicit connection between the Bayesian analysis of Bayesian optimal mechanism design, the dominant approach in economics, and the worst-case analysis of prior-free optimal mechanism design, the

[1] Computational payment schemes do not need the infrastructure required by micropayment schemes. One can interpret our results as analyzing the power of computational payments, which were first devised for spam-fighting, in a general mechanism design setting.

[2] We assume that valuations and burnt payments are measured in the same units. In other words, there is a known mapping between decreased service quality (e.g., additional delay) and lost value (e.g., dollars). This mapping can be different for different participants, but it must be publicly known and map onto $[0, \infty)$. See Section 6 for further discussion.

[3] Indeed, it follows from our results that in some settings lotteries are optimal (i.e., money-burning is useless); in others, Vickrey auctions are optimal; and sometimes, neither is optimal.

[4] Multi-unit auctions model symmetric situations, as when each agent seeks a path from a common source s to a common destination t; here, the number k of units equals the number of edges in a minimum s-t cut.

ubiquitous approach in theoretical computer science. Its goal is to fill a fundamental gap in prior-free optimal mechanism design methodology: the selection of an appropriate performance benchmark.

1. Characterize the Bayesian optimal mechanism for every i.i.d. valuation distribution.

2. Interpret the behavior of the symmetric, ex post incentive compatible, Bayesian optimal mechanism for every i.i.d. distribution on an arbitrary valuation profile to give a distribution-independent *benchmark*.

3. Design a single ex post incentive compatible mechanism that approximates the above benchmark on every valuation profile; the performance ratio of such a mechanism provides an upper bound on that of the optimal prior-free mechanism.

4. Obtain lower bounds on the best performance ratio possible in this framework by exhibiting a distribution over valuations such that the ratio between the expected value of the benchmark and the performance of the Bayesian optimal mechanism for the given distribution is large.

In hindsight, this approach has been employed implicitly in the context of (profit-maximizing) digital good auctions [19, 18]. However, the simplicity of the digital good auction problem obscures the importance of the first two steps, as the Bayesian optimal digital good auction is trivial: offer a posted price. For money-burning mechanisms, the benchmark we identify in Step 2 is not a priori obvious.

Further Related Work.

McAfee and McMillan study collusion among bidders in multi-unit auctions [26]. In a *weak cartel*, where the agents wish to maximize the cartel's total utility but are not able to make side payments amongst themselves, payments made to the auctioneer are effectively burnt. The optimization and incentive problem faced by the grand coalition in a multi-unit auction is similar to the auctioneer's problem in our money-burning setting; therefore, results for weak cartels follow from similar analyses to ours [26, 9].

Our characterization of Bayesian optimal money-burning mechanisms builds on analysis tools developed for profit maximization in Bayesian settings (see Myerson [28] and Riley and Samuelson [30]) that apply in general single parameter settings (see, e.g., [22]). Independently from our work, Chakravarty and Kaplan [5] describe the optimal Bayesian auction in multi-unit money-burning settings. Our work extends this analysis to general single-parameter agent settings with explicit focus on the case where the hazard rate is not monotone in either direction. Our paper is the first to study the relative power of money-burning mechanisms and mechanisms with or without transfers. It is also the first to consider prior-free money-burning mechanisms.

Our results that quantify the benefit of transfers have analogs in the price of anarchy literature, specifically in the standard (nonatomic) model of selfish routing (e.g. [31]). Namely, full efficiency is achievable in this model with general transfers, in the form of "congestion prices"; without transfers the outcome is a Nash equilibrium, with inefficiency measured by the price of anarchy; and with burnt transfers ("speed bumps" or other artificial delays) it is generally possible to recover some but not all of the efficiency loss at equilibrium [8].

There are several other studies that view transfers to an auctioneer as undesirable; however, these works are technically unrelated to ours. We already noted recent work on incentive-compatible interdomain routing without payments [15, 24]. Moulin [27] and Guo and Conitzer [21] independently studied how to redistribute the payments of the VCG mechanism in a multi-unit auction among the participants (using general transfers) to minimize the total payment to the auctioneer. Finally, as already mentioned, our prior-free techniques are related to recent work on profit maximization (e.g., [19, 16, 4]) and there is a related literature on the problem of cost minimization, a.k.a. frugality (e.g., [2, 34, 12, 23]).

2. BAYESIAN OPTIMAL MONEY BURNING

In this section we study optimal money-burning mechanism design from a standard economics viewpoint, where agent valuations are drawn from a known *prior distribution*. This will complete the first step of our template for prior-free optimal mechanism design.

Mechanism design basics.

We consider mechanisms that provide a good or service to a subset of n agents. The outcome of such a mechanism is an *allocation vector*, $\mathbf{x} = (x_1, \ldots, x_n)$, where x_i is 1 if agent i is served and 0 otherwise, and a *payment vector*, $\mathbf{p} = (p_1, \ldots, p_n)$. In this paper, the payment p_i is the amount of money that agent i must "burn". We allow the set of feasible allocation vectors, $\mathcal{X}$, to be constrained arbitrarily; for example, in a multi-unit auction with k identical units of an item, the feasible allocation vectors are those $\mathbf{x} \in \mathcal{X}$ with $\sum_i x_i \leq k$.

We assume that each agent i is *risk-neutral*, has a privately known valuation v_i for receiving service, and aims to maximize their (quasi-linear) utility, defined as $u_i = v_i x_i - p_i$. We denote the *valuation profile* by $\mathbf{v} = (v_1, \ldots, v_n)$.

Our mechanism design objective is to maximize the *residual surplus*, defined as

$$\sum_i (v_i x_i - p_i)$$

for a valuation profile $\mathbf{v}$, a feasible allocation $\mathbf{x}$, and payments $\mathbf{p}$. If the payments were transferred to the seller then the resulting *social surplus* would be $\sum_i v_i x_i$; however, in our setting the payments are burnt and the social surplus is equal to the residual surplus.

Bayesian mechanism design basics.

In this section, we assume that the agent valuations are drawn i.i.d. from a publicly known distribution with cumulative distribution function $F(z)$ and probability density function $f(z)$. We let $\mathbf{F}$ denote the joint (product) distribution of agent values. See Section 6 for a generalization to general product distributions.

We consider the problem of implementation in Bayes-Nash equilibrium. Agent i's strategy is a mapping from their private value v_i to a course of actions in the mechanism. The distribution on valuations $\mathbf{F}$ and a strategy profile induce a distribution on agent actions. These agent actions are in *Bayes-Nash equilibrium* if no agent, given their own valuation and the distribution on other agents' actions, can improve its expected payoff via alternative actions. By the *revelation principle* [28], we can restrict our attention to single round, sealed bid, *direct* mechanisms in which *truthtelling*, i.e., submitting a bid b_i equal to the private value v_i, is a Bayes-Nash equilibrium. It will turn out that there is always an optimal mechanism that is not only Bayesian incentive compatible but also *dominant strategy incentive compatible*, meaning truthtelling is an optimal agent strategy for every strategy profile of the other agents.

An *allocation rule*, $\mathbf{x}(\mathbf{v})$, is the mapping (in the truthtelling 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 probability that agent i is allocated when its 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 $p_i(v_i)$. *Positive transfers* from the mechanism to the agents are not allowed and we require *interim individual rationality* (i.e., that non-participation in the mechanism is an allowable agent strategy). The following lemma is the standard characterization of the allocation rules implementable by Bayesian incentive-compatible mechanisms and the accompanying (uniquely defined) payment rule.

LEMMA 2.1. *[28] Every Bayesian incentive compatible mechanism satisfies:*

1. *Allocation monotonicity: for all i and $v_i > v_i'$, $x_i(v_i) \geq x_i(v_i')$.*
2. *Payment identity: for all i and v_i, $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(v)dv$.*

Virtual valuations. Assume for simplicity that the distribution F has support $[a, b]$ and positive density throughout this interval. Myerson [28] defined "virtual valuations" and showed that they characterize the expected payment of an agent in a Bayesian incentive compatible mechanism.

DEFINITION 2.2 (VIRTUAL VALUATION FOR PAYMENT). *If agent i's valuation is distributed according to F, then its* virtual valuation for payment *is $\varphi(v_i) = v_i - \frac{1-F(v_i)}{f(v_i)}$.*

LEMMA 2.3. *[28] In a Bayesian incentive-compatible mechanism with allocation rule $\mathbf{x}(\cdot)$, the expected payment of agent i satisfies $\mathbf{E}_{\mathbf{v}}[p_i(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}[\varphi(v_i)x_i(\mathbf{v})]$.*

Myerson uses this correspondence to design optimal mechanisms for profit-maximization. The optimal mechanism for a given distribution is the one that maximizes the *virtual surplus* (for payment).

DEFINITION 2.4 (VIRTUAL SURPLUS). *For virtual valuation function $\varphi(\cdot)$ and valuations $\mathbf{v}$, the* virtual surplus *of allocation $\mathbf{x}$ is $\sum_i \varphi(v_i)x_i$.*

Our objective is to maximize the residual surplus, $\sum_i(v_i x_i(\mathbf{v}) - p_i(\mathbf{v}))$, which we can do quite easily using virtual valuations. To justify our terminology, below, notice that an agent's utility is $u_i(\mathbf{v}) = v_i x_i(\mathbf{v}) - p_i(\mathbf{v})$, and our objective of residual surplus maximization is simply that of maximizing the expected utility of the agents, $\mathbf{E}_{\mathbf{v}}\left[\sum_i u_i(\mathbf{v})\right]$. We define a *virtual valuation for utility* by simply plugging in the virtual valuation for payments into the equation that defines utility.

DEFINITION 2.5 (VIRTUAL VALUATION FOR UTILITY). *If agent i's valuation is distributed according to F, then its* virtual valuation for utility *is $\vartheta(v_i) = \frac{1-F(v_i)}{f(v_i)}$.*

This quantity is also known as the "information rent" or "inverse hazard rate function". Treating it as a virtual valuation of sorts, we generalize the theory of optimization by virtual valuations, beginning with the following lemma.

LEMMA 2.6. *In a Bayesian incentive-compatible mechanism with allocation rule $\mathbf{x}(\cdot)$, the expected utility of agent i satisfies $\mathbf{E}_{\mathbf{v}}[u_i(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}[\vartheta(v_i)x_i(\mathbf{v})]$.*

We can conclude from this that the Bayesian optimal mechanisms for residual surplus are precisely those that maximize the expected virtual surplus (for utility) subject to feasibility and monotonicity of the allocation rule. In other words, we should choose a feasible allocation vector $\mathbf{x}(\mathbf{v})$ to maximize $\sum_i \vartheta(v_i)x_i(\mathbf{v})$ for each $\mathbf{v}$, subject to monotonicity of $x_i(v_i)$. It is easy to see that if $\vartheta(\cdot)$ is monotone non-decreasing in v_i, then choosing $\mathbf{x}(\mathbf{v}) \in \operatorname{argmax}_{\mathbf{x}' \in \mathcal{X}} \sum_i \vartheta(v_i)x_i'$ results in a monotone allocation rule. Unfortunately $\vartheta(\cdot)$ is often not monotone non-decreasing; indeed, under the standard "monotone hazard rate" assumption, discussed further below, $\vartheta(\cdot)$ is monotone *in the wrong direction*.

Ironing. We next generalize an "ironing" procedure of Myerson [28] that transforms a possibly non-monotone virtual valuation function into an *ironed virtual valuation* function that is monotone; the optimization approach of the previous paragraph can then be applied to these ironed functions to obtain a monotone allocation rule. Further, the ironing procedure preserves the target objective, so that an optimal allocation for the ironed virtual valuations is equal to the optimal monotone allocation for the original virtual valuations.

DEFINITION 2.7 (IRONED VIRTUAL VALUATIONS [28]). *Given a distribution function $F(\cdot)$ with virtual valuation (for utility) 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, and extend to all of $[0, 1]$ by right-continuity.*
5. *Finally, $\bar{\vartheta}(z) = g(F(z))$.*

Step 4 of Definition 2.7 makes sense because G is convex function. Convexity of G also implies that g, and hence $\bar{\vartheta}$, is a monotone non-decreasing function.

The proof Myerson gives for ironing virtual valuations for payments extends simply to any other kind of virtual valuation including our virtual valuations for utility. We summarize this in Lemma 2.8, with the proof deferred to the full version.

LEMMA 2.8. *Let F be a distribution function with virtual valuation function $\vartheta(\cdot)$ and $\mathbf{x}(\mathbf{v})$ a monotone allocation rule. Define G, H, and $\bar{\vartheta}$ as in Definition 2.7. Then*

$$\mathbf{E}_{\mathbf{v}}[\vartheta(v_i)x_i(\mathbf{v})] \leq \mathbf{E}_{\mathbf{v}}\left[\bar{\vartheta}(v_i)x_i(\mathbf{v})\right], \tag{1}$$

with equality holding if and only if $\frac{d}{dv}x_i(v) = 0$ whenever $G(F(v)) < H(F(v))$.

Our main theorem now follows easily.

THEOREM 2.9. *Let F be a distribution function with virtual valuation function $\vartheta(\cdot)$. Define G, H, and $\bar{\vartheta}$ as in Definition 2.7. For valuation profiles drawn from distribution $\mathbf{F}$, the mechanisms that maximize the expected residual surplus are precisely those satisfying*

1. *$\mathbf{x}(\mathbf{v}) \in \operatorname{argmax}_{\mathbf{x}' \in \mathcal{X}} \sum_i \bar{\vartheta}(v_i){x_i}'$ for every $\mathbf{v}$; and*
2. *for all i, $\frac{d}{dv}x_i(v) = 0$ whenever $G(F(v)) < H(F(v))$.*

PROOF. First, there exists a mechanism that satisfies both of the desired properties. To see this, consider an allocation rule that maximizes $\sum_i \bar{\vartheta}(v_i)x_i(\mathbf{v})$ for every $\mathbf{v}$. Such a rule can without loss of generality be a function only of $\bar{\vartheta}(v_i)$ and not of v_i directly. At points v where $G(F(v)) < H(F(v))$, G is locally linear (since it is the convex hull of H) and hence $\bar{\vartheta}(v)$ is locally constant. Thus such an allocation rule will satisfy $\frac{d}{dv}x_i(v) = 0$ at all such points (for all i).

A mechanism that meets both conditions simultaneously maximizes the right-hand side of (1) while satisfying the inequality with equality. Lemmas 2.6 and 2.8 imply that such a mechanism maximizes the expected residual surplus and, conversely, that all optimal mechanisms must meet both conditions. $\square$

Theorem 2.9 shows that maximizing the ironed virtual surplus (for utility) is equivalent to maximizing expected residual surplus subject to incentive-compatibility. Different tie-breaking rules can yield different optimal mechanisms. With symmetric participants (that is, i.i.d. valuations) and a symmetric feasible region (e.g., k-item auctions), it is natural to consider symmetric mechanisms, and these will play a crucial role in our benchmark for prior-free money-burning mechanisms (see Definition 3.1).

Interpretations. To interpret Theorem 2.9, recall that the *hazard rate* of distribution F at v is defined as $\frac{f(v)}{1-F(v)}$. The *monotone hazard rate* (MHR) assumption is that the hazard rate is monotone non-decreasing and is a standard assumption in mechanism design (e.g. [28]). We will analyze this standard setting (MHR), the setting in which the hazard rate is monotone in the opposite sense (anti-MHR), and the setting where it is neither monotone increasing nor decreasing (non-MHR). Notice that the hazard rate function is precisely the reciprocal virtual valuation (for utility) function.

When the valuation distribution satisfies the MHR condition (Figure 1), the ironed virtual valuations (for utility) have a special form: they are constant with value equal to their expectation (proof omitted).

LEMMA 2.10. *For every distribution F that satisfies the monotone hazard rate condition, the ironed virtual valuation (for utility) function is constant with $\bar{\vartheta}(z) = \mu$, where μ denotes the expected value of the distribution.*

Therefore, under MHR the mechanism that maximizes the ironed virtual surplus is the one that *maximizes the ex ante expected surplus*, without asking for bids and without any transfers. For example, in a multi-unit auction with i.i.d. bidders, all agents are equal ex ante, and thus any allocation rule that ignores the bids and always allocates all k units (charging nothing) is optimal.

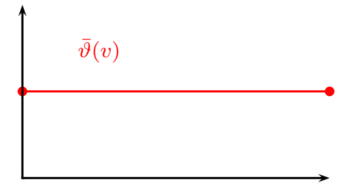

Figure 1: MHR (e.g., uniform): Lottery is optimal.

COROLLARY 2.11. *For agents with i.i.d. valuations satisfying the MHR condition, an optimal (symmetric) money-burning mechanism for allocating k units is a k-unit lottery.*

Suppose the distribution satisfies the anti-MHR condition (Figure 2) which implies that the virtual valuation (for utility) functions are monotone non-decreasing. The ironed virtual valuation function is then identical to the virtual valuation function. The i.i.d. assumption implies that all agents have the same virtual valuation function, so the agents with the highest virtual valuations are also the agents with the highest valuations. Therefore, an optimal money-burning mechanism for allocating k units assigns the units to the k agents with the highest valuations. This is precisely the allocation rule used by the k-unit Vickrey auction [35], so the truthtelling payment rule is that all winners pay the $k+1$st highest valuation.

COROLLARY 2.12. *For agents with i.i.d. valuations satisfying the anti-MHR condition, an optimal (symmetric) money-burning mechanism for allocating k units is a k-unit Vickrey auction.*

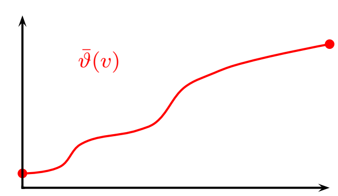

Figure 2: Anti-MHR (e.g., super-exponential): Vickrey is optimal.

To optimally allocate k units of an item in the non-MHR case (Figure 3), we simply award the items to the agents with the largest ironed virtual valuations (for utility). Skipping technical details, we have the following corollary.

COROLLARY 2.13. *For agents with i.i.d. non-MHR valuations, an optimal (symmetric) money-burning mechanism for allocating k units is an indirect k-unit Vickrey auction: for valuations in the range $R = [a, b]$ and subrange $R' \subset R$ on which $\bar{\vartheta}(v)$ has positive slope, it is the indirect mechanism where agents bid $b_i \in R'$ and the k agents with the highest bids win, with ties broken uniformly at random.*

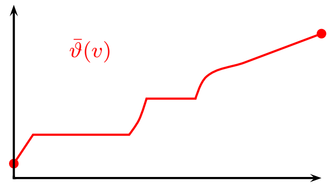

Figure 3: Non-MHR (e.g., bimodal): indirect Vickrey is optimal.

3. PRIOR-FREE MONEY-BURNING MECHANISM DESIGN

We now depart from the Bayesian setting and design near-optimal prior-free mechanisms for multi-unit auctions. Section 3.1 corresponds to the second step in our prior-free mechanism design template and leverages our characterization of Bayesian optimal mechanisms to identify a simple, tight, and distribution-independent performance benchmark. Section 3.3 gives a prior-free mechanism that, for every valuation profile, obtains expected residual surplus within a constant factor of this benchmark. This mechanism implements the third step of our design template. We consider lower bounds on the approximation ratio of all prior-free mechanisms (the final step of the template) in Section 4.

For ease of discussion the payment rules we describe in this section are for mechanism implementations that are dominant strategy incentive compatible for agents that are risk-neutral with respect to randomization in the mechanism, i.e., mechanisms that are *truthful in expectation*. All of these mechanisms have natural implementations with payment rules that make them dominant strategy incentive compatible for any fixed outcome of the mechanism's random decisions, i.e., mechanisms that are *truthful all the time*. In the computer science literature, discussion of these distinctions can be found in [1].

3.1 A Performance Benchmark for Prior-Free Mechanisms

Intuitively, our performance benchmark for a valuation profile is the maximum residual surplus achieved by a symmetric mechanism that is optimal for some i.i.d. distribution. The next definition formalizes the class of mechanisms that define the benchmark.

Definition 3.1 ($\mathrm{Opt}_{\mathbf{F}}$). *For an i.i.d. distribution* $\mathbf{F}$ *with ironed virtual valuation (for utility) function* $\bar{\vartheta}$*, the mechanism* $\mathrm{Opt}_{\mathbf{F}}$ *is defined as follows.*

1. *Given* $\mathbf{v}$*, choose a feasible allocation maximizing* $\sum_i \bar{\vartheta}(v_i)x_i$*. If there are multiple such allocations, choose one uniformly at random.*
2. *Let* $\mathbf{x}$ *denote the corresponding allocation rule, with* $x_i(\mathbf{v})$ *denoting the probability that player* i *receives an item given the valuation profile* $\mathbf{v}$*. Let* $\mathbf{p}$ *denote the (unique) payment rule dictated by Lemma 2.1.*
3. *Given valuations* $\mathbf{v}$ *and the random choice of allocation in the first step, charge each winner* i *the price* $p_i(\mathbf{v})/x_i(\mathbf{v})$ *and each loser 0.*

By Theorem 2.9, $\mathrm{Opt}_{\mathbf{F}}$ maximizes the expected residual surplus for valuations drawn from $\mathbf{F}$. Using Lemma 2.1, it is also incentive-compatible and ex post individually rational. It is symmetric provided the set of feasible allocations is symmetric (i.e., is a k-item auction). In this case, the first step awards the k items to the bidders with the top k ironed virtual valuations (for utility) with respect to the distribution $\mathbf{F}$, breaking ties uniformly at random.

Our benchmark is then:

$$\mathcal{G}(\mathbf{v}) = \sup_{\mathbf{F}} \mathrm{Opt}_{\mathbf{F}}(\mathbf{v}), \tag{2}$$

where $\mathrm{Opt}_{\mathbf{F}}(\mathbf{v})$ denotes the expected residual surplus (over the choice of random allocation) obtained by the mechanism $\mathrm{Opt}_{\mathbf{F}}$ on the valuation profile $\mathbf{v}$. This benchmark is, by definition, distribution-independent. As such, it provides a yardstick by which we can measure prior-free mechanisms: we say that a (randomized) mechanism *β-approximates the benchmark* $\mathcal{G}$ if, for every valuation profile $\mathbf{v}$, its expected residual surplus is at least $\mathcal{G}(\mathbf{v})/\beta$. Note the strength of this guarantee: for example, if a mechanism β-approximates the benchmark $\mathcal{G}$, then on any i.i.d. distribution it achieves at least a β fraction of the expected residual surplus of every mechanism. Naturally, no prior-free mechanism is better than 1-approximate; we give stronger lower bounds in Section 4.

Remark: Restricting attention in Definition 3.1 to optimal mechanisms that use symmetric tie-breaking rules is crucial for obtaining a tractable benchmark. For example, when $\mathbf{F}$ is an i.i.d. distribution satisfying the MHR assumption, Theorem 2.9 implies that *every* constant allocation rule that allocates all items (with zero payments) is optimal (recall Corollary 2.11). For a single-item auction and a valuation profile $\mathbf{v}$, say with the first bidder having the highest valuation, the mechanism that always awards the good to the first bidder and charges nothing achieves the full surplus. (Of course, this mechanism has extremely poor performance on many other valuation profiles.) As no incentive-compatible money-burning mechanism always achieves a constant fraction of the full surplus (see Proposition 5.1), allowing arbitrary asymmetric optimal mechanisms to participate in (2) would yield an unachievable benchmark.

3.2 Multi-Unit Auctions and Two-Price Lotteries

The definition of $\mathcal{G}$ in (2) is meaningful in general single-parameter settings, but appears to be analytically tractable only in problems with additional structure, symmetry in particular. We next give a simple description of this benchmark, and an even simpler approximation of it, for multi-unit auctions.

What does $\mathrm{Opt}_{\mathbf{F}}$ look like for such problems? When the distribution on valuations satisfies the MHR assumption, $\mathrm{Opt}_{\mathbf{F}}$ is a k-unit lottery (cf., Corollary 2.11). Under the anti-MHR assumption, $\mathrm{Opt}_{\mathbf{F}}$ is a k-unit Vickrey auction (cf., Corollary 2.12). We can view the k-unit Vickrey auction, ex post, as a k-unit $v_{(k+1)}$-lottery, where $v_{(k+1)}$ is the $k+1$st highest valuation, in the following sense.

Definition 3.2 (k-unit p-lottery). *The* k-unit p-lottery*, denoted* Lot_p*, allocates to agents with value at least* p *at price* p*. If there are more than* k *such agents, the winning agents are selected uniformly at random.*

One natural conjecture is that, ex post, the outcome of every mechanism of the form $\mathrm{Opt}_{\mathbf{F}}$ on a valuation profile $\mathbf{v}$ looks like a k-unit p-lottery for some value of p. For non-

MHR distributions $\mathbf{F}$, however, $\mathrm{Opt}_{\mathbf{F}}$ can assume the more complex form of a two-price lottery, ex post.

Definition 3.3 (k-unit (p,q)-lottery). *A k-unit (p,q)-lottery, denoted $\mathrm{Lot}_{p,q}$, is the following mechanism. Let s and t denote the number of agents with bid in the range (p,∞) and $(q,p]$, respectively.*

1. *If $s \geq k$, run a k-unit p-lottery on the top s agents.*
2. *If $s+t \leq k$, sell to the top $s+t$ agents at price q.*
3. *Otherwise, run a $(k-s)$-unit q-lottery on the agents with bid in $(q,p]$ and allocate each of the top s agents a good at the price dictated by Lemma 2.1: $\frac{k-s+1}{t+1}q + \frac{s+t-k}{t+1}p$.*

We now prove that for every i.i.d. distribution $\mathbf{F}$ and every valuation profile $\mathbf{v}$, the mechanism $\mathrm{Opt}_{\mathbf{F}}$ results in an outcome and payments that, ex post, are identical to those of a k-unit (p,q)-lottery.

Lemma 3.4. *For every valuation profile $\mathbf{v}$, there is a k-unit (p,q)-lottery with expected residual surplus $\mathcal{G}(\mathbf{v})$.*

Proof. By definition (2), we only need to show that, for every i.i.d. distribution $\mathbf{F}$ and valuation profile $\mathbf{v}$, $\mathrm{Opt}_{\mathbf{F}}(\mathbf{v})$ has the same outcome as a k-unit (p,q)-lottery.

Fix $\mathbf{F}$ and $\mathbf{v}$, and assume that $v_1 \geq \cdots \geq v_n$. Thus, $\bar{\vartheta}(v_1) \geq \cdots \geq \bar{\vartheta}(v_n)$. Recall by Definition 3.1 that $\mathrm{Opt}_{\mathbf{F}}$ maximizes $\sum_i \bar{\vartheta}(v_i)x_i$ and breaks ties randomly. Define $S = \{i : \bar{\vartheta}(v_i) > \bar{\vartheta}(v_{k+1})\}$, $T = \{i : \bar{\vartheta}(v_i) = \bar{\vartheta}(v_{k+1})\}$, $s = |S|$, and $t = |T|$. Assume we are in the more technical case that $0 < s < k < s+t$ (the other cases follow from similar arguments). It is easy to see that $\mathrm{Opt}_{\mathbf{F}}$ assigns a unit to each bidder in S and allocates the remaining $k-s$ units randomly to bidders in T. Let $q = \inf\{v : \bar{\vartheta}(v) = \bar{\vartheta}(v_{k+1})\}$ and $p = \inf\{v : \bar{\vartheta}(v) > \bar{\vartheta}(v_{k+1})\}$. The allocation is thus identical to a k-unit (p,q)-lottery. A straightforward argument (omitted) then shows that the payments are also correct. $\square$

As we have seen, mechanisms of the form $\mathrm{Opt}_{\mathbf{F}}$ can produce outcomes not equivalent to that of a single-price lottery. Our next lemma shows that k-unit p-lotteries give 2-approximations to k-unit (p,q)-lotteries. This allows us to relate the performance of single-price lotteries to our benchmark (Corollary 3.6), which will be useful in our construction of an approximately optimal prior-free mechanism in the next section.

Lemma 3.5. *For every valuation profile $\mathbf{v}$ and parameters k, p, and q, there is a p' such that the k-unit p'-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 $\mathrm{Lot}_{p,q}(\mathbf{v}) \leq \mathrm{Lot}_p(\mathbf{v}) + \mathrm{Lot}_q(\mathbf{v})$. We argue the stronger statement that each agent enjoys at least as large a combined expected utility in $\mathrm{Lot}_p(\mathbf{v})$ and $\mathrm{Lot}_q(\mathbf{v})$ as in $\mathrm{Lot}_{p,q}(\mathbf{v})$.

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 $0 < s < k < s+t$ as otherwise the k-unit (p,q) lottery is a single-price lottery. Each agent in T participates in a k-unit q-lottery in Lot_q and only a $(k-s)$-unit q-lottery in $\mathrm{Lot}_{p,q}$; its expected utility can only be smaller in the second case. Now consider $i \in S$. Writing $r = (k-s+1)/(t+1)$, we can upper bound the utility of an agent i in $\mathrm{Lot}_{p,q}$ by

$$v_i - rq - (1-r)p = (1-r)(v_i - p) + r(v_i - q) \leq (v_i - p) + \tfrac{k}{s+t}\cdot(v_i - q),$$

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

Corollary 3.6. *For every valuation profile $\mathbf{v}$, there is a k-unit p-lottery with expected residual surplus at least $\mathcal{G}(\mathbf{v})/2$.*

3.3 A Near-Optimal Prior-Free Money-Burning Mechanism

We now give a prior-free mechanism that $O(1)$-approximates the benchmark $\mathcal{G}$. This mechanism is motivated by the following observations. First, by Corollary 3.6, our mechanism only needs to compete with k-unit p-lotteries. Second, if many agents make significant contributions to the optimal residual surplus, then we can use random sampling techniques to approximate the optimal k-unit p-lottery. Third, if a few agents are single-handedly responsible for the residual surplus obtained by the optimal k-unit p-lottery, then the k-unit Vickrey auction obtains a constant fraction of the optimal residual surplus. The precise mechanism is as follows.

Definition 3.7 (Random Sampling Optimal Lottery). *With a set $S = \{1,\ldots,n\}$ of n agents and a supply of k identical units of an item, the* Random Sampling Optimal Lottery (RSOL) *is the following mechanism.*

1. *Choose a subset $S_1 \subset S$ of the agents uniformly at random, and let S_2 denote the rest of the agents. Let p_2 denote the price charged by the optimal k-unit p-lottery for S_2.*
2. *With 50% probability, run a k-unit p_2-lottery on S_1.*
3. *Otherwise, run a k-unit Vickrey auction on S_1.*

We have deliberately avoided optimizing this mechanism in order to keep its description and analysis as simple as possible.

Theorem 3.8. RSOL *$O(1)$-approximates the benchmark $\mathcal{G}$.*

In our proof of Theorem 3.8, we use the following "Balanced Sampling Lemma" of Feige et al. [13] to control the similarity between the random sample S_1 chosen by RSOL and its complement S_2.

Lemma 3.9 (Balanced Sampling Lemma [13]). *Let S be a random subset of $\{1,2,\ldots,n\}$. Let n_i denote $|S \cap \{1,2,\ldots,i\}|$. Then*

$$\mathbf{Pr}\big[n_i \leq \tfrac{3}{4}i \text{ for all } i \in \{1,2,\ldots,n\} \,\big|\, n_1 = 0\big] \geq \tfrac{9}{10}.$$

Proof. (of Theorem 3.8). Fix a valuation profile $\mathbf{v}$ with $v_1 \geq \cdots \geq v_n$ and a supply $k \geq 1$. For clarity, we make no attempt to optimize the constants in the following analysis.

We analyze the performance of RSOL only when certain sampling events occur. For $i = 1, 2$, let $\mathcal{E}_i$ denote the event that agent i is included in the set S_i. Clearly, $\mathbf{Pr}[\mathcal{E}_1 \cap \mathcal{E}_2] = 1/4$. Conditioning on $\mathcal{E}_1 \cap \mathcal{E}_2$, let $\mathcal{E}_3$ denote the event that the Balanced Sampling Lemma holds for the sample $S_1\backslash\{1\}$ when viewed as a subset of $\{2,3,\ldots,n\}$. Similarly, let $\mathcal{E}_4$ denote the event that the Balanced Sampling Lemma holds for the sample $S_2\backslash\{2\}$ when viewed as a subset of $\{1,3,\ldots,n\}$. By the Principle of Deferred Decisions and the Union Bound, $\mathbf{Pr}[\mathcal{E}_3 \cap \mathcal{E}_4 | \mathcal{E}_1 \cap \mathcal{E}_2] \geq 4/5$. Hence, $\mathbf{Pr}\left[\cap_{i=1}^4 \mathcal{E}_i\right] \geq 1/5$. We prove a bound on the approximation ratio conditioned on the event $\cap_{i=1}^4 \mathcal{E}_i$; since the mechanism always has nonnegative residual surplus, its unconditional approximation ratio is at most 5 times as large.

Let n_i and $\bar{n}_i$ denote $|S_1 \cap \{1,2,\ldots,i\}|$ and $|S_2 \cap \{1,2,\ldots,i\}|$, respectively. Since the event $\cap_{i=1}^4 \mathcal{E}_i$ holds, we have

$$n_i, \bar{n}_i \in \left[\tfrac{1}{6}i, \tfrac{5}{6}i\right] \tag{3}$$

for every $i \in \{2,3,\ldots,n\}$, and also $n_1 = 1$ and $\bar{n}_1 = 0$.

By Corollary 3.6, we only need to show that the expected residual surplus of the mechanism is at least a constant fraction of that of the optimal k-unit p-lottery for $\mathbf{v}$. For a subset T of agents and a price p, let $W(T,p)$ denote the residual surplus of the k-unit p-lottery for T. Letting n_i^T denote $|T \cap \{1,2,\ldots,i\}|$ and d_i denote $v_i - v_{i+1}$ for $i \in \{1,2,\ldots,n\}$ (interpreting $v_{n+1} = 0$), for every ℓ we obtain the following useful identity:

$$\begin{aligned} W(T, v_{\ell+1}) \quad = \quad & \frac{\min\{k, n_\ell^T\}}{n_\ell^T}\left(\sum_{i \in T\cap\{1,\ldots,\ell\}} v_i\right) - \\ & \min\{k, n_\ell^T\} \cdot v_{\ell+1} = \frac{\min\{k, n_\ell^T\}}{n_\ell^T} \sum_{i=1}^{\ell} n_i^T d_i. \end{aligned} \tag{4}$$

Let v_{ℓ^*+1} denote the optimal price for a k-unit p-lottery for $\mathbf{v}$, and note that $\ell^* \geq k$. By (5), the residual surplus of this optimal lottery is

$$W(S, v_{\ell^*+1}) = \frac{k}{\ell^*} \sum_{i=1}^{\ell} i d_i.$$

To analyze the expected residual surplus of RSOL, first suppose that it executes a k-unit p_2-lottery where $p_2 = v_{m+1}$ for some m. We then have

$$\begin{aligned} W(S_2, p_2) \quad & \geq \quad W(S_2, v_{\ell^*+1}) = \frac{\min\{k, \bar{n}_{\ell^*}\}}{\bar{n}_{\ell^*}} \sum_{i=1}^{\ell^*} \bar{n}_i d_i \\ & \geq \quad \frac{k}{\ell^*} \sum_{i=2}^{\ell^*} \frac{i}{6} d_i \geq \frac{W(S, v_{\ell^*+1})}{6} - d_1, \end{aligned}$$

where the first inequality follows from the optimality of p_2 for S_2, the first equality follows from (5), and the second inequality follows from (3). On the other hand, inequality (3) and a similar derivation shows that the price p_2 is nearly as effective for S_1:

$$\begin{aligned} W(S_1, p_2) \quad & = \quad \frac{\min\{k, n_m\}}{n_m} \sum_{i=1}^{m} n_i d_i \\ & \geq \quad \left(\frac{1}{5} \cdot \frac{\min\{k, \bar{n}_m\}}{\bar{n}_m}\right) \sum_{i=1}^{m} \frac{\bar{n}_i}{5} d_i \\ & = \quad \frac{W(S_2, p_2)}{25} \geq \frac{W(S, v_{\ell^*+1})}{150} - d_1. \end{aligned}$$

Finally, if the mechanism executes a k-unit Vickrey auction for S_1, then it obtains residual surplus at least $v_1 - v_2 = d_1$ (since the first agent is in S_1). Averaging the residual surplus from the two cases proves that RSOL $O(1)$-approximates $\mathcal{G}$. □

We can improve the approximation factor in Theorem 3.8 by more than an order of magnitude by modifying RSOL and optimizing the proof. Obtaining an approximation factor less than 10, say, appears to require a different approach.

4. LOWER BOUNDS FOR PRIOR-FREE MONEY-BURNING MECHANISMS

This section establishes a lower bound of 4/3 on the approximation ratio of every prior-free money-burning mechanism. This implements the fourth step of the prior-free mechanism design template outlined in the Introduction. Our proof follows from showing that there is a i.i.d. distribution $\mathbf{F}$ for which the expected value of our benchmark $\mathcal{G}$ is a constant factor larger than the expected residual surplus of an optimal mechanism for the distribution, such as $\mathrm{Opt}_{\mathbf{F}}$. This shows an inherent gap in the prior-free analysis framework that will manifest itself in the approximation factor of every prior-free mechanism.

Proposition 4.1. *No prior-free money-burning mechanism has approximation ratio better than 4/3 with respect to the benchmark $\mathcal{G}$, even for the special case of two agents and one unit of an item.*

We prove Proposition 4.1 by considering two agents with valuations drawn i.i.d. from a standard exponential distribution with density $f(x) = e^{-x}$ on $[0,\infty)$, and a single unit of an item. This distribution has constant hazard rate, so a lottery is an optimal mechanism, and it has expected (residual) surplus of this mechanism is 1. A calculation shows that the expected value of $\mathcal{G}(\mathbf{v})$ is 4/3, which proves the claimed lower bound.

For the special case of two agents and a single good, an appropriate mixture of a lottery and the Vickrey auction is a 3/2-approximation of the benchmark $\mathcal{G}(\mathbf{v})$. Determining the best-possible approximation ratio is an open question, even in the two agent, one unit special case.

Proposition 4.2. *For two bidders and a single unit of an item, there is a prior-free mechanism that 3/2-approximates the benchmark $\mathcal{G}$.*

Proof. Consider a valuation profile with $v_1 \geq v_2$. If we run a Vickrey auction with probability 1/3 and a lottery with probability 2/3, then the expected residual surplus is

$$\tfrac{1}{3}(v_1 - v_2) + \tfrac{2}{3}\left(\tfrac{v_1+v_2}{2}\right) \geq \tfrac{2}{3}\max\left\{\tfrac{v_1+v_2}{2}, v_1 - \tfrac{v_2}{2}\right\} = \tfrac{2}{3}\mathcal{G}(\mathbf{v}).$$

□

5. QUANTIFYING THE POWER OF TRANSFERS AND MONEY-BURNING

For the objective of surplus maximization, mechanisms with general transfers are clearly as powerful as money-burning mechanisms, which in turn are as powerful as mechanisms without money. This section quantifies the distance between the levels of this hierarchy by studying surplus approximation in multi-unit auctions. Precisely, we call a class of mechanisms *α-surplus maximizers* if, for every multi-unit auction problem, there is a mechanism in the class that obtains at least a $1/\alpha$ fraction of the full surplus for every valuation profile. For example, mechanisms with transfers are 1-surplus maximizers, because the VCG mechanism achieves full surplus in every multi-unit auction problem. Mechanisms without transfers are (n/k)-surplus maximizers, since the expected surplus of a k-unit lottery is k/n times the full surplus. One can show (details omitted) that mechanisms without transfers are not significantly better than $\Theta(n/k)$-surplus maximizers.

The interesting question is to identify the exact location of money-burning mechanisms between these two extremes: what is the potential benefit of implementing monetary transfers in a system that initially only supports money burning? We give a lower bound (proof omitted) and a matching upper bound, for all k and n.

PROPOSITION 5.1. *Money-burning mechanisms are* $\Omega(1+\log\frac{n}{k})$*-surplus maximizers in* k*-unit auctions.*

THEOREM 5.2. *Money-burning mechanisms are* $O(1+\log\frac{n}{k})$*-surplus maximizers in* k*-unit auctions.*

PROOF. Fix k and a valuation profile $\mathbf{v}$ with $v_1 \geq \cdots \geq v_n$. Assume for simplicity that both k and n are powers of 2. Our simple mechanism is as follows. First, choose a nonnegative integer j uniformly at random, subject to $k \leq 2^j \leq n$. Note that there are $1+\log_2(n/k)$ possible choices for j. Second, run a k-unit v_{2^j+1}-lottery, where we interpret v_{n+1} as zero.

Write $V^* = \sum_{i=1}^{k} v_i$ for the full surplus. For each j between $\log_2 k$ and $\log_2 n$, let R_j denote the residual surplus obtained by the mechanism for a given value of j. We claim that

$$\mathbf{E}[R_j \mid j \text{ is chosen}] \geq \begin{cases} \frac{V^*}{2} - \frac{k}{2}v_{k+1} & \text{if } j = \log_2 k \\ \frac{k}{2}\left(v_{2^{j-1}+1} - v_{2^j+1}\right) & \text{otherwise.} \end{cases}$$

When $j = \log_2 k$, the residual surplus is exactly $V^* - kv_{k+1} \geq (V^* - kv_{k+1})/2$. To justify the second case, note that k units will be randomly allocated amongst the top 2^j bidders at price v_{2^j+1}. Each of these goods is allocated to one of the top 2^{j-1} of these bidders with 50% probability, and the residual surplus contributed by such an allocation is at least $v_{2^{j-1}} - v_{2^j+1} \geq v_{2^{j-1}+1} - v_{2^j+1}$.

Let R denote the residual surplus obtained by our mechanism. The following derivation completes the proof:

$$\begin{aligned} \mathbf{E}[R] &= \tfrac{1}{1+\log_2(n/k)} \sum_{j=\log_2 k}^{\log_2 n} \mathbf{E}[R_j \mid j \text{ is chosen}] \\ &\geq \tfrac{V^*}{2(1+\log_2(n/k))}. \end{aligned}$$

□

Since the mechanism in Theorem 5.2 is prior-free, we obtain the same (tight) guarantee for every Bayesian optimal mechanism.

COROLLARY 5.3. *For every i.i.d. distribution* $\mathbf{F}$*, the expected residual surplus of the Bayesian optimal mechanism for* $\mathbf{F}$ *obtains an* $\Omega(1/(1+\log(n/k)))$ *fraction of the expected full surplus.*

Theorem 5.2 and Corollary 5.3 suggest that the cost of implementing money-burning payments instead of (possibly expensive or infeasible) general transfers is relatively modest, provided an optimal money-burning mechanism is used.

6. CONCLUSIONS

We phrased our analysis of the Bayesian setting in terms of feasible allocations (e.g., $\mathbf{x} \in \mathcal{X}$ if and only if $\sum_i x_i \leq k$ for the k-unit auction problem); however, it applies more generally to single-parameter agent problems where the service provider must pay an arbitrary cost $c(\mathbf{x})$ for the allocation $\mathbf{x}$ produced. Standard problems in this setting include fixed cost services, non-excludable public goods, and multicast auctions [14]. The solution to these problems is again to maximize the ironed virtual surplus, which in this context is the sum of the agents' ironed virtual valuations less the cost of providing the service, $\sum_i \bar{\vartheta}_i(v_i)x_i - c(\mathbf{x})$. This generalization also applies when the agents' valuations are independent but not identically distributed, i.e., agent i has ironed virtual valuation function $\bar{\vartheta}_i(\cdot)$.

THEOREM 6.1. *Given service cost* $c(\cdot)$ *and a valuation profile,* $\mathbf{v}$*, drawn from distribution* $\mathbf{F} = F_1 \times \cdots \times F_n$ *with ironed virtual valuation (for utility) function* $\bar{\vartheta}_i(\cdot)$ *for agent* i*, every mechanism with allocation rule satisfying*

1. $\mathbf{x}(\mathbf{v}) \in \operatorname{argmax}_{\mathbf{x}'} \sum_i \bar{\vartheta}_i(v_i)x_i - c(\mathbf{x}')$ *and*
2. $\frac{d}{dv_i}\bar{\vartheta}_i(v_i) = 0 \Rightarrow \frac{d}{dv_i}x_i(v_i) = 0$

is optimal with respect to expected residual surplus.

Our results for the Bayesian problem also extend beyond dominant strategy mechanisms. The well known *revenue equivalence* result [28] is popularly stated as: first price, second price (a.k.a., Vickrey), and all-pay auctions all achieve the same profit. Of course this applies to money burning as well. While this paper emphasized the dominant strategy "second price" optimal auction, there are also first-price and all-pay variants that achieve the same performance. The all-pay variant is especially interesting because of its potential usefulness for network problems. For example, in network routing, all agents could attach a proof of a computational payment to their packets. The routing protocol can then route the appropriate packets (depending on the amount of computational payment) and drop the rest. There is no need for a round of bidding, a round of transmitting the packets of winning agents, and a round of collecting payments.

One of our main results is in giving a benchmark based on the optimal mechanism for the symmetric setting of i.i.d. agents and k-unit auctions. Another main result is in approximating this benchmark with a prior-free mechanism. Can these techniques be generalized beyond symmetric settings? In particular, the notion that agents' private valuations may be paired with publicly observable *attributes* allowed for prior-free mechanisms to approximate Bayesian mechanisms for digital good auctions and non-identically distributed valuations [4]. Further, there has been some limited success in prior-free optimal mechanism design with structured costs or feasible allocations (e.g., [16] for multicast auctions and [23] for path auctions).

Our analyses and the prior-free template extend to k-unit auction problems beyond our objective of residual surplus. Imagine the k-unit auction in an i.i.d. Bayesian setting where the optimal solution is characterized by optimizing an ironed virtual value for some quantity other than utility. For example, the "virtual valuation for a 8% government sales tax", to optimize the value of the agents and mechanism less the tax deducted by government, would be

$\varphi(v) = 0.92v - 0.08\frac{1-F(v)}{f(v)}$. The optimal k-unit (p, q)-lottery is still the appropriate benchmark. Furthermore, as long as the optimal (p, q)-lottery makes use of prices p and q bounded above by the second highest bid, $v_{(2)}$, as in the money-burning context, then it is likely that our prior-free mechanism, RSOL, can be employed to approximate the benchmark. Notice that when applying this technique to "virtual valuations for payments", which are the appropriate notion of virtual valuations for the objective of profit maximization, the optimal k-unit (p, q)-lottery is simply a posted price at p. Furthermore, the optimal posted price might satisfy $p = v_{(1)} \gg v_{(2)}$. As it is not possible to approximate such a benchmark to within any constant factor, the prior-free digital goods auction literature excludes this possibility by defining the benchmark to be the profit of the optimal posted price $p \leq v_{(2)}$.

In our work there was an implicit, publicly known, exchange rate for money burnt. In network settings, where burnt payments correspond to degraded service quality or computational payments, the designer may not know each agent's relative disutility for such payments. This motivates considering the more general setting where agents have a private value for burnt money in addition to their private value for service. This moves the problem from a single-parameter setting to the far more challenging multi-parameter setting where there are few positive results.

There are additionally a few loose ends to tie up with the particular question of money-burning. For k-unit auctions, can we give tighter upper and lower bounds for prior-free money-burning mechanisms with a small number of agents? For general settings beyond i.i.d. distributions and k-unit auctions, can we quantify the power of money burning?

7. REFERENCES

[1] A. Archer, C. Papadimitriou, K. Tawar, and É. Tardos. An Approximate Truthful Mechanism for Combinatorial Auctions with Single Parameter Agents. In *SODA '03*.

[2] A. Archer and É. Tardos. Frugal path mechanisms. *ACM TALG*, 3(1):Article 3, 2007.

[3] K. Arrow. *Social Choice and Individual Values*. Wiley, 1951.

[4] M. Balcan, A. Blum, J. Hartline, and Y. Mansour. Mechanism design via machine learning. In *FOCS '05*.

[5] S. Chakravarty and T. R. Kaplan. Manna from heaven or forty years in the desert: Optimal allocation without transfer payments. Working paper, 2006.

[6] G. Christodoulou, E. Koutsoupias, and A. Nanavati. Coordination mechanisms. In *ICALP '04*, pages 345–357.

[7] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[8] R. Cole, Y. Dodis, and T. Roughgarden. How much can taxes help selfish routing? *JCSS*, 72(3):444–467, 2006.

[9] D. Condorelli. Weak cartels at standard auctions. Working paper, 2007.

[10] C. Dwork, A. Goldberg, and M. Naor. On memory-bound functions for fighting spam. In *CRYPTO' 03*, pages 426–444.

[11] C. Dwork and M. Naor. Pricing via processing or combating junk mail. In *CRYPTO '92*, pages 139–147.

[12] E. Elkind, A. Sahai, and K. Steiglitz. Frugality in path auctions. In *SODA '04*, pages 701–709.

[13] U. Feige, A. Flaxman, J. Hartline, and R. Kleinberg. On the Competitive Ratio of the Random Sampling Auction. In *WINE '05*, pages 878–886.

[14] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *JCSS*, 63(1):21–41, 2001.

[15] J. Feigenbaum, M. Schapira, and S. Shenker. Distributed algorithmic mechanism design. In Nisan et al. [29], chapter 14, pages 363–384.

[16] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Generalized competitive auctions. In *STOC '02*, pages 72–81.

[17] A. Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41:211–215, 1973.

[18] A. V. Goldberg, J. D. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55:242–269, 2006.

[19] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. In *SODA '01*, pages 735–744.

[20] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[21] M. Guo and V. Conitzer. Worst-case optimal redistribution of VCG payments. In *EC '07*, pages 30–39.

[22] J. Hartline and A. Karlin. Profit maximization in mechanism design. In Nisan et al. [29], chapter 13, pages 331–362.

[23] A. Karlin, D. Kempe, and T. Tamir. Beyond VCG: Frugality in truthful mechanisms. In *FOCS '05*, pages 615–626.

[24] H. Levin, M. Schapira, and A. Zohar. Interdomain routing and games. In *STOC '08*. To appear.

[25] D. Liu and J. Camp. Proof of work can work. In *Fifth Workshop on the Economics of Info. Sec.*, 2006.

[26] P. McAfee and J. McMillan. Bidding rings. *American Economic Review*, 82:579–599, 1992.

[27] H. Moulin. Efficient, strategy-proof and almost budget-balanced assignment. Working Paper, 2007.

[28] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

[29] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[30] J. Riley and W. Samuelson. Optimal auctions. *American Economic Review*, 71:381–92, 1981.

[31] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.

[32] M. Satterthwaite. Strategy-proofness and Arrow's condition: Existence and correspondence theorems for voting proceedures and social welfare functions. *Journal of Economic Theory*, pages 187–217, 1975.

[33] J. Schummer and R. V. Vohra. Mechanism design without money. In Nisan et al. [29], chapter 10, pages 243–265.

[34] K. Talwar. The price of truth: Frugality in truthful mechanisms. In *STACS '03*, pages 608–619.

[35] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance*, 16:8–37, 1961.