Money Burning and Implementation

Jason D. Hartline Tim Roughgarden

January 31, 2007

Abstract

In response to public disapproval for any electronic mail system that includes a charge for sending email, *computational spam fighting*, which was originally proposed in [4], has reemerged as an economic approach that could be used to prevent spam in email systems [3]. This approach requires email senders to solve a challenge that expends their computational resources but otherwise has no direct social benefit. This challenge is tantamount to requiring that the sender burn money. This paper explores the extent to which *money burning* may be useful in broader implementation contexts.

We consider the general problem of designing socially optimal, single round, sealed bid mechanisms when transfers made from the agents to the mechanism must be burnt (i.e., are a social loss). In these settings the socially optimal outcome is the one that maximizes the *marginal surplus*, that is, the sum of the agents' valuations minus the agents' payments (minus any social cost of the outcome). In the Bayesian setting where the agents' valuations are drawn independently, but not necessarily identically, from a known distribution, we give a concise characterization of the mechanism that maximizes the expected marginal surplus. From this characterization we observe that the socially optimal way to allocate a single item to agents with i.i.d. valuations is, depending on the distribution, to either assign the item to an arbitrary agent or to run a second-price-like auction. Furthermore, for non-identical distributions that satisfy the monotone hazard rate condition, the socially optimal mechanism (for any given social cost function) is the one that chooses the allocation to maximize the expected social surplus ex ante and requires no payments.

1 Introduction

We study implementation problems in which monetary transfers are not allowed. Consider the problem of implementing an economic solution to spam in electronic mail systems. A natural solution is to charge the sender of each email a fraction of a cent for each email sent. Such a scheme would eradicate spam as the cost of delivering a spam message to the number of recipients necessary to make a sale would be higher than the revenue from the sale. Unfortunately, such schemes are met with massive resistance by a public that believes that free email is their unalienable right. To circumvent the public disaster that would result from charging for email, the approach of *computational spam fighting* can be used [4] (see also [3, 8]). This approach uses cryptographic techniques to force the sender of email to expend their computational resources performing a worthless computation. From an economic point of view, this computational payment is equivalent to a monetary payment that is burnt (i.e., that is a social loss). This solution to fighting email spam motivates a number of questions on the role of money burning in implementation.

The role of transfers in implementation is fundamental. With the ability to transfer money between the agents and the mechanism, the Vickrey-Clarke-Groves mechanism (VCG) [14, 2, 7] implements the objective of social surplus maximization. With no ability for transfers of any kind, the impossibilities of Arrow [1], Gibbard [5], and Satterthwaite [13] show that only dictatorships can be implemented. The problem of social surplus maximization in the case where agents make payments that are burnt instead of transferred to the mechanism is an interesting middle ground between these settings. In particular, while payments can be levied to align the preferences of the agents with the social objective, such payments also degrade the this objective. It is clear that the social surplus implementable with money burning is at most that possible with transfers and at least that possible with no transfers. We consider the following questions in a general setting:

1. When can the social surplus of an implementation with money burning strictly exceed that of an implementation with no transfers?

2. How much better is implementation with transfers than implementation with money burning?

3. What is the optimal implementation with money burning (for maximizing the social surplus)?

We consider these questions in an abstract setting. A social planner must select an outcome $\mathbf{x}$ from some class of outcomes $\mathcal{X}$. Each outcome has a social cost $c(\mathbf{x})$. There is a set of n agents and each agent i has a valuation function over outcomes, v_i, that maps outcomes to non-negative real numbers. We assume that these valuation functions are drawn from a known joint distribution $\mathbf{F}$.

A direct revelation mechanism solicits bids in the form of valuation functions and selects an outcome $\mathbf{x}$ and an n-tuple of monetary amounts $\mathbf{p} = (p_1, \ldots, p_n)$ for each agent to burn. We refer to the p_i as *payments*. We assume that the agents have quasi-linear *utility* (so $u_i = v_i(\mathbf{x}) - p_i$ for agent i), and are risk-neutral utility maximizers. The *marginal surplus* of a mechanism when the true agent valuations are $\mathbf{v} = (v_1, \ldots, v_n)$, the outcome is $\mathbf{x}$, and the payments are $\mathbf{p}$ is

$$\sum_i (v_i(\mathbf{x}) - p_i) - c(\mathbf{x}).$$

If the payments were transferred to the seller then the social surplus would be $\sum_i v_i(\mathbf{x}) - c(\mathbf{x})$; however, in our setting the payments are burnt and the social surplus is equal to the marginal surplus.

As the payments of the agents are in the objective function our analysis relies on a concise characterization of the relationship between the payments and the social choice function. Such a characterization is possible in the restricted setting of *binary single-parameter* agents, where each agent's valuation function belongs to the following class. Assume that for agent i, there is a publicly known partition of $\mathcal{X}$ into *good outcomes* and *bad outcomes*. For all bad outcomes the agent's valuation is identically zero. For all good outcomes the agent's valuation is a constant non-negative real value. Thus, if we let x_i be an indicator for whether $\mathbf{x}$ is good or bad, then we can write the valuation of agent i as $v_i x_i$, where we abuse notation and let v_i represent agent i's valuation for any good outcome. For a set of n such agents our distribution $\mathbf{F}$ is simply a distribution on n-tuples of non-negative real numbers. For binary single-parameter agents we can assume that $\mathcal{X}$ is $\{0, 1\}^n$. We represent an infeasible allocation $\mathbf{x}$ by assigning it infinite social cost: $c(\mathbf{x}) = \infty$.

This setting, while specialized, is rich enough to express a number of fundamental mechanism design problems.

Single-item auction: We wish to allocate a single indivisible item among a set of agents. The social cost satisfies:

$$c(\mathbf{x}) = \begin{cases} 0 & \text{if } \sum_i x_i \leq 1; \\ \infty & \text{otherwise.} \end{cases}$$

Non-excludable public good: We are building a bridge that costs C. If the bridge is built then all agents may use it.

$$c(\mathbf{x}) = \begin{cases} C & \text{if } x_i = 1 \text{ for all } i; \\ 0 & \text{if } x_i = 0 \text{ for all } i; \\ \infty & \text{otherwise.} \end{cases}$$

Excludable public good: It costs C to publish an issue of an electronic journal, but only a limited set (the "subscribers") can access it.

$$c(\mathbf{x}) = \begin{cases} C & \text{if } \sum_i x_i \geq 1; \\ 0 & \text{otherwise.} \end{cases}$$

We consider implementation by a *Bayesian incentive compatible* mechanism. For binary single-parameter agents and a product distribution $\mathbf{F} = (F_1 \times \cdots \times F_n)$, Myerson [10] considers the related problem of optimizing the profit, $\sum_i p_i c(\mathbf{x})$, of the mechanism.[1] His approach is to characterize the expected payments of agents in an incentive-compatible mechanism in terms of the mechanism's allocation function. In particular, he shows that an agent's expected payment is equal to the expectation of their *virtual valuation*, $\varphi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ (with f_i denoting the density function of F_i). Using this characterization, there is a concise description of the optimal auction (for any social cost function $c(\cdot)$) as the one that maximizes the *virtual surplus*,[2]—the sum of the virtual valuations of the agents allocated minus the social cost. Myerson concludes that in the special case of single-item auctions with i.i.d. valuations that satisfy the *monotone hazard rate*, the optimal auction is a second-price auction with appropriate reserve price.

We analyze the objective of maximizing the marginal surplus, i.e., $\sum_i (v_i x_i - p_i) - c(\mathbf{x})$. Myerson's characterization of the expected payments holds whether or not the payments are burnt. We define the *virtual marginal valuation* of an agent to be the difference between their valuation and their virtual valuation, $\vartheta_i(v_i) = \frac{1-F_i(v_i)}{f_i(v_i)}$, and show that the mechanism that maximizes the expected marginal surplus is precisely the one that maximizes the virtual marginal surplus subject to monotonicity. These virtual marginal valuation functions are often not monotone non-decreasing—indeed, if the monotone hazard rate holds, they are monotone decreasing. However, we can apply the ironing procedure of Myerson to get an *ironed virtual marginal valuation function*, $\bar{\vartheta}_i(v_i)$. Our main theorem states that the mechanism that maximizes the social surplus when payments are burnt is precisely the one that maximizes the *ironed virtual marginal surplus*, i.e., $\sum_i \bar{\vartheta}_i(v_i) x_i - c(\mathbf{x})$.

We interpret this theorem for several interesting cases. If the monotone hazard rate condition holds then the ironed marginal valuation functions are constant and equal to the expected value of

[1] Myerson only analyzes the case of single-item auctions; however, all of his theorems extend mutatis mutandis to our general single-parameter settings.

[2] For virtual valuation functions that are not monotone, Myerson gives an "ironing" procedure for making them monotone.

the distribution, i.e., $\bar{\vartheta}_i(v_i) = \mu_i$. We conclude that under the monotone hazard rate assumption, for general social costs, $c(\cdot)$, the optimal mechanism is the one that maximizes the expected surplus ex ante ($\sum_i \mu_i x_i - c(\mathbf{x})$) and has no agent payments. Thus, under the monotone hazard rate condition, money burning offers no advantage for maximizing social surplus over implementation without transfers.

For single-item auctions with i.i.d. valuations, if the virtual marginal valuation functions are strictly increasing then the optimal mechanism is the second-price auction. Otherwise, the optimal mechanism can be viewed as an indirect mechanism that allocates the item to the highest bidder, but requires the bids to be in a particular subset of the valuation distribution's support. This looks like a second-price auction except that the gaps in the set of allowable bids result in potential for ties. These ties can be broken arbitrarily and the payments that induce truthful reporting adjusted appropriately. Of particular interest is the special case where the virtual marginal valuations are decreasing (i.e., monotone hazard rate) where, as described in the preceding paragraph, the optimal mechanism is a dictatorship.[3]

Related Work. This work is based heavily upon the optimal mechanism design literature [10, 11]. Our goal of quantifying the ability to optimize social surplus in implementation with money burning relative to that of implementation with transfers is related to work on the price of anarchy (See, e.g., [12]). There is also recent work that studies the case where transfers to the mechanism are burnt, but transfers between agents are not. For this problem, Moulin [9] gives a mechanism for single-item multi-unit unit-demand settings and shows that even in worst case settings the loss in social surplus due to money burning decreases exponentially with the number of agents. In games of complete information, for allocating a single item among a set of n agents (i.e., the single-item auction problem), there is a multi-stage game with unique subgame perfect equilibrium that implements the social surplus maximizing allocation (i.e., the agent with the highest value wins) [6].

2 Review of Optimal Mechanism Design

We consider the problem of implementation in Bayes-Nash equilibrium. By the revelation principle, it suffices to consider direct *Bayesian incentive compatible* mechanisms in which truthtelling is a Bayes-Nash equilibrium; though as we will see, the optimal mechanisms in our setting turn out to have truthtelling as a dominant strategy. Let $\mathbf{x}(\mathbf{v})$ and $\mathbf{x}(\mathbf{p})$ represent the allocation and payments selected by a mechanism in equilibrium when the actual valuations of the agents are $\mathbf{v}$. Given allocation rule $\mathbf{x}(\mathbf{v})$, let $x_i(v_i)$ be the probability that agent i is allocated when their valuation is v_i (given the randomization of the other agents' valuations). I.e., $x_i(v_i) = \mathbf{E}_{\mathbf{v}_{-i}}[x_i(v_i, \mathbf{v}_{-i})]$. Similarly define $p_i(v_i)$. Positive transfers from the mechanism to the agents are not allowed and we require ex interim individual rationality. It is well known (e.g. [10]) that under these conditions Bayesian incentive-compatibility implies that

1. Allocation monotonicity: for all i, and $v_i > {v_i}'$,

$$x_i(v_i) \geq x_i({v_i}').$$

[3] A dictatorship is a mechanism that chooses the most desirable outcome of one agent ignoring all other agents. In a single item auction, a dictatorship simply assigns the item to an arbitrary agent.

2. Payments: for all i and v_i,

$$p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(v) dv.$$

For example, the above conditions are satisfied by a *threshold* allocation and payment rule: for t_i that is a function of the bids, $\mathbf{b}_{-i}$, of all other agents, agent i wins when bidding above t_i and pays t_i, and loses when bidding below t_i and pays nothing. As payments are determined by the allocation rule, we often omit explicit discussion of payment rules.

Our analysis is based on the following definitions and theorems of Myerson [10]. We assume that each component F_i of the distribution has support (a_i, b_i), and has positive density throughout this interval.

Definition 2.1 (Virtual valuation) *If agent i's valuation is distributed according to F_i, then its* virtual valuation *is*

$$\varphi_i(v_i) = v_i - \tfrac{1-F_i(v_i)}{f_i(v_i)}.$$

Definition 2.2 (Virtual surplus) *If the agents' valuations are $\mathbf{v}$, then the* virtual surplus *of an allocation $\mathbf{x}$ is*

$$\textstyle\sum_i \varphi_i(v_i) x_i - c(\mathbf{x}).$$

Lemma 2.1 *In a Bayesian incentive-compatible mechanism, the expected payment of agent i is*

$$\mathbf{E}_{\mathbf{v}}[\varphi_i(v_i) x_i(\mathbf{v})] \,.$$

Theorem 2.2 *The expected profit of a Bayesian incentive compatible mechanism is*

$$\mathbf{E}_{\mathbf{v}}\left[\textstyle\sum_i \varphi_i(v_i) x_i(\mathbf{v}) - c(\mathbf{x}(\mathbf{v}))\right] \,.$$

In the case that the virtual valuation functions are monotone, Theorem 2.2 suggests a natural optimal mechanism. On input $\mathbf{b}$ select the allocation $\mathbf{x}$ that maximizes the virtual surplus, $\sum_i \varphi_i(b_i) x_i - c(\mathbf{x})$. When the virtual valuations are not monotone there is an *ironing* procedure that can be applied to obtain monotone $\bar{\varphi}_i(z)$ from $\varphi_i(z)$ such that maximizing $\sum_i \varphi_i(v_i) x_i - c(\mathbf{x})$ subject to monotonicity is equivalent (in expectation over $\mathbf{F}$) to maximizing $\sum_i \bar{\varphi}_i(v_i) x_i - c(\mathbf{x})$, where the any optimal allocation procedure is monotone. See Section 4 or [10] for details.

Theorem 2.3 *The dominant strategy mechanism that on input bids $\mathbf{b}$ outputs $\mathbf{x}$ to maximize*

$$\textstyle\sum_i \bar{\varphi}_i(b_i) x_i - c(\mathbf{x})$$

obtains the optimal seller profit in expectation.

3 Mechanism Design for Optimal Marginal Surplus

Our goal is to design a mechanism to maximize the marginal surplus, i.e., the social surplus less the payments made by the agents.

Definition 3.1 (marginal surplus) *The* marginal surplus *for valuations* $\mathbf{v}$*, allocation* $\mathbf{x}$*, and payments* $\mathbf{p}$ *is*

$$\sum_i (v_i x_i - p_i) - c(\mathbf{x}).$$

Given the equilibrium allocation $\mathbf{x}(\mathbf{v})$ the expected payment of an agent i is $\mathbf{E_v}[\varphi_i(v_i)x_i(\mathbf{v})]$. This motivates the following definition (derived by substituting the definition of $\varphi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ into the equation $v_i - \varphi_i(v_i)$).

Definition 3.2 (virtual marginal valuation) *The* virtual marginal valuation *of an agent with valuation* v_i *is*

$$\vartheta_i(v_i) = \frac{1 - F_i(v_i)}{f_i(v_i)}.$$

Theorem 3.1 *In a Bayesian incentive-compatible mechanism with allocation* $\mathbf{x}(\mathbf{v})$ *when agent valuations are* $\mathbf{v}$*, the expected marginal surplus is*

$$\mathbf{E_v}\left[\sum_i \vartheta_i(v_i) x_i(\mathbf{v}) - c(\mathbf{x}(\mathbf{v}))\right].$$

Just as Myerson's characterization gives a simple rule for maximizing the profit of a mechanism (following Theorem 2.3), this theorem suggests an approach for maximizing marginal surplus: on input bids $\mathbf{b}$, choose the outcome $\mathbf{x}$ to maximize $\frac{1-F_i(b_i)}{f_i(b_i)} x_i - c(\mathbf{x})$. For the special case of a single-item auction, this rule assigns the item to the agent with the highest virtual marginal valuation, $\frac{1-F_i(b_i)}{f_i(b_i)}$. Unfortunately, this approach does not generally yield truthful mechanisms.

Definition 3.3 (hazard rate) *The* hazard rate *of a distribution* F *at* z *is* $\frac{f(z)}{1-F(z)}$*. The* monotone hazard rate condition *requires that the hazard rate be monotone non-decreasing.*

When the monotone hazard rate condition holds, the virtual marginal valuation functions are monotone *non-increasing*. The single-item auction that assigns the item to the agent with the highest virtual marginal valuation is a non-monotone allocation rule and therefore cannot arise as a Bayes-Nash equilibrium of a mechanism.

Myerson addressed the same problem in his study of optimal auctions by defining an "ironing" procedure for removing the non-monotonicities of a virtual valuation function. He shows that maximizing virtual surplus by a monotone allocation is equivalent to maximizing the ironed virtual surplus. The same ironing procedure can be applied to virtual marginal valuations to arrive at *ironed marginal virtual valuation functions*, which we denote $\bar{\vartheta}_1(z),\ldots,\bar{\vartheta}_n(z)$. (See Section 4, below.)

Theorem 3.2 *The dominant strategy mechanism that on bids* $\mathbf{b}$ *chooses the outcome* $\mathbf{x}$ *to maximize*

$$\sum_i \bar{\vartheta}_i(b_i) x_i - c(\mathbf{x})$$

gives the optimal marginal surplus in expectation.

In Section 5 we interpret this theorem in several interesting contexts.

4 Ironing

Given a function $\vartheta(z)$ and distribution function $F(z)$ on support (a, b), the following ironing procedure results in a monotone non-decreasing function $\bar{\vartheta}(z)$. This procedure is identical to that given in [10] for ironing virtual valuation functions.

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 $G(q)$ to be the convex hull of $H(q)$.
4. Define $g(q)$ as the derivative of $G(q)$, where defined, and extend to all of $[0, 1]$ by right-continuity.
5. Finally, $\bar{\vartheta}(z) = g(F(z))$.

We next extend Myerson's proof of Theorem 2.3 to prove Theorem 3.2.

Let $\mathbf{x}(\mathbf{v})$ represent the allocation function chosen for valuations $\mathbf{v}$. Given allocation rule $\mathbf{x}(\mathbf{v})$, let $x_i(v_i)$ be the probability that agent i is allocated when their valuation is v_i (given the randomization of the other agent valuations). I.e., $x_i(v_i) = \mathbf{E}_{\mathbf{v}_{-i}}[x_i(v_i, \mathbf{v}_{-i})]$. We use ${x_i}'(v_i)$ to denote the derivative of $x_i(v_i)$ with respect to v_i. Recall that Bayesian incentive compatibility requires that $x_i(v_i)$ is monotone non-decreasing in v_i.

We seek a monotone allocation rule, $\mathbf{x}(\mathbf{v})$, that maximizes $\mathbf{E}_{\mathbf{v}}[\sum_i \vartheta_i(v_i)x_i(\mathbf{v}) - c(\mathbf{x}(\mathbf{v}))]$.

Lemma 4.1 *For every allocation rule $x_i(\mathbf{v})$,*

$$\mathbf{E}_{\mathbf{v}}[\vartheta_i(v_i)x_i(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}\left[\bar{\vartheta}_i(v_i)x_i(\mathbf{v})\right] - \int_a^b \left[H_i(F_i(v_i)) - G_i(F_i(v_i))\right] {x_i}'(v_i)dv_i.$$

Proof: First we use the fact that $\vartheta_i(v_i) = \bar{\vartheta}_i(v_i) + h_i(F_i(v_i)) - g_i(F_i(v_i))$ to conclude,

$$\mathbf{E}_{\mathbf{v}}[\vartheta_i(v_i)x_i(\mathbf{v})] = \mathbf{E}_{\mathbf{v}}\left[\bar{\vartheta}_i(v_i)x_i(\mathbf{v})\right] + \mathbf{E}_{\mathbf{v}}\left[\left(h_i(F_i(v_i)) - g_i(F_i(v_i))\right) x_i(\mathbf{v})\right]. \tag{1}$$

Since $\mathbf{F}$ is a product distribution, the second term satisfies

$$\begin{aligned}\mathbf{E}_{\mathbf{v}}[(h_i(F_i(v_i)) - g_i(F_i(v_i)))x_i(\mathbf{v})] &= \int_{\mathbf{v}} \left(h_i(F_i(v_i)) - g_i(F_i(v_i))\right) x_i(\mathbf{v})f_i(\mathbf{v})d\mathbf{v} \\ &= \int_a^b \left(h_i(F_i(v_i)) - g_i(F_i(v_i))\right) x_i(v_i)f_i(v_i)dv_i. \end{aligned} \tag{2}$$

Now, integrate by parts to obtain

$$\begin{aligned}\mathbf{E}_{\mathbf{v}}[(h_i(F_i(v_i)) - g_i(F_i(v_i)))x_i(\mathbf{v})] &= \left[H_i(F_i(v_i)) - G_i(F_i(v_i))\right] x_i(v_i)\Big|_a^b - \int_a^b \left[H_i(F_i(v_i)) - G_i(F_i(v_i))\right] {x_i}'(v_i)dv_i \\ &= \int_a^b \left[G_i(F_i(v_i)) - H_i(F_i(v_i))\right] {x_i}'(v_i)dv_i. \end{aligned} \tag{3}$$

Equation (3) follows from the fact that, as the convex hull of $H_i(\cdot)$ on interval $(0,1)$, $G_i(\cdot)$ satisfies $G_i(0) = H_i(0)$ and $G_i(1) = H_i(1)$. Combining this with equation (1) gives the lemma. □

We now give the proof of the main theorem of the paper.

Proof: (of Theorem 3.2) The optimal mechanism with money burning is the one that maximizes the expected virtual marginal surplus, $\mathbf{E}_{\mathbf{v}}[\sum_i \vartheta_i(v_i)x_i(\mathbf{v}) - c(\mathbf{x})]$, subject to having a monotone allocation rule. By Lemma 4.1 and linearity of expectation, this expectation is equal to

$$\mathbf{E}_{\mathbf{v}}\left[\sum\nolimits_i \bar{\vartheta}_i(v_i)x_i(\mathbf{v}) - c(\mathbf{v})\right] - \sum\nolimits_i \int_a^b \left[G_i(F_i(v_i)) - H_i(F_i(v_i))\right] x_i{}'(v_i)dv_i.$$

Clearly, we would like to maximize the first term in this difference at the same time as we minimize the second term. We show that the allocation rule $\mathbf{x}(\cdot)$ that maximizes the first term is monotone and minimizes the second term over all monotone allocation rules. Thus, the theorem is proved as the optimal mechanism is the one that maximizes the ironed marginal virtual surplus.

Assume we have $\mathbf{x}(\cdot)$ that maximizes the first term, i.e., the ironed virtual surplus (and breaks ties consistently). Since $\bar{\vartheta}_i(\cdot)$ is monotone non-decreasing, for all i, $x_i(\cdot)$ is monotone non-decreasing. For the next part, notice that the term in the integral is always non-negative for every monotone allocation rule:

- $[G_i(F_i(v_i)) - H_i(F_i(v_i))] \geq 0$ as $H_i(\cdot)$ is the convex hull of $G_i(\cdot)$, and
- $x_i{}'(v_i) \geq 0$ as $x_i(v_i)$ is monotone non-decreasing.

Since $G_i(\cdot)$ is the convex hull of $H_i(\cdot)$, for every v_i for which $G_i(F_i(v_i)) - H_i(F_i(v_i)) > 0$, it must be that $G_i(F_i(v_i))$ is locally a line and thus $g(F_i(v_i)) = \bar{\vartheta}_i(v_i)$ is locally constant around v_i. Since we assumed that $\mathbf{x}(\cdot)$ maximizes the ironed virtual surplus, the $x_i(v_i)$ is constant around v_i. Thus, $x_i{}'(v_i) = 0$. We conclude that for this $\mathbf{x}(\cdot)$ the term in the integral is identically zero and at its minimum over all monotone allocation rules. □

5 Interpretations

The monotone hazard rate assumption is standard in mechanism design. As such, we start by considering its implications for implementation with money burning. The following lemma is central to our analysis.

Lemma 5.1 *For every distribution F that satisfies the monotone hazard rate condition, the ironed virtual marginal valuation function is constant with $\bar{\vartheta}(z) = \mu$, where μ denotes the expected value of the distribution.*

Proof: Apply the ironing procedure of Section 4 to $\vartheta(z)$. The monotone hazard rate condition implies that $\vartheta(z)$ is monotone non-increasing. Since $F(z)$ is monotone non-decreasing so is $F^{-1}(q)$ for $q \in [0,1]$. Thus, $h(q) = \vartheta(F^{-1}(q))$ is monotone non-increasing. The integral $H(q)$ of the monotone non-increasing function $h(q)$ is concave. The convex hull $G(q)$ of the concave function $H(q)$ is a straight line. In particular, $H(q)$ is defined on the range $[0,1]$, so $G(q)$ is the straight line between $(0, H(0))$ and $(1, H(1))$. Thus, $g(q)$ is the derivative of a straight line and is therefore constant with value equal to the line's slope, namely $H(1)$. Thus, $\bar{\vartheta}(z) = H(1)$. It remains to show that $H(1) = \mu$. By definition,

$$H(1) = \int_0^1 \vartheta(F^{-1}(q))dq$$

Substituting $q = F(z)$, $dq = f(z)dz$, and the support of F as (a, b), we have

$$H(1) = \int_a^b \vartheta(z)f(z)dz.$$

Using the definition of $\vartheta(\cdot)$ and the definition of expectation for non-negative random variables gives

$$H(1) = \int_a^b (1 - F(z))dz = \mu.$$

□

We believe the following alternative proof of the latter half of the above lemma is interesting as well. In particular, it uses the fact that the mechanism that maximizes ironed virtual marginal surplus is optimal.

Lemma 5.2 *For every distribution F with expected valued μ, if its virtual marginal valuation function is constant, then its value equals μ.*

Proof: Consider a single agent with valuation distributed according to F and a social cost function equal to C if the agent is served and 0 otherwise. Interpreting Theorem 3.2 in this setting, the marginal surplus-maximizing mechanism allocates whenever the agent's ironed virtual marginal surplus $\bar{\vartheta}(z)$ is at least C. By assumption, $\bar{\vartheta}(z)$ is everywhere equal to some constant Z. The marginal surplus-maximizing mechanism thus either always allocates (if $Z \geq C$) or never allocates (if $Z < C$), and never charges payments. The expected marginal surplus of a mechanism of this form is either $\mu - C$ (if $Z \geq C$) or 0 (if $Z < C$). Optimality now dictates that $Z = \mu$. □

Our main result for distributions that satisfy the monotone hazard rate follows from Lemmas 5.1 and 5.2 and Theorem 3.2.

Corollary 5.3 *For a product distribution satisfying the monotone hazard rate, $\mathbf{F} = F_1 \times \cdots \times F_n$, and arbitrary costs $c(\cdot)$, the optimal mechanism chooses $\mathbf{x}$ ex ante to maximize*

$$\sum_i \mu_i x_i - c(\mathbf{x}),$$

where μ_i is the expected value of F_i. The payments are identically zero.

Corollary 5.3 implies that under the monotone hazard rate assumption, implementation with money burning offers no advantage over implementation with no transfers.

We now turn our attention to interpreting Theorem 3.2 and Corollary 5.3 for single-item auctions. The first result is a simple corollary of Corollary 5.3.

Corollary 5.4 *When the agents' valuations satisfy the monotone hazard rate condition and are i.i.d., the single-item auction with optimal marginal surplus is a dictatorship (i.e., an arbitrary agent is picked to receive the good).*

When the agents' valuations do not satisfy the monotone hazard rate and the ironed virtual marginal valuation functions are not constant, then analogous to Myerson's optimal auction, we simply award the item to the agent with the largest ironed virtual marginal valuation. If the virtual marginal valuations are strictly increasing then the optimal auction is the second-price auction. Otherwise, ties in ironed virtual marginal valuation space can be broken arbitrarily (as in a dictatorship mechanism) or randomly (as in a lottery mechanism). When there can be ties, the optimal auction can be expressed succinctly as a indirect mechanism that is dominant strategy and generalizes the second-price auction (apply the revelation principle to obtain a direct mechanism).

Corollary 5.5 *For general i.i.d. distributions, the single-item auction with maximum marginal surplus is an indirect version of the second-price auction: For valuations on the range* $R = [a, b]$*,* $\bar{\vartheta}^{-1}(q)$ *defined as* $\inf\{z \ : \ \bar{\vartheta}(z) = q\}$*, and* $R' \subseteq R$ *defined with* $v \in R'$ *if* $v = \bar{\vartheta}^{-1}(\bar{\vartheta}(v))$*; it is the indirect mechanism where agents bid* $b_i \in R'$ *and the agent with the highest bid wins, ties broken arbitrarily.*[4]

It is fairly simple to construct examples where the ironed marginal valuation functions are not constant. The most natural examples come from tail heavy distributions.

Example, a tail heavy distribution. Consider the i.i.d. distribution with $F(z) = 1 - z^{-2}$ and $f(z) = 2z^{-3}$ with support $[1, \infty)$. The virtual marginal valuation function is $\vartheta(v) = z/2$ and monotone. The single-item auction maximizing virtual marginal surplus awards the item to the agent with the highest valuation, i.e., it is the second-price auction.

Example, a finite support distribution. Consider two bidders, i.i.d. valuations in $[0, 10]$ with density function

$$f(z) = \begin{cases} 1/4 & z \in [0,2) \\ 1/16 & z \in [2,10]. \end{cases}$$

It is possible to calculate the ironed marginal virtual valuations as

$$\bar{\vartheta}(v) = \begin{cases} 3 & v \in [0,2) \\ 4 & v \in [2,10]. \end{cases}$$

We thus view the bidders as either having a high type ($v_i \in [2, 10]$) or a low type ($v_i \in [0, 2)$). Notice that both bidders are low with probability 1/4, but otherwise, at least one is high. To maximize the ironed virtual marginal surplus we simply allocate to a low type in the first case and a high type in the latter case. We break ties in favor of agent 1. The expected ironed virtual marginal surplus of this allocation procedure is $1/4 \times 3 + 3/4 \times 4 = 3 + 3/4 = 15/4$. The mechanism is this:

- if $b_2 < 2$, (happens with prob. 1/2)
  allocate to bidder 1, charge nothing. (marginal surplus = 7/2)
- if $b_2 \geq 2$ and $b_1 \geq 2$ (happens with prob. 1/4)
  allocate to bidder 1, charge \$2. (marginal surplus = 4)

[4]Implicit in this mechanism is a threshold payment for the winning bidder that is not, in general, equal to the bid value of the second highest bidder. The threshold payment rule has depends subtly on the tie-breaking procedure.

- if $b_2 \geq 2$ and $b_1 < 2$ (happens with prob. 1/4)
  allocate to bidder 2, charge \$2. (marginal surplus = 4)

The total expected marginal surplus is $1/2 \times 7/2 + 1/4 \times 4 + 1/4 \times 4 = 15/4$ and, as expected, it is equal to the expected ironed virtual marginal surplus.

Notice that the optimal mechanism with no transfers, i.e., the dictator mechanism, would always just pick agent 1. The marginal surplus is the expected valuation of agent 1 which is 7/2. Notice that $15/4 > 7/2$ as one would expect.

6 Conclusions

Now that we have an understanding of the mechanism that maximizes the marginal surplus, we can quantify the cost to society for imposing computational payments (money burning) instead of monetary payments (regular transfers). For instance, this could be measured as the ratio between the optimal social surplus and the optimal marginal surplus (as is popular in work studying the *price of anarchy*, see e.g., [12]). We are in the process of formalizing results along these lines.

References

[1] K. Arrow. *Social Chioce and Individual Values.* Wiley, 1951.

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

[3] C. Dwork, A. Goldberg, and M. Naor. On memory-bound functions for fighting spam. In *23rd Annual International Cryptology Conference*, pages 426–444, 2003.

[4] C. Dwork and M. Naor. Pricing via processing or combatting junk mail. In *12th Annual International Cryptology Conference*, pages 139–147, 1992.

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

[6] J. Glazer and A. Ma. Efficient allocation of a “prize”–king solomon’s dilemma. *Games and Economic Behavior*, 1(3):222–233, September 1989.

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

[8] D. Liu and J. Camp. Proof of work can work. In *Fifth Workshop on the Economics of Invormation Security*, 2006.

[9] H. Moulin. Strategy-proof and almost efficient assignment. Working Paper, 2006.

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

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

[12] T. Roughgarden and Eva Tardos. How bad is selfish routing? In *41st Annual Symposium on Foundations of Computer Science*, pages 93–102, 2000.

[13] 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.

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