Social Welfare of Selfish Agents: Motivating Efficiency for Divisible Resources

Rajiv T. Maheswaran and Tamer Başar *
Coordinated Science Laboratory, University of Illinois,
1308 W. Main Street,
Urbana, Illinois 61801-2307, USA
(maheswar, tbasar)@control.csl.uiuc.edu

Abstract

In today's landscape of distributed and autonomous computing, there is a challenge to construct mechanisms which can induce selfish agents to act in a way that satisfies a global goal. In the domain for the allocation of computational and network resources, proportionally fair schemes are commonly advocated. In this paper, we investigate the efficiency of the resulting equilibria in such systems. We then develop a method of generating an entire class of divisible auctions with minimal signaling and computation costs which maximize social welfare even though agents act solely to optimize their own utility.

1. Introduction

The development of the Internet and software agent technology has created a world of distributed systems composed of autonomous entities. However, these agents still interact and compete for scarce resources. Designers and managers of these systems often aim to satisfy certain global performance metrics which leads to the task of devising mechanisms that yield desired behavior from agents trying to optimize individual utilities.

We consider the domain of network and computational resources. Divisible auctions serve as a useful paradigm for transparent allocation of these services. In the realm of bandwidth allocation, economic modeling of packet transmission has been advocated [3] and proportional fairness [2] has been developed as notion for the basis of socially optimal networks. Proportional allocation schemes also have been proposed for decentralized processor scheduling [1, 7].

Unfortunately, these systems do not lead to efficient use of resources. In this paper, using social welfare as a metric, we analyze a subset of the space of divisible auctions generated as an expansion around proportionally fair systems with linear weights to see if the proposed systems are at local optima. This analysis leads to insight that allows us to find a class of mechanisms that can motivate selfish agents to produce an equilibrium state which is maximally efficient. Furthermore, these systems have the minimum signaling load (single-dimensional messages) and minimum computation requirement for allocation at the resource ($O(N)$).

The paper is organized as follows. In Section 2, we present the models of the spaces we are working in and obtain equilibrium properties to choose an appropriate subset of allocation rules to consider. In Section 3, we investigate the efficiency of mechanisms within this set. In Section 4, we engineer a method to determine cost rules which maximize social welfare for selfish agents. In Section 5, we present some concluding thoughts and areas for future work.

2. Generalized Models and Equilibrium

Auction mechanisms are characterized by an allocation rule $x(s)$ and a cost rule $c(s)$, where $s = [s_1 \cdots s_N]$ represents the signals from a population of N agents and $x_i(s)$ and $c_i(s)$ are respectively the allocation and cost to the i^{th} agent. We work in the signaling space where $s_i \in I\!R$ and $s_i \geq 0$. One subset of the space of auction mechanisms is the collection of those that can be characterized by the following allocation rule: $x_i(s) = w_i(s)/\left(\sum_{j=1}^{N} w_j(s) + \epsilon\right)$ where ϵ is a parameter controlled by the resource. Signals are translated to weights, denoted by the functions $\{w_i(\cdot)\}_{i=1}^{N}$, which determine the proportions of the allocation. We begin our analysis with the case for which $w_i(s) = w(s_i)$ and

* Research of the second author was supported in part by the NSF Grant ANI-031976.

$c_i(s) = c(s_i)$. These restrictions incorporate the notion of fairness where each agent is given the same weight and pays the same cost as any other agent who makes the same signal. It also removes the coupling of the signals away from the weight and cost functions and isolates the interaction in the allocation rule. We assume that the weights and costs are strictly increasing functions of their arguments. We also assume that a signal of zero will yield a weight and cost of zero as well. We consider this class of rules to be a reasonable and tractable initial expansion around the proportionally fair auction which is the "point" in mechanism space characterized by $w(s_i) = s_i$ and $c(s_i) = s_i$. It can be shown that we do not need to express both $w(s_i)$ and $c(s_i)$. By making the substitutions $t_i = c(s_i)$ and $\tilde{w}(t_i) := w(c^{-1}(t_i))$ (c is invertible if it is monotonically increasing), we can express this class of mechanisms with the rules $x_i(t) = \tilde{w}(t_i)/\left(\sum_{j=1}^{N} \tilde{w}(t_j) + \epsilon\right)$ and $c_i(t) = t_i$. With similar substitutions, we can equivalently express this class with the rules

$$x_i(s) = \frac{s_i}{\sum_{j=1}^{N} s_j + \epsilon} \qquad c_i(s) = c(s_i). \tag{1}$$

We choose to work with the characterization described in (1) where $c(s_i) \in C^2$ is a twice differentiable increasing function of s_i. We model the agents with quasilinear utility functions $U_i(s) = v_i(x_i(s)) - c_i(s)$ where v_i is a twice differentiable concave increasing function ($v_i'(\cdot) > 0,\ v_i''(\cdot) \le 0$). We have the derivatives $U_i'(s) = v_i'(x_i(s))x_i'(s) - c'(s_i)$ and $U_i''(s) = v_i''(x_i(s))[x_i'(s)]^2 + v_i'(x_i(s))x_i''(s) - c''(s_i)$ where $x_i'(s) = (s_{-i} + \epsilon)/((s_i + s_{-i} + \epsilon)^2 > 0$, and $x_i''(s) = -2(s_{-i} + \epsilon)/(s_i + s_{-i} + \epsilon)^3 < 0$. If $c_i''(s_i) \ge 0$, then we have $U_i''(s) < 0$. The strict concavity of the i^{th} agent's utility implies that it will have a unique optimal response to each opponent state $s_{-i} + \epsilon$. For the optimal response to be nonzero, we need the marginal utility when bidding zero to be positive. This occurs when $v_i'(0)/(s_{-i} + \epsilon) - c_i'(0) > 0 \ \Rightarrow\ v_i'(0)/c_i'(0) > s_{-i} + \epsilon$. The i^{th} agent's response can then be determined from

$$v_i'\left(\frac{s_i}{s_i + s_{-i} + \epsilon}\right)\frac{s_{-i} + \epsilon}{(s_i + s_{-i} + \epsilon)^2} - c'(s_i) = 0 \tag{2}$$

which yields the unique optimal s_i when facing $s_{-i} + \epsilon$.

Let us define the term *power*, denoted by p, to be the sum of signals, $p := \sum_j s_j + \epsilon$. The power, p, serves as a measure of demand for the resource and allows us to characterize agents' optimal responses with respect to a parameter which is identical for all agents at equilibrium. Let us call this characterization a demand function, $d(p)$, which captures an agent's allocation as a function of power when it uses the strategy obtained from (2). Thus, the demand function captures that $s_i = d(p)p$ is the optimal response to $s_{-i} + \epsilon = d(p)(1 - p)$. We now investigate the effects if $c''(s_i) < 0$.

Proposition 1 *If the cost function is concave, then there exists a valuation function for which the optimal response is not unique.*

Proof. By making the substitutions $p = \sum_j s_j + \epsilon$ and $x_i = s_i/(\sum_j s_j + \epsilon)$ into (2), we can express the first order necessary condition for the optimal response as $v'(x) = pc'(px)/(1 - x) =: f(x|p)$. The derivative of the RHS of the previous equation, $f(x|p)$, as a function of allocation x is $\left[(1 - x)p^2c''(px) + pc'(px)\right]/(1 - x)^2$. The sign of $f'(x|p)$ is determined by the quantity $c''(px) + c'(px)/(p(1 - x))$. If $c''(s) < 0$, then $c''(s) < -\delta,\quad \forall s \in [s_1, s_2]$, for some $\delta > 0$ sufficiently small and some $s_1, s_2 > 0$. Then, $c''(px) + c'(px)/(p(1 - x)) < -\delta + c'(\hat{s})/(p - s_2) < 0$ for p sufficiently large, where $\hat{s} = \arg\max_{s \in [s_1, s_2]} c'(s)$. Specifically, if $p > c'(\hat{s})/\delta + s_2$, we know that $f(x|p)$ is decreasing on $x \in [s_1/p, s_2/p]$. Clearly, there are many functions v with $v''(\cdot) < 0$, where $v'(x)$ will intersect $f(x|p)$ more than once, which implies that there is more than one extremal point. One can construct a variety of valuation functions where there are multiple optimal responses. For p sufficiently large, let $v'(x)$ be a decreasing function on the interval $[0, s_1/p]$, where $v'(s_1/p) = f(s_1/p\ |p)$, and a decreasing function on the interval $[s_2/p, 1]$, where $v'(s_2/p) = f(s_2/p\ |p)$. Then, if $v'(x) = f'(x|p) + \alpha \sin(2\pi k(px - s_1)/(s_2 - s_1))$ on $[s_1/p, s_2/p]$ and α is sufficiently small to ensure that $v'(x)$ is decreasing, we can have an arbitrary number of extremal points by choosing k appropriately. A graphical representation of multiple optimal responses is shown in Figure 1. ■

We restrict our analysis to allocation mechanisms described by (1) where the cost function $c(s_i)$ is convex. We denote this class of mechanisms by $\mathcal{C}$. The intuition behind convex cost functions is that agents who receive larger allocations (due to greater signals) pay a higher cost per unit resource obtained. These occur for strictly convex cost functions and are classified as *discriminatory price* auctions. Mechanisms in $\mathcal{C}$ with linear cost functions such as the proportionally fair auction are *uniform price* auctions. For games played by agents attempting to gain access using a resource allocated through a mechanism from $\mathcal{C}$, it is important to know whether we can obtain a unique operating point, i.e., a unique Nash equilibrium.

Proposition 2 *For every mechanism in $\mathcal{C}$, there is a unique Nash equilibrium.*

Proof. Making the substitutions $p = \sum_j s_j + \epsilon$ and $x_i = s_i/(\sum_j s_j + \epsilon)$ into (2), we can express the first

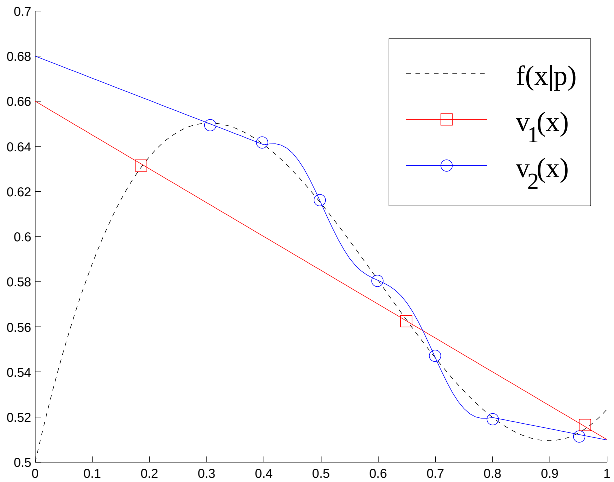

Figure 1. Graphical Representation of Multiple Extremal Points for Mechanisms with Concave Cost Functions

order necessary condition for the optimal response as

$$v_i'(x_i)(1-x_i) = pc_i'(px_i) \ . \tag{3}$$

Every pair (p, x_i) that satisfies the previous equation represents an optimal state for the i^{th} agent. We can interpret these states as demand functions (where x_i is a function of p). By treating the previous equation as an identity, we obtain $[v_i''(x_i)(1-x_i) - v_i'(x_i)] \ \frac{\partial x_i}{\partial p} = c_i'(px_i) + pc_i''(px_i)\left[x_i + p\frac{\partial x_i}{\partial p}\right]$ which implies $\frac{\partial x_i}{\partial p} = \frac{c_i'(px_i)+px_ic_i''(px_i)}{v_i''(x_i)(1-x_i)-v_i'(x_i)-p^2c_i''(px_i)}$. Because $c''(s_i) \geq 0$ for all mechanisms in $\mathcal{C}$, and the valuations are increasing concave functions, we have that $\partial x_i/\partial p < 0$. This implies that the demand functions $\{d_i(p)\}_{i=1}^N$ which characterize the optimal responses of agents are decreasing, where $d_i(p) := x_i(p)$ is obtained from the unique value of x_i which solves (3) for a particular p. We note that $d_i(0) = 1 \ \forall i$. Following the reasoning in [4,6], since all agents are characterized by decreasing demand functions, the total demand will be a decreasing function. The Nash equilibrium point is defined by total demand being one which occurs at only one level of demand p^*. Thus, there is a unique Nash equilibrium with signals $s_i = d_i(p^*)p^*$.■

Given that we have a class of auctions that yield the desirable property of a unique Nash equilibrium, a natural question is how we go about choosing a mechanism within $\mathcal{C}$. In the next sections, we consider this question with social welfare maximization as a metric.

3. Efficiency in $\mathcal{C}$

A common performance measure is the efficiency of the system. Whether allocating bandwidth or processing share, it is desirable to have the resource partitioned in a way that yields the greatest benefit to those accessing it. In economic terms, efficiency is also referred to as the social welfare of those participating in the allocation process. Social welfare is the sum of the valuations of allocations to all agents receiving service, i.e., $\sum_{i=1}^N v_i(x_i)$ where $v_i(x_i)$ are increasing concave functions. Given a scarce resource, the optimal allocations are obtained from the solution of the optimization problem $x^* = \arg\max_{x\in X} \sum_{i=1}^N v_i(x_i)$ where $X = \{x : x_i \geq 0, \sum_{i=1}^N x_i = 1\}$ which can be found by maximizing the Lagrangian $\mathcal{L} = \sum_{i=1}^N v_i(x_i) + \lambda\left(1 - \sum_{i=1}^N x_i\right) + \sum_{i=1}^N \mu_i x_i$. This is a classical optimization problem whose solution is discussed in [5]. Essentially, an optimal solution is an allocation of the entire available resource where the marginal valuations for agents with positive allocations are all equal. The value of the marginal valuations is the Lagrange multiplier. Agents that do not receive positive allocations are those whose highest marginal valuations are less than the value of the Lagrange multiplier. Mathematically, the optimal allocations are characterized as follows: $x_i > 0 \ \Leftrightarrow \ v_i'(x_i) = \lambda, \ x_i = 0 \ \Leftrightarrow \ v_i'(0) \leq \lambda$ where λ is chosen such that $\sum_{i:x_i>0} x_i = \sum_{i:x_i>0} v_i'^{-1}(\lambda) = 1$ where $v_i'^{-1}(\lambda)$ is the inverse of the i^{th} agent's valuation function. The intuition behind the solution is as follows. Given an allocation, if one agent (say the i^{th} agent) has a higher marginal valuation than another (say the j^{th} agent), the social welfare can be improved by marginally increasing x_i and marginally decreasing x_j. Thus, in an optimal allocation, all active agents should have identical marginal valuations.

Because the systems we are considering (network resources on the Internet, computational resources for software agents, etc.) involve autonomous decision making entities, we need to devise distributed methods to reach optimal allocations. Even if a central authority had the responsibility to partition the resource, it is doubtful that it would have access to the private valuation functions of all the agents. In this section, we investigate the effectiveness of market mechanisms in achieving efficient allocations. We begin with the class of mechanisms in $\mathcal{C}$ and show that linear cost functions outperform strictly convex cost functions with respect to social welfare. We then expand our rule space beyond $\mathcal{C}$ to devise maximum efficiency divisible auctions. We also assume that $\epsilon = 0$ (i.e., the entire resource is allocated) for mathematical and notational simplicity. Extensions to the case where $\epsilon > 0$ are straightforward.

Proposition 3 *Within the class of mechanisms in $\mathcal{C}$, social welfare is maximized when the cost functions are linear.*

Proof. The allocation for the i^{th} agent is determined by $v_i'(x_i) = pc'(px_i)/(1-x_i) =: f(p, x_i)$. We have $c''(\cdot) \geq 0$ for all cost functions in $\mathcal{C}$, which implies that $f(p, x_1) > f(p, x_2)$ if $x_1 > x_2$ for all $p > 0$. If $v_i'(0) \leq f(p, 0)$, then the i^{th} agent will not receive a positive allocation at power level p. Otherwise, the i^{th} agent will receive a unique positive allocation $x_i(p)$ which is the solution of $v_i'(x_i) = f(p, x_i)$. The uniqueness is because v_i' is a decreasing function of x_i and f is a strictly increasing function of x_i. We also have $f(p_1, x) > f(p_2, x)$ if $p_1 > p_2$ for all $x \in (0, 1)$. This implies that $x_i(p_1) < x_i(p_2)$ if $p_1 > p_2$. Given any cost function $c(s)$, we can find the equilibrium power level p by solving $\sum_i x_i(p) = 1$, from which we can obtain the equilibrium allocations $\{x_i(p)\}$.

For a linear cost function $c(s) = ks$, we have $f(p, x_i) = pk/(1-x_i)$. We note that every linear cost function yields the same allocations to participating agents, though the equilibrium power level might differ. Let us assume that for $c(s) = k_1 s$, the equilibrium power level is p_1. Then, if $c(s) = k_2 s$, $p_2 = p_1 k_1/k_2$ satisfies all the conditions for equilibrium with the same allocations as the case with $c(s) = k_1 s$. Alternatively, we can let $\tilde{p} = pk$, and obtain $x_i(\tilde{p})$ from the equations $v_i'(x_i) = \tilde{p}/(1-x_i) =: \tilde{f}(\tilde{p}, x_i)$. The equilibrium allocations for all linear cost functions $\{x_i(\tilde{p}^*)\}$ can be expressed in terms of $\tilde{p}^*$, which is the solution to $\sum_i x_i(\tilde{p}) = 1$.

Let us now consider a strictly convex cost function $c(s)$. Let us assume that for some $\hat{x}$ and $\hat{p}$, we have $\hat{p}c'(\hat{p}\hat{x})/(1-\hat{x}) = \tilde{p}^*/(1-\hat{x}) \;\Rightarrow\; \hat{p}c'(\hat{p}\hat{x}) = \tilde{p}^*$. Because $c''(\cdot) > 0$, we have $\forall x > \hat{x}$, $\hat{p}c'(\hat{p}x)/(1-x) > \hat{p}c'(\hat{p}\hat{x})/(1-x) = \tilde{p}^*/(1-x)$ and $\forall x < \hat{x}$, we have $\hat{p}c'(\hat{p}x)/(1-x) > \hat{p}c'(\hat{p}\hat{x})/(1-x) = \tilde{p}^*/(1-x)$ What we have shown is that if $f_c(p, x) := pc'(px)/(1-x)$ (obtained from a strictly convex function) intersects $\tilde{f}(\tilde{p}, x)$ at some point $\hat{x}$, then f_c will be larger than $\tilde{f}$ for $x > \hat{x}$ and f_c will be less than $\tilde{f}$ for $x < \hat{x}$. This is illustrated in Figure 2.

Let us now assume that we have an arbitrary heterogeneous agent population. Let $x_i^c(p)$ be obtained from the solution of $v_i'(x_i) = f_c(p, x_i)$ and $\tilde{x}_i$ be obtained from the solution of $v_i'(x_i) = \tilde{f}(\tilde{p}^*, x_i)$. Let $x_{\min} := \min\{\tilde{x}_i : \tilde{x}_i > 0\}$ be the smallest positive allocation and $x_{\max} := \max\{\tilde{x}_i\}$ be the largest allocation under a linear cost function. Since we are comparing

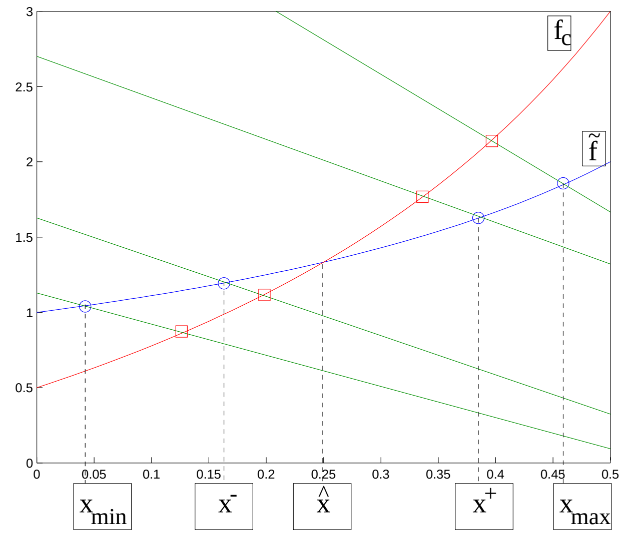

Figure 2. Illustration of $\tilde{f}$ and f_c Along with Marginal Valuations for Four Agents

the performance over all agent populations, we consider one for which $x_{\min} < x_{\max}$. If $p_{\min}$ is the power level at which $f_c(p_{\min}, x_{\min}) = p_{\min}c'(p_{\min}x_{\min})/(1-x_{\min}) = \tilde{p}^*/(1-x_{\min}) = \tilde{f}(\tilde{p}^*, x_{\min})$ then $f_c(p_{\min}, x) > \tilde{f}(\tilde{p}^*, x)$ for all $x > x_{\min}$, and $\sum_i x_i^c(p_{\min}) < \sum_i \tilde{x}_i = 1$. Similarly, if $p_{\max}$ is the power level at which $f_c(p_{\max}, x_{\max}) = p_{\max}c'(p_{\max}x_{\max})/(1-x_{\max}) = \tilde{p}^*/(1-x_{\max}) = \tilde{f}(\tilde{p}^*, x_{\max})$ then $f_c(p_{\max}, x) < \tilde{f}(\tilde{p}^*, x)$ for all $x < x_{\max}$, and $\sum_i x_i^c(p_{\max}) > \sum_i \tilde{x}_i = 1$.

This implies that the power level $\hat{p}$ for which $\sum_i x_i^c(\hat{p}) = 1$ must lie strictly between $p_{\min}$ and $p_{\max}$. Furthermore, $f_c(\hat{p}, x)$ intersects $\tilde{f}(\tilde{p}^*, x)$ at a point $\hat{x} \in (x_{\min}, x_{\max})$. For $x > \hat{x}$, $f_c(\hat{p}, x) > \tilde{f}(\tilde{p}^*, x) \;\Rightarrow\; \tilde{x}_i > x_i^c(\hat{p}) > \hat{x}$. For $x < \hat{x}$, $f_c(\hat{p}, x) < \tilde{f}(\tilde{p}^*, x) \;\Rightarrow\; \tilde{x}_i < x_i^c(\hat{p}) < \hat{x}$. If ΔSW is the difference in social welfare between the allocations for a linear cost function $(\tilde{x})$ and the allocations for a strictly convex cost function $(x^c(\hat{p}))$, we have $\Delta SW = \sum_i v_i(\tilde{x}_i) - \sum_i v_i(x_i^c(\hat{p})) = \sum_i \int_0^{\tilde{x}_i} v_i'(x)\, dx - \sum_i \int_0^{x_i^c(\hat{p})} v_i'(x)\, dx = \sum_{i:\tilde{x}_i > \hat{x}} \int_{x_i^c(\hat{p})}^{\tilde{x}_i} v_i'(x)\, dx - \sum_{i:\tilde{x}_i < \hat{x}} \int_{\tilde{x}_i}^{x_i^c(\hat{p})} v_i'(x)\, dx$. If $x^+ := \min\{\tilde{x}_i : \tilde{x}_i > \hat{x}\}$ and $x^- := \max\{\tilde{x}_i : \tilde{x}_i < \hat{x}\}$, then $\Delta SW > \sum_{i:\tilde{x}_i > \hat{x}} \int_{x_i^c(\hat{p})}^{\tilde{x}_i} \frac{\tilde{p}^*}{1-x^+}\, dx - \sum_{i:\tilde{x}_i < \hat{x}} \int_{\tilde{x}_i}^{x_i^c(\hat{p})} \frac{\tilde{p}^*}{1-x^-}\, dx = \frac{\tilde{p}^*}{1-x^+} \sum_{i:\tilde{x}_i > \hat{x}} [\tilde{x}_i - x_i^c(\hat{p})] - \frac{\tilde{p}^*}{1-x^-} \sum_{i:\tilde{x}_i < \hat{x}} [x_i^c(\hat{p}) - \tilde{x}_i]$. Since the sums of allocations in both cases are one, we have $\sum_{i:\tilde{x}_i > \hat{x}} [\tilde{x}_i - x_i^c(\hat{p})] = [\sum_{i:\tilde{x}_i > \hat{x}} \tilde{x}_i] - 1 + 1 - [\sum_{i:\tilde{x}_i > \hat{x}} x_i^c(\hat{p})] = [\sum_{i:\tilde{x}_i > \hat{x}} \tilde{x}_i] - [\sum_i \tilde{x}_i] + [\sum_i x_i^c(\hat{p})] - [\sum_{i:\tilde{x}_i > \hat{x}} x_i^c(\hat{p})] = \sum_{i:\tilde{x}_i < \hat{x}} [x_i^c(\hat{p}) - \tilde{x}_i] =: \alpha > 0$. Incorporating this into the bound on social welfare difference, we have $\Delta SW > \alpha \tilde{p}^* \left[1/(1-x^+) - 1/(1-x^-)\right] > 0$. ■

The optimality of linear cost functions (including the proportionally fair auction) among mechanisms in $\mathcal{C}$ gives credence to their use in addition to their practical benefits (ease of implementation, etc.). However, we still are unable to reach optimal efficiency by limiting ourselves to this class of allocation and cost rules.

Corollary 1 *There is no mechanism in $\mathcal{C}$ which optimally maximizes social welfare for all agent populations.*

Proof. We consider the case of two agents with valuation functions such that $v_1'(x) > v_2'(x)$ for all $x \in [0,1]$. We consider a mechanism with a linear cost function as it yields the optimal efficiency among all cost functions in $\mathcal{C}$. The allocation for Agent 1 is obtained from the solution to $v_1'(x_1) = p/(1-x_1)$ for some $p > 0$. This implies that $x_1 < 1$, which in turns implies that $x_2 > 0$. Since both agents are active, their marginal valuation functions must intersect $p/(1-x)$, which is an increasing function of x. This implies that the agents' marginal valuations are not equal at equilibrium unless $x_1 = x_2 = 1/2$, and that cannot occur because $v_1'(1/2) > v_2'(1/2)$. Since both agents are active and their marginal valuations are not equal, the social welfare can be improved and thus the mechanism is not optimal. ■

4. Designing Maximum Efficiency Divisible Auctions

The intuition behind why mechanisms in $\mathcal{C}$ cannot be optimally efficient lies in the equilibrium condition $v_i'(x_i) = pc'(px)/(1-x) =: f(p,x)$. Here, $f(p,x)$ plays the role of the Lagrange multiplier. For cost functions that are convex, $f(p,x)$ is an increasing function of x. For optimality we need a cost function that would yield a function $f(p,x)$ which was independent of x. This would make $f(p,x) = g(p)$ "flat" across $x \in [0,1]$ and the marginal valuation functions of all active agents would intersect f at the same value. We refer to $g(p)$ as the generator function. The generator function serves the purpose of the Lagrange multiplier in the social welfare maximization problem. Thus, if we are to use this generator function to obtain a maximum efficiency cost function for all agent populations, it must be able to span all viable values that a Lagrange multiplier might take, i.e., all nonnegative real numbers. Furthermore, we need the generator function to be one-to-one. Otherwise, an agent population whose optimal allocations occur at a particular Lagrange multiplier value could be reached at two different power levels, which indicates multiple equilibria. We now show that by using an appropriate generator function, we can construct cost functions that yield an equilibrium at which active agents have the same marginal valuation for their allocations.

Proposition 4 *Let $g(p)$ be a one-to-one function whose range space is the set of all nonnegative real numbers. Consider the mechanism $x_i(s) = s_i/(s_i + s_{-i})$, $c_i(s) = s_{-i} \int_0^{s_i} g(t+s_{-i})/(t+s_{-i})^2 dt$ where $s_{-i} = \sum_{j \neq i} s_j + \epsilon$. For an arbitrary agent population, any equilibrium under this mechanism will yield a solution where all active agents have the same marginal valuation.*

Proof. First, we explain the construction of the cost function. If we set the marginal valuation to be equivalent to the generator function, we have $v_i'(x_i) = pc'(px_i)/(1-x_i) = g(p) \Rightarrow c'(px_i) = g(p)(1-x_i)/p = g(p)[p(1-x_i)]/p^2$. At equilibrium, we have $px = s_i$, $p(1-x) = s_{-i}$, and $p = s_i + s_{-i}$. We realize that we cannot express the marginal cost solely as a function of s_i. However, we can express it as a function of s_i and s_{-i} as follows: $c_i'(s_i; s_{-i}) = g(s_i + s_{-i}) s_{-i}/(s_i + s_{-i})^2$. If we assume that this expression for marginal cost holds for all equilibrium solutions, we can integrate over s_i to obtain

$$c_i(s) = c(s_i; s_{-i}) = s_{-i} \int_0^{s_i} \frac{g(t+s_{-i})}{(t+s_{-i})^2}\, dt.$$

We now verify that this yields equal marginal valuations at equilibrium for active agents by analyzing agents' optimal responses. If an agent with utility $U_i(s) = v_i(x_i(s)) - c_i(s)$ has a positive allocation at equilibrium, its optimal signal is an extremal point obtained as a solution of $U_i'(s) = v_i'\left(s_i/(s_i+s_{-i})\right) s_{-i}/(s_i+s_{-i})^2 = c'(s_i; s_{-i})$. Substituting the marginal cost , we have $v_i'\left(s_i/(s_i+s_{-i})\right) s_{-i}/(s_i+s_{-i})^2 = g(s_i+s_{-i}) s_{-i}/(s_i+s_{-i})^2 \Rightarrow v_i'\left(s_i/(s_i+s_{-i})\right) = g(s_i+s_{-i}) = g(p)$. Thus, the marginal valuations of active agents for any set of bids that form an equilibrium solution are identical. In fact, the value of the marginal valuations is the output of the generator function at the equilibrium power level. Furthermore, for any inactive agent at equilibrium, we have $v_i'(0) \le s_{-i} c'(0; s_{-i}) = s_{-i}^2 g(0+s_{-i})/(0+s_{-i})^2 = g(s_{-i}) = g(p)$ which meets our conditions for a solution to the social welfare maximization problem. ■

We note that the key to these mechanisms is the s_i factor in the cost functions as it cancels the s_{-i} that appears when we take the partial derivative of $v_i(x_i(s))$ with respect to s_i. In essence, by making agents account for increased demand in their costs as well as the allocation, we are able to achieve maximum efficiency. In the following examples, we obtain cost functions associated with various generator functions.

Example 1 *We obtain the cost function generated by $g(p) = p$:* $c(s_i; s_{-i}) = s_{-i} \int_0^{s_i} 1/(t + s_{-i})\, dt = s_{-i} [\log(s_i + s_{-i}) - \log(s_{-i})] = s_{-i} \log(1 + s_i/s_{-i})$ □

Example 2 *We obtain the cost function generated by $g(p) = p^k$ where $k > 0, k \neq 1$:* $c_i(s) = s_{-i} \int_0^{s_i} (s_i + s_{-i})^{k-2}\, dt = (s_{-i}/k - 1) \left[(s_i + s_{-i})^{k-1} - s_{-i}^{k-1}\right]$ □

Example 3 *We obtain the cost function generated by $g(p) = p^2 e^p$:* $c_i(s) = s_{-i} \int_0^{s_i} e^{s_i + s_{-i}}\, dt = s_{-i} [e^{s_i + s_{-i}} - e^{s_{-i}}] = s_{-i} e^{s_{-i}} (e^{s_i} - 1)$ □

We see that we can generate a diverse set of cost functions that yield equal marginal valuations at equilibrium. We have cost functions that are both concave (Example 1) and convex (Example 3) in the agents' signals. All cost functions yield a cost of zero if the agent bids zero. The cost function with the simplest and most intuitive form may be that generated by $g(p) = p^2$, which yields $c(s) = s_i s_{-i}$. This states that an agent's cost depends linearly on both its own signal and the sum of signals of all other agents.

To this point, we have neglected to analyze the effect of the cost function on equilibrium. We know that if an equilibrium exists, it will maximize social welfare. The question that follows is what generator functions yield cost functions that lead to the existence of a unique equilibrium.

Proposition 5 *Let $g(p)$ be a one-to-one function whose range space is the set of all non-negative real numbers. Furthermore, let $g'(p)$ exist and be positive for all $p \geq 0$. Then, the mechanism using the cost function generated by $g(p)$ yields a unique equilibrium.*

Proof. We show this by obtaining demand functions and showing that they are decreasing functions of power level. We already know that $v'(x) = g(p)$ at equilibrium. Taking this as an identity, we have $v''(x)\frac{\partial x}{\partial p} = g'(p) \Rightarrow \frac{\partial x}{\partial p} = g'(p)/v''(x) < 0$ for all concave valuation functions. Since the demand functions are decreasing, we can apply similar reasoning from Proposition 2 to state that we have a unique Nash equilibrium. ■

To obtain a unique Nash equilibrium, we must limit ourselves to strictly increasing generator functions. Examples 1, 2, and 3 all satisfy this requirement. If we generated a cost function using $g(p) = p^k$ with $k \leq 0$, the resulting demand functions would not decrease and no equilibrium solution would exist. Generator functions that are both increasing and decreasing may lead to multiple equilibria.

Nevertheless, the set of strictly increasing functions is a large class from which we can obtain generator functions. We can obtain an infinite number of mechanisms which maximize social welfare.

5. Conclusion

We have addressed the efficiency of markets for divisible resources. We first looked at proportional schemes and showed that even though linear weighting is optimal in a particular subset of mechanism space, it is not maximally efficient. We then developed a method to generate an entire class of divisible auction rules that maximize social welfare. System designers can now address secondary performance measures and optimize over this class. Furthermore, all the mechanisms described here use one-dimensional signaling and take $O(N)$ steps to calculate the allocation. Thus, they have the lowest possible signaling and computational costs for a divisible auction. Because of transparency, efficiency, and overhead, one might surmise that the auctions presented here are ideal allocation mechanisms. However, we must still investigate the effects of cooperative phenomena and other higher degree responses to fully understand the consequences of implementing these systems. In addition, since equilibrium is left to the agents, learning stable decentralized negotiation schemes are important components that need to be developed.

References

[1] R. A. Gagliano, M. D. Fraser, and M. E. Schaefer. Auction allocation for computing resources. *Communications of the ACM*, 38(6):88–102, June 1995.

[2] F. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8(1):33–37, January 1997.

[3] J. K. MacKie-Mason and H. R. Varian. Pricing the internet. In B. Kahin and J. Keller, editors, *Public Access to the Internet*, pages 269–314. MIT Press, Cambridge, MA, 1995.

[4] R. T. Maheswaran and T. Başar. Auctions for divisible resources: Price functions, Nash equilibria and decentralized update schemes. In *AAMAS 2002: Workshop on Agent-Mediated Electronic Commerce IV*, Bologna, Italy, July 2002.

[5] R. T. Maheswaran and T. Başar. Coalition formation in proportionally fair divisible auctions. In *Autonomous Agents and Multi-Agent Systems 2003*, Melbourne, Australia, July 2003.

[6] R. T. Maheswaran and T. Başar. Nash equilibrium and decentralized negotiation in auctioning divisible resources. *Group Decision and Negotiation*, 12(5):361–395, September 2003.

[7] C. Waldspurger and W. E. Weihl. Stride scheduling: Deterministic proportional-share resource management. Technical Memorandum MIT/LCS/TM-528, MIT Laboratory for Computer Science, June 1995.