

Efficiency of Scalar-Parameterized Mechanisms

Ramesh Johari

Department of Management Science and Engineering, Stanford University, Stanford, California 94305,
ramesh.johari@stanford.edu

John N. Tsitsiklis

Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139,
jnt@mit.edu

We consider the problem of allocating a fixed amount of an infinitely divisible resource among multiple competing, fully rational users. We study the efficiency guarantees that are possible when we restrict to mechanisms that satisfy certain scalability constraints motivated by large-scale communication networks; in particular, we restrict attention to mechanisms where users are restricted to one-dimensional strategy spaces. We first study the efficiency guarantees possible when the mechanism is not allowed to price differentiate. We study the worst-case efficiency loss (ratio of the utility associated with a Nash equilibrium to the maximum possible utility), and show that Kelly's proportional allocation mechanism minimizes the efficiency loss when users are price anticipating. We then turn our attention to mechanisms where price differentiation is permitted; using an adaptation of the Vickrey-Clarke-Groves class of mechanisms, we construct a class of mechanisms with one-dimensional strategy spaces where Nash equilibria are fully efficient. These mechanisms are shown to be fully efficient even in general convex environments, under reasonable assumptions. Our results highlight a fundamental insight in mechanism design: when the pricing flexibility available to the mechanism designer is limited, restricting the strategic flexibility of bidders may actually improve the efficiency guarantee.

Subject classifications: game/group decisions: noncooperative, bidding/auctions; networks/graphs: theory; utility/preference: theory.

Area of review: Revenue Management.

History: Received January 2007; revisions received November 2007, March 2008, June 2008; accepted July 2008.

Published online in *Articles in Advance* May 6, 2009.

1. Introduction

We consider the problem of allocating a fixed amount of an infinitely divisible resource among multiple competing, fully rational users, a problem that has received extensive attention in recent literature on pricing of communication resources (Shenker et al. 1996, Falkner et al. 2000). We investigate a fundamental question: What is the best achievable efficiency when users are fully rational? We study the efficiency guarantees that are possible when we restrict ourselves to mechanisms that satisfy certain constraints motivated by large-scale communication networks.

Our starting point is a mechanism introduced by Kelly (1997). For the special case of a single resource in fixed supply, his scheme is easily described. Each user chooses a total payment, or *bid*. The resource is then allocated to users in proportion to their bids, and each user pays an amount equal to his bid; equivalently, a market-clearing price is set equal to the sum of the bids divided by the total amount of available resource.

This scheme has two desirable properties: First, the strategy space of the users is “simple,” i.e., one-dimensional; and second, the feedback required from the mechanism to the users is a single price per unit of resource. Both features are central requirements for a mechanism to be practical in the context of large networks. To keep the communication from the users to the network simple, it is

reasonable to restrict users to one-dimensional bids; we refer to mechanisms that restrict users to such bids as *scalar-parameterized mechanisms*. Further, in communication networks, mechanisms that use “aggregate feedback” are considered scalable (Shenker 1990); this is one motivation for mechanisms that choose only a single price. The Kelly mechanism is designed to incorporate both of these design principles simultaneously (see also Kelly et al. 1998 and Srikant 2004).

Kelly (1997) considered a model where each user is endowed with a concave, increasing utility that is a function of the received resource allocation. He showed that if users do not anticipate the effect of their bid choices on market-clearing prices, then at equilibrium, a fully efficient allocation is achieved—that is, an allocation where the aggregate utility of the users has been maximized. However, this analysis assumed that users are not fully rational and do not fully anticipate the effects of their bids on prices. By contrast, Johari and Tsitsiklis (2004) considered a game-theoretic model of the Kelly mechanism where users are price anticipating, and established that the aggregate utility at a Nash equilibrium is no worse than 75% of the maximal aggregate utility.

In this paper, we investigate the *optimality* of the result of Johari and Tsitsiklis (2004). Our results provide two major insights. First, we show that if we restrict attention

to a class of mechanisms where users only submit one-dimensional bids, and the mechanism only chooses a single price, then the Kelly mechanism exhibits the best possible efficiency guarantee. Formally, we consider a class of market-clearing mechanisms; in these mechanisms each user submits a scalar-parameterized demand function, and a single price is chosen to clear the market. Under reasonable assumptions, we show that the Kelly mechanism minimizes the worst-case efficiency loss at a Nash equilibrium when users are price anticipating. This result complements the work of Maheswaran and Basar (2004): they reached similar conclusions for a qualitatively different class of mechanisms involving differentiated prices that are functions of individual bids, as opposed to our single prices that are functions of all bids. See §2.1 for details.

Our second major insight is aimed at understanding the importance of the *single-price* assumption in the preceding result. To investigate this point, we preserve the assumption that bids are one-dimensional, but remove the restriction that the mechanism only chooses a single price. We ask: How much can we improve the efficiency guarantee if the mechanism is allowed to choose a different price for each user? Using a class of mechanisms inspired by the Vickrey-Clarke-Groves (VCG) approach (Vickrey 1961, Clarke 1971, Groves 1973), we show that if users are restricted to using one-dimensional bids, but the mechanism can price differentiate (i.e., set a different price for each user), then fully efficient Nash equilibria can be achieved—not only in a single-resource setting, but also in general convex environments. We call these mechanisms *scalar strategy VCG (SSVCG) mechanisms*.¹ We note that as a special case, SSVCG mechanisms contain an earlier class of mechanisms obtained by Yang and Hajek (2006a) and Maheswaran (2003); we further discuss this relationship in §§2.2 and 5.2.

Outline. The remainder of this paper is organized as follows. In §2, we discuss related work. In §3, we give the mathematical framework of our optimal mechanism design problem. We restrict attention to allocation of a single infinitely divisible resource among multiple users. Each user is endowed with a concave, strictly increasing, differentiable utility function, and makes strategic decisions to maximize his utility less payment.

In §3.1, we begin by embedding the Kelly mechanism into a broader class of market-clearing mechanisms. We make a simple but crucial observation: a user's bid θ establishes a relation $d = \theta/p$ between the quantity d allocated to the user, and the market-clearing price p per unit of the resource. This can be viewed as having the user submit a demand function $D(p, \theta)$, which specifies demand as a function of p . The user is, however, limited to a particular one-parameter family of demand functions, demand functions of the form $D(p, \theta) = \theta/p$; choosing a value for the parameter θ is all the freedom given to the user.

Although the Kelly mechanism assumes a specific form for $D(p, \theta)$, an entire family of market-clearing mechanisms

results if we consider alternative forms for $D(p, \theta)$. This added generality retains the simplicity of the Kelly mechanism: each user selects from a one-dimensional strategy space, and a single price is charged to all users. Formally, we restrict attention to *smooth market-clearing mechanisms* defined by a differentiable parameterized demand function $D(p, \theta)$.² In these mechanisms, each user r chooses a strategy $\theta_r \geq 0$. A price $p^* = p_D(\theta)$ is then chosen to *clear the market*, i.e., so that $\sum_r D(p^*, \theta_r) = C$, where C is the supply, the amount of available resource. Each user r then receives an allocation $D(p^*, \theta_r)$, and pays $p^* D(p^*, \theta_r)$. A smooth market-clearing mechanism is thus entirely specified by the parametric family of demand functions $D(p, \theta)$. We emphasize that although the class of demand functions is constrained in this way, no such constraints are imposed on the true utility functions of the users (other than concavity and monotonicity).

In §3.2, we recapitulate existing results regarding the Kelly mechanism. We note that the Kelly mechanism yields full efficiency when users are price taking, and guarantees an *efficiency* of 75%: i.e., the aggregate utility at a Nash equilibrium is no less than 75% of the maximal aggregate utility when users are price anticipating, and this bound is tight.

With this background, our objective is to maximize the market's efficiency. In particular, we wish to determine whether the 75% guarantee of the Kelly mechanism can be improved. In §4, we consider a class \mathcal{D} of smooth market-clearing mechanisms where: (1) users' payoffs are concave when users are *price anticipating*, i.e., when they fully anticipate the functional form of the market-clearing price $p_D(\theta)$; and (2) for each p , the range of $D(p, \theta)$, as θ varies over the nonnegative real numbers, is the entire interval $[0, \infty)$. We explain why these assumptions are reasonable, and give examples of mechanisms in \mathcal{D} beyond the Kelly mechanism.

We show in Theorem 1 that in this class of mechanisms, the Kelly mechanism minimizes the worst-case efficiency loss when users are price anticipating, and is the unique mechanism with this property. In the process, we also give a structural description of all mechanisms in this class: any mechanism satisfying the given assumptions must be of the form $D(p, \theta) = \theta/B(p)$, where B is concave, strictly increasing, and invertible on $(0, \infty)$. Notably, this structural characterization implies that any such mechanism is a *proportional allocation mechanism*: the allocation is always made in proportion to the bids, as is true for the Kelly mechanism. Thus, mechanisms in \mathcal{D} differ only in the prices they choose for any given vector of users' bids.

For any smooth market-clearing mechanism, users have one-dimensional strategy spaces, and the mechanism only chooses a single price. Because of these constraints, even the highest performance mechanism—the Kelly mechanism—suffers a positive efficiency loss. In §5, we consider the implications of removing the “single price” constraint, while retaining the constraint that strategy spaces must be one-dimensional. We show in §5.2 that if we consider

mechanisms with scalar strategy spaces, and allow the mechanism to choose a different price for each user, then in fact full efficiency is achievable at Nash equilibrium; this is shown via a class of mechanisms called SSVCG mechanisms (see above); we show that as a special case of this mechanism, one can recover the mechanism and main conclusions of Yang and Hajek (2006a) and Maheswaran (2003). The result involves adapting the well-known class of VCG mechanisms to the case of scalar strategy spaces. For SSVCG mechanisms, we are also able to consider far more general resource allocation environments. In particular, we show that for a very general class of convex resource allocation problems, if the utility of a user depends only on a scalar function of the vector of resources allocated to that user, then SSVCG mechanisms ensure that full efficiency is achievable at Nash equilibrium. We show that an efficient Nash equilibrium always exists, and we also give reasonable sufficient conditions under which all Nash equilibria are efficient. As an illustration, we apply our SSVCG framework to a standard multicommodity flow problem. We conclude with a discussion of extensions and open problems in §6.

2. Related Work

In this section, we review threads of the existing literature related to our work. One body of recent work concerns the characterization of efficiency loss in scalar-parameterized mechanisms; we compare this work with the results of §4. A second thread of the literature studies the design of scalar-parameterized mechanisms that achieve full efficiency; we compare this work with the results of §5.

2.1. Efficiency Loss of Scalar-Parameterized Mechanisms

In §4, we show that in the class \mathcal{D} of smooth market-clearing mechanisms, the Kelly mechanism uniquely minimizes the worst-case efficiency loss when users are price anticipating. The efficiency loss of the Kelly mechanism was first studied by Johari and Tsitsiklis (2004); subsequently, several others have studied efficiency loss of related resource allocation mechanisms, including Chen and Zhang (2008), Johari et al. (2005), Johari and Tsitsiklis (2006), Maheswaran (2003), Maheswaran and Basar (2004), Moulin (2008b), Sanghavi and Hajek (2004), and Yang and Hajek (2006a). More generally, several works in the recent literature aim at quantifying efficiency losses in games, e.g., for routing (Czumaj and Voecking 2002, Koutsoupias and Papadimitriou 1999, Mavronicolas and Spirakis 2001), traffic networks (Correa et al. 2004, Roughgarden and Tardos 2002), network design problems (Anshelevich et al. 2003, Fabrikant et al. 2003), and supply chain management (Roels and Perakis 2007); for a more comprehensive discussion on work in this area, the reader is referred to the survey text of Nisan et al. (2007).

The most closely related result to ours is presented by Maheswaran and Basar (2004). They consider mechanisms

where each user r chooses a bid w_r , and the allocation is made proportional to each user's bid. However, rather than assuming that every user pays w_r as in the Kelly mechanism, Maheswaran and Basar consider a class of mechanisms where the user pays $c(w_r)$, where c is a convex function.³ They show that in this class of mechanisms, the Kelly mechanism (i.e., a linear function c) achieves the minimal worst-case efficiency loss when users are price anticipating.

We note that, in general, in Maheswaran and Basar's model, different users are charged *different* per-unit prices for the resource they obtain (i.e., pricing is *nonuniform*), unless c is linear—exactly the Kelly mechanism. By contrast, in *every* mechanism in the class of mechanisms we consider, all users are charged the same per-unit price (i.e., the market-clearing price), where the price is a function of the vector of all bids, which depends on the specific mechanism chosen. Thus, our results prove optimality of the Kelly mechanism in a distinct regime from Maheswaran and Basar.

We also note that our work is substantially different in another way from Maheswaran and Basar because we do not postulate a priori that the mechanism must allocate the resource in proportion to users' bids; instead, this emerges as a *consequence* of our rather natural assumptions on the mechanisms in \mathcal{D} . Several other works on efficiency of resource allocation mechanisms, including Maheswaran and Basar (2004) and Yang and Hajek (2006a), assume a priori that allocations are made in proportion to users' bids. In this sense, our result lends a rigorous foundation to the intuition that proportional allocation mechanisms—those that allocate the resource in proportion to users' bids—yield a natural approach to resource allocation among competing users.

We conclude this section by noting that because the class \mathcal{D} induces constraints on the communication both *to* and *from* the mechanism, a strictly positive efficiency loss is inevitable. Our approach in §4 can be viewed as an optimization problem that aims at minimizing efficiency loss subject to communication constraints.

2.2. Fully Efficient Scalar-Parameterized Mechanisms

In §5, we relax the assumption that all users must pay the same per-unit price for the resource. We construct a class of scalar strategy VCG mechanisms that ensure that Nash equilibria are fully efficient. The literature in both economics and, more recently, in operations research has a significant body of work related to this problem. The seminal results of Vickrey (1961), Clarke (1971), and Groves (1973) established that if users can submit entire utility functions, and the mechanism can choose individualized prices (i.e., one price per user), then full efficiency can be achieved as a dominant strategy equilibrium. In our setting, however, because of the restriction to one-dimensional strategy spaces, it is not possible to achieve full efficiency in dominant strategies.

Reichelstein and Reiter (1988) presented a key general result in the theory of mechanism design with restricted strategy spaces. Their paper calculates the minimal strategy space dimension that allows for fully efficient Nash equilibria for a general class of economic models known as *exchange economies*. Their results apply to general convex environments such as those studied in §5.3; however, their mechanisms require a strategy space per user of dimension $(J-1) + J/(R(R-1))$, where J is the number of resources being allocated (e.g., the number of links in a network), and are “asymmetric,” in the sense that different users are treated differently, even if they all choose the same strategies. We are able to use a strategy space smaller than their bound, in part because each user’s utility depends only on a scalar function of the allocation to that user, and in part because we assume quasilinear preferences.

More recently, for the special case of multicommodity flow, Semret (1999) presented a class of mechanisms similar to the VCG class, but where each user has a two-dimensional strategy space; the author demonstrates existence of an approximately efficient Nash equilibrium, but does not establish conditions under which all equilibria are efficient. Maheswaran and Basar (2004) (as well as the earlier work, Maheswaran 2003) and Yang and Hajek (2006a) independently suggested a mechanism for allocation of a single infinitely divisible resource in fixed supply. They demonstrated that Nash equilibria of this mechanism are efficient. Remarkably, as we will see in §5, their mechanism can be recovered as a special case of the SSVCG mechanisms we develop.

We note here that simultaneously and independently of our work described in §5, Yang and Hajek (2007) proposed a class of VCG-like mechanisms with scalar strategy spaces; these results were first presented by Yang and Hajek (2006b). In some respects, their work is more specialized than ours: they restrict attention to network resource allocation settings and consider a less general class of mechanisms. However, their paper also considers various extensions not discussed here: “soft” capacity constraints, specified through penalty functions; revenue of SSVCG mechanisms; and dynamic algorithms that converge to efficient Nash equilibria.

Subsequent to our work (Johari and Tsitsiklis 2005) and the work of Yang and Hajek (2007), several papers have presented related constructions of mechanisms that use limited communication yet achieve fully efficient Nash equilibria (Dimakis et al. 2006, Stoenescu and Ledyard 2006, Moulin 2008a). Building on the work of Semret (1999) discussed above, Dimakis et al. (2006) establish that a VCG-like mechanism where agents submit a pair of strategies (requested price and quantity) can achieve fully efficient equilibrium for a related resource allocation game. Stoenescu and Ledyard (2006) consider the problem of resource allocation by building on the notion of minimal message spaces addressed in earlier literature on mechanism design, and propose a class of efficient mechanisms with restricted strategy spaces.

Finally, Moulin (2008a) considers the design of efficient scalar-parameterized mechanisms where the net monetary transfer from users is minimized.

Several other papers have considered the consequences of restricted communication in a discrete setting where complexity is measured in bits, rather than dimension. Blumrosen et al. (2007) consider a setting where each bidder in an auction for a single indivisible item can only transmit a fixed number of bits to the auctioneer. Fadel and Segal (2009) use a communication-complexity theoretic approach to study the number of bits necessary to achieve efficiency under a variety of notions of equilibrium. Both papers treat substantively different environments from those studied here, and of course the methods of analysis are substantially different as well.

3. Preliminaries

The basic model we study is similar to the one studied by Kelly (1997) and Johari and Tsitsiklis (2004). We consider a collection of R users bidding for a share of a finite, infinitely divisible resource of capacity C . We begin by describing the users. Each user r is endowed with a *utility function* $U_r(\cdot)$ that determines the monetary value of any resource allocation to user r . We let $\mathbf{U} = (U_1, \dots, U_R)$ denote the vector of utility functions. We make the following assumption, which will remain in force throughout the paper, unless explicitly mentioned otherwise.

ASSUMPTION 1. For each r , and over the domain $d_r \geq 0$, the utility function $U_r(d_r)$ is nonnegative, concave, strictly increasing, and continuous; and over the domain $d_r > 0$, $U_r(d_r)$ is continuously differentiable. Furthermore, the right directional derivative at 0, denoted $U'_r(0)$, is finite. We let \mathcal{U} denote the set of all utility functions satisfying these conditions.

Note that although we make rather strong differentiability assumptions, these are not essential to the argument; however, they ease the technical presentation. We call a triple (C, R, \mathbf{U}) , where $C > 0$, $R > 1$, and $\mathbf{U} \in \mathcal{U}^R$, a *utility system*; our goal is to design a resource allocation mechanism that has high efficiency for all utility systems.

We assume that utility is measured in monetary units; thus, if user r receives a rate allocation d_r , but must pay w_r , he receives a net payoff given by⁴

$$U_r(d_r) - w_r.$$

Given any vector of utility functions $\mathbf{U} \in \mathcal{U}^R$, our goal is to maximize aggregate utility, as defined in the following problem:

SYSTEM(C, R, \mathbf{U}):

$$\text{maximize } \sum_{r=1}^R U_r(d_r) \quad (1)$$

$$\text{subject to } \sum_{r=1}^R d_r \leq C, \quad (2)$$

$$\mathbf{d} \geq 0. \quad (3)$$

We will say that \mathbf{d} solves $SYSTEM(C, R, \mathbf{U})$ if \mathbf{d} is an optimal solution to (1)–(3).

3.1. Smooth Market-Clearing Mechanisms

In general, the utility system (C, R, \mathbf{U}) is unknown to the mechanism designer, so a mechanism must be designed to elicit information from the users. The following definition captures the types of market mechanisms that we study. Let $\mathbb{R}^+ = \{x \in \mathbb{R} \mid x \geq 0\}$.

DEFINITION 1. A differentiable function $D: (0, \infty) \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is said to define a *smooth market-clearing mechanism* if for all $C > 0$, for all $R > 1$, and for all nonzero $\boldsymbol{\theta} \in (\mathbb{R}^+)^R$, there exists a unique solution $p > 0$ to the following equation:

$$\sum_{r=1}^R D(p, \theta_r) = C. \quad (4)$$

We let $p_D(\boldsymbol{\theta})$ denote this solution.⁵

We can interpret a smooth market-clearing mechanism in terms of a market-clearing process. Each user r chooses a parameter θ_r , which specifies a demand function $p \mapsto D(p, \theta_r)$, and determines user r 's demand for the resource as a function of the price p of the resource. The mechanism then chooses a single price $p = p_D(\boldsymbol{\theta})$ so that aggregate demand equals the available capacity (cf. (4)). Each consumer r then receives a resource allocation of $D(p, \theta_r)$, and pays $pD(p, \theta_r)$. Note that (1) each user is restricted to a *one-dimensional* strategy space; and (2) the mechanism only uses a *single price* to allocate the resource; i.e., the mechanism does not price discriminate.

REMARK 1. Although our definition implicitly restricts the strategy θ_r of each user to \mathbb{R}^+ , the subsequent analysis can be adapted to hold even if the strategy space of each user is allowed to be $[c, \infty)$, where $c \in \mathbb{R}$.

REMARK 2. The market-clearing price is undefined if $\boldsymbol{\theta} = \mathbf{0}$. As we will see below, when we formulate a game between consumers for a given function D , we will assume that the payoff to all users is $-\infty$ if the composite strategy vector is $\boldsymbol{\theta} = \mathbf{0}$. Note that this is slightly different from the definition in Johari and Tsitsiklis (2004), where the payoff to a user with utility function U and strategy $\theta = 0$ is $U(0)$. We will discuss this distinction further later; we simply note for the moment that it does not affect our results.

We will study the performance of smooth market-clearing mechanisms through their equilibria. Two notions of equilibrium are relevant to our analysis: *competitive equilibrium* and *Nash equilibrium*. In a competitive equilibrium, users act as *price takers*; that is, they do not anticipate the effect of their actions on the market-clearing price. Conversely, at a Nash equilibrium, users are completely rational and anticipate the functional dependence of the market-clearing price $p_D(\boldsymbol{\theta})$ on their own strategic decisions. Nash

equilibrium is a *complete information* solution concept: implicitly, we assume that players' payoff functions are common knowledge. The formal definitions of competitive equilibrium and Nash equilibrium are as follows.

DEFINITION 2. Given a utility system (C, R, \mathbf{U}) and a smooth market-clearing mechanism D , we say that a nonzero vector $\boldsymbol{\theta} \in (\mathbb{R}^+)^R$ is a *competitive equilibrium* if, for $\mu = p_D(\boldsymbol{\theta})$, there holds for all r :

$$\theta_r \in \arg \max_{\bar{\theta}_r \geq 0} [U_r(D(\mu, \bar{\theta}_r)) - \mu D(\mu, \bar{\theta}_r)]. \quad (5)$$

DEFINITION 3. Given a utility system (C, R, \mathbf{U}) and a smooth market-clearing mechanism D , we say that a nonzero vector $\boldsymbol{\theta} \in (\mathbb{R}^+)^R$ is a *Nash equilibrium* if there holds for all r :

$$\theta_r \in \arg \max_{\bar{\theta}_r \geq 0} Q_r(\bar{\theta}_r; \boldsymbol{\theta}_{-r}), \quad (6)$$

where

$$Q_r(\theta_r; \boldsymbol{\theta}_{-r}) = \begin{cases} U_r(D(p_D(\boldsymbol{\theta}), \theta_r)) - p_D(\boldsymbol{\theta})D(p_D(\boldsymbol{\theta}), \theta_r) & \text{if } \boldsymbol{\theta} \neq \mathbf{0}, \\ -\infty & \text{if } \boldsymbol{\theta} = \mathbf{0}. \end{cases} \quad (7)$$

Notice that the payoff Q_r is $-\infty$ if the composite strategy vector is $\boldsymbol{\theta} = \mathbf{0}$ because in this case no market-clearing price exists.

One prominent example of a smooth market-clearing mechanism is the *Kelly mechanism*, studied in the context of communication networks by Kelly (1997); we discuss this mechanism in the next section.

3.2. The Kelly Mechanism

In this section, we define the *Kelly mechanism* and survey some of its basic properties. In the Kelly mechanism, each user r submits a payment (also called a *bid*) of w_r to the resource manager. Given the vector $\mathbf{w} = (w_1, \dots, w_r)$, the resource manager chooses an allocation $\mathbf{d} = (d_1, \dots, d_r)$. The mechanism does not price discriminate; thus, each user is charged the same price $\mu > 0$, leading to $d_r = w_r/\mu$. If the entire capacity C is allocated, then we expect the price μ to satisfy

$$\sum_r \frac{w_r}{\mu} = C.$$

The preceding equality can only be satisfied if $\sum_r w_r > 0$, in which case we have

$$\mu = \frac{\sum_r w_r}{C}. \quad (8)$$

In other words, if the manager chooses to allocate the entire resource and does not price discriminate between users,

then for every nonzero \mathbf{w} there is a *unique* price $\mu > 0$ that must be chosen by the network, given by the previous equation. The allocation $d_r(\mathbf{w})$ to user r is thus

$$d_r(\mathbf{w}) = \begin{cases} \frac{w_r}{\sum_s w_s} C & \text{if } w_r > 0, \\ 0 & \text{if } w_r = 0, \end{cases} \quad (9)$$

and the payoff to user r (when $\mathbf{w} \neq \mathbf{0}$) is $U_r(d_r(\mathbf{w})) - w_r$.

The Kelly mechanism is a smooth market-clearing mechanism. To see this interpretation, note that when a user chooses a total payment w_r , it is as if the user has chosen a demand function $D(p, w_r) = w_r/p$ for $p > 0$. The resource manager then chooses a price μ so that $\sum_r D(\mu, w_r) = C$, i.e., so that the aggregate demand equals the supply C . For the specific form of demand functions we consider here, this leads to the expression for μ given in (8). User r then receives an allocation given by $D(\mu, w_r)$, and makes a payment $\mu D(\mu, w_r) = w_r$.

The following remarks summarize the key insights about the Kelly mechanism.

REMARK 3. Notice that the Kelly mechanism allocates the resource *in proportion to the bids of the users* (from (9)). We refer to mechanisms of this form as *proportional allocation mechanisms*; they will play a key role in our analysis in §4.

REMARK 4. When users act as price takers, then given a price $\mu > 0$, each user r chooses $w_r \geq 0$ to maximize $U_r(w_r/\mu) - w_r$. For the Kelly mechanism, *a competitive equilibrium always exists, and the resulting allocation is fully efficient* (Kelly 1997); this result is a special case of the first fundamental theorem of welfare economics (Mas-Colell et al. 1995, Chapter 16).

REMARK 5. When users act as price anticipators, they maximize their payoff with the knowledge that the price will be set according to (8). Johari and Tsitsiklis (2004) analyze this setting through the Nash equilibria of the resulting game; they show that for any utility system, *at the unique Nash equilibrium of this game, the aggregate utility at the resulting allocation achieves at least 75% of the maximum possible aggregate utility*.

In the next section, we consider a broader class of mechanisms that includes the Kelly mechanism as a special case, and use this setting to investigate the optimality of the Kelly mechanism.

4. Optimality of the Kelly Mechanism

In this section, we ask the following type of fundamental question: Is a particular mechanism “desirable” among a class of mechanisms satisfying certain “reasonable” properties? Defining desirability is the simpler of the two tasks: we consider a mechanism to be desirable if it minimizes worst-case efficiency loss when users are price anticipating.

Importantly, we ask for this efficiency property *independent* of the utility functions of the market participants, and their number. That is, the mechanisms we seek are those that perform well under broad assumptions on the nature of the preferences of market participants. We will show that under a specific set of assumptions, the Kelly mechanism in fact minimizes the worst-case efficiency loss when users are price anticipating.

We now frame the specific class \mathcal{D} of market mechanisms we consider in this section, defined as follows.

DEFINITION 4. The class \mathcal{D} consists of all functions $D(p, \theta)$ such that the following conditions are satisfied:

1. D defines a smooth market-clearing mechanism (cf. Definition 1).

2. For all $C > 0$, and for all $U_r \in \mathcal{U}$, a user’s payoff is *concave if he is price anticipating*; that is, for all R , and for all $\theta_{-r} \in (\mathbb{R}^+)^R$, the function

$$U_r(D(p_D(\boldsymbol{\theta}), \theta_r)) - p_D(\boldsymbol{\theta})D(p_D(\boldsymbol{\theta}), \theta_r)$$

is concave in $\theta_r > 0$ if $\theta_{-r} = \mathbf{0}$, and concave in $\theta_r \geq 0$ if $\theta_{-r} \neq \mathbf{0}$.

3. For all $p > 0$, and for all $d \geq 0$, there exists a $\theta \geq 0$ such that $D(p, \theta) = d$.

We pause here to briefly discuss the conditions in the previous definition. The second one allows us to characterize Nash equilibria in terms of only first-order conditions. To justify this condition, we note that some assumption of quasiconcavity is generally used to guarantee existence of pure-strategy Nash equilibria (Fudenberg and Tirole 1991). The third condition ensures that given a price p and desired allocation $d \in [0, C]$, each user can make a choice of θ to precisely guarantee the allocation d . This is an “expressiveness” condition on the mechanism, which ensures all possible demands can be chosen at any market-clearing price. The following example gives a family of mechanisms that lie in \mathcal{D} .

EXAMPLE 1. Suppose that $D(p, \theta) = \theta p^{-1/\alpha}$, where $\alpha \geq 1$. It is easy to check that this class of mechanisms satisfies $D \in \mathcal{D}$ for all choices of α ; when $\alpha = 1$, we recover the Kelly mechanism. The market-clearing condition yields that $p_D(\boldsymbol{\theta}) = (\sum_r \theta_r / C)^{1/\alpha}$. Note that as a result, the allocation to user r at a nonzero vector $\boldsymbol{\theta}$ is

$$D(p_D(\boldsymbol{\theta}), \theta_r) = \frac{\theta_r}{\sum_s \theta_s} C.$$

In other words, regardless of the value of α , the market-clearing allocations are chosen proportional to the bids. This remarkable fact is a special case of a more general result we establish below: *all mechanisms in \mathcal{D} are proportional allocation mechanisms*; they differ only in the market-clearing price that is chosen.

Our interest is in the worst-case ratio of aggregate utility at any Nash equilibrium to the optimal value of *SYSTEM*. Formally, for $D \in \mathcal{D}$ we define a constant $\rho(D)$ as follows:

$$\rho(D) = \inf \left\{ \frac{\sum_{r=1}^R U_r(D(p_D(\boldsymbol{\theta}), \theta_r))}{\sum_{r=1}^R U_r(d_r^S)} \mid C > 0, R > 1, \right. \\ \left. \begin{array}{l} \mathbf{U} \in \mathcal{U}^R, \mathbf{d}^S \text{ solves } \text{SYSTEM}(C, R, \mathbf{U}), \\ \text{and } \boldsymbol{\theta} \text{ is a Nash equilibrium} \end{array} \right\}.$$

Note that because all $U \in \mathcal{U}$ are strictly increasing and non-negative, the aggregate utility $\sum_{r=1}^R U_r(d_r^S)$ is positive for any utility system (C, R, \mathbf{U}) with $C > 0$, and any optimal solution \mathbf{d}^S to *SYSTEM*. However, Nash equilibria may not exist for some utility systems (C, R, \mathbf{U}) ; in this case, we set $\rho(D) = 0$.

Our main result in this section is the following theorem.

THEOREM 1. *Given $D \in \mathcal{D}$:*

1. *For any utility system (C, R, \mathbf{U}) , there exists a competitive equilibrium $\boldsymbol{\theta}$. Furthermore, for any such $\boldsymbol{\theta}$, the resulting allocation \mathbf{d} given by $d_r = D(p_D(\boldsymbol{\theta}), \theta_r)$ solves *SYSTEM* (C, R, \mathbf{U}) .*

2. *There exists a concave, strictly increasing, differentiable, and invertible function $B: (0, \infty) \rightarrow (0, \infty)$ such that for all $p > 0$ and $\theta \geq 0$,*

$$D(p, \theta) = \frac{\theta}{B(p)}.$$

3. *$\rho(D) \leq 3/4$, and this bound is met with equality if and only if $D(p, \theta) = \Delta\theta/p$ for some $\Delta > 0$.*

Before continuing to the proof, we make the following key remark.

REMARK 6. Parts 1 and 2 of the theorem provide a *characterization* of the types of mechanisms allowed by the constraints that define \mathcal{D} . In particular, notice that for nonzero $\boldsymbol{\theta}$, the market-clearing condition (4) yields

$$C = \sum_{r=1}^R D(p_D(\boldsymbol{\theta}), \theta_r) = \frac{\sum_{r=1}^R \theta_r}{B(p_D(\boldsymbol{\theta}))}, \quad (10)$$

which implies that

$$D(p_D(\boldsymbol{\theta}), \theta_r) = \frac{\theta_r}{\sum_{s=1}^R \theta_s} C; \quad (11)$$

in other words, every mechanism in \mathcal{D} is a proportional allocation mechanism. As a result, we conclude that for a given vector $\boldsymbol{\theta}$, when the market clears, mechanisms in \mathcal{D} differ from the Kelly mechanism only in the market-clearing price—the allocation is the same. Part 3 of the theorem then characterizes the efficiency loss of mechanisms of this form.

We emphasize that the theorem here is distinguished in part from related work because the allocation rule (11) was not assumed in advance. Rather, the result here starts from a set of simple assumptions on the structure of mechanisms to be considered (the definition of the class \mathcal{D}), and uses them to *prove* that any mechanism in the class must lead to the allocation in (11).

PROOF. Throughout the proof, we fix a particular mechanism $D \in \mathcal{D}$.

Step 1. For all $R > 1$ and $\boldsymbol{\theta}_{-r} \in (\mathbb{R}^+)^{R-1}$, the functions $D(p_D(\boldsymbol{\theta}), \theta_r)$ and $-p_D(\boldsymbol{\theta})D(p_D(\boldsymbol{\theta}), \theta_r)$ are concave in $\theta_r > 0$ if $\boldsymbol{\theta}_{-r} = \mathbf{0}$, and concave in $\theta_r \geq 0$ if $\boldsymbol{\theta}_{-r} \neq \mathbf{0}$. This conclusion follows by considering linear utility functions with a very large and a very small slope, respectively. Let $U_r(d_r) = \alpha d_r$, with $\alpha > 0$. If we apply Condition 2 of Definition 4 with $\alpha \rightarrow \infty$, then it follows that $D(p_D(\boldsymbol{\theta}), \theta_r)$ is concave in θ_r for nonzero $\boldsymbol{\theta}$. Similarly, if we let $\alpha \rightarrow 0$, it follows that $-p_D(\boldsymbol{\theta})D(p_D(\boldsymbol{\theta}), \theta_r)$ is concave in θ_r for nonzero $\boldsymbol{\theta}$.

Step 2. A user's payoff is concave if he is price taking. In other words, we will show that for all $U \in \mathcal{U}$ and for all $p > 0$, $U(D(p, \theta)) - pD(p, \theta)$ is concave in θ . The key idea is to use a limiting regime where both the number of users and the capacity grow large, so that users that are price anticipating effectively become price taking. The details are deferred to the electronic companion, which is available as part of the online version that can be found at <http://or.pubs.informs.org/>.

Step 3. There exists a positive function B such that $D(p, \theta) = \theta/B(p)$ for $p > 0$ and $\theta \geq 0$. By Step 2, a user's payoff is concave when he is price taking. By choosing $U(d) = \alpha d$ where $\alpha > p$ (respectively, $\alpha < p$), it follows that $D(p, \theta)$ must be concave (respectively, convex) in θ for a given $p > 0$. Thus, for fixed $p > 0$, $D(p, \theta)$ is an affine function of θ . Condition 3 in Definition 4 then implies that the constant term must be zero, while the coefficient of the linear term is positive; thus, $D(p, \theta) = \theta/B(p)$ for some positive function $B(p)$.

Before continuing, we note that the previous step already implies the fact that *any mechanism $D \in \mathcal{D}$ must be a proportional allocation mechanism.* This follows from Remark 6.

Step 4. For all utility systems (C, R, \mathbf{U}) , there exists a competitive equilibrium, and it is fully efficient. This step relies critically on Condition 3 in Definition 4: given a price μ , a user can first determine his optimal choice of quantity, and then choose a parameter θ to express this choice.

Formally, fix a utility system (C, R, \mathbf{U}) . Let \mathbf{d} be an optimal solution to *SYSTEM* (C, R, \mathbf{U}) , and let μ be the corresponding Lagrange multiplier, i.e., μ satisfies the following optimality conditions:

$$U'_r(d_r) = \mu \quad \text{if } d_r > 0, \quad (12)$$

$$U'_r(0) \leq \mu \quad \text{if } d_r = 0. \quad (13)$$

For each r , choose $\theta_r = B(\mu)d_r$. Clearly, $\mu = p_D(\boldsymbol{\theta})$; furthermore, the optimality conditions (12)–(13) ensure that (5) holds for all r . Thus, $\boldsymbol{\theta}$ is a competitive equilibrium.

Finally, we check that every competitive equilibrium yields an allocation that is an optimal solution to *SYSTEM*

(C, R, \mathbf{U}) . By Condition 3 in Definition 4, note that for any price $\mu > 0$, we have

$$\max_{d_r \geq 0} (U_r(d_r) - \mu d_r) = \max_{\theta_r \geq 0} (U_r(D(\mu, \theta_r)) - \mu D(\mu, \theta_r)).$$

Thus, if $\boldsymbol{\theta}$ is a competitive equilibrium, then the optimality conditions (12)–(13) must hold with $\mu = p_D(\boldsymbol{\theta})$ and $d_r = D(\mu, \theta_r)$. We conclude that \mathbf{d} is an optimal solution to $SYSTEM(C, R, \mathbf{U})$. This establishes the first claim of the theorem.

Step 5. B is an invertible, differentiable, strictly increasing, and concave function on $(0, \infty)$. The technical details are deferred to the electronic companion. In the sequel, we let $\Phi: (0, \infty) \rightarrow (0, \infty)$ be the (necessarily differentiable) inverse of B .

Step 6. Let (C, R, \mathbf{U}) be a utility system. A vector $\boldsymbol{\theta} \geq 0$ is a Nash equilibrium if and only if at least two components of $\boldsymbol{\theta}$ are nonzero, and there exists a nonzero vector $\mathbf{d} \geq 0$ and a scalar $\mu > 0$ such that $\theta_r = \mu d_r$ for all r , $\sum_{r=1}^R d_r = C$, and the following conditions hold:

$$U_r'(d_r) \left(1 - \frac{d_r}{C}\right) = \Phi(\mu) \left(1 - \frac{d_r}{C}\right) + \mu \Phi'(\mu) \left(\frac{d_r}{C}\right) \quad \text{if } d_r > 0, \quad (14)$$

$$U_r'(0) \leq \Phi(\mu) \quad \text{if } d_r = 0. \quad (15)$$

In this case $d_r = D(p_D(\boldsymbol{\theta}), \theta_r)$, $\mu = \sum_{r=1}^R \theta_r / C$, and $\Phi(\mu) = p_D(\boldsymbol{\theta})$. The proof is developed using the first-order necessary and sufficient conditions for a vector $\boldsymbol{\theta}$ to be a Nash equilibrium; details are deferred to the electronic companion.

Step 7. Let (C, R, \mathbf{U}) be a utility system. Then, there exists a unique Nash equilibrium. Our approach is to demonstrate existence of a Nash equilibrium by finding a solution $\mu > 0$ and $\mathbf{d} \geq 0$ to (14)–(15), such that $\sum_{r=1}^R d_r = C$. The technical details are deferred to the electronic companion.

Step 8. If the function Φ is not linear, there exists a utility system (C, R, \mathbf{U}) such that at the unique Nash equilibrium $\boldsymbol{\theta}$, the aggregate utility is strictly less than $3/4$ of the maximal aggregate utility. Consider a utility system with the following properties. Let $C = 1$. Fix $\mu > 0$, and let $U_1(d_1) = Ad_1$, where $A > \Phi(\mu)$. We will choose the remaining utilities so that we obtain a solution to the Nash conditions (14)–(15) with market-clearing price $\Phi(\mu)$.

We start by calculating d_1 by assuming it is nonzero, and applying (14):

$$d_1 = \frac{(A - \Phi(\mu))C}{A - \Phi(\mu) + \mu \Phi'(\mu)}. \quad (16)$$

We then choose users $2, \dots, R$ to have identical linear utility functions, with slopes less than A . As we will see, this will be possible if R is large enough. Formally, let $d = (C - d_1)/(R - 1)$, and (cf. (14)) define

$$\alpha = \frac{\Phi(\mu)C + (\mu \Phi'(\mu) - \Phi(\mu))d}{C - d}. \quad (17)$$

Let $U_r(d_r) = \alpha d_r$ for $r = 2, \dots, R$. Note that as R increases to infinity, d tends to zero and α converges to $\Phi(\mu)$, which is less than A . This guarantees that d_1 must be nonzero at any Nash equilibrium, which then implies that d_1 is indeed given by (16). In turn, letting $d_r = d$ for $r = 2, \dots, R$, this implies that (d_1, \dots, d_R) and μ are a valid solution to (14)–(15), when users have utility functions U_1, \dots, U_R .

Now consider the limiting ratio of Nash aggregate utility to maximal aggregate utility, as $R \rightarrow \infty$, in which case, $d \rightarrow 0$, and $\alpha \rightarrow \Phi(\mu)$. Regardless of R , a solution to $SYSTEM(C, R, \mathbf{U})$ is to allocate the entire resource to user 1, so the maximal aggregate utility is AC . Thus, the limiting ratio of Nash aggregate utility to maximal aggregate utility becomes

$$\begin{aligned} & \frac{Ad_1 + \alpha(C - d_1)}{AC} \\ &= \frac{(A - \Phi(\mu))}{A - \Phi(\mu) + \mu \Phi'(\mu)} \\ &+ \left(1 - \frac{(A - \Phi(\mu))}{A - \Phi(\mu) + \mu \Phi'(\mu)}\right) \left(\frac{\Phi(\mu)}{A}\right). \end{aligned} \quad (18)$$

We will show that for a suitable choice of A and μ , with $\mu > 0$ and $A > \Phi(\mu)$, this value is less than $3/4$.

Given that the nonnegative and convex function Φ has been assumed to be nonlinear, it is straightforward to check that there exists some $\mu > 0$ for which $\psi = \mu \Phi'(\mu) / \Phi(\mu) > 1$. At this value of μ , the right-hand side of (18) can be rewritten in the form

$$\frac{(1 - x)^2}{1 + (\psi - 1)x} + x, \quad (19)$$

where $x = \Phi(\mu) / A$. By choosing $x = 1/2$, i.e., $A = 2\Phi(\mu)$, we see that that the ratio is indeed strictly smaller than $3/4$, establishing the claim.

The proof of the theorem is completed by observing that when Φ is linear, we have $D(p, \theta) = \Delta\theta/p$ for some $\Delta > 0$, recovering the Kelly mechanism. Thus, the worst-case efficiency ratio is known to be exactly equal to $3/4$ (Johari and Tsitsiklis 2004). \square


REMARK 7. In fact, it is possible to prove something much stronger than the previous theorem: for any mechanism $D \in \mathcal{D}$, we can explicitly characterize the exact worst-case efficiency loss when users are price anticipating. We have shown that all mechanisms D satisfying the conditions of the theorem must be of the form $D(p, \theta) = \theta/B(p)$. Recall from the proof that $\Phi(\mu)$ is the differentiable inverse of B . If we define

$$\psi^* = \max_{\mu > 0} \left[\frac{\mu \Phi'(\mu)}{\Phi(\mu)} \right] = \left(\inf_{p > 0} \left[\frac{pB'(p)}{B(p)} \right] \right)^{-1}, \quad (20)$$

we can then show that $\rho(D) = G(\psi^*)$, where the function G is given by

$$G(\psi) = \begin{cases} \frac{3}{4} & \text{if } \psi = 1, \\ \frac{2\psi^2 - 3\psi\sqrt{\psi} + \sqrt{\psi}}{(\psi - 1)^2\sqrt{\psi}} & \text{if } \psi > 1. \end{cases} \quad (21)$$

Figure 1. The function $G(\psi)$ defined in (21). Note that $G(\psi)$ is strictly decreasing, with $G(1) = 3/4$.



Details are available in Chapter 5 of Johari (2004); the function G is plotted in Figure 1. The proof involves (i) showing that examples constructed as in Step 8 of the proof of Theorem 1 in fact yield the worst case efficiency loss when users are price anticipating, and (ii) picking the worst possible values of ψ and x . For any ψ , the value $G(\psi)$ is just the minimum of expression (19) over $x \in (0, 1)$. Furthermore, G is strictly decreasing for $\psi \geq 1$, which means that the worst possible value of expression (19) is obtained by employing the largest-possible value of ψ . Note that the quantity $pB'(p)/B(p)$ is the *elasticity* of $B(p)$ (Varian 1992), a common measure of the degree of “nonlinearity” of a function. Thus, $1/\psi^*$ is the minimal elasticity of $B(p)$ over all $p > 0$. To summarize, the worst case efficiency loss of such a mechanism is governed by the degree of nonlinearity of $B(p)$, as measured through the quantity ψ^* defined in (20).

REMARK 8. We note one potentially undesirable feature of the family of market-clearing mechanisms considered: the payoff to user r is defined as $-\infty$ when the composite strategy vector is $\theta = 0$ (cf. (7)). This definition is required because when the composite strategy vector is $\theta = 0$, a market-clearing price may not exist. One possible remedy is to restrict attention instead to mechanisms where $D(p, \theta) = 0$ if $\theta = 0$ for all $p \geq 0$; in this case, we can define $p_D(\theta) = 0$ if $\theta = 0$, and let the payoff to user r be $U_r(0)$ if $\theta_r = 0$. This condition amounts to a “normalization” on the market-clearing mechanism. Furthermore, this modification now exactly captures the mechanism of Johari and Tsitsiklis (2004), where $Q_r(0; \mathbf{w}_{-r}) = U_r(0)$ for all $\mathbf{w}_{-r} \geq 0$. It is straightforward to show that this modification does not alter the conclusion of Theorem 1. The argument involves a modification of the class of mechanisms \mathcal{D} , which includes the requirement $D(p, 0) = 0$.

REMARK 9. We note a key desirable feature of the mechanisms considered in Theorem 1, particularly in the context of communication networks. In general, even though the strategy space of the users is one-dimensional, the market-clearing price $p_D(\theta)$ may have a complex dependence on the vector θ . However, under the conditions of Theorem 1, the market-clearing price is only a function of $\sum_s \theta_s$, so that the market-clearing process does not require identification of individual users interacting with the mechanism, or even the number of users. This is a desirable scaling property for market mechanisms to be deployed in large-scale networks (Shenker 1990).

REMARK 10. As discussed in §2.1, Maheswaran and Basar (2004) also considered a class of mechanisms that allocate the resource according to the bids made by users. However, for every mechanism they consider, the payment made by a user depends *only* on the bid of that user, whereas in our setting this condition holds only for the Kelly mechanism. Thus, even though we have shown that we can restrict attention to proportional allocation mechanisms in parts 1 and 2 of the theorem, the class of proportional allocation mechanisms we consider is distinct from the class considered by Maheswaran and Basar, and thus our proof of optimality of the Kelly mechanism is distinct as well.

5. The Vickrey-Clarke-Groves (VCG) Approach

The mechanisms we considered in the last section had several restrictions placed on them; chief among these are that (1) users are restricted to using “simple” strategy spaces; and (2) the mechanism uses only a single price to clear the market. One could consider lifting both restrictions: allow more complex strategies, with users perhaps declaring their entire utility function to the market; and also, allow price discrimination so that each user is charged a personalized per-unit price for the resource. The best-known solution employing both these generalizations is the Vickrey-Clarke-Groves (VCG) approach to eliciting utility information (Vickrey 1961, Clarke 1971, Groves 1973). We review VCG mechanisms in §5.1.

In this section, we are interested in deciding whether the same outcome can be realized while preserving restriction (1) above, but removing restriction (2): that is, can mechanisms with “simple” strategy spaces that employ price discrimination achieve full efficiency? In §5.2, we present a class of mechanisms, inspired by the VCG class, in which users only submit scalar strategies; we show that these SSVCG mechanisms have desirable efficiency properties. In particular, we establish existence of an efficient Nash equilibrium, and under an additional condition, we also establish that all Nash equilibria are efficient. We extend this result to general convex environments in §5.3, and discuss an application to multicommodity flow routing.

5.1. VCG Mechanisms

In the VCG class of mechanisms, the basic approach is to let the strategy space of each user r be the set \mathcal{U} of possible utility functions, as defined in Assumption 1, and structure the payments made by each user so that the payoff of each user r has the same form as the objective function in the problem *SYSTEM*. Thus, in a VCG mechanism, each user is simply asked to *declare* their utility function; of course, if the payments are not structured properly, there is no guarantee that individuals will make truthful declarations. For each r , we use \tilde{U}_r to denote the declared utility function of user r , and use $\tilde{\mathbf{U}} = (\tilde{U}_1, \dots, \tilde{U}_R)$ to denote the vector of declared utilities.

Suppose that user r receives an allocation d_r , but has to make a payment t_r ; we use the notation t_r to distinguish from the bid w_r of §3. Then, the payoff to user r is

$$P_r(t_r, d_r) = U_r(d_r) - t_r.$$

On the other hand, the social objective (1) can be written as

$$U_r(d_r) + \sum_{s \neq r} U_s(d_s).$$

Comparing the preceding two expressions, the most natural means to align user objectives with the social planner's objectives is to define the payment t_r as the negation of the sum of the utilities of all users other than r .

Formally, given a vector of declared utility functions $\tilde{\mathbf{U}}$, a VCG mechanism chooses the allocation $\mathbf{d}(\tilde{\mathbf{U}})$ as an optimal solution to *SYSTEM* for the declared utility functions $\tilde{\mathbf{U}}$. For simplicity, let $\mathcal{X} = \{\mathbf{d} \geq 0: \sum_r d_r \leq C\}$; this is the feasible region for *SYSTEM*(C, R, \mathbf{U}). Then, for a VCG mechanism, we have

$$\mathbf{d}(\tilde{\mathbf{U}}) \in \arg \max_{\mathbf{d} \in \mathcal{X}} \sum_r \tilde{U}_r(d_r). \quad (22)$$

The payments are structured so that

$$t_r(\tilde{\mathbf{U}}) = - \sum_{s \neq r} \tilde{U}_s(d_s(\tilde{\mathbf{U}})) + h_r(\tilde{\mathbf{U}}_{-r}). \quad (23)$$

Here h_r is an arbitrary function of the declared utilities of users other than r ; because user r cannot affect this term through the choice of \tilde{U}_r , he chooses \tilde{U}_r to maximize

$$P_r(d_r(\tilde{\mathbf{U}}), t_r(\tilde{\mathbf{U}})) = U_r(d_r(\tilde{\mathbf{U}})) + \sum_{s \neq r} \tilde{U}_s(d_s(\tilde{\mathbf{U}})).$$

Now note that given $\tilde{\mathbf{U}}_{-r}$, the above expression is bounded above by

$$\max_{\mathbf{d} \in \mathcal{X}} \left[U_r(d_r) + \sum_{s \neq r} \tilde{U}_s(d_s) \right].$$

However, because $\mathbf{d}(\tilde{\mathbf{U}})$ satisfies (22), user r can achieve the preceding maximum by truthfully declaring $\tilde{U}_r = U_r$. Because this optimal strategy does not depend on the utility functions ($\tilde{U}_s, s \neq r$) declared by the other users, we recover the well-known fact that in a VCG mechanism, truthful declaration is a dominant strategy for user r . This discussion is summarized in the following well-known proposition; see, e.g., Green and Laffont (1979).

PROPOSITION 1. *Consider a VCG mechanism defined according to (22) and (23). Then, declaring $\tilde{U}_r = U_r$ is a dominant strategy for each user r . Furthermore, under these strategies, the resulting allocation is efficient.*

For our purposes, the interesting feature of the VCG mechanism is that it elicits the true utility functions from the users, and in turn (because of the definition of $\mathbf{d}(\tilde{\mathbf{U}})$) chooses an efficient allocation. However, it does so with a high degree of required communication: Each user submits an entire utility function, essentially an infinite-dimensional object. In the next section, we explore a class of mechanisms inspired by the VCG mechanisms, but with limited communication requirements.

5.2. Scalar Strategy VCG Mechanisms

We now consider a class of mechanisms where each user's strategy is a submitted utility function (as in VCG mechanisms), except that users are only allowed to choose from a given single-parameter family of utility functions. One cannot expect such mechanisms to have efficient dominant strategy equilibria, and we will focus instead on the efficiency properties of the resulting Nash equilibria.

Formally, SSVCG mechanisms allow users to choose from a given family of utility functions $\bar{U}(\cdot; \theta)$, parameterized by $\theta \in (0, \infty)$.⁶ We make the following assumptions about this family.

ASSUMPTION 2. 1. *For every $\theta > 0$, the function $\bar{U}(\cdot; \theta): d \mapsto \bar{U}(d; \theta)$, defined for $d \geq 0$, is strictly concave, strictly increasing, continuous, and continuously differentiable for $d > 0$.⁷*

2. *For every $\gamma \in (0, \infty)$ and $d \geq 0$, there exists a $\theta > 0$ such that $\bar{U}'(d; \theta) = \gamma$.⁸*

Given $\boldsymbol{\theta}$, the mechanism chooses $\mathbf{d}(\boldsymbol{\theta})$ such that

$$\mathbf{d}(\boldsymbol{\theta}) = \arg \max_{\mathbf{d} \in \mathcal{X}} \sum_r \bar{U}(d_r; \theta_r). \quad (24)$$

Because $\bar{U}(\cdot; \theta_r)$ is strictly concave for each r , the solution $\mathbf{d}(\boldsymbol{\theta})$ is uniquely defined. (Note the similarity between (22) and (24).)

By analogy with expression (23), the monetary payment by user r is

$$t_r(\boldsymbol{\theta}) = - \sum_{s \neq r} \bar{U}(d_s(\boldsymbol{\theta}); \theta_s) + h_r(\boldsymbol{\theta}_{-r}). \quad (25)$$

Here h_r is a function that depends only on the strategies $\boldsymbol{\theta}_{-r} = (\theta_s, s \neq r)$ submitted by the users other than r . Although we do not advocate any particular choice of h_r , a natural candidate is to define $h_r(\boldsymbol{\theta}_{-r}) = \sum_{s \neq r} \bar{U}(d_s(\boldsymbol{\theta}_{-r}); \theta_s)$, where $\mathbf{d}(\boldsymbol{\theta}_{-r})$ is the aggregate utility-maximizing allocation excluding user r . This leads to a natural scalar strategy analogue of the Clarke pivot mechanism (Clarke 1971). Note that, of course, the per-unit price paid by each user will typically be different.

Given h_r , the payoff to user r is

$$P_r(d_r(\boldsymbol{\theta}), t_r(\boldsymbol{\theta})) = U_r(d_r(\boldsymbol{\theta})) + \sum_{s \neq r} \bar{U}(d_s(\boldsymbol{\theta}); \theta_s) - h_r(\boldsymbol{\theta}_{-r}).$$

A strategy vector $\boldsymbol{\theta}$ is a *Nash equilibrium* if no user can profitably deviate through a unilateral deviation, i.e., if for all users r , there holds

$$P_r(d_r(\boldsymbol{\theta}), t_r(\boldsymbol{\theta})) \geq P_r(d_r(\theta'_r, \boldsymbol{\theta}_{-r}), t_r(\theta'_r, \boldsymbol{\theta}_{-r})) \quad \text{for all } \theta'_r > 0. \quad (26)$$

We start with the following key lemma, proven using an argument analogous to Proposition 1.

LEMMA 1. *The vector $\boldsymbol{\theta}$ is a Nash equilibrium of an SSVCG mechanism if and only if for all r ,*

$$\mathbf{d}(\boldsymbol{\theta}) \in \arg \max_{\mathbf{d} \in \mathcal{X}} \left[U_r(d_r) + \sum_{s \neq r} \bar{U}(d_s; \theta_s) \right]. \quad (27)$$

PROOF. Fix a user r . Because θ_r does not affect h_r , from (26) user r will choose θ_r to maximize the following effective payoff:

$$U_r(d_r(\boldsymbol{\theta})) + \sum_{s \neq r} \bar{U}(d_s(\boldsymbol{\theta}); \theta_s). \quad (28)$$

The optimal value of the objective function in (27) is certainly an upper bound to user r 's effective payoff (28). Thus, given a vector $\boldsymbol{\theta}$, if (27) is satisfied for all users r , then (26) holds for all users r , and we conclude $\boldsymbol{\theta}$ is a Nash equilibrium.

Conversely, given a vector $\boldsymbol{\theta}$, suppose that (27) is not satisfied for some user r . We will show that $\boldsymbol{\theta}$ cannot be a Nash equilibrium. Because \mathcal{X} is compact, an optimal solution exists to the problem in (27) for user r ; call this optimal solution \mathbf{d}^* . The vector \mathbf{d}^* must satisfy the first-order optimality conditions (analogous to (12)–(13)), which only involve the first derivatives $U'_r(d_r^*)$ and $(\bar{U}'(d_s^*; \theta_s), s \neq r)$. Suppose now that user r chooses $\theta'_r > 0$ such that $\bar{U}'(d_r^*; \theta'_r) = U'_r(d_r^*)$. Then, \mathbf{d}^* also satisfies the optimality conditions for problem (24). Because $\mathbf{d}(\theta'_r, \boldsymbol{\theta}_{-r})$ is the unique optimal solution to (24) when the strategy vector is $(\theta'_r, \boldsymbol{\theta}_{-r})$, we must have $\mathbf{d}(\theta'_r, \boldsymbol{\theta}_{-r}) = \mathbf{d}^*$. Thus, we have

$$\begin{aligned} P_r(d_r(\boldsymbol{\theta}), t_r(\boldsymbol{\theta})) &< U_r(d_r^*) + \sum_{s \neq r} \bar{U}(d_s^*; \theta_s) - h_r(\boldsymbol{\theta}_{-r}) \\ &= U_r(d_r(\theta'_r, \boldsymbol{\theta}_{-r})) + \sum_{s \neq r} \bar{U}(d_s(\theta'_r, \boldsymbol{\theta}_{-r}); \theta_s) \\ &\quad - h_r(\boldsymbol{\theta}_{-r}) \\ &= P_r(d_r(\theta'_r, \boldsymbol{\theta}_{-r}), t_r(\theta'_r, \boldsymbol{\theta}_{-r})). \end{aligned}$$

(The first inequality follows from the assumption that (27) is not satisfied for user r .) We conclude that (26) is violated for user r , so $\boldsymbol{\theta}$ is not a Nash equilibrium. \square

The following corollary states that there exists a Nash equilibrium that is efficient. Furthermore, at this efficient Nash equilibrium, all users truthfully reveal their utilities in a local sense: Each user r chooses θ_r so that the declared marginal utility $\bar{U}'(d_r(\boldsymbol{\theta}); \theta_r)$ is equal to the true marginal utility $U'_r(d_r(\boldsymbol{\theta}))$.

COROLLARY 1. *For any SSVCG mechanism, there exists an efficient Nash equilibrium $\boldsymbol{\theta}$ defined as follows: Let \mathbf{d}^S be an optimal solution to $SYSTEM(C, R, \mathbf{U})$. Each user r chooses θ_r so that $\bar{U}'(d_r^S; \theta_r) = U'_r(d_r^S)$. The resulting allocation satisfies $\mathbf{d}(\boldsymbol{\theta}) = \mathbf{d}^S$.*

PROOF. By Assumption 2, each user r can choose θ_r so that $\bar{U}'(d_r^S; \theta_r) = U'_r(d_r^S)$. For this vector $\boldsymbol{\theta}$, it is clear that $\mathbf{d}(\boldsymbol{\theta}) = \mathbf{d}^S$ because the optimal solution to (24) is uniquely determined, and the optimality conditions for (24) involve only the first derivatives $\bar{U}'(d_r(\boldsymbol{\theta}); \theta_r)$. By the same argument, it also follows that \mathbf{d}^S is an optimal solution in (27). Because $\mathbf{d}(\boldsymbol{\theta}) = \mathbf{d}^S$, we conclude that (27) is satisfied for all r , and thus $\boldsymbol{\theta}$ is a Nash equilibrium. \square

We note that, as in classical VCG mechanisms, there can be additional, possibly inefficient, Nash equilibria, as the following example shows.

EXAMPLE 2. Consider a system with R identical users with strictly concave utility function U . Suppose that user 1 chooses θ_1 so that $\bar{U}'(C; \theta_1) > U'(0)$, and every other user r chooses θ_r so that $\bar{U}'(0; \theta_r) < U'(C)$. We then have $\mathbf{d}(\boldsymbol{\theta}) = (1, \mathbf{0})$, so that the entire resource is allocated to user 1. Because $U'(C) \leq U'(0)$, it follows that (27) is satisfied for all users r , and we conclude $\boldsymbol{\theta}$ is a Nash equilibrium. However, this Nash equilibrium is inefficient: The unique optimal solution to $SYSTEM(C, R, \mathbf{U})$ is symmetric, and allocates C/R units of the resource to each of the R users.

The equilibrium in the preceding example involves a “bluff.” User 1 declares such a high marginal utility at C that all other users concede. One way to preclude such equilibria is to enforce an assumption that guarantees participation. The next proposition assumes that at least two users have infinite marginal utility at zero allocation; this guarantees that all Nash equilibria are efficient.

PROPOSITION 2. *Suppose that Assumptions 1 and 2 hold, except that $U'_r(0) = \infty$ for at least two users r . Suppose that $\boldsymbol{\theta}$ is a Nash equilibrium. Then, $\mathbf{d}(\boldsymbol{\theta})$ is an optimal solution to $SYSTEM(C, R, \mathbf{U})$.*

PROOF. Note that Lemma 1 applies even if we allow $U'_r(0)$ to be infinite. Suppose, without loss of generality, that $U'_1(0) = U'_2(0) = \infty$. Let $\mathbf{d} = \mathbf{d}(\boldsymbol{\theta})$. We observe that users 1 and 2 must have positive allocations at equilibrium, from (27). The optimality conditions for (27) thus imply that for $i = 1, 2$, and for each $s \neq i$, we have

$$U'_i(d_i) = \bar{U}'(d_s; \theta_s) \quad \text{if } d_s > 0, \quad (29)$$

$$\geq \bar{U}'(d_s; \theta_s) \quad \text{if } d_s = 0. \quad (30)$$

In particular, $U'_1(d_1) = \bar{U}'(d_2; \theta_2)$ and $U'_2(d_2) = \bar{U}'(d_1; \theta_1)$. Furthermore, the optimality conditions for problem (24) yield $\bar{U}'(d_1; \theta_1) = \bar{U}'(d_2; \theta_2)$. From these relations, we first conclude that $U'_1(d_1) = \bar{U}'(d_1; \theta_1) = \mu$ for some scalar μ . Further, if we consider the optimality conditions for (27)

for an arbitrary user r , we obtain $U'_r(d_r) = \bar{U}'(d_1; \theta_1) = U'_1(d_1) = \mu$ if $d_r > 0$, and $U'_r(d_r) \leq \bar{U}'(d_1; \theta_1) = U'_1(d_1) = \mu$ if $d_r = 0$. However, these are exactly the sufficient conditions for optimality for \mathbf{d} for the problem $SYSTEM(C, R, U)$, with Lagrange multiplier μ . \square

Intuitively, for efficiency to hold, we need to have at least two actively “competing” users. In the previous result, this is guaranteed because the two users with infinite marginal utility at zero allocation will want a strictly positive rate in any equilibrium.

The results of this section demonstrate that by relaxing the assumption that the resource allocation mechanism must set a single price, we can in fact significantly improve upon the efficiency guarantee of Theorem 1. We note that Maheswaran and Basar (2004) (as well as the earlier work, Maheswaran 2003) and Yang and Hajek (2006a) have (independently of each other) presented a class of mechanisms for this resource allocation problem that operate as follows. Each user r chooses a “bid” $\theta_r \geq 0$. Define $T = \sum_s \theta_s$ and $T_{-r} = \sum_{s \neq r} \theta_s$. Their mechanism allocates in proportion to the bids, so $d_r^{MY}(\boldsymbol{\theta}) = \theta_r C / T$. Each user r makes a payment defined by

$$t_r^{MY}(\boldsymbol{\theta}) = \theta_s (\varphi(T) - \varphi(T_{-r})),$$

where φ is a strictly increasing, continuously differentiable function, such that, in addition, the function $u \mapsto u^2 \varphi'(u) / C$ is a strictly increasing, onto function from $[0, \infty)$ to $[0, \infty)$. Under these assumptions, define $\bar{U}(x; \theta) = -\theta \varphi(\theta/x)$. Then, it can be verified that: (1) \bar{U} satisfies Assumption 2; and (2) the corresponding SSVCG mechanism is the Maheswaran-Basar/Yang-Hajek mechanism, so that $\mathbf{d}(\boldsymbol{\theta}) = \mathbf{d}^{MY}(\boldsymbol{\theta})$ and $\mathbf{t}(\boldsymbol{\theta}) = \mathbf{t}^{MY}(\boldsymbol{\theta})$.⁹ Yang and Hajek have shown that as long as at least two users r have $U'_r(0) = \infty$, then all Nash equilibria are efficient; this matches the result of Proposition 2. A similar result is derived by Maheswaran and Basar for allocation of a single infinitely divisible resource, by assuming a single nonstrategic “virtual” user who always bids ϵ for the resource; Nash equilibria of the resulting mechanism become efficient as ϵ approaches zero. In conclusion, our class of SSVCG mechanisms includes the mechanism of Maheswaran (2003), Maheswaran and Basar (2004), and Yang and Hajek (2006a) as a special case.

5.3. General Convex Environments

Despite the fact that we have formulated SSVCG mechanisms for the specific choice of \mathcal{X} corresponding to allocation of a single resource, our results hold for far more general convex environments. Formally, we continue to assume that each user r is to receive an allocated quantity d_r . However, we now assume that the feasible resource allocations are determined by constraints involving the vector \mathbf{d} , as well as some auxiliary variables $\mathbf{y} = (y_{R+1}, \dots, y_{R+M})$. As an example, \mathbf{y} may be the path flows

in a multicommodity flow model, in which case d_r is the total rate allocated to user r . We make the following assumption about \mathcal{X} .

ASSUMPTION 3. *The set \mathcal{X} of feasible allocations \mathbf{d} is of the form*

$$\mathcal{X} = \{\mathbf{d} \geq \mathbf{0}: \text{there exists } \mathbf{y} \in \mathbb{R}^M, \mathbf{y} \geq \mathbf{0}, \\ \text{such that } g_j(\mathbf{d}, \mathbf{y}) \leq 0, j = 1, \dots, J\},$$

for some given functions g_j .

- Further: 1. *The set \mathcal{X} is compact;*
2. *For each j , the function g_j is convex and differentiable on an open set containing $(\mathbb{R}^+)^{R+M}$;*
3. *There exists a pair $(\bar{\mathbf{d}}, \bar{\mathbf{y}}) \in \mathbb{R}^{R+M}$ such that $g_j(\bar{\mathbf{d}}, \bar{\mathbf{y}}) < 0$ for all j .*

In the preceding assumption, M denotes the dimension of the space of auxiliary variables, and J the number of convex inequality constraints. Although we assume differentiability in condition (2), this is not crucial; we assume differentiability for technical simplicity.

REMARK 11. Condition (3) in the preceding assumption ensures that the Slater constraint qualification holds for all the optimization problems we consider; see, e.g., Bertsekas et al. (2003). We note that all our results apply to the more general setting where in addition to the convex inequality constraints, the set \mathcal{X} is constrained by *affine equality constraints* of the form $h_k(\mathbf{d}, \mathbf{y}) = 0$, $k = 1, \dots, K$. In this case, condition (3) in Assumption 3 is modified to require that there exists a pair $(\bar{\mathbf{d}}, \bar{\mathbf{y}})$ in the *relative interior* of the set $\{(\mathbf{d}, \mathbf{y}): h_k(\mathbf{d}, \mathbf{y}) = 0\}$, such that $g_j(\bar{\mathbf{d}}, \bar{\mathbf{y}}) < 0$ for all j ; with this modification the Slater constraint qualification once again applies. All our results continue to hold in this setting, if we simply represent the k th affine equality constraint by the two equivalent inequality constraints $h_k(\mathbf{d}, \mathbf{y}) \leq 0$ and $h_k(\mathbf{d}, \mathbf{y}) \geq 0$.

We define an efficient allocation as the solution to the following generalized definition of $SYSTEM$:

$SYSTEM$:

$$\text{maximize } \sum_{r=1}^R U_r(d_r) \quad (31)$$

$$\text{subject to } \mathbf{d} \in \mathcal{X}. \quad (32)$$

It is straightforward to check that under Assumption 3, nearly all the basic results established above continue to hold; in particular, Proposition 1, Lemma 1, and Corollary 1 continue to hold, with proofs nearly identical to their existing versions. Thus, (27) continues to provide a necessary and sufficient condition for the characterization of Nash equilibria, and from Corollary 1 we know that *there always exists at least one efficient Nash equilibrium*.

It is natural to search for an analog of Proposition 2. In particular, we would like a more general sufficient condition that guarantees that *all* Nash equilibria are efficient.

We now provide such a guarantee, under additional conditions on the structure of the constraints.

We first introduce some notation. Given a vector θ of user strategies, we denote by $\mathbf{z} = (\mathbf{d}, \mathbf{y})$ the optimal solution $\mathbf{d} = \mathbf{d}(\theta)$ to problem (24), together with optimal values of the auxiliary variables \mathbf{y} . (Recall that \mathbf{d} , the vector of resulting allocations, is uniquely determined by strict concavity of the objective function in (24).) Given such a solution, we let

$$\mathcal{P} = \{k \in \{1, \dots, R + M\} : z_k > 0\} \quad (33)$$

be the indices associated with variables whose value is nonzero, and we let

$$\mathcal{J} = \{j : g_j(\mathbf{z}) = 0\} \quad (34)$$

be the set of constraint indices that are binding. For each k , we define the vector $\mathbf{D}(k)$ as

$$\mathbf{D}(k) = \left(\frac{\partial g_j}{\partial z_k}(\mathbf{z}), j \in \mathcal{J} \right). \quad (35)$$

For example, if the constraints are of the form $\mathbf{A}\mathbf{z} - \mathbf{b} \leq \mathbf{0}$, for some matrix \mathbf{A} and some vector \mathbf{b} , and if all constraints are binding, then $\mathbf{D}(k)$ is the k th column of \mathbf{A} .

PROPOSITION 3. *Suppose that Assumptions 1–3 hold, except that $U_r'(0)$ is allowed to be infinite. Suppose that θ is a Nash equilibrium, $\mathbf{d} = \mathbf{d}(\theta)$ is the optimal solution to problem (24), and \mathbf{y} is an associated optimal vector of auxiliary variables. Define \mathcal{P} , \mathcal{J} , and $\mathbf{D}(k)$ according to (33), (34), and (35), respectively. Suppose furthermore that for every user $r \in \{1, \dots, R\}$, the vector $\mathbf{D}(r)$ is in the span of the vectors $\mathbf{D}(s)$, $s \neq r$, $s \in \mathcal{P}$. Then, \mathbf{d} is an efficient allocation.*

PROOF. Let θ be a Nash equilibrium, and let $\mathbf{z} = (\mathbf{d}, \mathbf{y})$ correspond to an optimal solution to problem (24). Because θ is a Nash equilibrium, \mathbf{z} is also an optimal solution to problem (27) for each $r \in \{1, \dots, R\}$.

To simplify the argument, we first assume that $\mathcal{J} = \{1, \dots, J\}$, so that all constraints are binding. Let $\lambda \geq 0$ and $\mu(r) \geq 0$ be vectors of Lagrange multipliers associated with the optimal solution \mathbf{z} for problem (24), and problem (27) for user r , respectively. Such Lagrange multipliers are guaranteed to exist by Assumption 3 because the Slater constraint qualification is satisfied; see, e.g., Bertsekas et al. (2003, §5.4). To simplify the presentation, we define $U_k(d) = 0$ for $k > R$ and all $d \geq 0$. Then, the optimality conditions for problem (24) are

$$\bar{U}'(z_k; \theta_k) = \lambda^\top \mathbf{D}(k) \quad \text{if } z_k > 0, k \in \{1, \dots, R\}, \quad (36)$$

$$\bar{U}'(z_k; \theta_k) \leq \lambda^\top \mathbf{D}(k) \quad \text{if } z_k = 0, k \in \{1, \dots, R\}, \quad (37)$$

$$U'_k(z_k) = \lambda^\top \mathbf{D}(k) \quad \text{if } z_k > 0, k \in \{R + 1, \dots, R + M\}, \quad (38)$$

$$U'_k(z_k) \leq \lambda^\top \mathbf{D}(k) \quad \text{if } z_k = 0, k \in \{R + 1, \dots, R + M\}. \quad (39)$$

Similarly, the optimality conditions for problem (27) for user r , $r \in \{1, \dots, R\}$, are

$$\bar{U}'(z_k; \theta_k) = \mu(r)^\top \mathbf{D}(k) \quad \text{if } z_k > 0, k \in \{1, \dots, R\}, k \neq r, \quad (40)$$

$$\bar{U}'(z_k; \theta_k) \leq \mu(r)^\top \mathbf{D}(k) \quad \text{if } z_k = 0, k \in \{1, \dots, R\}, k \neq r, \quad (41)$$

$$U'_k(z_k) = \mu(r)^\top \mathbf{D}(k) \quad \text{if } z_k > 0, k \in \{r, R + 1, \dots, R + M\}, \quad (42)$$

$$U'_k(z_k) \leq \mu(r)^\top \mathbf{D}(k) \quad \text{if } z_k = 0, k \in \{r, R + 1, \dots, R + M\}. \quad (43)$$

From (36), (38), (40), and (42), we deduce that for all $r \in \{1, \dots, R\}$ and $k \in \mathcal{P}$ such that $k \neq r$, we have

$$\mu(r)^\top \mathbf{D}(k) = \lambda^\top \mathbf{D}(k). \quad (44)$$

The preceding equality holds because z_k is positive for all $k \in \mathcal{P}$.

Fix a user $r \in \{1, \dots, R\}$. By assumption, there exists a vector of coefficients $(\alpha_{rk}, k \in \mathcal{P}, k \neq r)$ such that

$$\mathbf{D}(r) = \sum_{\{k \in \mathcal{P} : k \neq r\}} \alpha_{rk} \mathbf{D}(k). \quad (45)$$

Suppose first that $z_r = 0$. We have

$$\begin{aligned} U'_r(z_r) &\leq \mu(r)^\top \mathbf{D}(r) = \sum_{\{k \in \mathcal{P} : k \neq r\}} \alpha_{rk} \mu(r)^\top \mathbf{D}(k) \\ &= \sum_{\{k \in \mathcal{P} : k \neq r\}} \alpha_{rk} \lambda^\top \mathbf{D}(k) = \lambda^\top \mathbf{D}(r). \end{aligned}$$

The inequality follows from the optimality condition (43). The first equality follows from (45). The second equality follows from (44). The third equality uses (45) again. If instead $z_r > 0$, then by using (42), we see that the inequality above holds with equality. Combining this result with (38)–(39), we conclude that

$$U'_k(z_k) = \lambda^\top \mathbf{D}(k) \quad \text{if } z_k > 0, k \in \{1, \dots, R + M\},$$

$$U'_k(z_k) \leq \lambda^\top \mathbf{D}(k) \quad \text{if } z_k = 0, k \in \{1, \dots, R + M\}.$$

These are precisely the optimality conditions for the problem *SYSTEM* (with Lagrange multiplier vector λ), and therefore \mathbf{d} is an efficient allocation.

If $\mathcal{J} \neq \{1, \dots, J\}$, so that some constraints are not binding, the preceding argument continues to hold: We need only consider the constraints in \mathcal{J} because the Lagrange multipliers associated with nonbinding constraints are zero. \square

REMARK 12. For the case where the constraints defining \mathcal{R} are linear, of the form $\mathbf{A}\mathbf{z} - \mathbf{b} \leq \mathbf{0}$, the linear dependence condition of Proposition 3 is equivalent to the following: Given a Nash equilibrium and $r \in \{1, \dots, M\}$, if we perturb the r th component of \mathbf{z} , then it is always possible to feasibly modify the components z_k , $k \neq r$, $k \in \mathcal{P}$, so that those constraints that are binding (i.e., the constraints in \mathcal{J}) remain binding.

To gain some intuition for Proposition 3, suppose that we have no auxiliary variables, that the set \mathcal{X} is a polyhedron, and that \mathcal{X} has an extreme point with positive components. At that extreme point, and in the absence of degeneracy, there are exactly R constraints that are binding. In this case, the vectors $\mathbf{D}(k)$ are linearly independent, and therefore the linear dependence condition of Proposition 3 fails to hold. Intuitively, for the condition to hold, we need to have a number of actively “competing” users (i.e., users with positive d_r) which is greater than the effective number of resources (i.e., the number of binding constraints). (Notice that the violation of this condition is behind the failure observed in Example 2.)

The linear dependence condition is satisfied in a broad class of problems. The following presents an important example: multicommodity flow routing.

EXAMPLE 3 (MULTICOMMODITY FLOW). In this model, R users send traffic through a network of L links. Each link l has a given finite positive capacity, denoted $C_l > 0$. Each user r has available a collection $P(r)$ of paths on which she can send traffic; each path $p \in P(r)$ is identified with the set of links it goes through, so that $p \subset \{1, \dots, L\}$. Let $P = \bigcup_r P(r)$ and let $M = |P|$. By duplicating paths if necessary, we assume without loss of generality that each path $p \in P$ is associated with exactly one user. The set \mathcal{X} is then defined as follows:

$$\mathcal{X} = \left\{ \mathbf{d} \geq 0: \text{there exists } \mathbf{y} \geq 0 \text{ with} \right.$$

$$\left. \begin{aligned} \sum_{p \in P(r)} y_p = d_r, \quad r = 1, \dots, N; \quad \text{and} \\ \sum_{p \in P: l \in p} y_p \leq C_l, \quad l = 1, \dots, L \end{aligned} \right\}. \quad (46)$$

$$(47)$$

As discussed in Remark 11, \mathcal{X} satisfies a variation of Assumption 3, and therefore Proposition 3 still applies.

We now make the following assumptions. First, we assume that for all users r , we have $U'_r(0) = \infty$. Second, we assume that for each link l , there are two users $r_1(l)$ and $r_2(l)$ for which the path $\{l\}$, consisting of the single link l , is the only candidate path. The latter assumption guarantees sufficient competition at every link; informally, it requires that competition for resources from flows travelling a short distance is nonnegligible. (A similar assumption was used by Yang and Hajek 2007.) Under these assumptions, we now show that Proposition 3 applies, and thus establish that any resulting Nash equilibrium is efficient.

We only need to verify that the linear dependence condition in Proposition 3 is satisfied at a Nash equilibrium. To check this, let $\boldsymbol{\theta}$ be a Nash equilibrium, and let $\mathbf{z} = (\mathbf{d}, \mathbf{y})$ be a corresponding optimal solution to problem (24), so $\mathbf{d} = \mathbf{d}(\boldsymbol{\theta})$. We will use the equivalent statement of the linear dependence condition provided in Remark 12.

First note that because $U'_r(0) = \infty$, at a Nash equilibrium every variable d_r is nonzero. Furthermore, for those users

that have a single path p , the corresponding path variable y_p is also nonzero.

Let us consider perturbing a variable d_r , say by adding δ to it. Let $p \in P(r)$ be such that $y_p > 0$. Then, we can add δ to y_p , so that the binding constraint (46) for that particular user r remains binding. In addition, for every link l on the path p , there exists a user $r_i(l) \neq r$ that has a single path that only uses this link. By decreasing $d_{r_i(l)}$ and the associated component of \mathbf{y} by δ , the binding constraint (46) for $r_i(l)$ remains binding. Furthermore, if constraint (47) was binding, it remains binding. We conclude that the condition of Remark 12 holds, as required, and thus the result of Proposition 3 applies.

We conclude this section by noting several extensions. First, it is possible to construct a similar class of mechanisms when each user r has a utility function $U_r(\mathbf{x}_r)$ that depends on a vector of resources \mathbf{x}_r . In this case, user r needs a strategy space of the same dimension as the vector \mathbf{x}_r . The structure of the mechanism is then analogous to the one described in §5.2.

As for standard VCG mechanisms, the mechanisms we have proposed are certainly not *budget balanced*; that is, $\sum_r t_r(\boldsymbol{\theta}) \neq 0$ in general. However, one can consider a variety of forms for the functions h_r to mitigate this effect; a choice similar to the “pivot” mechanisms of Clarke (1971) results if $h_r(\boldsymbol{\theta}_{-r}) = -\sum_{s \neq r} \bar{U}(x_s(\boldsymbol{\theta}_{-r}); \theta_s)$, where $x_s(\boldsymbol{\theta}_{-r})$ is the optimal resource allocation if user r is removed from the system. In this mechanism all payments t_r are negative, and a user r makes a payment if and only if his presence affects the optimal resource allocation.

6. Discussion and Conclusions

This paper studies the efficiency achievable by scalar-parameterized mechanisms, particularly in allocation of a single infinitely divisible resource. In §4, we show our first main result: when a single price is used for all users, then within a class of reasonable market-clearing mechanisms the Kelly mechanism uniquely minimizes the worst case efficiency loss when users are price anticipating. An intermediate result also shows that allocating the resource in proportion to users’ bids (as in the Kelly mechanism) is a natural consequence of simple assumptions on the mechanism. In §5, we show that relaxing the single-price assumption actually allows us to significantly improve the performance of the mechanism: in our second main result, we demonstrate a class of mechanisms that exhibit efficient Nash equilibria. We also argue that these SSVCG mechanisms can be extended to much more general resource allocation environments, and provide conditions under which *all* Nash equilibria are efficient. We then successfully apply SSVCG mechanisms to a standard multicommodity flow setting.

Compared to the two major results described above, there is another possibility: we might relax the assumption of one-dimensional bids, while preserving the single-price restriction on the market mechanism. Consider,

for example, mechanisms where users submit *arbitrary* demand functions (giving demand as a function of price), and a single price is chosen to clear the market. Using an analog of the argument in Klemperer and Meyer (1989), it is straightforward to show that for the model we are considering, there exist games with arbitrarily inefficient Nash equilibria. The argument involves showing that essentially any feasible outcome (consisting of a price and allocation vector) can be sustained at a Nash equilibrium.

Our results, as well as those of earlier work such as Maheswaran (2003), Maheswaran and Basar (2004), and Yang and Hajek (2006a, 2007), suggest that there is an interesting interplay between the *strategic flexibility* granted to users, and the *pricing flexibility* granted to the mechanism. When users are allowed great flexibility in submitting their bids, then full efficiency is possible with price differentiation (the VCG approach, where users' bids consist of entire utility functions), whereas arbitrarily high inefficiency results if only a single price is allowed (as in the preceding paragraph). On the other hand, when users are allowed to use only one-dimensional signals, then with a single price the most efficient mechanism (i.e., the Kelly mechanism) suffers a nonzero efficiency loss, whereas with price differentiation full efficiency can again be achieved (as established for SSVCG mechanisms and their precursors in Maheswaran and Basar 2004 and Yang and Hajek 2007). Note, in particular, the critical insight from this stream of work that *when pricing flexibility is limited, restricting strategic flexibility may actually increase efficiency*. These results are summarized in Figure 2.

We briefly note several extensions and open directions. First, §4 considers only a single resource. For general network topologies, we can consider a game where users

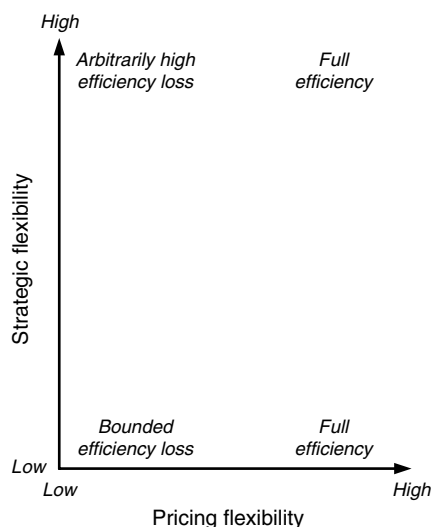
submit individual bids to each link in the network, similar to the approach taken by Johari and Tsitsiklis (2004). In this case, if each link is endowed with a smooth market-clearing mechanism satisfying the assumptions of §4, then the worst case efficiency loss will be minimized if each link uses the Kelly mechanism.

A second extension of the work in §4 applies to markets with supply function bidding. Motivated by current problems in market design for electric power systems, we consider a model where multiple producers compete to satisfy an inelastic demand. Demand for electricity, particularly in the short run, is characterized by low elasticity with respect to price, i.e., changes in price do not lead to significant changes in the level of demand; see, e.g., Stoft (2002, §1-7.3). A basic model for electricity market operation involves *supply function bidding* (Klemperer and Meyer 1989): Each generator submits a supply function expressing their willingness to produce electricity as a function of the market-clearing price. A single price is then chosen to ensure that supply matches the inelastic demand. In Johari (2004), we consider *restrictions* on the supply functions that can be chosen by the generators, and aim to design these restrictions so that nearly efficient allocations are achieved even if firms are price anticipating. We prove a characterization theorem, similar to Theorem 1, describing the best possible efficiency guarantees within a class of mechanisms in which the generators are restricted to submitting a supply function chosen from within a restricted, one-parameter family. Details of this work can be found in Johari (2004).

All the work in this paper is concerned with *quasilinear environments*, i.e., users' utilities are measured in monetary units. Further, our models assume that users have unlimited budgets. An interesting set of parallel issues arises if one relaxes these assumptions. Indeed, we note that the Kelly mechanism is quite similar to the celebrated Shapley-Shubik "trading post" mechanism (Shapley and Shubik 1977). In the Shapley-Shubik mechanism, users maximize utility (which may not be quasi-linear), and users are subject to a budget constraint. These games thus exhibit very different equilibria from the Kelly mechanism. An interesting open direction concerns evaluating the efficiency properties of the Shapley-Shubik mechanism and determining the extent to which the Shapley-Shubik mechanism is an "optimal" design. Zhang and others have considered notions of optimality such as envy-freeness in the setting where users have budget constraints; see Zhang (2005) and Feldman et al. (2005) for details.

We note that all our work considers efficiency from a *static* standpoint. This is natural as a starting point, and our positive results give proof-of-concept of potentially viable mechanism designs. However, a complete approach also requires attention to the dynamic setting. Several works have recently considered dynamics for convergence to efficient equilibrium. Yang and Hajek (2007) have shown that a form of myopic best-response dynamics can converge to an

Figure 2. Strategic flexibility vs. pricing flexibility.



Notes. The horizontal (respectively, vertical) axis increases with pricing flexibility (respectively, strategic flexibility), and the annotations reflect a summary of current understanding. This paper discusses efficiency when mechanisms have restricted strategic flexibility (i.e., one-dimensional bids).

efficient Nash equilibrium of the SSVCG mechanism in the multicommodity flow setting, whereas the ascending auction mechanism of Ausubel (2006) provides a mechanism with limited communication requirements and at least one efficient subgame-perfect equilibrium. The former analyzes a stylized model of dynamics assuming a certain model of user behavior; the latter proposes a dynamic mechanism design assuming fully rational user behavior. Clearly, a significant open direction is to develop a unified treatment of the efficiency of resource allocation mechanisms with restricted communication requirements in a *dynamic* setting.

Ultimately, however, the most intriguing questions raised by this work are captured by Figure 2. The results of this paper and earlier work in the area suggest that in designing efficient markets, strategic flexibility can be increased only if pricing flexibility is available. However, we have only studied this issue in the extremes depicted in Figure 2. In particular, we have not quantitatively characterized the *trade-offs* between strategic flexibility and pricing flexibility. For example, we have studied mechanisms with a single price per resource, and with perfect price discrimination. A natural question arises: What are the gains in efficiency if a two-tiered pricing system is used? How many prices are needed for a bounded efficiency loss guarantee if users have strategy spaces parameterized by a fixed number of dimensions? The answers to such questions remain interesting directions for future work.

7. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://or.journal.informs.org/>.

Endnotes

1. A similar result was independently and simultaneously presented by Yang and Hajek (2006b, 2007); see §2.2.
2. Some notion of smoothness is required to ensure that dimension is a reasonable measure of strategic flexibility; for example, in the absence of a differentiability requirement, an arbitrary amount of information can be communicated in a single scalar, e.g., using a space-filling curve.
3. We note that the authors' assumption that c is convex is largely to ensure existence of a Nash equilibrium; in our context, a similar role is played by the second condition in Definition 4.
4. We follow the terminology used in the communications literature. The economics literature frequently uses "valuation" and "utility" instead of our "utility" and "payoff," respectively.
5. Note that we suppress the dependence of this solution on C ; where necessary we will emphasize this dependence.
6. Note that, by contrast with §4, the choice of bid θ by a user indexes a utility function, rather than a demand function. However, this is not crucial: if a user with utility function U maximizes $U(d) - pd$ (i.e., the user acts as a price

taker), the solution yields the demand function $D(p) = (U')^{-1}(p)$. Up to an additive constant, the utility function and demand function can be recovered from each other. Thus, we could define SSVCG mechanisms where users submit demand functions from a parameterized class. We define our SSVCG mechanisms in terms of utility functions to maintain consistency with the standard definition of VCG mechanisms in the literature.

7. This is almost the same as Assumption 1, except that we require strict concavity, do not require nonnegativity, and allow $U'_r(0)$ to be infinite.

8. Because we do not assume differentiability with respect to θ , the only differentiation of \bar{U} is with respect to the first coordinate d , and $\bar{U}'(d; \theta)$ will always stand for the derivative with respect to d .

9. A small technical issue arises here because $\bar{U}'(0, \theta) = \infty$ for all $\theta > 0$, which violates the last part of Assumption 2; however, the development in this section can be adapted to cover this case.

Acknowledgments

The authors are grateful for helpful conversations with John Ledyard, Tim Roughgarden, Ilya Segal, Eva Tardos, and Robert Wilson. This research was supported by a National Science Foundation Graduate Research Fellowship, by the Defense Advanced Research Projects Agency under the Next Generation Internet Initiative, by the Army Research Office under grant DAAD10-00-1-0466, and by the National Science Foundation under grant 0428868. Portions of §4 are based on work in the Ph.D. thesis of Johari (2004). The results in §5 were first presented at the 2005 Allerton Conference on Communication, Control, and Computing, Monticello, Illinois (Johari and Tsitsiklis 2005).

References

- Anshelevich, E., A. Dasgupta, É. Tardos, T. Wexler. 2003. Near-optimal network design with selfish agents. *Proc. ACM Sympos. Theory Comput.* ACM, New York, 511–520.
- Ausubel, L. M. 2006. An efficient dynamic auction for heterogeneous commodities. *Amer. Econom. Rev.* **96**(3) 602–629.
- Bertsekas, D. P., A. Nedic, A. E. Ozdaglar. 2003. *Convex Analysis and Optimization*. Athena Scientific, Cambridge, MA.
- Blumrosen, L., N. Nisan, I. Segal. 2007. Auctions with severely bounded communication. *J. Artificial Intelligence Res.* **28** 233–266.
- Chen, Y.-J., J. Zhang. 2008. Design of price mechanisms for network resource allocation via price of anarchy. Submitted.
- Clarke, E. H. 1971. Multipart pricing of public goods. *Public Choice* **11** 19–33.
- Correa, J. R., A. S. Schulz, N. S. Moses. 2004. Selfish routing in capacitated networks. *Math. Oper. Res.* **29**(4) 961–976.
- Czumaj, A., B. Voecking. 2002. Tight bounds for worst-case equilibria. *Proc. ACM-SIAM Sympos. Discrete Algorithms*. ACM, New York, 413–420.
- Dimakis, A., R. Jain, J. Walrand. 2006. Mechanisms for efficient allocation in divisible capacity networks. *Proc. IEEE Conf. Decision and Control (CDC)*. IEEE, Piscataway, NJ, 1264–1269.

- Fabrikant, A., A. Luthra, E. Maneva, C. H. Papadimitriou, S. Shenker. 2003. On a network creation game. *Proc. 22nd Annual ACM Sympos. Principles of Distributed Comput.* ACM, New York, 347–351.
- Fadel, R., I. Segal. 2009. The communication cost of selfishness. *J. Econom. Theory*. Forthcoming.
- Falkner, M., M. Devetsikiotis, I. Lambadaris. 2000. An overview of pricing concepts for broadband IP networks. *IEEE Comm. Surveys* 3(2) 2–13.
- Feldman, M., K. Lai, L. Zhang. 2005. A price-anticipating resource allocation mechanism for distributed shared clusters. *Proc. ACM Conf. Electronic Commerce (EC)*. ACM, New York, 127–136.
- Fudenberg, D., J. Tirole. 1991. *Game Theory*. MIT Press, Cambridge, MA.
- Green, J. R., J.-J. Laffont. 1979. *Incentives in Public Decision-Making*. North-Holland, Amsterdam.
- Groves, T. 1973. Incentives in teams. *Econometrica* 41(4) 617–631.
- Johari, R. 2004. Efficiency loss in market mechanisms for resource allocation. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Johari, R., J. N. Tsitsiklis. 2004. Efficiency loss in a network resource allocation game. *Math. Oper. Res.* 29(3) 407–435.
- Johari, R., J. N. Tsitsiklis. 2005. Communication requirements of VCG-like mechanisms in convex environments. *Allerton Conf. Comm., Control, Comput.*, Monticello, IL.
- Johari, R., J. N. Tsitsiklis. 2006. A scalable network resource allocation mechanism with bounded efficiency loss. *IEEE J. Selected Areas Comm.* 24(5) 992–999. Special Issue on Price-Based Access Control and Economics for Communication Networks.
- Johari, R., S. Mannor, J. N. Tsitsiklis. 2005. Efficiency loss in a network resource allocation game: The case of elastic supply. *IEEE Trans. Automatic Control* 50(11) 1712–1724.
- Kelly, F. P. 1997. Charging and rate control for elastic traffic. *Eur. Trans. Telecomm.* 8 33–37.
- Kelly, F. P., A. K. Maulloo, D. K. H. Tan. 1998. Rate control for communication networks: Shadow prices, proportional fairness, and stability. *J. Oper. Res. Soc.* 49 237–252.
- Klemperer, P. D., M. A. Meyer. 1989. Supply function equilibria in oligopoly under uncertainty. *Econometrica* 57(6) 1243–1277.
- Koutsoupias, E., C. Papadimitriou. 1999. Worst-case equilibria. C. Meinel, S. Tison, eds. *Proc. 16th Annual Sympos. Theoret. Aspects Comput. Sci.*, Lecture Notes in Computer Science. Springer, New York, 404–413.
- Maheswaran, R. T. 2003. A game-theoretic analysis of agent-mediated resource allocation. Ph.D. thesis, University of Illinois at Urbana-Champaign, Champaign, IL.
- Maheswaran, R. T., T. Basar. 2004. Social welfare of selfish agents: Motivating efficiency for divisible resources. *Proc. IEEE Conf. Decision and Control (CDC)*. IEEE, Piscataway, NJ, 1550–1555.
- Mas-Colell, A., M. D. Whinston, J. R. Green. 1995. *Microeconomic Theory*. Oxford University Press, Oxford, UK.
- Mavronicolas, M., P. Spirakis. 2001. The price of selfish routing. *Proc. ACM Sympos. Theory Comput.* ACM, New York, 510–519.
- Moulin, H. 2008a. An efficient and almost budget-balanced cost sharing method. *Games and Econom. Behavior*. Forthcoming.
- Moulin, H. 2008b. The price of anarchy of serial, average, and incremental cost sharing. *Econom. Theory*. 36 379–405.
- Nisan, N., T. Roughgarden, E. Tardos, V. V. Vazirani, eds. 2007. *Algorithmic Game Theory*. Cambridge University Press, Cambridge, UK.
- Reichelstein, S., S. Reiter. 1988. Game forms with minimal message spaces. *Econometrica* 56(3) 661–692.
- Roels, G., G. Perakis. 2007. The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts. *Management Sci.* 53(8) 1249–1268.
- Roughgarden, T., É. Tardos. 2002. How bad is selfish routing? *J. ACM* 49(2) 236–259.
- Sanghavi, S., B. Hajek. 2004. Optimal allocation of a divisible good to strategic buyers. *Proc. IEEE Conf. Decision and Control (CDC)*. IEEE, Piscataway, NJ, 2748–2753.
- Semret, N. 1999. Market mechanisms for network resource sharing. Ph.D. thesis, Columbia University, New York.
- Shapley, L. S., M. Shubik. 1977. Trade using one commodity as a means of payment. *J. Political Econom.* 85(5) 937–968.
- Shenker, S. 1990. A theoretical analysis of feedback flow control. *Proc. ACM SIGCOMM*. ACM, New York, 151–160.
- Shenker, S., D. Clark, D. Estrin, S. Herzog. 1996. Pricing in computer networks: Reshaping the research agenda. *Telecomm. Policy* 20(3) 183–201.
- Srikant, R. 2004. *The Mathematics of Internet Congestion Control*. Birkhauser, Boston.
- Stoescu, T. M., J. Ledyard. 2006. A pricing mechanism which implements in Nash equilibria a rate allocation problem in networks. *Proc. IEEE Conf. Decision and Control (CDC)*. IEEE, Piscataway, NJ, 1270–1277.
- Stoft, S. 2002. *Power System Economics: Designing Markets for Electricity*. IEEE Press, Piscataway, NJ.
- Varian, H. R. 1992. *Microeconomic Analysis*. W. W. Norton and Co., New York.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* 16(1) 8–37.
- Yang, S., B. Hajek. 2006a. Revenue and stability of a mechanism for efficient allocation of a divisible good. Submitted.
- Yang, S., B. Hajek. 2006b. VCG-Kelly mechanisms for allocation of divisible goods: Adapting VCG mechanisms to one-dimensional signals. *40th Annual Conf. Inform. Sci. Systems*. Princeton, NJ, 1391–1396.
- Yang, S., B. Hajek. 2007. VCG-Kelly mechanisms for divisible goods: Adapting VCG mechanisms to one-dimensional signals. *IEEE J. Selected Areas Comm.* 25(6) 1237–1243.
- Zhang, L. 2005. The efficiency and fairness of a fixed budget resource allocation game. *Proc. 32nd Internat. Colloquium on Automata, Languages and Programming*. Springer, New York, 485–496.