
A Game Theoretic View of Efficiency Loss in Resource Allocation

Ramesh Johari¹ and John N. Tsitsiklis²

¹ Stanford University, Stanford, CA ramesh.johari@stanford.edu

² MIT, Cambridge, MA jnt@mit.edu

Dedicated to Pravin Varaiya on the occasion of his 65th birthday

Summary. Motivated by resource allocation problems in communication networks as well as power systems, we consider the design of market mechanisms for such settings which are robust to gaming behavior by market participants. Recent results in this work are reviewed, including: (1) efficiency loss guarantees for a data rate allocation mechanism first proposed by Kelly, both when link capacities are fixed and when they are elastic; (2) characterization of mechanisms that minimize the efficiency loss, within a certain class of “simple” mechanisms; (3) extensions to general networks; and (4) mechanism design for supply function bidding in electric power systems.

1 Introduction

This paper addresses a problem at the nexus of engineering, computer science, and economics: in large scale, decentralized systems, how can we efficiently allocate scarce resources among competing interests? On one hand, constraints are imposed on the system designer by the inherent architecture of any large scale system. These constraints are counterbalanced by the need to design mechanisms that efficiently allocate resources, even when the system is being used by participants who have only their own interests at stake.

We consider two main classes of resource allocation problems. First, we consider a setting where a resource in scarce supply must be allocated among multiple competing consumers. Second, we discuss a setting where multiple producers compete to satisfy a fixed demand. The former model is motivated by applications to communication networks, while the latter is motivated by electric power market design.

What goals might we have for markets in such settings? We would of course like the equilibria of mechanisms designed for such settings to be “desirable;” a common requirement is that equilibria should be Pareto efficient.

In other contexts, we want the equilibria to satisfy a predetermined notion of fairness; or we may wish the resulting vector of monetary transfers to satisfy certain properties, such as profit maximization for the market operator. Beyond such constraints on the properties at equilibrium, however, we are also concerned with the complexity of such mechanisms. In particular, we may desire mechanisms which have relatively low information overhead: the strategy spaces of the players should be “simple,” and the feedback from the market to the players should be “simple” as well. Often, such complexity issues arise in a discussion of the dynamic behavior of market mechanisms, in trying to determine whether equilibria are actually achieved over time by players.

In this paper, we will focus on efficiency of mechanisms which maintain low complexity, appropriately defined. We focus on efficiency primarily as a first test of feasibility. Traditionally, economics has focused on selection of efficient mechanisms because mechanisms with inefficient equilibria are less likely to be useful in practice. Indeed, the classical theory of mechanism design is largely devoted to determining when fully efficient equilibria can be guaranteed (see, e.g., Chapter 23 of [23] for an overview).

The landmark contribution of mechanism design is the Vickrey-Clarke-Groves class of mechanisms, which guarantee efficient allocations at dominant strategy equilibria [4, 11, 32]; unfortunately, implementing VCG mechanisms is generally a very complex proposition with many possible pitfalls [2, 27]. The task is further complicated by the fact that the VCG class of mechanisms are essentially the only class which guarantee fully efficient outcomes as dominant strategy equilibria [8]. Thus, to make progress, either the notion of equilibrium must be weakened, or some efficiency must be lost. Previous results in the economics literature have considered weakening the notion of equilibrium; for example, Maskin has shown that if we only consider Nash equilibria, efficiency can be guaranteed if certain conditions are satisfied by players’ characteristics [24]. However, no guidance is available as to how to design such mechanisms with low complexity.

We consider an alternate approach, by weakening of the requirement of efficiency. The basic technique we consider is one of restricting the strategy spaces of the players (either buyers or sellers). With the proper choice of restriction, we can achieve two goals simultaneously. First, by ensuring that strategy spaces are relatively simple, we can restrict attention to mechanisms with low complexity. Second, if strategies of players are restricted, we can reduce their opportunities to game the system; this will lead to provable bounds on efficiency loss at Nash equilibria.

In the remainder of the paper, we provide an overview of the progress made in our earlier work [17, 16, 15]. In Section 2, we consider a setting of multiple consumers and inelastic supply, motivated by rate allocation in communication networks. For a single link of fixed capacity, we investigate a resource allocation mechanism proposed by Kelly [18]. Network users choose bids, which denote the total amount they are willing to pay. A price is then chosen to clear the market; for the case of a single link, this allocation mech-

anism allocates fractions of the resource to the users in *proportion* to their bids. Kelly has previously shown that if users are *price taking*—that is, if they do not anticipate the effects of their actions on the market-clearing price—the resulting *competitive equilibrium* allocation is fully efficient. Our key result in this section is that when users are price anticipating, aggregate utility falls by no more than 25% relative to the maximum possible.

In Section 3, we consider the same basic mechanism as in Section 2, but now consider a setting where supply is elastic; this is the model considered by Kelly et al. [19]. In this case the link is characterized by a cost depending on the total allocated rate, rather than a fixed capacity. Again, Kelly et al. have previously shown that if users are price taking, this mechanism maximizes aggregate surplus (i.e., aggregate utility minus cost). For this setting we establish that when users are price anticipating, aggregate surplus falls by no more than approximately 34% relative to the maximum possible.

Sections 2 and 3 establish efficiency loss results for a specific market mechanism. In Section 4, we characterize the mechanism studied in Section 2 as the “best” choice of mechanism under reasonable assumptions. Formally, we show that in a class of market-clearing mechanisms satisfying certain simple mathematical assumptions and for which there exist fully efficient competitive equilibria, the mechanism of Section 2 uniquely minimizes efficiency loss when market participants are price anticipating. These results justify the attention devoted to understanding the particular market mechanism studied in Sections 2 and 3; furthermore, they clearly delineate conditions which must be violated if we hope to achieve higher efficiency guarantees than those provided by the results of Sections 2 and 3.

In Section 5, we summarize two further directions of research. First, in Section 5.1, we discuss the generalization of the models of Sections 2 and 3 to networks with arbitrary topology. We consider games where users submit individual bids to each link in the network. Such games are then proven to have the same efficiency loss guarantees as the single link games considered in Sections 2 and 3.

Next, in Section 5.2, motivated by power systems, we discuss a setting where multiple producers bid to satisfy an inelastic demand D . We consider a market mechanism where producers submit supply functions restricted to lie in a certain one-parameter family, and a market-clearing price is chosen to ensure that aggregate supply is equal to the inelastic demand. We establish that when producers are price anticipating, aggregate production cost rises by no more than a factor $1 + 1/(N - 2)$ relative to the minimum possible production cost, where $N > 2$ is the number of firms competing. Finally, we conclude with some open issues in Section 6.

2 Multiple Consumers, Inelastic Supply

Suppose R users share a communication link of capacity $C > 0$. Let d_r denote the rate allocated to user r . We assume that user r receives a *utility* equal to $U_r(d_r)$ if the allocated rate is d_r ; we assume that utility is measured in monetary units. We make the following assumptions on the utility function.

Assumption 1. For each r , over the domain $d_r \geq 0$ the utility function $U_r(d_r)$ is 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.

Given complete knowledge and centralized control of the system, it would be natural for the link manager to try to solve is the following optimization problem [18]:

$$\text{maximize } \sum_r U_r(d_r) \quad (1)$$

$$\text{subject to } \sum_r d_r \leq C; \quad (2)$$

$$d_r \geq 0, \quad r = 1, \dots, R. \quad (3)$$

Note that the objective function of this problem is the *aggregate utility*. Since the objective function is continuous and the feasible region is compact, an optimal solution $\mathbf{d} = (d_1, \dots, d_R)$ exists. If the functions U_r are strictly concave, then the optimal solution is unique, since the feasible region is convex.

In general, the utility functions are not available to the link manager. As a result, we consider the following pricing scheme for rate allocation. Each user r submits a payment (also called a *bid*) w_r to the link manager; we assume $w_r \geq 0$. Given the vector $\mathbf{w} = (w_1, \dots, w_r)$, the link manager chooses a rate allocation $\mathbf{d} = (d_1, \dots, d_r)$. We assume the manager treats all users alike—in other words, the link manager does not *price discriminate*. Each user is charged the same price $\mu > 0$, leading to $d_r = w_r/\mu$. We further assume the manager always seeks to allocate the entire link capacity C ; in this case, 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:

$$\mu = \frac{\sum_r w_r}{C}. \quad (4)$$

In other words, if the manager chooses to allocate the entire available rate at the link, and does not price discriminate between users, then for every nonzero \mathbf{w} there is a *unique* possible price $\mu > 0$, given by the previous equation.

We can interpret this mechanism as a *market-clearing* process by which a price is set so that demand equals supply. To see this interpretation, note that when a user submits a total payment w_r , it is as if the user has submitted a *demand function* $D(p, w_r) = w_r/p$ for $p > 0$. The demand function describes the rate that the user demands at any given price $p > 0$. The link 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 (4). User r then receives a rate allocation given by $D(\mu, w_r)$, and makes a payment $\mu D(\mu, w_r) = w_r$. This interpretation of the mechanism we consider here will be further explored in Section 4, where we consider other market-clearing mechanisms with the users submitting demand functions from a family parametrized by a single scalar.

In the remainder of the section, we consider two different models for how users might interact with this price mechanism. In Section 2.1, we consider a model where users do not anticipate the effect of their bids on the price, and provide a result, due to Kelly [18], on the existence of a competitive equilibrium. Furthermore, this competitive equilibrium leads to an allocation which is an optimal solution to (1)-(3). In Section 2.2, we change the model and assume users are price anticipating, and provide a result (due to Hajek and Gopalakrishnan [12]) on the existence and uniqueness of a Nash equilibrium. In Section 2.3, we then consider the loss of efficiency at this Nash equilibrium, relative to the optimal solution to (1)-(3).

2.1 Price Taking Users and Competitive Equilibrium

In this section, we consider a *competitive equilibrium* between the users and the link manager [23], following the development of Kelly [18]. A central assumption in the definition of competitive equilibrium is that each user does not anticipate the effect of their payment w_r on the price μ , i.e., each user acts as a *price taker*. In this case, given a price $\mu > 0$, user r acts to maximize the following payoff function over $w_r \geq 0$:

$$P_r(w_r; \mu) = U_r \left(\frac{w_r}{\mu} \right) - w_r. \quad (5)$$

The first term represents the utility to user r of receiving a rate allocation equal to w_r/μ ; the second term is the payment w_r made to the manager. Observe that since utility is measured in monetary units, the payoff is *quasilinear* in money [23].

We say that a pair (\mathbf{w}, μ) , with $\mathbf{w} \geq 0$ and $\mu > 0$, is a *competitive equilibrium* if users maximize their payoff as defined in (5), and the network “clears the market” by setting the price μ according to (4):

$$P_r(w_r; \mu) \geq P_r(\bar{w}_r; \mu) \quad \text{for } \bar{w}_r \geq 0, \quad r = 1, \dots, R; \quad (6)$$

$$\mu = \frac{\sum_r w_r}{C}. \quad (7)$$

Kelly shows in [18] that when users are price takers, there exists a competitive equilibrium, and the resulting allocation is an optimal solution to (1)-(3). This is formalized in the following theorem, adapted from [18].

Theorem 1 (Kelly, [18]). *Suppose that Assumption 1 holds. Then there exists a competitive equilibrium, i.e., a vector $\mathbf{w} = (w_1, \dots, w_R) \geq 0$ and a scalar $\mu > 0$ satisfying (6)-(7).*

In this case, the scalar μ is uniquely determined, and the vector $\mathbf{d} = \mathbf{w}/\mu$ is an optimal solution to (1)-(3). If the functions U_r are strictly concave, then \mathbf{w} is uniquely determined as well.

Theorem 1 shows that under the assumption that the users of the link behave as price takers, there exists a bid vector \mathbf{w} where all users have optimally chosen their bids w_r , with respect to the given price $\mu = \sum_r w_r/C$; and at this “equilibrium,” aggregate utility is maximized. However, when the price taking assumption is violated, the model changes into a game and the guarantee of Theorem 1 is no longer valid.

2.2 Price Anticipating Users and Nash Equilibrium

We now consider an alternative model where the users of a single link are price anticipating, rather than price takers. The key difference is that while the payoff function P_r takes the price μ as a fixed parameter in (5), price anticipating users will realize that μ is set according to (4), and adjust their payoff accordingly; this makes the model a game between the R players.

We use the notation \mathbf{w}_{-r} to denote the vector of all bids by users other than r ; i.e., $\mathbf{w}_{-r} = (w_1, w_2, \dots, w_{r-1}, w_{r+1}, \dots, w_R)$. Given \mathbf{w}_{-r} , each user r chooses w_r to maximize:

$$Q_r(w_r; \mathbf{w}_{-r}) = \begin{cases} U_r\left(\frac{w_r}{\sum_s w_s} C\right) - w_r, & \text{if } w_r > 0; \\ U_r(0), & \text{if } w_r = 0. \end{cases} \quad (8)$$

over all nonnegative w_r . The second condition is required so that the rate allocation to user r is zero when $w_r = 0$, even if all other users choose \mathbf{w}_{-r} so that $\sum_{s \neq r} w_s = 0$. The payoff function Q_r is similar to the payoff function P_r , except that the user anticipates that the network will set the price μ according to (4). A *Nash equilibrium* of the game defined by (Q_1, \dots, Q_R) is a vector $\mathbf{w} \geq 0$ such that for all r :

$$Q_r(w_r; \mathbf{w}_{-r}) \geq Q_r(\bar{w}_r; \mathbf{w}_{-r}), \quad \text{for all } \bar{w}_r \geq 0. \quad (9)$$

Hajek and Gopalakrishnan have shown that there exists a unique Nash equilibrium when multiple users share the link, by showing that at a Nash equilibrium it is *as if* the users are solving another optimization problem of the same form as the problem (1)-(3), but with “modified” utility functions. This is formalized in the following theorem, adapted from [12].

Theorem 2 (Hajek and Gopalakrishnan, [12]). *Suppose that $R > 1$, and that Assumption 1 holds. Then there exists a unique Nash equilibrium $\mathbf{w} \geq 0$ of the game defined by (Q_1, \dots, Q_R) , and it satisfies $\sum_r w_r > 0$.*

In this case, the vector \mathbf{d} defined by:

$$d_r = \frac{w_r}{\sum_s w_s} C, \quad r = 1, \dots, R, \quad (10)$$

is the unique optimal solution to the following optimization problem:

$$\text{maximize } \sum_r \hat{U}_r(d_r) \quad (11)$$

$$\text{subject to } \sum_r d_r \leq C; \quad (12)$$

$$d_r \geq 0, \quad r = 1, \dots, R, \quad (13)$$

where

$$\hat{U}_r(d_r) = \left(1 - \frac{d_r}{C}\right) U_r(d_r) + \left(\frac{d_r}{C}\right) \left(\frac{1}{d_r} \int_0^{d_r} U_r(z) dz\right). \quad (14)$$

Theorem 2 shows that the unique Nash equilibrium of the game is characterized as the solution to the above optimization problem. Other games have also profited from such relationships—notably traffic routing games, in which Nash-Wardrop equilibria can be found as solutions to a related global optimization problem. Roughgarden and Tardos use this fact to their advantage in computing efficiency loss for such games [28]; Correa, Schulz, and Stier Moses also use this relationship to consider routing games in capacitated networks [5]. Finally, we note that for the game presented here, several authors have derived results similar to Theorem 2 [7, 21, 22], though not as general.

2.3 Efficiency Loss

We let \mathbf{d}^S denote an optimal solution to (1)-(3), and let \mathbf{d}^G denote the unique optimal solution to (11)-(13). We now investigate the efficiency loss of this system; that is, the utility loss caused by the price anticipating behavior of the users. More precisely, we will compare the utility $\sum_r U_r(d_r^G)$ obtained when the users fully evaluate the effect of their actions on the price, and the maximum possible aggregate utility $\sum_r U_r(d_r^S)$. (We know, of course, that $\sum_r U_r(d_r^G) \leq \sum_r U_r(d_r^S)$, by definition of \mathbf{d}^S .) According to the following theorem, the worst case efficiency loss is exactly 25%; the proof may be found in [17].

Theorem 3. *Suppose that $R > 1$, and that Assumption 1 holds. Suppose also that $U_r(0) \geq 0$ for all r . If \mathbf{d}^S is any optimal solution to (1)-(3), and \mathbf{d}^G is the unique optimal solution to (11)-(13), then:*

$$\sum_r U_r(d_r^G) \geq \frac{3}{4} \sum_r U_r(d_r^S).$$

Furthermore, this bound is tight: for every $\varepsilon > 0$, there exists a choice of R , and a choice of (linear) utility functions U_r , $r = 1, \dots, R$, such that:

$$\sum_r U_r(d_r^G) \leq \left(\frac{3}{4} + \varepsilon\right) \left(\sum_r U_r(d_r^S)\right).$$

We provide some comments on the method for proving a result such as Theorem 3. The first step is to show that the worst case efficiency loss occurs when the utility functions belong to a certain finite-dimensional family; in the current context, it is the family of linear utility functions. Identifying the worst case utility functions amounts to minimizing an efficiency measure over all possible choices of the coefficients of the linear utility functions. It turns out that this minimization can be cast as a sequence of finite-dimensional nonlinear optimization problems (each problem in the sequence corresponding to a different number R of users), which can be studied analytically. In the context of Theorem 3, the worst efficiency loss corresponds to a link of capacity 1, where user 1 has utility $U_1(d_1) = d_1$, and all other users have utility $U_r(d_r) \approx d_r/2$. As $R \rightarrow \infty$, at the Nash equilibrium of the game, user 1 receives a rate $d_1^G = 1/2$, while the remaining users uniformly split the rate $1 - d_1^G = 1/2$ among themselves, yielding an aggregate utility of $3/4$.

We note that a similar bound was observed by Roughgarden and Tardos for traffic routing games with affine link latency functions [28]. They found that the ratio of worst case Nash equilibrium cost to optimal cost was $4/3$. However, it is questionable whether a relationship can be drawn between the two games; in particular, we note that while Theorem 3 holds even if the utility functions are nonlinear, Roughgarden and Tardos have shown that the efficiency loss due to selfish users in traffic routing may be arbitrarily high if link latency functions are nonlinear.

3 Multiple Consumers, Elastic Supply

In this section, we allow the supply of the scarce resource to be elastic, rather than fixed as in the previous section. Rather than being characterized by a capacity, we will characterize the resource through a *cost function* that gives the cost incurred by the resource as a function of the flow through it. We continue to assume that R users share a single communication link, and that user r receives a utility $U_r(d_r)$ if the allocated rate is d_r . We let $f = \sum_r d_r$ denote the total rate allocated at the link, and let $C(f)$ denote the cost incurred at the link when the total allocated rate is $f \geq 0$. We will assume that both U_r and C are measured in the same monetary units. A natural interpretation is that $U_r(d_r)$ is the monetary value to user r of a rate

allocation d_r , and $C(f)$ is a monetary cost for congestion at the link when the total allocated rate is f .

We continue to assume the utility functions U_r satisfy Assumption 1. In addition, we make the following assumption on the cost function C .

Assumption 2. *There exists a continuous, convex, strictly increasing function $p(f)$ over $f \geq 0$ with $p(0) = 0$, such that for $f \geq 0$:*

$$C(f) = \int_0^f p(z) dz.$$

Thus $C(f)$ is strictly convex and strictly increasing.

Given complete knowledge and centralized control of the system, it would be natural for the link manager to try to solve the following optimization problem[18]:

$$\text{maximize } \sum_r U_r(d_r) - C\left(\sum_r d_r\right) \tag{15}$$

$$\text{subject to } d_r \geq 0, \quad r = 1, \dots, R. \tag{16}$$

We refer to the objective function (15) as the *aggregate surplus*. This is the net monetary benefit to the economy consisting of the users and the single link. Since the objective function is continuous, and U_r increases at most linearly, while C increases superlinearly, an optimal solution $\mathbf{d}^S = (d_1^S, \dots, d_R^S)$ exists; since the feasible region is convex and C is strictly convex, if the functions U_r are strictly concave, then the optimal solution is unique.

We consider the following pricing scheme for rate allocation, a natural analogue of the mechanism presented in Section 2. Each user r submits a payment (or *bid*) of w_r to the resource manager. Given the composite vector $\mathbf{w} = (w_1, \dots, w_r)$, the resource manager chooses a rate allocation $\mathbf{d}(\mathbf{w}) = (d_1(\mathbf{w}), \dots, d_R(\mathbf{w}))$. We assume the manager treats all users alike—in other words, the network manager does not *price differentiate*. Thus the network manager sets a single price $\mu(\mathbf{w})$; we assume that $\mu(\mathbf{w}) = 0$ if $w_r = 0$ for all r , and $\mu(\mathbf{w}) > 0$ otherwise. All users are then charged the same price $\mu(\mathbf{w})$, leading to:

$$d_r(\mathbf{w}) = \begin{cases} 0, & \text{if } w_r = 0; \\ \frac{w_r}{\mu(\mathbf{w})}, & \text{if } w_r > 0. \end{cases}$$

Notice that, with this formulation, the rate allocated to user r is similar to the rate allocated to user r in the model of Section 2. The key difference in this setting is that the aggregate rate is not constrained to an inelastic supply; rather, associated with the choice of price $\mu(\mathbf{w})$ is an aggregate rate function $f(\mathbf{w})$, defined by:

$$f(\mathbf{w}) = \sum_r d_r(\mathbf{w}) = \begin{cases} 0, & \text{if } \sum_r w_r = 0; \\ \frac{\sum_r w_r}{\mu(\mathbf{w})}, & \text{if } \sum_r w_r > 0. \end{cases} \quad (17)$$

Let us assume for now that given a price $\mu > 0$, user r wishes to maximize the following payoff function over $w_r \geq 0$:

$$P_r(w_r; \mu) = U_r\left(\frac{w_r}{\mu}\right) - w_r. \quad (18)$$

The first term represents the utility to user r of receiving a rate allocation equal to w_r/μ ; the second term is the payment w_r made to the manager.

Notice that as formulated above, the payoff function P_r assumes that user r acts as a *price taker*; that is, user r does not *anticipate* the effect of his choice of w_r on the price μ , and hence on his resulting rate allocation $d_r(\mathbf{w})$. Informally, we expect that in such a situation the aggregate surplus will be maximized if the network manager sets a price equal to marginal cost; that is, if the price function satisfies:

$$\mu(\mathbf{w}) = p(f(\mathbf{w})). \quad (19)$$

The well-posedness of such a pricing mechanism is the subject of the following proposition.

Proposition 1. *Suppose Assumption 2 holds. Given any vector of bids $\mathbf{w} \geq 0$, there exists a unique pair $(\mu(\mathbf{w}), f(\mathbf{w})) \geq 0$ satisfying (17) and (19), and in this case $f(\mathbf{w})$ is the unique solution f to:*

$$\sum_r w_r = fp(f). \quad (20)$$

Furthermore, $f(\cdot)$ has the following properties: (1) $f(\mathbf{0}) = 0$; (2) $f(\mathbf{w})$ is continuous for $\mathbf{w} \geq 0$; (3) $f(\mathbf{w})$ is a strictly increasing and strictly concave function of $\sum_r w_r$; and (4) $f(\mathbf{w}) \rightarrow \infty$ as $\sum_r w_r \rightarrow \infty$.

Observe that we can view (20) as a market-clearing process. Given the total revenue $\sum_r w_r$ from the users, the link manager chooses an aggregate rate $f(\mathbf{w})$ so that the revenue is exactly equal to the aggregate charge $f(\mathbf{w})p(f(\mathbf{w}))$. Due to Assumption 2, this market-clearing aggregate rate is uniquely determined. Kelly et al. present two algorithms in [19] which amount to dynamic processes of market-clearing; as a result, a key motivation for the mechanism we study in this section is that it represents the equilibrium behavior of the algorithms in [19].

For the remainder of this section, we will assume that $\mu(\mathbf{w})$ is set according to the choice prescribed in Proposition 1, as follows.

Assumption 3. For all $\mathbf{w} \geq 0$, the aggregate rate $f(\mathbf{w})$ is the unique solution to (20): $\sum_r w_r = f(\mathbf{w})p(f(\mathbf{w}))$. Furthermore, for each r , $d_r(\mathbf{w})$ is given by:

$$d_r(\mathbf{w}) = \begin{cases} 0, & \text{if } w_r = 0; \\ \frac{w_r}{p(f(\mathbf{w}))}, & \text{if } w_r > 0. \end{cases} \quad (21)$$

Note that we have $f(\mathbf{w}) > 0$ and $p(f(\mathbf{w})) > 0$ if $\sum_r w_r > 0$, and hence d_r is always well defined.

In the remainder of this section, we consider two different models for how users might interact with this price mechanism. In Section 3.1, we consider a model where users do not anticipate the effect of their bids on the price, in which case there exists a competitive equilibrium. Furthermore, this competitive equilibrium leads to an allocation which is an optimal solution to (15)-(16). In Section 3.2, we change the model and assume users are price anticipating, in which case there exists a Nash equilibrium. Finally, Section 3.3 considers the loss of efficiency at Nash equilibria, relative to the optimal solution to (15)-(16).

3.1 Price Taking Users and Competitive Equilibrium

Kelly et al. show in [19] that when users are price takers, and the network sets the price $\mu(\mathbf{w})$ according to (17) and (19), the resulting allocation is an optimal solution to (15)-(16). This is formalized in the following theorem, adapted from [19].

Theorem 4 (Kelly et al., [19]). Suppose Assumptions 1, 2, and 3 hold. Then there exists a vector \mathbf{w} such that $\mu(\mathbf{w}) > 0$, and:

$$P_r(w_r; \mu(\mathbf{w})) = \max_{\bar{w}_r \geq 0} P_r(\bar{w}_r; \mu(\mathbf{w})), \quad r = 1, \dots, R. \quad (22)$$

For any such vector \mathbf{w} , the vector $\mathbf{d}(\mathbf{w}) = \mathbf{w}/\mu(\mathbf{w})$ is an optimal solution to (15)-(16). If the functions U_r are strictly concave, such a vector \mathbf{w} is unique as well.

Theorem 4 shows that with an appropriate choice of price function (as determined by (17) and (19)), and under the assumption that the users of the link behave as price takers, there exists a bid vector \mathbf{w} where all users have optimally chosen their bids w_r , with respect to the given price $\mu(\mathbf{w})$; and at this “equilibrium,” the aggregate surplus is maximized. However, when the price taking assumption is violated, the model changes into a game and the guarantee of Theorem 4 is no longer valid.

3.2 Price Anticipating Users and Nash Equilibrium

We now consider an alternative model where the users of a single link are price anticipating, rather than price taking, and play a game to acquire a share of the link. Throughout the remainder of this section as well as in Section 3.3, we will assume that the link manager sets the price $\mu(\mathbf{w})$ according to the unique choice prescribed by Proposition 1, as follows.

We adopt the notation \mathbf{w}_{-r} to denote the vector of all bids by users other than r ; i.e., $\mathbf{w}_{-r} = (w_1, w_2, \dots, w_{r-1}, w_{r+1}, \dots, w_R)$. Then given \mathbf{w}_{-r} , each user r chooses $w_r \geq 0$ to maximize:

$$Q_r(w_r; \mathbf{w}_{-r}) = U_r(d_r(\mathbf{w})) - w_r, \quad (23)$$

over nonnegative w_r . The payoff function Q_r is similar to the payoff function P_r , except that the user now anticipates that the network will set the price according to Assumption 3, as captured by the allocated rate $d_r(\mathbf{w})$. A *Nash equilibrium* of the game defined by (Q_1, \dots, Q_R) is a vector $\mathbf{w} \geq 0$ such that for all r :

$$Q_r(w_r; \mathbf{w}_{-r}) \geq Q_r(\bar{w}_r; \mathbf{w}_{-r}), \quad \text{for all } \bar{w}_r \geq 0. \quad (24)$$

The proof of the following proposition can be found in [16].

Proposition 2. *Suppose that Assumptions 1, 2, and 3 hold. Then there exists a Nash equilibrium \mathbf{w} for the game defined by (Q_1, \dots, Q_R) .*

3.3 Efficiency Loss

We let \mathbf{d}^S denote an optimal solution to (15)-(16), and let \mathbf{w} denote any Nash equilibrium of the game defined by (Q_1, \dots, Q_R) . We now investigate the associated efficiency loss. In particular, we compare the aggregate surplus $\sum_r U_r(d_r(\mathbf{w})) - C(\sum_r d_r(\mathbf{w}))$ obtained when the users fully evaluate the effect of their actions on the price, and the aggregate surplus $\sum_r U_r(d_r^S) - C(\sum_r d_r^S)$ obtained by choosing an allocation which maximizes aggregate surplus. According to the following theorem, the efficiency loss is no more than approximately 34%, and this bound is essentially tight; the proof can be found in [16].

Theorem 5. *Suppose that Assumptions 1, 2, and 3 hold. Suppose also that $U_r(0) \geq 0$ for all r . If \mathbf{d}^S is any optimal solution to (15)-(16), and \mathbf{w} is any Nash equilibrium of the game defined by (Q_1, \dots, Q_R) , then:*

$$\sum_r U_r(d_r(\mathbf{w})) - C\left(\sum_r d_r(\mathbf{w})\right) \geq (4\sqrt{2} - 5) \left(\sum_r U_r(d_r^S) - C\left(\sum_r d_r^S\right) \right). \quad (25)$$

In other words, there is no more than approximately a 34% efficiency loss when users are price anticipating.

Furthermore, this bound is tight: for every $\delta > 0$, there exists a choice of R , a choice of (linear) utility functions U_r , $r = 1, \dots, R$, and a (piecewise linear) price function p such that a Nash equilibrium \mathbf{w} exists with:

$$\sum_r U_r(d_r(\mathbf{w})) - C \left(\sum_r d_r(\mathbf{w}) \right) \leq (4\sqrt{2} - 5 + \delta) \left(\sum_r U_r(d_r^S) - C \left(\sum_r d_r^S \right) \right). \quad (26)$$

Let us remark here that, according to the proof of Theorem 5, the worst possible efficiency loss is achieved along a sequence of games where:

1. The price function p has the following form, with $b \rightarrow \infty$:

$$p(f) = \begin{cases} (2 - \sqrt{2})f, & \text{if } 0 \leq f \leq 1; \\ 2 - \sqrt{2} + b(f - 1), & \text{if } f \geq 1; \end{cases}$$

2. The number of users becomes large ($R \rightarrow \infty$); and
3. User 1 has linear utility with $U_1(d_1) = d_1$, and all other users r have linear utility with $U_r(d_r) = \alpha_r d_r$, where $\alpha_r \approx p(1) = 2 - \sqrt{2}$.

(Note that formally, we must take care that the limits of $R \rightarrow \infty$ and $b \rightarrow \infty$ are taken in the correct order; in particular, in the proof we first have $R \rightarrow \infty$, and then $b \rightarrow \infty$.) In this limit, we find that at the Nash equilibrium the aggregate allocated rate is 1, and the Nash equilibrium aggregate surplus converges to $4\sqrt{2} - 5$.

4 A Characterization Theorem

In this section we revisit the resource allocation problem of Section 2, and address the following question: can we identify a mechanism that minimizes the efficiency loss, in the presence of price anticipating users, within a class of mechanisms with certain desirable properties?

Formally, we consider a collection of users bidding to receive a share of a finite, infinitely divisible resource of capacity C . Each user has a utility function $U : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ (where $\mathbb{R}^+ = [0, \infty)$) that satisfies Assumption 1. More specifically, U belongs to the set \mathcal{U} utility functions defined by

$$\mathcal{U} = \{U : \mathbb{R}^+ \rightarrow \mathbb{R}^+ \mid U \text{ is continuous, strictly increasing, concave on } [0, \infty), \text{ and continuously differentiable on } [0, \infty), \text{ with } U'(0) < \infty\}.$$

We let R denote the number of users, and let $\mathbf{U} = (U_1, \dots, U_R)$ denote the vector of utility functions, where U_r is the utility function of user r . We call a pair (R, \mathbf{U}) , where $R > 1$ and $\mathbf{U} \in \mathcal{U}^R$, a *utility system*; our goal will be to design a resource allocation mechanism with attractive efficiency guarantees for all utility systems.

We assume once more that utility is measured in monetary units; thus, if user r receives a rate allocation d_r , but must pay w_r , his net net payoff is:

$$U_r(d_r) - w_r.$$

Given a utility system $\mathbf{U} \in \mathcal{U}^R$, the social objective is to maximize aggregate utility, as defined in the problem (1)-(3); we repeat that problem here, and refer to it as the problem $SYSTEM(C, R, \mathbf{U})$, to emphasize that the problem is specified by C , R , and the utility system (U).

$$\text{maximize } \sum_{r=1}^R U_r(d_r) \quad (27)$$

$$\text{subject to } \sum_{r=1}^R d_r \leq C; \quad (28)$$

$$\mathbf{d} \geq 0. \quad (29)$$

We will say that \mathbf{d} solves $SYSTEM(C, R, \mathbf{U})$ if \mathbf{d} is an optimal solution to (27)-(29), given the utility system (R, \mathbf{U}) .

In general, the utility system (R, \mathbf{U}) is unknown to the mechanism designer, so a mechanism must be designed to elicit information from the users. We will focus on mechanisms in which each user i submits a demand function, within a one-parameter family of admissible demand functions. In particular, each user has a one-dimensional strategic variable, denoted by θ_i .

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

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

We let $p_D(\boldsymbol{\theta})$ denote this solution, and refer to it as the market-clearing price.

Note that 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 mechanism D , we will assume that the payoff to all players is $-\infty$ if the composite strategy vector is $\boldsymbol{\theta} = \mathbf{0}$. Note that this is slightly different from the definition in Section 2, where the payoff to a player who submits $\theta = 0$ is set to zero. We will discuss this distinction further later; we simply note for the moment that it does not affect the results of this section.

Our definition of smooth market-clearing mechanisms generalizes the mechanism discussed in Section 2. We recall that in that development, each user submits a demand function of the form $D(p, \theta) = \theta/p$, and the link manager chooses a price $p_D(\boldsymbol{\theta})$ to ensure that $\sum_{r=1}^R D(p, \theta_r) = C$. Thus, for this mechanism, we have $p_D(\boldsymbol{\theta}) = \sum_{r=1}^R \theta_r / C$ if $\boldsymbol{\theta} \neq \mathbf{0}$. Another related example

is provided by $D(p, \theta) = \theta/\sqrt{p}$; in this case it is straightforward to verify that $p_D(\boldsymbol{\theta}) = (\sum_{r=1}^R \theta_r/C)^2$, for $\boldsymbol{\theta} \neq \mathbf{0}$.

We will further restrict attention to a particular class of mechanisms denoted \mathcal{D} , which we define as follows.

Definition 2. *The class \mathcal{D} consists of all functions $D(p, \theta)$ such that the following conditions are satisfied:*

1. *For all $C > 0$, D is a smooth market-clearing mechanism for C (cf. Definition 1).*
2. *For all $C > 0$, and for all $U_r \in \mathcal{U}$, the payoff of a price anticipating user is concave; that is, for all R , and for all $\boldsymbol{\theta}_{-r} \in (\mathbb{R}^+)^R$, 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 $\boldsymbol{\theta}_{-r} = \mathbf{0}$, and concave in $\theta_r \geq 0$ if $\boldsymbol{\theta}_{-r} \neq \mathbf{0}$.

3. *The demand functions are nonnegative; i.e., for all $p > 0$ and $\theta \geq 0$, $D(p, \theta) \geq 0$.*

The first condition requires a mechanism in \mathcal{D} to be a smooth market-clearing mechanism for any $C > 0$; in particular, the market-clearing price $p_D(\boldsymbol{\theta})$ must be uniquely defined for any $C > 0$. (Note that in the notation we suppress the dependence of the market-clearing price $p_D(\boldsymbol{\theta})$ on the capacity C .) The second condition allows us to characterize Nash equilibria in terms of only first order conditions; indeed, some such assumption needs to be in place in order to guarantee existence of pure strategy Nash equilibria [26]. Finally, the third condition is a normalization condition, which ensures that a user is never required to *supply* some quantity of the resource (which would be the case if we allowed $D(p, \theta) < 0$).

In order to state the main result of this section, we must define *competitive equilibrium* and *Nash equilibrium*. Given a utility system (R, \mathbf{U}) , a capacity $C > 0$, and a smooth market-clearing mechanism $D \in \mathcal{D}$, we say that a nonzero vector $\boldsymbol{\theta} \in (\mathbb{R}^+)^R$ is a competitive equilibrium if $\mu = p_D(\boldsymbol{\theta})$ satisfies:

$$\theta_r \in \arg \max_{\bar{\theta}_r \geq 0} [U_r(D(\mu, \bar{\theta}_r)) - \mu D(\mu, \bar{\theta}_r)], \quad \forall r. \quad (30)$$

Similarly, we say that a nonzero vector $\boldsymbol{\theta} \in (\mathbb{R}^+)^R$ is a Nash equilibrium if:

$$\theta_r \in \arg \max_{\bar{\theta}_r \geq 0} Q_r(\bar{\theta}_r; \boldsymbol{\theta}_{-r}), \quad \forall r, \quad (31)$$

where

$$Q_r(\theta_r; \boldsymbol{\theta}_{-r}) = \begin{cases} U_r(D(p_D(\boldsymbol{\theta}), \theta_r)) - p_D(\boldsymbol{\theta})D(p_D(\boldsymbol{\theta}), \theta_r), & \text{if } \boldsymbol{\theta} \neq \mathbf{0}; \\ -\infty, & \text{if } \boldsymbol{\theta} = \mathbf{0}. \end{cases} \quad (32)$$

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

Our interest is in the worst-case ratio of aggregate utility at any Nash equilibrium to the optimal value of $SYSTEM(C, R, \mathbf{U})$ (termed the “price of anarchy” by Papadimitriou [25]). Formally, for $D \in \mathcal{D}$ and a capacity $C > 0$ we define a constant $\rho(C, D)$ as follows:

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

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

The following theorem shows that among smooth market-clearing mechanisms for which there always exists a fully efficient competitive equilibrium, the mechanism proposed in Section 2 minimizes efficiency loss when users are price anticipating. The proof can be found in Chapter 5 of [15].

Theorem 6. *Let $D \in \mathcal{D}$ be a smooth market-clearing mechanism such that for all capacities $C > 0$ and utility systems (R, \mathbf{U}) , there exists a competitive equilibrium $\boldsymbol{\theta}$ such that $(D(p_D(\boldsymbol{\theta}), \theta_r), r = 1, \dots, R)$ solves $SYSTEM(C, R, \mathbf{U})$. Then for any capacity C and utility system (R, \mathbf{U}) , there exists a unique Nash equilibrium. Furthermore, $\rho(C, D) \leq 3/4$ for all $C > 0$ and all $D \in \mathcal{D}$, and this bound is met with equality if and only if $D(p, \boldsymbol{\theta}) = \Delta \boldsymbol{\theta}/p$ for some $\Delta > 0$.*

Theorem 6 suggests that the best efficiency guarantee we can hope to achieve is 75%, if we are restricted to market-clearing mechanisms with scalar strategy spaces. A key restriction in the mechanisms we consider is that a single price is chosen to clear the market. If the market designer is granted the latitude to *price discriminate* (i.e., to charge a different price to each user), better efficiency guarantees are possible. The most famous mechanisms which ensure such a guarantee are the Vickrey-Clarke-Groves class of mechanisms, for which fully efficient dominant strategy equilibria exist [4, 11, 32]. More recently, in a networking context, Sanghavi and Hajek [30] have shown that if users choose their payments (as in the Kelly mechanism), but the link manager is allowed to choose the allocation to users as an arbitrary function of the payments, it is possible to ensure no worse than a 13% efficiency loss. Furthermore, Yang and Hajek [34] have shown that if a mechanism allocates resources in proportion to the users’ strategies (i.e., user r receives a fraction $\theta_r / (\sum_{s=1}^R \theta_s)$ of the resource), then by using differentiated pricing, it is possible to guarantee arbitrarily small efficiency loss at the Nash equilibrium. The mechanisms proposed by both Sanghavi and Hajek [30] as well as Yang

and Hajek [34] require price discrimination, since the ratio of payment to allocation is not necessarily identical for all users (as must be the case in the market-clearing mechanisms studied here).

5 Further Directions

In addition to the results outlined above, several additional threads are included in this body of research. In this section, we describe two extensions: (1) resource allocation in general networks; and (2) a setting of multiple producers competing to satisfy an inelastic demand.

5.1 General Networks

The models presented in Sections 2 and 3 only consider resource allocation for a single link. We now consider extensions to the network case, following [16] and [17]. We consider networks consisting of a set of links; each user has a set of paths available through the network to send traffic, and each path uses a subset of the links. In a setting of inelastic supply, each link j is characterized by a fixed capacity C_j . In a setting of elastic supply, each link j is characterized by a cost function $C_j(\cdot)$. We continue to assume that each user r receives a utility $U_r(d_r)$ from a total rate allocation d_r ; however, note that in a network context d_r is the total rate delivered to user r across all paths available to user r through the network.

We extend the single link market mechanisms to multiple links by treating each link as a separate market. Thus we consider a game where each user requests service from multiple links by submitting an individual bid to each link. Links then allocate rates using the same scheme as in the single link model, and each user sends the maximum rate possible, given the vector of rates allocated from links in the network. Although this definition of the game is natural, we demonstrate that Nash equilibria may not exist in the setting of inelastic supply, due to a discontinuity in the payoff functions of individual players. (This problem also arises in the single link setting, but is irrelevant there as long as at least two players share the link.) To address the discontinuity, we extend the strategy space by allowing each user to request a nonzero rate without submitting a positive bid to a link, if the total payment made by other users at that link is zero; this extension is sufficient to guarantee existence of a Nash equilibrium. In the setting of elastic supply, Nash equilibria are always guaranteed to exist, without having to extend the strategy space. Finally, we show that in this network setting, if link capacities are inelastic then the total utility achieved at any Nash equilibrium of the game is no less than $3/4$ of the maximum possible aggregate utility; and if link supplies are elastic then the aggregate surplus achieved at any Nash equilibrium of the game is no less than a factor $4\sqrt{2} - 5$ of the maximal aggregate surplus. These

results extend the efficiency loss results from the single link setting to general networks.

The mechanisms we have studied require each user to submit a separate bid for each link that the user may use. An alternative mechanism had been proposed earlier by Kelly [18] whereby a user submits a single total payment, and the network determines both the rate allocations, as well as the divisions of the users' total payments among the links; in the single link case, this scheme reduces to that studied in Section 2. But Hajek and Yang [13] have shown that Kelly's mechanism can result in Nash equilibria in which the aggregate utility is an arbitrarily small fraction of the optimal aggregate utility. It remains an open problem whether there exists a network resource allocation mechanism in which each user submits a single number, representing total payment, and which has some nontrivial efficiency guarantees.

5.2 Multiple Producers, Inelastic Demand

The models presented thus far consider consumers competing for resources in scarce supply. 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., [31], Section 1-7.3. A basic model for electricity market operation involves *supply function bidding* [20]: 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.

Most previous work on supply function bidding has focused almost entirely on using the supply function equilibrium (SFE) framework of Klemperer and Meyer [20] for its predictive power. In such models, generators can submit nearly arbitrary supply functions; the Nash equilibria of the resulting game are used to give insight into expected behavior in current markets. In other words, by solving the SFE model for an appropriate set of assumptions, most previous work hopes to lend insight into the operation of power markets which require generators to submit complete supply schedules as bids [1, 6, 9, 10, 29, 33]. But because there may be a multiplicity of equilibria, an explicit understanding of efficiency losses in these games has not been developed. Papers such as the work of Rudkevich et al. [29] do suggest, however, that in the presence of inelastic demand, price anticipating behavior can lead to significant deviations from perfectly efficient allocations.

For this reason we take a different approach (see Chapter 4 of [15]). We consider *restrictions* on the supply functions which 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. Formally, we assume that each

firm n has a convex cost function C_n , as a function of the electricity generated. An efficient production vector minimizes the aggregate cost $\sum_n C_n(s_n)$, subject to the constraint that the total produced electricity $\sum_n s_n$ must equal the demand D . We then consider the following market. Each firm submits a supply function of the form $S(p, w) = D - w/p$, where D is the fixed (exogenous) demand and w is a nonnegative scalar chosen by the firm. The market then chooses a price so that aggregate supply is equal to demand.

If we assume that firms are price taking, it is possible to show that there exists a competitive equilibrium; furthermore, at this competitive equilibrium the resulting allocation minimizes aggregate production cost. If we assume instead that firms are price anticipating, we can establish existence of a Nash equilibrium and uniqueness of the resulting production vector, as long as more than two firms compete. Next, we consider the aggregate production cost at a Nash equilibrium relative to the minimal possible aggregate production cost. As long as more than two firms are competing, we show that the ratio of Nash equilibrium production cost to the minimal production cost is no worse than $1 + 1/(N - 2)$, where N is the number of firms in the market. Furthermore,, we demonstrate that this efficiency loss result carries over even to a setting where demand is inelastic but stochastically determined, by showing that in such an instance it is *as if* firms play a game with deterministic demand but different cost functions. Finally, a characterization theorem, similar to the one in Section 4, is also available, indicating that the mechanism under study has 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.

These results, which have been established in [15], suggest that market power can be controlled, and efficient allocations guaranteed, by restricting the supply functions available to generators in electricity markets. Restricted families of supply functions have also been considered elsewhere in the literature, e.g., in [3]. However, these models are typically used as approximations to unconstrained supply function bidding, and thus the resulting efficiency loss has not been studied. Still, this work leaves many open questions; in particular, the dynamics of power systems, together with their complex network structure, has not been captured in the models developed (in contrast to the telecommunications models previously discussed). Furthermore, the work described here depends on convexity assumptions on the cost functions of the producers, and such assumptions may generally not hold in electricity markets [14]. Finally, away from a Nash equilibrium, e.g., if some generators do not act rationally, the remaining generators may have to produce electricity at highly undesirable or even impossible levels. Addressing these types of questions is the subject of current research.

6 Open Issues

We have discussed the efficiency properties of Nash equilibria associated with certain resource allocation mechanisms. For the case where there is a single available resource (respectively, a single demand to be satisfied), the mechanisms involve the submission of a demand (respectively, a supply) function, which can be specified in terms of a single parameter, followed by market-clearing. In each case, we have provided a tight bound on the worst case efficiency loss. It remains to understand the worst case efficiency loss when mechanisms belonging to broader classes are considered. For example, in the context of Section 2 what efficiency guarantees are possible if users can choose a demand function from within a two-parameter family of demand curves?

Another research direction relates to the study of natural adjustment dynamics in the context of various mechanisms. Indeed, a desirable mechanism should not only have efficiency guarantees for the resulting Nash equilibria. It should also allow for simple adjustment algorithms whereby the different players can converge, in a stable manner, to such a Nash equilibrium.

7 Acknowledgment

This research was supported by the National Science Foundation under a Graduate Research Fellowship and grant ECS-0312921, by the Defense Advanced Research Projects Agency under the Next Generation Internet Initiative, and by the Army Research Office under grant DAAD10-00-1-0466. The material of Section 3 was reported in a paper jointly authored with Shie Mannor; the authors would also like to acknowledge helpful conversations with Frank Kelly.

References

1. E. J. Anderson and A. B. Philpott. Using supply functions for offering generation into an electricity market. *Operations Research*, 50(3):477–499, 2002.
2. L. M. Ausubel and P. Milgrom. The lovely but lonely Vickrey auction. 2004. In preparation.
3. R. Baldick, R. Grant, and E. P. Kahn. Theory and application of linear supply function equilibrium in electricity markets. *Journal of Regulatory Economics*, 2004. To appear.
4. E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:19–33, 1971.
5. J. R. Correa, A. S. Schulz, and N. Stier Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 2004. To appear.
6. C. J. Day, B. F. Hobbs, and J.-S. Pang. Oligopolistic competition in power networks: a conjectured supply function approach. *IEEE Transactions on Power Systems*, 17(3):597–607, 2002.

7. R. J. Gibbens and F. P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35:1969–1985, 1999.
8. J. R. Green and J.-J. Laffont. *Incentives in Public Decision-Making*. North-Holland Publishing Company, Amsterdam, The Netherlands, 1979.
9. R. J. Green. Increasing competition in the British electricity spot market. *Journal of Industrial Economics*, 44(2):205–216, 1996.
10. R. J. Green and D. M. Newbery. Competition in the British electricity spot market. *Journal of Political Economy*, 100(5):929–953, 1992.
11. T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.
12. B. Hajek and G. Gopalakrishnan. Do greedy autonomous systems make for a sensible Internet? Presented at the Conference on Stochastic Networks, Stanford University, 2002.
13. B. Hajek and S. Yang. Strategic buyers in a sum-bid game for flat networks. Submitted, 2004.
14. B. F. Hobbs, M. H. Rothkopf, L. C. Hyde, and R. P. O’Neill. Evaluation of a truthful revelation auction in the context of energy markets with nonconcave benefits. *Journal of Regulatory Economics*, 18(1):5–32, 2000.
15. R. Johari. *Efficiency loss in market mechanisms for resource allocation*. PhD thesis, Massachusetts Institute of Technology, 2004.
16. R. Johari, S. Mannor, and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game: the case of elastic supply. Publication 2605, MIT Laboratory for Information and Decision Systems, 2004.
17. R. Johari and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):407–435, 2004.
18. F. P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.
19. F. P. Kelly, A. K. Maulloo, and D. K. Tan. Rate control for communication networks: shadow prices, proportional fairness, and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
20. P. D. Klemperer and M. A. Meyer. Supply function equilibria in oligopoly under uncertainty. *Econometrica*, 57(6):1243–1277, 1989.
21. R. J. La and V. Anantharam. Charge-sensitive TCP and rate control in the Internet. In *Proceedings of IEEE INFOCOM*, pages 1166–1175, 2000.
22. R. T. Maheswaran and T. Basar. Nash equilibrium and decentralized negotiation in auctioning divisible resources. *Group Decision and Negotiation*, 12(5):361–395, 2003.
23. A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, Oxford, United Kingdom, 1995.
24. E. Maskin. Nash equilibrium and welfare optimality. *Review of Economic Studies*, 66(1):23–38, 1999.
25. C. Papadimitriou. Algorithms, games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing*, pages 749–753, 2001.
26. J. Rosen. Existence and uniqueness of equilibrium points for concave n -person games. *Econometrica*, 33(3):520–534, 1965.
27. M. H. Rothkopf, T. J. Teisberg, and E. P. Kahn. Why are Vickrey auctions rare? *Journal of Political Economy*, 98(1):94–109, 1990.
28. T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.

29. A. Rudkevich, M. Duckworth, and R. Rosen. Modeling electricity pricing in a deregulated generation industry: the potential for oligopoly pricing in a poolco. *Energy Journal*, 19(3):19–48, 1998.
30. S. Sanghavi and B. Hajek. Optimal allocation of a divisible good to strategic buyers. 2004. Submitted.
31. S. Stoft. *Power System Economics: Designing Markets for Electricity*. IEEE Press, Piscataway, New Jersey, 2002.
32. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.
33. N.-H. M. von der Fehr and D. Harbord. Spot market competition in the U.K. electricity industry. *Economic Journal*, 103(418):531–546, 1993.
34. S. Yang and B. Hajek. An efficient mechanism for allocation of a divisible good and its application to network resource allocation. 2004. In preparation.