

Parameterized Supply Function Bidding: Equilibrium and Efficiency

Ramesh Johari

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

John N. Tsitsiklis

Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology,
Cambridge, Massachusetts 02139, jnt@mit.edu

We consider a model where a finite number of producers compete to meet an infinitely divisible but *inelastic* demand for a product. Each firm is characterized by a production cost that is convex in the output produced, and firms act as profit maximizers. We consider a uniform price market design that uses *supply function bidding*: firms declare the amount they would supply at any positive price, and a single price is chosen to clear the market. We are interested in evaluating the impact of price-anticipating behavior both on the allocative efficiency of the market and on the prices seen at equilibrium. We show that by *restricting* the strategy space of the firms to parameterized supply functions, we can provide upper bounds on both the inflation of aggregate cost at the Nash equilibrium relative to the socially optimal level, as well as the markup of the Nash equilibrium price above the competitive level: as long as $N > 2$ firms are competing, these quantities are both upper bounded by $1 + 1/(N - 2)$. This result holds even in the presence of asymmetric cost structure across firms. We also discuss several extensions, generalizations, and related issues.

Subject classifications: supply function equilibrium; resource allocation; efficiency loss.

Area of review: Revenue Management.

History: Received October 2008; revisions received June 2009, November 2009; accepted January 2010.

1. Introduction

We consider a model where a finite number of producers compete to meet an infinitely divisible, inelastic demand for a product. Each firm is characterized by a production cost that is convex in the output produced, and firms act as profit maximizers. We study a simple market design question: given a fixed, inelastic demand, how should a market mechanism be designed to yield an efficient allocation of production across suppliers—that is, an allocation that minimizes production cost?

In this paper, we focus our attention on *uniform price market-clearing mechanisms* for the allocation problem. These are mechanisms that set a single per-unit price for the resource; firms compete by submitting *supply functions* that describe their desired production level as a function of price. A central clearinghouse then chooses a price that clears the market. Such mechanisms have been used to model competition in a range of industries, including energy markets, airline pricing, and contracts for management consulting services (Vives 2009). Uniform price mechanisms are interesting objects of study due to their simplicity. In particular, they are transparent and fair from the point of view of market participants: all agents are charged the same per-unit price and asked to supply exactly what they bid via their supply functions.

We have two related goals in this market design problem. First, we desire that such a mechanism does not exhibit a large “welfare loss”; i.e., we hope that the efficiency loss remains bounded when firms are price anticipating, regardless of the firms’ cost functions. Second, we wish to ensure that the price markup when firms are price anticipating is bounded relative to the competitive price level; such a bound ensures that the exercise of market power by the firms is mitigated.

Our task is complicated by a fundamental trade-off in mechanism design. On one hand, sufficient flexibility must be granted to the firms in declaring their supply functions to ensure that they can approximately declare their costs. On the other hand, as the strategic flexibility granted to firms increases, their temptation to misdeclare their cost increases as well. Indeed, although in principle arbitrary supply functions allow firms to declare all marginal cost information, in theory and practice we find that such strategic flexibility only encourages the exercise of market power.

Our paper sheds light on this trade-off by studying a parameterized class of supply functions that allow firms enough flexibility to communicate information about their production cost, yet not enough flexibility to enable them to exercise market power and cripple the performance of the overall market. In other words, we partially restrict the range of possible supply functions firms can declare

and demonstrate that the resulting market-clearing mechanism is nearly efficient. Our analysis lends credence to the hypothesis that *restricting the strategy space granted to firms can improve allocative efficiency*, as observed by several pieces of related work (see §2).

Before proceeding, we fix some terminology. We will assume that N firms compete to satisfy a fixed demand $D > 0$, where firm n has a convex, strictly increasing, and nonnegative production cost function $C_n(s_n)$. We assume that each firm n submits a supply function $S_n(p)$ to a central clearinghouse. The clearinghouse then clears the market by choosing a price p^* such that $\sum_n S_n(p^*) = D$, and firm n is asked to supply $S_n(p^*)$. For each firm, we also let $P_n(S)$ denote the inverse supply function, i.e., $P_n(S_n(p)) = p$. (These definitions are made informally, without regard to ensuring that market-clearing prices or inverse supply functions exist; we will consider these technical issues more carefully in the remainder of the paper.) We may reframe our objective as follows: are there any restrictions that can be placed on supply functions firms are allowed to submit that ensure both efficient market performance and a bounded markup above the competitive price level?

One can start by considering simple Bertrand and Cournot structures for the supply functions. Bertrand competition has the deficiency that equilibria may fail to exist when the marginal production cost of each firm is not linear (Shapiro 1989). On the other hand, Cournot competition is not well defined when the elasticity of demand is zero; and furthermore, if the price elasticity of demand is low, then it is straightforward to check that Cournot equilibria may have arbitrarily high welfare loss and price markup above competitive levels (Day et al. 2002).

Thus, neither perfectly vertical nor perfectly horizontal supply functions yield reasonable solutions in this setting. One is then led to consider *supply function equilibrium*. In such a model, the strategy of each firm is not limited to one scalar (either price or quantity), but rather consists of an entire function $S_n(p)$ describing the amount of the good a firm is willing to produce at any price p . The seminal work in the study of supply function equilibria is the paper of Klemperer and Meyer (1989); for further details, see §2. Unfortunately, one lesson of that line of literature is that the SFE framework is problematic from a market design standpoint; in general, there may exist highly inefficient equilibria.

In the remainder of the paper, for the resource allocation environment described above, we demonstrate that we can achieve a successful mechanism design by properly restricting the class of supply functions that firms are allowed to submit. We start by discussing related work in §2. In §3, we precisely define the market mechanism we consider; in particular, we assume that each firm submits a supply function of the form $S(p, w) = D - w/p$, where D is the demand and w is a nonnegative scalar chosen by the firm. The parameter w can be interpreted as the amount of revenue the supplier is willing to *forgo* from the total payout

pD that will be created when the market clears. The market manager then chooses a price so that aggregate supply is equal to demand. This is, of course, a somewhat unintuitive market design; however, our goal is to use this design as a vehicle to demonstrate the strong efficiency properties attainable when strategic flexibility is properly constrained.

We begin with a preliminary investigation of equilibria of the mechanism. For our mechanism we recover the fundamental theorem of welfare economics: when firms are price taking, there exists a competitive equilibrium, and the resulting allocation minimizes aggregate production cost. We next assume instead that firms are price anticipating, and establish existence and uniqueness of a Nash equilibrium as long as more than two firms compete.

Sections 4 and 5 present the key results of this paper: the former provides a theoretical bound on the welfare loss at the Nash equilibrium relative to the competitive outcome, and the latter provides a similar bound on the price markup at the Nash equilibrium relative to the competitive price. In §4 we show that as long as at least two firms are competing, 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. We emphasize that this result holds regardless of the cost functions of the firms (as long as they are convex)—thus, it provides a very strong competitive limit theorem without any assumptions of symmetry between firms. In §5, we show that the same bound applies to the Nash equilibrium price relative to the competitive price; we also provide a bound on the *Lerner index*. All bounds discussed in these two sections are tight.

In §6 we ask a design question: to what extent are the parameterized supply functions we have chosen “optimal?” We study a class of reasonable parameterized supply function mechanisms, using the metric of *worst-case welfare loss* at the Nash equilibrium. We note that the mechanism we have chosen is “optimal” in this sense. This result is closely related to a theorem derived by the authors in Johari and Tsitsiklis (2009) and first presented in (Johari 2004, Chapter 4).

We conclude by considering two extensions of the model, first to mitigate the possibility of negative supply by the firms (discussed in §7), and second to cover settings with stochastic demand (in §EC.2). Section 7 addresses the fact that the supply functions we have chosen allow for nonequilibrium outcomes in which a firm may be forced to buy the good rather than supply it. We show that this problem is generic to the design of parameterized supply function mechanisms, but we also provide a simple resolution using a “maximum liability” guarantee to each firm. Section 8 concludes. *Note:* Due to space constraints, all proofs can be found in the e-companion, which is available as part of the online version that can be found at <http://or.journal.informs.org/>.

2. Related Work

Our work is inspired by a recent line of literature that studies the efficiency guarantees possible in market design when the

strategy spaces of market participants are restricted (Chen and Zhang 2010; Johari et al. 2005; Johari and Tsitsiklis 2006; Maheswaran 2003; Maheswaran and Basar 2004; Moulin 2008; Sanghavi and Hajek 2004; Yang and Hajek 2006, 2007; Dimakis et al. 2006; Stoenescu and Ledyard 2006; Moulin 2010). The results of Johari and Tsitsiklis (2009) are most closely related to our paper: in that work, the efficiency of scalar-parameterized mechanisms is studied for a setting where buyers bid for a capacitated resource. As we discuss in §3, the mechanism we consider is related to the mechanisms studied in this literature; however, the efficiency loss bounds we derive are novel. In large part this distinction is due to the fact that in our paper we consider a setting where suppliers compete to meet demand. This is part of a growing literature that quantifies efficiency losses in a variety of game-theoretic environments; see Nisan et al. (2007) for a comprehensive survey.

One well-known approach to the market design problem is to use the Vickrey-Clarke-Groves (VCG) class of mechanisms (Clarke 1971, Groves 1973, Vickrey 1961). Because each firm's profit is quasilinear (i.e., the production costs are measured in monetary units), a VCG mechanism ensures that truthful reporting for each firm is a dominant strategy. However, there are several reasons why a VCG mechanism may not be desirable in practical settings. For example, there is no bound on the payment the market manager may have to make to the market participants; further, VCG mechanisms exhibit the implicit "unfairness" of providing a different price to different purveyors of the same good. See, e.g., Hobbs et al. (2000), Ausubel and Milgrom (2006), and Rothkopf et al. (1990) for extensive discussion of some of the shortcomings of the VCG mechanism. In part due to these shortcomings, VCG mechanisms are rarely observed in complex multiunit resource allocation settings such as power and electricity markets. It is worth noting that several recent papers have studied approaches to pricing divisible resources using VCG-like mechanisms with *scalar* strategy spaces; see, e.g., Maheswaran and Basar (2004), Yang and Hajek (2006, 2007), Johari and Tsitsiklis (2009). Similar approaches could be applied in our context to yield efficient or nearly efficient market mechanisms, although with attendant shortcomings analogous to standard VCG mechanisms.

Instead, as discussed in the introduction, we focus on uniform price market-clearing mechanisms; this model is closely related to the analysis of supply function equilibria (SFE). Grossman (1981) and Hart (1985) provide concrete examples of SFE models. In particular, Grossman's analysis shows that in the presence of fixed startup costs to the firms, it is possible for a supply function equilibrium to achieve full efficiency; however, in general it is difficult to guarantee that the number of supply function equilibria is small, and other inefficient supply function equilibria may exist. The seminal work in the study of supply function equilibria is the paper of Klemperer and Meyer (1989).

The authors begin by showing that in the absence of uncertainty, nearly any production allocation can be supported as a supply function equilibrium. They then show that if demand is uncertain, then the range of equilibria is dramatically reduced; and that in equilibrium, possible prices and allocations range between those achieved at Bertrand and Cournot equilibria.

The SFE framework is somewhat problematic from a market design standpoint. The original model of Klemperer and Meyer (1989) required that the different firms have identical cost functions; recent work has made progress in studying models with asymmetric firms with both affine (Baldick and Hogan 2001, 2006; Baldick et al. 2004) and nonlinear (Anderson and Hu 2008) supply functions. This literature primarily focuses on computational approaches to finding SFE (when they exist); as a result, general bounds on efficiency and price markups are not typically available—such properties are evaluated on a case-by-case basis. In summary, therefore, the complexity of the SFE model places restrictions on the types of environments that can be successfully analyzed, and does not yield a satisfactory answer to the market design questions raised above.

SFE models can be used to model a range of industries, including airline pricing and contracts for management consulting (Vives 2009), but the most prominent application of the SFE concept is to electricity markets. Many power markets actually operate in practice by having generators submit complete supply functions (see Green 1996 and Green and Newbery 1992 for the first applications of this approach). We do not aim to provide a comprehensive survey of the electricity market literature here; the reader is referred to Day et al. (2002) for a more complete list of references, and Wilson (2002) for an elegant discussion of some of the issues involved in power market design. We conclude by noting that electricity markets also typically exhibit inelastic demand, as assumed in our paper; see, e.g., Stoft (2002, §§1–7.3). In light of the short-run price inelasticity of demand, many short-term markets for generation today operate by setting a price for electricity so that the aggregate supply offered by generators meets the demand requirements of a given region.

As noted above, the SFE model of Klemperer and Meyer, and subsequent results on that model, treat a general model of supply function bidding where demand is stochastic. As a result, such models apply well to markets where accurate demand forecasts are not available when bids are submitted, or where demand may vary over the period where the bid is binding. For example, in some day-ahead electricity markets, generators are required to submit a single bid for the entire day, over which demand varies due to time-of-day effects. By contrast, our results on efficiency are developed in a setting with deterministic demand, and thus are applicable mainly to markets where demand forecasts are accurate at the time that bids are submitted. Such a model is reasonable, for example, in electricity markets where generators are allowed to rebid a different supply function for

each hour of the day, if we presume that fairly accurate hourly forecasts of demand are available in advance.

3. Preliminaries

We consider a model where $N \geq 2$ firms compete to satisfy an inelastic demand $D > 0$. Note that our model assumes that the demand D is deterministically known; we believe this is a reliable assumption in short-term markets (e.g., day-ahead), where prediction of demand based on historical models is likely to have very low variance. We consider a model where demand is stochastic in §EC.2 in the e-companion.

Let s_n denote the amount produced and supplied by firm n . We assume that firm n incurs a cost $C_n(s_n)$ when it produces s_n units; we assume that cost is measured in monetary units, and firms are profit maximizers. We make the following assumption on the cost functions C_n .

ASSUMPTION 1. For each n , the cost function $C_n(s_n)$ is continuous, with $C_n(s_n) = 0$ if $s_n \leq 0$. Over the domain $s_n \geq 0$, the cost function $C_n(s_n)$ is convex and strictly increasing.

Because demand is inelastic, it is clear that aggregate welfare maximization is equivalent to aggregate cost minimization, i.e., the following optimization problem:

$$\text{minimize } \sum_n C_n(s_n) \quad (1)$$

$$\text{subject to } \sum_n s_n = D; \quad (2)$$

$$s_n \geq 0, \quad n = 1, \dots, N. \quad (3)$$

Any solution \mathbf{s} to (1)–(3) is referred to as *efficient*.

We consider the following market mechanism for production allocation. Each firm n submits a *supply function* to the market manager, which gives (as a function of price) the amount the firm is willing to produce. In contrast to much of the literature on supply function equilibria (e.g., Klemperer and Meyer 1989), we consider the implications of restricting the space of supply functions that firms are allowed to choose from; in particular, we will assume that the supply functions are chosen from a parameterized family of supply functions.

Formally, the details of our market mechanism are as follows. We assume that firm n submits a parameter $w_n \geq 0$ to the market manager. The parameter indicates that at a price $p > 0$, firm n is willing to supply $S(p, w_n)$ units given by:

$$S(p, w_n) = D - \frac{w_n}{p}. \quad (4)$$

We then assume that the market manager chooses the price $p(\mathbf{w}) > 0$ to *clear the market*, i.e., so that $\sum_s S(p(\mathbf{w}), w_n) = D$. Such a choice is only possible if $\sum_n w_n > 0$, in which case:

$$p(\mathbf{w}) = \frac{\sum_n w_n}{(N-1)D}. \quad (5)$$

On the other hand, if $\sum_n w_n = 0$, then $S(p, w_n) = D$ for all n regardless of the value of p ; therefore, we fix the

following conventions:

$$S(0, 0) = D, \quad \text{and} \quad p(\mathbf{0}) = 0. \quad (6)$$

(This makes the function p continuous in \mathbf{w} .)

The parameter w_n may be interpreted as the revenue that firm n is willing to *forgo*; this follows because pD is the total pool of revenue when the price is p , and $pS(p, w_n) = pD - w_n$ is the revenue to firm n when the price is p . It is straightforward to observe various peculiarities of this mechanism. First, it does not make any provision for firms to submit capacity constraints; we will find that this does not affect equilibrium behavior, but of course out of equilibrium the inability to declare capacity could force a firm to deliver supply beyond its means. On the other hand, any static supply function bidding model (with deterministic demand) where firms are allowed to declare capacities will have a range of highly undesirable equilibria, where firms choose capacities that exactly total the desired demand. In such equilibria, firms choose supply functions that approach capacity as the price approaches infinity. As a result, these equilibria always have astronomical prices, and every measure of market performance can be shown to degrade in such a context.¹ We believe that extracting capacity information from firms in such settings will require a more complete dynamic model of market structure, and this remains an important research direction.

A second concern regarding the market mechanism we have designed here is that prospective suppliers may be forced to *purchase* the good at the outcome of the market. This can happen when \mathbf{w} is chosen so that $D - w_n/p(\mathbf{w}) < 0$ for some firm n ; nothing rules this out in the definition of the mechanism. Of course, such behavior cannot happen in equilibrium, but firms will be rightly nervous about agreeing to a market mechanism with such a property. We will discuss in §6 two responses to this issue. We will also provide a characterization of the mechanism we have chosen as the best-possible mechanism available in a certain class with reasonable properties.

We begin by considering a setting where firms act as price takers: given a price $\mu > 0$, a price-taking firm n acts to maximize the following profit function over $w_n \geq 0$:

$$P_n(w_n; \mu) = \mu S(\mu, w_n) - C_n(S(\mu, w_n)). \quad (7)$$

The first term represents the revenue to firm n when the price is μ and the firm supplies $S(\mu, w_n)$ units; the second term represents the cost to the firm of producing $S(\mu, w_n)$ units. Observe that because cost is measured in monetary units, the payoff is *quasilinear* in money.

It is straightforward to show that under our assumptions, when firms are price takers, there exists a *competitive equilibrium*, and the resulting allocation is efficient (i.e., an optimal solution to (1)–(3)). In a competitive equilibrium, firms maximize their payoff, and the price is chosen according to (5) to clear the market. For details on this result, see Theorem EC.1 in the e-companion.

In contrast to the price-taking model, we now consider an oligopoly model where the firms are price anticipating instead. Price-anticipating firms will realize that μ is set according to $\mu = p(\mathbf{w})$ from (5), and adjust their payoff accordingly. We use the notation \mathbf{w}_{-n} to denote the vector of strategies of firms other than n ; i.e., $\mathbf{w}_{-n} = (w_1, w_2, \dots, w_{n-1}, w_{n+1}, \dots, w_N)$. Given \mathbf{w}_{-n} , each firm n chooses w_n to maximize:

$$Q_n(w_n; \mathbf{w}_{-n}) = p(\mathbf{w})S(p(\mathbf{w}), w_n) - C_n(S(p(\mathbf{w}), w_n)) \quad (8)$$

over nonnegative w_n . If we substitute for $p(\mathbf{w})$ from (5) and for $S(p, w_n)$ from (4), we have:

$$Q_n(w_n; \mathbf{w}_{-n}) = \begin{cases} \frac{\sum_m w_m}{N-1} - w_n - C_n \\ \cdot \left(D - \left(\frac{w_n}{\sum_m w_m} \right) (N-1)D \right), & \text{if } w_n > 0; \\ \frac{\sum_{m \neq n} w_m}{N-1} - C_n(D), & \text{if } w_n = 0. \end{cases} \quad (9)$$

The payoff function Q_n is similar to the payoff function P_n , except that the firm anticipates that the network will set the price μ according to $\mu = p(\mathbf{w})$ from (5). The following theorem shows that there exists a unique Nash equilibrium allocation when $N > 2$ firms compete, by showing that at a Nash equilibrium it is as if the firms are solving another optimization problem of the same form as the aggregate cost minimization problem (1)–(3), but with “modified” cost functions.

THEOREM 1. *Assume that $N \geq 2$, and suppose that Assumption 1 is satisfied. If $N = 2$, then no Nash equilibrium exists for the game defined by (Q_1, \dots, Q_N) . On the other hand, if $N > 2$, then there exists a Nash equilibrium $\mathbf{w} \geq 0$ of the game defined by (Q_1, \dots, Q_N) , and it satisfies $\sum_n w_n > 0$. For any Nash equilibrium \mathbf{w} , the vector \mathbf{s} defined by $s_n = S(p(\mathbf{w}), w_n)$ is the unique optimal solution to the following optimization problem:*

$$\text{minimize } \sum_n \hat{C}_n(s_n) \quad (10)$$

$$\text{subject to } \sum_n s_n = D; \quad (11)$$

$$s_n \geq 0, \quad n = 1, \dots, N, \quad (12)$$

where

$$\hat{C}_n(s_n) = \left(1 + \frac{s_n}{(N-2)D} \right) C_n(s_n) - \frac{1}{(N-2)D} \int_0^{s_n} C_n(z) dz. \quad (13)$$

We note the following corollary for later reference: at a Nash equilibrium, one firm produces the entire supply only if it was efficient to do so.

COROLLARY 1. *Assume that $N > 2$, and suppose that Assumption 1 is satisfied. Suppose that \mathbf{w} is a Nash equilibrium $\mathbf{w} \geq 0$ of the game defined by (Q_1, \dots, Q_N) such that $S(p(\mathbf{w}), w_n) = D$, and $S(p(\mathbf{w}), w_m) = 0$ for $m \neq n$. Then the Nash equilibrium production vector is efficient.*

We note that the use of “modified” cost functions in the proof of Theorem 1 is similar to the use of *potential functions* to study Nash equilibria in some games (Monderer and Shapley 1996). However, the objective function (10) is not a potential function for the game defined by (Q_1, \dots, Q_N) : whereas each payoff function is determined by the vector \mathbf{w} of *strategic choices* of other players, the objective function (10) depends only on the resulting production vector \mathbf{s} . As a result, it is straightforward to check that (10) is neither an ordinal nor exact potential function.

The most closely related results concern the *proportional allocation mechanism* for allocation of a single divisible good among strategic buyers, studied by Kelly (1997), La and Anantharam (2000), Hajek and Gopalakrishnan (2002), Maheswaran and Basar (2003), and Johari and Tsitsiklis (2004). Observe that the quantity $w_n/p(\mathbf{w})$ is the “relief” provided to firm n relative to the total demand D , because firm n is asked to supply $D - w_n/p(\mathbf{w})$ when the market clears. Suppose we define a virtual divisible good of $(N-1)D$ units that we interpret as the total relief that will be granted to firms. At the market-clearing price, if $w_n > 0$, firm n ’s relief is:

$$\frac{w_n}{p(\mathbf{w})} = \frac{w_n}{\sum_m w_m} (N-1)D.$$

If $w_n = 0$, then firm n receives no relief: it must supply the entire demand D . Thus, total relief is allocated to the firms in proportion to their bids. Further, we can define a utility function $U_n(x_n)$ for firm n as a function of the relief, by simply negating the cost at the supply $D - x_n$:

$$U_n(x_n) = -C_n(D - x_n).$$

(Recall that $C_n(s_n) = 0$ for $s_n \leq 0$.) The payoff to firms when they are price anticipating may then be written:

$$Q_n(w_n; \mathbf{w}_{-n}) = \begin{cases} U_n \left(\frac{w_n}{\sum_m w_m} (N-1)D \right) \\ - \left(1 - \frac{1}{N-1} \right) w_n - h \left(\sum_{m \neq n} w_m \right), & \text{if } w_n > 0; \\ U_n(0) - h \left(\sum_{m \neq n} w_m \right), & \text{if } w_n = 0, \end{cases}$$

where $h_n(W) = W/(N-1)$. Note that the last term in the utility does not depend on firm n ’s bid.

Thus, the game firms play when they are price anticipating is strategically equivalent to one where firm n submits a bid w_n ; receives a share $w_n/(\sum_m w_m)$ of the total relief $(N - 1)D$ if $w_n > 0$, and receives zero relief if $w_n = 0$; and pays $(1 - 1/(N - 1))w_n$. This is essentially a proportional allocation mechanism, except that the payment made is scaled by the constant $1 - 1/(N - 1)$. Our proof technique for existence and uniqueness of competitive and Nash equilibria is therefore essentially equivalent to similar proofs for the proportional allocation mechanism (such as those by Hajek and Gopalakrishnan 2002 and Maheswaran and Basar 2003). Our presentation is most closely related to Johari and Tsitsiklis (2004), where existence and uniqueness of Nash equilibria is established via the use of modified utility functions.

Despite the close connection of our supply function bidding game and the proportional allocation mechanism, efficiency loss results for the proportional allocation mechanism from Johari and Tsitsiklis (2004) cannot be directly translated to efficiency loss analysis of our supply function bidding game. Note that the utility functions defined above are negative, whereas positivity is essential to the efficiency loss analysis by Johari and Tsitsiklis (2004). A strategically equivalent positive utility function is $U_n(x_n) = C_n(D) - C_n(D - x_n)$. However, the ratio of Nash equilibrium aggregate utility to the maximum possible aggregate utility is *not* directly related to the ratio of Nash equilibrium aggregate cost to the minimum possible aggregate cost. This difference accounts for differing efficiency loss ratios in our subsequent development. Further, our efficiency loss analysis relies heavily on the specific form of the modified cost function \hat{C}_n ; this function is not equivalent to the modified utility function used in the proof of existence and uniqueness of Nash equilibria for the proportional allocation mechanism.

We conclude by discussing a brief example to illustrate computation of the equilibrium using Theorem 1. Our emphasis is on both the level of inefficiency possible at the equilibrium, as well as the worst-case markup above the competitive price level that may be observed in equilibrium; in subsequent sections we study these quantities analytically.

EXAMPLE 1. Assume that N firms compete to meet demand $D = 1$. We assume that all firms have linear cost functions: $C_n(s_n) = \alpha_n s_n$, where $\alpha_n > 0$. In this case, a straightforward calculation reveals that:

$$\hat{C}_n(s_n) = \alpha_n s_n \left(1 + \frac{s_n}{2(N - 2)} \right).$$

A key fact we exploit is provided by Step 7 of the proof of Theorem 1: if \mathbf{s} is an optimal solution to (10)–(12), with Lagrange multiplier ρ corresponding to the constraint (11), and we define $w_n = (D - s_n)\rho$ for all n , then \mathbf{w} is a Nash equilibrium of the game with allocation \mathbf{s} and price $p(\mathbf{w}) = \rho$. Indeed, as seen in Step 8 of the proof of

Theorem 1, if $0 \leq s_n < D$ for all n , then all Nash equilibria are obtained in this way.

We first consider a scenario where $\alpha_n = \alpha > 0$ for all firms n . Using Theorem 1, by symmetry we conclude that at the Nash equilibrium, all firms n produce the same quantity $s_n = 1/N$; further, it follows that the Lagrange multiplier corresponding to the constraint (11) must be:

$$\rho = \alpha \left(1 + \frac{1}{N(N - 2)} \right).$$

If we let $w_n = (1 - s_n)\rho$, then \mathbf{w} is a Nash equilibrium with allocation \mathbf{s} and price $p(\mathbf{w}) = \rho$. In particular, observe in this case that the price markup above the competitive price is of order $1/N^2$; of course, the allocation is efficient because all firms have the same cost function.

Now suppose that, instead, one firm has cost coefficient $\alpha_1 = \alpha > 0$, whereas all other firms $n > 1$ have cost coefficient $\alpha_n = \beta > \alpha$; thus, firm 1 is a “dominant” firm. We consider two possibilities: either (1) $\beta \geq \alpha(1 + 1/(N - 2))$, or (2) $\beta \leq \alpha(1 + 1/(N - 2))$.

In the first case, it follows that the optimal solution to (10)–(12) has $s_1 = D$, and $s_2, \dots, s_N = 0$. Let μ be any value such that $\alpha(1 + 1/(N - 2)) \leq \mu \leq \beta$. If we define $w_1 = 0$ and $w_2, \dots, w_N = \mu D$, it can be verified that \mathbf{w} is a Nash equilibrium, and the resulting price is μ . Note that at any equilibrium the allocation is efficient in this case (as firm 1 produces the entire demand), but the worst-case price markup above the efficient level is bounded only by the cost coefficient β of the inefficient firms, and does not necessarily decrease with N . Nevertheless, the best-case price markup $(1 + 1/(N - 2))$ approaches unity as $N \rightarrow \infty$.

On the other hand, suppose that $\beta \leq \alpha(1 + 1/(N - 2))$. Then all firms have $s_n > 0$ in the optimal solution to (10)–(12). It can be verified that in this case, the Lagrange multiplier corresponding to constraint (11) is:

$$\rho = \alpha \left(1 + \frac{1}{N - 2} \right) \left(\frac{N - 1}{1 + (N - 1)\alpha/\beta} \right).$$

Because $s_n > 0$ for all n , it follows from Step 8 of the proof of Theorem 1 that ρ is the price in the unique Nash equilibrium of the game. In particular, observe that the price markup above the competitive level is of order $1/N$ in this case. Finally, it can be shown that in the Nash equilibrium:

$$s_1 = (N - 1)^2 \left(\frac{1}{1 + (N - 1)\alpha/\beta} \right) - N + 2.$$

In general, therefore, the allocation in this case is inefficient. However, as $\beta/\alpha \rightarrow 1 + 1/(N - 2)$, the demand served by firm 1 approaches unity.

4. Welfare Loss

We let \mathbf{s}^* denote an efficient production vector and let \mathbf{s} denote the production vector at a Nash equilibrium.

We now ask: what is the welfare loss at the Nash equilibrium? To answer this question, we must compare the cost $\sum_n C_n(s_n)$ with the cost $\sum_n C_n(s_n^*)$. (We know, of course, that $\sum_n C_n(s_n) \geq \sum_n C_n(s_n^*)$ by definition of s^* .) The following theorem provides an explicit bound for the welfare loss.

THEOREM 2. *Assume that $N > 2$, and suppose that Assumption 1 is satisfied. If s^* is any efficient production vector, and s is the production vector at a Nash equilibrium, then,*

$$\sum_n C_n(s_n) \leq \left(1 + \frac{1}{N-2}\right) \sum_n C_n(s_n^*). \quad (14)$$

Furthermore, this bound is tight: for every $\epsilon > 0$ and $N > 2$, there exists a choice of cost functions C_n , $n = 1, \dots, N$, such that:

$$\sum_n C_n(s_n) \geq \left(1 + \frac{1}{N-2} - \epsilon\right) \sum_n C_n(s_n^*). \quad (15)$$

The preceding theorem shows that in the worst case, aggregate cost rises by no more than a factor $1 + 1/(N-2)$ when firms are price anticipating. Furthermore, this bound is essentially tight. We note that the increase in aggregate production cost, a factor of $1/(N-2)$, approaches zero as the number of firms N grows large, even though the firms are price anticipating. This is a form of a competitive limit theorem (Mas-Colell et al. 1995). Competitive limit theorems have been extensively studied, especially in the context of strategic market games (Giraud 2003, Shapley and Shubik 1977, Dubey and Shubik 1978). However, such results typically use a *replication* approach: starting from K firms, an economy of NK firms is created by replicating each firm N times. As a result, no single firm can be dominant in the limit. By contrast, our result holds even if only a small number of firms continue to remain dominant as $N \rightarrow \infty$; i.e., we do not require any symmetry constraints on the cost functions of the firms. In general, in an industry with one large firm and many small firms, we do not expect to achieve full efficiency; nevertheless, the mechanism described in this paper ensures that this is the case.

In Johari and Tsitsiklis (2004), for the proportional allocation of a divisible good among strategic buyers, it is shown that the worst-case efficiency loss when buyers are price anticipating is 25% of the maximum possible aggregate utility. The worst-case efficiency loss there is obtained when one buyer receives a constant fraction of the resource, and the remainder of the resource is split into vanishingly small slices among a large collection of buyers. However, in our setting, the total “relief” available to firms scales as $(N-1)D$, and each firm can only obtain a relief at most equal to D (cf. the discussion in §3). This feature of the mechanism intuitively limits the efficiency loss one firm can create, and provides insight into the decay of efficiency loss to zero as $N \rightarrow \infty$.

5. Price Bounds

In this section we prove two upper bounds on the Nash equilibrium price; because the demand is fixed and inelastic, these will also be upper bounds on the revenue to the firms at a Nash equilibrium. We restrict attention to settings where at least two firms are active at the Nash equilibrium; by Corollary 1, this is guaranteed if any efficient allocation calls for at least two active producers.

The following theorem compares the price at a Nash equilibrium to the price at a competitive equilibrium.

THEOREM 3. *Assume that $N > 2$, and suppose that Assumption 1 is satisfied. Suppose that $\mathbf{w} \geq 0$ is a Nash equilibrium of the game defined by (Q_1, \dots, Q_N) such that $S(p(\mathbf{w}), w_n) > 0$ and $S(p(\mathbf{w}), w_m) > 0$ for at least two firms m, n , $m \neq n$. Let \mathbf{w}^* and μ^* denote a competitive equilibrium: Then,*

$$p(\mathbf{w}) \leq \left(1 + \frac{1}{N-2}\right) \mu^*.$$

The preceding theorem shows that the Nash equilibrium price is no more than a factor $1 + 1/(N-2)$ above the competitive price level. Note that because the total payment to the firms is $p(\mathbf{w})D$ at the Nash equilibrium, and μ^*D at the competitive equilibrium, the same bound holds for the total payment made to the firms.

We can use a similar approach to bound the *Lerner index*, which we define as follows at a strategy vector \mathbf{w} :

$$L(\mathbf{w}) = \max_r \left[\frac{p(\mathbf{w}) - \partial C_r^+(S(p(\mathbf{w}), w_r))/\partial s_r}{p(\mathbf{w})} \right].$$

(We have chosen to define the Lerner index using the right directional derivative of the cost; of course, if the cost functions are differentiable, this choice is inconsequential.) The Lerner index is commonly used to measure the price markup above competitive levels in oligopolies. We bound the index in the following corollary.

COROLLARY 2. *Assume that $N > 2$, and suppose that Assumption 1 is satisfied. Suppose that $\mathbf{w} \geq 0$ is a Nash equilibrium of the game defined by (Q_1, \dots, Q_N) such that $S(p(\mathbf{w}), w_n) > 0$ and $S(p(\mathbf{w}), w_m) > 0$ for at least two firms m, n , $m \neq n$. Then,*

$$L(\mathbf{w}) \leq \frac{1}{N-1}.$$

The bounds of this section, together with the welfare loss bounds of the previous section, strongly characterize the performance of the mechanism we have proposed. We note that as in the comment at the end of the preceding section, our bounds on prices show convergence to the competitive price level as the number of firms approaches infinity. This limit holds even in the absence of symmetry between the firms.

6. A Characterization Theorem

In this section, we provide a positive characterization of the market mechanism studied in this paper. Our presentation is inspired by a related result derived for a mechanism where consumers submit demand functions to a market that allocates a resource with inelastic supply, as reported in Johari and Tsitsiklis (2009). In that paper, we showed that among all mechanisms satisfying certain reasonable assumptions, the bidding mechanism considered minimizes the worst-case efficiency loss at a Nash equilibrium. In this paper, we present a corresponding result that demonstrates that within a reasonable class of market mechanisms that use parameterized supply functions, the one studied in this paper yields the lowest-possible worst-case welfare loss at Nash equilibrium. Because the argument is closely related to that in Johari and Tsitsiklis (2009), we omit the details of the proof; they can be found in Johari (2004). Nevertheless, as we describe below, this characterization result differs from the result obtained in Johari and Tsitsiklis (2009).

For the purposes of this section, we let \mathcal{C} denote the set of all cost functions that satisfy Assumption 1; i.e.,

$$\mathcal{C} = \{C: \mathbb{R} \rightarrow \mathbb{R}^+ \mid C \text{ is continuous on } \mathbb{R}, \text{ strictly increasing, and convex on } \mathbb{R}^+, \text{ and } C(s) = 0 \text{ for } s \leq 0\}. \tag{16}$$

We begin by defining the class of market mechanisms we will study in this section.

DEFINITION 1. Given $D > 0$ and $N > 1$, the class $\mathcal{S}(D, N)$ consists of all differentiable functions $S: (0, \infty) \times [0, \infty) \rightarrow \mathbb{R}$ such that:

1. For all nonzero $\mathbf{w} \in (\mathbb{R}^+)^N$, there exists a unique market-clearing price, i.e., a unique solution $p > 0$ to the following equation:

$$\sum_{n=1}^N S(p, w_n) = D.$$

We let $p_S(\mathbf{w})$ denote this solution.

2. For all $C \in \mathcal{C}$, a firm’s payoff is concave if the firm is price taking; that is, for all $p > 0$ the function

$$w \mapsto pS(p, w) - C(S(p, w))$$

is concave for $w \geq 0$.

3. For all $C_n \in \mathcal{C}$, a firm’s payoff is concave if the firm is price anticipating; that is, for all $\mathbf{w}_{-n} \in (\mathbb{R}^+)^{N-1}$, the function

$$w_n \mapsto p_S(\mathbf{w})S(p_S(\mathbf{w}), w_n) - C_n(S(p_S(\mathbf{w}), w_n))$$

is concave in $w_n > 0$ if $\mathbf{w}_{-n} = \mathbf{0}$, and concave in $w_n \geq 0$ if $\mathbf{w}_{-n} \neq \mathbf{0}$.

4. The function S is uniformly less than or equal to D ; i.e., for all $p > 0$ and $w \geq 0$, $S(p, w) \leq D$.

5. The function $S(p, \cdot)$ has range containing $[0, D]$; i.e., for all $p > 0$ and for all $x \in [0, D]$, there exists a $w \geq 0$ such that $S(p, w) = x$.

Given any $S \in \mathcal{S}(D, N)$, we interpret the resulting market mechanism as follows: each firm n chooses a parameter w_n , thus specifying a supply function $S(\cdot, w_n)$, and the market clears according to $p_S(\mathbf{w})$. This determines the production, and hence the profit, of each firm—just as for the mechanism developed in §3.

We now briefly discuss each of the assumptions in Definition 1. The first assumption ensures that a *unique* market-clearing price exists; without this assumption, firms may not have a unique prediction of the market outcome given a strategic decision they make. The second and third conditions ease characterization of equilibria in terms of only first-order conditions. The second condition allows us to characterize competitive equilibria in terms of only first-order conditions, as in the proof of Theorem EC.1. The third condition allows us to characterize Nash equilibria in terms of only first-order conditions, a property we exploited in the proof of Theorem 1; indeed, at least quasiconcavity is generally used to guarantee existence of pure-strategy Nash equilibria in games with continuous action spaces (Fudenberg and Tirole 1991). The fourth condition forces the declared supply to be no larger than D ; this is a reasonable assumption when demand is deterministically known in advance, as we have assumed. Finally, the fifth condition is a “full range” assumption: it ensures that firms always have a strategic choice available to declare any supply between $[0, D]$, given the eventual market-clearing price p . This condition would be necessary to ensure that welfare maximization is possible; otherwise, for some choices of cost functions, no choice of \mathbf{w} could achieve the efficient outcome.

This class of market mechanisms generalizes the supply function interpretation of the mechanism discussed in §3. According to (4), each firm submits a supply function of the form $S(p, w) = D - w/p$, and the resource manager chooses a price $p_S(\mathbf{w})$ to ensure that $\sum_{n=1}^N S(p, w_n) = D$. Thus, for this mechanism we have $p_S(\mathbf{w}) = \sum_{n=1}^N w_n / ((N - 1)D)$ if $\mathbf{w} \neq \mathbf{0}$. Another possible mechanism is given by $S(p, w) = D - w/\sqrt{p}$; it is straightforward to verify that $p_S(\mathbf{w}) = [\sum_{n=1}^N w_n / ((N - 1)D)]^2$ if $\mathbf{w} \neq \mathbf{0}$.

Our interest is in the worst-case ratio (over $C_1, \dots, C_N \in \mathcal{C}$) of aggregate cost at any Nash equilibrium to minimal aggregate cost, defined as the solution to (1)–(3). Formally, for $S \in \mathcal{S}(D, N)$ we define a constant $\rho(D, N, S)$ as follows:

$$\rho(D, N, S) = \sup \left\{ \frac{\sum_{n=1}^N C_n(S(p_S(\mathbf{w}), w_n))}{\sum_{n=1}^N C_n(s_n)} \mid C_i \in \mathcal{C} \text{ for all } i, \right. \\ \left. \mathbf{s} \text{ solves (1)–(3), and } \mathbf{w} \text{ is a Nash equilibrium} \right\}.$$

Note that because any cost function satisfying Assumption 1 is strictly increasing and nonnegative, and $D > 0$, $\sum_{n=1}^N C_n(s_n)$ is strictly positive at any optimal solution \mathbf{s} to (1)–(3). However, Nash equilibria may not exist for some cost function choices C_1, \dots, C_N ; in this case we set $\rho(D, N, S) = \infty$. We have the following theorem; as discussed above, the proof is related to that for Theorem 1 in Johari and Tsitsiklis (2009), and full details of the argument can be found in Johari (2004), Chapter 5, Theorem 5.9.

THEOREM 4. *Assume $D > 0$ and $N > 1$. Fix $S \in \mathcal{S}(D, N)$. Then:*

1. *For any choice $C_1, \dots, C_N \in \mathcal{C}$, there exists a competitive equilibrium, and the resulting production vector solves (1)–(3) (i.e., it is efficient).*

2. $\rho(D, 2, S) = \infty$.

3. *If $N > 2$, then there exists a concave, strictly increasing, differentiable, and invertible function $B: (0, \infty) \rightarrow (0, \infty)$ such that for all $p > 0$ and $w \geq 0$,*

$$S(p, w) = D - \frac{w}{B(p)}.$$

4. *For $N > 2$, $\rho(D, N, S) \geq 1 + 1/(N - 2)$, and this bound is met with equality if and only if $S(p, w) = D - \Delta w/p$ for some $\Delta > 0$.*

We now briefly comment on the relationship between this result and the result obtained in Johari and Tsitsiklis (2009), where the authors consider the allocation of a divisible good among strategic buyers. It is shown that the proportional allocation mechanism (cf. §3) minimizes the worst-case efficiency loss when users are price anticipating, among all market-clearing mechanisms satisfying a set of conditions similar (though weaker) than those in Definition 1. In some sense, that theorem obtains a slightly stronger result than our result in this paper: in particular, Johari and Tsitsiklis (2009) obtain an explicit characterization of all mechanisms in the class they study; in part, this is due to the fact that they study mechanisms that are well defined for any quantity of resource and number of users. By contrast, in our setting the mechanisms we consider explicitly depend on the amount of resource available. Further, because our efficiency loss result in Theorem 2 depends explicitly on the number of firms, we consider a class of mechanisms in Definition 1 where the number of firms is fixed. Both of these features require a slightly different argument than that in Johari and Tsitsiklis (2009), although of course the results share a common heritage.

The preceding theorem establishes that, at least within the class $\mathcal{S}(D, N)$, the choice $S(p, w) = D - w/p$ that we have studied in this paper is in some sense “optimal:” it minimizes the worst-case welfare loss over all mechanisms in $\mathcal{S}(D, N)$. Note that in the third result of the theorem, we find that all mechanisms in $\mathcal{S}(D, N)$ are actually *independent* of N . Thus, our chosen mechanism from (4) is actually worst-case optimal for any number of firms N .

This result has two major caveats. First, there are strong concavity restrictions imposed for mathematical tractability; removing these remains a direction for future research. Second, and perhaps more undesirable from a practical standpoint, is the eventual deduction that any mechanism in $\mathcal{S}(D, N)$ has the possibility of “negative supply” out of equilibrium—i.e., a firm being asked to purchase the good rather than to supply it. (This follows from the characterization in part 3 of the theorem.) We address this issue in the following section.

7. Negative Supply

One undesirable feature of the parameterized supply functions we have chosen is that they allow for nonequilibrium outcomes in which a firm may have a negative supply. Although we have shown that the supply of each firm is nonnegative at both the competitive equilibrium and at the Nash equilibrium, nonequilibrium bidding may lead to negative supply to some firms. In practical terms, this implies that firms may be asked to buy the good when the market clears, even if they have no intention or ability to do so.

A natural question, therefore, is whether a reasonable parametric class of supply functions can be designed with properties similar to those we have already proven, but where the supply functions are uniformly nonnegative. Our first result, in §7.1, will demonstrate this is not possible in general. However, if we enforce a finite “maximum liability” for each firm, then we can use the same mechanism as that described earlier in the paper to achieve good market performance; we discuss such an approach in §7.2.

7.1. Nonnegative Parameterized Supply Functions

In this section, we will consider parameterized supply functions S that are uniformly nonnegative, and for which the payoffs to market participants are concave when they are price anticipating. These are formalized in the following definition; recall the definition of \mathcal{C} in (16).

DEFINITION 2. Given $D > 0$ and $N > 0$, the class $\mathcal{S}^+(D, N)$ consists of all differentiable functions $S: (0, \infty) \times [0, \infty) \rightarrow \mathbb{R}$ such that:

1. For all nonzero $\mathbf{w} \in (\mathbb{R}^+)^N$, there exists a unique solution $p > 0$ to the following equation:

$$\sum_{n=1}^N S(p, w_n) = D.$$

We let $p_S(\mathbf{w})$ denote this solution.

2. For all $p > 0$ and all $w \geq 0$, $S(p, w) \geq 0$.

3. For all $C_n \in \mathcal{C}$, a firm’s payoff is concave if the firm is price anticipating; that is, for all $\mathbf{w}_{-n} \in (\mathbb{R}^+)^{N-1}$, the function:

$$w_n \mapsto p_S(\mathbf{w})S(p_S(\mathbf{w}), w_n) - C_n(S(p_S(\mathbf{w}), w_n))$$

is concave in $w_n > 0$ if $\mathbf{w}_{-n} = \mathbf{0}$, and concave in $w_n \geq 0$ if $\mathbf{w}_{-n} \neq \mathbf{0}$.

Note that Condition 1 in the preceding definition is identical to Condition 1 in Definition 1, and Condition 2 in the preceding definition is identical to Condition 3 in Definition 1. The key difference is that in addition to these two conditions, we only have one simple constraint: the parameterized supply functions must be nonnegative. We have the following result.

PROPOSITION 1. *Fix $N > 0$ and $D > 0$. Let $S \in \mathcal{S}^+(D, N)$. Let $\mathbf{C} = (C_1, \dots, C_N)$, and $\bar{\mathbf{C}} = (\bar{C}_1, \dots, \bar{C}_N)$ be two collections of cost functions such that $C_n, \bar{C}_n \in \mathcal{C}$ for all n . If \mathbf{w} is a Nash equilibrium when the firms have costs given by \mathbf{C} , then \mathbf{w} is also a Nash equilibrium when the firms have costs given by $\bar{\mathbf{C}}$.*

The preceding result is disconcerting: clearly, any mechanism satisfying the conditions of the proposition can have arbitrarily high welfare loss, as well as an arbitrarily high markup above the competitive price, when firms are price anticipating. In the next section, we consider an alternate method by which we can address the problem of negative supply.

7.2. Maximum Liability

In this section we consider a simple modification to the basic model that protects firms from large payments due to having to buy some of the product (i.e., if they have negative supply) at a nonequilibrium outcome.

We continue to assume that $S(p, w_n)$ and the market-clearing price $p(\mathbf{w})$ are defined as before, i.e., $S(p, w_n) = D - w_n/p$. We fix a *maximum liability* $W > 0$, such that no firm will ever have to pay more than W when the market is cleared. Thus, if $p(\mathbf{w})S(p(\mathbf{w}), w_n) < -W$, then firm n only pays W to the market manager. Formally, the payoff of firm n now becomes:

$$\bar{Q}_n(w_n; \mathbf{w}_{-n}) = \max\{-W, p(\mathbf{w})S(p(\mathbf{w}), w_n)\} - C_n(S(p(\mathbf{w}), w_n)). \quad (17)$$

One interpretation of this game is as follows. Each firm submits a “deposit” of W to the market manager. The game is then played as before, and the market manager clears the market. At the resulting allocation, any required payment higher than W by a firm is forgiven.

We have the following proposition.

PROPOSITION 2. *Assume that $N > 2$ and $W > 0$. Suppose also that Assumption 1 is satisfied. Then \mathbf{w} is a Nash equilibrium of the game defined by (Q_1, \dots, Q_N) if and only if \mathbf{w} is a Nash equilibrium of the game defined by $(\bar{Q}_1, \dots, \bar{Q}_N)$.*

While this extension to the game is appealing from a market implementation point of view, we must be careful in interpreting the preceding result. Suppose that \mathbf{w} is a strategy vector where $p(\mathbf{w})S(p(\mathbf{w}), w_n) < -W$; in particular, $S(p(\mathbf{w}), w_n) < 0$. In this case we will have $\sum_{m \neq n} S(p(\mathbf{w}), w_m) > D$ —that is, the remaining firms will

be producing *excess supply*. In an economy with free disposal, this does not pose any problem, but if free disposal fails, then the market mechanism is problematic.² Furthermore, we note that if the maximum liability rule is implemented, then the market operator effectively subsidizes the misalignment of supply and demand; the total payment received from the demand side of the market is insufficient to compensate those suppliers that deliver positive supply. In general, then, addressing the possibility of negative supply when \mathbf{w} is out of equilibrium remains an important implementation-dependent issue.

8. Conclusion

We have considered a resource allocation problem where multiple suppliers compete to meet an inelastic demand. We present a novel investigation of a fundamental issue in mechanism design: how much strategic flexibility should players be given to ensure successful market performance? In our model, we restrict attention to uniform price market-clearing mechanisms. We demonstrate in this paper that by using a properly chosen class of supply functions parameterized by a single scalar, both high efficiency and low price markups can be guaranteed.

In addition to the analysis carried out here, we have extended the model and market mechanism to include the possibility that demand may be stochastic. Due to space constraints, we have deferred discussion of this extension to §EC.2 in the e-companion. We demonstrate there that the welfare loss result of §4 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.

9. 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. Indeed, as Joskow (2001) notes, this behavior was observed quite forcefully in the California electricity markets during the crisis in 2000.
2. In the context of electricity markets, such a situation indicates a misalignment of supply and demand, and can induce instability in the power grid. If this situation arises in the *day-ahead* market, then in principle supply and demand might be properly aligned using the real-time markets prior to actual delivery; in practice, however, electricity markets never force generators to serve as power sinks.

Acknowledgments

Portions of this paper first appeared in the Ph.D. thesis of Johari (2004). The authors gratefully acknowledge helpful conversations with John Ledyard, Shmuel

Oren, Andrew Philpott, Jong Shi-Pang, Pravin Varaiya, and Robert Wilson. This work was supported in part by the National Science Foundation under a Graduate Research Fellowship and under grant ECS-0312921.

References

- Anderson, E. J., X. Hu. 2008. Finding supply function equilibria with asymmetric firms. *Oper. Res.* **56**(3) 697–711.
- Ausubel, L. M., P. Milgrom. 2006. The lovely but lonely Vickrey auction. P. Cramton, Y. Shoham, R. Steinberg, eds. *Combinatorial Auctions*. MIT Press, Cambridge, MA, 17–40.
- Baldick, R., W. Hogan. 2001. Capacity constrained supply function equilibrium models of electricity markets: Stability, non-decreasing constraints, and function space iterations. Working Paper PWP-089, University of California Energy Institute, Berkeley.
- Baldick, R., W. Hogan. 2006. Stability of supply function equilibria: Implications for daily vs. hourly bids in a poolco market. *J. Regulatory Econom.* **30**(2) 119–139.
- Baldick, R., R. Grant, E. P. Kahn. 2004. Theory and application of linear supply function equilibrium in electricity markets. *J. Regulatory Econom.* **25**(2) 143–167.
- Chen, Y.-J., J. Zhang. 2010. Design of price mechanisms for network resource allocation via price of anarchy. *Math. Programming*, ePub ahead of print May 12.
- Clarke, E. H. 1971. Multipart pricing of public goods. *Public Choice* **11** 19–33.
- Day, C. J., B. F. Hobbs, J.-S. Pang. 2002. Oligopolistic competition in power networks: A conjectured supply function approach. *IEEE Trans. Power Systems* **17**(3) 597–607.
- Dimakis, A., R. Jain, J. Walrand. 2006. Mechanisms for efficient allocation in divisible capacity networks. *Proc. IEEE Conf. Decision Control (CDC)*, 1264–1269.
- Dubey, P., M. Shubik. 1978. The noncooperative equilibria of a closed trading economy with market supply and bidding strategies. *J. Econom. Theory* **17** 1–20.
- Fudenberg, D., J. Tirole. 1991. *Game Theory*. MIT Press, Cambridge, MA.
- Giraud, G. 2003. Strategic market games: An introduction. *J. Math. Econom.* **39** 355–375.
- Green, R. J. 1996. Increasing competition in the British electricity spot market. *J. Indust. Econom.* **44**(2) 205–216.
- Green, R. J., D. M. Newbery. 1992. Competition in the British electricity spot market. *J. Political Econom.* **100**(5) 929–953.
- Grossman, S. J. 1981. Nash equilibrium and the industrial organization of markets with large fixed costs. *Econometrica* **49**(5) 1149–1172.
- Groves, T. 1973. Incentives in teams. *Econometrica* **41**(4) 617–631.
- Hajek, B., G. Gopalakrishnan. 2002. Do greedy autonomous systems make for a sensible Internet? *Conf. Stochastic Networks, Stanford University, Stanford, CA*.
- Hart, O. D. 1985. Imperfect competition in general equilibrium: An overview of recent work. J. A. Kenneth, S. Honkapohja, eds. *Frontiers of Economics*. Basil Blackwell, Oxford, UK, 100–149.
- Hobbs, B. F., M. H. Rothkopf, L. C. Hyde, R. P. O'Neill. 2000. Evaluation of a truthful revelation auction in the context of energy markets with nonconcave benefits. *J. Regulatory Econom.* **18**(1) 5–32.
- 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. 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., J. N. Tsitsiklis. 2009. Efficiency of scalar-parameterized mechanisms. *Oper. Res.* **57**(4) 823–839.
- 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.
- Joskow, P. L. 2001. California's electricity crisis. *Oxford Rev. Econom. Policy* **17** 365–388.
- Kelly, F. P. 1997. Charging and rate control for elastic traffic. *Eur. Trans. Telecomm.* **8** 33–37.
- Klemperer, P. D., M. A. Meyer. 1989. Supply function equilibria in oligopoly under uncertainty. *Econometrica* **57**(6) 1243–1277.
- La, R. J., V. Anantharam. 2000. Charge-sensitive TCP and rate control in the Internet. *Proc. IEEE INFOCOM*, 1166–1175.
- Maheswaran, R. T. 2003. A game-theoretic analysis of agent-mediated resource allocation. Ph.D. thesis, University of Illinois at Urbana-Champaign, Urbana-Champaign.
- Maheswaran, R. T., T. Basar. 2003. Nash equilibrium and decentralized negotiation in auctioning divisible resources. *Group Decision and Negotiation* **12**(5) 361–395.
- Maheswaran, R. T., T. Basar. 2004. Social welfare of selfish agents: Motivating efficiency for divisible resources. *Proc. IEEE Conf. Decision and Control (CDC)*, 1550–1555.
- Mas-Colell, A., M. D. Whinston, J. R. Green. 1995. *Microeconomic Theory*. Oxford University Press, Oxford, UK.
- Monderer, D., L. S. Shapley. 1996. Potential games. *Games Econom. Behav.* **14**(1) 124–143.
- Moulin, H. 2008. The price of anarchy of serial, average, and incremental cost sharing. *Econom. Theory* **36**(3) 379–405.
- Moulin, H. 2010. An efficient and almost budget-balanced cost sharing method. *Games and Econom. Behavior* **70**(1) 104–131.
- Nisan, N., T. Roughgarden, E. Tardos, V. V. Vazirani, eds. 2007. *Algorithmic Game Theory*. Cambridge University Press, Cambridge, UK.
- Rothkopf, M. H., T. J. Teisberg, E. P. Kahn. 1990. Why are Vickrey auctions rare? *J. Political Econom.* **98**(1) 94–109.
- Sanghavi, S., B. Hajek. 2004. Optimal allocation of a divisible good to strategic buyers. *Proc. IEEE Conf. Decision and Control (CDC)*, 2748–2753.
- Shapiro, C. 1989. Theories of oligopoly behavior. R. Schmalensee, R. D. Willig, eds. *Handbook of Industrial Organization*, Vol. 1. Elsevier Science, Amsterdam, 329–414.
- Shapley, L. S., M. Shubik. 1977. Trade using one commodity as a means of payment. *J. Political Econom.* **85**(5) 937–968.
- Stoenescu, T. M., J. Ledyard. 2006. A pricing mechanism which implements in Nash equilibria a rate allocation problem in networks. *45th IEEE Conf. Decision and Control, San Diego*, 1270–1277.
- Stoft, S. 2002. *Power System Economics: Designing Markets for Electricity*. IEEE Press, Piscataway, NJ.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* **16**(1) 8–37.
- Vives, X. 2009. Strategic supply function competition with private information. Cowles Foundation Discussion Paper Number 1736, Yale University, New Haven, CT.
- Wilson, R. 2002. Architecture of power markets. *Econometrica* **70**(4) 1299–1340.
- Yang, S., B. Hajek. 2006. Revenue and stability of mechanism for efficient allocation of a divisible good. Submitted.
- 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.