

A Smart Market for Industrial Procurement with Capacity Constraints

Jérémie Gallien

Sloan School of Management, Massachusetts Institute of Technology, 50 Memorial Drive, Cambridge, Massachusetts 02142,
jgallien@mit.edu

Lawrence M. Wein

Graduate School of Business, Stanford University, 518 Memorial Way, Stanford, California 94306, lwein@stanford.edu

We address the problem of designing multi-item procurement auctions for a monopsonistic buyer in capacity-constrained environments. Using insights from classical auction theory, we construct an optimization-based auction mechanism (“smart market”) relying on the dynamic resolution of a linear program minimizing the buyer’s cost under the suppliers’ capacity constraints. Suppliers can modify their offers in response to the optimal allocation corresponding to each set of bids, giving rise to a dynamic competitive bidding process. To assist suppliers, we also develop a bidding-suggestion device based on a myopic best-response (MBR) calculation that solves an associated optimization problem. Assuming linear costs for the suppliers, we study within a game-theoretic framework the sequence of bids arising in this smart market. Under a weak behavioral assumption and some symmetry requirements, an explicit upper bound for the winning bids is established. We then formulate a complete behavioral model and solution methodology based on the MBR rationale and show that the bounds derived earlier continue to hold. We analytically derive some structural and convergence properties of the MBR dynamics in the simplest nontrivial market environment, which suggests further possible design improvements, and investigate bidding dynamics and incentive compatibility issues via numerical simulations.

Key words: procurement auctions; smart markets; iterative bidding mechanism; best response

History: Accepted by Fangruo Chen and Stefanos A. Zenios, special issue editors; received October 24, 2000.

This paper was with the authors 12 months for 3 revisions.

1. Introduction

1.1. Background and Motivation

Online business-to-business auctions are becoming increasingly widespread because their demand revelation properties are superior to that of fixed-price mechanisms, and the Internet maintains low transaction costs. A May 2000 study by Forrester Research, Inc. (Kafka et al. 2000) estimated that the total annual volume of transactions realized through online business-to-business auctions in the United States would reach \$746 billion by 2004; even if this estimate is far too high, given the particularly bullish stock market environment when it was made, it does point to a profound, ongoing transformation of procurement practices. An important resulting challenge faced by online procurement service providers is to design auction mechanisms adapted to complex procurement environments.

An example of such a mechanism—and the one that inspired our focus on capacity constraints—is the bidding process used by FreeMarkets, Inc. (Rangan 1999). FreeMarkets organizes bidding events for large

industrial buyers typically seeking to auction off procurement contracts for specified quantities of several component types (or lots). In the mechanism currently used, the various component types are ordered, and each component is effectively auctioned off in a sequential manner. More precisely, binding bids on selling prices are solicited for each component type from competing suppliers, which can observe the value of the lowest bid submitted thus far. When a prespecified time limit is reached for the first component type in the list (after perhaps some overtime period triggered by an activity-based criteria), bidding closes for that component type, competition turns to the next component type, and this process repeats itself until the list of component types the buyer requires is exhausted. However, in contrast to traditional auctions, the actual allocation of procurement contracts is only decided by the buyer after the entire bidding event is completed, and the lowest bidder may not be awarded the contract for a component. The reason for postponing the allocation decision is to allow the buyer to take into account some of the nonprice factors mentioned earlier.

Although this mechanism appears to be an industry-wide practice as of summer 2000 (it is also used by FreeMarkets' competitor eBreviate), several sources suggest directions for improvement. In particular, a consequence of the delayed-allocation decision is that suppliers' bids are based on very limited competitive information, which appears to be especially damaging in environments characterized by supplier capacity constraints. During our meeting with FreeMarkets, cofounder Sam Kinney hypothesized that some suppliers grew nervous during the bidding event about receiving a total allocation exceeding their production capacity, especially if they had submitted relatively low bids on one or several of the components auctioned early on (Kinney 1999). As a result, these suppliers would withhold their bids on later components, thereby hampering competition. Indeed, some buyers also seem to be aware of this issue; for example, Jack Porter, a purchasing manager at Caterpillar, made the following comment during an interview on his experience as a FreeMarkets customer: "This concept only applies in those areas where there is a great deal of capacity" (Janhke 1998).

Motivated by these observations, our objective is to design and analyze a stylized online multi-item procurement auction mechanism adapted to supply environments with production capacity constraints. Following a review of the relevant literature, we describe in §2 the market environment under study, the proposed auction mechanism, and the specifics of the game formulation. In §3, we discuss our solution methodology, derive some convergence bounds, and investigate bidding dynamics and incentive compatibility issues. Concluding remarks are provided in §4.

1.2. Literature Review

Multi-item auctions have been one of the most active research areas in microeconomics lately (see §10 of Klemperer 1999). They are also becoming an increasingly popular topic in operations management, operations research, and computer science. The crux of the associated outcome prediction problem lies in the interdependence among the various objects being auctioned off, which arises in practice from factors such as capacity constraints, budget constraints, economies of scale/scope, and legal restrictions. In our model, capacity constraints are the only source of interdependence across component types, and there are no interdependencies among different units of the same component type because component production costs are assumed to be linear. Our discussion of the multi-item auction literature is organized into two parts: The first is relevant to the design of our auction mechanism, and the second addresses outcome analysis.

Our auction format belongs to the family of market mechanisms known as smart markets, which are

exchange institutions supported by a computer executing an optimization algorithm to solve the allocation problem associated with each given set of bids. Although this approach goes back at least as far as Stanley et al. (1954), the combined developments of computational power and network communications enabled the extension of Stanley et al.'s static allocation approach to a dynamic market design problem. While the dynamic approach has been considered in a wide variety of market environments (see McCabe et al. 1991 and Gallien 2000 for surveys), most papers in the smart-market literature only contain mechanism descriptions, problem formulations, and/or experimental results, as opposed to analytical outcome predictions (exceptions are discussed below). In contrast, we study analytically the iterative competitive process that results from the real-time linear programming solution generating competitive information feedback to the bidders, who can then respond by modifying their bids, thereby creating a new instance of the linear program (LP).

A key goal in designing an auction mechanism is to strike a satisfactory balance between facilitating competition and preventing collusion. To this end, we employ an ascending-auction format (motivated by the discussion by Cramton 1998) and try to take advantage of the design flexibility allowed by organizing the auction on a distributed network of computers. An important issue, largely overlooked in the smart-market literature, impacting the facilitation of competition is that bidders confronted with a multi-item smart-market mechanism in an ascending auction face formidable bidding decisions, even in a complete information setting (see Rothkopf et al. 1998 for details on the similarity of the complexity of the bidder's problem and the "winner determination" problem). When the bidders are given incomplete information to limit collusion, their computational abilities are likely to be quickly overwhelmed in all smart markets of interest; our mechanism attempts to mitigate this problem.

To motivate our mechanism design, we discuss three possible approaches (there may be others) in the multi-item setting. One approach is to auction the items sequentially. In this case, the apparent simplicity of the auction and information feedback given to bidders inhibits competition, as discussed in the context of FreeMarkets' practices in §1.1; see Elmaghraby (2002) for a study of the (difficult) sequencing problem that arises in the sequential setting.

The other two approaches employ simultaneous auctions. In a simultaneous price-setting auction, the auctioneer announces individual prices for all items in each round, soliciting from the bidders the quantities they are willing to produce at these prices, and using this information to update the prices for the next

round. This is the approach adopted by Demange et al. (1986), Wellman et al. (2001), Milgrom (2000), and Parkes and Ungar (2000), although in market environments where the allocation variables are discrete. The auction interpretation of the algorithm described in Bertsekas and Castanon (1989) provides a price-setting mechanism adapted to the continuous transportation problem in our paper. However, this price-setting approach has several drawbacks. Posted unit prices for the same item may vary across bidders (Bertsekas and Castanon 1989, Parkes and Ungar 2000), which may fuel perceptions of inequity from suppliers. This type of mechanism also imposes a formidable computational burden on suppliers, who (in settings with demand for multiple items) need to evaluate in real time their utility over the set of all subsets of items being traded; there is some evidence that this so-called valuation computation problem, which is associated with the possibility of package bidding, can significantly hamper efficiency (Larson and Sandholm 2001). Finally, because suppliers compete through price offers in most traditional request for quotes (RFQ) processes, requiring them instead to propose quantities in response to given prices may be too much of a deviation from current industrial practices. In fact, actual implementations of price-setting auctions seem to be extremely rare, despite the relatively large number of papers dedicated to this type of iterative auction in the literature.

For these reasons, we investigate in this paper a third option: A simultaneous quantity-setting auction whereby suppliers submit independent price bids on all items in each round, and the mechanism computes the resulting potential quantity allocation in an iterative manner. In our setting, where capacity constraints are the only source of interdependency across items and they are known to the auctioneer (we discuss the incentive compatibility issues associated with this assumption in §3.3.4), our optimization-based allocation engine effectively alleviates the need for package bidding by taking these constraints into account. We note that Emptoris and IBM both have smart-market software supporting simultaneous quantity-setting auctions. In this setting, we believe that the potential allocation quantity and even the shadow prices associated with the procurement quantity constraints do not constitute sufficient competitive feedback information to aid the bidders' decisions.¹ Consequently, we propose a particular form of competitive information feedback called myopic best response (MBR). This feature allows each supplier to enter its production

costs into a private calculation device that will compute on the supplier's behalf the bids maximizing the potential payoff in the next round, under the assumption that the competitors' bids remain unchanged. The novelty of our MBR feedback is that it fully accounts for the optimization problem that will be solved by the buyer in the next round when performing this calculation for the bidder. More specifically, the MBR feedback mechanism can be viewed as an oracle that is given access—for the sake of performing its computations—to some competitive data that are otherwise hidden to suppliers; we note that bidding decision-support systems are used in practice by at least one company, Logistics.com, when organizing smart markets for truckload carrier selection.

We now turn to the relevant literature on outcome prediction. Because the computational requirements associated with the assumption of full rationality are particularly daunting in this game under uncertainty, we follow instead the approach developed in the theory of tatonnement and learning in games, further discussed in §3.1. Under a “weak” behavioral assumption defined in §3.2 and a condition on the minimum level of competition, we derive an upper bound on the winning bids at the end of the auction, which to our knowledge is the first explicit outcome prediction as a function of the model parameters in the smart-market literature. We also propose a complete solution concept (see §3.1) based on the natural adaptation to our quantity-setting context of the MBR or “straightforward” bidding strategies assumed in the price-setting mechanisms of Demange et al. (1986), Wellman et al. (2001), Milgrom (2000), Parkes and Ungar (2000), and of the price-updating rules in Bertsekas and Castanon (1989). These authors establish explicit relationships in their respective market environments between the final bids and the minimum equilibrium bids, or between the final allocation and the surplus-maximizing (i.e., efficient) allocation. The key to most of these convergence results is the interpretation of these mechanisms as primal/dual algorithms, a connection that is extensively discussed in Bikhchandani et al. (2002). In contrast, we consider the traditional perspective of a revenue-maximizing auctioneer; although revenue-maximizing auctions are efficient if there is a resale market after the auction (Ausubel and Cramton 1999), this is unlikely to occur in procurement settings. From the revenue-maximizing viewpoint, we derive some upper bounds on the final bids, and hence on the buyer's final procurement cost, in terms of the specific data of the procurement problem under study (production costs, capacity constraints) and investigate how they are affected by the initial bids. However, we do not present any results related to the (sub)optimality of our mechanism with respect to revenue or efficiency maximization.

¹ Bidding below the current shadow price on a given component does not imply any guarantee for the corresponding allocation quantity.

2. Mechanism Description and Game Definition

2.1. Market Environment

The initial situation giving rise to our market design problem is a monopsony where a buyer (typically a large manufacturing corporation) wants to purchase some specified quantity of m component types (namely, q_j units of each component type j in $\{1, \dots, m\}$). Facing this opportunity are n suppliers ready to compete on selling prices for the award of the corresponding procurement contracts. We assume that the only parameters left to specify in these contracts are the total quantity along with the corresponding selling price for each component type. In particular, the rules by which payment terms, delivery schedules, insurance provisions, etc. will be determined are clearly defined from the outset and accepted by suppliers. In addition to the buyer's quantity requirements and his low cost objective, the other consideration relevant to the allocation of components is that each supplier has a limited overall production capacity across all component types. In the present study, we assume that the capacity limitation of each supplier $i \in \{1, \dots, n\}$ can be appropriately described by a linear model specifying a total amount c_i of production resource available (e.g., machine hours or man hours), and the amounts a_{ij} of this resource needed to produce one unit of each component type j . We also assume these capacity constraints to be exogenous, so that suppliers do not have the option of acquiring additional production capacity as part of the competitive interaction under study (see Budde and Göx 1999 for an analysis of a procurement auction with endogenous capacity decisions). Although we do not specify the duration nor delivery schedule of the procurement contracts being auctioned, note that our specific set of assumptions (deterministic quantities, fixed supply capacity, limited number of allocation factors) would likely correspond in practice to short-term contracts. Finally, note that this market environment is formally equivalent to the setting where a firm wants to buy a single type of item for each one of its m manufacturing facilities (facility j requires q_j units) from a set of n capacitated suppliers with varying production and transportation costs.

The remainder of this section is organized as follows. The core of the proposed optimization-based allocation mechanism is presented in §2.2, along with some practical implementation issues. Section 2.3 focuses on the suppliers' payoff function, which is subsequently used both to derive outcome predictions in §3 and to generate more efficient competitive information feedback. The design of this information feedback is discussed in §2.4, which more generally

summarizes the information structure underlying this competitive interaction.

2.2. Proposed Mechanism

Bidding in our mechanism takes place through a discrete set of bidding rounds, denoted by $t \in \mathbb{N}$ ($t = 0$ is the initial round). More precisely, each supplier i is invited in each round t of the auction to simultaneously bid on a unit selling price $b_{ij}(t)$ for all component types j in $\{1, \dots, m\}$.

Each bid $b_{ij}(t)$ corresponds to an offer in round t by supplier i to sell *any* quantity between 0 and q_j of component type j at a unit price given by the value of the bid (see §2.3 for a discussion of how this relates to the suppliers' objective). The first constraint on these bidding decisions is a classical nonreneging rule (i.e. $b_{ij}(t) \leq b_{ij}(t')$ for $t' \leq t$) imposed to ensure the efficiency of this mechanism: Increasing a bid submitted in some previous round would amount to reneging a prior offer to sell, and thus threaten the stability of the price formation process. The second constraint is a common multiple rule (i.e., $b_{ij}(t) \in \epsilon\mathbb{N}$ for some $\epsilon > 0$) that further restricts the set of admissible bids to a regular lattice with granularity ϵ , or " ϵ -grid." This rule both enforces a minimum bid decrement frequently imposed in practice (see Aeppel 1999) and prevents strategic collusion signals using the lower digits of the bids (e.g., Kelly and Steinberg 2000), as long as the grid is sufficiently coarse (Rothkopf and Harstad 1994). In addition, the common multiple rule resolves a technical definition problem associated with the solution concept we introduce in §3.3 (see the discussion of the arg max operator in §2.4).

The key feature of our auction format is that the overall capacity constraint of each supplier will always be satisfied by all possible allocation outcomes, which seems desirable from both the suppliers' and the buyer's perspective, as discussed in §1. This is achieved in two steps: First, the capacity constraints $((a_{ij})_{j \in \{1, \dots, m\}}, c_i)$ of each supplier i are estimated *before* the beginning of the auction. The parameters of these capacity constraints can be either assessed by an auditor during a precertification/qualification of supplier i prior to the auction, or merely provided by supplier i . In the latter case, the incentive compatibility issues linked to this capacity revelation become particularly important and are discussed in §3.3.4. Second, once all the bids have been submitted for round t during the auction, the mechanism calculates the quantity allocation $x_{ij}(t)$ of component type j to supplier i , minimizing the buyer's total procurement cost, subject to both the quantity requirement constraint for each component type and the capacity constraint of each supplier (estimated during the first step mentioned above). This allocation $\mathbf{x}(t) = (x_{ij}(t))_{i \in \{1, \dots, n\}, j \in \{1, \dots, m\}}$ is obtained by solving in each round t the following LP,

referred to in the remainder of this paper as $AE[\mathbf{b}(t)]$ (for “allocation engine”):

$$\begin{aligned} \min_{x_{ij}} \quad & \sum_{i=1}^n \sum_{j=1}^m b_{ij}(t) x_{ij} \\ \text{s.t.} \quad & \sum_{j=1}^m a_{ij} x_{ij} \leq c_i \quad \forall i, \\ & \sum_{i=1}^n x_{ij} = q_j \quad \forall j, \\ & x_{ij} \geq 0 \quad \forall (i, j). \end{aligned} \quad (1)$$

Note that the complete specification of the allocation mechanism must also include a selection rule discriminating among multiple optimal solutions to $AE[\mathbf{b}(t)]$ when needed. Although we do not impose in general any particular selection rule, the convergence results presented in §3.2 require the optimal solution $\mathbf{x}(t)$ to be an extreme point of the feasible set of $AE[\mathbf{b}(t)]$ (which seems to be a desirable feature for the buyer, because solutions corresponding to extreme points tend to involve a smaller number of suppliers, and is easy to achieve in practice, because some optimization algorithms such as the simplex method naturally generate extreme point solutions). Following the practice of eBay and other online auction sites, we propose that the tie-breaking rule among optimal extreme-point solutions be based on the time at which bids are submitted. Hence, the allocation engine $AE[\cdot]$ along with this optima selection rule defines a function that unambiguously associates a potential allocation $\mathbf{x}(t) = (x_1(t), \dots, x_n(t))$ to any given set of bids $\mathbf{b}(t)$, which we can summarize with the notation $\mathbf{x}(t) = AE[\mathbf{b}(t)]$.

To encourage competition among suppliers across bidding rounds, our proposed mechanism displays on each supplier i 's screen at the end of each round t its own private potential allocation $\mathbf{x}_i(t) = (x_{i1}(t), x_{i2}(t), \dots, x_{im}(t))$; we refer to this process as private potential allocation feedback. Each supplier i can then respond to this competitive information in the following round by submitting a new set of bids $\mathbf{b}_i(t+1)$ and may in particular lower its bids on the component types for which its potential allocation $\mathbf{x}_i(t)$ is deemed unsatisfactory. We defer until §2.4 a more complete description of the information structure in this mechanism.

The relevance of the private potential allocation feedback as a competitive signal—and hence its ability to generate bidding competition—follows from the activity-based termination rule that we impose. More precisely, the bidding round iterations described above terminate after the first round $T \geq 1$, where no supplier modifies any of its bids from the previous

round, i.e., when $\mathbf{b}(T) = \mathbf{b}(T-1)$. At the time of termination, the allocation $\mathbf{x}_i(T)$ calculated for each supplier i by the allocation engine becomes final (hence the adjective “potential” used to qualify $\mathbf{x}(t)$ for $t \leq T$, because the termination round of the auction is not known a priori), and the final unit prices paid by the buyer to supplier i for the component types for which it receives a positive allocation are given by $\mathbf{b}_i(T)$, so that the buyer's final procurement cost resulting from the auction is $\sum_{i=1}^n \sum_{j=1}^m b_{ij}(T) x_{ij}(T)$. Note that this objective function does not take into account the fixed cost of establishing and managing a vendor relationship, so that nothing prevents this auction mechanism from selecting a large number of suppliers and generating many fragmented orders for the same component type in the final allocation. More generally, while the stylized auction mechanism just described is well suited to an analytical study, the allocation environment it captures may be too simplistic in practical situations. Some possible generalizations, including a bid adjustment system allowing the buyer to incorporate relative preferences among suppliers (arising, for example, from incumbent switching costs), are briefly discussed in §4.

2.3. Bidders' Payoff Function

Our key assumption here is that each supplier i has a fixed unit production cost v_{ij} for each component type j , so that the supplier's payoff (or utility) function in this market mechanism is described by $\Pi_i(\mathbf{b}(T)) \equiv \sum_{j=1}^m (b_{ij}(T) - v_{ij}) x_{ij}(T)$, where T is the terminal bidding round as previously defined and $(b_{ij}(T), x_{ij}(T))$ are related through $\mathbf{x}(T) = AE[\mathbf{b}(T)]$. Consequently, our model does not capture the possible economies of scale that suppliers may experience when they receive large orders (setup times, production learning curve, etc.)—nor does it capture the possible diseconomies of scale associated with, for example, the necessary purchase of additional production capacity. One scenario where the linear cost assumption may be realistic is a market for several metal-stamped components (e.g., body parts) in the automotive industry, where the existing stamps/press matrix are provided by the auto manufacturer to the selected supplier(s); in this case, the equipment would be used solely to provide the auto manufacturer's work, and the setup times would be relatively small. However, cost linearity is a strong assumption, and we refer to Gallien and Wein (2002) for a discussion of the design of improved smart-auction mechanisms allowing suppliers to take (dis)economies of scale into account when they formulate their bids.

Finally, neither our payoff function nor our allocation mechanism captures the possible synergies (or incompatibilities) across component types, because each supplier's production cost for one component

type is also independent of its allocation of the other component types. In procurement situations in which such (dis)economies of scope are particularly important, we believe that some alternate allocation mechanisms (e.g., a combinatorial auction, see Rothkopf et al. 1998) may be more appropriate.

2.4. Information Structure

In this section, we describe the information structure underlying our proposed market mechanism, namely which entity knows what information. We also show how the suppliers' payoff function presented in the previous section can be used to complete the competitive information feedback given to bidders during the auction and lead them to make more efficient bidding decisions.

Private information is held by three types of agents: the suppliers, the buyer, and the mechanism/auctioneer. We assume here that the auctioneer is trusted by the suppliers to be respectful of its confidentiality commitments. This is legitimized in our view by the current fierce competition between online procurement auctioneers (e.g., Freemarkets and eBreviate), where reputations for integrity standards seem to play a particularly important role. This, along with the technological medium used to support this auction (a distributed network of computers), effectively justifies the decoupling between the buyer and the mechanism/auctioneer that is inherent in our mechanism.

We first discuss the information structure relative to the market environment. Because the quantities $(q_j)_{j \in \{1, \dots, m\}}$ of each component type required by the buyer motivate in large part the suppliers' interest in the auction, it seems natural that they should be common knowledge to all participants. However, a key feature enabled by a computer network running the auction is that the suppliers need not know how many competitors they have, let alone their identities (for studies of traditional auctions with unknown or uncertain number of bidders, see Harstad et al. 1990 and references therein). Because this feature offers some protection against possible collusive behaviors, we have designed our auction mechanism to take advantage of it. In adherence to this design philosophy, suppliers are not informed by the mechanism of their competitors' capacity constraints $((a_{ij})_{j \in \{1, \dots, m\}}, c_i)$ either. However, this capacity information is required by the mechanism to ensure that all potential allocations generated satisfy these constraints (see §3.3.4 for a discussion of the incentive compatibility issues associated with the revelation of capacity by the suppliers).

The only dynamic competitive information feedback proposed so far is the private potential allocation feedback $\mathbf{x}_i(t)$ communicated to each supplier i

during the auction but hidden from competitors. In particular, because we seek to limit the possibilities of collusion and strategic behavior, our mechanism does not send any direct information to suppliers during the auction about competitors' bids (see Cramton 1998). Using private potential allocation feedback alone, however, would probably lead to an inefficient price-formation process because suppliers would lack sufficient information to rationally decide on which component types to lower their bids, and by how much. As a result, suppliers in this situation would most likely engage in exploratory bidding behaviors, whereby they would try to identify the minimal bid decrease leading to some improvement of their potential allocation through multiple rounds of minimum bid decrement trials.

The option that we have investigated to prevent this involves giving each supplier in every round the opportunity to have its MBR bid calculated by a local bidding suggestion device. In other words, the suggested bid to supplier i is the one that maximizes the potential payoff function in the next round, assuming that all competitors' bids remain unchanged from the current round; i.e.,

$$\mathbf{b}_i^*(t+1) \equiv \sup \left[\arg \max_{\substack{0 \leq \mathbf{w} \leq \mathbf{b}_{-i}(t) \\ \mathbf{w} \in (\epsilon\mathbb{N})^m}} \Pi_i(\mathbf{w}, \mathbf{b}_{-i}(t)) \right],$$

where $\Pi_i(\cdot)$ is the payoff function defined in §2.3, \mathbf{w} represents the decision variables associated with the bids of supplier i in the following round, and $\mathbf{b}_{-i}(t)$ is the standard game-theoretic notation for the bids of all supplier i 's competitors in the current round. Note that the $\arg \max$ operator acts upon a finite set because of both the nonreneging and the common multiple rules, so that the output of this operator is always nonempty even though $\Pi_i(\cdot)$ is not necessarily continuous. Moreover, the \sup operation over the $\arg \max$ set, provided it is associated with a total order over $(\epsilon\mathbb{N})^m$, ensures that the MBR bid is uniquely defined. With the lexicographic order, for example, suppliers will not receive a suggestion to lower their bids unless their prospective potential payoff when doing so is strictly greater than their current potential payoff.

Note that the MBR bid-suggestion device has access to considerably more information (competitors' bids and capacity constraints) than its user, which should constitute an attractive feature from the suppliers' standpoint. However, a key information requirement of the MBR calculation is that each supplier i willing to use this device needs to provide the value of its own production costs $(v_{ij})_{j \in \{1, \dots, m\}}$ via the auction software interface. Because the production costs of each supplier are sensitive and private information, it is essential for the auctioneer to credibly promise

Table 1 Summary of the Information Structure

Information	Notation	Information available to			
		Supplier i	Other suppliers	Mechanism auctioneer	Buyer
Number and identity of competing suppliers	n	no	no	yes	yes
Component quantities required	$(q_j)_j$	yes	yes	yes	yes
Capacity constraint of supplier i	c_i	yes	no	yes	yes
Production cost of supplier i	$(a_{ij})_j$ $(v_{ij})_j$	yes	no	yes	no
Bids of supplier i	$\mathbf{b}_i(t)$	yes	no	yes	yes
Potential allocation of supplier i	$\mathbf{x}_i(t)$	yes	no	yes	yes
Best-response suggestion of supplier i	$\mathbf{b}_i^*(t)$	yes	no	yes	no

that these costs, if entered into the MBR calculation device, will not be disclosed to the buyer or to any other supplier. We also observe that in end-consumer online auction sites (Yahoo!, eBay, etc.), users routinely enter their maximum bid for an item into some software agent (known as proxy bidding, bidding agent, bidding elf, etc.) that then submits bids on their behalf whenever needed and up to this specified limit (Lucking-Reiley 2000). Because these maximum bid limits are the exact analog of the production costs in our procurement auction setting (and more generally correspond to bidders' valuations in the traditional auction theory literature), we feel that the MBR bidding suggestion feature is likely to be implementable in practice; for a more persuasive argument in support of this statement, the reader is referred to our analysis in §3.3.4 of the incentive compatibility issues linked to the revelation by suppliers of their costs.

To conclude this section, we summarize in Table 1 the information structure underlying our proposed auction mechanism. We are now ready to turn to the analysis of our proposed market mechanism.

3. Analysis

3.1. Methodology and Solution Concept

The basic premise underlying our methodology is that bidders will be less than fully rational when confronted with the smart-auction mechanism described earlier. While this assumption simplifies the analysis, we share with other researchers (Chapter 1, Samuelson 1997) the belief that it is also motivated by behavioral considerations. In this dynamic game of incomplete information, the classic solution concepts associated with the requirement of full rationality predict that bidders will form probabilistic beliefs about all relevant uncertain information in the game, update these beliefs in a consistent manner during the auction, and derive and execute in every round a bidding

strategy maximizing the expected value of their final payoff function accordingly (Chapter 8, Fudenberg and Tirole 1996 and references therein). Because the uncertain information includes not only the competitors' production costs but also their capacity constraints, their bids, and even their number, a rational Bayesian expected utility maximizer would have to solve an overwhelmingly complex stochastic dynamic program with imperfect information. Our solution concept assumes instead a deterministic, less than fully rational behavioral model, ultimately specifying for each supplier a mechanistic relation between the information received in each round and its subsequent bidding decision. This type of model is known in the literature (e.g., Samuelson 1997, Fudenberg and Levine 1998) as a nonequilibrium adjustment or tatonnement process.

In §3.2, we first examine the implications of our so-called weak behavioral assumption on the limiting bids of the corresponding process. In §3.3, we investigate the properties of the bidding process arising when each bidder truthfully enters its costs into the MBR device described earlier and follows its bidding suggestions at every round; we use the limiting bids of this process as our solution concept.

The concept of MBR thus plays a role in our study on two related but distinct levels: as an information feedback and bidding decision-support mechanism (described in §2.4), and as part of the behavioral assumptions underlying our solution concept. While one could envision the MBR bidding device as a mandatory requirement built into the auction mechanism, we believe that such a heavy-handed approach would encounter considerable resistance from bidders in practice. Rather, because MBR has several relatively realistic features among all possible behavioral models (see further discussion) and is also somewhat amenable to analysis, it seems legitimate to study its implications. But we point out that in practice actual bidding decisions may deviate from the MBR bidding suggestions; e.g., suppliers could find out from the MBR bids which component seems most profitable, then place a different bid value on it than suggested. In general, the actual outcome of the auction in practice could thus be quite different from the prevision provided by the MBR solution concept.

We conclude this subsection with a detailed discussion of the MBR behavioral assumption. Although this model seems close to the classical Cournot adjustment process (Chapter 1, Fudenberg and Levine 1998), there remain a few important differences between them. First, the information structure of the game to which they apply differs, because the duopoly competition studied by Cournot is a game of common knowledge/full information, whereas our auction format corresponds to a game under uncertainty. Moreover, payoffs in our mechanism are not obtained at the end of

every round, as is implicitly the case in the Cournot adjustment process, but only at the end of the auction. Consequently, the objective function used in each supplier's MBR calculation only represents its potential payoff in the next round and only coincides with the actual payoff in the round immediately preceding the end of the auction. The MBR behavioral model thus implicitly assumes the belief that the auction will terminate in the next round, justifying the use of the adjective "myopic." Finally, because of the non-renegeing bidding rule, decision spaces in different rounds are not independent from each other, as in the Cournot adjustment. More precisely, each bid constrains the future bidding decisions of the supplier that submitted it, and this constraint is more stringent as the value of this bid is lowered.

These specific features provide some justification for applying the MBR model to the game under study. In the classical case with no uncertainty, the Cournot adjustment process has been sharply criticized for the players' naiveness and/or the reasoning inconsistencies it implies. While it relies in every round on the assumption that the competitors' actions will remain constant in the next round (strategy stability), this very assumption is observably proven wrong in every round until equilibrium is reached (Fudenberg and Levine 1998). However, in the information structure underlying the game we study, bidders cannot observe their competitors' bids directly, and the indirect process by which they could try to infer these bids—based on the history of their own bids, private allocation feedback, and MBR bidding suggestions—seems particularly complex. As a consequence, the assumption of strategy stability seems somewhat plausible, in sharp contrast to more classic games where past actions are common knowledge.

Moreover, suppliers in this auction will be particularly cautious not to submit low bids prematurely, because of the lock-in effect induced by the nonrenewing rule. In this setting, strategy stability could be perceived as a wise assumption on which to base bidding decisions, even when also perceived as probably wrong. This is because incorrectly assuming that the competition will be more aggressive in the next round may lead to lower bids than necessary, thus wasting some margin and restricting the space of future possible actions. On the other hand, the activity-based termination rule ensures that a supplier that underestimates its competitors' aggressiveness in a given round by assuming strategy stability will not be penalized: If this supplier ends up lowering one of its current bids, then the auction will continue for at least one more round, giving that supplier the opportunity to adapt its bids to the actual bidding strategy of his competitors. If this supplier does not modify any of its bids, then the auction may terminate, proving

the strategy stability assumption correct and therefore the bidding decision optimal, or it may continue, leaving the supplier the option to further modify its bids. Despite these attractive features, however, we believe that the predictive power of the MBR solution concept could turn out in practice to be quite weak, even if it seems legitimate to study as a benchmark model.

3.2. Weak Behavioral Assumption and Convergence Bounds

To establish some convergence bounds, we only rely initially on the weak behavioral assumption that a supplier receiving in some round a private allocation feedback of zero for all component types will lower at least one of its bids, unless doing so would generate a negative profit. Formally, this means that if $x_i(t) = 0$ and there exists $j \in \{1, \dots, m\}$ such that $b_{ij}(t) > v_{ij} + \epsilon$, then there exists $j \in \{1, \dots, m\}$ such that $b_{ij}(t+1) \leq b_{ij}(t) - \epsilon$.

We begin by stating an ex post result (i.e., relying on an observation of the final allocation made a posteriori) that relates the final winning bids and the order statistics of the production costs.

PROPOSITION 1 (EX POST BOUND). *Let T be the final round of the auction and let $v_{1:n}^j, \dots, v_{n:n}^j$ denote the order statistics of (v_{1j}, \dots, v_{nj}) , $j \in \{1, \dots, m\}$, and define $P^C \equiv \{i \in \{1, \dots, n\}, x_i(T) = 0\}$. Provided that $|P^C| \geq 1$ and under the weak behavioral assumption, we have*

$$x_{ij}(T) > 0 \Rightarrow b_{ij}(T) \leq v_{n-|P^C|+1:n}^j + \epsilon \quad \forall j \in \{1, \dots, m\}. \quad (2)$$

The proof of Proposition 1 (see the appendix) exploits linear programming duality and complementary slackness; we strongly suspect that it can be extended to a more general class of optimization-based auctions. In other words, Proposition 1 states that the highest selected bid on each component type at the end of the auction cannot be larger than some order statistic of the production costs for that component type (modulo the minimum bid decrement ϵ), and the greater the number of rejected suppliers, the lower the rank of this order statistic. Moreover, it seems natural that the upper bound in (2) involves the number of suppliers eventually receiving no allocation, because the only behavioral assumption required to derive this result concerns precisely the suppliers receiving a null private-allocation feedback. Proposition 1 generalizes to the case of a multi-item and optimization-based procurement auction the well-known result (Vickrey 1961) that in a single-item English ascending auction, the price eventually obtained by the seller is equal to the second highest valuation among the bidders (modulo the minimum

bid increment). Although the information structure is slightly different in our mechanism, the set of parameters $m = 1$, $q_1 = 1$, c_i unbounded, and $a_{ij} = 1$ for all i and j , when the allocation selected is always an extreme point, corresponds to the classic auction of a single indivisible object with no capacity constraints. In this case, there are $|P^C| = n - 1$ losing suppliers, so that the bound provided by Proposition 1 is indeed $v_{2:n}^1 + \epsilon$.

Although Proposition 1 is an ex post result, it can be used to prove an ex ante result (i.e., applicable before the auction outcome is observed). To state this result in a compact form, we first assume without loss of generality that the suppliers are numbered by increasing order of capacity, i.e., $c_1 \leq c_2 \leq \dots \leq c_n$. Next, for all $j \in \{1, \dots, m\}$ we define $a_j = \max_{i \in \{1, \dots, n\}} a_{ij}$. We can then conveniently define the maximal load number p as the smallest integer such that $\sum_{i=1}^{p+1} c_i > \sum_{j=1}^m a_j q_j$ if $\sum_{i=1}^n c_i > \sum_{j=1}^m a_j q_j$, or as $p = n$ otherwise. The observation that the maximal load number provides an upper bound on the number of suppliers loaded to their capacity in any feasible allocation is key to the following proposition (the proof is in the appendix).

PROPOSITION 2 (EX ANTE BOUND). *Assume that the optimal solution selected by the allocation engine $AE[\cdot]$ is always an extreme point and that the maximal load number p satisfies $n - p \geq m + 1$. Then, under the weak behavioral assumption, we have*

$$x_{ij}(T) > 0 \quad \Rightarrow \quad b_{ij}(T) \leq v_{p+m+1:n}^j + \epsilon, \quad (3)$$

and the buyer's final procurement cost is bounded by

$$\sum_{i=1}^n \sum_{j=1}^m b_{ij}(T) x_{ij}(T) \leq \sum_{j=1}^m q_j (v_{p+m+1:n}^j + \epsilon). \quad (4)$$

This last result is noteworthy in that it provides a performance guarantee for the buyer *before* the auction takes place, even though the behavioral assumptions required are particularly mild. As in (2), the upper bound (4) generalizes in the setting of a multidimensional optimization-based procurement auction the relation between winning bids and valuations in simpler auctions. However, it shows much more explicitly how this bound for the final bids depends on the primary data of the problem, namely the suppliers' capacity constraints and production costs and the buyer's requirements. In particular, the larger the buyer's required quantities in relation to the suppliers' available capacities, the looser the bound on the final bids; hence, this expression provides quantitative support for the intuition that the buyer's market power decreases as supply resources become more scarce. Likewise, the applicability condition $n - p \geq m + 1$ can be interpreted as a minimum requirement on the level of competition in the auction (via the number of participating suppliers) or on the buyer's

market power. Finally, the upper bound (4) helps reveal the relationship between the buyer's final procurement cost and its required quantities; note that in this model with linear production costs, an increase in the component market size may potentially increase the cost per unit because of the competitive effects.

3.3. Myopic Best-Response Analysis

We now assume that every supplier uses in each round the MBR calculation device and exactly follows its bidding suggestions, i.e., $\mathbf{b}_i(t+1) = \mathbf{b}_i^*(t+1)$ for all $i \in \{1, \dots, n\}$ and all $t \geq 0$, where $\mathbf{b}_i^*(t+1)$ is the expression defined in §2.4. Note that the MBR bidding model almost satisfies the weak behavioral assumption, the exception being the suppliers with no current potential allocation that may still not benefit from lowering any of their bids (e.g., when all competing bids are lower than his production costs), because in such cases the sup operator used in the definition of $\mathbf{b}_i^*(\cdot)$ will cause these suppliers' bids to remain constant. Nevertheless, the convergence bounds established in §3.2 remain valid under the MBR behavioral model, as can be quickly checked from their respective proofs.

3.3.1. Convergence of MBR Bidding Sequences.

From a mathematical standpoint, the MBR behavioral model amounts to a recursive relation between consecutive values of the bidding sequence $(\mathbf{b}(t))_{t \in \mathbb{N}}$ of the form $\mathbf{b}(t+1) = \mathcal{F}[\mathbf{b}(t)]$, $t \geq 0$. Because $\mathcal{F}[\cdot]$ is a well-defined function (see §2.4), the entire bidding history of the auction is completely characterized by the initial bids $\mathbf{b}(0)$ and this recursive relation. A property of the MBR bidding sequences essential to the theoretical validity of our solution concept is that they always converge in finite time (they are non-decreasing by the nonreneging rule, bounded from below, and may only take a finite set of values by the common multiple rule). This can be formalized by the following proposition.

PROPOSITION 3. *Let $(\mathbf{b}(t))_{t \in \mathbb{N}}$ be an MBR bidding sequence defined by a set of initial bids $\mathbf{b}(0)$ and the recursive relation $\mathbf{b}(t+1) = \mathcal{F}[\mathbf{b}(t)]$. There exists an integer $T \geq 0$ such that $\mathbf{b}(t) = \mathbf{b}(T)$ for all $t \geq T$.*

By the recursive definition of the MBR bidding sequence $(\mathbf{b}(t))_{t \in \mathbb{N}}$, its limit $\mathbf{b}(T)$ is necessarily a fixed point of the relation $\mathcal{F}[\cdot]$; i.e., $\mathbf{b}(T) = \mathcal{F}[\mathbf{b}(T)]$. Conversely, for every fixed point \mathbf{b} of $\mathcal{F}[\cdot]$, there obviously exists an MBR bidding sequence converging to \mathbf{b} (e.g., $\mathbf{b}(0) = \mathbf{b}$). The solution concept we propose to predict the final outcome in our market mechanism given an initial set of bids $\mathbf{b}(0)$ is precisely the limit $\mathbf{b}(T)$ of the MBR bidding sequence $(\mathbf{b}(t))_{t \in \mathbb{N}}$. This is the natural adaptation to our setting of the auction-specific local Nash equilibrium defined in Bykowsky et al. (2000) in the context of the simultaneous ascending auction.

3.3.2. Impact of Initial Bids. Although Propositions 1 and 2 show that the dependence of the auc-

tion equilibrium on the initial bids should be relatively limited when the level of competition among suppliers is sufficiently high, these results do not apply in market environments characterized by supply resource scarcity. To gain some insights into the impact of initial bids when the level of competition is low, we turn to the analysis of the MBR bidding dynamics in the simplest nontrivial market environment in which our smart-market mechanism can be applied: the 2×2 symmetric case, where $n = 2, m = 2, q_1 = q_2 = q, a_{ij} = 1 \forall (i, j)$, and $c_1 = c_2 = c$, but v_1 need not equal v_2 . The only interesting choices of parameters in this environment satisfy $1 < \theta \leq 2$, where $\theta = 2q/c$, which is the number of suppliers required to cover the total procurement needs.² For this range of θ , the buyer has no choice but to use a positive amount of capacity from both suppliers.

The following proposition shows that even with such scarce supply capacity, when the initial bids are sufficiently close to each other and are greater than the componentwise maximum of the production costs, they have relatively little impact on the final outcome. Moreover, it also shows that the particular information structure used in our mechanism may still enable some bidding competition. Its proof is omitted for the sake of brevity but appears in Gallien and Wein (2002).

PROPOSITION 4. Let

$$d(\mathbf{b}_1, \mathbf{b}_2) = \max(|b_{21} - b_{11}|, |b_{22} - b_{12}|).$$


In the symmetric 2×2 auction and under the MBR bidding model, there exists a selection rule discriminating among multiple optimal solutions to the allocation engine AE[...] such that

$$d(\mathbf{b}_1(0), \mathbf{b}_2(0)) \leq \epsilon \Rightarrow d(\mathbf{b}_i(T), \mathbf{v}_1 \vee \mathbf{v}_2) \leq \left(\frac{3q - c}{c - q} + 1 \right) \epsilon, \quad i \in \{1, 2\}.$$

A legitimate question at this point is to determine to what extent the final outcome is affected when the initial bids are further apart. Although we do not entirely resolve this issue here, it is revealing to plot (see Figure 1) for a particular choice of parameters ($c, q, \mathbf{v}_1, \mathbf{b}_1(0)$) the stability zones for supplier 1, i.e., the subset of supplier 2's bidding space such that supplier 1 would not move in the next round should $\mathbf{b}_2(0)$ belong to it.

Figure 1 shows the existence of premature equilibria, arising when the initial bids are too far apart, so that the increased order volume the higher bidder may receive by outbidding his competitor does not compensate for the margin reduction that would be incurred. For example, if supplier 2's bid was to

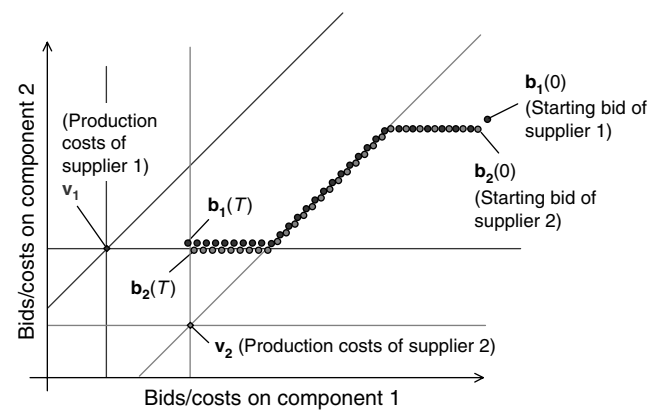
Figure 1 Supplier 1's Stability Zones (Shaded Areas)



lie in the lower-left quadrant of the stability zone, supplier 1 would not move in the following round by definition. Depending on supplier 2's production costs, it will either not move (thereby achieving a premature equilibrium) or try to switch the component type for which it received the larger volume. If these observations also apply to environments with more suppliers and component types, then there are implications for mechanism design: Whenever supply capacity is suspected to be scarce, it seems in the buyer's interest to have initial bids as close as possible, which could be achieved, for example, by enforcing a common entry bid. Hence, as in the cases of Propositions 1 and 2, Proposition 4 can be viewed as a generalization to a multi-item smart-market setting of a well-known result for the independent private-values auction: The reserve price only plays an important role when there are a few heterogeneous bidders.

3.3.3. Bidding Dynamics. When the initial bids are sufficiently close, the proof of Proposition 4 also reveals the structure of the bidding dynamics in the symmetric 2×2 auction according to the MBR model. This structure is illustrated in Figure 2, which depicts

Figure 2 Bidding Dynamics in the Symmetric 2×2 Auction




² In the case $\theta \leq 1$, the allocation of each component type is entirely independent from the other, and the case $\theta > 2$ is infeasible.

a typical MBR bidding sequence. Here, the initial bids are below the 45° lines going through the production costs of both suppliers, which we refer to as the margin switching lines for a reason that will soon become clear. As long as the bids remain in this area of the bidding plane, both suppliers try to out-bid their competitor primarily on component type 1 (so as to receive the full order quantity of this component), because it is the one with the higher relative margin. Thus, the bidding sequence moves horizontally, until it hits supplier 2's margin switching line, at which point supplier 2 starts competing primarily for component type 2. The bidding sequence then moves downward along this line until it hits the horizontal line going through supplier 1's cost for component type 2, at which point this supplier stops competing altogether on this component type (because its margin is reduced to zero). The bidding sequence starts moving horizontally again, before reaching an equilibrium in the neighborhood of $v_1 \vee v_2$, as predicted by Proposition 4.

Because of the combinatorial complexity involved (the number of possible bidding strategies for each supplier corresponds to the number of extreme points in the feasible polyhedron of $AE[\mathbf{b}]$), it seems hard to perform a similar analytical study of the MBR bidding dynamics in market environments with more suppliers and/or more component types. Instead, we undertook some numerical experiments simulating competitive interactions under the MBR behavioral model. The two experiments with eight suppliers and two component types represented in Figures 3 and 4 are fairly typical of the results we have obtained (see Gallien 2000 for more examples). These figures show each supplier's bidding sequence in the auction as a color-coded dotted line (a color version of these figures is available from the authors) and the corresponding production costs as diamonds. In addition, the set of connected squares is the sequence of dual prices associated in each round with the quantity requirement constraints in the formulation of the allocation engine $AE[\cdot]$ (referred to below as market prices according to the classical interpretation of LP duality theory). By design, these two experiments have the exact same parameters (capacity, required quantities, initial bids), with the exception of the production costs, which are more homogeneous across component types for the experiment³ reported in Figure 4.

In Figure 3, all suppliers except one exhibit a specialized pattern of bidding, where they compete by decreasing their bids for only one of the two components. However, in Figure 4, there are many more "diagonal" bidding paths; i.e., competition takes place


Figure 3 MBR Dynamics in 8×2 Auction with Heterogeneous Costs



for more suppliers along the two component axes. The primary driving force behind these different bidding patterns seems to be the relative position of each supplier's production costs in each round with respect to the market prices. More specifically, as the auction evolves, the path formed by these market prices partitions the suppliers' production costs into three categories. Suppliers with production costs roughly proportional to the market prices tend to compete for both components, while suppliers with production costs clearly above or below the path of market prices typically only compete for the one component on which their current relative margin is the highest (as defined by the relative position of their production costs and the current market prices). This interpretation explains why more homogeneous production costs across components (as in Figure 4) lead to a competition pattern that is more mixed.

3.3.4. Incentive Compatibility. In this subsection we numerically investigate the incentive compatibility issues linked to the revelation of suppliers' capacity and cost information in the auction mechanism

Figure 4 MBR Dynamics in 8×2 Auction with Homogeneous Costs



³ For each supplier, production costs in Figure 4 have been obtained from those of Figure 3 by a translation along the off diagonal.

Table 2 Revelation Incentives Experiments: Heterogeneous Costs

	Supplier #	1	2	3	4	5	6
	Capacity	35	45	50	30	42	60
	Production costs	9	11	6	10	5	8
Benchmark parameters	Initial bids	2	3	5	7	8	10
		16	13	12	16	14	16
		12	15	11	13	15	16
	Final bids	16	13	9	11	8	9
		8	8	8	8	9	15
	Final allocation	0	0	38	0	42	0
		35	45	0	0	0	0
	Final payoff	210	225	114	0	126	0
	Capacity input distortion (%)	Resulting final payoff (% change from benchmark)					
Capacity revelation incentives	+25	*262.5 (+25)	225 (+0)	76 (−33)	0 (n/a)	*105 (−17)	0 (n/a)
	+20	*252 (+20)	225 (+0)	76 (−33)	0 (n/a)	*100.8 (−20)	0 (n/a)
	+15	*241.5 (+15)	225 (+0)	76 (−33)	0 (n/a)	*96.6 (−23)	0 (n/a)
	+10	*231 (+10)	225 (+0)	114 (+0)	0 (n/a)	*138.6 (+10)	0 (n/a)
	+5	*220.5 (+5)	225 (+0)	114 (+0)	0 (n/a)	*132.3 (+5)	0 (n/a)
	−5	199.5 (−5)	213.75 (−5)	142.5 (+25)	0 (n/a)	120 (−5)	0 (n/a)
	−10	189 (−10)	202.5 (−10)	135 (+18)	0 (n/a)	120 (−5)	0 (n/a)
	−15	178.5 (−15)	191.25 (−15)	127.5 (+12)	0 (n/a)	120 (−5)	0 (n/a)
	−20	168 (−20)	180 (−20)	120 (+5)	0 (n/a)	120 (−5)	0 (n/a)
	−25	157.5 (−25)	168.75 (−25)	112.5 (−1)	0 (n/a)	120 (−5)	0 (n/a)
	Costs input distortion (%)	Resulting final payoff (% change from benchmark)					
Costs revelation incentives	+25	210 (+0)	0 (−100)	114 (+0)	0 (n/a)	120 (−5)	0 (n/a)
	+20	210 (+0)	0 (−100)	114 (+0)	0 (n/a)	120 (−5)	0 (n/a)
	+15	210 (+0)	225 (+0)	114 (+0)	0 (n/a)	120 (−5)	0 (n/a)
	+10	210 (+0)	225 (+0)	114 (+0)	0 (n/a)	120 (−5)	0 (n/a)
	+5	210 (+0)	225 (+0)	114 (+0)	0 (n/a)	126 (+0)	0 (n/a)
	−5	210 (+0)	225 (+0)	114 (+0)	0 (n/a)	126 (+0)	0 (n/a)
	−10	210 (+0)	225 (+0)	114 (+0)	0 (n/a)	126 (+0)	0 (n/a)
	−15	210 (+0)	225 (+0)	114 (+0)	0 (n/a)	126 (+0)	0 (n/a)
	−20	210 (+0)	225 (+0)	100 (−12)	0 (n/a)	126 (+0)	0 (n/a)
	−25	210 (+0)	225 (+0)	100 (−12)	0 (n/a)	126 (+0)	0 (n/a)

Note. An asterisk (*) before an entry indicates an infeasible allocation where the corresponding supplier is overloaded.

with MBR bidding. We simulate two scenarios: Each has six bidders with the same production technology (i.e., $a_{ij} = 1 \forall i, j$), two components with identical demand $q_1 = q_2 = 80$, a common bid multiple $\epsilon = 1$, and a tie-breaking rule that favors lower-indexed suppliers. The first scenario corresponds to relatively heterogeneous costs and the second to relatively homogeneous costs. For each scenario and each given set of initial bids, we simulate the MBR bidding sequence and the final outcome for the benchmark case where each supplier truthfully inputs its private capacity and cost information, and for the cases where each supplier misrepresents in varying proportions either costs or total capacity, all other parameters (including the other suppliers' information input) remain unchanged. The parameters and results of these experiments are summarized in Tables 2 (heterogeneous costs) and 3 (homogeneous costs). Note that some of the interpretations later in this section rely on a more detailed presentation of these results

(including final allocations and bids), which can be obtained from the authors on request.

Capacity Revelation Incentives. Analysis of the results in Tables 2 and 3 reveals the relative importance of *strategic* incentive effects, caused by the impact of a capacity statement on the other suppliers' bidding behavior, versus *nonstrategic* incentive effects, caused by a supplier's capacity statement on its own final margin or awarded quantity. In all cases where the percentage change (from baseline) in a supplier's final payoff is equal to the percentage by which its capacity input was distorted (e.g., supplier 1 in both tables), only nonstrategic effects occurred: The final bids of all suppliers were unchanged, and the change in the supplier's payoff was entirely due to its final awarded quantity. An important observation in Tables 2 and 3 is that a supplier cannot benefit by overstating capacity without becoming overloaded and incurring an observable default (i.e., refusing to accept all of the work it has been awarded). Because capacity over-

Table 3 Revelation Incentives Experiments: Homogeneous Costs

Supplier #		1	2	3	4	5	6
Benchmark parameters	Capacity	35	45	50	30	42	60
	Production costs	7	8	6	8	7	9
		4	6	5	9	6	9
	Initial bids	16	13	12	16	14	16
		12	15	11	13	15	16
	Final bids	9	10	9	9	8	10
		7	7	8	13	7	10
	Final allocation	0	0	38	0	42	0
	35	45	0	0	0	0	
Final payoff	105	45	114	0	42	0	
Capacity input distortion (%)		Resulting final payoff (% change from benchmark)					
Capacity revelation incentives	+25	*131.25 (+25)	45 (+0)	*125 (+10)	0 (n/a)	30 (−29)	0 (n/a)
	+20	*126 (+20)	45 (+0)	*120 (+5)	0 (n/a)	30 (−29)	0 (n/a)
	+15	*120.75 (+15)	45 (+0)	*115 (+1)	0 (n/a)	30 (−29)	0 (n/a)
	+10	*115.5 (+10)	*49.5 (+10)	114 (+0)	0 (n/a)	*46.2 (+10)	0 (n/a)
	+5	*110.25 (+5)	*47.25 (+5)	114 (+0)	0 (n/a)	*44.1 (+5)	0 (n/a)
	−5	100 (−5)	42.75 (−5)	114 (+0)	0 (n/a)	39.9 (−5)	0 (n/a)
	−10	94.5 (−10)	40.5 (−10)	114 (+0)	0 (n/a)	37.8 (−10)	0 (n/a)
	−15	89.25 (−15)	38.25 (−15)	114 (+0)	0 (n/a)	35.7 (−15)	0 (n/a)
	−20	112 (+7)	36 (−20)	114 (+0)	0 (n/a)	33.6 (−20)	0 (n/a)
	−25	78.75 (−25)	33.75 (−25)	112.5 (−1)	0 (n/a)	31.5 (−25)	0 (n/a)
Costs input distortion (%)		Resulting final payoff (% change from benchmark)					
Costs revelation incentives	+25	105 (+0)	90 (+100)	114 (+0)	0 (n/a)	0 (−100)	0 (n/a)
	+20	105 (+0)	6 (−87)	114 (+0)	0 (n/a)	0 (−100)	0 (n/a)
	+15	105 (+0)	45 (+0)	114 (+0)	0 (n/a)	0 (−100)	0 (n/a)
	+10	105 (+0)	45 (+0)	114 (+0)	0 (n/a)	42 (+0)	0 (n/a)
	+5	105 (+0)	45 (+0)	114 (+0)	0 (n/a)	42 (+0)	0 (n/a)
	−5	105 (+0)	48 (+7)	114 (+0)	0 (n/a)	42 (+0)	0 (n/a)
	−10	105 (+0)	48 (+7)	114 (+0)	0 (n/a)	42 (+0)	0 (n/a)
	−15	105 (+0)	48 (+7)	114 (+0)	0 (n/a)	42 (+0)	0 (n/a)
	−20	105 (+0)	48 (+7)	100 (−12)	0 (n/a)	42 (+0)	0 (n/a)
	−25	105 (+0)	48 (+7)	100 (−12)	0 (n/a)	42 (+0)	−120 (n/a)

Note. An asterisk (*) before an entry indicates an infeasible allocation where the corresponding supplier is overloaded.

loading is eventually observable by definition, one should be able to prevent it by imposing a heavy penalty in the event of misdeliveries and/or quality problems; e.g., Freemarkets keeps instances of capacity overloading to a minimum by suspending from subsequent bidding events overloaded suppliers that default after an auction (Wnorowski 2000). It is this punitive action that led to the observed impact of overload fear on suppliers' bidding behavior, which in turn motivated the design of our smart-market approach (see §1).

A more subtle type of nonstrategic effect occurs when a supplier's optimal response to its own distorted capacity input is different than in the benchmark case. This can be observed in Table 2 with supplier 5, which sees its payoff decrease by about 5% (from 126 to 120) for all capacity understatements. This is because lowering its capacity of 42 to less than 40 does not leave enough volume to make up for an

aggressive final bid of 8 and a resulting margin of 3; the supplier instead chooses to secure a higher final margin of 4 with a final bid of 9, which makes up for the lower final quantity of 30, as supplier 3 is loaded to capacity at 50 for this component (the threshold occurs with a capacity statement of 40 because $4 \times 30 = 3 \times 40$).

The strategic effects linked to capacity revelation in Tables 2 and 3 frequently correspond to the phenomenon of demand reduction described in Ausubel and Cramton (1998). For example, supplier 3 in Table 2 benefits from understating his capacity by at least 3% to less than 48.5 because this reduces the bidding aggressiveness of supplier 5 for the same component 1; supplier 5 can then get more than $80 - 48.5 = 31.5$ units with a higher margin of 4 (against 3 previously), which more than compensates for the fact that supplier 5 is no longer loaded up to capacity of 42 because $4 \times 31.5 = 3 \times 42$. As a result,

supplier 3 is then loaded to stated capacity with the same high margin as before—a volume increase in the awarded quantity of up to 27.6%. Finally, when the costs are homogeneous, as in Table 3, one can see yet another type of strategic incentive (incentive of supplier 1 to understate its capacity by 20%), linked to a substitution effect across components.

The experimental results described above reveal that understating capacity may either profit or penalize a supplier, depending on the particular situation. Although it does not seem easy, one could try to construct a mapping from the suppliers' information sets during the auction (described in §2.4) to the likelihood of a payoff increase through understating capacity. However, the possibility of such inference is considerably reduced when, as is the case in the mechanism we propose, suppliers are required to report their capacity before the auction begins. We thus believe that in most cases, a successful capacity misrepresentation by suppliers would require more information (e.g., the other suppliers' capacity constraints), rationality, and risk taking than would typically be observed in practice.

Costs Revelation Incentives. Of the 120 experiments we conducted, the final payoff of a supplier misrepresenting its production costs increased in 6 cases (5%), decreased in 15 cases (12.5%), and remained unchanged in 99 cases (82.5%). In addition, all payoff increases corresponded to the same supplier in the same scenario, which may indicate a somewhat pathological case.

Besides one instance of a supplier understating its costs so much that its payoff ended up being negative (supplier 6 in Table 3), all misrepresentation penalties resulted from two main types of effects. A "cut-off" effect occurs when a supplier overstates its costs above a certain threshold, forcing the MBR bidding agent to stop bidding early (under the false impression that the supplier's margin would become negative otherwise), when the bids are still too high to ensure any allocation. This effect can be observed with supplier 2 in Table 2, when the costs are distorted by more than 20%, and with suppliers 2 and 5 in Table 3, when they inflate their costs by 20% and more than 15%, respectively. A more subtle effect occurs when misrepresented costs bias the bidding agent's trade-off assessment between a higher margin with a lower volume against a lower margin with a higher volume. This effect can follow from both over- and under-representations, e.g., supplier 3 when understating its costs by more than 20% in Tables 2 and 3, and supplier 4 when overstating its costs by more than 10% in Table 2.

All payoff increases linked to cost distortion occurred with supplier 2 in Table 3 and were linked to substitution effects across components. The payoff

increase when the supplier understates his costs by at least 5% occurs because the cost on component 1 is greater than on component 2, so that the margin on component 1 is proportionally increased more, resulting in an inversion of the components' relative attractiveness: The MBR bidding agent bids more aggressively on component 1 and less aggressively on component 2, and through a strategic effect supplier 5 finds it beneficial to do the opposite. However, supplier 2 then benefits from both its tie rule advantage and the fact that supplier 5 has a slightly lower capacity, so that supplier 2 not only gets the 42 units of component 1 previously allocated to supplier 5 as part of this component swap, but also 3 units of component 2 at a slightly higher margin, accounting for the 7% payoff increase. An intriguing 100% payoff increase occurs in the same scenario when supplier 2 over-represents his costs by 25%. This 25% cost distortion produces a less-aggressive bidding behavior, thus softening the strategic response of other bidders and leaving supplier 2 with a higher final margin (2 instead of 1). At the same time, supplier 2's bidding remains just aggressive enough so that it still gets the same final allocation of 45 units. A fortuitous consequence of the bidding agent's myopia, this effect seems dangerous to speculate on, as an additional cost inflation by just 7% to 33% results in supplier 2 losing its entire allocation and a lower-cost inflation by just 1% to 24% results in an 87% payoff reduction.

4. Conclusion

Partly motivated by the observed inefficiencies of current online procurement auction mechanisms when suppliers' capacity constraints are stringent, this paper investigates an alternative mechanism designed specifically for such environments. It relies on an estimation of these capacity constraints prior to the bidding event, which enables the use of an optimization-based allocation engine (justifying the term "smart market"). For every set of bids submitted during the auction, this engine dynamically computes the allocation of procurement contracts minimizing the buyer's total cost under both the buyer's quantity requirement and the suppliers' capacity constraints. The bidders' information sets in this mechanism are designed to achieve a middle ground between facilitating competition and preventing collusion, exploiting some of the possibilities offered by the use of a distributed network of computers to support the auction. They include in particular an original MBR bidding-suggestion device, taking as an input a supplier's production costs and delivering as an output the set of bids maximizing the supplier's potential profit in the next round under the assumption that the competitors' bids remain constant.

The solution concept we propose for this auction mechanism is based on tatonnement theory, whereby

we first formulate less than fully rational behavioral models for the bidders and then examine the implications of these models on the dynamics and the convergence properties of the resulting bidding sequences. Under a weak behavioral assumption loosely characterizing the reactions of potentially rejected suppliers and a minimum competition-level condition, we construct upper bounds for the bids of the winning suppliers at the end of the auction, providing in turn an upper bound on the buyer's final procurement cost. These bounds are shown to carry over to the complete behavioral model, in which the suppliers follow the MBR bid suggestions. Some additional structural properties of the MBR bidding sequences in a symmetric (except for production costs) market with two bidders and two component types suggest that when supplier capacity is scarce (i.e., competition is low), the buyer may find it beneficial to impose a common entry bid to discourage premature equilibria generated by initial bids that are far apart. These results can all be viewed as generalizations of classic results to the setting of a multi-item smart-market environment. Finally, numerical simulations of MBR sequences in more complex market configurations imply that suppliers would find it difficult to profit by distorting their private capacity and cost information.

We believe this work constitutes a useful basic framework toward the development of electronic trade systems enabling real-time complex industrial transactions. Within the scope of our model, natural steps in this direction include more complex capacity constraints, nonlinear production costs and bid structures, requirements on the supply base size, and incumbent switching costs; this last generalization is particularly important, because these switching costs (which could, more generally, incorporate the bidder's perceptions of the suppliers' reputations) are the primary reason why Freemarkets makes its allocation decisions after—rather than during—its bidding events. In the framework of our model, one way for the buyer to account for its relative preferences across suppliers is to adjust the actual bids b_{ij} by an additive factor $\alpha_{ij} > 0$ for the purpose of computing potential allocations. That is, while the selling prices resulting from the auction still correspond to the actual bids b_{ij} submitted by the winners, the potential and final allocations of components are now calculated by solving $\mathbf{x}(t) = AE[\mathbf{b}(t) - \boldsymbol{\alpha}]$; the higher α_{ij} , the greater supplier i 's advantage relative to his competitors for the allocation of component type j . Some of our results can be generalized to include these adjusted bids. For example, the right side of the ex ante bound in Proposition 2 becomes $[\mathbf{v} - \boldsymbol{\alpha}]_{m+p+1:n}^j + \epsilon + \alpha_{ij}$, where $[\mathbf{v} - \boldsymbol{\alpha}]_{k:n}^j$ denotes the k th order statistic of $(v_{1j} - \alpha_{1j}, \dots, v_{nj} - \alpha_{nj})$. In other words, modifying the bids for allocation purposes through the mechanism just

described roughly amounts to modifying the production costs by the same adjustment factors, as one would intuitively expect.

Unfortunately, this adjustment model does not allow us to capture fixed switching costs (i.e., independent of the quantity allocated). From a practical standpoint, this particular extension along with most of the others previously mentioned can be incorporated by using more sophisticated integer programming formulations for the allocation engine $AE[\cdot]$ in (1) (see Gallien and Wein 2002 for details), although we have not derived theoretical results corresponding to Propositions 1–4 in these more complex settings. Another challenge would be to incorporate into our allocation engine more criteria specific to industrial procurement, such as delivery performance, quality, and insurance provisions, leading perhaps to the design of a multidimensional bidding mechanism. Such complex and dynamic trading environments will require the design of decision-support tools available to bidders, for example by generalizing our MBR bid computational device. From a methodological perspective, we believe that formulating behavioral models associated with these decision-support tools is a fruitful approach to predicting various mechanism outcomes.

An online companion (Gallien and Wein 2002) to this paper is available at <http://mansci.pubs.informs.org/ecompanion.html>.

Acknowledgments

The authors thank Sam Kinney and Doug Wnorowski for sharing details about the operations of FreeMarkets; Jim Orlin for valuable suggestions leading to the weakening of some of the initial assumptions for the convergence results reported in §3; and Yashan Wang, Sridhar Tayur, John Serman, Erik Brynjolfsson, Nelson Repenning, Rachel Croson, Robert Porter, Glenn Ellison, and Andreas Schulz for provocative discussions and/or suggestions of relevant references. They are also grateful to Ozlem Ergun for reviewing some conjectures, and the referees and Mike Rothkopf for helpful suggestions on an earlier draft. This research was supported in part by the Singapore-MIT Alliance and a Ph.D. fellowship from the Sloan School of Management at MIT.

Appendix

A.1. PROOF OF PROPOSITION 1. Let $\mathbf{z}(T)$ denote the value in round T of the optimal dual variables associated with the quantity constraints in $AE[\mathbf{b}(T)]$ and \mathbf{e}_m an m -dimensional vector with all its components equal to 1. Then, we have

$$i \in P^C \Rightarrow \mathbf{v}_i \geq \mathbf{z}(T) - \epsilon \mathbf{e}_m. \quad (5)$$

This is because $\mathbf{b}_i(T) \geq \mathbf{z}(T)$ for any supplier $i \in P^C$, so that if we had $v_{ij} < z_j(T) - \epsilon$ for some j , then supplier i could increase its profit (currently equal to zero) by bidding $b_{ij}(T+1) = z_j(T) - \epsilon < b_{ij}(T)$. But $b_{ij}(T+1) = b_{ij}(T)$ by definition of T , and therefore $z_j(T) - \epsilon \leq v_{ij} \forall j$. This last statement applies to at least $|P^C|$ suppliers, so that

$$z_j(T) \leq v_{n-|P^C|+1:n}^j + \epsilon. \quad (6)$$

Finally, $x_{ij}(T)$ is the dual variable associated with the constraint $z_j(T) \leq b_{ij}(T) + a_{ij}y_i$ of the dual to $AE[\mathbf{b}(T)]$, where $y_i \geq 0$ is the dual variable associated with the capacity constraint of supplier i in $AE[\mathbf{b}(T)]$. Therefore, $x_{ij}(T) > 0$ implies $b_{ij}(T) \leq z_j(T)$, which combined with (6) completes the proof. \square

A.2. PROOF OF PROPOSITION 2. The proof amounts to finding a lower bound for $|P^C|$ (or equivalently, an upper bound for $|P|$), and then invoking Proposition 1 in the cases where this bound allows us to claim that $|P^C| \geq 1$. In any feasible allocation, the number of suppliers that are fully loaded is clearly smaller than the maximal load factor p , because of the quantity requirement in (1). We now claim that in any extreme point of the polytope of feasible solutions to (1), the number of suppliers that are not fully loaded but still receive a positive allocation is smaller than m . This is because the allocation would otherwise include at least two nonfully loaded suppliers receiving a positive allocation of the same product and would therefore not be an extreme point (as it could be written as a convex combination with positive weights of two distinct feasible solutions). Therefore, among all the suppliers receiving a positive allocation of each product, at most one can be nonfully loaded, and the total number of nonfully loaded suppliers is smaller than m . This completes the proof of $|P| \leq p + m$, from which (3) follows immediately by Proposition 1. The bound on the buyer's final procurement cost is a direct consequence of (3). \square

References

- Aeppel, T. 1999. Bidding for e-nuts and e-bolts on the net. *Wall Street Journal* (March 12) B1.
- Ausubel, M., P. Cramton. 1998. Demand reduction and inefficiency in multi-unit auctions. Working paper 96-07, University of Maryland, College Park, MD.
- Ausubel, M., P. Cramton. 1999. The optimality of being efficient. Working paper, Department of Economics, University of Maryland, College Park, MD.
- Berstekas, D. P., D. A. Castanon. 1989. The auction algorithm for the transportation problem. *Ann. Oper. Res.* **20** 67–96.
- Bikhchandani, S., S. de Vries, J. Schummer, R. V. Vohra. 2002. Linear programming and Vickrey auctions. B. Dietrich, R. Vohra, eds. *Mathematics and the Internet: E-Auction and Markets*. Springer-Verlag, New York, 75–115.
- Budde, J., R. F. Göx. 1999. The impact of capacity costs on bidding strategies in procurement auctions. *Rev. Accounting Stud.* **4**(1) 5–13.
- Bykowsky, M. M., R. J. Cull, J. O. Ledyard. 2000. Mutually destructive bidding: The FCC auction design problem. *J. Regulatory Econom.* **17** 205–228.
- Cramton, P. 1998. Ascending auctions. *Eur. Econom. Rev.* **42** 745–756.
- Demange, G., D. Gale, M. Sotomayor. 1986. Multi-item auctions. *J. Political Econom.* **94** 863–872.
- de Vries, S., R. Vohra. 2003. Combinatorial auctions: A survey. *INFORMS J. Comput.* **15**(3) 284–309.
- Elmaghraby, W. 2002. The importance of ordering in sequential auctions. *Management Sci.* **49**(5) 673–682.
- Fudenberg, D., D. K. Levine. 1998. *The Theory of Learning in Games*. MIT Press, Cambridge, MA.
- Fudenberg, D., J. Tirole. 1996. *Game Theory*. MIT Press, Cambridge, MA.
- Gallien, J. 2000. Optimization-based auctions and stochastic assembly replenishment policies for industrial procurement. Ph.D. thesis, Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA.
- Gallien, J., L. M. Wein. 2002. A smart market for industrial procurement with capacity constraints: Appendix. <http://mansci.pubs.informs.org/ecompanion.html>.
- Harstad, R. M., J. Kagel, D. Levin. 1990. Equilibrium bid functions for auctions with an uncertain number of bidders. *Econom. Lett.* **33** 35–40.
- Jahnke, A. 1998. How bazaar-purchasing strategies. *CIO Web Bus. Magazine*, http://www.cio.com/archive/webbusiness/080198_freemarkets_content.html.
- Kafka, S. J., B. Temkin, L. Wegner. 2000. *B2B auctions go beyond price*. Forrester Research, Inc., Cambridge, MA.
- Kelly, F., R. Steinberg. 2000. A combinatorial auction with multiple winners for universal service. *Management Sci.* **46** 586–596.
- Kinney, S. 1999. Personal communication, February 18.
- Klemperer, P. D. 1999. Auction theory: A guide to the literature. *J. Econom. Surveys* **13** 227–286.
- Larson, K., T. Sandholm. 2001. Costly valuation computation in auctions: Deliberation equilibrium. *Theoretical Aspects of Reasoning about Knowledge (TARK)*. Siena, Italy, 169–182.
- Lucking-Reiley, D. 2000. Auctions on the Internet: What's being auctioned, and how? *J. Indust. Econom.* **48** 227–252.
- McCabe, K. A., S. J. Rassenti, V. L. Smith. 1991. Smart computer-assisted markets. *Science* **254** 534–538.
- Milgrom, P. 2000. Putting auction theory to work: The simultaneous ascending auction. *J. Political Econom.* **108** 245–272.
- Parkes, D. C., L. H. Ungar. 2000. Iterative combinatorial auctions: Theory and practice. *Proc. 17th National Conf. Artificial Intelligence (AAAI-00)*, The AAAI Press, Menlo Park, CA, 74–81.
- Rangan, V. K. 1999. FreeMarkets OnLine. Harvard Business School Case 9-598-109, Harvard University, Cambridge, MA.
- Rothkopf, M. H., R. M. Harstad. 1994. On the role of discrete bid levels in oral auctions. *Eur. J. Oper. Res.* **74** 572–581.
- Rothkopf, M. H., A. Pekec, R. M. Harstad. 1998. Computationally manageable combinatorial auctions. *Management Sci.* **44** 1131–1147.
- Samuelson, L. 1997. *Evolutionary Games and Equilibrium Selection*. MIT Press, Cambridge, MA.
- Stanley, E. D., D. P. Honig, L. Gainen. 1954. Linear programming in bid evaluation. *Naval Res. Logist. Quart.* **1** 48–54.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* **16** 8–37.
- Wellman, M. P., W. E. Walsh, P. R. Wurman, J. K. MacKie-Mason. 2001. Auction protocols for decentralized scheduling. *Games Econom. Behavior* **35** 271–303.
- Wnorowski, D. M. 2000. Personal communication. (July 12).