

Combinatorial and Quantity-Discount Procurement Auctions Benefit Mars, Incorporated and Its Suppliers

Gail Hohner • John Rich • Ed Ng
Mars, Incorporated, Winnersh Triangle, United Kingdom

Grant Reid
Mars, Incorporated, McLean, Virginia

Andrew J. Davenport • Jayant R. Kalagnanam • Ho Soo Lee • Chae An
IBM T.J. Watson Research Center, PO Box 218, Yorktown Heights, New York 10598
gail.hohner@eu.ibm.com • davenport@us.ibm.com • jayant@us.ibm.com
leehs@us.ibm.com • chae@us.ibm.com

Simple auctions neglect the complex business constraints required by strategic sourcing. The Mars-IBM team created a procurement auction Web site (www.number1traders.com) that enables buyers to incorporate complex bid structures (such as bundled all-or-nothing bids and quantity-discounted bids) and business constraints into strategic-sourcing auctions. Outcomes in such auctions must lead to win-win solutions to sustain long-term relationships between procurer and suppliers. These factors are as important or more important than price. The Mars procurement auction Web site supports several alternatives to simple auctions that help match its needs as procurer and the capabilities of suppliers by incorporating optimal bid selection subject to constraints based on business rules in a dynamic environment. The ability to consider geographic, volume, and quality factors helps both parties. Feedback from participant suppliers has highlighted the benefits of time efficiency, transparency, and fairness. Although they reflect just one side of the benefits ledger, the monetary benefits to Mars (a \$14 billion company) and to its suppliers are significant.

(Industries: agriculture, food. Games/group decisions: bidding, auctions.)

Mars is one of the world's largest privately owned businesses (\$14 billion in sales). Its philosophy, culture, and style set it apart from other organizations. From its earliest beginnings making buttercream candies sold door to door, Mars operations have grown to include global businesses in food, pet care, drinks vending, and electronic automated payment systems. Today the company produces many top brands of confectionery, pet food, and rice, including Mars[®], M&M's[®], Snickers[®], Whiskas[®], Pedigree[®], and Uncle Ben's[®]. To survive, prosper, and grow in its competitive markets, it must develop steady brand loyalty by providing consistently high quality to the

consumer. To prosper, Mars must seek the best value for its money in obtaining its inputs. The food and pet-food industries run at much lower margins than many manufacturing or service industries, making the control of input costs vitally important. Mars needs to maintain a base of long-term and reliable suppliers and at the same time to design procurement mechanisms to obtain the best value for its money without adversely affecting its long-term relationships with suppliers. This is the central focus of procurement at Mars.

Mars bases its worldwide business in over 100 countries on five continents on five core principles: quality, efficiency, responsibility, mutuality, and freedom. It

calls its employees associates and spreads responsibility out into the organization. In procurement, its five principles guide the firm's management of supplier relationships and its negotiation protocols.

Mars' Business Environment and Supplier Relationships

Mars relies on a small number of suppliers for each material it procures. Small supply pools arise by necessity as well as by design. For example, many agricultural inputs are available from a limited number of origins, a limited number of brokers, or under tariff regimes that limit available supplies. Integrated supply pools with highly valued suppliers reinforce a corporate culture based on mutually shared benefits. Buyers responsible for developing new sources and contract conditions (including price) maintain relationships with suppliers. Single buyers are responsible for large portfolios often covering two or three product groups for a large region (for example, all of Europe) or nearly all inputs required for a specific production site. A buyer may be responsible for up to 50 relationships.

Mars and its suppliers agree on contract conditions and prices in a number of ways. Mars has no one-size-fits-all solution. Mars contracts with private businesses, brokers, traded agricultural markets, monopolies, cartels, and governments. It uses different buying techniques to handle a huge number of purchasing situations. Negotiation and sealed-bid tendering are the most common. Procurement processes based on many parallel one-to-one negotiations or on single sealed-bid tenders have a number of fundamental problems:

(1) One-to-one negotiations prevent the procurement division from fully leveraging the competitive environment across suppliers to negotiate prices. Single sealed-bid tenders are even more static than one-to-one negotiations.

(2) Status quo methods give the firm no effective way of taking advantage of synergies or economies of scale that suppliers may have for parts of the total business.

(3) In negotiations, buyers spend disproportionate amounts of time determining prices and quantities,

which they could do far more efficiently using auctions.

(4) The end-point of the negotiation process is somewhat arbitrary (based on how much time the buyer invests to obtain future savings). The lack of transparency in such negotiations can cause disgruntled suppliers. For example, a supplier may discover after negotiations that the buyer has contracted with a price it would have been willing to better. The buyer may have accepted an offer without checking all suppliers in the pool.

Electronic auctions have emerged as a popular mechanism for conducting negotiations. Procurement auctions take the form of reverse auctions with a single buyer and a set of precertified suppliers negotiating within the context of a private exchange. Mars wanted to harness the efficiencies of auction mechanisms for procurement while capturing some of the complexities of the bid structures it allowed and the constraints arising from its strategies for maintaining supplier relationships.

Auctions are generally thought to promote market competition and make the purchasing process efficient but have two disadvantages: (1) they are too reliant on price, making them a brute-force way of managing relationships, and (2) they are inappropriate when the firm wants to control business volumes or the number of suppliers. For any auction system to gain wide acceptance, it would have to address these two disadvantages.

Problem Definition

The goal of this project was to enable Mars buyers worldwide to run auctions whose design complimented the overall business strategy. A typical buyer has a wide range of materials to procure. Some of these are very straightforward, simple purchases with a single contract with a sole source (for example, for office supplies). Our goal was not to service this type of purchase. Many auction suppliers in the market handle such purchases. Our objective was to support the strategic purchases. Strategic purchases are typically characterized by (1) small and fairly static supply pools, (2) long-term relationships, and (3) significant business integration. The contracts in strategic purchases typically

are of high value, are renewed quarterly or annually, and require the use of special price-negotiation schemes that incorporate appropriate business practices. Typically, bids in these settings have the following properties:

(1) The transaction volume is large, and the suppliers provide volume discounts that they specify as a curve with a quantity range associated with each price level (for example, \$1,000 per unit up to 100 units, \$750 per unit over 100 units), and

(2) Often the suppliers make all-or-nothing bids on a set of items, offering a special discounted price on a bundle (for example, \$150 for 30 units of item A and

It takes a long time to change the way people think!

20 units of item B) and will not sell the items partially or separately.

Two auction mechanisms incorporate these bid structures and a set of business requirements central to strategic procurement:

(1) Supply-curve auctions are specifically tailored to industries in which volumetric discounts are common, for example, bulk chemicals and agricultural commodities. In a supply-curve auction, suppliers specify these bids as curves with a quantity range associated with each price level (for example, \$1,000 per unit up to 100 units and \$750 per unit over 100 units). Such auctions may deal with one product or many.

(2) Combinatorial auctions are multilot auctions that allow bids for combined lots with prices that are conditional on winning the entire lot. Combinatorial auctions are ideal for situations in which synergies exist between lots (for example, specifically located freight routes). Suppliers provide all-or-nothing bids on a set of items, offering a special discounted price on a bundle. Bids in this auction can be overlapping (for example, bid 1 for A, bid 2 for A + B, bid 3 for A + B + D). These auctions are also useful for situations in which the firm wants to buy small volumes of similar but not identical materials at once. The combinations allow it to aggregate business to a level sufficient to interest suppliers.

Business Rules

After receiving bids in such auctions, a Mars' buyer must identify the set of bids that minimizes total procurement cost subject to the following business rules:

(1) The number of winning suppliers must be at least a minimum number to avoid Mars depending too heavily on just a few suppliers.

(2) The number of winning suppliers must be at most a maximum number to avoid the administrative overhead of managing a large number of suppliers.

(3) The maximum amount procured from each supplier must be bounded to limit exposure to a single supplier.

The minimum amount to be procured from each supplier must be bounded to avoid receiving economically inefficient orders (for example, less than a full truck load), and if two or more sets of bids could win, the buyer must pick the set that arrived first. (This is imperative if the suppliers are to see the solution as fair.) Although it might seem computationally expedient, approximate solutions are unacceptable because the difference between an approximate solution and the real solution can change how much and exactly which business a single supplier receives. For example, a supplier allocated business in the optimal solution might get nothing in an approximate solution. Such occurrences, if made public, would destroy the credibility of an auction mechanism. Identifying the cost-minimizing bid set subject to these business rules is a hard optimization problem; a typical problem with 15 items and 10 suppliers is usually too complex to do by hand and is beyond most buyers' programming ability.

Previous Work

Researchers have studied procurement problems from both the seller's and buyer's points of view (Stanley et al. 1954, Kim and Huang 1988, Lee and Rosenblatt 1986, Narsimhan and Stoyloff 1986). The use of linear programming for evaluating bids is almost as old as the theory of linear programming itself (Stanley et al. 1954). The inventory-control and lot-sizing literature includes studies of various discounting schemes, such as unit discounts (Silverman and Peterson 1979), incremental quantity discounts and carload quantity discounts (Jucker and Rosenblatt 1985, Lee and Rosenblatt

1986). More recently, Katz et al. (1994) studied business-volume discounts (across several commodities). The focus of all this work was to build decision models for selecting bids. The discounting models we developed are similar to those of Sadrian and Yoon (1994) but differ in three ways: (1) the discounts are based on quantities for each item, (2) the total-allocation constraint on each supplier binds the allocation across different items into a single large optimization problem, and (3) a minimum number of suppliers is specified that further constrains the choice of winners. Another qualitatively new practice we introduced is the use of bundled bids in a reverse auction across different items, which leads to the set-covering problem. Combinatorial auctions have been proposed for forward auctions, such as the FCC spectrum rights auctions (Rothkopf et al. 1998, Park and Rothkopf 2001). Traditionally, buyers have used sequential negotiations or auctions when procuring multiple items. For suppliers with cost complementarities, this could lead to exposure problems. Suppliers who could provide an overall low cost for multiple items might refrain from bidding aggressively because their cost structures are contingent on winning subsequent auctions. Their reluctance to bid aggressively could lead to inefficient outcomes.

However, the more fundamental difference we introduced is the iterative negotiation process for both the bundled and quantity-discounted bids. Gallien and Wein (2001) discussed iterative approaches for procurement. However, they treated bid evaluation as a linear-programming problem (as did Stanley et al. 1954), which allows fractional allocation of a demanded lot to a supplier. We were restricted to integral allocations and hence needed to solve an integer-programming problem. This led to a reverse-procurement auction that introduces an entirely new dynamic (competition among suppliers, transparency, and efficiency) into the negotiation process. The success of this approach makes it a forerunner of widespread change in procurement practice.

None of the commercial products available today is able to provide iterative optimization and conform to all of Mars' business-rule constraints in a dynamic environment. Even those with more sophisticated optimization offerings do not support all of Mars' business

rules, provide quick solutions, and prioritize bids according to the time received.

Mars' Design for Procurement Auctions

Complex bid structures necessitated an iterative auction design, largely because completely specifying the cost structure using bundled bids or supply-curve bids can result in an exponentially large bid set. An iterative format also (1) induces competition among suppliers, (2) allows suppliers to correct their bids using information learned during the process, and (3) elicits bids incrementally so that the suppliers do not have to specify all of their preferences. We wanted to provide a fair negotiation process that did not squeeze the suppliers to a point of unprofitability, and that would help Mars' long-term relationships with them. The iterative mechanisms had to terminate at a win-win outcome for the buyer and the suppliers. The iterative auction design used by Mars is illustrated in Figure 1.

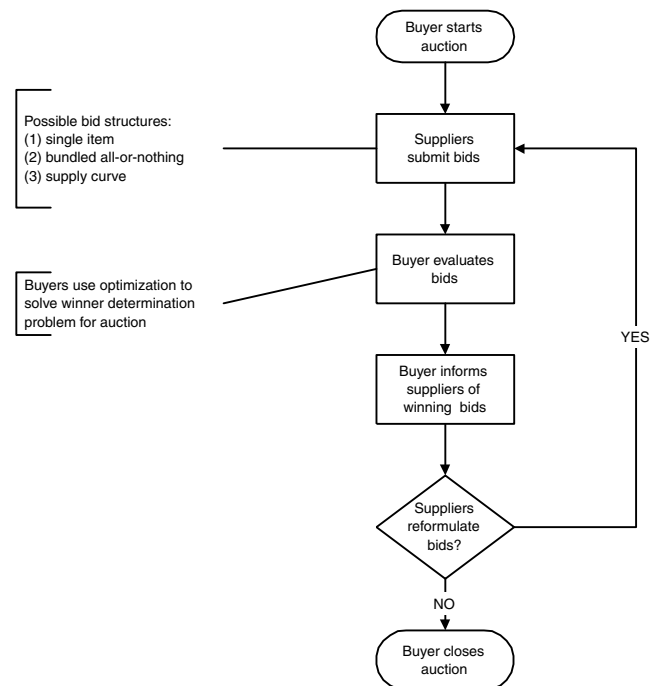


Figure 1: This is the process flow for an iterative, multiround procurement auction. Both the auctions Mars runs have termination criteria to stop when they receive no new bids for a specified time.

Bid Submission

The company's existing procurement practice was to post a request for quote for multiple items, giving the demand for each item (Table 1). Mars also specifies different types of responses (or bids). Most negotiations are based on a simple-bid structure where the supplier provides a unit price for each item along with the minimum and maximum quantity it is able to provide. The supply-curve bids give the unit price as a function of quantity. With bundled bids, in contrast, suppliers provide only a single price for a bundled offer.

Bid Evaluation

Ideally, every time Mars received a new bid, it would evaluate the bids to identify the provisional winners. However, the complexity of identifying the set of winning bids (often referred to as winner determination) depends on the bid structure. For two classes of bid structures (all-or-nothing bids and supply-curve bids), the problem of determining the winners is NP-hard. In addition, the introduction of business rules further complicates these problems, and the integer-programming formulations for solving the winner-determination problems must incorporate the business rules as side constraints. (We provide integer-programming formulations for determining the winner for each bid type in the Appendix.) As a result, it is difficult to solve the winner-determination problem

to identify the provisional winners as every new bid comes in. In designing the system, we compromised by allowing two minutes for each evaluation, accumulating the new bids within this interval. The winner-determination engine therefore had to evaluate bids in two minutes. We chose two minutes to allow enough time for computation under most situations but not unduly slow the progress of the auction.

We developed the bid-evaluation engine as an independent optimization module in C++ and then integrated it into the auction platform. The engine uses IBM's optimization subroutine library (OSL) as the LP/IP Solver. The implementation of this framework was done using IBM's e-commerce platform—WebSphere Commerce Suite 4.1, which provides the infrastructure support for the Number1Traders Web site.

Feedback on Clearing Prices

After every provisional allocation with an iterative auction format, Mars gives the nonwinning suppliers some feedback to help them reformulate their bids. For simple single-item auctions, the suppliers are usually provided with a *clearing price* at which supply for each item equals demand. However, for the multi-item auctions, no clearing prices exist for each item. We would develop clearing prices for bundles of items, but we could then have to report an exponential number of prices. However, in any iteration only a small number of bundles (at most as large as the number of items) have a provisional allocation, and only on these bundles would we need to report clearing prices. Over the span of the entire auction, the buyer must keep a list of all the bundles allocated in any round. In the worst case, this list can become large, but in practice we find that a small set of bundles provides cost complementarities and the bidding centers around this set.

Bid Reformulating

Suppliers use the feedback provided in each round to reformulate their bids. A rational supplier would reformulate its bids so as to maximize its outcome over all possible bid strategies of the other suppliers. In reality, developing such a complete contingent-bid strategy is impractical for several reasons: (1) Suppliers do not have complete information about the cost structures of the other bidders and as a result must compute

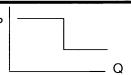
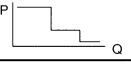
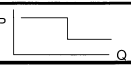
	Simple Bid	Bundled Bid	Supply Curve	
Item 1	Price Qty, D_1	\$/unit (Low, Hi)	Q_1	
Item 2	Price Qty, D_2	\$/unit (Low, Hi)	Q_2	
Item N	Price Qty, D_N	\$/unit (Low, Hi)	Q_N	
Price			P	

Table 1: In this example request for quote and bid types, the items are shown in rows and give the price and quantity demand associated with each item. The table shows different bid structures as columns.

their plans based on their expectations to incorporate the informational uncertainty, and (2) the computational burden of developing such a contingent plan is far greater than that of solving the winner-determination problem and hence would take longer. We find that the suppliers act in a straightforward fashion, trying to locally optimize their bids based on the feedback in the current round. An associated issue concerning supplier behavior is whether they implicitly collude. The auctions run so far do not indicate collusion. The bid price seems to move down quickly, and bidder behavior does not seem to have a (perceptible) pattern (such as sequential bidding). For some items purchased, cartels openly engage in price fixing. However, such markets are not suited for auctions.

In reformulating these bids, suppliers must follow the rules Mars imposes on subsequent rounds, such as minimum price decrement. For items or bundles for which Mars has made an allocation, nonwinning suppliers must offer a new price that is less than the current price by a fixed decrement. Such decrements are used to ensure rapid convergence.

Combinatorial Auctions

It can be advantageous to aggregate demand over several locations and plants to form larger transactions. Also, suppliers may make gains in efficiency that permit them to provide discounted bids on a bundle (for example, two sizes of identically printed packaging that the supplier can print without changing printing press inks). Such a discounted price would apply only if the buyer accepted the entire bid.

For example, the firm might send out an RFQ for packaging in three formats to four suppliers (Table 2). The suppliers could provide bundled all-or-nothing bids for two or three of these items, at a price shown in the "Seller bid price" row. By introducing simple decision variables (x_1 , x_2 , x_3 , and x_4) for the four bids, we can formulate an optimization problem that can be solved optimally using integer programming to minimize total procurement cost while ensuring that the demand for each item is satisfied. The optimal supply solution may oversatisfy demand. If there are no holding costs, this might be acceptable or even desirable.

Items	Suppliers			
	S1	S2	S3	S4
1,000 large Mars® brand display boxes	1	0	1	1
800 small Mars® brand display boxes	0	1	1	1
800 small M&Ms® brand display boxes	1	1	1	0
Seller bid price	\$150	\$125	\$300	\$125
Decision variable	x_1	x_2	x_3	x_4

Table 2: This example of a combinatorial auction shows the bids provided by four suppliers.

Bid Submission

With all-or-nothing bids, the supplier would have to provide exponentially many bids (with respect to the number of items in the auction) to completely describe its cost structure. For example, in the simple case of three items (Figure 2), each supplier could provide $2^3 = 8$ bids over all possible combinations. For 10 items, this could lead to over 1,000 bids per supplier. Even if suppliers could determine complete sets of combinatorial bids, they would probably be unwilling to provide this information.

Iterative schemes permit suppliers to make only those bids that maximize their profits at the current price levels for the items. They do not need to report their entire cost structures.

Winner Determination

The computational complexity of the winner-determination problem for combinatorial auctions is, even without side constraints, an NP-hard optimization problem. Each supplier is usually allowed more than one bid, and as the number of items increases, the number of bids can get quite large. The combinatorial auction can be formulated as a set-covering problem with side constraints (arising from the business rules). The formulation differs from the conventional formulation in that the side constraints make even determining the existence of a feasible solution NP-complete. We model such problems as integer-programming problems and solve them using OSL. Integer-programming techniques have proven to be effective in solving problems with 500 items and up to 5,000 bids.

Feedback for Bid Reformulation

For combinatorial auctions, the winning prices are provided by the buyer on bundles of items specified as winning bids in some round of the auction. In subsequent rounds, the bids suppliers make must be at least some fixed decrement δ lower than the current

The payback on Mars' investment was less than a year.

winning price on the bundle. The advantage of this design is that one can derive an extended integral formulation for the set-covering problem with side constraints that allows for dual prices on bundles. When suppliers bid to maximize their profits in the current round, the outcome is that each winning supplier is at an optimum at the given prices, leading to a price equilibrium.

Bid Submission for Volume-Discount Auctions

In making supply-curve bids, sellers specify the prices they charge for various quantities of an item. For instance, a supplier may charge \$100 per ton for up to 30 tons of sugar, but \$50 per ton for more than 30 tons. Bids take the form of supply curves, specifying prices per unit of an item within particular quantity intervals. Such price schedules can also be represented by an infinite set of all-or-nothing bundled bids across multiple items.

Winner Determination for Volume-Discount Auctions

In general, when multiple suppliers provide such volume-discount bids, the problem of determining the winning bids and the amount to buy from each supplier is a difficult optimization problem that can be modeled as an integer program and solved using commercial integer-programming software. In addition, the various business rules of the buyer can be captured as side constraints within the mathematical formulation. We based our solution approach on modeling the problem as a variation of the multiple-choice knapsack

problem (Martello and Toth 1984) (Appendix). By using customized knapsack covers, we improved the performance of the integer-programming solvers for these problems for up to 40 suppliers and 30 items. For larger problems we developed heuristics based on column-generation-based heuristics that provide approximate solutions to within one percent of optimal (Ladanyi et al. 2001).

Feedback for Bid Reformulation

In volume-discount auctions, suppliers state prices for individual items. In each round, Mars reports the highest price paid for any lot to the suppliers. This convention is similar to traditional negotiation practice, and suppliers find this feedback easy to use in reformulating their bids.

The Impact of Business Rules

In procurement settings, buyers must consider the business rules used in sourcing decisions. Mars' business rules reflect its practices, and buyers use them to constrain the allocations they can consider.

The business rules complicate the evaluation of bids because they must be incorporated into the integer-programming formulations as side constraints (Appendix). These side constraints affect the cost of procurement and often force Mars to accept higher costs than it would in a constraint-free solution (Figure 2).

Figure 2 presents the results of solving the winner-determination problem (using mixed-integer programming) for a randomly generated volume-discount auction with 10 suppliers and 20 lots. The size of the supply pool is varied, between five and 20 suppliers. This is enforced by setting both the minimum and maximum number of suppliers allowed in the auction to be the size of the supply pool required. Each point on the x-axis represents where we solved the winner-determination problem for a particular size of supply pool. (With less than five suppliers, the available bids could not satisfy the entire demand of the buyer.) The optimal (lowest-cost) solution for this auction occurs for a supply pool of 13 suppliers. With 13 suppliers, we also see the shortest CPU time required to solve the corresponding winner-determination problem. As we




Figure 2: Side constraints for the volume-discount auction, such as the exact number of winning suppliers, affect (a) the total procurement cost to Mars, and (b) the CPU time required to solve the winner-determination problem.

vary the size of the supply pool away from the size required to obtain the optimal procurement cost (which here is 13 suppliers), the total procurement cost increases significantly. The constraint on the minimum quantity awarded to each supplier has a significant effect on how much deterioration we observe as we increase the size of the required supply pool. In practice, the buyers have often specified that the size of the supply pool must be fixed or perturbed as little as possible from the status quo. This has the effect of limiting the possible savings that can be achieved using these types of auctions. Often the size of the status quo supply pool is larger than that required to find the optimal procurement cost, which also has a negative impact on the CPU time required to solve the winner-determination problem for these auctions (Figure 2).

Desirable Properties for the Auction Design

Two main properties are desirable for procurement auctions, fairness and optimality, and we have incorporated them into the auction design for Mars.

Auctions are a competitive mechanism for allocating the firm's business to suppliers. It is very important that the bidders perceive Mars' auctions to be fair. If they did not, bidders might refuse to participate in the auctions. We tried to ensure optimality and win-win

outcomes in designing the auctions to make sure they were fair.

Given a set of bids, the winner-determination engine determines the cost-minimizing bid set and thereby chooses the winning suppliers. This optimization problem is hard to solve computationally. Using heuristics to solve the problem might yield an approximate solution that was within less than one percent of the optimal-cost solution. However, such an approximate solution might call for an allocation drastically different from the optimal. A supplier who would have an allocation in the optimal solution might not get any allocation in the approximate solution. Suppliers who failed to get an allocation because of the approximate nature of the algorithm would see the auction as unfair. Therefore, the winner-determination engine must provide optimal allocations. The integer-programming solver we developed and tuned for efficiency provides optimal solutions within two minutes response time.

Time Stamps

In multiround auctions, the buyer must decide how to treat sets of bids made in different rounds for the same items at the same price. For example, a supplier might create a combinatorial auction to purchase some quantities of items *A*, *B*, and *C*. In the first round of the

auction, Supplier 1 bids \$100 for items *A*, *B*, and *C*, and Supplier 2 bids \$30 for item *A*. The solution to the winner-determination problem for this round is that Supplier 1 wins with its bid because it supplies the entire demand. During the second round, Supplier 3 enters the auction and bids \$70 for items *B* and *C*. This winner-determination problem has two potential solutions: either Supplier 1's bid to supply all three items or the combination of bids by Suppliers 2 and 3, which together would supply all three items. In either case, the buyer would pay \$100 for its total demand. From a fairness point of view, these two solutions are not equally desirable. In a multiround auction, new bids should be deemed winning bids only if they result in the buyer paying a lower price for the items in the auction. In simple auctions, the usual rule governing such situations is that the earlier of the identically priced bids for the same items is preferred. This rule is straightforward to enforce in a simple, single-item or multi-item forward or reverse auction. In combinatorial and volume-discount auctions, this rule becomes harder to enforce, because the number of possible solutions to the winner-determination problem may be exponential with respect to the number of bids placed in the auction.

In this project, we gave each bid a numeric time stamp, representing the time it was accepted in the auction. We then had two objectives for the winner-determination problem for both types of auctions:

- (1) To determine all the sets of bids that would minimize the buyers' total procurement cost, and
- (2) From these sets of bids, to select the set that would minimize the sum of the time stamps for the bids in this set.

To solve the winner-determination problem with this new objective, we formulate two integer-programming problems. The first problem we solve is the winner-determination problem (for either the combinatorial or volume-discount auction) without considering time stamps.

The second integer-programming problem differs from the first in two ways:

- (1) We add a constraint stating that the demand should be bought at the minimum cost determined in the solution to the first problem.
- (2) The objective for the second problem does not

consider price at all, because the new constraint will ensure the minimum price. Instead we seek to minimize the sum of the time stamps of the candidate winning bids in solving the winner-determination problem.

In practice, we find that it takes far longer (four to 20 times) to solve the second integer-programming problem than to solve the first.

Deployment Issues

It takes a long time to change the way people think! Coping with large geographical distances, varying business environments, and different cultures made the implementation part of this project at least as hard as the technical part. Mars buyers are spread over five continents and more than one hundred countries. To launch this project we had to inform all the buyers about the Number1Traders.com Web site, what it does and how to use it, and convince them of the benefits so they would change the way they ran their businesses.

Training Across the Net and Around the World

To inform buyers quickly about the basic ideas underlying the new system, we developed a half-day course on electronic auctions. In 2001, we trained over 300 buyers worldwide through classes offered in North America, Europe, Australia, and Asia. We trained another 150 via video conferences and Net meetings. Four buyers interested in auctions volunteered to become buyer-experts and took on additional responsibilities.

The training classes were only the beginning of the education and support we provided. Buyers required personal assistance the first time they ran an auction. Gail Hohner and the buyer-experts helped first-time users to choose buying opportunities appropriate for auctions, to design the first auction, and to train suppliers in the supply pool. They made themselves available for technical support during the auctions and documented the results so that other Mars buyers could learn from the experience.

Changing the Business Process

The Number1Traders Web site has been available to Mars buyers since September 2001. The growth in usage has been steady. Buyers see the efficiency of the process as a prime reason for adopting the system. As the number of buyers who have had positive experiences using the system increases, word-of-mouth has become an effective internal advertisement.

Still, we should not underestimate the changes buyers must make to adopt a new purchasing practice. With bilateral negotiations, buyers maintain all the social aspects of the supply relationship throughout the process. They have many occasions for both formal and informal communication with suppliers. Although the buyer and supplier meet at other times during the year, the majority of the face-to-face meeting time is concentrated around the contract negotiation. When they change to an auction process, buyers must continue to work on integrating suppliers into their businesses and on maintaining informal social contact (because it strengthens the buyer-supplier relationship and provides a vector for information); however, they must avoid negotiating price or quantity allocations. Buyers have a very difficult time making this separation. They must not negotiate prices and quantities, the task on which they formerly spent most of their time.

Suppliers have generally reacted favorably to these auctions. They have commented on some of their advantages. Many suppliers value the transparency of the auction process. In negotiations, they had no information on competing suppliers' prices. In auctions, the prices they offer are compared to those of other suppliers. Also suppliers now control their last bids rather than having the buyer end the process at an arbitrary point. Suppliers value their time. When the system becomes very efficient (as it does when auctions have been run many times), suppliers have been very happy with the process. Suppliers view these auctions as more equitable than any other auction mechanism they have encountered because many suppliers win business, not just one. These auctions also allow suppliers to present their unique sales propositions (for example, discounts reflecting unique synergies). Schut Cees of Thermphos commented, "We have participated in auctions that have just been price gouging

exercises. This is the most equitable design we have ever seen."

Procurement auctions have received a lot of press recently, and spectacular savings have been reported for reverse English auctions. But where do those savings come from? In many cases, they come from squeezing suppliers to the point of (or beyond) zero profit. Consequently, a strong backlash is developing in some industries against using auctions at all. We find that our approach to auctions draws our supplier pool closer rather than alienating it. No system is perfect, however, and not everybody is happy, but generally the response has been very positive.

Impact

The auctions have yielded consistent cost savings. We don't mean that suppliers are selling to us below cost. The efficiencies come from matching supplier capabilities and the company's needs and thus increasing suppliers' margins, part of which are to provide Mars with savings. We have found that when buyers were willing to change the size of the supplier pool or shift large amounts of business, the auctions yielded greater savings. The payback on Mars' investment was less than a year.

To determine savings, we compared the buyers' predictions made one or two days before the auction as to the best outcome possible from a negotiation. Such a prediction incorporates the market knowledge to date and is the best way we know of making a comparison. Buyers do not make such predictions solely as benchmarks for auctions but for use in Mars' production-planning system. Mars assesses its buyers on the quality of their predictions and their understanding of the markets they cover. Any systematic distortion would come from buyers' optimism about what they could negotiate. In such cases, we were conservative in estimating savings. Once they are integrated into the business process, auctions take much less time than negotiations. Mars buyers, as a result, have more time to align businesses and to seek synergistic value from suppliers.

It usually takes a day or two to design, set up, and train suppliers for auctions the first time they are run. As we repeat an auction, training times for suppliers

drop from one hour for first-time users to less than 10 minutes for repeat users. Auctions have never taken more time than negotiations. In the most significant reported time savings, a 40-minute auction replaced a price-only negotiation process that had lasted over two weeks and required the buyer to make nine separate air trips to finalize only the prices and volumes of the contract.

Adoption has been encouraging. Many more buyers have expressed interest in the process, but they must wait until their next contract cycle (which can be a year away). By the end of 2002, we expect to have conducted nearly 60 auctions.

The Future of Combinatorial and Supply Curve Auctions at Mars

Mars has started new business units (for example, www.Freight-Traders.com) that utilize auctions, and the Number1Traders.com site has been recognized in the 2001 "Mars, Incorporated, Make the Difference Awards" for innovation that changes the business process. Clearly, combinatorial and supply-curve auctions provide value. Mars is investing money, time, and personnel to make these tools available to its buyers and to ensure that they understand them. The only question now is under what circumstances and how broadly can buyers use auctions? Learning to use these new tools will take time and experimentation.

Executive summaries of Edelman award papers are presented here. The complete article was published in the INFORMS journal *Interfaces* [2003, 33:1, 23-35]. Full text is available by subscription at <http://www.extenza-eps.com/extenza/contentviewing/viewJournal.do?journalId=5>