Marginal Bidding: An Application of the Equimarginal Principle to Bidding in TAC SCM

Tyler Odean
todean@cs.brown.edu

Department of Computer Science, Brown University, Box 1910, Providence, RI 02912

Abstract. We present a fast and effective bidding strategy for the Trading Agent Competition in Supply Chain Management (TAC SCM). In TAC SCM, simulated computer manufacturers compete to procure component parts from suppliers and then sell the assembled products to customers in reverse auctions. One of the many sub-problems implicit in the competition is deciding how many computers to sell of each type and at what prices to sell them. We propose a greedy solution, Marginal Bidding, which performs comptetively with a computationally intensive integer linear programming solution on small problem instances. Moreover, because of the computational efficiency Marginal Bidding can handle larger problem instances and therefore can more effectively reason about predicted future demand.

1.1 Introduction

A supply chain consists of a network of autonomous agents engaged in procurement of raw materials for, manufacturing and distribution of finished products. The Trading Agent Competition in Supply Chain Management (TAC SCM) is a simulated computer manufacturing market in which six independent software agents compete to maximize the profits from their supply chain. In this paper we study the TAC SCM bidding problem, in which agents compete in reverse auctions to sell computers, balancing the desire to maximize revenue per order by keeping the bid prices high and the need to maximize orders won by keeping the bid prices low. Ideally these decisions must be made in the context of future conditions: in a bull market it may be advantageous to conserve production capacity for future demand, in a bear market it may be more desirable to bid aggressively and claim a larger share of current demand to be fulfilled with future production.

To solve the bidding problem, we first model the game-theoretic problem as a simpler decision-theoretic problem of stochasticity. We then further simplify the stochastic decision problem to a tractable approximation called *expected bidding* as presented in Benish *et al.* [1]. We reduce expected bidding to a generalization of the classic knapsack problem called *nonlinear knapsack problem* (NLK). Then we apply the Equimarginal Principle—which states that revenue is maximized among possible uses of a resource when the return on the last unit of the resource is the same across all areas of use—to the problem of expected bidding and propose a greedy solution which we call Marginal Bidding.

We advocate for Marginal Bidding in this paper because it scales linearly with the number of days in the problem and hence can more easily solve an N-day extension of expected bidding than a traditional ILP solution.

To analyze the performance of our heuristics designed for the TAC SCM, we built a simulator that generates decision theoretic simplifications of the game-theoretic problems TAC SCM agents face, including bidding. Using our simulator we compared several variants of Marginal Bidding to an ILP solution.

We show that certain variations of Marginal Bidding can compute bids faster than our ILP solution; hence, incorporating a Marginal Bidder into a TAC SCM agent would allow for more time to be spent on other decision problems (e.g., procurement). Moreover, this speedup enables Marginal Bidders to reason about future demand as well as current demand, and hence achieve greater revenues when knowledge of the future is valuable. While the gains to be realized by reasoning about future demand in TAC SCM appear modest, we demonstrate that more substantial gains can be realized under more volatile or seasonal conditions that generate more extreme market swings.

This paper is organized as follows. We begin by describing the Equimarginal Principle of marginal utility theory, originally posited by Gossen in the mid 1800's. We note that this principle can be applied to solve the nonlinear knapsack problem. Then, we present a discretization technique coupled with a greedy algorithm to approximately solve an NLK. Next, we formalize TAC SCM bidding as an N-day recursive stochastic program, and argue that expected bidding, a 1-day deterministic approximation, can be reduced to solving an instance of the NLK. Then, we present Marginal Bidding, a heuristic for solving an N-day extension of expected bidding that incorporates the aforementioned discretization technique and greedy approach to solving the NLK. Finally, we compare experimentally the performance of two heuristics, Marginal Bidding and an ILP, in simulations of the TAC SCM bidding problem.

1.2 The Equimarginal Principle

Prussian Economist H. H. Gossen posited the two fundamental laws of utility, the Equimarginal Principle and the Law of Diminishing returns. The Law of Diminishing Marginal Returns simply states:

> The amount of any pleasure is steadily decreasing as we continue until the last saturation is reached.

Gossen's corresponding law of utility maximization, The Equimarginal Principle states:

> If a man is free to choose among several pleasures but has not time to afford them all to their full extent, then in order to maximize the sum of his pleasures he must engage in them all to at least some extent before

enjoying the largest one fully, so that the amount of each pleasure is the same at the moment when it is stopped; and this however different the absolute magnitude of the various pleasures may be.

The Equimarginal Principle applies to problems in which a limited resource such as time, capital or labor must be distributed among two or more independent uses. Taken alongside the Law of Diminishing Marginal Returns, it is easy to see that an optimal solution must imply that the marginal returns for each of the possible allocations are equal. Indeed, if they were not, a better allocation could be achieved by redistributing a unit of resource from a use with smaller marginal returns to one with larger. Gossen's claim, that equivalent marginal returns imply an optimal solution is less obvious - for a proof see Mas-Colell *et al.* [6] (Theorem M.K.3 on page 961).

1.2.1 The Nonlinear Knapsack Problem

The problems domains over which the Equimarginal Principle operates are fundamentally similar to knapsack problems. In traditional knapsack problems, we are presented with a finite, positive capacity $C > 0$, and a set of n items with a designated value v_i and weight w_i. The objective is to maximize the sum of the values of selected items subject to the constraint that the weight of selected items may not exceed the knapsack capacity. More formally:

$$\max_x \sum_{i=1}^{n} v_i x_i \tag{1.1}$$

$$\text{s.t.} \quad \sum_{i=1}^{n} w_i x_i \leq C \tag{1.2}$$

In the economics problems discussed, the decision is not which items to take, but what quantity $q_i \geq 0$ to take of each item (or alternatively, how much of the limited resource to devote to each use), where in general the value of a item or use depends on the quantity of capacity devoted to it. The result is a knapsack problem with a potentially non-linear objective function, i.e. a *nonlinear knapsack problem*. Again, more formally:

$$\max_x \sum_{i=1}^{n} f_i(x_i) x_i \tag{1.3}$$

$$\text{s.t.} \quad \sum_{i=1}^{n} g_i(x_i) \leq C \tag{1.4}$$

Where f_is are value functions, g_is are cost functions, and the knapsack capacity C is reinterpreted as a budget, B.

In a typical NLK instance, the x_is are unbounded and positive, the f_is are real-valued, concave and nondecreasing and the g_is are real-valued, convex and

Inputs:
discretization factor K
value functions f_i
cost functions g_i
Outputs:
a vector q of quantities consumed, one per use

1. for each use i
   a) initialize $q_i = 0$
   b) insert i with priority $\nu_i(s_i) = \frac{f_i(s_i)}{g_i(s_i)}$ into a priority queue Q
2. for $t = 1$ to K
   a) pop off of Q a use j with the highest priority
   b) increment q_j by s_j
   c) insert j into Q with priority $\nu_j(q_j + s_j) = \frac{f_j(q_j+s_j)-f_j(q_j)}{g_j(q_j+s_j)-g_j(q_j)}$
3. return q

Fig. 1.1. A FPTAS for NLK. The algorithm runs in time $O(\frac{1}{\epsilon} n \log n)$.

nondecreasing. A concave value function implies non-increasing marginal values and a convex cost function implies non-decreasing marginal costs. When we divide non-increasing marginal values by non-decreasing marginal costs the result is diminishing marginal returns. Hence, by the Equimarginal Principle total value of an NLK is maximized when the marginal returns are equal across all possible areas of use. i.e.

$$\frac{f_1(x_1)}{g_1(x_1)} = \ldots = \frac{f_i(x_i)}{g_i(x_i)} = \ldots = \frac{f_n(x_n)}{g_n(x_n)}$$

The subset of NLKs for which f is quadratic and g is linear can be solved exactly in polynomial time (see, for example, Tarasov et al [11]). Our approach however can be applied more generally to any arbitrary nondecreasing concave value and convex cost functions. In this section we present a Fully Polynomial Time Approximate Scheme (FPTAS) for solving NLKs. To discretize, we assume that the resource can be spent in K equal units of size $k = \frac{B}{K}$. We then formulate a special 0/1 knapsack problem where the items are of values v_{ij}, the marginal return of the jth unit of the ith use, and constant weights equal to k. Because the weights are constant, these instances of knapsack can be solved by greedily consuming units of various uses in sorted order by value from highest to lowest until the budget is exhausted. By virtue of our assumptions about NLKs we know that the marginal returns will diminish and therefore we know $v_ij \leq v_{i,j-1}$ and that our greedy solution will never consume the jth unit before consuming the $j-1$th, so it is easy to reconstruct a solution to the original NLK.

This intuition leads to our FTPAS algorithm for NLKs:

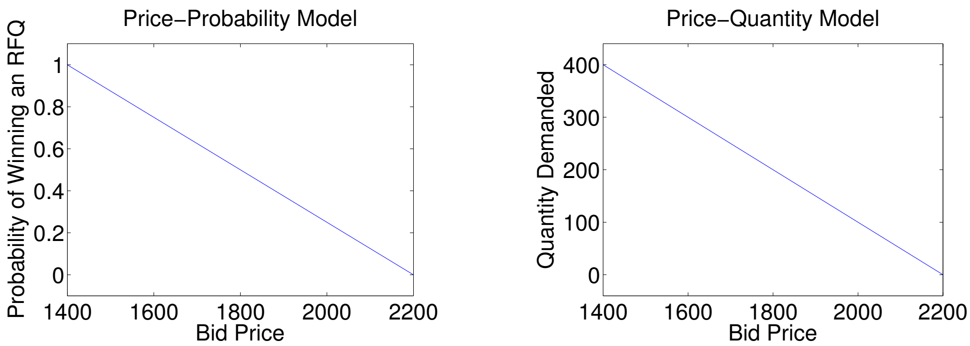

Fig. 1.2. (a) Sample price-probability model. (b) Sample price-quantity model.

Our approximation yields a solution with $\epsilon = \frac{2n}{K}$ and runtime $O(\frac{1}{\epsilon} n \log n)$[1]. For a proof of these claims, see Greenwald, et al. [4] In the next sections we will present Bidding in the TAC SCM and define a tractable approximation of it called expected bidding. We will reduce expected bidding to an NLK satisfying the usual f_i and g_i requirements, and we will apply our greedy algorithm and compare it to our ILP solution.

1.3 Bidding in TAC SCM

In TAC SCM six software agents compete in a simulated personal computer market. Each agent can manufacture any of sixteen different *stock keeping units* (SKUs). The agents then compete to sell the SKUs in reverse auctions to a common pool of customers. More specifically, the agents receive identical RFQs from a customer specifying a SKU type, a quantity, a due date, a penalty rate for late orders and a reserve price beyond which the customer will be unwilling to purchase. Agents submit offers representing the price at which they would be willing to satisfy a given RFQ. Customers then award the contract to the agent who has offered the lowest price.

1.3.1 Price Probability Models

In a marketplace with indistinguishable products, a seller hoping to adjust its market share can do so only by changing its price. Such a seller is likely to gather relevant historical data for use in predicting the market shares that correspond to various price settings. Following Benisch *et al.* [1], we assume that this prediction task has already been completed, and the agent is already endowed with a

[1] Although mentioned here as justification for the algorithms presented, the theoretical proofs surrounding our algorithm for NLKs are largely the work of Victor Naroditsky and are not part of our contribution to the project.

price-probability model that reports the probability of winning an order for each possible bid on current and future RFQs.

Rather than specifying a price-probability model for each individual RFQ, we partition the set of RFQs according to their defining characteristics so that we can obtain a richer set of price-probability models (we are assuming that models built using more data can make more accurate predictions). In TAC SCM, a natural partitioning of the set of RFQs is by SKU type and due date. We refer to each element of such a partition as a *market segment*.

Figure 1.2(a) depicts the price-probability model defined by this equation:

$$p(x) = \frac{2200 - x}{800} \qquad 1400 \leq x \leq 2200 \tag{1.5}$$

This model asserts that a bid of 2200 has no chance of winning (it is the reserve price above which there is no demand), whereas a bid of 1400 is guaranteed to win (it is the price below which there is no supply). In between, at a price of 1800, say, a bid wins with probability 0.50. Price-probability models need not be linear, but can incorporate whatever techniques necessary to model the likelihood of a bid price being the lowest offered in a market segment.

1.3.2 Expected Bidding

The N-Day stochastic bidding problem is formulated as a recursive stochastic program in Appendix A of [llncs]. A tractable approximation of the one day stochastic bidding problem is *expectedbidding* where an offer on a bid for quantity q with probability of p deterministically results in a partial order of quantity of pq. Collapsing the stochastic content of a price-probability model into a deterministic partial order model is done by scaling the price probability model by the quantity demanded in a given market segment. The ensuing models are called *price-quantity* models and provide a mapping from price to expected market share within the segment.

The objective in expected bidding is to maximize the expected revenue over the space of possible bids subject to the capacity constraint and given *price quantity* models $h_i(x_i)$ for each segment i. More formally:

$$\max_{x} \sum_{i=1}^{n} f_i(x_i) x_i \tag{1.6}$$

$$\text{s.t.} \quad \sum_{i=1}^{n} c_i f_i(x_i) \leq C \tag{1.7}$$

where x_i is the bid price for RFQs in segment i. Note that expected bidding is an NLK with $f_i = h_i(x_i)$ and $g_i = c_i h_i(x_i)$. Assuming that $h_i(x_i)$ is invertible, i.e. that each bid price uniquely maps to an expected market share, we can equivalently calculate a desired quantity from each market segment and bid the corresponding bid price.

1.4 A Greedy Algorithm

1.4.1 Marginal Bidding

An important simplification in the approximation of expected bidding is that true TAC SCM bidding spans a range of days. In this section, we describe how to extend the greedy algorithm into the multi-day setting. The result we call *Marginal Bidding*. Since there is no upper bound on the number of days possible in an arbitrary multi-day bidding problem, the extension of the greedy solution to multiple days requires an additional parameter, the *window size* W, to define the number of days within the multi-day problem the greedy algorithm will consider.

The algorithm is presented in detail in figure [?]. At the conceptual level, it fulfills orders in order of non-increasing revenues-per-cycle, next it schedules production of different SKU types greedily in order of unit marginal returns, finally it bids the prices associated with the quantity desired from that given market segment. For simplicity, the pseudocode presented does not address component restrictions. Extending the algorithm to incorporate those constraints is not complex however - simply add as inputs to the bidder a current component inventory and expected component arrivals. Then as each potentially scheduled increment is considered, reject it if the components are available by the scheduled production date. Then after each component is scheduled and decrement the anticipated component supplies appropriately, backwards from the production date.

Scheduling The marginal bidder is, in general, agnostic to it's production scheduling strategy. That is, we could schedule using any arbitrary heuristic for planning production that we desired. Two obvious examples are *as soon as possible*, which ensures that the most profitable products are scheduled for production with the highest priority, and *as late as possible*, ensuring as much flexibility as possible to respond to changes in market conditions. However, scheduling early can cause an agent to select a suboptimal production schedule by occupying early production capacity that could be used to address less profitable but more immediate contracts, and scheduling late can cause an agent to begin with empty or near empty production schedules in some conditions - a risky bet that the market predictions will continue to hold. Our Marginal Bidder uses a hybrid approach, scheduling order production as soon as possible, to ensure against late delivery and defaulting penalties, and scheduling expected market share production as late as possible to hedge against shifts in the market climate.

1.4.2 A caveat

Notice that we implicitly vary K across market segments. The reason for this is that because we partition our market into segments that are not necessarily of equal size. When unit sizes are large with respect to given market segment, it can interfere with the algorithm's ability to target that segment with an appropriate level of granularity. More conceptually, we specify the unit size by market percentage, rather than by number of requisite cycles. Although this allows our algorithm to appropriately fine-tune their bidding even to smaller market segments, it also renders the theoretical guarantees presented earlier in the paper inapplicable. Therefore it behooves us to empirically compare performance of our greedy approximation to a performance benchmark, in this case an Integer Linear Programming solution to the expected bidding problem.

1.5 Experiments

In this section, we report on experiments designed to compare the performance of four bidding algorithms with varying abilities to reason about the future, an ILP bidding heuristic (see Benisch *et al.* [1]) and three variations on the Marginal Bidding heuristic developed in this paper. We expect the Marginal Bidders to compute bids faster than the ILP, and we expect this speed to enable them to consider larger windows into the future, which should lead to higher revenues than the ILP under some market conditions (and never lead to lower revenues). We test these conjectures on instances of TAC SCM bidding in a simulator we built that tests individual agents in isolation by generating decision-theoretic simplifications of the game-theoretic problems TAC SCM agents face.

1.5.1 Test Suite

We tested an integer linear programming solution with a 1 day window (ILP), meaning it did not reason about any future demand beyond the current RFQs and outstanding orders arriving each day. We compared this ILP with three variations of the marginal bidder: a marginal bidder with a 17-day[2] window (MB-17), a marginal bidder with a full-game window (MB-Full), and a marginal bidder with a hybridization of the two that considers the full game window, but does so at a coarser granularity as it reasons further into the future in order to keep its run time in check (MB-Coarse).

The 17 Day (MB-17) and full-game (MB-Full) bidders partition demand (i.e., the set of current and future RFQs) into market segments by SKU type and due date, and the size of a unit in each market segment is 1 product. The hybrid full-game bidder (MB-Coarse) also divides demand up by SKU type and due date. For the first 17 days, it considers each due date separately, but beyond the initial 17 days it divides demand into increasingly larger chunks, whose due-date ranges grow by powers of 1.8.[3] For the coarse bidder, each market segment's unit size is 1 product multiplied by the number of days in that segment.

Since each TAC SCM day is 15 seconds, and a bidding policy is one of many decisions an agent must make each day, it may not be wise for an agent to

[2] We chose 17 as the default window size because it is the last day on which a current RFQ with the latest possible due date can be filled in TAC SCM.

[3] For example, SKUs due on days 18-19 are grouped together ($1.8 \approx 2$), as are SKUs due on days 20-22 ($1.8^2 \approx 3$), and days 23-28 ($1.8^3 \approx 6$), and so on.

allot too much of its daily run time to bidding alone. We thus study a likely TAC SCM situation in which the bidder is only given 5 seconds to formulate its daily bidding policy. The full-game Marginal Bidder often requires more than 5 seconds per day to compute its policy, so it is not a feasible TAC SCM bidder, but we include it in this discussion for benchmarking purposes.

In order to reach a reasonable solution within the allotted 5 seconds, the ILP dynamically calculates an appropriate degree of discretization using a formula that was empirically determined to minimize the ILP's distance from optimality within a 5 second window. The equation for the number of price points is 2300/(# of RFQs + # of Orders). An ILP with a run time of up to 15 seconds and additional price points was also tested, but did not yield significant gains.[4]

1.5.2 Experimental Design

Recall that in TAC SCM each agent submits its bids to a reverse auction, so that an RFQ is awarded to the agent that bids the lowest price below the reserve price. Using our simulator, we tested our bidding algorithms in isolation, not against other bidding agents, as would be the case in a true reverse-auction setting. The simulator simply awarded contracts by transforming each offer into an order with a certain probability, namely that which is associated with the bid price under the price-probability model for the relevant market segment. Hence, we simulated the stochastic bidding problem, although our heuristic solutions are approximate solutions to the expected bidding problem.

In our experiments, agents were endowed with perfect price prediction: i.e., the various price-probability models (one per market segment per simulation day) were shared between the agent and the simulator. Regarding demand, the number of customer RFQs of each SKU type scheduled to arrive each day was broadcast before the simulations began. Then, on each simulation day, the agents received a set of current RFQs whose quantities and due dates were sampled from the distributions outlined in the TAC SCM game specifications, and they assumed that the quantity and due date associated with each of the future RFQs were the means of the same distributions.[5] Reserve prices were also known to the agents; they were built in to the price-probability models.

We tested our bidders by running 25 simulations of 100 day games under three families of market conditions: (i) constant: i.e., conditions on one day are reflective of the conditions on the next; (ii) gradually changing conditions; and (iii) sudden shifts, including demand or price shocks. Under the non-constant conditions we examine situations of rising demand and price. Falling demand and price conditions are not presented, but produce similar results.

[4] An ILP with a 2-day window was also tested, as was one with a 17-day window and constrained capacity (2000 cycles on day 1 and 2000 cycles on days 2 through 17). Again, these variants did not yield significant gains.

[5] The reason for drawing a distinction between the quality of the predictions of the number of RFQs of each SKU type and their attributes is: the former is somewhat predictable in TAC SCM—it is dependent on history (see, for example, Kiekintveld *et al.* [5])—while the latter is not.

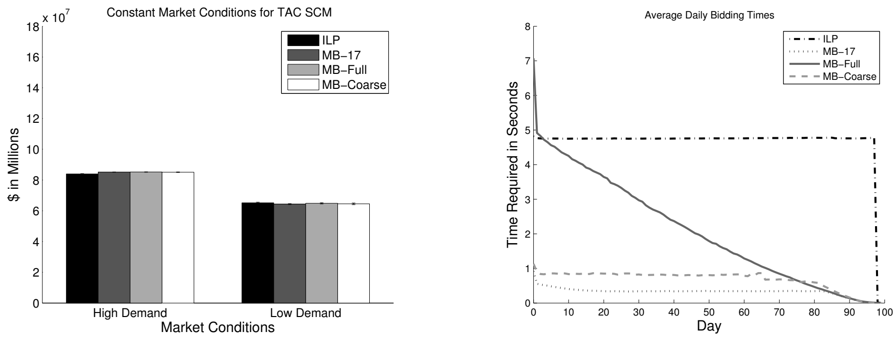

Fig. 1.3. (a) Revenue from deliveries under constant market conditions. (b) Average daily bidder times in high demand conditions. Low demand bidder times were similar.

For simplicity, in these simulations we assume infinitely many components. Introduction of component constraints does not appear to significantly alter the relative performance of our bidders.

1.5.3 Constant Conditions

In our first set of market conditions, we compare the bidders under constant demand and price. Presented here are steady conditions of high demand, defined as 20 RFQs per SKU type per day, which is the maximum possible according to the TAC SCM game specification, and low demand, defined as 5 RFQs per SKU type per day, the lowest possible. Prices in this experimental setup range linearly from 50% to 125% of the SKU base price.

Under such conditions, we should expect to see no particular advantage to planning for the future, since an optimal solution to the entire game can be contructed by concatenating a sequence of optimal solutions, one per day, computed for each day in isolation. Indeed, in terms of revenue, all the bidders are competitive with one another under these conditions (see Figure 1.3(a)). Note however that MB-17 and MB-Coarse arrive at their solutions an order of magnitude faster than the ILP or the MB-Full bidding algorithms (see Figure 1.3(b)).

1.5.4 Shifting Conditions

More interestingly, market conditions can change over the course of a TAC SCM game, either steadily as in a market adjustment or suddenly as in a demand or price shock. In our next experimental setup, demand is initialized to 5 RFQs per SKU per day, and prices range linearly from 50% to 75% of the SKU base price. We then considered shifts to 20 RFQs per SKU per day and prices ranging from 100% to 125% of the base prices by day 50. These shifts are representative of the magnitude of changes an agent might observe while playing a typical TAC SCM game. These changing market conditions were tested both as steady

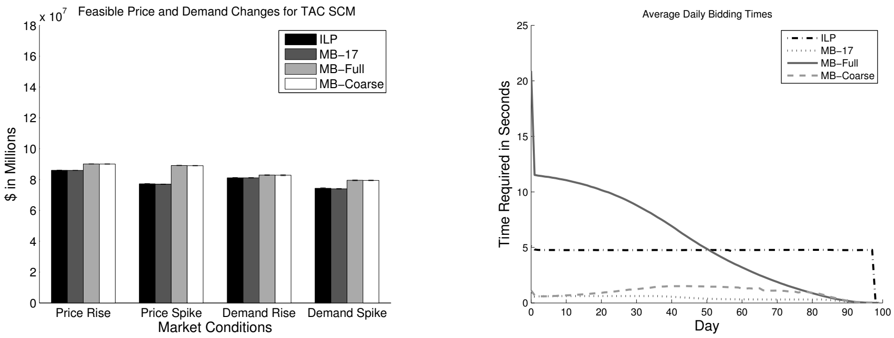

Fig. 1.4. (a) Revenue from deliveries under feasible SCM market conditions. (b) Average daily bidder times in Price Rise conditions. Other market conditions had similar run times.

linear accumulations from day 1 to day 100 and as abrupt surges on day 50. In our price-shifting simulations demand is held constant; in our demand-shifting simulations price is held constant.

As expected, those bidders with more extensive knowledge of the future (MB-Full, MB-Coarse) are able to exploit the mid-game surges by dedicating production from today to future demand when conditions are more favorable. Bidders with a shorter window (ILP, MB-17) are unable to plan far enough ahead to take advantage of the upcoming shifts, and hence accumulate less revenue over the course of the game. In addition to the additional revenue gained by exploiting its knowledge of the future, the MB-Coarse bidder continues to run in substantially less time than the ILP. See Figure 1.4.

The advantages of a larger window are more pronounced under those market conditions in which the shift in demand or price comes as a sudden spike rather than as a steady rise. When demand or prices rise gradually, even an agent with a small window is aware that tomorrow's market conditions are slightly more profitable than today's, and can reserve some inventory for future sales. However, when demand or price spikes suddenly, an agent is not aware of more desirable future market conditions until the spike falls within its window.

Because one of our simplifying assumptions for these simulations is that agents have perfect models of future demand and price, it is encouraging that MB-Coarse performs just as well as MB-Full. Their similar performance suggests that the benefits of looking into the future may still be realized by agents with more realistic but less accurate models.

1.5.5 Extreme Conditions

Within the context of TAC SCM, the previous experimental setup characterizes shifts from one extreme set of realistic conditions to another, and the gains resulting from knowledge of the future are modest. However, it is easy to envision

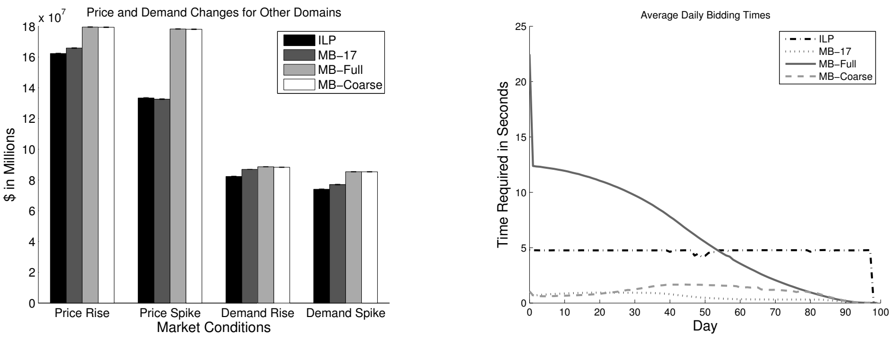

Fig. 1.5. (a) Revenue from deliveries in extreme market conditions. (b) Average daily bidder times in Price Rise conditions. Other market conditions had similar run times.

markets that are more naturally volatile or are subject to large seasonal trends in demand. The greater the extent to which market conditions vary across time, the greater the opportunity for bidders able to consider a larger window into the future to earn greater profits. In order to demonstrate this effect, we present a second set of simulations assuming shifting market conditions, but the shifts are of greater magnitudes. In particular, demand surges from 5 to 40 RFQs per SKU per day, and price rises from $[50\%, 75\%]$ to $[200\%, 250\%]$ of the base prices, again as both an interpolated steady rise and as an overnight jump.

With no significant changes in run time (compare Figures 1.4(b) and 1.5(b)), the marginal bidders are able to exploit the extreme changes in market conditions, and in particular the bidders with larger windows (MB-Full and MB-Coarse) are able to earn even greater profits (see Figure 1.5(a)). Also of interest is the relative impact of demand changes versus price changes. We observe a more pronounced impact when considering knowledge of the future under price-changing conditions for two reasons.

First, because of capacity constraints, an agent can only produce a limited quantity of each product on each day. Hence, an increase in demand does not necessarily translate into an increase in the number of finished products. So even if a demand shift results in higher prices, revenues need not increase substantially, particulary in comparison to the revenue increase associated with a price increase (see Figure 1.6(a)). If the magnitude of the price shifts in our experiments were reduced, or if production capacity were increased, stockpiling products until a demand shift could be as worthwhile as stockpiling products until a price shift.

The second factor that mitigates the advantage of knowledge of the future in conditions of shifting demand is the relatively flat slopes of our quantity-price curves. With flatter slopes, the difference in revenue between prices on the initial curve and prices on the curve after a demand shift is small (Figure 1.6(b)). Thus it matters less if the agent stockpiles products for the future, and in turn it matters less if the agent has any knowledge of the future. If the quantity-price curves had steeper slopes, knowledge of the future in conditions of shifting demand would likely prove more valuable than our current experiments suggest.

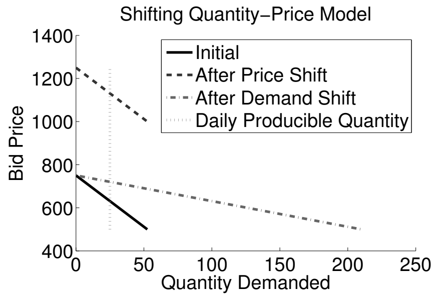

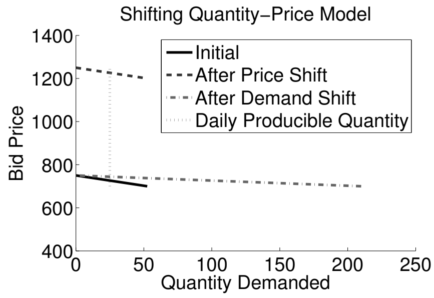

Fig. 1.6. (a) Sample quantity-price models before any shift, after a price shift, and after a demand shift. To illustrate the constraining effects of production capacity, also shown is a sample daily producible quantity. In our experiments, price shifting conditions result in higher revenues than demand shifting conditions, and thus knowledge of a future price shift is more valuable than knowledge of a future demand shift. (b) For quantity-price models with flat slopes, predicting future demand is not very important.

1.6 Related Work

Researchesr at the Cork Constraint Computation Center implemented an ILP approach to bidding within a constraint based agent, which they called Foreseer [3]. Similar to the expected bidder we compare our marginal bidder to, Foreseer maximizes profit over the space of possible bid prices subject to capacity, component and reserve price constraints.

Researchers at CMU reduced probabilistic pricing (a problem somewhat similar to TAC SCM bidding) to an NLK under the assumption of diminishing marginal returns, and present an ϵ-optimal solution to their problem over arbitrary value functions [2]. The efficiency of their method, however, is dependent on normally distributed customer valuations (an analog of price probability models). Our efficiency is independent of the form the price probability models take.

Finally and most similarly, the TacTex team have developed a greedy bidder with many similarities to the marginal bidder presented here [7]. A few distinctions do exist, however: TacTex initializes bids to the reserve prices and then iteratively reduces bids according to some selection mechanism until production capacity is reached or profit is no longer increasing. The selection mechanism relies on a heuristic to determine whether production capacity or component availability is the limiting resource, and selects by profit-per-cycle or $\frac{\Delta Profit}{\Delta Probability}$ respectively.

1.7 Discussion and Conclusion

In this paper we describe a technique for solving NLKs by converting them into discrete simple allocation problems that can be solved greedily. Although more complicated algorithms with better runtimes are known, our simple incremental solution allows us to easily incorporate more general acceptance conditions such as scheduling and component constraints.

It remains to be seen whether our Marginal Bidding approach can be extended to handle interdependent uses, where devoting resources to one use can affect the marginal return of another. Interdependencies arise naturally in procurement because components are shared among SKU types.

Despite the game-theoretic nature of bidding in TAC SCM, our focus here was simply on a decision-theoretic (stochastic) optimization problem, not on game-theoretic equilibrium calculations. The enormity of the decision space in TAC SCM renders game-theoretic strategic analysis intractable with current technology. It remains to be seen whether an effective game-theoretic approach can be developed to exploit strategic opportunities in the TAC SCM game.

Finally, our ultimate desire is to extend the Marginal Bidding algorithm to accomplish marginal procurement as well, incorporating knowledge about future supply and demand conditions. Because the ILP considers each RFQ as a separate decision variable, its complexity grows rapidly as a function of the number of RFQs. By reasoning about SKUs in collective market segments, the Marginal Bidders avoid this complexity and appear to be more readily extensible to the procurement problem.

1.8 Acknowledgments

The work presented in this paper would not have been possible without programmatic support from Eric Sodomka, theoretical work from Victor Naroditsky and Warren Schudy and guidance from Amy Greenwald. We are also grateful to NSF Career Grant #IIS-0133689, caffeinated beverages and the color mauve.

References

1. M. Benisch, A. Greenwald, I. Grypari, R. Lederman, V. Naroditskiy, and M. Tschantz. Botticelli: A supply chain management agent. In *Third International Conference on Autonomous Agents and Multiagent Systems*, volume 3, pages 1174–1181, July 2004.
2. Michael Benisch, James Andrews, and Norman Sadeh. Pricing for customers with probabilistic valuations as a continuous knapsack problem. In *ACM International Conference Proceeding Series*, volume 156, pages 38 – 46, 2006.
3. David A. Burke, Kenneth N. Brown, Onur Koyuncu, Brahim Hnich, and Armagan Tarim. Foreseer: A constraint based agent for tac-scm. In *Trading Agent Competition - Supply Chain Management, IJCAI '05, Edinburgh, Scotland*, August 2005.

4. Amy Greenwald, Victor Naroditskiy, Tyler Odean, Mauricio Ramirez, Eric Sodomka, Joe Zimmerman, and Clark Cutler. Marginal bidding: An application of the equimarginal principle to bidding in tac scm. Technical Report CS-08-04, Brown University, 2008.
5. C Kiekintveld, MP Wellman, S Singh, J Estelle, Y Vorobeychik, V Soni, and M Rudary. Distributed feedback control for decision making on supply chains. In *Fourteenth International Conference on Automated Planning and Scheduling*, page To Appear, 2004.
6. A. Mas-Colell, M.D. Whinston, and J.R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
7. David Pardoe and Peter Stone. TacTex-03: A supply chain management agent. *SIGecom Exchanges*, 4(3):19–28, 2004.