An Analysis of the 2004 Supply Chain Management Trading Agent Competition

Christopher Kiekintveld, Yevgeniy Vorobeychik, and Michael P. Wellman

University of Michigan
Artificial Intelligence Laboratory
Ann Arbor, MI 48109-2110 USA
ckiekint@umich.edu

Abstract. We present and analyze results from the 2004 Trading Agent Competition supply chain management scenario. We identify behavioral differences between the agents that contributed to their performance in the competition. In the market for components, strategic early procurement remained an important factor despite rule changes from the previous year. We present a new experimental analysis of the impact of the rule changes on incentives for early procurement. In the finals, novel strategy designed to block other agent's access to suppliers at the start of the game was pivotal. Some agents did not respond effectively to this strategy and were badly hurt by their inability to get crucial components. Among the top three agents, average selling prices in the market for finished goods were the decisive difference. Our analysis shows that supply and demand were key factors in determining overall market prices, and that some agents were more adept than others at exploiting advantageous market conditions.

1 Introduction

The Trading Agent Competition (TAC) provides an international forum for researching the design and analysis of automated trading agents. A new scenario in supply chain management (TAC/SCM) debuted in 2003 [1]. We will not describe the scenario here, but direct the reader to the game specification for details [2]. Studying the outcomes of these competitions is a valuable exercise that helps us to better understand the strengths and weaknesses of current approaches. Here we present and analyze results from the final round of the 2004 tournament. Our primary objective is to determine the important behavioral factors that distinguished the agents and contributed to their relative performance.

We start by presenting the main results from the 2004 finals. In Section 3 we consider procurement strategies. We discuss the role of strategic early procurement, and present new experimental analysis of the effect of the rule changes on early procurement. A novel blocking strategy was used in the final round, and we discuss the problems this cause for some of the agents. In Section 4 we consider the PC sales market. Agents had widely varying average selling prices (ASPs) for PCs, and this was a deciding factor between the top three agents. Further analysis identifies four factors strongly correlated with market ASPs, and reveals that the agents differed in their ability to target profitable markets. We conclude with a summary of the strengths and weaknesses of each agent.

Table 1. Average scores and breakdowns from the TAC-04/SCM final round (in millions of dollars). Margin is the raw difference between revenue and supply costs.

Agent	Score	Revenue	Supplies	Margin	Storage	Penalty	Interest
FreeAgent	10.28	99.06	-80.94	18.12	-7.14	0	-0.61
Mr.UMBC	8.65	94.14	-76.49	17.65	-8.39	0	-0.61
UMTac-04	6.52	83.59	-67.21	16.37	-8.97	0	-0.88
Botticelli	0.44	25.83	-23.74	2.09	-0.39	-1.16	0
Deep Maize	-5.12	56.24	-48.32	7.92	-9.59	-2.49	-0.95
SouthamptonSCM	-10.41	71.37	-69.25	2.12	-11.34	0	-1.20

2 2004 Final Round Results

The 2004 competition started with qualifying and seeding rounds lasting two weeks each; these rounds were used primarily for development and testing. The top 24 agents participated in a three day tournament at AAMAS-04. A quarterfinal round eliminated 12 agents and a semi-final round eliminated 6 more. The surviving 6 agents played 14 games in the final round. Here we focus on analyzing the results of the final round, as it represents direct competition between the strongest agents. In Table 1 we see that the top three agents were very close, with scores in a narrow range of $4M. These agents had much higher raw margins than the bottom three agents, and generally higher transaction volume. Among the top three, small differences in storage costs and raw margins determined the ordering. SouthamptonSCM had supply volume comparable to the top finishers, but high storage costs and low revenue indicate possible sales problems. Botticelli and Deep Maize both transacted substantially less volume than the top finishers.

3 Agent Procurement in Supplier Markets

3.1 A History of Early Procurement

Understanding the market for PC components TAC-04/SCM is aided by discussion of results from the first competition. During the early rounds of the 2003 competition agent designers noticed that there were strong incentives to procure large quantities of components on the very first day of the simulation, day 0.[1] By the end of the seeding rounds most agents were making very large component purchases at the start of the game, before any information about customer demand was available. In games with low demand this could lead to large losses for all agents, as in one semi-final heat where all agents purchased aggressively and finished with negative average profits. In the other semi-final heat and the final round, Deep Maize surprised the field with a novel *preemptive* strategy that blocked the other agents from making large day-0 purchases.[2]

---

[1] Available supplier capacity was at a maximum, and prices were at a minimum.

[2] Essentially, the agent requested most of the supplier capacity for the entire game and the suppliers reserved it until the next day. Any requests considered after this request generated useless offers. The agent accepted a partial quantity from this request.

Post-tournament analysis showed that aggressive early procurement was a rational strategy despite the potential for negative profits, but that the presence of a preemptive agent could potentially improve profits for the entire field by knocking the agents out of the undesirable equilibrium [3].

While these strategic interactions were interesting, the extreme emphasis on early procurement detracted from other research problems, including factory scheduling [4], optimizing customer bids [5], and dynamically managing inventory in response to new information. The random order in which suppliers considered requests also introduced a "lottery effect," where these random outcomes had a strong effect on the overall outcome of the game [6]. Several changes were made to the specification for the 2004 competition; these were intended to reduce the incentives for day-0 procurement. The changes included modifications to the supplier pricing policy, segmentation of the customer markets, and the addition of storage costs.

3.2 Early Procurement in TAC-04/SCM

During the TAC-04/SCM qualifying round day-0 procurement remained very high, despite the rule changes. In response, the GameMaster increased storage costs fivefold for the remaining rounds. Even this did not dampen the day-0 purchasing; the number of components ordered based on day-0 requests actually *increased* by 14% from 2003 to 2004 (in games with no blocking strategies). This sequence of events raises questions about the impact of the rule changes (especially storage costs) on agent behavior. Do higher storage costs actually reduce incentives for early procurement, as suggested by intuitive arguments? Was the high level of early procurement observed in TAC-04/SCM a rational response to the new rules? Could any level of storage costs have reduced day-0 procurement to an acceptable level?

We address these questions with a systematic exploration of the relationship between storage costs and day-0 procurement Conceptually, each setting of storage costs induces a different game between the agents. Game theory suggests that stable profiles (e.g. Nash equilibria) are likely to be played when rational, self-interested agents compete in games. In general, we model the strategic interactions between a mechanism designer and participants as a two-stage game. The designer moves first by setting the mechanism parameter θ (e.g. storage costs), and all the participants observe θ and move simultaneously thereafter (e.g. selecting a day-0 procurement quantity). We refer to game between the participants in the second stage as game *induced* by θ:

$$\Gamma_\theta = [I, \{S_i\}, \{u_i(s, \theta)\}].$$

Suppose the goal of the designer is to optimize some welfare function $W(\cdot)$. Let $\{s^*(\theta)\}$ be the set of Nash equilibria of Γ_θ. Here we define $W(s^*(\theta), \theta) = \inf\{W(s, \theta) : s \in \{s^*(\theta)\}\}$. Alternatively, if one has a probability distribution over the Nash equilibria given θ, it may be natural to take the expectation of W instead: $W(s^*(\theta), \theta) = E_{s \in s^*}[W(s, \theta)]$.[3] If there are no Nash equilibria of Γ_θ (a possibility for infinite games),

---

[3] For example, such a distribution could be derived from analysis of evolutionary dynamics, as in [7].

let $\bar{s}^*$ be the set of strategy profiles with the lowest benefit to deviation for any agent and define $W(s^*(\theta), \theta) = \inf\{W(s, \theta) : s \in \{\bar{s}^*(\theta)\}\}$. The designer's optimization problem is then

$$\max_{\theta \in \Theta} W(s^*(\theta), \theta).$$

To analyze the effect of storage costs in the TAC/SCM game we consider the correspondence between storage costs and equilibrium outcomes. The aggregate quantity of day-0 procurement in these stable profiles yields an estimate of the behavior we would expect to see in actual games. In the TAC/SCM domain, the designer's problem is to minimize aggregate day-0 procurement. In the notation above, the designer maximizes $W(s^*(\theta), \theta)$, defined by $\mathbf{I}\{\sup\{\phi(s^*(\theta))\} \leq \alpha\}$, where $\phi(s) = \sum_{i=1}^{6} s_i$ and s_i is each agent's day-0 procurement choice.

The full strategy space in TAC/SCM is very complex, but for the purposes of this analysis we define a restricted space that allows agent to select only day-0 purchase quantity multiplier. We implemented this strategy space by parameterizing our tournament agent, Deep Maize, with a multiplier on its day-0 requests.[4] Players select a multiplier from the set $\{0, 0.3, 0.6, \ldots, 1.5\}$. This strategy space defines the induced game Γ_θ. The payoffs for this game are not directly known, but we can obtain estimates by simulating games on the TAC server. We must do this for each setting of storage costs we wish to investigate, and collecting the samples is very time-consuming. Fortunately, the questions we would like to answer are high-level and we can gather evidence about them using approximate methods. Instead of requiring exact equilibrium solutions, we aim to find regions of the profile space that are likely to be stable using the notion of ϵ-Nash equilibrium, where agents cannot gain more than a small benefit ϵ by deviating to a different strategy. We also use two different techniques to approximate sets of stable profiles without sampling the full profile space.

The first method approximates payoff functions of the game using supervised learning. We tried three different learning techniques from those introduced in [8]: quadratic regression (QR), locally weighted average (LWA), and locally weighted linear regression (LWLR). For quadratic regression, it is possible to directly compute equilibria of the learned game analytically. For the other methods, we applied replicator dynamics [9] to a discrete approximation of the learned game. The second method uses directed search to find stable profiles. Given a partial game matrix we can compute a bound on the epsilon for each profile that we have sample data for; this bound is the maximum benefit for deviating to any profile in the data set. The current set of profiles with the best ϵ-bounds is the set of *candidate equilibria*. We employed a "best-first" search that always samples unexplored deviations from a candidate equilibria. The idea is to confirm or refute the stability of promising individual profiles without requiring the full game matrix to be sampled. A limitation of this approach is that it cannot rule out the existence of additional equilibria in the set of profiles that have not been sampled.

We gathered data for storage costs in the set $\{0, 50, 100, 150, 200\}$. An initial data set was generated by sampling 10 randomly generated profiles for each storage cost setting (playing 5–10 games for each profile). We then performed 12–32 iterations of

---

[4] Deep Maize requested a total of 11800 components for each combination of supplier and product, spread out over different due dates.

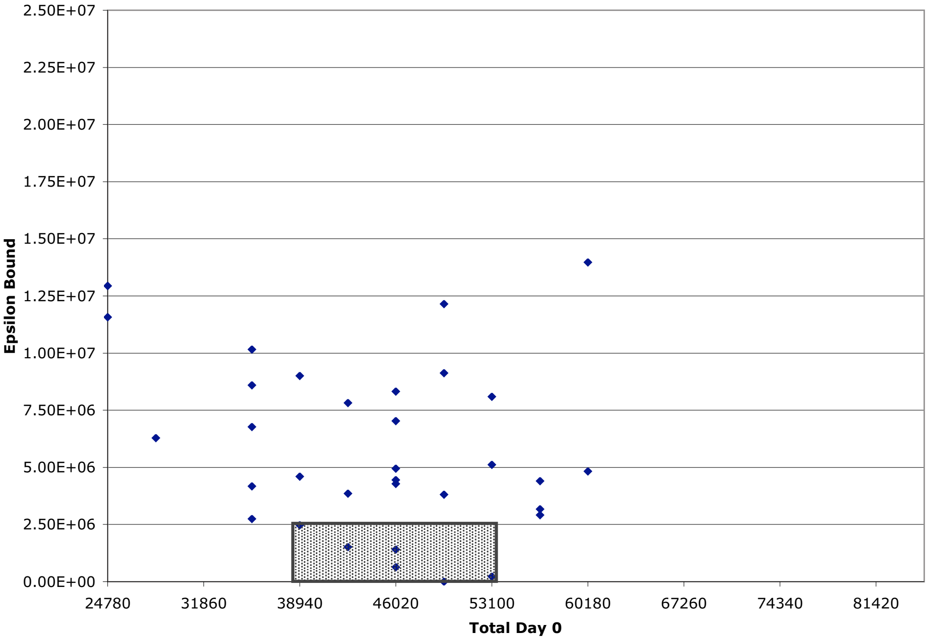

Fig. 1. Profile data for storage costs of 100% annually. The plot shows the ϵ-bound for explored profiles against the aggregate level of day-0 purchasing (per supplier/component) for all strategies in the profile. The dark box represents the region for which there are known profiles with ϵ-bounds less than \$2.5M.

the best-first search procedure for each setting of storage cost. We ran a total of 2670 games over 6 months, exploring approximately 10% of the total profile space for these discrete parameter and strategy settings.

Figure 1 shows a plot of the data annual storage costs of 100% annually (the mean storage cost setting from the 2004 tournament). Each point represents the ϵ-bound for a sampled profile, plotted against the aggregate day-0 procurement in the profile. Note that many different profiles have the same aggregate procurement. The dark box shows the region with the most stable (lowest-ϵ) profiles. This region yields a predicted range for the total day-0 procurement induced by this storage cost setting. To calibrate, an aggregate day-0 procurement of 3.0 corresponds to an expected commitment of 1/3 of the total supplier capacity for the entire game.

Figure 2 shows results for a range of settings of the storage cost parameter. The SearchMin and SearchMax lines correspond to the endpoints of the region define like the gray region in Figure 1. The other three lines indicate approximate equilibria found by the three learning methods, trained on the initial 10 randomly-generated profiles for each storage cost setting. It is encouraging that the results obtained using very different methods (learning and directed search) have the same qualitative structure. This experimental evidence supports the initial intuition that day-0 procurement should decrease with higher storage costs; all of the methods show this relationship. There is also

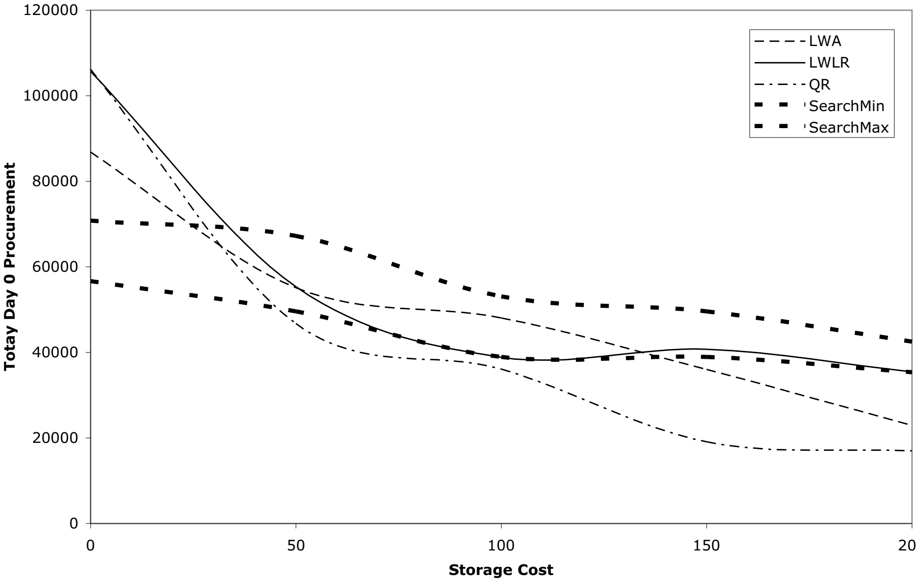

Fig. 2. Several different estimates for the correspondence between storage costs and aggregate day-0 purchases in equilibrium. Estimates from three learning methods are shown, along with an interval estimate from the best-first search algorithm. SearchMin is the minimum day-0 level for a profile with ϵ-bound less than 2.5M, and SearchMax is the corresponding maximum.

evidence that high levels of day-0 procurement were a rational response by agents to the new specification. The minimum prediction of any method for the storage cost setting used in the final round was approximately 3.0, and the maximum was considerably higher at approximately 4.5. The observed levels during the tournament were somewhat above even the high estimate given by our methods, but it seems clear that undesirably high levels of early purchasing are rational.

We also considered whether any setting of storage costs could have resulted in a desirable outcome for day-0 purchasing. To test this, we attempted to find a setting that would yield equilibrium outcomes with aggregate procurement less than 2.0 (still higher than we would want in practice). Linear extrapolation of the SearchMax line predicts that this should occur for a storage cost setting of 320%. However, futher simulations resulted in an estimated outcome range of 2.7–3.3 for this profile, only slightly lower than the estimates for storage costs of 200%. There appears to be very little benefit to additional increases of storage costs beyond 200%. Furthermore, agent profits were almost always negative for storage costs of 320%, so additional increases would be undesirable even if day-0 procurement could eventually be reduced to acceptable levels.

Our analysis suggests that the changes to the game rules did have the desired effect to some extent, but that this effect was not as large as anticipated. In games as complex as TAC/SCM it is very difficult to asses the effects of potential rule changes. In principle, the techniques used here provide ways of gathering additional data to assess the impact of design decisions in games with important strategic interactions.

Table 2. A breakdown of the percentages of CPU and non-CPU components ordered in response to requests sent on days 0, 1, and 2+, with respect to the total quantity ordered by each agent.

Agent	CPU 0	CPU 1	CPU 2+	Other 0	Other 1	Other 2+
FreeAgent	0.61	0.05	0.34	0	0.05	0.95
Mr.UMBC	0.57	0.16	0.27	0.29	0.47	0.24
UMTac-04	0.85	0.15	0.00	0.39	0.60	0.01
Botticelli	0	0	1	0	0	1
Deep Maize	0.90	0.01	0.09	0.46	0.02	0.52
SouthamptonSCM	0.81	0	0.19	0.54	0.38	0.07

3.3 A Blocking Strategy in TAC-04/SCM

The overall levels of day-0 procurement observed in TAC-04/SCM were similar to those observed in 2003, so it is perhaps not surprising that a blocking strategy again proved pivotal in the final round. However, the specific preemptive tactic employed by Deep Maize in 2003 was no longer useful due to changes in the supplier's pricing formula. The Deep Maize preemptive strategy relied on accepting partial fulfillment offers at low prices to purchase some cheap components on day 0, while still blocking some of the opponent's requests. In the 2004 rules these partial offers have very high prices, so a blocking agent must either pay high prices or try to purchase components later, despite a bad reputation with suppliers.

Blocking and purchasing later is a somewhat risky strategy, but FreeAgent used this tactic in the final round. [5] The strategy had a novel twist that mitigated some of the risks: FreeAgent only blocked requests for non-CPU components and purchased large quantities of CPU components on day 0 along with the other agents. This is significant because CPUs cost much more than any of the other components (on average, the CPU represents half of the total cost of components for a PC). The strategy locked in relatively low prices for the components with the highest base prices, but risked paying relatively higher prices for the cheaper components in order to disrupt the other agents' procurement strategies. Table 2 illustrates the effect of this blocking strategy on the distribution of purchases over the game for each agent. Most of the agents, including FreeAgent, procured the majority of their CPU components on day 0 to take advantage of the low prices. FreeAgent used its requests for "other" (non-CPU) components on day 0 to block opponents' requests. Consequently, it ordered no components of these types on day 0, and the quantities purchased by the other agents were reduced.

One of the crucial disparities between the agents was how they reacted to the new market environment created by FreeAgent's blocking strategy. Mr.UMBC, UMTac-04, and SouthamptonSCM had backup strategies that procured large quantities of components again on the next simulation day. Since most of the supplier capacity was uncommitted at this point, the agents secured reasonably low prices for these backup orders and were in much the same situation as if they had ordered all of their components on day 0 as originally planned. Botticelli made large day-0 requests like the rest

---

[5] Mr.UMBC also submitted large blocking requests for some types non-CPU components. However, these RFQs were relatively low in the priority ordering assigned by the agent and likely did not come into play.

Table 3. Average total quantities ordered and prices paid for components.

Agent	Ordered	CPU Price	Other Price
FreeAgent	40199	0.60	0.88
Mr.UMBC	42557	0.61	0.71
UMTac-04	40738	0.56	0.61
Botticelli	11598	0.74	0.70
Deep Maize	26902	0.58	0.65
SouthamptonSCM	39350	0.58	0.71

Table 4. Inventory management statistics. CPU/Other is the ratio of CPUs ordered to the number of other components ordered. The ratio needed to produce PCs is 0.33 (1 CPU:3 Other).

Agent	CPU/Other	Unsold CPU	Unsold Other	Daily Inv.	Ave. Delivery Day
FreeAgent	0.34	3516	5956	30677	95
Mr.UMBC	0.31	713	15769	48877	85
UMTac-04	0.33	1832	2588	41975	94
Botticelli	0.33	475	1497	1317	160
Deep Maize	0.40	12153	10200	31026	88
SouthamptonSCM	0.33	13651	43922	67147	69

of the agents, but chose not to accept any of the offers and waited until much later in the game to purchase supplies and start production. FreeAgent and Deep Maize did not come back with large requests immediately, but purchased additional components in smaller chunks throughout the rest of the game.

Additional details about overall procurement and pricing are in Table 3. Four of the agents (including the top three) purchased approximately the same number of components overall; Botticelli and Deep Maize purchased significantly fewer components. Prices paid for CPUs are comparable for all agents except Botticelli, which did not procure any cheap CPUs on day 0. The prices paid for non-CPU components show more disparity. UMTac-04 paid the lowest prices due to a large day-1 purchase. FreeAgent paid *very* high prices for non-CPU components. This stands in contrast to the much lower prices paid by Deep Maize, despite the similar approach these agents took in purchasing additional inventory over the duration of the game. FreeAgent seems to have disregarded the prices paid for these components, in part to compensate for its strategic maneuver on day 0.

3.4 Inventory Management

Agents' procurement strategies had important implications for inventory management. The two lowest-scoring agents in particular were crippled by difficulties in managing inventory that are at least partially attributable to FreeAgent's blocking tactic. There are a number of striking numbers in table 4, which lists inventory management statistics. The first is Botticelli's very late average delivery date for components. This agent effectively sat out most of the game after declining early component purchases. Equally striking are the very large unsold inventories held by Deep Maize and Southampton-

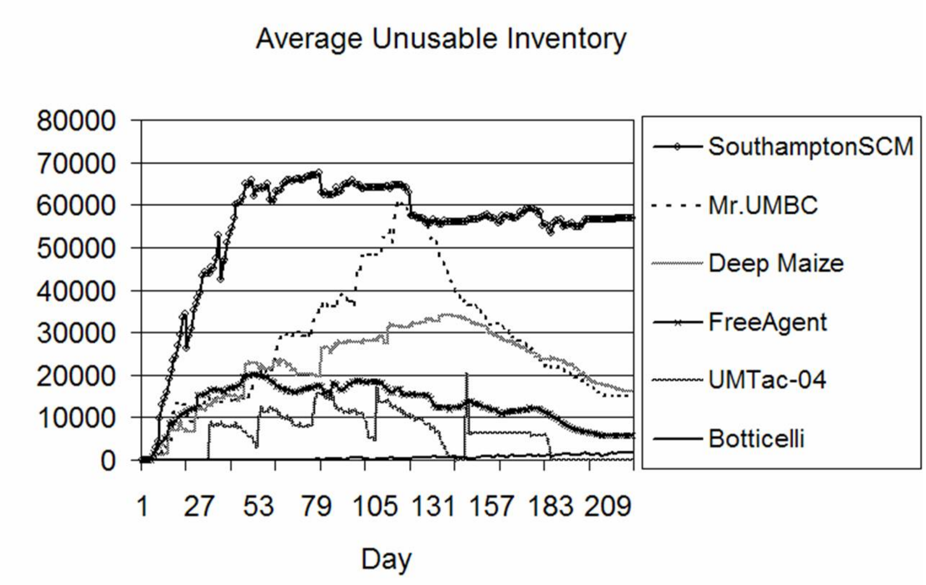

Fig. 3. The average cumulative difference between components delivered and components that could potentially be assembled into finished PCs for each agent. High differences indicate that the agent is missing one or more complementary components (e.g. the agent has no memory components).

SCM at the end of the game. Mr.UMBC also had a substantial unsold inventory, but this inventory was almost exclusively composed of much cheaper non-CPU components.

Figure 3 gives a more detailed breakdown of inventory management over the course of the game. The plot shows the difference between the number of components delivered and the number that could possibly be assembled for sale, emphasizing management of the complementarities between components. The bulk of the unsold inventories could not have been sold due to not having the right combinations of components. The unsold inventory problem was not very severe for Mr.UMBC since it was composed of cheaper components; this may actually have been a deliberate hedge against production down time.

Deep Maize and SouthamptonSCM both purchased large quantities of some components, but had difficulties obtaining enough complementary components to allow full production. SouthamptonSCM had large differences in the orders it placed early on, and did not compensate for these disparities later in the game. The imbalances for Deep Maize were not quite as large as those for SouthamptonSCM, and it was able to mitigate them to some extent by procuring additional components throughout the game. However, the agent was very selective about the prices paid for these additional components (note the low price paid for non-CPUs from Table 3). That FreeAgent was very successful with a similar strategy that paid much higher prices for these compo-

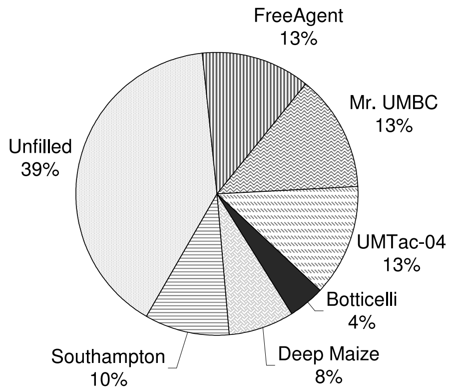

Fig. 4. Market share in the SCM finals.

nents suggests that Deep Maize was probably too selective about prices for certain key components.

A final point about inventory management is that FreeAgent had a substantially lower average daily inventory than the next two agents, despite comparable overall quantities purchased. This appears to be the result of ordering non-CPU inventory over many days, which allowed the agent to divide quantities into more requests and distribute them move evenly over time. This game FreeAgent some advantage in storage costs over these two agents, and helped to offset the higher prices paid by the agent for non-CPU components.

4 Agent Sales Behavior

4.1 Basic Sales Data

We now attend to activity in the PC sales market, starting with overall market share in Figure 4 (raw sales numbers are listed in Table 5). Market shares mostly parallel the total quantities of components ordered, as given in Table 3. The major exception is SouthamptonSCM, which purchased similar quantities to the top three agents but took a much smaller share of the customer market. This reflects the large amount of unsold inventory for this agent noted in the previous section. We also note that almost 40% of the total customer demand was unmet. Some of this is unavoidable since agents must build inventory at the start of the game before they can sell PCs. However, there does seem to be significant opportunity for agents to expand market share by filling unmet demand.

More detailed information on sales activity is given in Table 5. The top three agents sold very nearly the same number of PCs, and won an almost identical fraction of

Table 5. Customer sales statistics for each agent.

Agent	PCs sold	PCs bid on	Percent Won	ASP
FreeAgent	54660	166293	0.33	0.90
Mr.UMBC	56748	172517	0.33	0.83
UMTac-04	56341	167746	0.34	0.74
Botticelli	17290	28581	0.60	0.81
Deep Maize	33166	124485	0.27	0.92
SouthamptonSCM	41798	93936	0.44	0.85

their bids. However, they had strikingly different average selling prices (ASPs). This difference in ASPs is one of the major reasons that the three agents finished in the order they did. That the three agents bid on almost identical fractions of the market and win similar fractions of bids suggests that this difference in ASPs is due to targeting different types of markets. Deep Maize had the highest ASP of any agent and the lowest winning percentage, suggesting that this agent made systematically higher offers to customers. Conversely, Botticelli won a very high fraction of its bids with a relatively low ASP, suggesting systematically lower bids. SouthamptonSCM won a high fraction of bids with mid-range ASPs, but bid on a much smaller fraction of the market than the other three agents with similar component purchases.

4.2 Market Behavior

To better understand why the agents had different ASPs, we consider features that differentiate markets along the dimension of ASP. For this analysis we consider a “market” to be the set of customer requests for a PC type on a simulation day. We identify four factors that are strongly correlated with overall market ASPs. Figure 5 shows these relationships as scatter plots, with superimposed lines representing binned averages. Except for simulation day, all of these factors are measures of supply and demand motivated by basic economic principles. The simulation day is important primarily due to start- and end-game effects.

Plot 5(a) shows the relationship between prices and simulation day. Prices start very high early in the game as agents build inventory and decrease over time. At the end of the game prices can fall very low as agents try to recover some value for excess inventory. The second factor, shown in 5(b), is market demand (i.e., total quantity requested). Prices increase as demand increases, with any demand level less than 100 occasionally subject to very low ASPs. Plot 5(c) is a measure of bid density calculated this by summing the number of bids for each individual PC and dividing by the total number of PCs requested. ASPs fall approximately linearly as bid density increases. ASPs also fall approximately linearly as manufacturer PC inventory increases, as seen in 5(d).

We ran linear regressions to test the strength of these relationships with ASPs (all of the plots suggest a linear relationship, so this is a reasonably approximation). The individual R^2 values were 0.29 for simulation day, 0.13 for market demand, 0.42 for bid density, and 0.44 for PC inventory. A multiple regression using all four factors yields an R^2 value of 0.65, with all coefficients significantly different from 0. We could

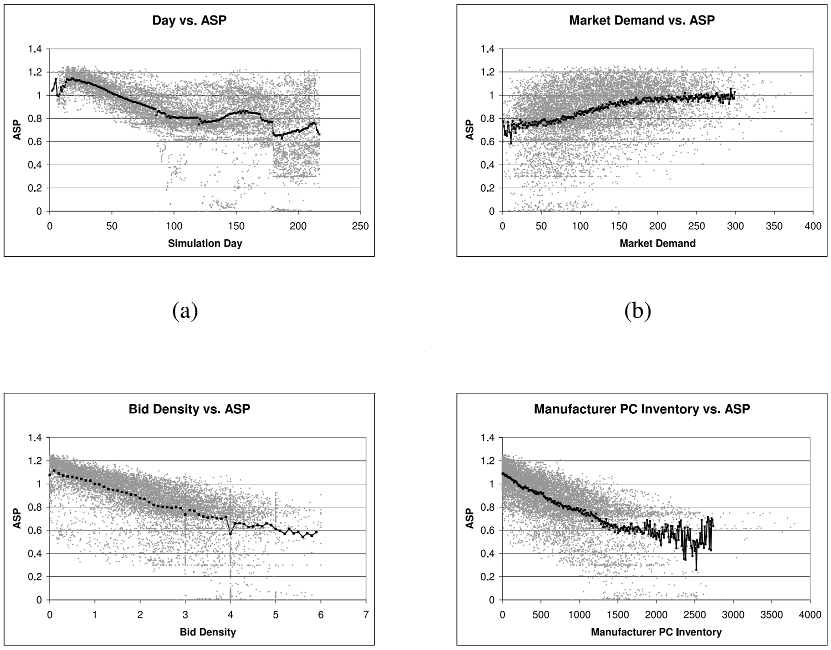

(a) (b)

(c) (d)

Fig. 5. Plots showing the relationships between four variables and market ASPs in SCM-04. 10000 randomly selected points are shown for each variable (out of approximately 45000 total). The dark lines on each plot are the average ASPs at each value for discrete variables or within small bins for continuous variables.

certainly improve this model by considering additional factors (e.g. reserve prices, lead times, smoothed market demand). However, a simple linear fit to these four variables is sufficient to explain 65% of the variance in ASPs. This surprisingly powerful result that speaks to the power of the basic forces of supply and demand in the SCM markets.

We take this analysis a step further by considering the features of the markets that individual agents bid and sold PCs in. Table 6 gives average values for each of these four factors for each of the PCs sold and PCs bid on by each agent. The two agents with the highest ASPs were Deep Maize and FreeAgent. These agents sold PCs earlier in the game, in markets with the lowest PC inventory, highest demand, and low bid densities; all of these are associated with high market ASPs. UMTac-04 had the lowest ASP by a wide margin. This agent sold later than most agents in markets with low demand and high PC inventory levels; all of these are associated with low market ASPs. The other agents have mid-range ASPs, and seem to have a mix of factors working for and against them. For instance, SouthamptonSCM sells exceptionally early, but in lower demand markets with relatively high inventory levels.

Table 6. Average values of four factors correlated with market ASPs given for the PCs each agent won and bid on. Day is the simulation day, Demand is market demand, Bids is the bid density (defined in text), and PC inv is average manufacturer PC inventory.

Agent	Sold, Day	Bid, Day	Sold, Demand	Bid, Demand
FreeAgent	122	130	160	147
Mr.UMBC	124	129	150	148
UMTac-04	135	126	143	140
Botticelli	163	164	153	146
Deep Maize	122	123	168	158
SouthamptonSCM	105	112	143	136

Agent	Sold, Bids	Bid, Bids	Sold, PC Inv	Bid, PC Inv
FreeAgent	2.89	3.39	588	832
Mr.UMBC	2.58	3.26	761	813
UMTac-04	2.90	3.48	933	925
Botticelli	3.75	3.99	730	822
Deep Maize	2.40	3.35	551	675
SouthamptonSCM	3.38	3.71	836	939

5 Conclusions: Agent Performance

We have presented and analyzed data on many aspects of the TAC-04/SCM final round, which has revealed interesting details about the different strategies employed by the agents. We can now offer a reasonably compelling explanation for why the agents finished in the order they did. The following is a brief synopsis of the important points about each agent, starting from the lowest-scoring agent and working up.

SouthamptonSCM bought far more components than it sold, leaving it with large amounts of unsold inventory and high storage costs. The agent had reasonably high ASPs and winning percentages in the customer market, but bid on a very low fraction of customer orders. The underlying problem was that the agent was not successful at getting all of the complementary components needed for production of finished products.

Deep Maize had very high ASPs in the customer market. However, it also had inventory management problems with complementary components. The agent was often left with insufficient non-CPU inventory to manufacture and sell aggressively, and had large numbers of unsold CPU components. The agent paid low average prices for non-CPU components, so it may have been too selective about prices for these components.

Botticelli did not make any supply purchases at the start of the game, and did not enter the market substantially until very late in the game. By this time there was too little time remaining in the game and the customer market was too competitive to allow for large profits.

UMTac-04 compared favorably to FreeAgent on virtually every metric we considered except average daily inventory and customer market ASP. This agent had the lowest ASP of any agent. It sold later in the game in markets with relatively low demand, high PC inventory; all of these factors correlate with low market prices.

Mr UMBC had a better ASP than UMTac-04, but lower than FreeAgent. It sold in lower demand and higher inventory markets than FreeAgent. This agent also had a sizable amount of non-CPU inventory unsold at the end of the game and higher daily inventories than FreeAgent.

FreeAgent made some interesting strategic choices for the final round. It opted to block other agents only on non-CPU inventory, while acquiring CPUs at low prices on day 0. It then purchased additional components at very high prices during the rest of the game. It was able to compensate to some extent for these high prices through more evenly spaced deliveries of components and lower resulting storage costs. The major strength of this agent compared to the others was an ability to consistently sell high volumes of PCs throughout game at high prices by targeting profitable markets.

Acknowledgments

We thank SICS and the rest of the TAC/SCM organizers and participants for another interesting opportunity to analyze trading agent strategy. This work was supported in part by NSF grant IIS-0205435 and the DARPA REAL strategic reasoning program.

References

1. Arunachalam, R., Sadeh, N.M.: The supply chain trading agent competition. Electronic Commerce Research and Applications **4** (2005) 63–81
2. Arunachalam, R., Sadeh, N., Eriksson, J., Finne, N., Janson, S.: The supply chain management game for the trading agent competition 2004. Technical Report CMU-CS-04-107, Carnegie Mellon University (2004)
3. Wellman, M.P., Estelle, J., Singh, S., Vorobeychik, Y., Kiekintveld, C., Soni, V.: Strategic interactions in a supply chain game. Computational Intelligence **21** (2005) 1–26
4. Benisch, M., Greenwald, A., Naroditskiy, V., Tschantz, M.: A stochastic programming approach to scheduling in TAC SCM. In: Fifth ACM Conference on Electronic Commerce. (2004)
5. Pardoe, D., Stone, P.: TacTex-03: A supply chain management agent. SIGecom Exchanges **4**(3) (2004) 19–28
6. Ketter, W., Kryzhnyaya, E., Damer, S., McMillan, C.: Analysis and design of supply-driven strategies in TAC SCM. In: AAMAS-04 Workshop on Trading Agent Design and Analysis. (2004)
7. Walsh, W.E., Das, R., Tesauro, G., Kephard, J.O.: Analyzing complex strategic interactions in multi-agent systems. In: AAAI-02 Workshop on Game Theoretic and Decision Theoretic Agents. (2002)
8. Vorobeychik, Y., Wellman, M.P., Singh, S.: Learning payoff functions in infinite games. In: Ninth International Joint Conference on Artificial Intelligence. (2005) 977–982
9. Friedman, D.: Evolutionary games in economics. Econometrica **59**(3) (1991) 637–666