

Strategic Interactions in the TAC 2003 Supply Chain Tournament

Joshua Estelle, Yevgeniy Vorobeychik, Michael P. Wellman, Satinder Singh,
Christopher Kiekintveld, and Vishal Soni

University of Michigan
Artificial Intelligence Laboratory
Ann Arbor, MI 48109-2110 USA
{ jestelle, yvorobey, wellman, baveja, ckiekint, soniv }@umich.edu

Abstract. The TAC 2003 supply-chain game presented automated trading agents with a challenging strategic problem. Embedded within a complex stochastic environment was a pivotal strategic decision about initial procurement of components. Early evidence suggested that the entrant field was headed toward a self-destructive, mutually unprofitable equilibrium. Our agent, **Deep Maize**, introduced a preemptive strategy designed to neutralize aggressive procurement, perturbing the field to a more profitable equilibrium. It worked. Not only did preemption improve **Deep Maize**'s profitability, it improved profitability for the whole field. Whereas it is perhaps counterintuitive that action designed to prevent others from achieving their goals actually helps them, strategic analysis employing an empirical game-theoretic methodology verifies and provides insight about this outcome.

1 Introduction

Like classic computer games, multiagent research competitions [Stone, 2002] present well-defined problems for testing and comparing AI techniques and systems. The annual Trading Agent Competition (TAC) series provides a forum for research on strategic market behavior, and has led to several promising concepts and methods for implementing strategies in such domains [Wellman et al., 2003].

The TAC Supply Chain Management (TAC/SCM) scenario [Sadeh et al., 2003] defines a complex six-player game with severely incomplete and imperfect information, and high-dimensional strategy spaces. Like the real supply-chain environments it is intended to model, the TAC/SCM game presents participants with challenging decision problems in a context of great strategic uncertainty. This paper is a case study of a strategic issue that arose in the first TAC/SCM tournament. We present our reasoning about the issue, and our effort to perturb the environment from an “equilibrium” we considered undesirable, to another more profitable domain of operation. We recount the experience as it played out in the competition, and analyze the outcome of this naturalistic experiment. We then perform a more controlled experimental analysis of the issue, applying empirical game-theoretic methods to produce compelling results, narrow in scope but arguably accounting well for strategic interactions.

2 TAC/SCM Game

In the TAC/SCM scenario,¹ six agents representing PC (personal computer) assemblers operate in a common market environment, over a simulated year. The environment constitutes a *supply chain*, in that agents trade simultaneously in markets for supplies (PC components) and the market for finished PCs. Agents may assemble for sale 16 different models of PCs, defined by the compatible combinations of the four component types: CPU, motherboard, memory, and hard disk.

Figure 1 diagrams the basic configuration of the supply chain. The six agents (arrayed vertically in the middle of the figure) procure components from the eight suppliers on the left, and sell PCs to the entity representing customers, on the right. Trades at both levels are negotiated through a *request-for-quote* (RFQ) mechanism, which proceeds in three steps:

1. Buyer issues RFQs to one or more sellers.
2. Sellers respond to RFQs with *offers*.
3. Buyers accept or reject offers. An accepted offer becomes an *order*.

The suppliers and customer implement fixed negotiation policies, defined in the game specification, and discussed in detail below where applicable.

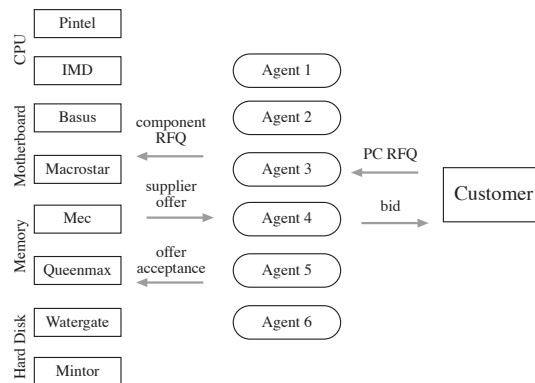


Fig. 1. TAC/SCM supply chain.

The game runs for 220 simulated days. On each day, the agent may receive offers and component delivery notices from suppliers, and RFQs and offer acceptance notifications from customers. It then must make several decisions:

1. What RFQs to issue to component suppliers.
2. Given offers from suppliers (based on the previous day's RFQs), which to accept.

¹ For complete details of the game rules, see the specification document [Arunachalam et al., 2003]. This is available at <http://www.sics.se/tac>, as is much additional information about TAC/SCM and TAC in general.

3. Given component inventory and factory capacity, what PCs to manufacture.
4. Given inventory of finished PCs, which customer orders to ship.
5. Given RFQs from customers, to which to respond and with what offers.

In the simulation, the agent has 15 seconds to compute and communicate its daily decisions to the game server. At the end of the game, agents are evaluated by total profit, with any outstanding component or PC inventory valued at zero.

As we describe below, a key stochastic feature of the game environment is level of demand for PCs. The underlying demand level is defined by an integer parameter Q (called RFQ_{avg} in the specification document [Arunachalam et al., 2003, Section 6]). Each day, the customer issues a set of \hat{Q} RFQs, where \hat{Q} is drawn from a Poisson distribution with mean value defined by the parameter Q for that day. Since the order quantity, PC model, and reserve price are set independently for each customer RFQ, the number of RFQs serves as a sufficient statistic for the overall demand, which in turn is a major determinant of the potential profits available to the agents.

The demand parameter Q evolves according to a given stochastic process. In each game instance, an initial value, Q_0 , is drawn uniformly from [80,320]. If Q_d is the value of Q on day d , then its value on the next day is given by [Arunachalam et al., 2003, Section 6]:

$$Q_{d+1} = \min(320, \max(80, \tau_d Q_d)), \quad (1)$$

where τ is a trend parameter that also evolves stochastically. The initial trend is neutral, $\tau_0 = 1$, with subsequent trends updated by a perturbation $\epsilon \sim U[-0.01, 0.01]$:

$$\tau_{d+1} = \max(0.95, \min(1/0.95, \tau_d + \epsilon)). \quad (2)$$

In a given game, the demand may stay at predominantly high or low levels, or oscillate back and forth. The probabilistic behavior of Q figures importantly in our analysis, as presented in Section 5.5 below.

3 Deep Maize

The University of Michigan’s entry in TAC-03/SCM is an agent called **Deep Maize** [Kiekintveld et al., 2004a,b]. The agent is organized in modular functional units controlling procurement, manufacturing, and sales. Its behavior is based on distributed feedback control, in that it acts to maintain a reference zone of profitable operation. To coordinate the distributed modules, **Deep Maize** employs aggregate price signals, derived from a market equilibrium analysis and continual Bayesian demand projection. The design of **Deep Maize** optimizes for performance in the steady-state, with little explicit attention to transient or end-game behaviors.

In the present study we focus on one pivotal feature of **Deep Maize**’s strategy, described in full detail below. We thus defer specifics of the rest of our agent’s strategy to our other reports (which in turn do not address the strategic analysis presented here).

4 Day-0 Procurement Strategies

A close examination of the game rules suggests that procurement of components at the very beginning of the game (*day-0 procurement*) may be a pivotal strategic issue. This was indeed borne out by the behavior observed in preliminary rounds of the tournament, as discussed below. In this section, we explain the reason for expecting day-0 procurement to be so significant, and its ramifications for Deep Maize and other agents.

4.1 Supplier Pricing

In the TAC/SCM market, suppliers set prices for components based on an analysis of their available capacity. Conceptually, there exist separate prices for each type of component, from each supplier. Moreover, these prices vary over time: both the time that the deal is struck, and time that the component is promised for delivery.

The TAC/SCM component catalog [Arunachalam et al., 2003, Figure 3] associates every component c with a *base price*, b_c . The correspondence between price and quantity for component supplies is defined by the suppliers' pricing formula [Arunachalam et al., 2003, Section 5.5]. The price offered by a supplier at day d for an order to be delivered on day $d + i$ is

$$p_c(d + i) = b_c - 0.5b_c \frac{\kappa_c(d + i)}{500i}, \quad (3)$$

where $\kappa_c(j)$ denotes the cumulative capacity for c the supplier projects to have available from the current day through day j . The denominator, $500i$, represents the *nominal capacity* controlled by the supplier over i days, not accounting for any capacity committed to existing orders.

Supplier prices according to Eq. (3) are date-specific, depending on the particular pattern of capacity commitments in place at the time the supplier evaluates the given RFQ. A key observation is that component prices are never lower than at the start of the game ($d = 0$), when $\kappa_c(i) = 500i$ and therefore $p_c(i) = 0.5b_c$, for all c and i .² As the supplier approaches fully committed capacity ($\kappa_c(d + i) \rightarrow 0$), $p_c(d + i)$ approaches b_c .

In general, one would expect that procuring components at half their base price would be profitable, up to the limits of production capacity. Customer reserve prices range between 0.75 and 1.25 the base price of PCs, defined as the sum of base prices of components. Therefore, unless there is a significant oversupply, prices for PCs should easily exceed the component cost, based on day-0 prices.

An agent's procurement strategy must also take into account the specific TAC/SCM RFQ process. Each day, agents may submit up to 10 RFQs, ordered by priority, to each supplier. The suppliers then repeatedly execute the following, until all RFQs are exhausted: (1) randomly choose an agent, (2) take the highest-priority RFQ remaining

² As discussed below, this creates a powerful incentive for early procurement, with significant consequences for game balance. In retrospect, the supplier pricing rule was generally considered a design flaw in the game, and has been substantially revised for the 2004 TAC/SCM tournament.

on its list, (3) generate a corresponding offer, if possible. In responding to an RFQ, if the supplier has sufficient available capacity to meet the requested quantity and due date, it offers to do so according to its pricing function. If it does not, the supplier instead offers a partial quantity at the requested date and/or the full quantity at a later date, to the best of its ability given its existing commitments. In all cases, the supplier quotes prices based on Eq. (3), and reserves sufficient capacity to meet the quantity and date offered.

4.2 Implications of Aggressive Day-0 Procurement

From the discussion above, it would appear advantageous to any agent that it attempt to procure a large number of components on day 0. We call this strategy *aggressive day-0 procurement*, or simply *aggressive*. From each agent's perspective, the main effect of being aggressive is on its own component procurement profile. If every agent is aggressive, however, it can significantly change the character of the game environment.

An aggressive day-0 procurement commits to large component orders before overall demand over the game horizon is known. This leaves agents with little flexibility to respond to cases of low demand, except by lowering PC prices to customers. Since component costs are sunk at the beginning, there is little to keep prices from dropping below (ex ante) profitable levels.

As more agents procure aggressively, several factors make aggressiveness even more compelling. The aggressive agents reserve significant fractions of supplier capacity, thus reducing subsequent availability and raising prices, according to their pricing function (3). A natural response might induce a "race" dynamic, where agents issue day-0 RFQs in increasingly large chunks, ultimately requesting all components they expect to be able to use over the entire game horizon. Not only does this exacerbate the risk of locking in aggregate oversupply, it also produces a less interleaved and more unbalanced distribution of components, especially at the beginning of the game. This in turn can prevent many agents from being able to acquire key components needed for particular PC models until relatively far into the production year.

For all these reasons, the aggressive strategy is appealing to individual agents, yet potentially quite damaging for the agent pool overall. We considered this situation particularly bad for our agent, given that it was designed for high performance in the steady state [Kiekintveld et al., 2004a]. **Deep Maize** devotes a considerable effort toward developing accurate demand projections, and thus is quite responsive to actual demand conditions. If most of the game's component procurement is up front, we never reach a steady state, and the ability to respond to demand conditions is much less relevant.

The **Deep Maize** development team therefore decided not to employ aggressive day-0 procurement in the preliminary rounds, instead treating it just like any other day. We did not really expect that others would miss the opportunity, but did not want to encourage or accelerate it.

5 TAC-03 Tournament

The twenty agents who participated in the TAC-03/SCM tournament are listed in Table 1. The table presents average scores from each of three preliminary rounds, mea-

Table 1. TAC-03/SCM tournament participants, and their performance in preliminary rounds. Results from the qualifying rounds are weighted, seeding rounds are unweighted.

Agent	Affiliation	Average Profit (\$M)		
		Qualifying	Seeding 1	Seeding 2
TacTex	U Texas	33.65	32.66	32.97
RedAgent	McGill U	15.09	24.57	29.52
Botticelli	Brown U	13.88	17.29	28.03
Jackaroo	U Western Sydney	14.89	35.55	19.23
WhiteBear	Cornell U	-3.17	13.57	16.50
PSUTAC	Pennsylvania State U	-120.0	15.52	15.25
HarTAC	Harvard U	12.41	4.19	10.72
UMBCTAC	U Maryland Baltimore Cty	-13.94	30.16	10.23
Sirish		-109.4	-0.17	8.27
Deep Maize	U Michigan	1.85	0.45	7.49
TAC-o-matic	Uppsala U	0.22	1.79	7.07
RonaX	Xonar GmbH	-0.92	9.24	4.29
MinneTAC	U Minnesota	10.88	6.56	-0.32
Mertacor	Aristotle U Thessaloniki	9.29	-0.38	-3.53
zepp	Poli Bucharest	-24.83	-7.80	-5.46
PackaTAC	N Carolina State U	-5.11	-25.67	-5.71
Socrates	U Essex	-48.94	-3.31	-6.84
Argos	Bogazici U	3.65	-4.24	-8.43
DummerAgent		-8.08	-20.56	—
DALhard	U Tulsa	-11.36	-39.05	—

sured in millions of dollars of profit. Results from the semifinal and final rounds are presented in Section 5.3 below.

Two seeding rounds were held during the periods 7–11 and 14–18 July,³ with each agent playing 60 and 66 games, respectively. Two agents were eliminated based on scores and/or inactivity after Seeding Round 1. The remaining 18 agents advanced to the semifinals, with assignment to heats based on standing in Seeding Round 2. The semifinals and finals were held live at IJCAI-03, 11–13 August in Acapulco, Mexico, each round consisting of nine games in one day. Semifinal 1 heat 1 (S1H1) comprised agents seeded 1–6 and 16–18, and the 7–15 seeds played in S1H2. The top six teams from each S1 heat (9 games) proceeded to the second semifinal round. S2H1 comprised teams ranked 1–3 in S1H1, and those ranked 4–6 in S1H2. The top three in S1H2 played, along with the second three in S1H1, in S2H2. The top three from each of S2H1 and S2H2 then entered the finals on 13 August. Further details about the TAC-03 tournament are available at <http://www.sics.se/tac>.

³ An earlier “qualifying” round spanned 16–27 June, but this was mainly for debugging and no agents were eliminated.

5.1 Evolution of Day-0 Policies in Preliminary Rounds

As we expected, competition entrants noticed the individual advantages of aggressive day-0 procurement. Early in the qualifying rounds we noticed **Jackaroo**'s distinct saw-tooth shaped profits, indicating a steady increase in wealth with large periodic drops corresponding to large deliveries of supplies. This pattern was the result of large supply orders placed early in the game (over the first seven days, not just day 0) for delivery at regular intervals [Zhang et al., 2004].

Based on our subsequent analysis of early game logs,⁴ we can identify **TacTex** [Par-doe and Stone, 2004] as the first to employ an aggressive day-0 strategy in competition. In their very first qualifying round game, **TacTex** requested 8000 of each component from each supplier. Although we have found many agents performed mild day-0 procurement during the qualifying rounds, **TacTex** was more aggressive, earlier—likely a factor in their supremacy this first round.

Throughout the first seeding round, more agents began using increasingly aggressive day-0 procurement strategies. In particular we noticed the successful agents **TacTex**, **Botticelli**, **RedAgent**, **UMBCTAC**, and **Jackaroo** ordering very large quantities on day 0 and very little later in the game. Interestingly, there was no discussion of the issue on the TAC/SCM message boards, possibly because entrants recognized its strategic sensitivity. By the second seeding round it was obvious that the majority of agents were using aggressive strategies. In particular, we verified that all the agents that placed higher than **Deep Maize** in the second seeding round (see Table 1) employed aggressive day-0 procurement.

While observing the increase in aggressiveness, we compiled detailed dossiers describing the day-0 strategies of other agents. We hoped to use this data to understand how widespread the use of day-0 procurement had become, and to understand how it was affecting the dynamics of the game.

5.2 Deep Maize Preemptive Strategy

After much deliberation, we decided that the only way to prevent the disastrous rush toward all-aggressive equilibrium was to *preempt* the other agents' day-0 RFQs. By requesting an extremely large quantity of a particular component, we would prevent the supplier from making reasonable offers to subsequent agents, at least in response to their requests on that day. Our premise was that it would be sufficient to preempt only day-0 RFQs, since after day 0 prices are not so especially attractive.

The **Deep Maize** preemptive strategy operates by submitting a large RFQ to each supplier for each component produced. The preemptive RFQ requests 85000 units—representing 170 days' worth of supplier capacity—to be delivered by day 30. See Figure 2. It is of course impossible for the supplier to actually fulfill this request. Instead, the supplier will offer us both a partial delivery on day 30 of the components they can offer by that date (if any), and an earliest-complete offer fulfilling the entire quantity

⁴ The TAC/SCM game server records all agent actions (e.g., RFQs, manufacturing, bids) along with supplier and customer behavior, and releases the log files after each game instance is complete.

(unless the supplier has already committed 50 days of capacity). With these offers, the supplier reserves necessary capacity. This has the effect of preempting subsequent RFQs, since we can be sure that the supplier will have committed capacity at least through day 172. (The extra two days account for negotiation and shipment time.) We will accept the partial-delivery offer, if any (and thereby reject the earliest-complete), giving us at most 14000 component units to be delivered on day 30, a large but feasible number of components to use up by the end of the game.



Fig. 2. Deep Maize’s preemptive RFQ.

The TAC/SCM designers anticipated the possibility of preemptive RFQ generation, (there was much discussion about it in the original design correspondence), and took steps to inhibit it. The designers instated a reputation mechanism, in which refusing offers from suppliers reduces the priority of an agent’s RFQs being considered in the future. Even with this deterrent, we felt our preemptive strategy would be worthwhile. Since most agents were focusing strongly on day 0, priority for RFQ selection on subsequent days might not turn out to be crucial.

5.3 Tournament Story

Having developed the preemptive strategy, we still faced the question of when to deploy it. Based on our performance in preliminaries, we were reasonably confident that we could make the top six out of nine in S1H2 without resorting to preemption, and instead chose to implement a moderate form of aggressive day-0 procurement. As expected, other agents actually scaled up their day-0 procurement, and consequently, Deep Maize did not put on a very strong showing in this round. Fortunately, fourth place was sufficient to advance to the next round.

Table 2 presents results for the top twelve agents after Semifinal 1. Network problems at the competition venue caused difficulties for agents running locally—Jackaroo and HarTAC, in particular.⁵

After the first semifinal closed, the next few hours were filled with a great deal of hustle as the team activated the preemptive strategy that would be played the next day. These hours were also filled with anxiety. We had only intuition about the effect of preemptive strategy on Deep Maize and other agents, but had never had a chance to test it against other competitors. On the other hand, we could hardly wait to see the

⁵ The problems did not affect the majority of agents communicating over the Internet from entrants’ home institutions to the servers in Sweden.

⁶ The score of HarTAC in Semifinal 2 was adversely affected by one game in which it experienced connectivity problems and lost \$364M. Omitting this game would boost their average profit to \$8.46M.

Table 2. Results for twelve agents participating in the second semifinal and final rounds.

Agent	Average Profit (\$M)		
	Semifinal 1	Semifinal 2	Final
RedAgent	12.75 (H1)	25.09 (H1)	11.61
Deep Maize	10.51 (H2)	15.28 (H1)	9.47
TacTex	1.85 (H1)	-15.54 (H2)	5.02
Botticelli	5.69 (H1)	-4.83 (H2)	3.33
PackaTAC	18.31 (H1)	8.70 (H1)	-1.68
WhiteBear	5.26 (H1)	-9.58 (H2)	-3.45
PSUTAC	17.81 (H1)	-1.56 (H1)	—
TAC-o-matic	-1.24 (H2)	-13.50 (H1)	—
Sirish	15.86 (H2)	-20.21 (H2)	—
MinneTAC	13.92 (H2)	-24.98 (H2)	—
UMBCTAC	10.78 (H2)	-29.91 (H2)	—
HarTAC ⁶	2.59 (H2)	-32.95 (H1)	—

“unexpected” dramatic change in **Deep Maize** behavior in the arena with presumably the three best agents (since we did not place very highly in the first round, we would play the top three placing agents from the other heat).

In the morning of 12 August, the **Deep Maize** team stood waiting by the computer screen as the second round of semifinals began. As day 29 rolled around, everyone held their breath, releasing it when the first large delivery of components dropped in. Once we saw distinct manifestations of the preemptive strategy, we began to wonder how other agents would react. Our suspense did not last long: soon after the game’s midpoint, a comment emerged in the TAC game chatroom: “why we can’t get hard disks? How server handle purchase RFQs? is the administrator around!!!?” Apparently, one agent at least was taking for granted that its day-0 requests would be fulfilled.

At the end of S2H1, **Deep Maize** came in second behind the eventual tournament winner, **Red Agent** [Keller et al., 2004], followed closely by **PackaTAC** [Dahlgren, 2003]. These agents, it turned out, were relatively resilient to the preemptive strategy, as they did not excessively rely on day-0 procurement, but adaptively purchased components throughout the game.

Although none had anticipated it explicitly, it turned out that most agents playing in the finals were individually flexible enough to recover from day-0 preemption. By preempting, it seemed that **Deep Maize** had leveled the playing field, but **RedAgent**’s apparent adaptivity in procurement and sales [Keller et al., 2004] earned it the top spot in the competition rankings.

5.4 Analysis

Did **Deep Maize**’s preemption strategy work? We can first examine whether it had its intended direct effect, namely, to reduce the number of components ordered at the very beginning of the game. Table 3 presents, for each tournament round, the number of components ordered on day 1 (based on day-0 RFQs). Each value represents a total

Table 3. Effect of preemption on day 1 component orders and average profits.

	S1H1	S1H2	S2H1	S2H2	Finals
(DM?, P?, N)	-, -, 9	DM, -, 9	DM, P, 8	-, -, 9	DM, P, 16
components	59390	46989	27377	70744	27172
avg profits	2.97	-3.05	7.02	-17.51	4.05

over delivery dates and agents, averaged over the 16 supplier-component pairs. Above the component numbers we indicate whether **Deep Maize** played in that round (DM), whether it employed preemption (P), and the number of games. Note that this data includes one game in S2H1 and two in the finals in which **Deep Maize** failed to preempt due to network problems. It does exclude one anomalous S2H1 game, in which **HarTac** experienced connectivity problems, to wildly distorting effect.

From the table, it is clear that the preemptive day-0 strategy had a large effect. The difference is most dramatic in Semifinal 2, where the heat with **Deep Maize** preempting saw an average of 27377 components committed on day 1, as compared to 70744 in the heat without **Deep Maize**.

The tournament results also indicate that preemption was successful. The fact that **Deep Maize** performed well overall is suggestive, though of course there are many other elements of **Deep Maize** contributing to its behavior. Evidence that the preemptive strategy in particular was helpful can be found in the results from Semifinal 1, where **Deep Maize** did not preempt and ended up in fourth place. This was sufficient for advancing in the tournament, but clearly not as creditable as its second place showing in the finals, among the (presumably) top agents in the field.

We can conclude, then, that preemption helped **Deep Maize**. How did it affect the rest of the field? Table 3 also suggests a positive relation between preemption and profits averaged over *all* agents. Again, the contrast is greatest between S2H1 and S2H2. In the heat without **Deep Maize**, it appears that competition among aggressive agents led to an average *loss* of \$17.51M. With **Deep Maize** preempting in S2H1, profits are a healthy \$7.02M per agent. Preemption was also operative in the finals, and profits there were also positive. That it is preemption and not **Deep Maize** per se is supported by examination of Semifinal 1, in which the heat without our agent appears to be substantially more profitable on average.

Pooling all of these semifinal and final games, we compared average profits for games with and without preemption. Games with preemption averaged \$3.97M in profits, compared to a loss of \$4.02M in games without preemption. Given the small dataset and large variance, this difference is only marginally statistically significant ($p = .09$).

Drawing inferences from tournament results is complicated by the presence of many varying and interacting factors. These include details of participating agents, and random features of environment, in particular the level of demand. To test the influence of demand, we measured the overall demand level for each game, \bar{Q} , defined as the average number of customer RFQs per day. Figure 3 presents a scatterplot of the tournament games, showing \bar{Q} and per-agent profits for each. We distinguish the games with and without preemption, and for each class, fit a line to the points. The linear fit

was quite good for the games with preemption ($R^2 = 0.84$), capturing somewhat less of the variance for the games without ($R^2 = 0.66$).

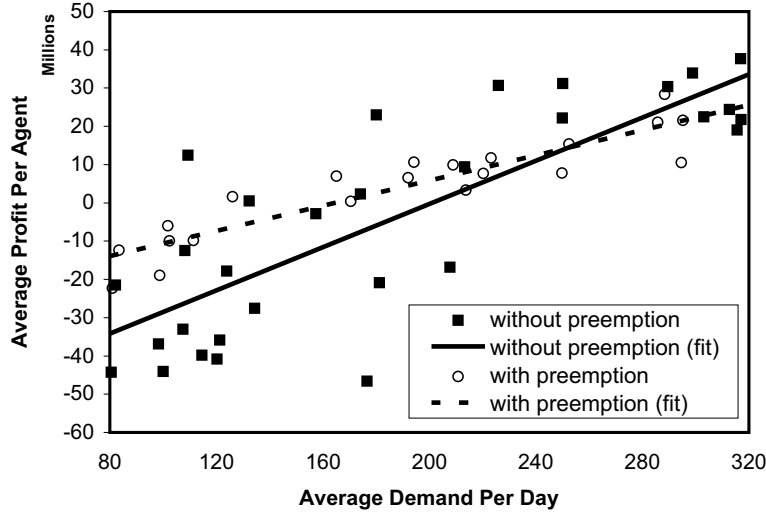


Fig. 3. Profits versus \bar{Q} in TAC-03 tournament games. The lines represent best fits to data from games with and without preemption.

As seen in the figure, with or without preemption, demand clearly exhibits a significant ($p < 10^{-6}$) relation to profits. The relation is attenuated by preemption, and indeed the revealed trend indicates that preemption is beneficial when demand is low, and detrimental in the highest-demand games. This is what we would expect, given that the primary effect of preemption is to inhibit early commitment to large supplies. Given the apparently important influence of demand, we developed a more elaborate mechanism to control for demand in our analysis of tournament games as well as our post-competition experiments.

5.5 Demand Adjustment

Given a sufficient number of random instances, the problem of variance due to stochastic demand would subside, as the sample means for outcomes of interest would converge to their true expectations. However, for TAC/SCM, sample data is quite expensive, as each game instance takes approximately one hour. Therefore, datasets from tournaments and even offline experiments will necessarily reflect only limited sampling from the distribution of demand environments.

To address this issue, we can calibrate a given sample with respect to the known underlying distribution of demand (\bar{Q}). Our approach is closely related to the standard method of variance reduction by conditioning [Ross, 1997, Section 11.6.2]. Given a specification for the expectation of some game statistic y as a function of \bar{Q} , its overall

expectation accounting for demand is given by

$$E[y] = \int_{\bar{Q}} E[y|\bar{Q}] \text{Pr}(\bar{Q}) d\bar{Q}. \quad (4)$$

Although we do not have a closed-form characterization of the density function $\text{Pr}(\bar{Q})$, we do have a specification of the underlying stochastic demand process. From this, we can generate Monte-Carlo samples of demand trajectories over a simulated game.⁷ We then employ a kernel-based density estimation method using Parzen windows [Duda et al., 2000] to approximate the probability density function for \bar{Q} . This distribution is shown in Figure 4. Its mean is 196, with a standard deviation of 77.4. Note that much of the probability is massed at the extremes of demand, with a skew toward the low end. The tendency toward the extremes comes from the combination of trend (τ) momentum and bounding of Q . The skew toward the low end comes from the fact that the trend is multiplicative, so the process tends to transition more rapidly while at the higher levels of demand.




Fig. 4. Probability density for average RFQs per day (\bar{Q}).

Given this distribution, we define *demand-adjusted profit* (DAP) as the expected profit, adjusted for demand. We calculate this by substituting the per-agent profit for y in Eq. (4). Using this formula requires an estimate for profits as a function of \bar{Q} , which we obtain by linear regression from the sample data. The two lines in Figure 3 thus represent our estimates for profits given \bar{Q} for the two sets of TAC-03 tournament

⁷ We could also use historical game data, but simulating Eqs. (1) and (2) is much faster. The 200,000 data points we generated for our density estimate would take 22.8 years of game simulation time to produce.

games. Although the linearity assumption introduces bias, for limited samples this is compensated by the reduction in variance due to adjusting for \bar{Q} .

From the linear model of profits given \bar{Q} , we can obtain a summary comparison of overall profits with and without preemption. For the TAC-03 games without preemption, DAP was $-\$1.41\text{M}$. Preemption increased DAP to $\$5.20\text{M}$. Thus, we find that on average, **Deep Maize**'s preemptive strategy improved not only its own profits, but those of the other agents as well. These results are corroborated by controlled experiments described below.

6 Game-Theoretic Model

Although the tournament results presented above are illuminating, it is difficult to support general conclusions due to the many contributing factors and differences among agents. To isolate the effect of preemption on the key strategic variable (aggressiveness of day-0 procurement), we developed a stylized game-theoretic model, then calibrated it using simulation experiments. Our results are summarized here; see the extended version of this paper [Wellman et al., 2004] for our detailed analysis.

As noted at the outset, TAC/SCM defines a six-player game of incomplete and imperfect information, with an enormous space of available strategies. The game is *symmetric* [Gintis, 2000], in that agents have identical action possibilities, and face the same environmental conditions. In our stylized model, we restrict the agents to two strategies, differing only in their approach to day-0 procurement. Both are implemented as variants of **Deep Maize**. In strategy A (aggressive), the agent requests large quantities of components from every supplier on day 0. In strategy B (baseline), the agent treats day 0 just like any other day, issuing requests according to its usual policy of serving anticipated demand and maintaining a buffer inventory [Kiekintveld et al., 2004a].

To calibrate our models, we ran 30 or more simulated games for each strategy profile (i.e., combination of A and B), with and without the presence of an agent playing the preemptive strategy, P. For each sample, we collected the average profits for the As and Bs, as well as the demand level, \bar{Q} . We derive the demand-adjusted payoff (DAP) for each strategy, using the method described in Section 5.5. Our findings are as follows.

Aggressiveness has a negative effect on total profits. We regressed total DAP for each profile on the number of aggressive agents in that profile. For games without preemption, the linear relationship was quite strong ($p = 0.0018$, $R^2 = 0.88$), with each A in the profile subtracting $\$20.9\text{M}$ from total profits, on average. With a preemptive agent, the effect was insignificant ($p = 0.54$, $R^2 = 0.10$).

Preemption neutralizes aggressive procurement. In non-preemptive profiles, the raw difference in average profits between aggressive and baseline agents was on the order of $\$10\text{M}$, as compared to $\$1\text{M}$ for the preemptive profiles. Moreover, the average variance across agents for non-preemptive profiles was an order of magnitude larger than the average variance for preemptive profiles.

The expected behaviors obtain in equilibrium. Our observations about the game’s propensity to promote aggressive procurement were consistent with the 2003 tournament results, but does this actually constitute rational behavior? From our empirical payoff function, we can derive pure-strategy Nash equilibria, providing one way to answer this question. As seen in Figure 5, the unique pure Nash equilibrium profile without preemption comprises four As and two Bs. A similar analysis with preemption reveals three equilibria, with zero, two, or four As, respectively. The differences are much smaller in this case, and the statistical differences much less significant. This is consistent with our finding above that preemption neutralizes the difference between A and B. Without preemption, a predominance of As is expected.




Fig. 5. DAP payoffs for strategy profiles, without preemption. Arrows indicate for each column, whether an agent in that profile would prefer to stay with that strategy (arrow head), or switch (arrow tail). Solid black arrows denote statistically significant comparisons.

Symmetric equilibria confirm these findings. For the game without preemption, the unique symmetric mixed-strategy equilibrium plays strategy A with probability 0.82. With preemption, there are two equilibria, at probabilities 0.03 and 0.99.

Preemption increases average profits for everybody. Analysis of the mixed-strategy equilibrium of the game without preemption reveals that the expected payoff (equal for A and B, by definition) is a loss of \$9.59M. With preemption, the two equilibria have expected payoffs of \$5.92M and \$7.01M, respectively.

To evaluate the degree to which preemption neutralizes the difference between A and B, we can identify an ϵ^* for each game such that *any* mixed strategy is a symmetric ϵ -Nash equilibrium at $\epsilon = \epsilon^*$. A profile is ϵ -Nash if no agent can improve its payoff by more than ϵ by deviating from its assigned strategy. For games without preemption, ϵ^* is \$10.6M. With preemption, ϵ^* is only \$0.97M. This provides a bound on how much it can matter to make the right choice about aggressiveness, given a symmetric set of other agents.

Preemption obtains in equilibrium. When agents are allowed to choose among all three strategies (A, B, and P), some will choose to preempt. Among the 28 distinct strategy profiles, there are four pure-strategy equilibria, which have 1–3 preemptors. We have also identified a symmetric mixed-strategy equilibrium, in which agents preempt with probability 0.58.

7 Conclusion

The TAC supply-chain game presented automated trading agents (and their designers) with a challenging strategic problem. Embedded within a highly-dimensional stochastic environment was a pivotal strategic decision about initial procurement of components. Our reading of the game rules and observation of the preliminary rounds suggested to us that the entrant field was headed toward a self-destructive, mutually unprofitable equilibrium of chronic oversupply. Our agent, **Deep Maize**, introduced a preemptive strategy designed to neutralize aggressive procurement. It worked. Not only did preemption improve **Deep Maize**'s profitability, it improved profitability for the whole field. Whereas it is perhaps counterintuitive that actions designed to prevent others from achieving their goals actually helps them, strategic analysis explains how that can be the case.

Investigating strategic behavior in the context of a research competition has several distinct advantages. First, the game is designed by someone other than the investigator, avoiding the kinds of bias that often doom research projects to success. Second, the entry pool is uncontrolled, and so we may encounter unanticipated behavior of individual agents and aggregates. Third, the games are complex, avoiding many of the biases following from the need to preserve analytical or computational tractability. Fourth, the environment model is precisely specified and repeatable, thus subject to controlled experimentation. We have exploited all of these features in our study, in the process developing a repertoire of methods for empirical game-theoretic analysis, which we expect to prove useful for a range of problems.

There is no doubt that this form of study also has several limitations, for example in justifying generalizations beyond the particular environment studied. Nevertheless, we believe that the methods developed here provide a useful complement to the kinds of (a priori) stylized modeling most often pursued in game-theoretic analysis, and to the non-strategic analyses typically applied to simulation environments.

Acknowledgments

Thanks to the TAC/SCM organizers and participants. At the University of Michigan, **Deep Maize** was designed and implemented with the additional help of Shih-Fen Cheng, Thede Loder, Kevin O'Malley, and Matthew Rudary. Daniel Reeves assisted our equilibrium analyses. This research was supported in part by NSF grant IIS-0205435.

Bibliography

- Raghu Arunachalam, Joakim Eriksson, Niclas Finne, Sverker Janson, and Norman Sadeh. The TAC supply chain management game. Technical report, Swedish Institute of Computer Science, 2003. Draft Version 0.62.
- Erik Dahlgren. PackaTAC: A conservative trading agent. Master's thesis, Lund University, 2003.
- Richard O. Duda, Peter E. Hart, and David G. Stork. *Pattern Classification*. Wiley-Interscience, second edition, 2000.
- Herbert Gintis. *Game Theory Evolving*. Princeton University Press, 2000.
- Philipp W. Keller, Félix-Olivier Duguay, and Doina Precup. RedAgent: Winner of TAC SCM 2003. *SIGecom Exchanges*, 4(3):1–8, 2004.
- Christopher Kiekintveld, Michael P. Wellman, Satinder Singh, Joshua Estelle, Yevgeniy Vorobeychik, Vishal Soni, and Matthew Rudary. Distributed feedback control for decision making on supply chains. In *Fourteenth International Conference on Automated Planning and Scheduling*, Whistler, BC, 2004a.
- Christopher Kiekintveld, Michael P. Wellman, Satinder Singh, and Vishal Soni. Value-driven procurement in the TAC supply chain game. *SIGecom Exchanges*, 4(3):9–18, 2004b.
- David Pardoe and Peter Stone. TacTex-03: A supply chain management agent. *SIGecom Exchanges*, 4(3):19–28, 2004.
- Sheldon M. Ross. *Introduction to Probability Models*. Academic Press, sixth edition, 1997.
- Norman Sadeh, Raghu Arunachalam, Joakim Eriksson, Niclas Finne, and Sverker Janson. TAC-03: A supply-chain trading competition. *AI Magazine*, 24(1):92–94, 2003.
- Peter Stone. Multiagent competitions and research: Lessons from RoboCup and TAC. In *Sixth RoboCup International Symposium*, Fukuoka, Japan, 2002.
- Michael P. Wellman, Shih-Fen Cheng, Daniel M. Reeves, and Kevin M. Lochner. Trading agents competing: Performance, progress, and market effectiveness. *IEEE Intelligent Systems*, 18(6):48–53, 2003.
- Michael P. Wellman, Joshua Estelle, Satinder Singh, Yevgeniy Vorobeychik, Christopher Kiekintveld, and Vishal Soni. Strategic interactions in a supply chain game. Technical report, University of Michigan, 2004.
- Dongmo Zhang, Kanghua Zhao, Chia-Ming Liang, Gonelur Begum Huq, and Tze-Haw Huang. Strategic trading agents via market modelling. *SIGecom Exchanges*, 4(3):46–55, 2004.