Strategic Procurement in TAC/SCM: An Empirical Game-Theoretic Analysis

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 supply-chain game presented automated trading agents with a challenging strategic problem. Embedded within a high-dimensional 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

The TAC Supply Chain Management (TAC/SCM) scenario [9] defines a complex six-player game with severely incomplete and imperfect information, and high-dimensional strategy spaces.[1] 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. In the full paper [11], we recount the experience as it played out in the competition, and analyze the outcome of this naturalistic experiment. Here we report on 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.

The direct result of this study is validation of the insight behind our particular policy for strategic procurement in the TAC/SCM game. Our experimental analysis verifies that the predominant patterns we observed among agents in the tournament reflect strategically rational behavior for this issue. It also confirms the surprising phenomenon in which a tactic designed to preempt the actions of others actually leads to global welfare improvements. More broadly, we view this exercise as illustrating a general approach by which agent designers can reason through pivotal strategic issues in a principled way, despite computational and analytical intractability of their environments.

1 We assume familiarity with TAC/SCM-03 game rules. For details, see the specification document [1].

2. Deep Maize

The University of Michigan's entry in TAC-03/SCM is an agent called Deep Maize [5]. 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 below.

3. Day-0 Procurement Strategies

Close examination of the game rules reveals that there are major advantages to procuring a large number of components on day 0. Indeed, all agent developers noticed this, and pursued this approach to some degree in the tournament. 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. 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 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, which was designed for high performance in the steady state [5]. Deep Maize devotes considerable effort to 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.

3.1. 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 1. 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 (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.

Figure 1. Deep Maize's preemptive RFQ.

In the situation that our preemptive RFQ gets considered after the supplier has committed 50 days of production to other agents, we will not receive an offer, and our preemption is unsuccessful. For this reason, we also submit backup RFQs of 35000 to be delivered on day 50, and 15000 to be delivered on day 70.

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. This is accomplished by adjusting agent i's selection probability π_i as follows [1, Section 5.1]:

$$\begin{aligned} \text{weight}_i &= \max\left(0.5, \frac{\text{QuantityPurchased}_i}{\text{QuantityRequested}_i}\right), \\ \pi_i &= \frac{\text{weight}_i}{\sum_x \text{weight}_x}. \end{aligned}$$

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.

3.2. TAC-03 Tournament

In a nutshell, our tactic succeeded in the 2003 TAC/SCM tournament. The preemptive strategy had its intended effect of inhibiting day-0 procurement, and enabled Deep Maize to place second in the final round. The most interesting result, however, was that by preempting aggressive procurement, the strategy improved profitability for the entire field of agents, on average. For a detailed analysis of this phenomenon in the official TAC event, see the full version of our report [11].

4. Game-Theoretic Model

4.1. Normal-Form Model Structure

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* [3], in that agents have identical action possibilities, and face the same environmental conditions. In our stylized model, we restrict the agents to three strategies, differing only in their approach to day-0 procurement. Each strategy is implemented as a variant of Deep Maize. By basing the strategies on a particular agent, we clearly cannot capture the diversity of approaches to all aspects of TAC/SCM. Fixing much of the behavior, however, enables our focus on the particular issue of strategic procurement.

In strategy A (aggressive), the agent requests large quantities of components from every supplier on day 0. The specific day-0 RFQs issued correspond to aggressive day-0 policies we observed for actual TAC-03/SCM participants. We encoded four of these as RFQ quantity lists:

1. (4250,5000,5000,2500,1250), based on TacTex [7].
2. (3000,3000,3000,3000,3000), based on UMBCTAC.
3. (4000,3000,8000), based on HarTac.
4. (1672,1672,1672,1672,1672), and double this for CPU components, based on Botticelli [2].

Strategy A randomly selects among these at the beginning of each game instance.

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 [5]. Strategy P (preemptive) is actually Deep Maize as we ran it in the tournament, with preemptive day-0 procurement as described above.

We consider three versions of this game in our analysis. The first is an *unpreempted* six-player game, where agents are restricted to playing A or B. The second is a five-player game, with the sixth place taken up by a fixed agent playing strategy P. We refer to this as the *single-preemptor* game. The third is the full six-player game where agents are allowed to play any of the three strategies A, B, or P.

Since the three strategies incorporate specified policies for conditioning on private information, we represent the game in normal form. By symmetry there are only seven distinct profiles for the unpreempted game, corresponding to the number j of agents playing A, $0 \leq j \leq 6$. There are six distinct profiles for the single-preemptor game, and a total of twenty-eight for the full game (including the thirteen from the more restricted games). Payoffs for each profile represent the expected profits for playing A, B, or P, respectively, given the other agents, with expectation taken over all stochastic elements of the game environment.

4.2. Demand Adjustment

It is apparent to any observer of TAC/SCM games that performance is highly sensitive to underlying demand. This is the case with or without preemption, although the relation is attenuated by preemption. Given that the primary effect of preemption is to inhibit early commitment to large supplies, we would expect that preemption should be beneficial when demand is low, and detrimental in the highest-demand games. This is indeed what we observed in the 2003 tournament. Given the high variability of demand, and its apparently important influence, we developed a more elaborate mechanism to control for demand in our analysis of tournament games as well as our post-competition experiments.

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. (55 minutes of game simulation time, plus a few minutes for pre- and post-game processing) 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 [8, 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

$$E[y] = \int_{\bar{Q}} E[y|\bar{Q}] \Pr(\bar{Q}) d\bar{Q}. \tag{1}$$

Although we do not have a closed-form characterization of the density function $\Pr(\bar{Q})$, 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. We then employ a kernel-based density estimation method using Parzen windows [4] to approximate the probability density function for $\bar{Q}$. This distribution is shown in Figure 2. 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.

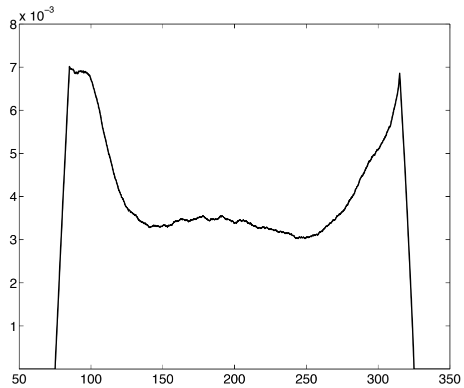

Figure 2. Probability density for average RFQs per day ($\bar{Q}$).

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. (1). Using this formula requires an estimate for profits as a function of $\bar{Q}$, which we obtain by linear regression from the sample data. Although the actual relationship is not linear, the fitted line represents an unbiased estimate of the mean. For limited samples, adjusting for $\bar{Q}$ in this manner indeed produces a substantial reduction in variance.

4.3. Simulation Results

To estimate our game's expected payoff function, we sampled an average of around 30 game instances for each strategy profile—834 in total. For each sample, we collected the average profits for the As, Bs, and Ps, as well as the demand level, $\bar{Q}$. We then used the demand-adjustment method described above to derive DAP for each strategy, which we take as its payoff in that profile.

From this data, we verify that increasing the prevalence of aggressiveness degrades aggregate profits. We show that inserting a single preemptive agent neutralizes the effect of aggressiveness, diminishing the incentive for aggressive procurement, and ameliorating its negative effects. Moreover, the presence of a preemptor tends to improve performance *for all agents* in profiles containing a preponderance of agents playing A. We then study the equilibrium behavior of each of the three versions of the game. From the empirical game models, we derive asymmetric pure-strategy equilibria, as well as symmetric mixed-strategy equilibria, for each of the games.[2] Comparison of the features of equilibrium behavior in the respective games confirms our findings about the effects of strategic preemption.

To test our hypothesis that aggressive strategy has a negative effect on total profits, we regressed total DAP for each profile on the number of aggressive agents in the profile. For profiles without preemption, the linear relationship was quite strong ($p = 0.0018$, $R^2 = 0.88$), with each A in the profile subtracting \$20.9M from total profits, on average.

In the single-preemptor game, the effect of number of aggressive agents on average profits was statistically insignificant, explaining little variance ($p = 0.54$, $R^2 = 0.10$). In unpreempted profiles with four or more aggressive players, agents playing either strategy would benefit substantially (at least \$6.5M on average) from one of the others (either type) switching to play P. Thus, preemption appears to eliminate the detrimental effect that aggressive agents exert on total profits, and for individual profits as well compared to profiles with a predominance of strategy A.

We also confirmed that preemption levels the playing field, as the difference in average profits between aggressive and baseline agents was on the order of \$10M for the unpreempted profiles, as compared to \$1M for the single-preemptor case. Examining the variance across agents in each particular game, we observe that average variance for unpreempted profiles was an order of magnitude larger than that for profiles with preemption. See Table 1.

4.4. Pure Strategy Equilibria

A pure-strategy Nash equilibrium is a strategy profile such that no agent can improve its payoff by changing strategies, assuming all other agents play according to the profile. We identify pure strategy Nash equilibria for both two-strategy games, as well as the full three-strategy game.

4.4.1. Two-Strategy Games In a two-strategy ($\{$A,B$\}$) symmetric game, a profile is defined by the number of As. Profile $0 \leq i \leq N$ is a Nash equilibrium if and only if:

1. the payoff to A in i exceeds the payoff to B in $i - 1$ (or $i = 0$), *and*
2. the payoff to B in i exceeds that to A in $i + 1$ (or $i = N$).

2 It can be shown that for any N-player two-strategy symmetric game, there must exist at least one equilibrium in pure strategies, and there also must exist at least one symmetric equilibrium (pure or mixed) [3].

Profile	Variance	Baseline DAP	Aggress. DAP	Preempt. DAP
AAAAAA	8.37E+15	n/a	–12.26	n/a
AAAAAB	8.29E+15	–9.03	–13.58	n/a
AAAABB	1.14E+16	–10.78	–1.47	n/a
AAABBB	9.70E+15	–10.36	9.73	n/a
AABBBB	6.79E+15	–2.47	19.16	n/a
ABBBBB	2.61E+15	–1.83	13.28	n/a
BBBBBB	1.84E+13	8.17	n/a	n/a
PAAAAA	6.38E+14	n/a	6.86	9.98
PAAAAB	6.36E+14	7.23	8.84	10.29
PAAABB	6.50E+14	3.62	5.15	9.04
PAABBB	7.53E+14	6.06	7.34	10.96
PABBBB	1.18E+15	4.19	5.75	11.41
PBBBBB	1.03E+15	6.06	n/a	13.64
PPAAAA	5.90E+14	n/a	3.67	5.45
PPAAAB	4.24E+14	5.11	4.71	6.69
PPAABB	4.28E+14	4.55	4.70	6.71
PPABBB	4.95E+14	1.74	2.57	4.46
PPBBBB	9.03E+14	4.78	n/a	7.31
PPPAAA	2.49E+14	n/a	7.41	7.30
PPPAAB	1.75E+14	5.76	5.84	6.32
PPPABB	1.99E+14	10.10	10.14	10.08
PPPBBB	3.51E+14	3.76	n/a	4.30
PPPPAA	2.33E+14	n/a	2.26	1.50
PPPPAB	2.13E+14	6.98	7.24	6.16
PPPPBB	2.87E+14	5.77	n/a	5.69
PPPPPA	1.43E+14	n/a	7.74	6.64
PPPPPB	2.04E+14	5.46	n/a	4.39
PPPPPA	1.19E+14	n/a	n/a	4.14

Table 1. Payoffs by strategy profile.

We consider the games defined by DAP payoffs, as well as raw average profits. The full set of DAP payoffs are provided in Table 1. As we ran our simulations, we observed that DAP results anticipated those we would obtain from raw averages after collecting more samples. In a more systematic trial, we found that DAP estimates exhibit lower mean-squared-error compared to sample means, for a range of subsample sizes going well beyond what we could collect for each profile. Given our relatively small datasets, therefore, we have greatest confidence in the DAP results. An advantage of the raw averages is that we have associated variance measures, enabling statistical hypothesis testing.

Let iA denote the profile with no preemption, and i agents playing A (the rest playing B). Whether we define payoffs by DAP or raw averages, the unique pure-strategy Nash equilibrium is 4A. That this is an equilibrium for DAP payoffs can be seen by comparing adjacent columns in the bar chart of Figure 3. 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, as discussed below. Profile 4A is the only one with only in-pointing arrows.

Let PiA denote the profile with a preemptive agent, and i As. In the game with preemption, we find several pure-

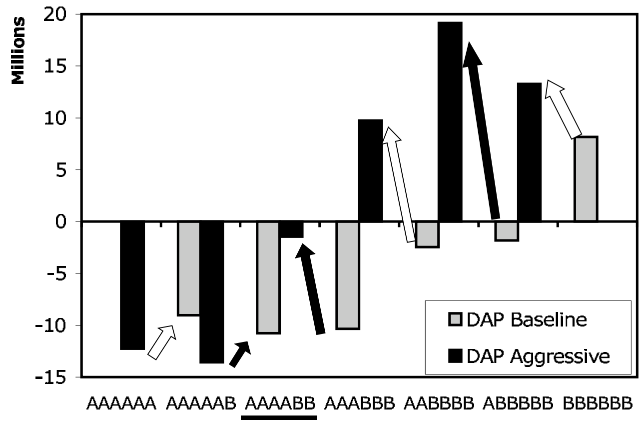

Figure 3.
DAP payoffs, unpreempted profiles.

strategy Nash equilibria. P4A and P2A are equilibria under either DAP or raw profits, and P0A is an equilibrium in the game defined by DAP only. The DAP comparisons are illustrated by Figure 4.

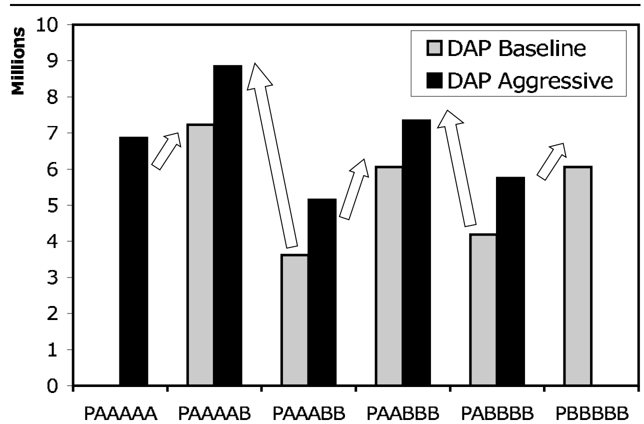

Figure 4. DAP payoffs, with preemption.

To assess the robustness of these equilibria, we conducted statistical tests. For each relevant comparison we performed two-sample t-tests, using average profits, assuming unequal variance. The p-values are presented in Table 2. Whereas some of the comparisons in the unpreempted game indicate significant differences, none of those in preemptive profiles are particularly significant. Thus, the equilibria we found should be considered suspect, or weak equilibria at best. Since payoffs in the preemptive games have much lower variance, if anything we would expect significant differences to show up earlier. This is consistent with our finding above that the preemptive agent neutralizes the difference between strategies A and B. In that respect, identifying an equilibrium is less important in this context.

Comparison	P-value
AAAAAA → AAAAAB	0.6595
AAAAAB → AAAABB	0.0020
AAAABB ← AAABBB	0.0246
AAABBB ← AABBBB	0.5123
AABBBB ← ABBBBB	0.0001
ABBBBB ← BBBBBB	0.2879
PAAAAA → PAAAAB	0.7678
PAAAAB ← PAAABB	0.2294
PAAABB → PAABBB	0.4413
PAABBB ← PABBBB	0.3436
PABBBB ← PBBBBB	0.3845

Table 2. Statistical significance.

Regardless of which equilibrium is played, both A and B agents are clearly better off in the single-preemptor game. In all its equilibria, all agents earn over \$6M profit. In the unpreempted game equilibrium (4A), in contrast, all profits are negative, with the B agents losing over \$10M each.

4.4.2. Full Three-Strategy Game Our analysis of the two-strategy games confirms our hypothesis that introducing a single preemptive agent neutralizes the effect of aggressiveness and moves equilibrium play toward a more profitable space. The success of preemption, however, raises the question whether an incentive to preempt will create a similar mutually destructive competition among preemptors. To check this, we can perform the same kind of equilibrium analysis in the three-player game, where agents are allowed to choose strategy P. The twenty-eight profiles are arrayed in Figure 5, with arrows indicating the transitions between profiles induced by agents switching strategies.

The four pure-strategy Nash equilibria of this game are indicated in bold: PAAAAB, PPBBBB, PPPAAA, and PPPABB. Although the average scores vary across equilibria, in every case the A and B players earn substantial profit, unlike the unpreempted case. Indeed, there exists only one unpreempted profile (2A) from which an A would not deviate, and no unpreempted profiles where playing B is stable.

We also note that as more agents adopt a preemptive strategy, the difference in performance among strategies diminishes. Almost all the comparisons between preemptive profiles are statistically insignificant, as can be seen by the thin arrows in Figure 5. One way to quantify the indifference between strategies given preemption is to consider the ϵ-Nash equilibria. A profile is ϵ-Nash if no agent can improve its payoff by more than ϵ by deviating from its assigned strategy. In Figure 5, we display for each profile the minimum ϵ that would render it an ϵ-Nash equilibrium. For example, although PPPPPA is not an equilibrium, agents can gain at most \$0.34M by deviating from their assigned strategies. Among the 21 preemptive profiles, 17 of them are ϵ-Nash equilibria at an ϵ of \$5.38M or less. In contrast, none of the unpreempted profiles are ϵ-equilibria at that level.

4.5. Mixed Strategy Equilibria

Although the pure-strategy equilibria are interesting, we might consider *symmetric equilibria* more natural, given the symmetry of the game and its lack of identifying roles [6]. In order to identify a symmetric equilibrium, we need in general to consider mixed strategies.

4.5.1. Two-Strategy Games Let N be the total number of strategies in the profile (in our context, $N = 6$ without preemption, and $N = 5$ when we include a single fixed preemptive agent). Define DAP(X, j) as the DAP of strategy X (A or B) when j agents out of N play strategy A. If k agents each independently choose whether to play A with probability α (henceforth, "play α"), then the probability that exactly i will choose A is given by

$$\Pr(\alpha, i, k) = \binom{k}{i} \alpha^i (1-\alpha)^{k-i}.$$

Let $V(A, \alpha)$ denote the DAP of an agent playing A when the remaining agents play α:

$$V(A, \alpha) = \sum_{i=0}^{N-1} \Pr(\alpha, i, N-1)\text{DAP}(A, i+1).$$

Similarly, we define DAP values for playing B or α, respectively, in the setting where others play α:

$$\begin{aligned} V(B, \alpha) &= \sum_{i=0}^{N-1} \Pr(\alpha, i, N-1)\text{DAP}(B, i), \\ V(\alpha, \alpha) &= \alpha V(A, \alpha) + (1-\alpha) V(B, \alpha). \end{aligned}$$

We plot these values of playing A, B, or α in response to α, for the two games, in Figure 6. A necessary and sufficient condition for a symmetric mixed-strategy equilibrium is

$$V(A, \alpha) = V(B, \alpha).$$

Therefore, we can identify such equilibria by the points in these figures where the curves intersect. For the game without preemption, we have a single symmetric mixed-strategy equilibrium, at $\alpha = 0.82$. When the preemptive agent is present, we find two symmetric mixed-strategy equilibria: $\alpha = 0.03$ and $\alpha = 0.99$.

The expected payoff for the equilibrium strategy (necessarily equal for A and B) of the game without preemption is

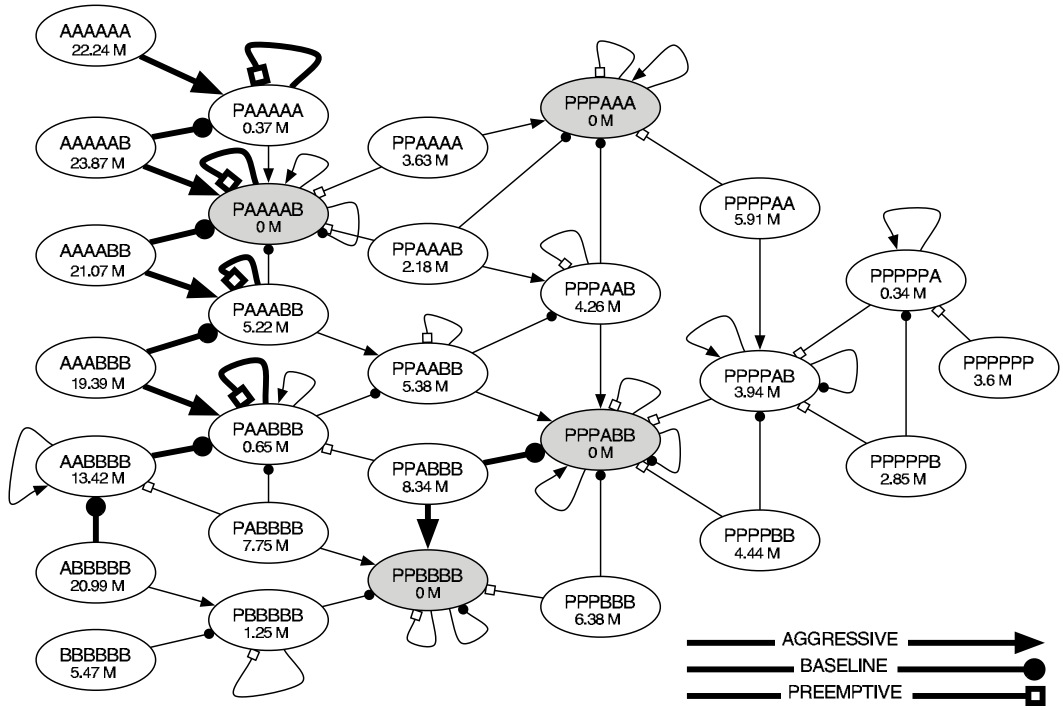

Figure 5.
Profiles for the full three-strategy game, with arrows indicating a desire by the associated agent to change its strategy. Bold arrows indicate statistically significant comparisons. Values specified for each profile represent the minimum ϵ such that the profile constitutes an ϵ-Nash equilibrium.

a *loss* of \$9.59M. With a single preemptor, the two equilibria have expected payoffs of \$5.92M and \$7.01M, respectively. The preemptive agent itself also does well, earning profits of \$13.3M and \$9.99M in the respective equilibria.

Although we have no direct way to perform a statistical hypothesis test using demand-adjusted values, a conservative option is to compare the mean DAP scores using the variance of the raw averages. In this instance, DAP for the two preemptive equilibria exceed that of the non-preemptive equilibrium at p-values less than 0.0001.

Inspection of Figure 6 confirms our prior finding that preemption reduces the difference between A and B strategies. One way to quantify this is to identify an ϵ^* for each game such that *any* mixed strategy is a symmetric ϵ-Nash equilibrium at $\epsilon = \epsilon^*$. In our context, ϵ^* is therefore the maximum payoff difference between playing the best-response strategy, and playing α:

$$\epsilon^* = \max_{\alpha} \left(\max_{X \in \{A,B\}} V(X, \alpha) - V(\alpha, \alpha) \right).$$

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.

4.5.2. Full Three-Strategy Game We were also able to derive a symmetric mixed-strategy equilibrium for the full three-strategy game, using replicator dynamics [10]. In equilibrium, agents play A with probability 0.23, B with probability 0.19, and P with 0.58. The expected payoff for this mixed strategy is \$5.78M. This is not quite as good as the environment allowing only a single preemptor, but of course much better than the unpreempted situation.

5. 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

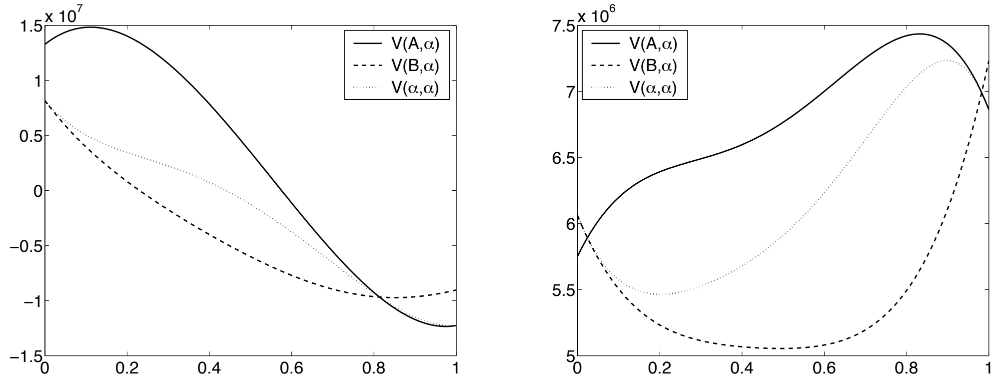

Figure 6.
Response to mixed strategy α, for the unpreempted (left) and single-preemptor (right) games. Note that the payoff scale is an order of magnitude wider in the left graph.

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 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.

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.

References

[1] R. Arunachalam, J. Eriksson, N. Finne, S. Janson, and N. Sadeh. The TAC supply chain management game. Technical report, Swedish Institute of Computer Science, 2003. Draft Version 0.62.

[2] M. Benisch, A. Greenwald, I. Grypari, R. Lederman, V. Naroditskiy, and M. Tschantz. Botticelli: A supply chain management agent designed to optimize under uncertainty. *SIGecom Exchanges*, 4(3):29–37, 2004.

[3] S.-F. Cheng, D. M. Reeves, Y. Vorobeychik, and M. P. Wellman. Notes on equilibria in symmetric games. In *AAMAS-04 Workshop on Game-Theoretic and Decision-Theoretic Agents*, New York, 2004.

[4] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. Wiley-Interscience, second edition, 2000.

[5] C. Kiekintveld, M. P. 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*, Whistler, BC, 2004.

[6] D. M. Kreps. *Game Theory and Economic Modelling*. Oxford University Press, 1990.

[7] D. Pardoe and P. Stone. TacTex-03: A supply chain management agent. *SIGecom Exchanges*, 4(3):19–28, 2004.

[8] S. M. Ross. *Introduction to Probability Models*. Academic Press, sixth edition, 1997.

[9] N. Sadeh, R. Arunachalam, J. Eriksson, N. Finne, and S. Janson. TAC-03: A supply-chain trading competition. *AI Magazine*, 24(1):92–94, 2003.

[10] P. Schuster and K. Sigmund. Replicator dynamics. *Journal of Theoretical Biology*, 100:533–538, 1983.

[11] M. P. Wellman, J. Estelle, S. Singh, Y. Vorobeychik, C. Kiekintveld, and V. Soni. Strategic interactions in a supply chain game. Technical report, University of Michigan, 2004.