Trading Agents Competing: Performance, Progress, and Market Effectiveness*

Michael P. Wellman, Shih-Fen Cheng,
Daniel M. Reeves, and Kevin M. Lochner

University of Michigan
Artificial Intelligence Laboratory
Ann Arbor, MI 48109-2110 USA
{wellman,chengsf,dreeves,klochner}@umich.edu

September 5, 2003

Abstract

Since the year 2000, the annual trading agent competition has provided a forum for designers to evaluate programmed trading techniques in a challenging market scenario in competition with other design groups. After three years of apparent progress, we attempt to evaluate the trading competence of competition participants, in the 2002 tournament and over time. Although absolute measure of individual performance is difficult to assess, relative measures, and measures of the market performance overall are more amenable to direct analysis. We quantify the effectiveness of the TAC travel market in terms of allocative efficiency, finding improvement within and between tournaments. By comparison with alternative allocation benchmarks, we can calibrate this efficiency, and identify opportunities for further gain from trade.

1 Introduction

One of the primary motivations for producing the original Trading Agent Competition (TAC) was to spur research on a common problem, thus enabling researchers to compare techniques and build on each others' ideas [Wellman and Wurman, 1999]. Working on a shared problem coordinates attention on particular issues (among the many of interest in the trading domain), and facilitates communication of methods and results by fixing a set of assumptions and other environment settings. An inspiring example for the TAC developers was the 1990 Santa Fe Double Auction Tournament

*This is a revised and extended version of a paper appearing in the IJCAI'03 workshop on Trading Agent Design and Analysis (TADA).

[Rust et al., 1993], which produced many valuable insights about the Continuous Double Auction (CDA) mechanism, and strategies for bidding in CDAs. Even though the CDA had been well studied before, the focused effort and attention catalyzed by the competition led to substantial cross-fertilization of ideas and perspectives. The book generated from that event [Friedman and Rust, 1993] remains a seminal reference for CDAs.

A multi-year event like TAC offers the further prospect of learning from shared experience. A goal of repeating the competition was to observe the progress of trading agents, in effect accelerating the evolution of an adapted population of traders. Now that we have three years of experience in the TAC series, it is appropriate to examine whether this objective has been fulfilled. Moreover, given all the effort devoted to TAC participation, it is incumbent on us to exploit the TAC data for what they are worth, and see what conclusions we might draw about the efficacy of trading agents and particular ideas about trading strategy.

By one measure, TAC has certainly succeeded in spurring research. Over a dozen publications reporting on the competitions, specific agents, techniques employed, and analyses have appeared to date in archival journals, refereed conferences, and magazines.[1] The TAC "literature" thus represents an uncommonly rich corpus of documentation on trading strategy and behavior for a particular complex environment. Many of the accounts include specific analyses or experiments involving agents from multiple developers, or variants on a particular agent inspired by techniques reportedly employed by others [Greenwald, 2003b, Stone et al., 2002, Wellman et al., 2002]. Such efforts augment the anecdotal evidence from entrants that each successive year the lessons and approaches presented previously are incorporated in new and improved agent designs.

In this paper, we present some data bearing on the assessment of performance and progress of trading agents, as reflected in the TAC series to date. We employ data from the actual TAC tournaments, as well as some post-competition experimentation. Our analysis is based almost entirely on outcomes (profits and allocations), with very little direct accounting for specific agent techniques. Though we offer some conclusions, our investigation also raises further questions. A definitive assessment of agent competence must await further studies, perhaps based on succeeding years of TAC experience.

2 TAC Classic Game

[*Note to editors: This section might best be presented as a sidebar.*]

The "classic" TAC game[2] presents a travel-shopping task, where traders assemble flights, hotels, and entertainment into trips for a set of eight probabilistically generated clients. Clients are described by their preferred arrival and departure days (pa and pd), the premium (*hp*) they are willing to pay to stay at the "Towers" (T) hotel rather than "Shanties" (S), and their respective values (e_1, e_2, e_3) for three different types of entertainment events. The agents' objective is to maximize the value of trips for their

[1] `http://auction2.eecs.umich.edu/researchreport.html`

[2] So called to distinguish the original TAC market game from newer variants, such as the supply chain game introduced for TAC-03 [Sadeh et al., 2003].

clients, net of expenditures in the markets for travel goods. The three categories of goods are exchanged through distinct market mechanisms.

Flights. A feasible trip includes air transportation both ways, comprising an inflight day i and outflight day j, $1 \leq i < j \leq 5$. Flights in and out each day are sold independently, at prices determined by a stochastic process. The initial price for each flight is $\sim U[250, 400]$, and follows a random walk thereafter with an increasingly upward bias.

Hotels. Feasible trips must also include a room in one of the two hotels for each night of the client's stay. There are 16 rooms available in each hotel each night, and these are sold through ascending 16th-price auctions. Agents submit bids for various quantities, specifying the price offered for each additional unit. When the auction closes, the units are allocated to the 16 highest offers, with all bidders paying the price of the lowest winning offer. Each minute, the hotel auctions issue *quotes*, indicating the 16th- (*ASK*) and 17th-highest (*BID*) prices among the currently active unit offers [Wurman et al., 1998]. Starting at minute four, one of the hotel auctions is selected at random to close, with the others remaining active and open for bids.

Hotel bidders are also subject to a "beat-the-quote" rule [Wurman et al., 2001], requiring that any new bid offer to purchase at least one unit at a price of $ASK + 1$, and at least as many units at $ASK + 1$ as the agent was previously winning at *ASK*.

Entertainment. Agents receive an initial random allocation of entertainment tickets (indexed by type and day), which they may allocate to their own clients or sell to other agents through continuous double auctions [Friedman and Rust, 1993]. The entertainment auctions issue *BID* and *ASK* quotes representing the highest outstanding buy and lowest sell offer, respectively, and remain open for buying and selling throughout the 12-minute game duration. A client may sell tickets that it does not own, but must pay a penalty of 200 per ticket for any "short sales" not covered by the end of the game.

A feasible client trip r is defined by an inflight day in_r, outflight day out_r, hotel type (H_r, which is 1 if T and 0 if S), and entertainment types (E_r, a subset of $\{1, 2, 3\}$). The value of this trip is given by

$$v(r) = 1000 - 100(|pa - in_r| + |pd - out_r|) + hp \cdot H_r + \sum_{i \in E_r} e_i. \qquad (1)$$

At the end of a game instance, the TAC server calculates the optimal allocation of trips to clients for each agent, given final holdings of flights, hotels, and entertainment. The agent's game score is its total client trip utility, minus net expenditures in the TAC auctions.

3 TAC-02 Tournament Results

Average scores for the sixteen agents that played in the final and semifinal rounds are posted in Table 1. The third column represents an adjustment to final-round scores to correct for favorability of client preference assignments, calculated according to the formula we developed for TAC-01 [Wellman et al., 2003a]. See `http://www.sics.se/tac` for a list of participant affiliations and team leaders, as well as results from preliminary rounds. Complete game logs are available, as they are for the previous TAC events. Brief agent descriptions have been collected by Greenwald [2003a].

Agent	Affiliation	Scores		
		Semifinals	Finals	Adj Finals
ATTac	AT&T Research (et al.)	H1: 3137	—	—
cuhk	Chinese U Hong Kong	H2: 3266	3069	3045
kavayaH	Oracle India	H1: 3200	3099	3039
livingagents[3]	Living Systems AG	H1: 3310	3181	3161
PackaTAC	N Carolina State U	H2: 3250	—	—
PainInNEC	NEC Research (et al.)	H1: 2193	—	—
RoxyBot	Brown U	H2: 3160	—	—
sics	Swedish Inst Comp Sci	H2: 3146	—	—
SouthamptonTAC	U Southampton	H1: 3397	3385	3337
Thalis	U Essex	H2: 3199	3246	3210
tniTac	Poli Bucharest	H1: 3108	—	—
TOMAhack	U Toronto	H2: 2843	—	—
tvad	Technion	H1: 2724	—	—
UMBCTAC	U Maryland Baltimore Cty	H1: 3208	3236	3291
Walverine	U Michigan	H2: 3287	3210	3277
whitebear	Cornell U	H2: 3324	3413	3479

Table 1: TAC-02 seeded agents, and their average scores during the semifinals (14 games) and finals (32 games).

Although we agree with those who have cautioned against focusing excessively on ranked results in the context of research competitions [Stone, 2002], tournament results provide an important source of information about agent quality. Agents are presumed to act to maximize expected score, and so all else equal, an increase in score reflects an improved agent. If several agents improve, however, this may or may not lead to higher total scores for those agents. Whereas some kinds of improvement unambiguously increase total agent surplus (e.g., fewer wasted flights, better allocation of entertainment), others may reduce the value retained by agents (e.g., smarter agents may be more effective at competing away the consumer surplus) or even the total system surplus (e.g., deadweight loss due to strategic behavior).

One way to measure progress over time is to track benchmark levels of performance by keeping some agents constant. For example, in the CADE ATP (automated theorem proving) series of competitions [Sutcliffe, 2001], the best systems from a given year typically enter unchanged in the next year's event (along with improved versions, of course). This provides a direct measure, in comparable terms, of the relative performance across years. In a game setting, where other agents are part of the environment, it is not strictly fair to judge an agent with respect to a different field. Nevertheless, it

[3] The score of livingagents was adversely affected by missing two games. Discounting these would have led to an average score of 3393.

can be quite instructive to observe the implications of such transplants.

The 2001 tournament [Wellman et al., 2003a] included two calibrating agents in the seeding round. ATTac-2000 [Stone et al., 2001] represented the highest-scoring agent from the TAC-00 finals. To account for the rule changes between TAC-00 and TAC-01, ATTac-2000 was modified with a one-line edit causing it to place all of its bids before the first hotel closure as opposed to during the last minute of the game. We also included dummy_buyer, the agent provided by the Michigan TAC team in 2001 to play in test games that do not have a full slate of agents. Whereas most of the other agents' behaviors were modified between (and during) the qualifying and seeding round, the dummy was left unchanged. Not surprisingly, we observed substantial deterioration in the dummy's standing as the preliminary rounds progressed.

The 2002 tournament did not explicitly insert calibrating agents, but the fact that twelve of the participating teams from 2001 also entered agents in TAC-02 provides some natural calibration. In particular, the two top-scoring agents in TAC-01, livingagents [Fritschi and Dorer, 2002] and ATTac-2001, participated with essentially unchanged agents in TAC-02. As shown in the table, livingagents did quite well, assuming we ignore the bug that caused it to skip two games. ATTac was top scorer in the TAC-02 seeding rounds, but then was eliminated in the semifinals. One possible explanation is simply that the agent experienced technical difficulties due to a change of computational environments between the seeding and semifinal rounds.[4] Another more substantive possibility is that prices during the preliminary rounds in 2002 (which ATTac uses as training data) were not sufficiently representative of the final rounds. We suspect that the decrease in relative performance also reflects a general increase in competence of the other agents in the field. Interestingly, it may well be that because livingagents in some respects benefits by playing along with effective and adaptive agents [Wellman et al., 2003a], it may be more robust with respect to improvements in the rest of the field.

The two top-scoring agents in TAC-02, whitebear and SouthamptonTAC [He and Jennings, 2002], also contended in TAC-01. These agents reportedly evolved from their 2001 designs, improved through adopting refined classifications of game environments [He and Jennings, 2003], and through extensive experimentation and parameter tuning [Vetsikas and Selman, 2003].

Of the eight other repeat entries:

- Three represent complete reimplementations of the corresponding TAC-01 entries, by essentially different agent designers (Thalis, sics, UMBCTAC).
- Two represent significant redesigns by the same essential designers (RoxyBot, cuhk).
- Three represent incremental or unknown changes by the same or related designers (PainInNEC, BigRed, harami).

[4] Peter Stone, personal communication.

4 TAC Market Efficiency

Another gauge of agent effectiveness is how well they allocate travel goods, *in the aggregate*, through their market interactions. This is an indirect measure, at best, since the objective of each agent is to maximize its own surplus, not that of the overall system—comprising all agents plus the TAC seller. Nevertheless, such a *social welfare* analysis can provide a benchmark, and shed light on the allocation of resources through an economy of interacting software agents.

4.1 Market Efficiency in the TAC-02 Tournament

We can measure aggregate effectiveness by comparing actual TAC market allocations with ideal global allocations, calculated centrally assuming knowledge of all client preference information. Consider the total group of 64 clients, and the set of available resources: 16 hotel rooms of each type per day, plus 8 entertainment tickets of each type per day. The global optimizer calculates the allocation of resources maximizing total client utility, net of expenditures on flights assuming they are available at their initial prices. We take initial prices to be the relevant inherent cost (exogenously determined, independent of TAC agent demand) of flights, treating the expected stochastic increase of flights during the game as a cost of decision delay that would be avoided by the idealized optimizer. Note that the global optimization completely neglects hotel and entertainment prices, as these are endogenous to the TAC market. Monetary transfers affect the distribution of surplus across TAC buyers and sellers, but not the total amount. We formulate the optimization problem as an integer linear program, and solve it using CPLEX.

The average idealized net utility, per client, in the various rounds of the TAC tournament as determined by global optimization is reported under the heading "Global" in Table 2. Average net utility achieved in the actual TAC games (also neglecting hotel and entertainment expenditures, but counting actual payments for flights) is reported under "TAC Market".

Round	Games	Global	TAC Market	TAC (%)
Qualify	390	618	415	67.0
Seeding	1045	618	470	75.7
Semi-Final	28	608	534	87.7
Final	32	609	542	89.1

Table 2: The efficiency of the TAC market compared to the global optimum.

As seen in the table, we found that the TAC market achieved 89% of the optimal value, on average, over the 32 games of the TAC-02 finals. There was a steady improvement from the qualifying round (67% optimal), seeding round (76%), and semifinals (88%). All of these differences are significant ($p < 0.01$), except the small increment from semifinals to finals.[5]

[5]Henceforth, all assertions of statistical significance are with respect to the 0.01 level, unless otherwise

It is difficult to assess this effectiveness in absolute terms, so we provide a couple of benchmarks for comparison.

1. Uniform hotel and entertainment. We distribute the hotel rooms and entertainment evenly across the eight agents, then optimize each agent's allocation to clients. This approach yields 95.2% of the globally optimal value on average. (Allocation values were significantly better than the market in every round.)

2. Uniform hotel, endowed entertainment. The relative average value drops to 85.4% if we distribute only the hotels, leaving agents with their original endowment of entertainment. (This value represents a significant improvement to the market in qualifying and seeding rounds, with the market significantly better in the finals.)

It is perhaps surprising that simply dividing the goods uniformly achieves such a high fraction of the available surplus—better than the market if entertainment is included in the distribution. One reason that the uniform distribution is relatively so effective is that the agents are *ex ante* symmetric, with independent and identically distributed clients. Potential gains from trade are thus not so great for hotels. Second, a direct allocation avoids the significant obstacles posed to agents pursuing their allotments individually through the market. Agents face substantial risk (price uncertainty, exposure due to complementarities, unknown hotel closing patterns), and this necessarily entails some loss in expected allocation quality. For example, the set of available hotels is sufficient to obtain trips for all clients (albeit shortened from desired lengths), and given a definite allocation the agent can optimize for its clients accordingly. With uncertainty, the agents may plan for longer trips than are jointly feasible, and thus wind up wasting flights, hoarding hotel rooms (to hedge), or resorting to suboptimal fallback trip options. A preliminary analysis indicates that this uncertainty is indeed a significant cause of misallocation in TAC play. Developing a precise quantititative characterization of the sources of loss is a subject for future work.

4.2 Comparisons

Given that agent programmers are actively debugging and developing their agents during the preliminary rounds, it seems fair to assume that agent competence improves in succeeding rounds of the tournament. The selection of best performers for the semi-finals and then the finals naturally amplifies this effect. Thus, the progressive improvement in market efficiency observed in Table 2 coincides with individual agent progress. We performed the same global optimization analysis for the TAC-01 finals (24 games), and found a market efficiency of 85.7%. Though better than the TAC-02 qualifying and seeding rounds, the TAC-01 finalists did not allocate resources as well as TAC-02 finalists ($p = 0.024$), or even the TAC-02 semi-finalists ($p = 0.097$). This indirectly (to the extent that overall market efficiency aligns with individual agent success) confirms our conjecture that the 2002 agents were on the whole more competent than their predecessors.

specified.

It is also potentially interesting to compare market effectiveness across different configurations of agents. Our initial explorations employ versions of Michigan's TAC-02 entry, Walverine [Cheng et al., 2003]. In 173 games with Walverine playing all eight agent slots, the market achieved 88.5% efficiency, a result not statistically distinguishable from that of the actual pool of TAC-02 finalists.

Of course, Walverine, like the others, aims to promote its own profit, not overall efficiency per se. We discuss the general tradeoff between individual and social welfare in Section 5 below, including an experiment based on variations of Walverine.

In further work, we intend to evaluate additional agent configurations, including further Walverine variants as well as agents developed by other groups.[6] One interesting question (at least as another benchmark) is what level of overall efficiency can be attained by agents actually designed with this objective in mind. It should be possible to achieve at least the level of our uniform-hotel-and-entertainment benchmark (95.2%), since it is straightforward to construct bidding policies that implement a uniform distribution of all goods.

4.3 Entertainment Trading Efficiency

One component of market effectiveness amenable to separate analysis is entertainment trading. Entertainment goods are initially distributed as endowments to the agents, who exchange among themselves through CDAs to reach a final allocation. Although the value of entertainment to agents depends on their choice of trip dates, it is possible to characterize with reasonable accuracy the gains from trade specifically attributable to the entertainment component of the TAC market.

To measure entertainment trading efficiency, we simply compare the aggregate "fun bonus" component of trip utility in the globally optimal allocation, with that attained in the actual TAC market. Efficiency percentages for the various game sets are presented in Table 3, which also repeats the overall efficiency numbers for convenient reference.

One interesting result is that the entertainment performance in the TAC-02 finals was virtually the same as in the TAC-01 finals, despite the significant improvement in market performance overall. This suggests that the strategic progress was focused on hotel and flight strategies, which certainly agrees with the entrants' reports of their concentrations of effort.

The data on entertainment performance can also help to calibrate estimates of potential gains from agent improvements. We reported above on two benchmarks based on uniform hotel distributions, one with no entertainment trading and one with uniform entertainment allocation. These define a potential gain from trading to uniformity of approximately 60.5 per client (484 per agent) given the fixed uniform allocation of hotels. In our process of training Walverine's entertainment strategy [Cheng et al., 2003], we observed a difference of roughly 478 on average between a policy of not trading entertainment (average fun bonus 1019), and that of a representative hand-coded strategy (that of livingagents, average fun bonus 1497). In contrast, the average fun bonus

[6] Performing our analysis requires only data describing client preferences and final allocations. We welcome any game logs other researchers would be willing to submit for our efficiency analysis, covering whatever profiles of their own agent (and variants) they may have investigated. We also invite participation in further experiments involving mixtures of agents from multiple groups.

Round	Games	TAC (%)	Entertainment (%)
02 Qualify	390	67.0	71.1
02 Seeding	1045	75.7	79.0
02 Semi-Final	28	87.7	83.0
02 Final	32	89.1	85.3
01 Final	24	85.7	85.5
all Walverine	173	88.5	85.4
non-shading Walverine	55	89.4	85.1
shading equilibrium	470	89.2	85.6

Table 3: The efficiency of the TAC market compared to the global optimum—overall and specifically with respect to entertainment.

in the global optimal allocation runs around 1677. The fact that the observed gain from trade in training[7] approaches the gain from uniform entertainment distribution in the uniform-hotel case suggests that the remaining benefit from entertainment trading is equal to the difference between uniform and optimal. We evaluated this by calculating a third benchmark, based on global optimization of entertainment subject to a uniform allocation of hotels to agents. The result is 92.7 per agent greater than the value with uniform entertainment allocation. Although all of these measurements are conditioned on the uniform hotel allocations, we believe that the *relative entertainment values* are likely to be robust to any reasonable fixed hotel allocation.

5 Individual versus Social Welfare

As emphasized above, measuring market efficiency is not at all the same thing as measuring the effectiveness of individual agents. Some factors in agent competence, such as choosing optimal trips and avoiding wasted flights and hotels, correspond directly to improvements in overall efficiency. Other factors, such as reducing demand for hotels so as to capture a larger fraction of surplus from sellers, will tend to reward the individual agent at the expense of social welfare. Understanding the extent that this tradeoff operates in the TAC game will shed light on the relationship between market efficiency measures and the assessment of agent performance.

We can study this question by identifying particular strategy components that would be *expected* to detract from overall welfare, and measuring the welfare loss between the socially and individually optimal policies. We consider behavior individually optimal if it is part of a Bayes-Nash equilibrium (BNE) for the game. Although deriving BNE for TAC is not remotely tractable, we have found that it is feasible to characterize restricted TAC BNE with respect to highly constrained subsets of allowable strategies.

Specifically, we investigate the strategic behavior of hotel bid *shading*, where the

[7] The fun bonus realized in the TAC-02 finals turned out to be somewhat lower than expected based on training observations, for both Walverine and livingagents, and apparently the rest of the field as well except for whitebear [Cheng et al., 2003].

agent offers to buy rooms at prices lower than its marginal values for the units. Walverine determines bid prices as part of a decision-theoretic optimization of expected surplus, which takes into account the probability of not obtaining a good even though its price is below the agent's valuation [Cheng et al., 2003].[8] Though it benefits the individual agent by design, such shading runs counter to the goal of market efficiency, as the market does not generally have available faithful signals of the relative value of goods to the various agents.[9]

To evaluate this effect, we defined a variant "non-shading" strategy implemented by Walverine with its optimal bidding procedure (i.e., shading) turned off. Agents in this version bid their true marginal values for every hotel room. Our hypothesis was that this would improve social welfare, but sacrifice individual profits. We in fact found, in a trial of 55 games with all non-shading Walverines, the market achieved 89.4% efficiency which is better ($p = 0.06$) than all Walverines. The actual effect of varying strategy, however, will in general depend on the strategies of other agents. Therefore, we employed an evolutionary search approach to find a restricted BNE for this game [Wellman et al., 2003b].

We begin by running a series of TAC games, distributed over each of the nine possible profiles of shading and non-shading agents. Averaging the scores for each profile (adjusted for client preference favorability) yields an expected payoff for each strategy in each profile. We can then calculate a BNE for the restricted game using any available technique. Applying replicator dynamics [Taylor and Jonker, 1978],[10] we identify a symmetric mixed strategy BNE where agents shade with probability 0.11, and refrain from shading with probability 0.89. Note that this result applies specifically to the version of shading employed in Walverine, alternative shading policies incorporated in other agent strategies may produce varying outcomes.

The predominance of truthful bidding in equilibrium demonstrates that this policy has advantages in a population with substantial shaders. It is not dominant, however, which means that as the population approaches all non-shaders, there is benefit to shading. In equilibrium, the payoffs to shading and non-shading agents are the same, each a best response to the given mixture. The average client-adjusted payoffs for all shading, all non-shading, and the BNE mixed strategy are 3339, 3155, and 3209, respectively. The corresponding market efficiencies are 88.5%, 89.4%, and 89.2%. Playing the equilibrium strategy results in an average payoff gain of 53 per agent per game but a loss of 46 in social welfare compared to all non-shaders.

[8] We refer to this as "optimal bidding", although the optimization embodies several simplifying assumptions. In fact we suspect that there is much room for improvement, and are actively working on this component for future versions of Walverine.

[9] If all agents shaded proportionally, then the relative offer prices would still provide the relevant information to the market. In general, however, the price reductions do not cancel out in this way, as the bidder's optimization includes many agent-specific contextual factors.

[10] Gambit was unable to find a symmetric BNE after some hours of cpu time, finding only the asymmetric equilibrium in which five agents shade and three do not.

6 Comment

The foregoing analysis provides some evidence for competence and progress in TAC traders. Since our measures are all indirect (e.g., measuring market efficiency rather than absolute agent performance), however, definitive conclusions are not justified. More compelling demonstrations of progress and competence might be based on further calibration studies, systematic search in strategy spaces, and attribution of allocation suboptimality among its many possible causes (e.g., agent suboptimality, and inherent risk—including the cost of its rational management). Further benchmarks, capturing less ideal conditions, may prove useful in this regard.

Another natural question not at all addressed by this work is how well TAC agents fare compared to what human traders could do? We are aware of no evidence that humans would be particularly adept at a TAC-like trading task. One of the few studies comparing human and computer traders (in an abstract CDA scenario) did not reflect very favorably on the humans [Das et al., 2001].

Acknowledgments

We benefited much from discussions with the rest of the Michigan TAC Team, and from comments by the anonymous reviewers. This research was supported in part by NSF grants IIS-9988715 and IIS-0205435.

References

Shih-Fen Cheng, Evan Leung, Kevin M. Lochner, Kevin O'Malley, Daniel M. Reeves, and Michael P. Wellman. Walverine: A Walrasian trading agent. In *Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Melbourne, 2003.

Rajarshi Das, James E. Hanson, Jeffrey O. Kephart, and Gerald Tesauro. Agent-human interactions in the continuous double auction. In *Seventeenth International Joint Conference on Artificial Intelligence*, pages 1169–1176, Seattle, WA, 2001.

Daniel Friedman and John Rust, editors. *The Double Auction Market*. Addison-Wesley, 1993.

Clemens Fritschi and Klaus Dorer. Agent-oriented software engineering for successful TAC participation. In *First International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Bologna, 2002.

Amy Greenwald. The 2002 trading agent competition: An overview of agent strategies. *AI Magazine*, 24(1):83–91, 2003a.

Amy Greenwald. Bidding under uncertainty in simultaneous auctions. In *IJCAI-03 Workshop on Trading Agent Design and Analysis*, Acapulco, 2003b.

Minghua He and Nicholas R. Jennings. SouthamptonTAC: Designing a successful trading agent. In *Fifteenth European Conference on Artificial Intelligence*, pages 8–12, Lyon, 2002.

Minghua He and Nicholas R. Jennings. SouthamptonTAC: An adaptive autonomous trading agent. *ACM Transactions on Internet Technology*, 3:218–235, 2003.

John Rust, John Miller, and Richard Palmer. Behavior of trading automata in a computerized double auction market. In Friedman and Rust [1993], pages 155–198.

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.

Peter Stone, Michael L. Littman, Satinder Singh, and Michael Kearns. ATTac-2000: An adaptive autonomous bidding agent. *Journal of Artificial Intelligence Research*, 15:189–206, 2001.

Peter Stone, Robert E. Schapire, János A. Csirik, Michael L. Littman, and David McAllester. ATTac-2001: A learning, autonomous bidding agent. In *Agent-Mediated Electronic Commerce IV*, volume 2153 of *Lecture Notes in Computer Science*. Springer-Verlag, 2002.

Geoff Sutcliffe. The CADE-17 ATP system competition. *Journal of Automated Reasoning*, 27:227–250, 2001.

P. Taylor and L. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40:145–156, 1978.

Ioannis A. Vetsikas and Bart Selman. A principled study of the design tradeoffs for autonomous trading agents. In *Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 473–480, Melbourne, 2003.

Michael P. Wellman, Amy Greenwald, Peter Stone, and Peter R. Wurman. The 2001 trading agent competition. *Electronic Markets*, 13:4–12, 2003a.

Michael P. Wellman, Jeffrey K. MacKie-Mason, Daniel M. Reeves, and Sowmya Swaminathan. Exploring bidding strategies for market-based scheduling. In *Fourth ACM Conference on Electronic Commerce*, San Diego, 2003b.

Michael P. Wellman, Daniel M. Reeves, Kevin M. Lochner, and Yevgeniy Vorobeychik. Price prediction in a trading agent competition. Technical report, University of Michigan, 2002.

Michael P. Wellman and Peter R. Wurman. A trading agent competition for the research community. In *IJCAI-99 Workshop on Agent-Mediated Electronic Trading*, Stockholm, August 1999.

Peter R. Wurman, William E. Walsh, and Michael P. Wellman. Flexible double auctions for electronic commerce: Theory and implementation. *Decision Support Systems*, 24:17–27, 1998.

Peter R. Wurman, Michael P. Wellman, and William E. Walsh. A parametrization of the auction design space. *Games and Economic Behavior*, 35:304–338, 2001.