

Bidding Heuristics for Simultaneous Auctions: Lessons from TAC Travel

Amy Greenwald and Victor Naroditskiy

Department of Computer Science, Brown University
{amy, vnarodit}@cs.brown.edu

Seong Jae Lee

Department of Computer Science, University of Washington
seongjae@u.washington.edu

Abstract

We undertake an experimental study of heuristics designed for the Travel division of the Trading Agent Competition. Our primary goal is to analyze the performance of the sample average approximation (SAA) heuristic, which is approximately optimal in the decision-theoretic (DT) setting, in this game-theoretic (GT) setting. To this end, we conduct experiments in four settings, three DT and one GT. The relevant distinction between the DT and the GT settings is: in the DT settings, agents' strategies do not affect the distribution of prices. Because of this distinction, the DT experiments are easier to analyze than the GT experiments. Moreover, settings with normally distributed prices, and controlled noise, are easier to analyze than those with competitive equilibrium prices. In the studied domain, analysis of the DT settings with possibly noisy normally distributed prices informs our analysis of the richer DT and GT settings with competitive equilibrium prices. In future work, we plan to investigate whether this experimental methodology—namely, transferring knowledge gained in a DT setting with noisy signals to a GT setting—can be applied to analyze heuristics for playing other complex games.

Introduction

In the design of autonomous trading agents that buy and sell goods in electronic markets, a variety of interesting computational questions arise. One of the most fundamental is to determine how to bid on goods being auctioned off in separate markets when the agent's valuations for those goods are highly interdependent (i.e., complementary or substitutable). The Trading Agent Competition (TAC) Travel division was designed as a testbed in which to compare and contrast various approaches to this problem (Wellman *et al.* 2001). We partake in an empirical investigation of heuristics designed for bidding in the simultaneous auctions that characterize TAC in a simplified TAC-like setting.

At a high-level, the design of many successful TAC agents (for example, Walverine (Cheng *et al.* 2005), RoxyBot (Greenwald and Boyan 2004 & 2005) and ATTac (Stone *et al.* 2003)) can be summarized as: Step 1: *predict*, i.e., build a model of the auctions' clearing prices; Step 2: *optimize*, i.e., solve for an (approximately) optimal set of bids,

given this model. This paper is devoted to the study of bidding, that is, the optimization piece of this design. We assume that agents are given price predictions in the form of a black box from which they can sample a vector of predicted prices; such samples are called *scenarios*. Because finding an optimal solution to the bidding problem is not generally tractable, our study centers around a series of *heuristics* that construct bids based on approximations or simplifications. We subject these heuristics to experimental trials within a simplified version of the TAC domain that we find more amenable to experimental study than the full-blown TAC Travel game.

TAC Travel Game

In this section, we briefly summarize the TAC game. For more details, see <http://www.sics.se/tac/>.

A TAC Travel agent is a simulated travel agent whose task is to organize itineraries for a group of clients to travel to and from TACTown. The agent's objective is to procure "desirable" travel goods as inexpensively as possible. An agent desires goods (i.e., it earns utility for procuring them) to the extent that they comprise itineraries that satisfy its clients' preferences.

Travel goods are sold in simultaneous auctions:

- Flights are sold by the "TAC seller" in dynamic posted-pricing environments. No resale is permitted.
- Hotel reservations are also sold by the "TAC seller," in multi-unit ascending call markets. Specifically, 16 hotel reservations are sold in each hotel auction at the 16th highest price. No resale is permitted.
- Agents trade tickets to entertainment events among themselves in continuous double auctions.

Flights and hotel reservations are complementary goods: flights do not garner utility without complementary hotel reservations; nor do hotel reservations garner utility without complementary flights. Tickets to entertainment events, e.g., the Boston Red Sox and the Boston Symphony Orchestra, are substitutable.

Clients have preferred departure and arrival dates, and a penalty is subtracted from the agent's utility for allocating packages that do not match clients' preferences exactly. For example, a penalty of 200 (100 per day) is incurred when a client who wants to depart Monday and arrive on Tuesday is

assigned a package with a Monday departure and a Thursday arrival. Clients also have hotel preferences, for the two types of hotels, “good” and “bad.” A client’s preference for staying at the good rather than the bad hotel is described by a *hotel bonus*, utility the agent accumulates when the client’s assigned package includes the good hotel.

Bidding Heuristics

Our test suite consists of six marginal-utility-based and two sample average approximation heuristics. We present a brief description of these heuristics here. Interested readers are referred to Wellman, Greenwald, and Stone (2007) for more detailed explanations.

Marginal-Utility-Based Heuristics

In a second-price auction for a single good, it is optimal for an agent to simply bid its independent value on that good (Vickrey 1961). In simultaneous auctions for multiple goods, however, bidding is not so straightforward because it is unclear how to assign independent values to interdependent goods. Perfectly complementary goods (e.g., an inflight and outflight for a particular client) are worthless in isolation, and perfectly substitutable goods (e.g., rooms in different hotels for the same client on the same day) provide added value only in isolation. Still, an agent might be tempted to bid on each good its *marginal utility* (MU), that is, the incremental value of obtaining that good relative to the collection of goods it already owns or can buy. Many reasonable bidding heuristics (e.g., Greenwald and Boyan (2004, 2005), Stone *et al.* (2003)) incorporate some form of marginal utility bidding.

Definition 1 Given a set of goods X , a valuation function $v : 2^X \rightarrow \mathbb{R}$, and bundle prices $q : 2^X \rightarrow \mathbb{R}$. The *marginal utility* $\mu(x, q)$ of good $x \in X$ is defined as:

$$\mu(x) = \max_{Y \subseteq X \setminus \{x\}} [v(Y \cup \{x\}) - q(Y)] - \max_{Y \subseteq X \setminus \{x\}} [v(Y) - q(Y)]$$

Consistent with TAC Travel, we assume additive prices: that is, in the above equation, the bundle pricing function q returns the sum of the predicted prices of the goods in Y .

Our heuristics actually sample a set of scenarios, not a single vector of predicted prices. We consider two classes of marginal utility heuristics based on how they make use of the information in the scenarios.

Bidding Heuristics that Collapse Available Distributional Information The following heuristics collapse all scenarios into a single vector of predicted prices, namely the average scenario, and then calculate the marginal utility of each good assuming the other goods can be purchased at the average prices.

StraightMU bids the marginal utility of each good.

TargetMU bids marginal utilities only on the goods in a target set of goods. The target set is one that an agent would optimally purchase at the average prices.

TargetMU* is similar to TargetMU, but calculates marginal utilities assuming only goods from the target set are available. This results in higher bids.

Bidding Heuristics that Exploit Available Distributional Information The heuristics discussed thus far collapse the distributional information contained in the sample set of scenarios down to a point estimate, thereby operating on approximations of the expected clearing prices. The heuristics described next more fully exploit any available distributional information; they seek bids that are effective across multiple scenarios, not in just the average scenario.

AverageMU calculates the marginal utilities of all goods, once per scenario, and then bids the *average* MU of each good in each auction.

BidEvaluator evaluates K candidate bidding policies on a fixed set of E sample scenarios. The policy that earns the highest total score is selected.

BidEvaluator generates its candidates by making successive calls to the TargetMU heuristic, each time sending it a different scenario to use as its predicted prices.

BidEvaluator* is identical to BidEvaluator, except that its candidate bidding policies are generated by calling TargetMU* instead of TargetMU.

Sample Average Approximation

The problem of bidding under uncertainty—how to bid given a distributional model of predicted prices—is a stochastic optimization problem. The objective is to select bids that maximize the expected value of the difference between the value of the goods the agent wins and the cost of those goods. Formally,

Definition 2 [Stochastic Bidding Problem] Given a set of goods X , a (combinatorial) valuation function $v : 2^X \rightarrow \mathbb{R}$, and a distribution f over clearing prices $p \in \mathbb{R}^X$, the *stochastic bidding problem* is defined as:

$$\max_{b \in \mathbb{R}^X} \mathbb{E}_{p \sim f} [v(\text{Win}(b, p)) - \tilde{p}(\text{Win}(b, p))] \quad (1)$$

Here, $x \in \text{Win}(b, p)$ if and only if $b(x) \geq p(x)$, and $\tilde{p} : 2^X \rightarrow \mathbb{R}$ is the *additive* extension of $p \in \mathbb{R}^X$, that is, the real-valued function on bundles defined as follows: $\tilde{p}(Y) = \sum_{x \in Y} p(x)$, for all $Y \subseteq X$.

Sample average approximation (SAA) is a standard way of approximating the solution to a stochastic optimization problem, like bidding under uncertainty. The idea behind SAA is simple: (i) generate a set of sample scenarios, and (ii) solve an approximation of the problem that incorporates only the sample scenarios.

Technically, the TAC Travel bidding problem, in which the goal is to maximize the difference between the value of allocating travel packages to clients and the costs of the goods procured to create those packages, is a stochastic program with integer recourse (Lee, Greenwald, & Naroditskiy

2007). Using the theory of large deviations, Ahmed and Shapiro (2002) establish the following: the probability that an optimal solution to the sample average approximation of a stochastic program with integer recourse is in fact an optimal solution to the original stochastic program approaches 1 exponentially fast as the number of scenarios $S \rightarrow \infty$. Given time and space constraints, however, it is not always possible to sample sufficiently many scenarios to make any reasonable guarantees about the quality of a solution to the sample average approximation.

Our default implementation of SAA which we call SAABottom always bids one of the sampled prices. However, given a set of scenarios, SAA is indifferent between bidding the highest sampled price or any amount above that price: in any case SAA believes it will win in all scenarios. Consequently, we do not know exactly how much SAA is willing to pay when it bids the highest sampled price. In the settings with imperfect price prediction or when SAA is given too few scenarios, it may be desirable to bid above the highest sampled price to increase the chances of winning. For this reason, we introduce a modified SAA heuristic—SAATop—in which bids equal to the highest sampled price are replaced with the “maximum” bid. In general, this bid is the most the agent is willing to pay. In our domain, this maximum is the sum of the utility bonus (300; see Footnote 2) and, for good hotels, the largest hotel bonus among the agent’s clients’.

Experiments in TAC Travel-like Auctions

We consider four experimental settings: normally distributed prices in two decision-theoretic settings, one with perfect and another with imperfect prediction; and competitive equilibrium (CE) prices in a decision-theoretic setting with perfect prediction and a game-theoretic setting with typically imperfect prediction.

Our experiments were conducted in a TAC Travel-like setting, in which nearly all the standard rules apply.¹ Most notably, we simplified the dynamics of the game. In TAC, flights and entertainment tickets are available continuously at time-varying prices, and hotel auctions close one at a time, providing opportunities for agents to revise their bids on other hotels. In this work, we focus on one-shot auctions. More specifically, we assume all hotels close after one round of bidding.

To reduce variance, we eliminated entertainment trading and simplified flight trading by fixing flight prices at zero.^{2,3}

¹For a detailed description of the TAC Travel rules, visit <http://www.sics.se/tac>.

²Since we fixed flight prices at zero (instead of roughly 700 for round trip tickets), we adjusted the utility bonus for constructing a valid travel package down from 1000 to 300. That way, our simulation scores fall in the same range as real game scores.

³Initially, we ran experiments with flight prices fixed at 350, which is the value close to the average flight price in the TAC Travel game. However the resulting one-shot setting was not interesting as flight tickets represented a very high sunk cost and the dominant hotel bidding strategy was to bid very high on the hotels that would complement the flights in completing travel packages.

We built a simulator of the TAC server, which can easily be tailored to simulate numerous experimental designs. Our simulator is available for download at <http://www.sics.se/tac/showagents.php>.

Each trial in an experiment (i.e., each simulation run) proceeded in five steps:

1. The agents predict hotel clearing prices in the form of *scenarios* - samples from the predicted distribution of clearing prices.
 - In the settings where prices are normally distributed, the scenarios were sampled from given distributions of predicted prices.
 - In the settings characterized by competitive equilibrium prices, scenarios were generated by simulating simultaneous ascending auctions, as described in Lee *et al.* (2007).
2. The agents construct bids using price information contained in the scenarios and submit them.
3. The simulator determines hotel clearing prices, and bids that are equal to or above those clearing prices are deemed winning bids.
 - In the *decision-theoretic* settings, the clearing prices were sampled from given distributions of clearing prices.
 - In the *game-theoretic* setting, each hotels’ clearing price was set to the 16th highest bid on that hotel.
4. Agents pay clearing prices for the hotels they win. They use the hotels and free flight tickets to create packages for their clients, based on which they earn the corresponding utilities.
5. Each agent’s final score is the difference between its utility and its cost.

The first two steps in the above sequence correspond to the prediction and optimization steps typical of autonomous bidding agents. To carry out step 2, the agents employ heuristics from a test suite that includes the eight bidding heuristics detailed in Wellman, *et al.*(2007), and summarized above.

Regarding price prediction in step 1, hotel price predictions were perfect in our first and third experimental setups and imperfect in our second and fourth. In the first two, hotel prices were predicted to be normally distributed; in the second two, hotel prices were predicted to be competitive equilibrium prices. Our first three experimental setups were decision-theoretic; the fourth was game-theoretic. In the second setup, we simply tweaked the normal distribution of predicted prices to generate a similar, but distinct, normal distribution of clearing prices. In the fourth setup, the game-theoretic setting, clearing prices were dictated by the outcome of 16th price auctions. Our experimental design is summarized in Table 1. All setups, with all settings of the parameters (μ , σ , and λ), were run for 1000 trials.

Heuristic Parameter Settings

The parameter settings we chose for the heuristics are shown in Table 2. Breaking down a TAC agent’s work into two

Experimental Design

Normally-Distributed Prices	Perfect Prediction	DT
Normally-Distributed Prices	Imperfect Prediction	DT
CE Prices	Perfect Prediction	DT
CE Prices	Imperfect Prediction	GT

Table 1: Price predictions were either normally distributed or competitive equilibrium prices. Moreover, they were sometimes perfect and sometimes imperfect. Three of the four experimental setups were decision-theoretic (DT); only the fourth was game-theoretic (GT).

key steps—price prediction and optimization—the column labeled SG lists the scenario generation (i.e., CE price prediction) times; the column labeled BC lists the bid construction (i.e., optimization) times. The rightmost column lists total runtimes. The goal in choosing these parameter settings was to roughly equalize total runtimes across agents in TAC games.

All experiments were run on AMD Athlon(tm) 64 bit 3800+ dual core processors with 2GB of RAM. All times are reported in seconds, averaged over 1000 games. The machines were not dedicated, which explains why generating 50 scenarios could take anywhere from 8.7 to 9.4 seconds, on average. Presumably, all the heuristics (but most notably, AverageMU, the variants of BidEvaluator, and the SAA heuristics) could benefit from higher settings of their parameters.

Agent	E	S	K	SG	BC	Total
TMU	–	50	–	9.4	1.0	10.4
TMU*	–	50	–	9.0	1.1	10.1
BE	15	–	25	7.0	5.3	12.3
BE*	15	–	25	7.0	4.7	11.7
AMU	–	15	–	2.3	10.2	12.5
SMU	–	50	–	8.7	1.5	10.2
SAABottom	–	50	–	8.8	1.7	10.5
SAATop	–	50	–	9.0	1.6	10.6

Table 2: Parameter Settings. E is the number of evaluations, S is the number of scenarios, and K is the number of candidate bidding policies.

We optimized the heuristics that bid only on the goods in a target set to bid ∞ on all flights in that set; they do not bother to calculate the marginal utilities of their desired flights.⁴ This helps explain why the bid construction phase within TargetMU and TargetMU* is so fast. StraightMU, and hence AverageMU, are also optimized to stop computing marginal utilities once a good’s marginal utility hits zero.

⁴Note that we ran many more experiments than those reported here. In particular, flight prices were not always zero (e.g., see Footnote 3). Indeed, in many cases it was sensible for the various heuristics to make informed decisions about how to bid on flights.

Multiunit Marginal Utility

TAC Travel auctions are multi-unit auctions. For bidding in multi-unit auctions, we extend the definition of marginal utility, originally defined for a single copy of each good, to handle multiple copies of the same good. The marginal utility of the first copy of a good is calculated assuming that no other copies of the good can be had; the marginal utility of the second copy of a good is calculated assuming that the first copy is on hand but that no other copies can be had; and so on.

We assume the set of goods X contains J goods, with K_j copies of each good $1 \leq j \leq J$.

Definition 3 Given a set of goods X , a valuation function $v : 2^X \rightarrow \mathbb{R}$, and a pricing function $q : 2^X \rightarrow \mathbb{R}$. The *marginal utility* $\mu(x_{jk}, X, v, q)$ of the k th copy of good j is given by:

$$\max_{Y \subseteq X \setminus \{x_{j1}, \dots, x_{jk}\}} [v(Y \cup \{x_{j1}, \dots, x_{jk}\}) - q(Y)] - \max_{Y \subseteq X \setminus \{x_{j1}, \dots, x_{j,k-1}\}} [v(Y \cup \{x_{j1}, \dots, x_{j,k-1}\}) - q(Y)]$$

In words, the marginal utility of the k th copy of good j is simply the difference between the value of an optimal set of goods to buy, assuming x_{j1}, \dots, x_{jk} cost 0 and $x_{j,k+1}, \dots, x_{jN}$ cost ∞ , and the value of an optimal set of goods to buy, assuming $x_{j1}, \dots, x_{j,k-1}$ cost 0 and x_{jk}, \dots, x_{jN} cost ∞ .

Our agent implementations of the marginal-utility-based agents employ this definition.

Decision-Theoretic Experiments with Perfect Distributional Prediction

Our first experimental setup is decision-theoretic, with prices determined exogenously. Each agent is endowed with perfect distributional information, so that it constructs its bids based on samples drawn from the true price distribution. Under these conditions, it is known that the SAA-based heuristics bid optimally in the limit as $S \rightarrow \infty$ (Ahmed & Shapiro 2002). The purpose of conducting experiments in this setting was twofold: (i) to evaluate the performance of the SAA-based heuristics with only finitely many scenarios; and (ii) to evaluate the performance of the MU-based heuristics relative to that of the SAA-based heuristics. We find that both the SAA-based heuristics and certain variants of the MU-based heuristics (primarily, TargetMU* and BidEvaluator*) perform well assuming low variance, but that the SAA-based heuristics and AverageMU outperform all the other heuristics assuming high variance.

Setup

Hotel prices were drawn from normal distributions with means⁵ $\bar{\mu} = (150, 150, 150, 150, 250, 250, 250, 250)$ constant across experiments and standard deviations $\sigma \in \{0, 20, 40, 60, 80, 100\}$ varying across experiments.

⁵In this, and all, hotel price vectors, the first four numbers refer to the price of the bad hotel on days 1 through 4, respectively, and the second four numbers refer to the price of the good hotel on days 1 through 4, respectively.




Figure 1: Mean Scores. Decision-theoretic setting with perfect distributional prediction.

Results

Figure 1 depicts the mean scores earned by each agent in each experiment: i.e., for each setting of σ .

The SAA-based agents perform better than most of the agents as variance increases. They gain an advantage by submitting low bids on more goods than necessary in an attempt to win only the goods that are cheap. We refer to this strategy as *hedging*. We see that the SAA agents employ hedging because the number of bids they place increases, their average bids decrease, and the number of hotels they win remains constant as the variance increases. The number of low-priced hotels increases with the variance making hedging especially effective when variance is high.

Recall that target bidders (TargetMU, TargetMU*, BidEvaluator, and BidEvaluator*) bid only on goods in their target set, i.e. they do not hedge. Consequently, failing to win one of the requisite hotels results in not being able to complete a package (most packages are for one-night stays as extending the stay for an extra day is likely to be more expensive than incurring the penalty for deviating from client’s preferences). TargetMU and BidEvaluator win fewer and fewer hotels as the variance increases, and hence complete fewer and fewer packages. At the same time the average cost of hotels they win decreases. The agents’ scores have a slight upward trend as the benefit from lower cost outweighs the loss from completing fewer packages.

BidEvaluator bids on more hotels than TargetMU when variance is 100. Recall that BidEvaluator chooses the best of K bidding policies. Bidding policies that bid on more hotels score higher because they hedge, implicitly. For example, a policy that bids to reserve two nights for a client may earn a higher score than a policy that bids to reserve one night as the reservation for two nights can be used to create two separate one-night packages if some of the other bids fail.

TargetMU* and BidEvaluator*, the main rivals of the SAA-based agents, do not perform well in this setting. Just like TargetMU and BidEvaluator, TargetMU* and Bid-

Evaluator* bid only on target goods. When variance is low ($\sigma = 20$), bidding high on target good is a good strategy as evidenced by TargetMU**s and BidEvaluator**s good performance. As variance increases the agents fail to win some of the target goods. In fact when variance is 100, TargetMU* submits 5.8 bids but wins only 4.8 while BidEvaluator* submits 7.4 bids and wins only 5.2. The average cost of hotels that TargetMU* and BidEvaluator* do win is 50% higher than the prices the SAA-based agents pay per hotel.

Interestingly, AverageMU’s strategy happens to be very close to hedging when variance is high. StraightMU submits a lot of bids too but unlike AverageMU does not perform well. StraightMU’s bids are higher than AverageMU’s resulting in more purchased hotels and higher average hotel cost. The increase in cost that StraightMU incurs compared to AverageMU is not compensated by the increase in utility that extra hotels bring.

In conclusion, the SAA-based agents and AverageMU with their hedging strategy outperform the other agents when variance is high.

Decision-Theoretic Experiments with Imperfect Distributional Prediction

In our second decision-theoretic experimental setup, the agents construct their bids based on samples drawn from a normal distribution that resembles, but is distinct from, the true distribution. Our intent here is to evaluate the agents’ behavior in a controlled setting with imperfect predictions, in order to inform our analysis of their behavior in the game-theoretic setting, where predictions are again imperfect. We find that SAATop performs worse than TargetMU*, and BidEvaluator* at low variance, but outperforms most of the other agents at high variance.

Setup

In these experiments, the *predicted* price distributions were normal with mean values $\bar{\mu} = (150, 150, 150, 150, 250, 250, 250, 250)$, whereas the *clearing* price distributions were normal with mean values $\bar{\mu} + \lambda$. That is, the mean of each predicted distribution differed by λ from the true mean. For example, for $\lambda = -40$, predicted prices were sampled from normal distributions with $\bar{\mu} = (150, 150, 150, 150, 250, 250, 250, 250)$, and clearing prices were sampled from normal distributions with $\bar{\mu} = (110, 110, 110, 110, 210, 210, 210, 210)$. Hence, negative values of λ implied “overprediction.” Similarly, positive values of λ implied “underprediction.” The λ parameter varied as follows: $\lambda \in \{-40, -30, -20, -10, 0, 10, 20, 30, 40\}$. We chose as standard deviations of the distributions a low setting ($\sigma = 20$) and a high setting ($\sigma = 80$).

In the low (and similarly in the high) deviation experiments the strategies of the agents did not change with λ because the agent received the same predictions for all values of λ . Experiments in this setting evaluate the strategies from the perfect prediction setting with $\sigma = 20$ and $\sigma = 80$ under different distributions of clearing prices as controlled by the values of λ .

Results

Low Variance: $\sigma = 20$ The results assuming low variance are shown in Figure 2(a).

Recall from the perfect prediction experiments that the strategy of bidding high on the goods from a target set is as good as hedging when variance is low. In particular, **TargetMU*** and **BidEvaluator*** perform as well as the SAA-based agents. We will see that hedging is not a good strategy in the low-variance setting with imperfect prediction while bidding high on the goods in a target set works fairly well.

In an attempt to hedge, the SAA-based agents submit twice as many bids as **TargetMU**, **TargetMU***, **BidEvaluator**, and **BidEvaluator***. The strategy of the SAA-based agents is to bid low hoping to win approximately half the bids. Because predictions are not perfect, the SAA-based agents win too many hotels when prices are lower than expected and too few hotels when prices are higher than expected. Not surprisingly, **SAATop**, which bids higher than its counterpart, performs worse than **SAABottom** when prices are lower than expected and better than **SAABottom** when the opposite is true.

When there is a high degree of overprediction and variance is low, (e.g., when $\lambda = -40$ and $\sigma = 20$), clearing prices are very likely to be below predicted prices. Since **TargetMU** always bids at least the predicted price, it is likely to win all the hotels it expects to win in this setting, and hence performs well. Consequently, **TargetMU***, **BidEvaluator**, and **BidEvaluator*** all perform well. In contrast, when prices are often lower than expected, **AverageMU** and **StraightMU** win too many goods and thus incur high unnecessary costs.

As λ increases from -40 to -10 , **AverageMU** and **StraightMU** win fewer unnecessary hotels, which improves their scores. But once λ reaches 0, they fail to win enough hotels, and their utilities decrease as λ increases to 40. **TargetMU** and **BidEvaluator** encounter the same difficulty.

TargetMU* and **BidEvaluator*** bid higher than **TargetMU** and **BidEvaluator**; hence, underprediction affects the former pair less than the latter pair.

To summarize, in the low-variance setting **TargetMU***'s and **BidEvaluator***'s strategy of bidding high on a target set of goods is more robust to imperfect predictions than the strategy of the SAA-based agents that involves some hedging.

High Variance: $\sigma = 80$ The results assuming high variance are shown in Figure 2(b).

As we observed in the experiments with perfect prediction and $\sigma = 80$, hedging allowed the SAA-based agents to dominate. We are going to see that hedging is effective in the high-variance setting even when predictions are not perfect.

The SAA-based agents submit over four times as many bids as **TargetMU**, **TargetMU***, **BidEvaluator**, and **BidEvaluator***. In contrast to the setting with low variance, high overprediction ($\lambda = -40$) does not cause the SAA-based agents to overspend on hotels. In the high-variance setting the SAA-based agents' bids are 40% lower than in

the low-variance setting ($\sigma = 20$) and only one-third of the bids are winning bids.

Similarly, the SAA-based agents perform much better in the high underprediction ($\lambda = 40$) setting when variance is high than when variance is low. In the high-variance setting with underprediction the SAA-based agents win at least as many hotels as the high bidding **TargetMU*** and **BidEvaluator*** agents. Although the SAA-based agents bid half the price that **TargetMU*** and **BidEvaluator*** bid, a much higher number of bids that the SAA-based agents submit combined with high variance results in a similar number of winning bids.

Performance of the other agents is similar to their performance in the setting with perfect prediction. **TargetMU**, **TargetMU***, **BidEvaluator**, and **BidEvaluator*** do not hedge and perform poorly in under and over prediction settings. **Target** bidders often fail to win some of the target hotels even in the overprediction setting. **AverageMU** submits a lot of low bids resulting in a well-hedged strategy and the scores that are as high as SAA's for some values of λ . As before, **StraightMU** wins too many hotels.


In contrast to the setting with low variance and imperfect predictions, the SAA-based agents' hedging strategy works well when there is high variance.

Experiments with Competitive Equilibrium Prices


In contrast with our first two experimental settings, in which the hotel clearing prices and their corresponding predictions are exogenously determined and hence independent of any game specifics, in our second two experimental settings, both hotel clearing prices and predictions are determined endogenously (i.e., based on features of each game instance). Specifically, following Walverine (Cheng *et al.* 2005), hotel clearing prices and their corresponding predictions are taken to be approximate *competitive equilibrium* (CE) prices. CE prices are prices at which supply equals demand when all market participants act as price-taking profit maximizers (Mas-Colell, Whinston, & Green 1995). CE prices need not exist, and likely do not in many of the games studied here. Still, we approximate CE prices as follows: in a market inhabited by its own eight clients and eight randomly sampled clients per competitor, each agent generates a scenario by simulating simultaneous ascending auctions (i.e., increasing prices by some small increment until supply exceeds demand; see Lee *et al.* (2007) for details); the resulting prices form a scenario.

Setup

In this context, where hotel price predictions are (roughly) competitive equilibrium prices, we designed two sets of experiments: one decision-theoretic and one game-theoretic. In the former, hotel clearing prices are also the outcome of a simulation of simultaneous ascending auctions, but depend on the actual clients in each game, not some random sampling like the agents' predictions. (Our simulator is more informed than the individual agents.) In the latter, hotel clearing prices are determined by the bids the agents submit. As



(a) $\sigma = 20$



(b) $\sigma = 80$

Figure 2: Mean Scores. Imperfect prediction.

in TAC Travel, the clearing price is the 16th highest bid (or zero, if fewer than 16 bids are submitted). Note that hotel clearing prices and their respective predictions are not independent of one another in these experiments.

In these experiments games are played with a random number of agents drawn from a binomial distribution with $n = 32$ and $p = 0.5$, with the requisite number of agents sampled uniformly with replacement from the set of eight possible agent types. The agents first sample the number of competitors from the binomial distribution, and then generate scenarios assuming the sampled number of competitors, resampling that number to generate each new scenario.

Decision-Theoretic Experiments

Marginal frequency distributions of CE prices in these experiments have means (109, 126, 126, 107, 212, 227, 227, 210) and standard deviations (47, 37, 37, 46, 50, 41, 41, 49). Standard deviation in this setting is close to 40 making this setting similar to the one with perfect prediction and $\sigma = 40$. The mean hotel prices are approximately 20% lower in this CE setting but we do not expect the difference in mean hotel prices to have a strong effect on the ranking of the agents

and attribute the differences in relative results to the different structure of prices: unlike the setting with normally distributed prices, CE prices are not independent.

SAATop, SAABottom, TargetMU*, and BidEvaluator* are among the best agents in this CE setting (see Figure 3). However, StraightMU and especially AverageMU perform poorly. AverageMU and StraightMU submit more bids and win more hotels than the other agents, but cannot create as many packages as the top-scoring agents. This is because (i) CE prices of substitutable goods are similar, and (ii) marginal utilities of substitutable goods are similar. As a result, AverageMU and StraightMU bid almost the same amount on all substitutable goods and either win or lose all of them.

SAA-based agents employ some hedging but do not perform significantly better than the non-hedging heuristics TargetMU* and BidEvaluator*.

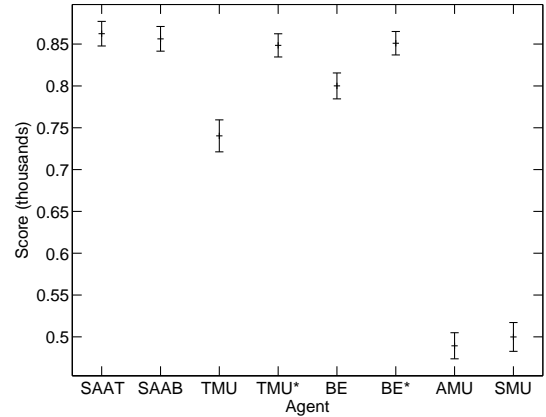


Figure 3: Mean scores and confidence intervals. Decision-theoretic setting with CE price prediction.

Game-Theoretic Experiments

The predicted prices are the same as in the decision-theoretic experiments with CE prices and 32 agents: means (109, 126, 126, 107, 212, 227, 227, 210) and standard deviations (47, 37, 37, 46, 50, 41, 41, 49). Marginal frequency distributions of clearing prices have means (91, 98, 100, 91, 198, 186, 187, 197) and standard deviations (41, 33, 32, 40, 50, 56, 54, 50). L1-norm of the difference between mean price vectors is 197. Predicted prices are slightly higher (by about 20) than the clearing prices. This is similar to the decision-theoretic setting with overprediction ($\lambda = -20$) and medium deviation (between 20 and 80).

Indeed, we find that the results in this setting (see Figure 4) are similar to the results in the decision-theoretic setting with imperfect prediction and high variance: $\lambda = -20$ and $\sigma = 80$ (see the ranking of agents for $\lambda = -20$ in Figure 2(b)). The ranking of non-SAA agents is almost the same in both settings. A notable exception is AverageMU, which performs much worse in the game-theoretic setting for the reasons described above. SAATop and SAABottom are the best agents in this setting, with SAABottom

performing slightly better.




Figure 4: Mean scores and confidence intervals. Game-theoretic setting with CE price prediction.

Summary and Discussion of Experimental Results

In our experiments, we evaluated the performance of various bidding heuristics in simultaneous auctions. Based on our findings, we summarize the performance of the heuristics analyzed as follows:

- SAATop and SAABottom perform well in all settings except for the setting with imperfect prediction and low variance. SAATop and SAABottom are especially effective in high-variance settings because they are able to take advantage of hedging opportunities.
- TargetMU and BidEvaluator are competitive only in the settings with low variance and high overprediction. BidEvaluator outperforms TargetMU in high-variance settings.
- TargetMU* and BidEvaluator* perform well in the settings with low variance.
- AverageMU performs well in the settings with independent prices.
- StraightMU performs worse than the other heuristics.

We can also make the following observations about the various bidding behaviors:

- SAABottom, SAATop, and AMU place low bids on many goods, intending to win whatever sells at cheap prices. These heuristics incur high penalties for not satisfying their clients' precise preferences.
- TargetMU, TargetMU*, BidEvaluator, and BidEvaluator* place higher bids but on fewer goods, namely those for which their clients have clear preferences. These heuristics incur lower penalties, but risk alienating some clients, by not allocating them any travel packages at all.⁶

⁶No penalty is incurred when a client is not allocated any package at all. (Of course, no utility is awarded either.)

The performance of SAA is known to approach optimality as the number of scenarios approaches ∞ in decision-theoretic settings. We investigated the viability of two SAA heuristics with only finitely-many scenarios in both decision-theoretic and game-theoretic settings. Our first and third experimental settings (with normally distributed and competitive equilibrium prices, assuming perfect price prediction) established the viability of these heuristics in decision-theoretic settings with only finitely-many scenarios. Our fourth experimental setting established the viability of these heuristics (again, with only finitely-many scenarios, but in addition) in a rich game-theoretic setting.

Related Work

The test suite considered here is far from exhaustive. In this section, we mention several heuristics that were not included in our study—some TAC-specific; some more general—and the reasons for their exclusion.

The creators of the ATTac agent (Stone *et al.* 2003) propose using AverageMU for TAC hotel bidding. ATTac also employs distributional information about hotel prices to determine the benefit of postponing flight purchases until hotel prices are known; this additional functionality, while certainly of interest, is not applicable to the one-shot auction setting studied here.

WhiteBear's (Vetsikas & Selman 2003) TAC hotel bids are computed by taking a weighted average of the current price and the marginal utility of each hotel. The particular weights, which were fine-tuned based on historical competition data, varied with time. In a one-shot setting, WhiteBear's strategy essentially reduces to TargetMU: it is too risky to bid anything lower.

SouthamptonTAC (He & Jennings 2003) and MertaCor (Toulis, Kehagias, & Mitkas 2006) focus on hotel price prediction, and do not thoroughly analyze bidding. SouthamptonTAC uses fuzzy reasoning to predict how hotel prices change during the game.

Unlike the heuristics studied in this paper, Walverine's (Cheng *et al.* 2005) bidding strategy incorporates some game-theoretic reasoning. Specifically, Walverine analytically calculates the distribution of marginal utilities of the other agents' clients and bids a best-response to this distribution. The authors implicitly assume that the other agents bid marginal utilities (i.e., act decision-theoretically) and only their agent bids a best-response (i.e., acts game-theoretically). We learned from the study reported in this paper that SAA can be a successful bidding heuristic in certain markets. Following Walverine's line of thought, we can imagine bidding a best-response to a distribution of SAA bids. However, if this bidding strategy were successful, we would have to assume that other agents would act game-theoretically as well; that is, they would also play a best-response to a distribution of SAA bids. We may then seek a fixed point of this process. This line of inquiry could be fascinating, but any approach based on this insight of Walverine's warrants a detailed study of its own.

Aside from TAC Travel there is a rich literature on bidding in other settings. We reference a few papers here, highlighting some of the settings that have been studied. We are

not aware of any papers that address the problem of bidding in multiple one-shot auctions for both complementary and substitutable goods. Gerding *et al.* (2006) describes a strategy for bidding in simultaneous one-shot second-price auctions selling perfect substitutes. Bye, Priest, & Jennings (2002) consider the decision-theoretic problem of bidding in multiple auctions with overlapping closing times. Their model treats all goods as indistinguishable (i.e., winning any n goods results in utility $v(n)$). Krishna & Rosenthal (1996) characterize a symmetric equilibrium for the case of one-shot simultaneous auctions with indistinguishable complementary goods (i.e. $v(n) \geq nv(1)$).

Conclusion

The primary purpose of this work was to show that using as much distributional information as possible is an effective approach to bidding in TAC Travel-like one-shot simultaneous auctions. Most TAC Travel agents used point price predictions or employed little distributional information about prices in constructing their bids. Some of the difficulties with using distributional price predictions include the inaccuracy of and the high computational cost of optimizing with respect to distributional predictions. We showed experimentally that the SAA heuristic, which uses more distributional information than the other heuristics in our test suite, is one of the best heuristics in the GT setting.

The underlying research question motivating this line of inquiry was: how can we facilitate the search for heuristics that perform well against a variety of competing agents in complex games? Analyzing the performance of an individual agent in a game-theoretic setting is complicated because each agent's performance is affected by the strategies of the others, and can vary dramatically with the mix of participants. Others tackling this problem in the TAC Travel domain have employed more direct game-theoretic analysis techniques based on equilibrium computations (e.g., Vetsikas *et al.* (2007) and Jordan, Kiekintveld, & Wellman (2007)). In contrast, we first used systematic decision-theoretic analysis to help us understand some of the intrinsic properties of our bidding heuristics, before attempting any game-theoretic analysis. We found that certain properties of the heuristics that may have been hard to identify in game-theoretic settings, such as how they perform in conditions of over- vs. under-prediction, carried over from our DT to our GT settings.

In summary, the methodology advocated in this paper for analyzing game-theoretic heuristics is this: first, evaluate the heuristic in DT settings with perfect and imperfect predictions; and second, measure the accuracy of the agent's predictions in GT experiments and use the corresponding DT analysis to inform the analysis of the GT results. It remains to test this methodology in other complex games, such as TAC SCM (Arunachalam & Sadeh 2005).

Acknowledgments

This research was funded by NSF Grant #IIS-0133689 and the Samsung Foundation of Culture

Appendix

We include game statistics collected from our experiments with CE prices to illustrate the type of data we used in our analyses. Statistics for other settings and further data can be found in Lee (2007).

References

- Ahmed, S., and Shapiro, A. 2002. The sample average approximation method for stochastic programs with integer recourse. *Optimization Online*, <http://www.optimization-online.org>.
- Arunachalam, R., and Sadeh, N. M. 2005. The supply chain trading agent competition. *Electronic Commerce Research and Applications* 4(1):66–84.
- Bye, A.; Priest, C.; and Jennings, N. R. 2002. Decision procedures for multiple auctions. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, 613–620. New York, NY, USA: ACM.
- Cheng, S.; Leung, E.; Lochner, K.; K.O'Malley; Reeves, D.; Schwartzman, L.; and Wellman, M. 2005. Walverine: A Walrasian trading agent. *Decision Support Systems* 39(2):169–184.
- Gerding, E.; Dash, R.; Yuen, D.; and Jennings, N. R. 2006. Optimal bidding strategies for simultaneous vickrey auctions with perfect substitutes. In *8th Int. Workshop on Game Theoretic and Decision Theoretic Agents*, 10–17.
- Greenwald, A., and Boyan, J. 2004. Bidding under uncertainty: Theory and experiments. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, 209–216.
- Greenwald, A., and Boyan, J. 2005. Bidding algorithms for simultaneous auctions: A case study. *Journal of Autonomous Agents and Multiagent Systems* 10(1):67–89.
- He, M., and Jennings, N. R. 2003. Southamptontac: An adaptive autonomous trading agent. *ACM Trans. Interet Technol.* 3(3):218–235.
- Jordan, P. R.; Kiekintveld, C.; and Wellman, M. P. 2007. Empirical game-theoretic analysis of the tac supply chain game. In *AAMAS '07: Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, 1–8. New York, NY, USA: ACM.
- Krishna, V., and Rosenthal, R. W. 1996. Simultaneous auctions with synergies. *Games and Economic Behavior* 17(1):1–31. available at <http://ideas.repec.org/a/eee/gamebe/v17y1996i1p1-31.html>.
- Lee, S.; Greenwald, A.; and Naroditskiy, V. 2007. Roxybot-06: An (SAA)² TAC travel agent. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, 1378–1383.
- Lee, S. J. 2007. Comparison of bidding algorithms in simultaneous auctions. B.S. honors thesis, Brown University, <http://list.cs.brown.edu/publications/theses/ugrad/>.

Agent	SAAT	SAAB	TMU	TMU*	BE	BE*	AMU	SMU
Score	863	859	742	847	799	850	488	499
Utility	1935	1838	1420	1811	1619	1830	1805	1902
Cost	1071	979	677	964	819	980	1316	1402
Penalty	383	348	251	309	237	331	229	234
# of Clients without a Package	1.60	2.01	3.40	2.20	2.89	2.09	2.48	2.28
# of Hotel Bids	11.0	10.7	6.3	6.3	5.8	6.8	42.2	35.5
Average Hotel Bid	270	187	187	292	233	280	106	125
Total Hotel Bonus	400	389	292	383	324	390	378	422
# of Hotels Won	6.7	6.2	4.6	5.8	5.1	5.9	9.3	9.6
# of Unused Hotels	0.0	0.0	0.0	0.0	0.0	0.0	1.2	1.3
Average Hotel Cost	159.6	159.0	147.4	166.5	160.5	166.0	140.9	145.6

Decision-theoretic setting

Agent	SAAT	SAAB	TMU	TMU*	BE	BE*	AMU	SMU
Score	981	999	899	954	938	948	652	617
Utility	2057	2007	1708	1885	1796	1904	2386	2397
Cost	1075	1007	808	931	857	955	1734	1779
Penalty	437	418	287	345	272	375	338	358
# of Clients without a Package	1.13	1.33	2.53	1.90	2.34	1.75	0.64	0.60
# of Hotel Bids	10.8	10.7	6.3	6.3	5.9	6.8	42.1	35.5
Average Hotel Bid	274	188	187	292	234	281	106	125
Total Hotel Bonus	435	426	354	400	371	404	518	536
# of Hotels Won	7.6	7.1	5.5	6.1	5.7	6.3	12.5	12.5
# of Unused Hotels	0.0	0.0	0.0	0.0	0.0	0.0	1.9	1.8
Average Hotel Cost	142.4	142.4	147.8	152.6	151.5	152.8	138.3	142.8

Game-theoretic setting

Table 3: Game statistics for experiments with CE prices

Mas-Colell, A.; Whinston, M.; and Green, J. 1995. *Microeconomic Theory*. New York: Oxford University Press.

Stone, P.; Schapire, R.; Littman, M.; Csirik, J.; and McAllester, D. 2003. Decision-theoretic bidding based on learned density models in simultaneous, interacting auctions. *Journal of Artificial Intelligence Research* 19:209–242.

Toulis, P.; Kehagias, D.; and Mitkas, P. 2006. Mertacor: A successful autonomous trading agent. In *Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, 1191–1198.

Vetsikas, I. A., and Selman, B. 2003. A principled study of the design tradeoffs for autonomous trading agents. In *AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, 473–480. New York, NY, USA: ACM.

Vetsikas, I. A., and Selman, B. 2007. Bayes-nash equilibria for mth price auctions with multiple closing times. *SIGecom Exch.* 6(2):27–36.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16:8–37.

Wellman, M.; Wurman, P.; O'Malley, K.; Banger, R.; Lin, S.; Reeves, D.; and Walsh, W. 2001. A Trading Agent Competition. *IEEE Internet Computing*.

Wellman, M. P.; Greenwald, A.; and Stone, P. 2007. *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*. MIT Press.

Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition. MIT Press.