

Heuristics for the Deterministic Bidding Problem

AMY GREENWALD and VICTOR NARODITSKIY

Brown University

We study a suite of heuristics that were designed for bidding in the simultaneous auctions that characterize the Trading Agent Competition (TAC) Travel Game. At a high-level, the design of many successful TAC agents can be summarized as: (i) *predict*: build a model of the auctions' clearing prices, and (ii) *optimize*: solve for an (approximately) optimal set of bids, given this model. We focus on the optimization piece of this design.

Categories and Subject Descriptors: I.2.11 [**Artificial Intelligence**]: Intelligent agents

Additional Key Words and Phrases: simultaneous auctions, bidding heuristics, trading agents

1. INTRODUCTION

Simultaneous auctions, which arise naturally on websites such as `ebay.com` and `amazon.com`, are forums on which to trade many goods simultaneously. Such auctions present a challenge to bidders, particularly when complementary and substitutable goods are on sale. Complementary goods are goods whose values are superadditive: *i.e.*, for goods x and y , $v(x) + v(y) \leq v(xy)$. For example, a flash, a tripod, and a case complement a camera, since an agent does not desire any of the former if she does not acquire the latter. Substitutable goods are goods whose values are subadditive: *i.e.*, for goods x and y , $v(x) + v(y) \geq v(xy)$. For example, a Canon and an Olympus are substitutes, since an agent desires one or the other, but not both.

In contrast to *combinatorial auctions*, in which bids may be placed for combinations of goods (e.g., “camera and flash for \$295”), in simultaneous auctions, separate bids are placed for each individual good. In combinatorial auctions, the NP-hard problem of choosing a set of winning bids that maximizes revenue—the so-called winner determination problem—falls in the hands of the auctioneer. In simultaneous auctions, however, the complexity burden falls upon the bidders.

In this paper, we study heuristics that were designed for bidding in the simultaneous auctions that characterize the Trading Agent Competition (TAC) Travel Game [10]. 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 travel goods that satisfy its clients' preferences as inexpensively as possible. Travel goods are sold in simultaneous auctions, as follows:

- flights are sold by the “TAC seller” in dynamic posted-pricing environments; no resale is permitted
- hotel reservations are 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 are not useful to a client without the complementary hotel reservations; nor are hotel reservations useful to a client without the complementary flights. Tickets to entertainment events, e.g., the Boston Red Sox and the Boston Symphony Orchestra, are substitutable. Similarly, travel packages themselves are substitutes: e.g., arriving on Monday and departing on Tuesday vs. arriving on Monday and departing on Wednesday.

At a high-level, the design of many successful TAC agents (for example, Wolverine [1], RoxyBot [5], and ATTac [9]) can be summarized as:

- (1) *predict*: build a model of the auctions' clearing prices
- (2) *optimize*: solve for an (approximately) optimal set of bids, given this model

This paper is devoted to the study of the optimization piece of this design, which we model as the problem of bidding in “pseudo-auctions.” We define a pseudo-auction as an idealized auction setting in which (i) there is only one bidder, and (ii) prices are specified by an exogenous model, that is, a model in which the bidder's price predictions are independent of its bidding strategy. Given such a model, the bidder faces the *bidding problem*: what is its utility-maximizing set of bids?

In the present paper, we assume the agent builds a *deterministic* model of the auctions' prices: that is, there is no noise in the agent's price predictions; rather its predictions are point estimates. This assumption gives rise to the *deterministic bidding problem*. In our analysis, we focus on the deterministic *second-price* bidding problem, in which the payment rule is: “pay the predicted price.”¹ This problem is an abstraction of the problem of bidding in TAC Travel auctions, in which—under appropriate assumptions—agents can be viewed as price-takers.

Analytically, we study a set of heuristics in the context of a (single-unit) decision-theoretic bidding problem. Specifically, we derive the class of bidding heuristics that solves the deterministic second-price bidding problem optimally. We also prove that the marginal-value-based bidding heuristic implemented in RoxyBot-2000 [5], and RoxyBot-2000*, a slight variant of RoxyBot-2000, are both instances of this class. The classic marginal value bidding heuristic itself, however, fails to solve this problem optimally in general, as noted in Greenwald and Boyan [4]. Nonetheless, we identify the special set of circumstances in which bidding marginal values is optimal.

Experimentally (see Greenwald *et al.* [6]), we embed these heuristics in TAC Travel agents and evaluate their success in the (multi-unit) game-theoretic bidding problem that characterizes TAC Travel flight and hotel auctions.² We find that RoxyBot-2000's bidding heuristic dominates the others in our test set. Based on both our analytical and experimental results, we conclude that RoxyBot-2000's bidding heuristic is effective in that it performs well in a decision-theoretic setting when prices are given—equivalently, under the assumption of perfect price prediction—and it performs well in a game-theoretic setting where price predictions are imperfect.

¹For a cursory treatment of the deterministic first-price bidding problem, see Greenwald *et al.* [6], Appendix A.

²To reduce variance, we disregard entertainment tickets in our experiments.

It is not our contention that solutions to the decision-theoretic bidding problem are generally applicable as solutions to the game-theoretic bidding problem—the contrast between the conclusions of our analytical and experimental studies rule out this possibility.³ We do believe, however, that the study of a related decision-theoretic bidding problem can inform the design of artificially intelligent agents that face a game-theoretic bidding problem.

2. BIDDING IN SIMULTANEOUS PSEUDO-AUCTIONS

In this paper, we study a suite of heuristics for bidding in simultaneous auctions for complementary and substitutable goods. We develop these heuristics in the context of an optimization problem faced by an agent that is bidding in simultaneous “pseudo-auctions.”

There are two defining features of pseudo-auctions: (i) there is only one bidder, and (ii) the auction’s prices are specified (i.e., predicted) by an exogenous model, that is, a model in which the bidder’s price predictions are independent of its bidding strategy.

The winner determination rule in a pseudo-auction is: “win by bidding at least the predicted price.” In a first-price pseudo-auction the payment rule is “pay the winning bid price,” whereas in a second-price pseudo-auction the payment rule is “pay the predicted price.”

Implicit in our definitions of the first-price (see Greenwald *et al.* [6], Appendix A) and second-price bidding problems is the assumption that these problems characterize the optimization problem faced by an agent bidding in first- and second-price “sealed-bid” pseudo-auctions.

Given an instance of a bidding problem, a *bidding heuristic* searches for a suitable $b \in \mathbb{R}^X$, that is, a function from a set of goods X to bids $b(x) \in \mathbb{R}$ (equivalently,⁴ a bid vector).

The extension of a real-valued function $q : X \rightarrow \mathbb{R}$ on goods to a real-valued function $\tilde{q} : 2^X \rightarrow \mathbb{R}$ on bundles (i.e., sets of goods) is called *linear* if and only if $\tilde{q}(Y) = \sum_{x \in Y} q(x)$ for all $Y \subseteq X$.

DEFINITION 2.1. [Pseudo-Auction Winner Determination Rule] Given a set of goods X and a pricing function $p : X \rightarrow \mathbb{R}$, $\text{Winnings}(X, p, b) \subseteq X$ is the set of goods the agent wins by bidding $b \in \mathbb{R}^X$: i.e.,

$$x \in \text{Winnings}(X, p, b) \text{ if and only if } b(x) \geq p(x) \quad (1)$$

DEFINITION 2.2. Given a set of goods X , a valuation function $v : 2^X \rightarrow \mathbb{R}$, and a distribution f over pricing functions $p : X \rightarrow \mathbb{R}$, the *second-price bidding problem* is defined as follows:

$$\text{2ndSIM}(X, v, f) = \max_{b \in \mathbb{R}^X} \mathbb{E}_{p \sim f} [v(\text{Winnings}(X, p, b)) - \tilde{p}(\text{Winnings}(X, p, b))] \quad (2)$$

The deterministic second-price bidding problem is the special case of the second-price bidding problem in which the distribution f is a Dirac δ function: i.e., a

³RoxyBot-2000* is optimal in our analytic framework, but is suboptimal in our experimental framework.

⁴If Z is finite, $\mathbb{R}^Z = \{f : Z \rightarrow \mathbb{R}\}$ is isomorphic to $\mathbb{R}^{|Z|}$.

distribution with all its mass at a single point. In other words, prices are certain, not uncertain.

DEFINITION 2.3. Given a set of goods X , a valuation function $v : 2^X \rightarrow \mathbb{R}$, and a pricing function $p : X \rightarrow \mathbb{R}$, the *deterministic second-price bidding problem* is defined as follows:

$$\text{2ndDET}(X, v, p) = \max_{b \in \mathbb{R}^X} [v(\text{Winnings}(X, p, b)) - \tilde{p}(\text{Winnings}(X, p, b))] \quad (3)$$

Since prices are fully specified in the deterministic bidding problem, the key decision an agent faces is which goods to buy. But the problem of deciding which goods to buy is “the acquisition problem” [3]. Indeed, the deterministic bidding problem reduces to the acquisition problem.

DEFINITION 2.4. 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 *acquisition problem* is defined as follows:

$$\text{ACQ}(X, v, q) = \max_{Y \subseteq X} (v(Y) - q(Y)) \quad (4)$$

Given a solution to the acquisition problem, to solve the deterministic second-price bidding problem, an agent can bid an amount greater than or equal to the specified price on each good it wishes to acquire and an amount strictly less than the specified price on any good it does not desire.

THEOREM 2.5. *Given a set of goods X , a valuation function $v : 2^X \rightarrow \mathbb{R}$, and a pricing function $p : X \rightarrow \mathbb{R}$, the following bidding heuristic, which returns $b^* \in \mathbb{R}^X$, solves $\text{2ndDET}(X, v, p)$ optimally:*

- (1) select an optimal acquisition $A^* \in \arg \text{ACQ}(X, v, \tilde{p})$
- (2) bid

$$b^*(x) \in \begin{cases} [p(x), \infty) & \text{if } x \in A^* \\ (-\infty, p(x)) & \text{otherwise} \end{cases} \quad (5)$$

In particular, first solving for an optimal acquisition A^* and then bidding $p(x)$ on all goods $x \in A^*$ and bidding $-\infty$ otherwise is an optimal heuristic in the deterministic second-price bidding problem. This heuristic is also optimal in the deterministic first-price bidding problem (see Greenwald *et al.* [6], Appendix A).

3. AN ANALYSIS OF MARGINAL VALUES

In the preceding section, we derived an optimal class of bidding heuristics for the deterministic bidding problem. If an agent is clairvoyant—i.e., if it can predict the auctions’ clearing prices perfectly—then it can (and should) bid using any heuristic in this class, whenever the acquisition problem is computationally feasible. The TAC Travel acquisition problem, for one, is NP-hard [3]. Even more egregious, typical agents are not clairvoyant. Hence, it may be reasonable for agents to employ alternative bidding heuristics.

In this section, we broaden our study of bidding heuristics by investigating two classic strategies: bidding independent and marginal values. We illustrate the performance of these heuristics (and some of the complexity of bidding in simultaneous

auctions on complementary and substitutable goods) on a series of numerical examples. Also in this section, an analysis of marginal values leads to a “characterization theorem,” which completely characterizes the relationship between marginal values and prices, assuming linear prices.

3.1 Bidding Independent and Marginal Values

One straightforward bidding heuristic is Heuristic IV, for independent values: For each good in each auction, bid its independent value. Unfortunately, with heuristic IV, an agent can fail to win goods it wishes it had won, when goods are complements, and can succeed at winning goods it wishes it had not won, when goods are substitutes.

DEFINITION 3.1. Given a set of goods X and a valuation function $v : 2^X \rightarrow \mathbb{R}$, the *independent value* of a good $x \in X$ is given by: $\iota(x) = v(\{x\})$.

EXAMPLE 3.2. Suppose an agent values a camera and flash together at 500, but values either good alone at 1. Also, suppose these two goods are sold separately in two simultaneous auctions, and the clearing prices are 200 for the camera and 100 for the flash. If the agent were to bid only its independent values (1), it would lose both goods, obtaining utility of 0 rather than $500 - 200 - 100 = 200$. This outcome is undesirable: the agent fails to win goods it wishes it had won. \square

EXAMPLE 3.3. Now suppose an agent values a Canon AE-1 at 300 and a Canon A-1 at 200, but values both cameras together at only 400. Also, suppose these two goods are sold separately in two simultaneous auctions, and the clearing prices are 275 for the AE-1 and 175 for the A-1. If the agent were to bid its independent values, it would win both goods, obtaining utility of $400 - 450 = -50$. This outcome is undesirable: the agent wins goods it wishes it had not won. \square

A natural alternative to Heuristic IV is Heuristic MV, for marginal value: For each good, bid its *marginal value*. Unfortunately, even with heuristic MV, an agent can succeed at winning goods it wishes it had not won—in particular, when goods are substitutes—although an MV agent never fails to win goods it wishes it had won (see Theorem 3.8).

DEFINITION 3.4. 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 value* $\mu(x) \equiv \mu(x, X, v, q)$ of good $x \in X$ is defined as follows:

$$\mu(x) = \max_{Y \subseteq X \setminus \{x\}} [v(Y \cup \{x\}) - q(Y)] - \max_{Y \subseteq X \setminus \{x\}} [v(Y) - q(Y)] \quad (6)$$

Intuitively, the marginal value of x is simply the difference between the value of an optimal acquisition, assuming x costs 0, and the value of an optimal acquisition, assuming x costs ∞ .

EXAMPLE 3.5. Consider once again the setup of Example 3.2. Given both the camera and flash together, the agent’s value is 500; but either one of these components without the other is valued at only 1. If the clearing prices of the camera and flash are 200 and 100, respectively, then bidding according to MV, $(500 - 100) - (0 - 0) = 400$ on the camera and $(500 - 200) - (0 - 0) = 300$ on the flash, the agent wins both goods, as desired. \square

EXAMPLE 3.6. Consider once again the setup of Example 3.3, where an agent values a Canon AE-1 at 300 and a Canon A-1 at 200, and both cameras together at 400. If the clearing prices of the camera and flash are 275 and 175, respectively, then bidding according to MV, $(300 - 0) - (200 - 175) = 275$ on the camera and $(200 - 0) - (300 - 275) = 175$ on the flash, the agent wins both goods. As in Example 3.3, this is not the desired outcome: the agent win goods it wishes it had not won. \square

Example 3.5 shows that, for complementary goods, the MV heuristic can be an effective means of solving the deterministic bidding problem, in spite of the well-documented *exposure* problem. Specifically, an agent suffers from an exposure problem if it bids more on a good than its independent value of that good [7; 8]. As noted here, an agent can also suffer from an exposure problem if it bids more on a set of goods than its combinatorial value of that set of goods. Indeed, in Example 3.6, which concerns substitutable goods, the MV heuristic suffers from its exposure. To further illustrate this point, we reproduce the following example from Greenwald and Boyan [4], which shows that MV's performance on the second-price deterministic bidding problem can be arbitrarily bad.

EXAMPLE 3.7. Consider a set of $N > 1$ goods that are up for auction simultaneously. Assume that the agent attributes the value 2 to one or more of these goods and that the price of each good is 1. Bidding marginal values amounts to bidding 1 on each good. In the worst case, the MV heuristic obtains utility $2 - N < 1$. In contrast, any heuristic that bids 1 on exactly one good obtains utility $2 - 1 = 1$. \square

Although the marginal value bidding heuristic does not solve the deterministic second-price bidding problem optimally in general, in what follows, we derive the special set of circumstances in which bidding marginal values is optimal. Our derivation follows from Theorem 2.5 and an immediate corollary of our “characterization theorem,” which completely characterizes the relationship between marginal values and prices, assuming linear prices.

3.2 Characterization Theorem

Throughout this section, we assume we are given a set of goods X , a valuation function $v : 2^X \rightarrow \mathbb{R}$, and a pricing function $p : X \rightarrow \mathbb{R}$.

Our main theorem, which generalizes Greenwald [2], states the following: if x is contained in an optimal acquisition, then either x is contained in all optimal acquisitions, in which case its marginal value is strictly greater than its price, or x is not contained in all optimal acquisitions, in which case its marginal value is exactly equal to its price (as in Examples 3.6 and 3.7); otherwise, if x is not contained in any optimal acquisition, then its marginal value is strictly less than its price.

THEOREM 3.8. *Assume prices are linear. If $A_1^*, \dots, A_n^* \subseteq X$ are all the optimal solutions to the acquisition problem $ACQ(X, v, \tilde{p})$, then for all goods $x \in X$,*

- (1) $\mu(x) > p(x)$ if and only if $x \in \bigcap_{i=1}^n A_i^*$
- (2) $\mu(x) = p(x)$ if and only if $x \in \bigcup_{i=1}^n A_i^*$ but $x \notin \bigcap_{i=1}^n A_i^*$
- (3) $\mu(x) < p(x)$ if and only if $x \notin \bigcup_{i=1}^n A_i^*$

Theorem 3.8 holds under the assumption that prices are linear: i.e., $\tilde{p}(Y \cup \{x\}) = \tilde{p}(Y) + p(x)$. Without this assumption, this theorem is no longer valid, as we show in these two counterexamples:

- Suppose $\tilde{p}(Y \cup \{x\}) < \tilde{p}(Y) + p(x)$. Imagine two goods a and b , with $v(\{a\}) = v(\{b\}) = 0$, $v(\{a, b\}) = 16$, $p(\{a\}) = p(\{b\}) = 10$, and $p(\{a, b\}) = 15$. Both goods are components of the optimal acquisition. However, the marginal value of each good is less than its price: $\mu(a) = \mu(b) = (16 - 10) - 0 = 6$.
- Suppose $\tilde{p}(Y \cup \{x\}) > \tilde{p}(Y) + p(x)$. Imagine two goods a and b , with $v(\{a\}) = v(\{b\}) = 0$, $v(\{a, b\}) = 25$, $p(\{a\}) = p(\{b\}) = 10$, and $p(\{a, b\}) = 30$. The optimal acquisition is the empty set. However, the marginal value of each good is greater than its price: $\mu(a) = \mu(b) = (25 - 10) - 0 = 15$.

When the optimal acquisition is unique, the marginal value of a good is strictly greater than its price if and only if the good is in the optimal acquisition; otherwise, the marginal value of the good is strictly less than its price. This corollary of Theorem 3.8 is immediate.

COROLLARY 3.9. *Assume prices are linear. If $A^* \subseteq X$ is the unique solution to the acquisition problem $ACQ(X, v, \tilde{p})$, then*

- (1) $\mu(x) > p(x)$ if and only if $x \in A^*$, and
- (2) $\mu(x) < p(x)$ if and only if $x \notin A^*$.

Finally, we characterize the relationship between $\mu(x, A^*)$ and $\mu(x, X)$, that is, the marginal utility of a good x relative to an optimal acquisition A^* vs. the marginal utility of a good x relative to the set of all goods X . Intuitively, the marginal value of a good that is in an optimal acquisition cannot decrease if the set of available goods is restricted to include precisely the goods in that acquisition. Analogously, the marginal value of a good that is not in an optimal acquisition cannot increase if the set of available goods is restricted to include precisely the goods in that acquisition.

PROPOSITION 3.10. *If $A^* \subseteq X$ is an optimal solution to the acquisition problem $ACQ(X, v, q)$, then*

- $\mu(x, A^*) \geq \mu(x, X)$, for all $x \in A^*$, and
- $\mu(x, A^*) \leq \mu(x, X)$, for all $x \notin A^*$ (i.e., $x \in X \setminus A^*$),

where $q : 2^X \rightarrow \mathbb{R}$ is an arbitrary pricing function.

4. A TEST SUITE OF BIDDING HEURISTICS

We now articulate the inner workings of four select bidding heuristics by presenting their pseudocode. **StraightMV** is an implementation of the marginal value bidding heuristic. **SecondBot** generalizes the class of bidding heuristics that solves the deterministic second-price bidding problem optimally. **FirstBot**, a bidding heuristic that solves the deterministic *first*-price bidding problem optimally, **RoxyBot-2000**, and a slight variant, **RoxyBot-2000***, are all instances of **SecondBot**.

We argue that **FirstBot**, **RoxyBot-2000**, and **RoxyBot-2000*** all solve the deterministic second-price bidding problem optimally, assuming linear prices. We also

establish that **StraightMV** is optimal whenever the solution to the acquisition problem is unique, again, assuming prices are linear.

Also in this section, we work through an example of the deterministic second-price bidding problem, and compare the performance of these four heuristics on this problem *without assuming clairvoyance*—that is, the agents optimize with respect to imperfect price predictions.

StraightMV (see Algorithm 1) is an implementation of the classic marginal value bidding heuristic. It bids the marginal value of each good in each auction, given as input predictions of the auctions' clearing prices. **StraightMV** calculates $|X|$ marginal values; hence, it solves $2|X|$ acquisition problems.

Algorithm 1 StraightMV(X, v, p)

```

1: for  $j = 1$  to  $|X|$  do
2:    $\text{bids}_j \leftarrow \mu(x_j, X, v, p)$ 
3: end for
4: return bids

```

THEOREM 4.1. *Bidding marginal values is optimal in the deterministic second-price bidding problem whenever the solution to the acquisition problem is unique, assuming prices are linear.*

PROOF. The proof follows immediately from Theorem 2.5 and Corollary 3.9. \square

EXAMPLE 4.2. Modifying the setup of Example 3.3 once again, suppose an agent values a Canon AE-1 at 305, a Canon A-1 at 200, and both cameras together at 400. Also suppose the agent predicts the clearing price of the Canon AE-1 and the Canon A-1 to be 275 and 175, respectively, while in reality, the clearing prices of the cameras are uniformly distributed in the ranges [265, 285] and [165, 185].

Given these predictions, the marginal value of the Canon AE-1 is $(305 - 0) - (200 - 175) = 280$, while the marginal value of the Canon A-1 is $(200 - 0) - (305 - 275) = 170$. **StraightMV** bids precisely these marginal values: 280 on the Canon AE-1 and 170 on the Canon A-1. But by bidding marginal values, a **StraightMV** agent is likely to win either too many substitutes or too few complements.

Specifically, the probability of winning the Canon AE-1 is $\frac{280-265}{285-265} = 0.75$, while the probability of winning the Canon A-1 is $\frac{170-165}{185-165} = 0.25$. Consequently, the probability of winning too many substitutes (both cameras) is $(.75)(.25) = .1875$, as is the probability of winning too few complements (neither camera). \square

Whereas **StraightMV** can win too many substitutes assuming either perfect or imperfect price prediction, none of the following three heuristics ever wins too many substitutes. Moreover, whereas **StraightMV** can win too few complements assuming imperfect price prediction, in turn, each of the next three heuristics wins more and more complements.

SecondBot **SecondBot** (see Algorithm 2) first solves for an optimal acquisition $A^* \in \arg \text{ACQ}(X, v, \tilde{p})$, and then bids on the goods in A^* according to some function

g . We study three instances of **SecondBot**, corresponding to three choices of the bid function $g(x) \equiv g(x, X, v, p, A^*)$.

- FirstBot**: $g = p$
- RoxyBot-2000**: $g = h$ where $h(x) = \mu(x, X, v, p)$, for $x \in X$
- RoxyBot-2000***: $g = h^*$ where $h^*(x) = \mu(x, A^*, v, p)$, for $x \in X$

Note the distinction between **RoxyBot-2000** and **RoxyBot-2000***: the former calculates marginal values with respect to the set of goods X , whereas the latter calculates marginal values with respect to the optimal acquisition A^* .

These three instances of **SecondBot** place progressively higher and higher bids:

- FirstBot** bids $p(x)$ on all goods $x \in A^*$
- RoxyBot-2000** bids $\mu(x, X, v, p) \geq p(x)$ on all goods $x \in A^*$ (by Theorem 3.8)
- RoxyBot-2000*** bids $\mu(x, A^*, v, p) \geq \mu(x, X, v, p)$ on all goods $x \in A^*$ (by Proposition 3.10)

FirstBot solves only 1 acquisition problem. In the worst case (when $A^* = X$), these versions of **RoxyBot** calculate $|X|$ marginal values, solving $2|X| + 1$ acquisition problems in total. In practice (e.g., in TAC Travel games), however, they calculate far fewer marginal values than **StraightMV**.

THEOREM 4.3. *FirstBot, RoxyBot-2000, and RoxyBot-2000* are optimal bidding heuristics in the deterministic second-price bidding problem, assuming prices are linear.*

PROOF. The proof follows immediately from Theorem 2.5, Theorem 3.8, and Proposition 3.10. \square

Algorithm 2 **SecondBot**(X, v, g, p)

```

1:  $A^* \in \arg \text{ACQ}(X, v, \tilde{p})$ 
2: for  $j = 1$  to  $|X|$  do
3:   if  $x_j \in A^*$  then
4:      $\text{bids}_j \leftarrow g(x_j)$ 
5:   else
6:      $\text{bids}_j \leftarrow 0$ 
7:   end if
8: end for
9: return bids
    
```

EXAMPLE 4.4. Since **SecondBot** bids only on the goods in a single acquisition, it cannot win too many substitutes, but still it may win too few complements. Continuing Example 4.2, the values of only the Canon AE-1 or only the Canon A-1 are $305 - 275 = 30$ and $200 - 175 = 25$, respectively. Hence, the Canon AE-1 alone is the unique optimal acquisition.

FirstBot bids the predicted price (275) on the Canon AE-1 and nothing on the Canon A-1, which yields a 50% chance of winning too few complements (i.e., losing both cameras). **RoxyBot-2000** bids its marginal value (280) on the Canon AE-1

and nothing on the Canon A-1, which yields a 25% chance of winning too few complements (i.e., losing both cameras).

RoxyBot-2000* assumes the Canon A-1 is not available. Under this assumption, the marginal value of the Canon AE-1 is $305 - 0 = 305$. This is the only bid RoxyBot-2000* submits. Since $305 \geq 285$, the upper bound on the clearing price of the Canon AE-1, RoxyBot-2000* wins neither too many substitutes nor too few complements in this example. \square

5. CONCLUSION

Based on our analytical and related experimental results (see Greenwald *et al.* [6]), we conclude that RoxyBot-2000's bidding heuristic is effective in that it performs well in a decision-theoretic setting when prices are given—equivalently, under the assumption of perfect price prediction—and it performs well in a game-theoretic setting where price predictions are imperfect. However, this bidding heuristic, which operates on (deterministic) price point estimates, does not explicitly plan for uncertainty in the auction dynamics. A heuristic that would be superior in this respect would optimize with respect to noisy (i.e., stochastic) models of estimated clearing prices. Indeed, embedded in RoxyBot-2006, the top-scoring agent in TAC-2006, is such a bidding heuristic.

6. ACKNOWLEDGMENTS

This research was supported by NSF Career Grant #IIS-0133689.

REFERENCES

- S. Cheng, E. Leung, K. Lochner, K.O'Malley, D. Reeves, L. Schwartzman, and M. Wellman. Walverine: A Walrasian trading agent. *Decision Support Systems*, page To Appear, 2004.
- A. Greenwald. Bidding marginal utility in simultaneous auctions. In *International Joint Conference on Artificial Intelligence Workshop on Trading Agent Design and Analysis*, August 2003.
- A. Greenwald. Bid determination in simultaneous auctions. Technical Report 16, Brown University, 2005.
- A. Greenwald and J. Boyan. Bidding under uncertainty: Theory and experiments. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, pages 209–216, July 2004.
- A. Greenwald and J. Boyan. Bidding algorithms for simultaneous auctions: A case study. *Journal of Autonomous Agents and Multiagent Systems*, 10(1):67–89, 2005.
- A. Greenwald, V. Naroditskiy, and S. Lee. Heuristics for the deterministic bidding problem: Lessons from tac travel. Technical Report 11, Brown University, 2006.
- J. Ledyard, D. Porter, and A. Rangel. Experiments testing multiobject allocation mechanisms. *Journal of Economics and Management Strategy*, 6:639–665, 1997.
- C. Plott. Experiments testing multiobject allocation mechanisms. *Journal of Economics and Management Strategy*, 6:605–638, 1997.
- P. Stone, R. Schapire, M. Littman, J. Csirik, and D. McAllester. Decision-theoretic bidding based on learned density models in simultaneous, interacting auctions. *Journal of Artificial Intelligence Research*, 19:209–242, 2003.
- M. Wellman, P. Wurman, K. O'Malley, R. Bangera, S. Lin, D. Reeves, and W. Walsh. A Trading Agent Competition. *IEEE Internet Computing*, April 2001.