# Heuristics for the Deterministic Bidding Problem: Lessons from TAC Travel

Amy Greenwald, Victor Naroditskiy and Seong Jae Lee

Department of Computer Science
Brown University
Providence, Rhode Island 02912

**CS-06-15**
June 2006

# Heuristics for the Deterministic Bidding Problem: Lessons from TAC Travel

**Amy Greenwald, Victor Naroditskiy, and Seong Jae Lee**
*Department of Computer Science, Brown University, Providence, RI 02912*
{amy,vnarodit,sjlee}@cs.brown.edu

## Abstract

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; (ii) *optimize*: solve for an (approximately) optimal set of bids, given this model. We focus on the optimization piece of this design.

Analytically, we address the (decision-theoretic) deterministic bidding problem. We derive the class of bidding heuristics that solves this problem optimally. In particular, we prove that the marginal-value-based bidding heuristic implemented in RoxyBot–2000—one of the top-scoring agents in TAC 2000—is an instance of this class. Moreover, we identify the special set of circumstances in which bidding marginal values themselves is also optimal.

Experimentally, we embed these heuristics in TAC Travel agents and evaluate their success in the game-theoretic bidding problem that characterizes TAC Travel auctions. We find that RoxyBot–2000's bidding heuristic dominates the others in our test set. We conclude that RoxyBot–2000's bidding heuristic is effective in that it performs well in a decision-theoretic setting assuming perfect price prediction and a game-theoretic setting where price predictions are imperfect.

**Keywords:** 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 [12]. 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 16*th* 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, Walverine [1], RoxyBot [6], and ATTac [10]) 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 [6], 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 [5]. Nonetheless, we identify the special set of circumstances in which bidding marginal values is optimal.

Experimentally, 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.

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.

1. For a cursory treatment of the deterministic first-price bidding problem, see Appendix A.
2. To reduce variance, we disregard entertainment tickets in our experiments.
3. RoxyBot–2000* is optimal in our analytic framework, but is suboptimal in our experimental framework.

## 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 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 $q : X \to \mathbb{R}$ on goods to a real-valued function $\tilde{q} : 2^X \to \mathbb{R}$ 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 $p : X \to \mathbb{R}$, $\text{Winnings}(X, p, b) \subseteq X$ is the set of goods the agent wins by bidding $b \in \mathbb{R}^X$: i.e.,

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

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

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

The deterministic second-price bidding problem is the special case of the second-price bidding problem in which the distribution $f$ is a Dirac $\delta$ function: i.e., a 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 $v : 2^X \to \mathbb{R}$, and a pricing function $p : X \to \mathbb{R}$, the *deterministic second-price bidding problem* is defined as follows:

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

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" [4]. Indeed, the deterministic bidding problem reduces to the acquisition problem.

**Definition 2.4** Given a set of goods $X$, a valuation function $v : 2^X \to \mathbb{R}$, and a pricing function $q : 2^X \to \mathbb{R}$, the *acquisition problem* is defined as follows:

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

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.

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

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

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

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

**Proof:** It suffices to show that the function $b^* \in \mathbb{R}^X$ maximizes the function $\text{2ndDET}(X, v, p, b) = v(\text{Winnings}(X, p, b)) - \tilde{p}(\text{Winnings}(X, p, b))$. For all $b \in \mathbb{R}^X$,

$$\begin{aligned} \text{2ndDET}(X, v, p, b^*) &= v(\text{Winnings}(X, p, b^*)) - \tilde{p}(\text{Winnings}(X, p, b^*)) & (6) \\ &= v(A^*) - \tilde{p}(A^*) & (7) \\ &\geq v(A) - \tilde{p}(A) & (8) \\ &= v(\text{Winnings}(X, p, b)) - \tilde{p}(\text{Winnings}(X, p, b)) & (9) \\ &= \text{2ndDET}(X, v, p, b) & (10) \end{aligned}$$

Equation 7 follows from the definition of $b^*$ and the winner determination rule. Equation 8 follows from the fact that the acquisition $A^*$ is optimal. ■

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 Appendix A).

## 3. An Analysis of Marginal Values

In the preceding section, we derived an optimal class of bidding heuristics for the deterministic second-price 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 [4]. 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 on a series of examples in which *we imagine the agent is clairvoyant* (equivalently, we assume posted prices). We observe that they do not solve the deterministic second-price bidding problem optimally. In later sections, when we drop the clairvoyance assumption, we find that the marginal value bidding heuristic outperforms some bidding heuristics that solve this problem optimally.

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

Perhaps the most 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 $v : 2^X \to \mathbb{R}$, 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. □

**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 also undesirable: the agent wins goods it wishes it had not won. □

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 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 $v: 2^X \to \mathbb{R}$, and a pricing function $q: 2^X \to \mathbb{R}$. The *marginal value* $\mu(x) \equiv \mu(x, X, v, q)$ of good $x \in X$ is defined as follows:

$$\mu(x) = \max_{Y \subseteq X \setminus \{x\}} [v(Y \cup \{x\}) - q(Y)] - \max_{Y \subseteq X \setminus \{x\}} [v(Y) - q(Y)] \tag{11}$$

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 $\infty$.

**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. □

**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. □

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. An agent suffers from its exposure if it bids more on a good than its independent value of that good [7, 9]. As noted here, an agent can also suffer from another kind of exposure in which 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 [5], 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 $2-1=1$. □

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 $v : 2^X \to \mathbb{R}$, and a pricing function $p : X \to \mathbb{R}$.

Our main theorem, which generalizes Greenwald [3], 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 $A_1^*, \ldots, A_n^* \subseteq X$ 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 $x \notin \bigcup_{i=1}^n A_i^*$*

Our proof of Theorem 3.8 relies on the following observation:

**Observation 3.9** *Assume prices are linear. The following equalities are equivalent: for all $x \in X$,*

$$\begin{aligned} & \mu(x) = p(x) && (12)\\ \textit{iff} \quad & \max_{Y \subseteq X \setminus x} [v(Y \cup \{x\}) - \tilde{p}(Y)] - \max_{Y \subseteq X \setminus x} [v(Y) - \tilde{p}(Y)] = p(x) && (13)\\ \textit{iff} \quad & \max_{Y \subseteq X \setminus x} [v(Y \cup \{x\}) - \tilde{p}(Y \cup \{x\})] = \max_{Y \subseteq X \setminus x} [v(Y) - \tilde{p}(Y)] && (14) \end{aligned}$$

*The same holds when the equal signs are replaced by (weak or strict) inequality signs.*

**Proof of Theorem 3.8 [Claim 1]:** Generally speaking,

$$\text{ACQ}(X, v, \tilde{p}) = \max_{Y \subseteq X} [v(Y) - \tilde{p}(Y)] \geq \max_{Y \subseteq X \setminus x} [v(Y) - \tilde{p}(Y)] \tag{15}$$

because the set of feasible solutions is smaller on the right hand side of Equation 15. Moreover,

$$\text{ACQ}(X, v, \tilde{p}) = \max_{Y \subseteq X} [v(Y) - \tilde{p}(Y)] \geq \max_{Y \subseteq X \setminus x} [v(Y \cup \{x\}) - \tilde{p}(Y \cup \{x\})] \tag{16}$$

In Equation 16, it is as if the optimal solution to the acquisition problem is forced to include $x$.

[$\Longrightarrow$] If $x \notin \bigcap_{i=1}^n A_i^*$, then

$$\text{ACQ}(X, v, \tilde{p}) = \max_{Y \subseteq X \setminus x} [v(Y) - \tilde{p}(Y)]$$

because the maximum value can be attained without $x$. Hence, Equation 15 collapses, so that

$$\max_{Y \subseteq X \setminus x} [v(Y \cup \{x\}) - \tilde{p}(Y \cup \{x\})] \leq \max_{Y \subseteq X \setminus x} [v(Y) - \tilde{p}(Y)]$$

[$\Longleftarrow$] Since $x \in \bigcap_{i=1}^n A_i^*$, the relationship in Equation 15 is strict, because otherwise there would exist an optimal acquisition with $x$ and another without $x$: i.e.,

$$\mathrm{ACQ}(X, v, \tilde{p}) > \max_{Y \subseteq X \setminus x} [v(Y) - \tilde{p}(Y)]$$

Moreover, Equation 16 collapses because $x$ is in all optimal solutions (it is not a hindrance to include $x$ in an arbitrary optimal acquisition):

$$\mathrm{ACQ}(X, v, \tilde{p}) = \max_{Y \subseteq X \setminus x} [v(Y \cup \{x\}) - \tilde{p}(Y \cup \{x\})]$$

Hence,

$$\max_{Y \subseteq X \setminus x} [v(Y \cup \{x\}) - \tilde{p}(Y \cup \{x\})] > \max_{Y \subseteq X \setminus x} [v(Y) - \tilde{p}(Y)]$$

■

**Proof of Theorem 3.8 [Claim 2]:** For all $Y \subseteq X$, define the utility $u(Y) = v(Y) - \tilde{p}(Y)$.

[$\Longrightarrow$] By assumption, $\mu(x) = p(x)$. Hence, by Observation 3.9, $A_1 \cup \{x\}$ and $A_2$ have the same utility, where

$$A_1 \in \arg\max_{Y \subseteq X \setminus x} [v(Y \cup \{x\}) - \tilde{p}(Y \cup \{x\})] \tag{17}$$

$$A_2 \in \arg\max_{Y \subseteq X \setminus x} [v(Y) - \tilde{p}(Y)] \tag{18}$$

Note that $x \in A_1 \cup \{x\}$ but $x \notin A_2$. An optimal acquisition either contains $x$ or it does not. Hence, one of $A_1 \cup \{x\}$ or $A_2$ must be optimal, but since they have the same utility, both of them are optimal. We have constructed an optimal acquisition with $x$ and an optimal acquisition without $x$. Therefore, $x \in \bigcup_{i=1}^n A_i^*$ but $x \notin \bigcap_{i=1}^n A_i^*$.

[$\Longleftarrow$] By assumption, $x \in \bigcup_{i=1}^n A_i^*$ but $x \notin \bigcap_{i=1}^n A_i^*$. Hence, there exists $A_i^* \neq A_j^*$ such that $x \in A_i^*$ but $x \notin A_j^*$. Now

$$x \in A_i^* \quad \Rightarrow \quad u(A_i^*) = \max_{Y \subseteq X \setminus x} [v(Y \cup \{x\}) - \tilde{p}(Y \cup \{x\})] \tag{19}$$

$$x \notin A_j^* \quad \Rightarrow \quad u(A_j^*) = \max_{Y \subseteq X \setminus x} [v(Y) - \tilde{p}(Y)] \tag{20}$$

$$u(A_i^*) = u(A_j^*) \quad \text{iff} \quad \max_{Y \subseteq X \setminus x} [v(Y \cup \{x\}) - \tilde{p}(Y \cup \{x\})] = \max_{Y \subseteq X \setminus x} [v(Y) - \tilde{p}(Y)] \tag{21}$$

This last equation is equivalent to $\mu(x) = p(x)$ by Observation 3.9. ■

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 3.10** *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 acquistion. 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 acquistion.

**Proposition 3.11** *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 $q : 2^X \to \mathbb{R}$ is an arbitrary pricing function.*

**Proof:** For all $x \in A^*$,

$$\begin{aligned}
\mu(x, A^*) &= \max_{Y \subseteq A^* \setminus \{x\}} [v(Y \cup \{x\}) - q(Y)] - \max_{Y \subseteq A^* \setminus \{x\}} [v(Y) - q(Y)] && (22)\\
&= \max_{Y \subseteq X \setminus \{x\}} [v(Y \cup \{x\}) - q(Y)] - \max_{Y \subseteq A^* \setminus \{x\}} [v(Y) - q(Y)] && (23)\\
&\geq \max_{Y \subseteq X \setminus \{x\}} [v(Y \cup \{x\}) - q(Y)] - \max_{Y \subseteq X \setminus \{x\}} [v(Y) - q(Y)] && (24)\\
&= \mu(x, X) && (25)
\end{aligned}$$

Equation 23 follows from the fact that $A^*$ is optimal and $x \in A^*$. Equation 24 follows from the fact that $A^* \subseteq X$ so that $\max_{Y \subseteq X \setminus \{x\}} f(Y) \geq \max_{Y \subseteq A^* \setminus \{x\}} f(Y)$, for all functions $f : 2^X \to \mathbb{R}$.

For all $x \notin A^*$,

$$\begin{aligned}
\mu(x, A^*) &= \max_{Y \subseteq A^* \setminus \{x\}} [v(Y \cup \{x\}) - q(Y)] - \max_{Y \subseteq A^* \setminus \{x\}} [v(Y) - q(Y)] && (26)\\
&= \max_{Y \subseteq A^* \setminus \{x\}} [v(Y \cup \{x\}) - q(Y)] - \max_{Y \subseteq X \setminus \{x\}} [v(Y) - q(Y)] && (27)\\
&\leq \max_{Y \subseteq X \setminus \{x\}} [v(Y \cup \{x\}) - q(Y)] - \max_{Y \subseteq X \setminus \{x\}} [v(Y) - q(Y)] && (28)\\
&= \mu(x, X) && (29)
\end{aligned}$$

Equation 27 follows from the fact that $A^*$ is optimal and $x \notin A^*$. Equation 28 follows from the fact that $A^* \subseteq X$ so that $\max_{Y \subseteq X \setminus \{x\}} f(Y) \geq \max_{Y \subseteq A^* \setminus \{x\}} f(Y)$, for all functions $f : 2^X \to \mathbb{R}$. ∎

## 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$(X, v, p)$

```
for j = 1 to |X| do
    bids_j ⇐ μ(x_j, X, v, p)
end for
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 3.10. ■

**Example 4.2** A (TAC) travel agent is deciding what to bid on hotels for a client for whom it has already purchased flights. The client's willingness to pay for travel packages including the good and bad hotel, respectively, are 1055 and 1000. Flights alone are worthless to the client.

Suppose the agent predicts the clearing price of the good hotel to be 80 and the clearing price of the bad hotel to be 30, while in reality, the clearing price of the good and bad hotels are uniformly distributed in the ranges $[70, 90]$ and $[20, 40]$, respectively.

Given its predictions, the marginal value of the good hotel is $(1055 - 0) - (1000 - 30) = 85$, while the marginal value of the bad hotel is $(1000 - 0) - (1055 - 80) = 25$. StraightMV bids precisely these marginal values: 85 on the good hotel and 25 on the bad hotel.

By bidding marginal values, a StraightMV agent is likely to win either too many substitutes or too few complements. The probability of winning the good hotel is $\frac{85-70}{90-70} = 0.75$, while the probability of winning the bad hotel is $\frac{25-20}{40-20} = 0.25$. Consequently, the probability of winning too many substitutes (both hotels) is $(.75)(.25) = .1875$, as is the probability of winning too few complements (neither hotel). □

Whereas StraightMV can win too many substitutes assuming either perfect or imperfection 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 $A^* \in \arg \mathrm{ACQ}(X, v, \tilde{p})$, 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 3.11)

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 3.11. ■

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

```
A* ∈ arg ACQ(X, v, p̃)
for j = 1 to |X| do
  if x_j ∈ A* then
    bids_j ⇐ g(x_j)
  else
    bids_j ⇐ 0
  end if
end for
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 the travel packages with the good and bad hotels are $1055 - 80 = 975$ and $1000 - 30 = 970$, respectively. Hence, the travel package with the good hotel is the unique optimal acquisition.

FirstBot bids the predicted price (80) on the good hotel and nothing on the bad hotel, which yields a 50% chance of winning too few complements (i.e., losing both hotels). RoxyBot–2000 bids its marginal value (85) on the good hotel and nothing on the bad hotel, which yields a 25% chance of winning too few complements (i.e., losing both hotels).

RoxyBot–2000* assumes the bad hotel is not available. Under this assumption, the marginal value of the good hotel is $1055 - 0 = 1055$. This is the only bid RoxyBot–2000* submits. Since $1055 \geq 90$, the upper bound on the clearing price of the good hotel, RoxyBot–2000* wins neither too many substitutes nor too few complements in this example. □

## 5. Bidding in TAC Travel Auctions

Having conducted an analytic study of the deterministic second-price bidding problem, and, in doing so, having developed a test suite of bidding heuristics, we now describe an experimental study in which we embedded these heuristics in TAC Travel agents and played TAC Travel games. There are three key differences between our analytical and experimental setups:

- the former is decision-theoretic, while the latter is game-theoretic
- the former is a single-unit second-price design, while the latter is a $k$-unit $k$th price design (NB: if $k = 1$, the latter is a first-price design rather than a second-price design)

- in the former, prices are given, which amounts to an assumption of perfect price prediction by an agent; in the latter, an agent's price predictions are likely to be imperfect

Still, we conducted these experiments to shed some light on the efficacy of our test suite of bidding heuristics, which are optimal or near-optimal in a decision-theoretic problem, in a related game-theoretic bidding problem. It is not our contention, however, that solutions to a decision-theoretic bidding problem are generally applicable as solutions to a game-theoretic bidding problem.

### 5.1 Multi-unit Auctions

TAC Travel auctions are multi-unit auctions. For bidding in multi-unit auctions, we extend the definition of marginal value, originally defined for a single copy of each good, to handle multiple copies of the same good. The marginal value of the first copy of a good is calculated assuming that no other copies of the good can be had; the marginal value 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 5.1** Given a set of goods $X$, a valuation function $v : 2^X \to \mathbb{R}$, and a pricing function $q : 2^X \to \mathbb{R}$. The *marginal value* $\mu(x_{jk}) \equiv \mu(x_{jk}, X, v, q)$ of the $k$th copy of good $j$ is:

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

In words, the marginal value of the $k$th copy of good $j$ is simply the difference between the value of an optimal acquisition, assuming $x_{j1},\ldots,x_{jk}$ cost 0 and $x_{j,k+1},\ldots,x_{jN}$ cost $\infty$, and the value of an optimal acquisition, assuming $x_{j1},\ldots,x_{j,k-1}$ cost 0 and $x_{jk},\ldots,x_{jN}$ cost $\infty$.

### 5.2 Experimental Setup

In our experiments, we embedded in TAC Travel agents our test suite of bidding heuristics, generalized to bid marginal values in multi-unit auctions. These agents played numerous TAC Travel games, bidding only in flight and hotel auctions. We disregarded entertainment ticket auctions to reduce the variance in the agents' scores. This feature of our experimental design was implemented by modifying the TAC–2004 Classic Java Server [2].

Recall that there are two key architectural components of TAC Travel agents, price prediction and optimization. In our experimental design, we vary the optimization component of the agents, but we fix the price prediction component. The tâtonnement process is a method of computing competitive equilibrium prices in a market [11]. All the agents in our experiments predict the hotel auctions' clearing prices via the tâtonnement process, just as it was implemented in Walverine–2003 [1].

### 5.3 Experimental Results

In this section, we report the results of two experiments with our test suite of bidding heuristics in TAC Travel games. For both experiments, we report the average score earned by each agent, along with the corresponding average utilities and costs. In addition, we report the results of paired $t$-tests, in which we consider all possible pairings of the agents' scores, utilities, and costs. In our statistical tests, the null hypothesis is: "there is no difference between the means."

**Four-agent Experiment** In our first set of experimental games (1200 of them), we pitted two copies of each of our four heuristics against one another—StraightMV, RoxyBot–2000, RoxyBot–2000*, and FirstBot. The agents' average scores, utilities, and costs in these games are listed in Table 1. In addition, 95% confidence intervals (for scores only) are plotted in Figure 1 (left). These intervals were computed based on 1200 independent observations; scores were averaged across agent type in each game to account for any game dependencies.

| Rank | Agent | Score | Utility | Cost |
|---|---|---|---|---|
| 1 | RoxyBot–2000* | 2848 | 8329 | 5481 |
| 2 | RoxyBot–2000 | 2842 | 8278 | 5436 |
| 3 | StraightMV | 2644 | 8321 | 5677 |
| 4 | FirstBot | 2372 | 7932 | 5560 |

Table 1: Four-agent Experiment: Average Scores, utilities, and costs for each bidding heuristic.

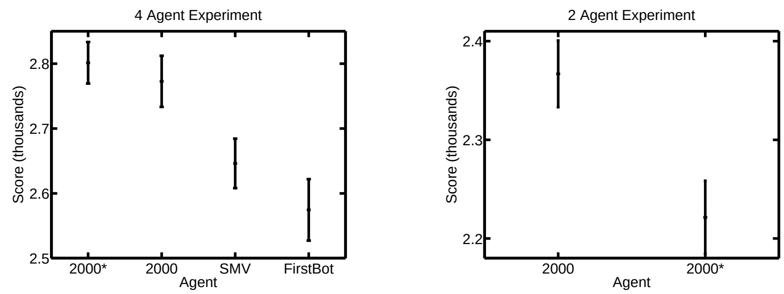


Figure 1: Four- and two-agent experiments: scores, with 95% confidence intervals.

Scorewise, RoxyBot–2000* ever so slightly outperforms RoxyBot–2000, which in turn outperforms StraightMV, which outperforms FirstBot. These latter differences are statistically significant (see Figure 1). It is interesting to note that StraightMV obtains a higher utility on average than RoxyBot–2000 ($\sim$ 40), but it does so at a substantially higher cost ($\sim$ 240). These differences are also statistically significant (not shown), and arise because StraightMV bids on all goods, not only the goods in a single optimal completion. Finally, FirstBot is unsuccessful because it places too many losing bids, foregoing too many essential hotels. It makes up for its losses by buying excess flights, driving up not only flight costs but travel penalties as well, since its ultimate allocation of travel packages to clients is not entirely satisfactory.

**Two-agent Experiment** In our second set of experimental games (2400 of them, providing 2400 independent observations as above), we pitted four copies of RoxyBot–2000 and RoxyBot–2000* against one another. The average scores, utilities, and costs in these games are depicted in Table 2. The differences between utilities and costs in this experiment are both statistically significant (not shown). Figure 1 (right) depicts 95% confidence intervals for scores only.

| Rank | Agent | Score | Utility | Cost |
|---|---|---|---|---|
| 1 | RoxyBot–2000 | 2367 | 8140 | 5773 |
| 2 | RoxyBot–2000* | 2236 | 8295 | 6059 |

Table 2: Two-agent experiment: Average scores, utilities, and costs for each bidding heuristic.

This time, RoxyBot–2000 outperforms RoxyBot–2000*. This outcome can be explained as follows. Embedded in both RoxyBot–2000 and RoxyBot–2000* are optimal bidding heuristics for the deterministic (second-price) bidding problem. TAC hotel bidding is not a second-price auction, however; it is a $k$th-price $k$-unit auction. As such, an agent's own bid for a good can determine the price of that good. Thus, the bidders have an incentive to shade their bids downward. But RoxyBot–2000* in effect shades its bids *upward*! In so doing, it wins more goods than RoxyBot–2000 so that it obtains a higher utility (by 155), but this increase is achieved at an even higher cost (286 more).

This two-agent experiment serves to highlight the distinction between the idealized deterministic bidding problem, and the actual TAC environment. Both RoxyBot–2000 and RoxyBot–2000* are optimal bidding heuristics in the former, but RoxyBot–2000 is empirically superior in the latter.

## 6. Discussion

A problem with marginal-value-based bidding heuristics is the bids they place can be incompatible with multi-unit auction mechanisms. Indeed marginal-value-based bidding is incompatible with the TAC Travel hotel auction mechanism! This incompatibility stems from the fact that the marginal values of additional copies of a good need not be diminishing when values are superadditive, as the following example demonstrates.

**Example 6.1** Several copies of two goods, $x$ and $y$, are up for auction. Imagine an agent bidding on behalf of two clients: the agent's valuation is the sum of the client's values.

The goods are perfect complements for client 1: $v_1(\{x, y\}) = V > 0$, while client 1 values all other subsets at zero. The goods are perfect substitutes for client 2: $v_2(\{x\}) = v_2(\{y\}) = V > 0$, while client 2 values all other subsets at zero. The agent owns one copy of $y$.

The marginal value of the first copy of $x$ is zero (one copy of $x$ does not change the value the agent can obtain), but the marginal value of the second copy of $x$ is strictly positive (a second copy of $x$ increase the agent's valuation from $V$ to $2V$).

The TAC Travel hotel auction mechanism processes bids of the form $(0, V)$ as a bid of $V$ on the first copy of a good and 0 on the second copy of a good. But, in this example, the agent is willing to pay $V$ for the second—not the first—copy of good $x$. $\square$

## 7. Conclusion

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. 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 [8], the top-scoring agent in TAC–2006, is such a bidding heuristic.

## 8. Acknowledgments

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

## References

[1] S.F. Cheng, E. Leung, K.M. Lochner, K.O'Malley, D.M. Reeves, L.J. Schvartzman, and M.P. Wellman. Walverine: A Walrasian trading agent. *Decision Support Systems*, 39(2):169–184, 2005.

[2] J. Eriksson, N. Finne, and S. Janson. The sics TAC classic server and agentware. In preparation, 2005.

[3] A. Greenwald. Bidding marginal utility in simultaneous auctions. In *International Joint Conference on Artificial Intelligence Workshop on Trading Agent Design and Analysis*, August 2003.

[4] A. Greenwald. Bid determination in simultaneous auctions. Technical Report 16, Brown University, 2005.

[5] 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.

[6] Amy Greenwald and Justin Boyan. Bidding algorithms for simultaneous auctions: A case study. *Journal of Autonomous Agents and Multiagent Systems*, 10(1):67–89, 2005.

[7] J. Ledyard, D. Porter, and A. Rangel. Experiments testing multiobject allocation mechanisms. *Journal of Economics and Management Strategy*, 6:639–665, 1997.

[8] S.J. Lee, A. Greenwald, and V. Naroditskiy. Roxybot-06: An $(\text{SAA})^2$ TAC travel agent. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 1378–1383, January 2007.

[9] Charles Plott. Laboratory experimental testbeds: Application to the pcs auction. *Journal of Economics and Management Strategy*, 6:605–638, 1997.

[10] P. Stone, R.E. Schapire, M.L. Littman, J.A. 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.

[11] L. Walras. *Eléments d'économie politique pure*. L. Corbaz, Lausanne, 1874.

[12] M.P. Wellman, P.R. Wurman, K. O'Malley, R. Bangera, S. Lin, D. Reeves, and W.E. Walsh. A Trading Agent Competition. *IEEE Internet Computing*, April 2001.

## Appendix A. First-Price Simultaneous Pseudo-Auctions

In this appendix, we derive the class of bidding heuristics that solves the deterministic first-price bidding problem optimally. In particular, we observe that 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 first-price bidding problem.

**Definition A.1** Given a set of goods $X$, a valuation function $v : 2^X \to \mathbb{R}$, and a distribution $f$ over pricing functions $p : X \to \mathbb{R}$, the *first-price bidding problem* is defined as follows:

$$\text{1stSIM}(X, v, f) = \max_{b \in \mathbb{R}^X} \mathbb{E}_{p \sim f} \left[ v(\text{Winnings}(X, p, b)) - \tilde{b}(\text{Winnings}(X, p, b)) \right] \tag{30}$$

The deterministic first-price bidding problem is the special case of the first-price bidding problem in which the distribution $f$ is a Dirac $\delta$ function: i.e., a distribution with all its mass at a single point. In other words, prices are certain, not uncertain.

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

$$\text{1stDET}(X, v, p) = \max_{b \in \mathbb{R}^X} \left[ v(\text{Winnings}(X, p, b)) - \tilde{b}(\text{Winnings}(X, p, b)) \right] \tag{31}$$

**Theorem A.3** *Given a set of goods $X$, a valuation function $v : 2^X \to \mathbb{R}$, and a pricing function $p : X \to \mathbb{R}$, the following bidding heuristic, which returns $b^* \in \mathbb{R}^X$, solves 1stDET$(X, v, p)$ optimally:*

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

2. *bid* $b^*(x) = p(x)$ *if and only if* $x \in A^*$

**Proof:** It suffices to show that the function $b^* \in \mathbb{R}^X$ maximizes the function $\text{1stDET}(X, v, p, b) = v(\text{Winnings}(X, p, b)) - \tilde{b}(\text{Winnings}(X, p, b))$.

First, observe that an agent should never bid more than $p(x)$ on any good $x \in X$, since by doing so, it cannot earn a higher valuation but it would pay a higher cost.

Now, for all $b \in \mathbb{R}^X$ such that $b(x) \in (-\infty, p(x)]$ for all $x \in X$,

$$\begin{aligned}
\text{1stDET}(X, v, p, b^*) &= v(\text{Winnings}(X, p, b^*)) - \tilde{b}^*(\text{Winnings}(X, p, b^*)) && (32)\\
&= v(A^*) - \tilde{b}^*(A^*) && (33)\\
&= v(A^*) - \tilde{p}(A^*) && (34)\\
&\geq v(\text{Winnings}(X, p, b)) - \tilde{p}(\text{Winnings}(X, p, b)) && (35)\\
&= v(\text{Winnings}(X, p, b)) - \tilde{b}(\text{Winnings}(X, p, b)) && (36)\\
&= \text{1stDET}(X, v, p, b) && (37)
\end{aligned}$$

Equation 33 follows from the definition of $b^*$ and the winner determination rule. Equation 34 follows from the definition of $b^*$. Equation 35 follows from the fact that $A^*$ is an optimal acquisition. Equation 36 follows from the fact that $b(x) = p(x)$ if $x \in \text{Winnings}(X, p, b)$ (if $x \in \text{Winnings}(X, p, b)$, then $b(x) \geq p(x)$; but $b(x) \leq p(x)$ and thus $b(x) = p(x)$). ■