

Bidding Marginal Utility in Simultaneous Auctions

Amy Greenwald

Department of Computer Science
Brown University, Box 1910
Providence, RI 02912
amy@cs.brown.edu

One of the key challenges autonomous bidding agents face is to determine how to bid on complementary and substitutable goods—*i.e.*, goods with combinatorial valuations—in simultaneous, or parallel, auctions (SAs). Complementary goods are goods with superadditive valuations: $v(A\bar{B}) + v(\bar{A}B) < v(AB)$. Substitutable goods are goods with subadditive valuations: $v(A\bar{B}) + v(\bar{A}B) > v(AB)$. The simple bidding strategy “for each good x in auction x , bid its (independent) valuation” is inapplicable in this framework.

Rather than attempt to reason about the independent valuations of goods, bidding agents that operate in this setting can reason about *marginal* valuations, or the valuation of a good x relative to a set of goods X . If an agent holds the goods in X , it can ask questions such as: “what is the marginal benefit of buying x ?” or, “what is the marginal cost of selling x ?” To do so, the agent must reason about the *set* of goods $X \cup \{x\}$ or $X \setminus \{x\}$, relative to the set X —the valuations of which are well-defined.

In reasoning about sets of goods bidding agents may pose and solve questions such as the following [Greenwald and Boyan, 2001]:

1. “Given only the set of goods I already hold, what is the maximum valuation I can attain, by arranging my individual goods into sets of goods?”
2. “Given the set of goods I already hold, and given market prices and *supply*, on what set of additional goods should I place *bids* so as to maximize my utility: *i.e.*, valuation less purchase costs?”
3. “Given the set of goods I already hold, and given market prices, *supply*, and *demand*, on what set of goods should I place bids or *asks* so as to maximize my utility: *i.e.*, valuation plus proceeds less costs?”

The third, and most general, of these problems, which we call *completion*, provides a foundation for bidding strategies in simultaneous single- and double-sided auctions. The second problem, which we term *acquisition*, provides a foundation for bidding strategies in single-sided auctions. The first problem, *allocation*, is a special case of the others in which all goods cost either 0 or ∞ , and all goods are worth either 0 or $-\infty$. These so-called bid determination problems are formally defined in [Greenwald and Boyan, 2001], where it is argued that completion can be reduced to acquisition.

1 Bidding Principles

This paper advocates the following bidding principle for agents participating in simultaneous auctions for complementary and substitutable goods: “coherent” marginal utility bidding: *i.e.*, bidding marginal utilities on a coherent subset of goods. To validate this principle, we compare two combinatorial bidding strategies inspired by ATTAC [Stone *et al.*, 2002] and ROXYBOT [Greenwald and Boyan, 2001], two agents in the Trading Agent Competition (TAC) (see [Wellman *et al.*, 2002]):

1. one strategy computes the marginal utility of each good independently, much like ATTAC
2. one strategy computes coherent marginal utilities—it solves the acquisition problem before bidding marginal utilities—much like ROXYBOT

It is established that approach 2 outperforms approach 1 in certain environments. In a longer version of this paper, we experiment with stochastic versions of these algorithms, and establish that approach 2 also outperforms approach 1 in uncertain environments. The insights gained from this study about the design of autonomous bidding agents are applicable beyond the scope of TAC: *e.g.*, eBay is home to numerous SAs for complementary and substitutable goods.

2 Bidding Under Certainty

In constructing bidding policies—mappings from goods to prices—for multi-unit SAs, such as those that characterize TAC, we propose the following natural breakdown of an agent’s bidding decisions:

1. how many copies of each good do I want?
2. for the goods I want, how much am I willing to pay?

One straightforward approach to answering these questions, which was employed by ATTAC [Stone *et al.*, 2002], is for the agent to skip question 1, and simply compute how much it is willing to pay for each copy of each good. An alternative approach, employed by ROXYBOT, is to explicitly answer question 1 before question 2; in this way, the agent is certain to bid on a “coherent” set of goods: *i.e.*, a set of goods which together comprise an optimal solution to the acquisition problem.

In this short paper, we compare these two approaches under two assumptions: (i) prices are exogenously determined: *i.e.*, we ignore the impact of agent behaviors on prices; and (ii) prices are known with certainty. These assumptions are clearly applicable when prices are posted. In an auction setting, we interpret these assumptions as follows: the payment rule is “pay the known price;” the winner determination rule is “win by bidding equal to or more than the price.” It is as if an agent is capable of perfectly predicting all other agents’ bids, as well as the impact of its own bids on clearing prices.

We begin by investigating ATTAC’s approach. One formula for determining willingness to pay is to compute marginal utilities (MUs). This direct MU approach to answering the aforementioned bidding questions explicitly answers only question 2, but implicitly answers question 1: a willingness to pay 0 suggests that an agent wants 0 additional copies of that good.

As alluded to earlier, computing MUs depends on solving the acquisition problem. Let us introduce some notation, and formally define acquisition and MU. Let X denote a set of goods; let $v(X)$ denote the (combinatorial) valuation of X ; let $p(x)$ denote the price of $x \in X$ and $p(X) = \sum_{x \in X} p(x)$; finally, let $u(X) = v(X) - p(X)$ denote the utility of X .

Definition 2.1 Given a set of goods X , a combinatorial valuation function $v : 2^X \rightarrow \mathbb{R}^+$, and an exogenous pricing mechanism $p : X \rightarrow \mathbb{R}^+$, an optimal solution to the acquisition problem is a subset $A^* \subseteq X$ s.t.

$$A^* \in \arg \max_{Y \subseteq X} v(Y) - p(Y) \quad (1)$$

Definition 2.2 Given a set of goods X , the marginal utility of good x is defined as follows: $\mu(x|X) = u(X^*) - u(Y^*)$, where

$$X^* \in \arg \max_{Y \subseteq X \cup \{x\}} v(Y) - p(Y \setminus \{x\}) \quad (2)$$

and

$$Y^* \in \arg \max_{Y \subseteq X \setminus \{x\}} v(Y) - p(Y) \quad (3)$$

In words, the marginal utility of good x , relative to the set X , is simply the difference between the utility of X , assuming x costs 0, $u(X|p(x) = 0)$, and the utility of X , assuming x costs ∞ , $u(X|p(x) = \infty)$.

As is evident from the definition of MU, computing marginal utilities requires two calls to a combinatorial optimization solver. If an agent explicitly answers question 1 before question 2, it might determine that it is not necessary to compute marginal utilities for all copies of all goods (although in the worst case it solves one additional optimization problem). ROXYBOT explicitly answers question 1, before question 2, and usually achieves computational savings over direct MU calculation. But the real motivation behind ROXYBOT’s approach is that reasoning about individual goods, rather than coherent sets of goods, when valuations are combinatorial, is suboptimal. In the following example, ROXYBOT’s approach is optimal, but the direct MU approach is not.

Example 2.3 Given goods x, y, z , with combinatorial valuations as follows: $v(xyz) = v(xy) = v(yz) = 500$ and $v(x) = v(y) = v(z) = v(xz) = 0$. Each good in isolation is of no value, but together with y , either x or z is worth 500. Assume all goods are priced equivalently at 100. The optimal sets of goods in this example are xy and yz , each of which yields utility of 300.

Now let us compute marginal utilities: $\mu(x|xyz) = u(xyz|p(x) = 0) - u(xyz|p(x) = \infty) = 400 - 300 = 100$; $\mu(y|xyz) = u(xyz|p(y) = 0) - u(xyz|p(y) = \infty) = 400 - 0 = 400$; and $\mu(z|xyz) = u(xyz|p(z) = 0) - u(xyz|p(z) = \infty) = 400 - 300 = 100$. Thus, a MU bidder bids on all three goods, and wins all three goods, obtaining utility $v(xyz) - (p(x) + p(y) + p(z)) = 200$.

A bidder that first determines the set of goods it wants bids on either x and y , or y and z , but not both. Computing marginal utilities for (say) x and y yields: $\mu(x|xy) = u(xy|p(x) = 0) - u(xy|p(x) = \infty) = 400 - 0 = 400$ and $\mu(y|xy) = u(xy|p(y) = 0) - u(xy|p(y) = \infty) = 400 - 0 = 400$. This bidder earns the optimal utility $v(xy) - (p(x) + p(y)) = 300$. \square

The main observations reported in this paper are the following: (i) for all goods x in an optimal acquisition, the price of x does not exceed the marginal utility of x ; (ii) for all goods x *not* in an optimal acquisition, the marginal utility of x does not exceed the price of x . It follows that bidding marginal utilities on precisely those goods in an optimal acquisition is an optimal bidding policy, if prices are exogenous (*e.g.*, an agent’s individual bidding behavior does not impact prices), and prices are known with certainty.

Proposition 2.4 *If $A^* \subseteq X$ is an optimal solution to the acquisition problem, then $\mu(x|A^*) \geq p(x)$, $\forall x \in A^*$, and $\mu(x|A^*) \leq p(x)$, $\forall x \notin A^*$.*

ROXYBOT’s strategy answers questions 1 and 2 in turn, and bids marginal utilities on all the goods it wants. This strategy is optimal, under our assumptions, by the previous observation. In particular, ROXYBOT wins all the goods it wants by bidding marginal utilities. Of course, bidding $p(x)$, or ∞ , $\forall x \in A^*$, are also optimal policies in this setting. But bidding the marginal utility of good x is a reasonable heuristic, since it is in fact an optimal bidding policy if the prices of all goods other than x are exogenous and certain.

Acknowledgments

This research was supported by NSF Career Grant #IIS-0133689. The author also acknowledges Maureen Hurtgen and the funding of CRA’s DMP.

References

- [Greenwald and Boyan, 2001] Amy Greenwald and Justin Boyan. Bidding algorithms for simultaneous auctions: A case study. In *Proceedings of Third ACM Conference on Electronic Commerce*, 115-124 2001.
- [Stone *et al.*, 2002] Peter Stone, Robert E. Schapire, János A. Csirik, Michael L. Littman, and David McAllester. ATTAC-01: A learning, autonomous bidding agent. In *Workshop on Agent Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, 2002.
- [Wellman *et al.*, 2002] Michael P. Wellman, Amy Greenwald, Peter Stone, and Peter R. Wurman. The 2001 Trading Agent Competition. In *Proceedings of the Fourteenth Innovative Applications of Artificial Intelligence Conference*, pages 935-941, July 2002.