One-Sided Decision Support for Competitive Bidding

Jesus Rios*, David Rios Insua†, David Banks‡

Abstract

Applications in counterterrorism and corporate competition have led to the development of new methods for the analysis of decision-making when there are intelligent opponents and uncertain outcomes. This field is sometimes called adversarial risk analysis. In this paper, we illustrate a general framework developed for supporting a decision maker in a problem with intelligent opponents through a simple price-sealed bid auction case.

Key words : adversarial risk analysis, price sealed-bid auctions, decision support, games with incomplete information, common prior assumption critique

1 Introduction

Applications in counterterrorism and corporate competition have led to the development of new methods for the analysis of decisions when there are intelligent opponents and uncertain outcomes. This field represents a combination of statistical risk analysis and game theory, and is sometimes called *adversarial risk analysis* (ARA). Prevalent methodologies are based on game theory, decision analysis, or conventional risk analysis, emphasizing separate aspects of the analysis. Rios Insua et al (2008) describe a unified framework for the analysis of decisions under uncertainty in presence of intelligent adversaries. This is an asymmetric prescriptive/descriptive approach in the spirit of Raiffa's (2002) approach to games. The key issue in this framework is the assessment of the probabilities of our adversaries' possible actions. It assumes that adversaries are expected utility maximizers, and, therefore, the probabilities on the adversary's actions stem from our uncertainty

*ERCIM fellow, University of Luxembourg. jrios@samsi.info
†Statistics and Operations Research Department, University Rey Juan Carlos, Spain.
‡Department of Statistical Science, Duke University, USA

about the adversary's decision problem. In the simplest case, when the structure of our adversary's decision problem is trivially known, we just need to express our uncertainty about our adversary's probabilities and utilities.

2 One-Sided Decision Support for Price Sealed-Bid Auctions

In this paper, we apply the ARA framework to analyze a simple, and concrete auction situation. We first provides the non adversarial case which allows us to introduce the basic problem for one decision maker. Later, we move to the adversarial case. It will illustrate the application of the adversarial risk analysis framework to simultaneous decision problems. This is the case in which the analysis of the decision problem of our adversary depends on our own decision problem, and the assessment of probabilities on the adversary's actions requires a way of getting around of this infinity regress.

2.1 Single Bidder, One Bid, and Unknown Reservation Price

Suppose Daphne is bidding for a certain object. She is the only bidder, but the owner has set a secret reservation price v below which the object will not be sold. Daphne does not know v, and expresses her uncertainty as a distribution $F(v)$. Daphne's utility function in money is u_D and her personal valuation of the auctioned object is v_D. Her choice set is $\mathcal{D} = \mathbb{R}^+$ and her expected utility for a bid of $d \in \mathcal{D}$ is $u_D(v_D - d)\ \mathbb{P}[d > V]$. Thus, in a standard decision analytic computation, see Raiffa (2002), Dahpne should maximize her expected utility by bidding $d^* = \text{argmax}_{d \in \mathcal{D}}\ u_D(v_D - d)F(d)$.

2.2 Two Sealed Bids, the Highest Bid Wins

Suppose now that Daphne and Apollo are bidding against each other. Each one knows his own valuation of the auctioned object (reservation price) but does not know the reservation price of the other. Each bidder submits one bid in a sealed envelop without knowing the other's bid, the winner being that with the highest bid. This is a simultaneous decision making situation in which bidders are uncertain about the other bidder's reservation values and hence each other's utilities. The influence diagram in Figure 1 represents this case.

Harsanyi's (1967) approach here would lead to the solution concept of Bayes-Nash equilibrium in games with incomplete information, which is based on the common prior assumption, entailing in this case that players need to disclose, inter alia, their true beliefs about the other player's reservation price. Thus, Daphne probabilistic assessment of

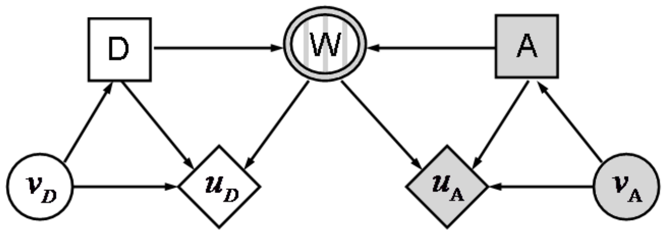

Figure 1: ID of the sealed bid auction problem

Apollo's reservation price and Apollo's probabilistic assessment of Daphne 's reservation price would be common knowledge. Only under this assumption, it is possible to compute a prediction for that game.

We are, however, interested in supporting Daphne who knows her value v_D of the object and has a (private) probabilistic assessment of Apollo's valuation $v_A \sim V_A$. Daphne has to decide which is her bid d. If this is bigger than Apollo's bid a, she wins obtaining a utility $u_D(v_D - d)$; if it is smaller $(d < a)$, she gets 0, as reflected in Figure 2(a). Therefore, similarly to the previous section, the problem she has to solve is

$$\max_d \; u_D(v_D - d) \; \mathbb{P}_D(d > A|d).$$

The key issue is assessing Daphne's probability $\mathbb{P}_D(d > A|d)$ of winning for each of her possible bids. This assessment would be based on her prediction about Apollo's bid, represented by the probability distribution $\pi_D(a)$, as $\mathbb{P}_D(d > A|d) = \int_{-\infty}^{d} \pi_D(a)da$. Sometimes, it may be possible to assess $\pi_D(a)$ based on past statistical data describing previous behavior of bidders, as in Keefer et al (1991), or use a noninformative distribution.

Another possibility would be to assess that probability judgementally through an analysis of Apollo's bidding decision problem from Daphne's perspective, as shown in Figure 2(b). To simplify the discussion assume that both Apollo and Daphne are risk neutral in the range of interest, so that their utility functions u_D and u_A are linear. Then, Apollo would aim at solving

$$\max_a (v_A - a) \; \mathbb{P}_A(a > D|a).$$

Daphne would need to know the solution of such problem, but she may not solve it as she is uncertain about v_A and $\mathbb{P}_A(a > D|a)$. Her beliefs about v_A were modelled through V_A. The assessment of $\mathbb{P}_A(a > D|a)$ requires Daphne's elicitation of Apollo's distribution π_A on her bid d, as $\mathbb{P}_A(a > D|a) = \int_{-\infty}^{a} \pi_A(d) \; \mathrm{d}d$. At this step of the analysis, again, we could base the assessment of $\pi_A(d)$ directly on the data that, she thinks, are available to him, expert opinions or, possibly, a combination of both. Should this kind of information not be available we could use a non informative prior to describe $\pi_A(d)$.

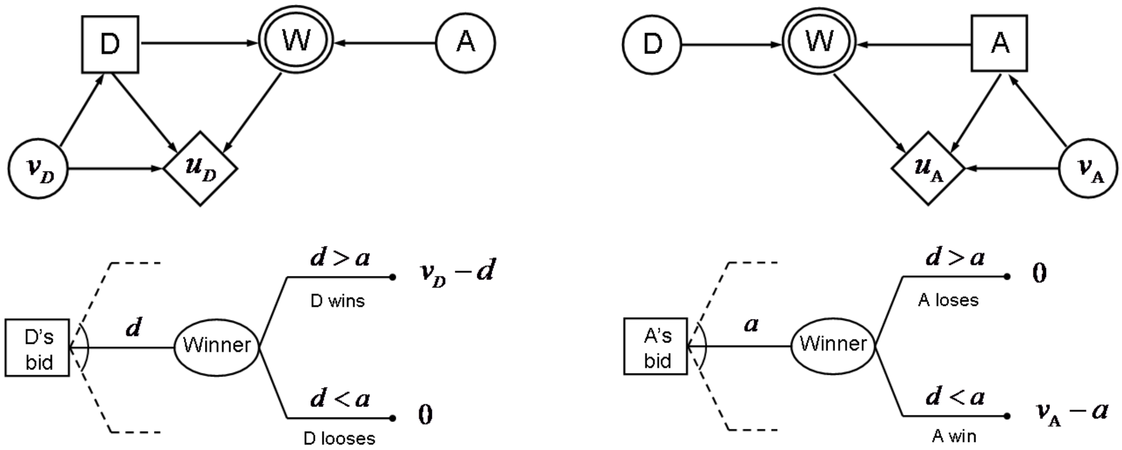

Figure 2: Auction analysis from Daphne's perspective. ID and decision tree representations

Going in deeper detail in the ARA analysis can help Daphne, as well, to elicit $\pi_A(d)$ through the identification of the relevant variables that affect Apollo's guess about her bid. To wit, Daphne will think that Apollo will analyze her problem as in Figure 2(a) and solve

$$\max_d \, (\hat{v}_D - d) \, \mathbb{P}_D(d > \hat{A}|d),$$

where $\hat{v}_D$ represents Apollo's estimation of Daphne's valuation v_D. Since it is unknown to Daphne in her analysis of Apollo's analysis of her problem, she will assign to $\hat{v}_D$ a distribution $\hat{V}_D$. Daphne's distribution about Apollo's bid, but now elicited from Apollo's perspective, will be $\hat{A} \sim \hat{\pi}_D(\hat{a})$. Thus, implicitly, the elicitation of $D \sim \pi_A(d)$ depends on $\hat{V}_D$ and $\hat{A}$, both assessed by Daphne. Should we go ahead one more step in this sort of analysis, we would see that, in turn, the assessment of $\hat{A}$ would depend on (i) $\hat{V}_A$ representing the random variable that Daphne thinks Apollo uses to represent Daphne's estimation of his valuation v_A; and (ii) a distribution over $\hat{D}$ representing what Daphne thinks Apollo thinks... of her bid.

To avoid an infinite regress, we shall stop here and use a heuristic approach in helping Daphne to assess the distribution $D \sim \pi_A(d)$ based on the relevant $\hat{V}_D$ and $\hat{V}_A$ distributions, that we have identified, disregarding the rest of variables associated with the thinking-about-what-the-other-is-thinking-about kind of analysis, whose meaning in practice renders of very difficult interpretation by most people.

One possibility for this heuristic approach is to assume that Apollo expects Daphne to make a bid d which is a function $f(\hat{v}_D, \hat{v}_A)$, with some uncertainty around it. We shall explore the case in which Apollo expects Daphne to bid

$$min(\alpha\ \hat{v}_D, \beta\ \hat{v}_A), \tag{1}$$

where $\alpha, \beta \in (0,1)$, $\hat{v}_D$ is Apollo's estimation of her valuation v_D and $\hat{v}_A$ is Daphne's estimation of his valuation v_A as thought by Apollo. The intuitive interpretation of Equation (1), which simplifies Daphne's analysis of how Apollo would think about her problem, is that Daphne's bidding behaviour consists of submitting a bid which guarantees (i) a profit in terms of a proportion $(1-\alpha)$ of her valuation of the object $\hat{v}_D$, and (ii) that she will not overbid Apollo's bid, which is also assessed as a percentage (β) of his object valuation $\hat{v}_A$.

Daphne's confidence about her assessment of the distributions over $(\alpha, \hat{v}_D, \beta, \hat{v}_A)$ as well as her uncertainty about the accuracy of the proposed heuristic for her analysis of Apollo's problem, could be incorporated through a hierarchical model with a new parameter σ modelling this confidence, e.g. arguing that $\pi_A(d \mid \alpha, \hat{v}_D, \beta, \hat{v}_A, \sigma)$ is a normal distribution with mean $min(\alpha\ \hat{v}_D, \beta\ \hat{v}_A)$ and standard deviation σ, truncated on $[0, \hat{v}_D]$, with $(\alpha, \hat{v}_D, \beta, \hat{v}_A, \sigma) \sim \Pi_A$. As usual, as σ gets larger we will be closer to the case in which we use a non informative distribution to model $\pi_A(d)$.

We have thus reduced Daphne's ARA about D with probability density function $\pi_A(d)$ to the elicitation of the distribution Π_A representing her uncertainty over the above parameters. All these quantities can be obtained directly from Daphne through her analysis of Apollo's problem in the form of probability distributions acknowledging her uncertainty about them, yielding a specific distribution D. Note that Daphne's distribution Π_A might reflect her opinions about e.g. how Apollo uses his valuation v_A of the object to assess his $\hat{v}_D$, and, therefore, Π_A might be correlated with V_A.

In general, we can use a Monte Carlo approach to estimate $\mathbb{P}_D(d > A|d)$, which would then consist of running for $i = 1, \ldots, n$ iterations

1. Draw $v_A^i \sim V_A$
2. Draw $\omega^i = (\alpha^i, \hat{v}_D^i, \beta^i, \hat{v}_A^i, \sigma^i) \sim \Pi_A \mid v_A^i$
3. Set $D_i \mid \omega^i \sim N\left(\min\left(\alpha^i\ \hat{v}_D^i, \beta^i\ \hat{v}_A^i\right), \sigma^i\right)$ truncated on $[0, \hat{v}_D^i]$, with $\pi_A^i(d_i|\omega^i)$ its probability density function
4. Solve
$$a_i^* = \text{argmax}_a (v_A^i - a)\ \mathbb{P}_{\pi_A^i}(D_i < a|a, \omega^i),$$
where
$$\mathbb{P}_{\pi_A^i}(D_i < a|a, \omega^i) = \int_{-\infty}^{a} \pi_A^i(d_i|\omega^i)\ \mathrm{d}d_i.$$

and then use $\{a_i^*,\ i = 1, \ldots, n\}$ to approximate Daphne's probability of winning conditional on her bid d through

$$\hat{\mathbb{P}}_D(d > A^*|d) = \#\{d > a_i^*\}/n.$$

The whole approach to support Daphne in choosing her sealed bid would then be as follows (recall we assume risk neutrality and, thus, we do not assess u_D from Daphne, neither u_A describing her beliefs about Apollo's utility function).

1. Assess v_D from Daphne, her object valuation
2. Assess $F = (V_A, \Pi_A)$ describing Daphne's uncertainties in her analysis of Apollo's problem
3. Estimate $\mathbb{P}_D(d > A|d)$ through $\hat{\mathbb{P}}_D(d > A^*|d)$ using Monte Carlo simulation as above
4. Solve
$$d^* = \arg\max_d\ (v_D - d)\ \hat{\mathbb{P}}_D(d > A^*|d).$$

Note that the elicitation of the proposed model for Daphne's bid D whose probability density function is $\pi_A(d)$, could be simplified by just asking Daphne for point estimates for the quantities $(\alpha, \hat{v}_D, \beta, \hat{v}_A, \sigma)$, in which case we would not be acknowledging her confidence in her numerical assessments or in the accuracy of (1) as heuristic representation of bidding behaviour.

3 A Numerical Example

We illustrate this heuristic approach with a numerical example. Recall we are supporting Daphne to find her maximum expected utility bid $d^* = \text{argmax}_{d \in \mathcal{D}}\ (v_D - d)\ \mathbb{P}_D(d > a|d)$. To do so, we obtain from her (who may possibly be assisted by a group of subject matter experts) the following judgmental assessments.

- $v_D = 100$: Value of the object for her.
- $v_A \sim V_A$: Daphne believes that Apollo's object valuation must be in a range between 60 (min) and 90 (max), and most likely is 80 (mode). Base on this information we decide to fit a triangular distribution whose mode, minimum and maximum correspond with the values indicated by Daphne, see Figure 3.

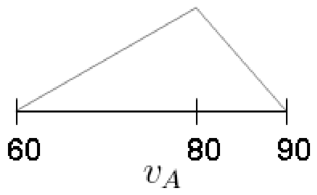

Figure 3: Daphne's assessment of V_A

- $\hat{v}_D \sim \hat{V}_D$: Apollo's estimation of Daphne's valuation v_D. Daphne believes that Apollo thinks her valuation of the auctioned object is around 100 units higher than his (v_A), with an error range between -5 and 5 units. She also consider that all possible errors in this range are equally likely. That is,

$$\hat{V}_D \mid v_A = v_A + 100 + U(-5, 5).$$

- $\hat{v}_A \sim \hat{V}_A$: Daphne's estimation of v_A as thought by Apollo. Daphne believes that Apollo thinks that she believes that his valuation of the object is 50 units lower than his (v_A), with an error range between -5 and 5 units. She also consider all this possible errors equally likely. That is,

$$\hat{V}_A \mid v_A = v_A - 50 + U(-5, 5).$$

- $1-\alpha$ and $1-\beta$: Profit value proportions used by Apollo in his analysis of her bidding problem when he thinks about how she analyzes bidding behavior. She believes these parameters are

$$1 - \alpha = 1 - \beta = 0.3 + U(-0.05, 0.05).$$

- $\sigma = 1$: Our assessment of Daphne's confidence on our heuristic model and second order parameter assessment.

From these assessments provided by Daphne, we can determinate the heuristic distribution representing the bidding behavior of Daphne as thought by Apollo, elicited from the perspective of Daphne:

$$D \mid v_A \sim N\left(\min\left(\alpha\, \hat{V}_D, \beta\, \hat{V}_A\right), \sigma\right).$$

We note that this distribution is a truncated normal distribution on $[0, \hat{v}_D]$, given $(\alpha, \hat{v}_D, \beta, \hat{v}_A, \sigma)$.

Now, for each $i = 1, \ldots, n$, we simulate $v_A^i \sim V_A$ and $\omega^i \sim \Pi_A \mid v_A^i$, set $D_i \mid \omega^i = N\left(\min\left(\alpha^i\, \hat{v}_D^i, \beta^i\, \hat{v}_A^i\right), \sigma^i\right)$, truncated on $[0, \hat{v}_D^i]$, and solve Apollo's optimization problem

$$a_i^* = \text{argmax}_{a \in \mathcal{A}} (v_A^i - a)\ \mathbb{P}_A(a > D_i | a, \omega^i).$$

Figure 4(a) illustrates the optimization problem solved at one of the iterations, including Apollo's expected utility for each of his possible bids along with his optimal bid $a_i^* = 23.4$ at that iteration. After $n = 1000$ iterations we obtain a sample $\{a_i^*,\ i = 1, \ldots, 1000\}$ from $A^* = \text{argmax}_{a \in \mathcal{A}}(V_A - a)\mathbb{P}_A(a > D|a)$, which represents Daphne's predictive distribution on Apollo's bid. Its probability density function $\pi_D(a)$ has been estimated from the obtained sample through a kernel density estimator shown in Figure 4(b).

Finally, we solve Daphne's decision problem

$$d^* = \text{argmax}_{d \in \mathcal{D}}(v_D - d)\ \hat{\mathbb{P}}_D(d > A^*|d),$$

where $\hat{\mathbb{P}}_D(d > A^*|d)$ is the proportion of simulated a_i^* below a given d: $\#\{d > a_i^*\}/1000$.

Figure 4(c) plots our Monte Carlo estimation of Daphne's expected utility for each of her possible bids. We can see that her maximum expected utility bid is $d^* = 30.6$ with (estimated) expected utility of 68.5.

4 The noninformative case

Should Daphne not be able to provide us with the necessary information to assess a probability density function $\pi_A(d)$ for D, or she does not feel confident with our heuristic model, we would use a noninformative prior to describe D distribution. We run now the same numerical example than before but with $\sigma = 100$, as a numerical approximation of the noninformative case. Figure 5 summarizes the results, in parallel to the previous description of our numerical example. We note that when $\sigma = 100$, $\pi_A^i(d_i|\omega^i) \approx U(0, v_D^i)$ at each iteration i. Therefore, Apollo's expected utility for each of his possible bids a will be $(v_A^i - a)\, a/v_D^i$, in correspondence with the parabola shown in Figure 5(a). This result is independent of $\hat{V}_A$, α and β, and consequently Daphne's maximum expected utility bid as well, as expected when a noninformative distribution is used to describe D.

5 Conclusions

We have applied ARA to support one bidder in a two-person price sealed-bid auction, a simultaneous decision making problem in each each bidder makes a bid without knowing the other's bid. We use a Bayesian decision analysis approach to support one bidder against another, which is different to the standard approach used in the game theory literature for auction modeling. To do so, we have propose an heuristic to avoid the infinity regress kind of thinking required to analyzed this kind of problems from a Bayesian perspective.

Real problems are extremely complex. For this reason, we have focused on a simple example that illustrate our general formulation to deal with the ARA problem. This paper focused on a two-person game, but the discussion directly extends to n-person games. However, when there are n players, the number of possible analysis increases combinatorially, quickly posing computational challenges. Also, implementation issues include the study of how to facilitate the elicitation of a valuable judgmental input from supported bidder. Other issues within ARA that require more research include other types of auctions, the case in which the number of agents is uncertain as well as the incorporation of cooperative issues within the analysis.

Acknowledgments

We acknowledge support for this research from SAMSI, the Spanish Ministry of Science and Education and the E-DEMOCRACIA-CM program.

References

Harsanyi, J. 1967. Games with incomplete information played by Bayesian players. *Management Science*, 14, 159–182.

Keefer, D., F. Beckley and H. Back. 1991. Development and use of a modeling system to aid a major oil company in allocating bidding capital. *Operations Research*, 39, 28–41.

Raiffa, H. 2002. *Negotiation Analysis*. Harvard University Press: Cambridge, Massachusetts.

Rios Insua, D., J. Rios and D. Banks. 2008. Adversarial risk analysis. *Technical Report*.

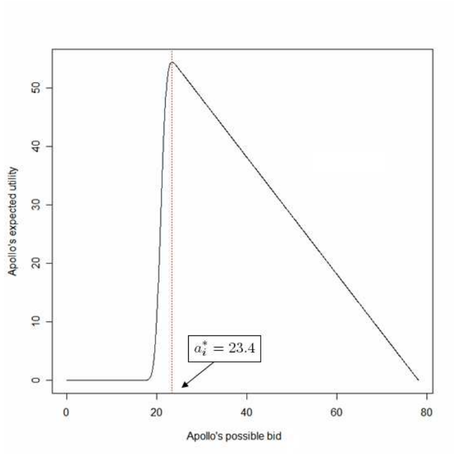

(a) Solving Apollo's optimization problem at the i-th iteration

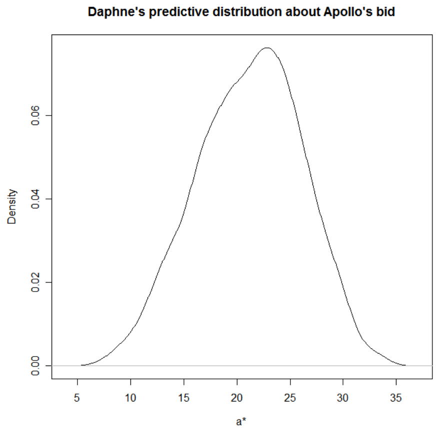

(b) Kernel density estimation of $\pi_D(a)$

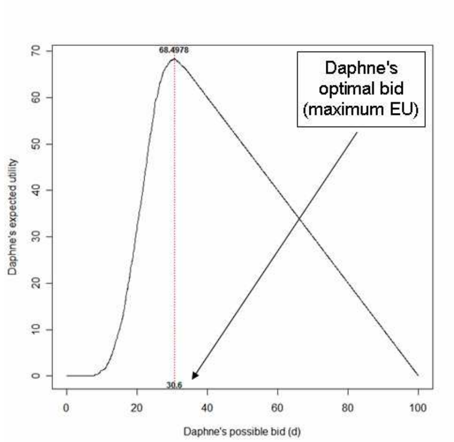

(c) Daphne's expected utility function

Figure 4: Numerical Example

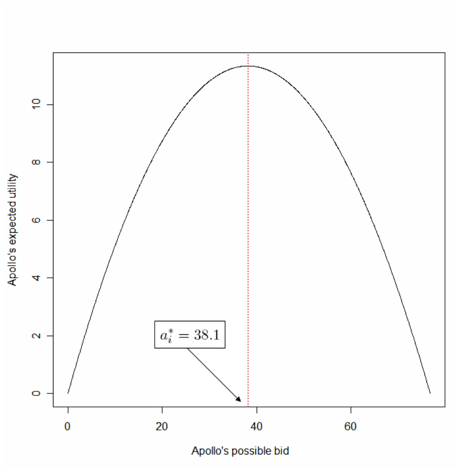

(a) Solving Apollo's optimization problem at the i-th iteration

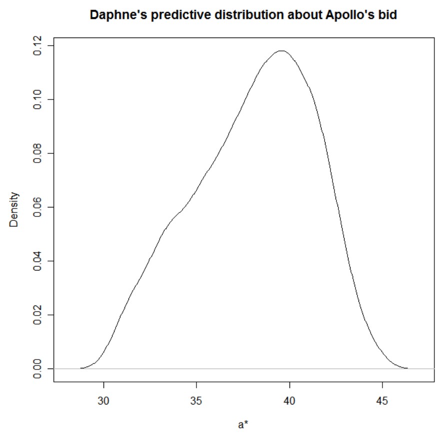

(b) Kernel density estimation of $\pi_D(a)$

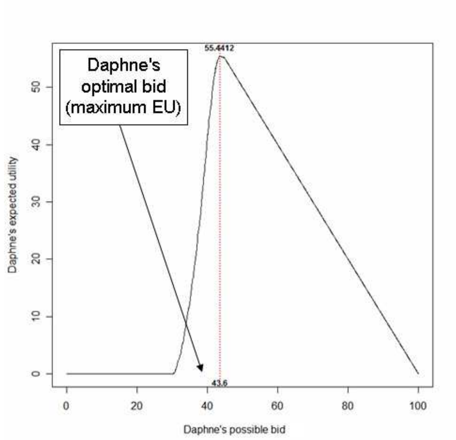

(c) Daphne's expected utility function

Figure 5: Noninformative case ($\sigma = 100$)