Optimal Value of Information Based Elicitation During Negotiation

Yasser Mohammad
AIST, Tokyo, Japan
Assiut University, Egypt
y.mohammad@aist.go.jp

Shinji Nakadai
NEC-AIST Collaboration Laboratory
Tokyo, Japan
s-nakadai@az.jp.nec.com

ABSTRACT

Autonomous agents engaging in automatic negotiations on behalf of humans or institutions are usually assumed to have full knowledge of the utility function for the actors they represent. In many cases, these utility functions are difficult to know apriori for every possible outcome of the negotiation. Moreover, it may not be necessary for the agent to know the utility of outcomes that are never offered or considered during the negotiation. State-of-the-art approaches to utility elicitation during negotiation assume that the agent can ask *questions* from a predefined countable set to reduce its uncertainty about the utility function. This paper extends that body of work by lifting the countability assumption providing an optimal algorithm for selecting the best outcome and utility level about which to ask the actor. The paper reports the results of comparing the proposed algorithm with state-of-the-art algorithms using both synthetic and realistic negotiation scenarios. These evaluations support the applicability of the proposed approach.

KEYWORDS

Autonomous Negotiation; Utility Elicitation; Probabilistic Inference

ACM Reference Format:
Yasser Mohammad and Shinji Nakadai. 2019. Optimal Value of Information Based Elicitation During Negotiation. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 9 pages.

1 INTRODUCTION

Automatic negotiation is attracting more attention from the research community in recent years especially given the rise of AI, machine learning systems and the Internet of Things (IoT) that promise to automate most repetitive aspects of our lives. The main advantage of automatic negotiation in this context is providing the means for multiple actors to organize their behavior to achieve win-win deals allowing for better use of resources and providing an advantage for the society as a whole. Recent applications of automatic negotiation include permission management in IoT systems [2], Wi-Fi channel assignment [7], providing feedback for student negotiation skills [13], and agriculture supply chain support [10].

Most of this work assumes that the negotiation agent has perfect knowledge of the utility function of the person/entity it is representing during the negotiation. While this can be the case in some limited cases, in many real-world scenarios; it is difficult and uncommon to have perfect apriori knowledge of this utility function. Utility elicitation from people always involves some elicitation *cost* (bother cost) [3]. Trying to pin the utility value of every possible outcome before the negotiation starts entails asking too many questions leading to high *bother cost.*

Utility elicitation is studied extensively in the decision support community [8, 11, 15]. Most of that work focuses on the problem of eliciting the utility function of an actor for several possible outcomes of the decision process. Negotiation adds a new complexity to this problem because – during negotiation – it is not enough to know the utility value of some outcome for the user in order to propose/accept it, but it is also essential to judge the probability that this outcome/offer is also acceptable by the partner(s) in the negotiation.

Both negotiation and preference elicitation have been studied for decades. For example, Rubinstein proved his perfect-game equilibrium for a simplified version of the alternating offers negotiation protocol in 1982 based on even earlier work [19]. In the same year, Llewellyn et al. studied the use of standard gambles for utility elicitation in the medical domain [16].

More recently, studies start to appear that focus on the combined problem of utility elicitation *during a negotiation* [3, 4, 17, 18].

Baarslag and Gerdeng proposed the optimal elicitation algorithm [3] is based on Pandora's Rule [20]. It assumes that the actor (user) can be queried to provide *exact* utility values for different outcomes. This is achievable using several possible elicitation strategies that do not require the assignment of a numeric value to any outcome by the user. Nevertheless, this form of *deep* elicitation for each outcome considered is time-consuming and would lead to high levels of elicitation bother to the user that can sometimes be avoided. For example, if a single query assigned a probability distribution to an outcome that made it *dominated* by another there is no need to keep eliciting this dominated outcome until a numeric utility assignment is achieved. A *shallow* version of this algorithm that uses a heuristic to avoid offering outcomes that turn out to have low utilities at the beginning of the negotiation was recently proposed [18]. One advantage of this approach though is the reliance on an easy to calculate *aspiration level* for the judgment of the expected utility for outcomes rejected by the partner(s). It is also efficient enough to be usable when time-pressure is an issue for the negotiation.

The *Optimal Query Agent (OQA)* [4] avoids the problem of deep elicitation by assuming that a predefined set of possible queries are available that can reduce the uncertainty in the probability distribution of utility values for a given outcome. The system selects the optimal query at each point as the one maximizing the value of information which is the difference between the *expected expected utility* if the answer to the query is known compared with it if the

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

answer is not known. This approach has several advantages: Firstly, it relies on principled probabilistic reasoning. Secondly, it is possible to extend to cases where the outcome space has some structure linking the marginal utility distributions for different outcomes. Finally, this approach does not assume any specific form for the queries given to the user. The main disadvantage of this approach is its high computational cost: $O\left(n_q n_o\right)$ for n_q questions and n_o outcomes.

Mohammad and Nakadai [17] proposed an efficient value of information based algorithm (FastVOI) that was shown to achieve the same utilities as OQA while reducing the computational complexity to $O\left(n_q \log n_o\right)$.

Both OQA and FastVOI assume that the questions that can be asked about the utility of different outcomes are predefined and countable. Moreover, the time complexity of both approaches is linear in the number of these predefined questions. This paper extends the value-of-information approach by removing the need for predefined questions at the expense of having all questions being about a single outcome. The main contribution of this work is a proof of optimality for the proposed method for selecting the cut-off utility value to ask about based on the value-of-information principle. Moreover, the paper reports extensive evaluation experiments comparing the proposed algorithm and state of the art elicitation during negotiation algorithms.

The rest of the paper is organized as follows: Section 2 gives an introduction to negotiation and details the negotiation protocol used in this paper. Section 3 gives the formal representation of negotiation used in this paper and Section 4 presents the elicitation during negotiation problem and details the value-of-information approach in the context of negotiation. Section 5 defines the elicitation during negotiation with uncountable single-outcome questions. Section 6 details the proposed algorithm for solving it and Section 7 evaluates two variants of it. Section 8 discusses the limitations of the proposed approach and provides directions for future research. The paper is then concluded.

2 PROBLEM SETTING

A negotiation session is conducted between multiple agents representing self-interested actors over a set of issues. Issues can have discrete or continuous values. Every possible assignment of a value to each issue is called an *outcome* and during negotiation, it may also be called an *offer* or *proposal*. If an agreement is reached the agreed upon outcome is called a *contract*. Each actor (which can be a human but not be) is assumed to be self-interested with some internal *utility function* that assigns a numeric value (assumed to be normalized to the range 0 to 1 in this work) to every possible outcome. The actor wishes to maximize the utility she receives from the negotiation through the behavior of its representing agent. Every actor also has a predefined *reserved value* that she gets if the negotiation was broken either due to timing out or explicitly by one of the agents.

Negotiation sessions are conducted in *rounds* in which different outcomes are offered/judged by the agents according to some negotiation *protocol*. Negotiation protocols can be moderated – with a moderator-agent that facilitates the interaction – or unmoderated. Negotiation sessions can end in an agreement (contract) or be broken without agreement. Negotiation sessions are usually time-limited, to provide an incentive for the agents to concede. The session is broken (times-out) automatically if agents did not reach an agreement within a predefined real-time or number of rounds limit and a discount factor may be used to encourage express agreement.

Negotiations are carried out between agents according to a protocol. Several negotiation protocols have been proposed over the years. They can either be mediated [12] or un-mediated [1]. This work utilizes the un-mediated Alternating Offers Protocol (AOP) [1] but is not limited to it.

AOP works as follows: An ordering of the agents is defined which we assume — without loss of generality — is the same as the agent index i. The first agent starts the negotiation by *offering* an outcome ω^0 which is visible to all other agents. The next agent either *accepts* the offer, ends the negotiation, or proposes a new offer. This process is continued until one of the following stopping criteria is met:

- **Agreement:** All agents accept an offer. In this case, this offer is declared as the agreed upon *contract*.
- **Timeout:** A predefined number of offer exchanges/rounds (N) or a preset number of seconds T has passed since the beginning of the negotiation.
- **Failure:** Some agent ends the negotiation when it has the chance to respond to some offer.

3 FORMAL REPRESENTATION

A negotiation scenario is defined – in this paper – by a tuple ($\Upsilon \equiv (A, T, N, R, \Omega, \{\tilde{u}_a\}, \{u^0_{\omega a}\})$ where:

$A \in \mathbf{I}^+ - \{1\}$**:** Number of agents/actors

$T \in \mathbf{R} \bigcup \infty$**:** The allowed time of the negotiation.

$N \in \mathbf{I} \bigcup \infty$**:** The allowed number of rounds of the negotiation.

$R : A \longrightarrow [0, 1)$**:** Reserved value for agent a. We define $r_a \equiv R(a)$ $\forall 1 \leq a \leq A$.

$\Omega \equiv \{\omega_j\}$**:** Possible outcomes (can be countable or uncountable) where $|\Omega| \geq 2$

$\tilde{u}_a : \Omega \longrightarrow [0, 1]\ \forall 1 \leq a \leq A$**:** Utility function of **actor** a with corresponding cummulative distribution $\tilde{U}_a$.

$u^0_{\omega a} : [0, 1] \longrightarrow [0, 1]\ \forall 1 \leq a \leq A \wedge \omega \in \Omega$**:** The probability distribution of utility values for outcome ω as know to **agent** a with corresponding cummulative distribution $U_{\omega a}$.

To simplify the notation we abuse the above definitions of $R, \tilde{u}_a, u_a$) by allowing subscripts to indicate the outcome and allowing j_ω to stand for ω_j. Moreover, the outcome order may stand for the outcome in subscripts. This means that the following are equivalent:

$$X_a\left(\omega_j\right) \equiv X_{\omega_j a} \equiv X_{j_\omega a} \equiv X_{ja}$$

The agents receive their utility values at the end of the negotiation as follows:

- **Agreement on contract ω after n rounds:** Each agent a receives a utility of $\tilde{u}_a(\omega) - C_a$ where C_a is the total cost incurred by the agent a during the negotiation.
- **Timeout at time t:** Each agent a receives a utility of $r_a - C_a$.
- **Failure at time t:** Each agent a receives a utility of $r_a - C_a$.

Hereafter, when the agent considered is clear from the context, the subscript a will be dropped.

4 ELICITATION WITH PREDEFINED QUESTIONS

The standard elicitation during negotiation problem assumes that a countable set of predefined questions are available to the agent. An elicitation during negotiation scenario with predefined questions Ξ_{finite} is defined by the tuple (Υ, Q_a, Q_a^0) where:

- Υ A negotiation scenario.
- $Q_a : \{q_{al}\}$ is the set of questions/queries available to the agent where: $q_{al} \equiv \left(\left\{r_{al}^i\right\}, c_{al}\right)$ is one question defined as a tuple of a set of possible replies $\left\{r_{al}^i\right\}$ and a cost c_{al} for asking the question q_{al}.

Defining, $P : [0,1]^m \to [0,1]$ as a joint probability distribution over m real-valued variables with unity domains, a reply can then be defined as: $r_{al}^i : P \to P$ (i.e. a mapping from a joint probability distribution to another). It is assumed that the output joint distribution has lower spread (i.e. variance) for at least one variable.

Elicitation during round n is the process of selecting one query q_a^n from Q_a, presenting it to the actor a, receiving a reply r_a^n then applying the transformation defined by that reply to the joint utility distribution:

$$u_a^n \leftarrow r_a^n\left(u_a^n\right).$$

We assume that all members of a negotiation scenario with elicitation Ξ_{finite} are constant for the duration of the negotiation except $\hat{u^n}_a$. A numeric upper script indicates the round number (n) if not otherwise specified and will be dropped when it is clear from the context.

We will assume some arbitrary ordering J on the outcomes Ω where $j_\omega \leq |\Omega|$ is the index of outcome ω in that ordering and $\omega_j \in \Omega$ is the outcome associated with index j in that ordering.

4.1 Value of Information Based Solution

The Value-of-information (VOI) based algorithm for preference elicitation was first proposed by Chajewska [6] in the context of decision theory. Baarslag and Kaisars [4] proposed a modified version explicitly designed for elicitation during negotiation. Mohammad and Nakadai [17] lately proposed a more efficient variant. This section provides a brief description of this approach using the formulation presented in section 4 as it is the base of the proposed algorithm.

All VOI variants assume that of an acceptance model $\mathcal{M} : \Omega \to [0,1]$ is available to the agent. We use the same subscript/superscript rules defined in section 3 with the acceptance model.

A policy $\Psi \equiv \left\langle \omega_j | \omega_j \in \Omega \right\rangle$ of length D is an ordered sequence of outcomes (Ψ_i is the outcome at index i in Ψ). For the rest of this section, it is assumed that $D = N - n$ (i.e. policy length is exactly the remaining number of rounds) if not otherwise stated.

Given an elicitation scenario Ξ_{finite} and an acceptance model $\mathcal{M}$, the expected utility of following a policy Ψ will be:

$$EU\left(\Psi|\mathcal{M}, \tilde{u}\right) = \sum_{k=1}^{|\Psi|} \mathcal{M}\left(\Psi_k\right) \tilde{u}\left(\Psi_k\right) \prod_{j=1}^{k-1}\left(1 - \mathcal{M}\left(\Psi_j\right)\right). \quad (1)$$

Because the agent has no access to the true utility of the actor $\tilde{u}$, it can only calculate the expectation of EU which is known as the *expected expected utility* [4, 5]:

$$EEU\left(\Psi|\mathcal{M}, u\right) = \int_u EU\left(\Psi|\mathcal{M}, u\right). \quad (2)$$

The order of integration and differentiation in Equation 2 can be exchanged leaving to the following definition for EEU:

$$EEU\left(\Psi|\mathcal{M}, u\right) = \sum_{k=1}^{|\Psi|} \mathbb{E}\left[u_{\Psi_k}\right] \mathcal{M}\left(\Psi_k\right) \prod_{j=1}^{k-1}\left(1 - \mathcal{M}\left(\Psi_j\right)\right), \quad (3)$$

where, $\mathbb{E}[u]$ is the expected value of u.

An optimal policy π is a policy that maximizes EEU with respect to a given acceptance model and utility distribution:

$$\pi|\mathcal{M}, u \equiv \arg\max_{\Psi} EEU\left(\Psi|\mathcal{M}, u\right) \quad (4)$$

The optimal policy of length 1 ($\langle\omega_*\rangle$) is easy to calculate:

$$\omega_* \equiv \arg\max_{\omega} \mathbb{E}\left[u_\omega\right] \mathcal{M}\left(\omega\right)$$

Baarslag and Kaisars [4] provided an efficient greedy algorithm for calculating $\pi|\mathcal{M}, u$ that incrementally finds the optimal policy of length l given the one of length $l-1$.

Given $\mathcal{M}, u$, the elicitation process of the Optimal Query Agent (OQA) [4] proceeds by calculating the optimal policy π with the expected expected utility under this policy (eeu^*), then for each question $q \in Q$, an *optimal alternative policy* π_{qr} for each possible answer r is calculated with the associated eeu_{qr} as:

$$\pi_{qr} = \arg\max_{\Psi} EEU\left(\Psi|\mathcal{M}, r(u)\right), \quad (5)$$

The expected expected utility of that question is then calculated as the weighted average of expected expected utility of following the optimal policy after getting each answer with the associated cost subtracted:

$$eeu_q = \sum_{r,p,c \in q} p \times eeu_{qr} - c, \quad (6)$$

where p, r, c are the probability of reply, the reply and its associated cost respectively. This leads to the following definition of the *value of information* for asking question q:

$$VOI(q) = \sum_{r \in \text{answers}(q)} p_r \times eeu_{qr} - eeu^* \quad (7)$$

where p_r is the probability of getting answer r and can either be given or estimated from the utility distribution u.

The VOI algorithm asks the question q^* that maximizes eeu_q as long as $eeu_{q^*} > eeu^* + c$ which means that asking this question would entail a positive value of information.

Recently, Mohammad and Nakadai [17] proposed the FastVOI algorithm for efficient selection of the optimal question q^* by incremental calculation of eeu^*.

To simplify the notation, we define the cummulative sum and product at element i in Equation 3 as:

$$P_i \equiv \prod_{j=1}^{i-1} 1 - \mathcal{M}\left(\pi_j\right) = \prod_{j=1}^{i-1} 1 - M_k, \quad (8)$$

$$S_i \equiv \sum_{k=1}^{i} \mathcal{M}\left(\omega_k\right) \mathbb{E}\left[u_k\right] \prod_{j=1}^{k-1} 1 - \mathcal{M}\left(\pi_j\right) = \sum_{k=1}^{i} M_k \mu_k P_k \quad (9)$$

where $M_k \equiv \mathcal{M}(\pi_k)$, and $S_0 = 0$ and P is defined up to and including P_{n+1} and $\mu_k \equiv \mathbb{E}\left[u_{\Psi_k}\right]$ is the expectation of the utility distribution u_ω where ω is the k'th element of Ψ.

Given these definitions, it is clear that $EEU = S_m$ (see Equations 3, 9) and that the list of outcomes sorted by $\{\mu_k\}$ constitute the optimal policy π [17].

5 PROPOSED PROBLEM

In this paper, we define the *elicitation with uncountable single outcome questions problem* $Xi_{infinite}$ by the tuple: (Υ, Q_a) where Υ is a negotiation scenario and Q_a is the uncountable set of all possible questions of the form *"Is $\tilde{u}(\omega) > x$?"* for $x \in (0,1)$ and $\omega \in \Omega$. For the rest of this paper, this kind of question will be represented as: "$\omega?x$". It has two possible answers: *yes* represented as "$\omega \succ x$", or *no* represented as "$\omega \prec x$".

Research literature about preference elicitation in decision support inspired the choice of this type of question [8, 11, 15]. Research about effective elicitation strategies goes back for decades and is carried out under names like search procedures [15], utility elicitation procedures [9], and preference elicitation strategies. There are two general kinds of elicitation strategies in the literature: probability equivalence (PE) and certainty equivalence (CE) methods. Both rely on the notion of a *gamble*. A gamble is defined by a tuple (ω^*, ω_*, p) which means getting outcome ω^* with probability p and ω_* with probability $1-p$. In the special case when $p = 1$, the gamble can be represented by its only possible outcome ω^*.

In both PE and CE methods, the agent is asked to compare two gambles ω and (ω^*, ω_*, p) (i.e., getting ω with certainty or the gamble between ω_* and ω^* with probability p of getting the later). For PE methods, the actor is asked to specify a value for p to make the two gambles equivalent given $\omega^*, \omega_*, \omega$ while for CE methods; she is asked to estimate a value for ω that makes the gambles equivalent given ω^*, ω_*, p.

Different strategies within the PE and CE frameworks differ in how they select these gambles. The most widely used PE method utilizes *extreme gambles* where ω^* and ω_* are chosen to be the absolute and worst outcomes respectively and are assigned a utility of one and zero. These gambles are sometimes called *standard gambles* [6]. Other possibilities include *adjacent gambles* and assorted gambles [9]. CE methods like fractile and chaining methods differ in how they select the endpoints (ω^* and ω_*) as well as the probabilities used in the gambles. For a recent survey, please refer to [9]. *The proposed question form is equivalent to a single PE standard gamble.*

Given the above definition, the solution to $\Xi_{infinite}$ is a selection of an outcome ω^* and cutoff utility level x_{ω^*} defining the question with the maximum value of information. Each such question will have precisely two answers (*yes* and *no*).

6 PROPOSED SOLUTION

Given that we know the initialized optimal policy π, Mohammad and Nakadai [17] showed that the change in EEU if the expected utility of outcome π_k changed from μ_k to μ_k' – without a change to the acceptance model – can be found as:

$$\delta EEU = \begin{cases} M_k P_k \mu_k' - M_k P_k \mu_k & j = k \\ M_k \left(P_j \mu_k' - P_k \mu_k + S_{j-1} - S_{k-1}\right) & j < k \\ \frac{M_k}{1-M_k}\left(P_{j+1}\mu_k' - P_k u_k + S_j - S_{k-1}\right) & j > k \end{cases}, \quad (10)$$

where j is the new location of π_k in the optimal policy after the change.

Let u_k^+, u_k^- be the new utility distribution for outcome π_k given the question q: *"Is $\tilde{u}(\pi_k) > x$?"* for $0 < x < 1$ after receiving answers *yes* and *no* respectively. The following expectation ($\mu_k|_{\pi_k?x}$, μ_k^-, μ_k^+) can then be defined:

$$\begin{aligned} \mu_k^- &= \int_0^1 u_k^-(x)\,dx \\ \mu_k^+ &= \int_0^1 u_k^+(x)\,dx \\ \mu_k|_{\pi_k?x} &= \mu_k^- U_{\pi_k}(x) + \mu_k^+\left(1 - U_{\pi_k}(x)\right) \end{aligned} \quad (11)$$

where $\mu_k|_{\pi_k?x}$ is the expected value for the expected utility of outcome π_k *before* receiving the answer (not to be confused with the EEU of the question).

Proving the following Lemma (Appendix A.1) is easy.

LEMMA 6.1. *Before receiving an answer to the question $\pi_k?x$, the agent can expect no change in the expected utility of that outcome after receiving the answer. (Section 8):*

$$\forall x \in (0,1), \pi_k \in \pi \qquad \mu_k|_{\pi_k?x} = \mu_k.$$

Define $\delta EEU|_{\pi_k?x}$, $\delta EEU|_{\pi_k \succ x}$, $\delta EEU|_{\pi_k \prec x}$ as the change in expected expected utility for the optimal policy after asking *"Is $\tilde{u}(\pi_k) > x$?"*, and after getting *yes* and no as answers. Define $\pi|_{\pi_k \succ x}$, $\pi|_{\pi_k \prec x}$ as the optimal policy after getting the "*yes*" and "*no*" answers respectively and j^+, j^- as the new indices of π_k in $\pi|_{\pi_k \succ x}$, $\pi|_{\pi_k \prec x}$ after getting the corresponding answers.

From Equation 10 and the fact that $\mu_k^+ \geq \mu_k \geq \mu_k^-$, we get:

$$\delta EEU|_{\pi_k \succ x} = \begin{cases} M_k P_k \left(\mu_k^+ - \mu_k\right) & j = k \\ M_k \left(P_j \mu_k^+ - P_k \mu_k + S_{j-1} - S_{k-1}\right) & j < k \end{cases}, \quad (12)$$

$$\delta EEU|_{\pi_k \prec x} = \begin{cases} M_k P_k \left(\mu_k^- - \mu_k\right) & j = k \\ \frac{M_k}{1-M_k}\left(P_{j+1}\mu_k^- - P_k \mu_k + S_j - S_{k-1}\right) & j > k \end{cases}, \quad (13)$$

Solving the *elicitation with uncountable single outcome questions* problem (Section 5) amounts to finding the solution to the following optimization problem:

$$\begin{aligned} k^*, x_{\pi_k}^* &= \arg\max_{k,x} \quad VOI(\pi_k?x), \\ s.t. \quad & 1 \leq k \leq |\pi| \;\wedge\; \alpha_k < x < \beta_k \;\wedge\; VOI(\pi_{k^*}?x_{\pi_k}^*) \geq c, \end{aligned} \quad (14)$$

where c is the cost for asking a question, ($\alpha_k \equiv \arg\max_x u_{\pi_k}(x) \neq 0$, $\beta_k \equiv \arg\min_x u_{\pi_k}(x) \neq 0$) are the limits of u_{π_k}'s support, and:

$$VOI(\pi_k?x) = \delta EEU|_{\pi_k \succ x}\left(1 - U_{\pi_k}(x)\right) + \delta EEU|_{\pi_k \prec x} U_{\pi_k}(x). \quad (15)$$

6.1 General Solution

This section provides a general solution to the problem represented by Equation 14 assuming no specific form for u_k.

LEMMA 6.2. *For any value of x, if both answers to the question "$\pi_k?x$" do not change the index of π_k in the optimal policy, the VOI of this question is exactly zero.*

$$j^+ = j^- = k \rightarrow VOI(\pi_k?x) = 0$$

Lemma 6.2 follows directly from Lemma 6.1 (see Appendix A.2). From Equations 12, 13 and Lemma 6.2, there are only three cases that need to be considered when solving the optimization problem defined in Equation 15 for every value of k, namely: $j^+ < k = j^-$, $j^+ < k < j^-$, $j^+ = k < j^-$.

The objectives $f(x) = VOI(\pi_k?x)$ of the three optimization problems can be formulated – by plugging the appropriate terms from Equations 12 and 13 into Equation 15 – as:

Problem 1 The $j^+ < k$ and $k < j^-$ case:

$$f(x) = \left(1 - U_{\pi_k}(x)\right)\left(z + M_k P_{j^+} \mu_k^+\right) + U_{\pi_k}(x)\left(y + M_k P_{j^-+1} \mu_k^-\right)$$

Problem 2 The $j^+ < k$ and $k = j^-$ case:

$$f(x) = \left(1 - U_{\pi_k}(x)\right)\left(z + M_k P_{j^+} \mu_k^+\right) + U_{\pi_k}(x) M_k P_k \left(\mu_k^- - \mu_k\right)$$

Problem 3 The $j^+ = k$ and $k < j^-$ case:

$$f(x) = \left(1 - U_{\pi_k}(x)\right) M_k P_k \left(\mu_k^+ - \mu_k\right) + U_{\pi_k}(x)\left(y + M_k P_{j^-+1} \mu_k^-\right)$$

where $y \equiv \frac{S_{k-1} - S_k + M_k (S_j - S_{k-1})}{1 - M_k}$ and $z \equiv S_{k-1} - S_k + M_k \left(S_{j-1} - S_{k-1}\right)$.

The above analysis immediately suggests the following algorithm: For each outcome π_k, solve the three problems (Problem i above for $i \in \{1, 2, 3\}$) and find the optimal values ($x_{\pi_k}^i$ for Problem i) value and corresponding VOI values voi_k^i. Let i^* be the solution index leading to maximal VOI value (i.e. $i^* \equiv \arg\max_i voi_k^i$).

The optimal cutoff utility value to ask about for outcome π_k is then $x_{\pi_k}^{i^*}$. Append $\pi_k?x_{\pi_k}^{i^*}$ to the *viable question set* if $x_{\pi_k}^{i^*} \geq c$ and $\alpha_k < x_{\pi_k}^{i^*} < \beta_k$.

The question to ask (solution to Equation 14) is then $\pi_{k^*}?x_{\pi_{k^*}}^{i^*}$ where π_{k^*}, and $x_{\pi_{k^*}}^{i^*}$ are the outcome and corresponding optimal cutoff utility hat correspond to the maximum $voi_k^{i^*}$ in the viable question set.

To find the optimal solution, all combinations of j^+ and j^- must be checked leading to $O\left(|\Omega|^2\right)$ operations. It is possible to find an approximate solution by stopping the search for different j^- values for the each j^+ at the first solution above the cost and halting the search for different j^+ values with the same condition. Section 7 evaluates the accuracy of this approximation.

6.2 Efficient Solution (Uniform Distribution)

The algorithm outlined in Section 6.1 gives a general solution the the proposed elicitation problem (Section 5) for any utility distribution. Nevertheless, solving the three optimizations involved can prove difficult especially for utility distributions that do not lead to convex optimization problems. In this section, we provide a closed-form solution when u_ω^0 is the uniform distribution $\mathcal{U}$.

Practically, a uniform distribution is a natural representation of uncertainty when only limits on the utility value are known. Secondly, the uniform distribution can represent complete ignorance $\mathcal{U}(0, 1)$ which models the case when the agent starts with *no* knowledge of the actor's utility function.

Moreover, computationally the uniform distribution assumption simplifies the problem considerably and is guaranteed by just assuming it for the initial utility distributions u_ω^0. Consider again the form of the questions used in this paper. Answers to $\omega?x$ correspond to multiplying u_ω with a step function. This means that, if u_ω^0 was a uniform distribution, u_ω^n, will always be a uniform distribution independent of the questions asked. Geometrically, the uniform distribution assumption on u_ω^0 amounts to assuming that the real utility function $\tilde{u}_\omega$ lies somewhere in a hyper-rectangle with faces parallel to the axes representing utility values for different outcomes. Asking any question of the form $\omega?x$ removes a part of this hyper-rectangle leaving another hyper-rectangle with faces parallel to the axes in which the true utility function $\tilde{u}_\omega$ still lies.

Under the assumption of a uniform distribution, the three optimization problems can be reduced to the following quadratic problem with linear constraints (see Appendix A.3 for derivation):

$$\begin{aligned} x_{\pi_k}^* &= \arg\max_x \quad ax^2 + bx + c, \\ s.t.\ \alpha_k &< x < \beta_k \ \wedge \ VOI(\pi_{k^*}?x_{\pi_k}^*) > c, \end{aligned} \tag{16}$$

where the values of a, b, c can be found using Table 1 for Problems 1, 2, 3.

LEMMA 6.3. *The unconstrained version of Equation 16 has a single and finite global maximum $x_{\pi_k}^+$ which is shown in Table 1 for Problems 1, 2, 3 (see Appendix A.3 for the proof).*

Based on Lemma 6.3, we know that the objective function will have its unconstrained global maximum within the range (α_k, β_k) or it will be monotonically increasing or decreasing within that range.

The solution to Equation 16 can then be found in three steps: Firstly, find the solution to the unconstrained version $x_{\pi_k}^+$ and the corresponding VOI (Table 1). Secondly, set the tentative optimal solution $x_{\pi_k}^{++}$ as:

$$x_{\pi_k}^{++} = \begin{cases} \alpha_k & x_{\pi_k}^+ \leq \alpha_k \\ x_{\pi_k}^+ & \beta_k > x_{\pi_k}^+ > \alpha_k \\ \beta_k & x_{\pi_k}^+ \geq \beta_k \end{cases}$$

Finally, set $x_{\pi_k}^* = x_{\pi_k}^{++}$ if the two constraints of Equation 16 are satisfied, otherwise, the problem has no solutions. Notice that for our problem, the cases where $x_{\pi_k}^{++} \in \{\alpha_k, \beta_k\}$ correspond to asking questions for which the answer is definitely "*yes*" and "*no*" respectively which is never a good idea. That is why these questions are not included in Q_a and only the case $\beta_k > x_{\pi_k}^+ > \alpha_k$ will actually lead to a solution.

The elicitation algorithm, in this case, will be the same as the one reported in Section 6.1 using the optimization algorithm detailed in this section for solving Problems 1, 2, 3.

7 EVALUATION

This section compares the proposed algorithm against baseline and state-of-the-art elicitation during negotiation algorithms. Two

Table 1: Parameters of the quadratic problem (Equation 16) and the solution to the unconstrained version using the assumption $u^0_{\pi_k} = \mathcal{U}(\alpha_k, \beta_k)$ (where $y = \frac{S_{k-1} - S_k + M_k(S_j - S_{k-1})}{1 - M_k}$, $z = S_{k-1} - S_k + M_k\left(S_{j-1} - S_{k-1}\right)$)

Problem	a	b	c	$x^+_{\pi_k}$	$VOI(\pi_k ? x^+_{\pi_k})$
Problem 1	$\frac{M_k\left(P_{j^-+1} - P_{j^+}\right)}{2(\beta_k - \alpha_k)}$	$\frac{y-z}{\beta_k - \alpha_k}$	$\frac{2z\beta_k + M_k P_{j^+} \beta_k^2 - 2y\alpha_k - M_k P_{j^-+1} \alpha_k^2}{2(\beta_k - \alpha_k)}$	$\frac{z-y}{M_k\left(P_{j^-+1} - P_{j^+}\right)}$	$c - a\left(x^+_{\pi_k}\right)^2$
Problem 2	$\frac{M_k\left(P_k - P_{j^+}\right)}{2(\beta_k - \alpha_k)}$	$\frac{2z - M_k P_k(\alpha_k + \beta_k)}{2(\alpha_k - \beta_k)}$	$\frac{\beta_k\left(2z + M_k P_{j^+} \beta_k + M_k P_k \alpha_k\right)}{2(\beta_k - \alpha_k)}$	$\frac{2z + M_k P_k(\beta_k - \alpha_k)}{2M_k\left(P_k - P_{j^+}\right)}$	$c - a\left(x^+_{\pi_k}\right)^2$
Problem 3	$\frac{M_k(P_{j^-+1} - P_k)}{2(\beta_k - \alpha_k)}$	$\frac{2y + M_k P_k(\alpha_k + \beta_k)}{2(\beta_k - \alpha_k)}$	$\frac{\alpha_k(2y + M_k P_k \beta_k + M_k P_{j^-+1} \alpha_k)}{2(\alpha_k - \beta_k)}$	$\frac{2y + M_k P_k(\alpha_k + \beta_k)}{2M_k(P_k - P_{j^-+1})}$	$c - a\left(x^+_{\pi_k}\right)^2$

baseline algorithms are used: A *full knowledge* algorithm that knows the utility function of the actor (called *ideal* as it represents the ideal case) and a *random* elicitation algorithm that selects an outcome ω and cutoff utility value x_ω randomly to construct questions for the user. Two state-of-the-art algorithms are used: the VOI based *FastVOI* [17] agents; and *Pandora* [3] which casts the elicitation problem as a Pandora's box problem and uses an optimal algorithm for solving it. The optimal and approximate variants of the proposed system are called *Optimal*, and *Approximate* hereafter.

7.1 Realistic Scenarios

The first set of scenarios that we used for comparison are taken from the ANAC competition [14] which started in 2010 and is still running annually. The main advantage of using scenarios from this competition is that they were designed to represent real-world situations and have a large variation in relative competitiveness and other attributes. We used ten bilateral negotiation scenarios from 2012 and 2013 competitions, namely: *Fifty-fifty* (a zero sum scenario with 11 outcomes), *Laptop* (A scenario with high number of win-win outcomes with 27 outcomes), *Flight-booking* (An integrative scenario having both good and bad possibilities with 36 outcomes), *Barter* (80 outcomes), *Outfit* (128 outcomes), *Itex vs Cyperess* (180 outcomes), *Dog Chasing* (270 outcomes), *House Keeping* (384 outcomes), *England/Zimbabwe* (576 outcomes), and *Ice-cream* (1001 outcomes).

In the first experiment, each scenario was run 10 times with the negotiation agent assigned to one of the two sides randomly with 21 different elicitation costs from 0 to 0.1 (210 runs). The initial utility distributions for all outcomes were uniform over the total utility range (0 to 1). The other agent accepted any of the outcomes in the top 25% of its utility value with a probability proportional to that utility. The elicitation agent did not know this information but had an acceptance model that assigned a probability of acceptance equaling the utility for the opponent. This simplified opponent and opponent model were chosen following [4].

Fig. 1 shows the results of this experiment. Fig. 1-a shows the effect of elicitation cost on the final utility obtained by the agent. VOI based algorithms achieve higher utilities than the *random* for all costs and *Pandora* for all costs, and all of these differences are statistically significant except at costs less than 0.010. Moreover, the proposed approach outperforms *FastVOI* at most of the costs. The difference is statistically significant only for costs less than 0.035 (t-test with Bonferroni's conservative multiple-comparisons correction). Comparing the mean utility obtained at all costs reveals a similar picture with the proposed algorithm (both variants) achieving higher utilities than *random* ($p < 1e-10, t = 32.898$) *Pandora* ($p < 1e-10, t = 18.375$), and *FastVOI* ($p < 1e-10, t = 7.0846$) and lower than *ideal* ($p < 1e-10, t = 40.301$). The same test failed to find a difference in the utility achieved by the optimal and approximate variants of the proposed algorithm ($p = 0.304, t = 1.03$). Separately analyzing each scenario, gives the same pattern of performance except for the smallest scenario *Fifty-fifty* for which the improvement above *FastVOI* was not statistically significant for the approximate variant.

Fig. 1-b compares the elicitation time of different algorithms. It is clear that VOI based methods are much slower than *random* and *Pandora* with the approximate variant of the proposed method (*Approximate*) being faster than *FastVOI* for low elicitation costs but more time-consuming than it for costs above 0.04. The exact version of the proposed method (*Optimal*) is always slower than *FastVOI*.

Fig. 1-c compares the number of questions asked by each elicitation algorithm. The proposed method generates fewer questions than other algorithms at low elicitation costs. This difference disappears once the elicitation cost reaches 0.015 after which, the proposed method tends to ask more questions leading to slightly higher utilities (Fig. 1-a) compared with *FastVOI*. It is also possible to directly compare the value of information per question for the proposed method and *FastVOI* (Fig. 1-d). In line with the theoretical analysis, the proposed method chooses questions with higher value of information compared with *FastVOI* at all costs. Fig. 1-e shows the total value of information of all questions, and it is clear , again, that the proposed method outperforms *FastVOI* on this front. The difference between the exact and approximate variants is clearly insignificant in terms of utility, and value of information.

7.2 Randomly Generated Scenarios

Realistic scenarios differed in both the number of outcomes and relative competitiveness of the negotiation which makes it difficult to assess the effects of problem size ($|\Omega|$) on the performance. The second experiment used randomly generated negotiation scenarios to evaluate this effect. The setup was identical to the first experiment but with utility functions sampled randomly from uniform distributions over a variable number of outcomes (10, 50, 100, 500, and 1000). A hard time-limit of $12 \times |\Omega|$ seconds was imposed on all sessions. Fig. 2 shows the results of this experiment.

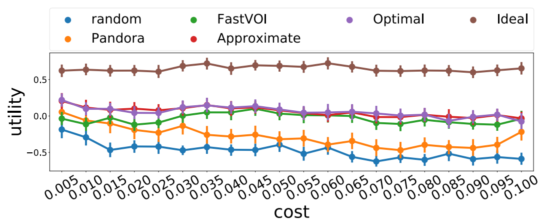

(a) Utility

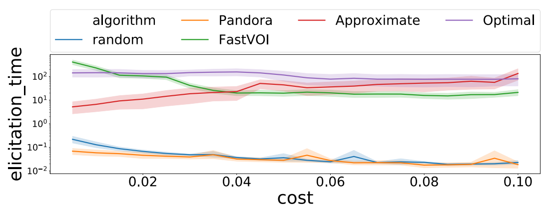

(b) Total Elicitation Time

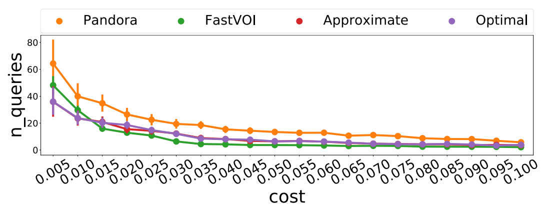

(c) N. Questions

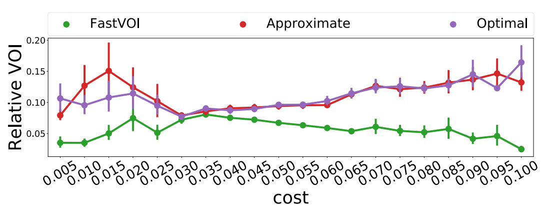

(d) VOI Per Question

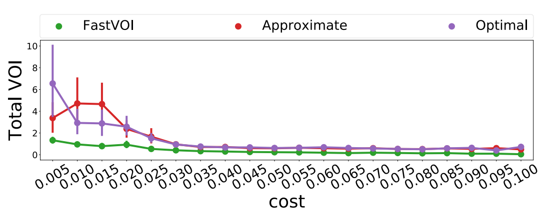

(e) Total VOI

Figure 1: The performance of different elicitation algorithms on realistic negotiation scenarios.

Fig. 2 reports the results of this experiment. With small outcome spaces (10 outcomes), all elicitation algorithms behaved similarly achieving performance comparable to *ideal* (except for *random*). With an increased number of outcomes, the proposed algorithms could consistently achieve higher accuracies compared with other

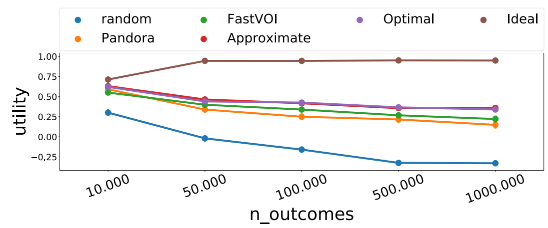

(a) Utility

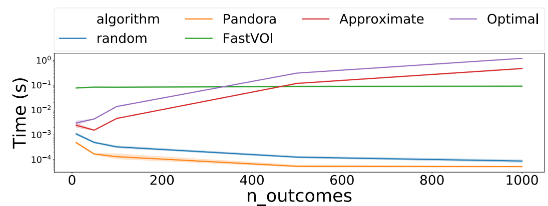

(b) Elicitation Time/Outcome

Figure 2: The performance of different elicitation algorithms on randomly generated negotiation scenarios.

elicitation algorithms (differences were statistically significant for all settings except at 10 and 50 outcomes). Again, there was no discernible difference in the utility obtained by the exact and approximate variants of the proposed algorithm (*Optimal* and *Approximate* in Fig. 2).

Fig. 2-b compares the elicitation time per outcome of different methods. VOI based methods were, in general, more than three orders of magnitude slower than *random* and *Pandora* with the proposed algorithm being faster than *FastVOI* at outcome spaces with less than 500 outcomes and slower than it for outcome spaces with 1000 outcomes. Among VOI based methods, only *FastVOI* scaled well with the number of outcomes. The proposed method has a clear quadratic behavior with *Approximate* providing a constant speedup compared to *Optimal*.

8 DISCUSSION AND LIMITATIONS

The results reported in Section 7 support the theoretical analysis of Section 6. The proposed method can ask more *relevant* questions to the user (Fig. 1-d, Fig. 1-e) with higher expected value of information. Moreover, this improvement in the questions asked is reflected as a *slight* increase in the actual utility obtained by the agent after the negotiation compared with other elicitation algorithms (Fig. 1-a, Fig. 2-a). The increase in utility is small especially for realistic and small synthetic problems. The price paid for removing the need to pre-define the questions and this improved performance is an increase in the required elicitation time for problems with large outcome spaces (Fig. 2-b) or high elicitation cost (Fig. 2-b). The approximate variant of the proposed algorithm was shown to achieve almost the same utility, *total VOI*, and *relative VOI* compared to the exact version at double the speed.

One limitation of the solution presented in this paper — and almost all elicitation during negotiation in the literature [3, 4, 18] — is the inability to use the internal structure of the utility function to guide the elicitation process. Another problem with the proposed methodology is that it implicitly assumes that the opponent will not walk away from the negotiation at any step.

Handling these limitations and analyzing the sensitivity of the proposed method to variations in opponent modeling quality, the initial uncertainty level, the bidding and acceptance strategies among other factors will be directions for future research.

CONCLUSIONS

Utility elicitation during negotiation is an interesting and important research direction that is starting to gain more attention from researchers in automatic negotiation recently. This paper presents a value-of-information based solution to this problem that does not require prior knowledge of a discrete set of questions but can automatically select questions with optimal VOI value from all possible standard gambles. The paper presents a proof of optimality and both a general algorithm valid for any disjoint probability distribution over utility values and an efficient version for the special case of uniform priors. The proposed algorithm was evaluated against the baseline and state of the art elicitation algorithms and was shown to select better questions and achieve higher utilities at the expense of increased processing time for large outcome spaces.

APPENDIX A

A.1 Proof of Lemma 6.1

After receiving an answer to the question $\pi_k?x$, the agent should update u_{π_k} as follows:

In case of a *no*, the new distribution is u_k^- where:

$$u_k^-(y) = \begin{cases} \frac{u_{\pi_k}(y)}{U_{\pi_k}(x)} & 0 \le y \le x \\ 0 & x < y \le 1 \end{cases}.$$

The division by the cumulative distribution U at x is required to normalize the distribution u_k^+.

Similarly, in case of a *yes* answer, the new distribution is u_k^+ where:

$$u_k^+(y) = \begin{cases} \frac{u_{\pi_k}(y)}{1-U_{\pi_k}(x)} & 0 \le y \le x \\ 0 & x < y \le 1 \end{cases}.$$

From Equation 11:

$$\begin{aligned} \mu_k|_{\pi_k?x} &= U_{\pi_k}(x)\int_0^1 u_k^-(x)\,dx\hat{+}\left(1-U_{\pi_k}(x)\right)\int_0^1 u_k^+(x)\,dx \\ &= U_{\pi_k}(x)\int_0^x u_k^-(x)\,dx\hat{+}\left(1-U_{\pi_k}(x)\right)\int_x^1 u_k^+(x)\,dx \\ &= U_{\pi_k}(x)\int_0^x \frac{u_{\pi_k}(x)}{U_{\pi_k}(x)}dx \\ &\quad +\left(1-U_{\pi_k}(x)\right)\int_x^1 \frac{u_{\pi_k}(x)}{1-U_{\pi_k}(x)}(x)\,dx \\ &= \int_0^x u_{\pi_k}(x) = \int_x^1 u_{\pi_k}(x) = u_{\pi_k} \quad \square \end{aligned}$$

A.2 Proof of Lemma 6.2

The case considered is $j^+ = k = j^-$. Substituting from Equations 12 and 13 into Equation 15:

$$\begin{aligned} VOI(\pi_k?x) &= \delta EEU|_{\pi_k>x}\left(1-U_{\pi_k}(x)\right) + \delta EEU|_{\pi_k<x}U_{\pi_k}(x) \\ &= \left(1-U_{\pi_k}(x)\right)M_kP_k\left(\mu_k^+ - \mu_k\right) \\ &\quad +U_{\pi_k}(x)\,M_kP_k\left(\mu_k^- - \mu_k\right) \end{aligned}$$

Let $z \equiv U_{\pi_k}(x)$,

$$\begin{aligned} VOI(\pi_k?x) &= M_kP_k\left(\mu_k^+ - \mu_k - z\mu_k^+ + z\mu_k + z\mu_k^- - z\mu_k\right) \\ &= M_kP_k\left((1-z)\,\mu_k^+ + z\mu_k^- - \mu_k\right) \\ &= M_kP_k\left(\mu_k|_{\pi_k?x} - \mu_k\right) \end{aligned}$$

From Lemma 6.1:

$$VOI(\pi_k?x) = M_kP_k\left(\mu_k|_{\pi_k?x} - \mu_k\right) = 0 \quad \square$$

A.3 Proof of Lemma 6.3

In this appendix, we show the proof for the case of Problem 1.

We start by the general form of Problem 1 after letting $\lambda \equiv U_{\pi_k}(x)$ for brevity:

$$f(x) = (1-\lambda)\left(z + M_kP_{j^+}\mu_k^+\right) + \lambda\left(y + M_kP_{j^-+1}\mu_k^-\right). \tag{17}$$

For a uniform distribution:

$$\lambda = \frac{x-\alpha_k}{\beta_k-\alpha_k}, \qquad (1-\lambda) = \frac{\beta_k - x}{\beta_k-\alpha_k}.$$

We know that $u_{\pi_k} = \mathcal{U}(\alpha_k, \beta_k)$ which leads directly to:

$$\mu_k^+ = \int_x^1 \frac{u_{\pi_k}(x)}{1-\lambda}dx = \frac{\beta_k + x}{2},$$

$$\mu_k^- = \int_x^1 \frac{u_{\pi_k}(x)}{\lambda}dx = \frac{\alpha_k + x}{2}.$$

Substituting into Equation 17:

$$\begin{aligned} f(x) &= (1-\lambda)\left(z + M_kP_{j^+}\mu_k^+\right) + \lambda\left(y + M_kP_{j^-+1}\mu_k^-\right) \\ &= \frac{\beta_k - x}{\beta_k-\alpha_k}\left(z + M_kP_{j^+}\frac{\beta_k+x}{2}\right) + \frac{x-\alpha_k}{\beta_x-\alpha_k}\left(y + M_kP_{j^-+1}\frac{\alpha_k+x}{2}\right) \\ &= \frac{2z\beta_k - 2zx + M_kP_{j^+}\left(\beta_k^2-x^2\right) + 2yx - 2y\alpha_k + M_kP_{j^-+1}\left(x^2-\alpha_k^2\right)}{2(\beta_k-\alpha_k)} \end{aligned}$$

With simple algebraic manipulation, we arrive at Equation 16 and derived the values for a, b, c in the first row of Table 1.

Now consider the a terms in Table 1: The denominator $(2(\beta_k - \alpha_k))$, and M_k are always positive. Moreover, for any two product terms P_i, P_j; $i < j \rightarrow P_i > P_j$. This can easily be seen from the definition of product terms in Equation 8 because $1 - \mathcal{M}(\pi_j) \le 1$. By construction, $j^+ < k < j^- < j^- + 1$, which leads to:

$$a < 0$$

$f(x)$ is a quadratic function with a strictly negative coefficient of x^2 which means that it opens down. Therefore, the unconstrained optimization problem $\arg\max_x f(x)$ has a single global maximum which can be found by completing the square:

$$f(x) = ax^2 + bx + c = (x-h)^2 + k,$$

where $h = \frac{-b}{2a}$ and $k = c - ah^2$.

For Problem 1:

$$h = \frac{-b}{2a} = -1 \times \frac{y-z}{\beta_k - \alpha_k} \times \frac{2(\beta_k-\alpha_k)}{2M_k\left(P_{j^-+1} - P_{j^+}\right)} = \frac{z-y}{M_k\left(P_{j^-+1} - P_{j^+}\right)}$$

The global maximum for $f(x)$ will then occur at $x_{\pi_k}^+ = h$ at which the value of information $VOI\left(\pi_k?x_{\pi_k}^+\right) = f\left(x_{\pi_k}^+\right) = k$). That proves Lemma 6.3 for Problem 1. The proof for Problems 2 and 3 follow exactly the same steps and will not be reported for lack of space. $\square$

REFERENCES

[1] Reyhan Aydoğan, David Festen, Koen V Hindriks, and Catholijn M Jonker. 2017. Alternating offers protocols for multilateral negotiation. In *Modern Approaches to Agent-based Complex Automated Negotiation*. Springer, 153–167.

[2] Tim Baarslag, Alper T Alan, Richard Gomer, Muddasser Alam, Charith Perera, Enrico H Gerding, et al. 2017. An automated negotiation agent for permission management. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 380–390.

[3] Tim Baarslag and Enrico H Gerding. 2015. Optimal Incremental Preference Elicitation during Negotiation.. In *IJCAI*. 3–9.

[4] Tim Baarslag and Michael Kaisers. 2017. The value of information in automated negotiation: A decision model for eliciting user preferences. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 391–400.

[5] Craig Boutilier. 2003. On the foundations of expected expected utility. In *IJCAI*, Vol. 3. Citeseer, 285–290.

[6] Urszula Chajewska, Daphne Koller, and Ronald Parr. 2000. Making rational decisions using adaptive utility elicitation. In *AAAI/IAAI*. 363–369.

[7] Enrique De La Hoz, Ivan Marsa-Maestre, Jose Manuel Gimenez-Guzman, David Orden, and Mark Klein. 2017. Multi-agent nonlinear negotiation for Wi-Fi channel assignment. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1035–1043.

[8] Ward Edwards and F Hutton Barron. 1994. SMARTS and SMARTER: Improved simple methods for multiattribute utility measurement. *Organizational behavior and human decision processes* 60, 3 (1994), 306–325.

[9] Jorge González-Ortega, Vesela Radovic, and David Ríos Insua. 2018. Utility elicitation. In *Elicitation*. Springer, 241–264.

[10] Wenjing Guo, Wenfeng Li, Ye Zhong, Gabriël Lodewijks, and Weiming Shen. 2016. Agent-based negotiation framework for agricultural supply chain supported by third party logistics. In *Computer Supported Cooperative Work in Design (CSCWD), 2016 IEEE 20th International Conference on*. IEEE, 584–589.

[11] Vu Ha and Peter Haddawy. 1997. Problem-focused incremental elicitation of multi-attribute tility models. In *Proceedings of the Thirteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc., 215–222.

[12] Takayuki Ito, Hiromitsu Hattori, and Mark Klein. 2007. Multi-issue Negotiation Protocol for Agents: Exploring Nonlinear Utility Spaces.. In *IJCAI*, Vol. 7. 1347–1352.

[13] Emmanuel Johnson, Jonathan Gratch, and David DeVault. 2017. Towards An Autonomous Agent that Provides Automated Feedback on Students' Negotiation Skills. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 410–418.

[14] Catholijn M Jonker, Reyhan Aydogan, Tim Baarslag, Katsuhide Fujita, Takayuki Ito, and Koen V Hindriks. 2017. Automated Negotiating Agents Competition (ANAC).. In *AAAI*. 5070–5072.

[15] Leslie A Lenert, Daniel J Cher, Mary K Goldstein, Merlynn R Bergen, and Alan Garber. 1998. The effect of search procedures on utility elicitations. *Medical Decision Making* 18, 1 (1998), 76–83.

[16] Hilary Llewellyn-Thomas, Heather J Sutherland, R Tibshirani, A Ciampi, JE Till, and NF Boyd. 1982. The measurement of patients' values in medicine. *Medical Decision Making* 2, 4 (1982), 449–462.

[17] Yasser Mohammad and Shinji Nakadai. 2018. FastVOI: Efficient Utility Elicitation During Negotiations. In *International Conference on Principles and Practice of Multi-Agent Systems (PRIMA)*. Springer, 560–567.

[18] Yasser Mohammad and Shinji Nakadai. 2018. Utility Elicitation During Negotiation with Practical Elicitation Strategies. In *IEEE SMC*.

[19] Ariel Rubinstein. 1982. Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society* (1982), 97–109.

[20] Martin L Weitzman. 1979. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society* (1979), 641–654.