

Learning preferences in electronic negotiations

Rudolf Vetschera
University of Vienna
Austria
Rudolf.Vetschera@univie.ac.at

Abstract

To create an integrative solution in a bargaining problem, negotiators need to have information about the each other's preferences. Empirical negotiation research therefore requires methods to measure the extent to which information about preferences is available during a negotiation. We propose such a method based on the domain criterion, which was originally developed for sensitivity analysis in decision making. Our method provides indices for the amount of preference information that can be inferred both in negotiations achieving a compromise and in failed negotiations. To test the external validity of our proposed measures, we conduct an empirical study which shows that the proposed measures exhibit the positive relationships to the success of negotiations as well as the efficiencies of outcomes that would be expected according to negotiation theory.

1 Introduction

One important aim in negotiations is to achieve a win-win, or integrative solution, which improves the position of both sides beyond the status quo (Sebenius, 1992). According to the widely used Dual Concern model of negotiations (Blake & Mouton, 1964; Pruitt, 1983; Thomas, 1992; Carnevale & Pruitt, 1992), these solutions can be achieved if negotiators have a high concern for both their own and their opponent's outcomes. Consequently, it is necessary for negotiators to develop an understanding of the goals and preferences of their opponents (Keeney & Raiffa, 1991; Sebenius, 1992; Kersten, 2001).

Knowledge about the opponent's preferences is therefore an important element in negotiations. Its impact on the success of negotiations has been studied in the literature since the early 1990s. Thompson and Hastie (1990) studied the different types of judgment errors negotiators might make about the preferences of their opponents, and the learning processes that take place during negotiations to reduce these errors. They distinguished between two types of errors: the "Fixed Sum Error" and the "Incompatibility Error". The fixed sum error is the erroneous assumption that negotiations are inherently a zero-sum game, which is caused by the failure to recognize differences in the importance of issues to the parties. In the incompatibility error, the parties believe to have opposite preferences in an issue, while in fact the preferences about an issue are identical, and there is no conflict at all.

The empirical research of Thompson and Hastie (1990) led to several important results. Many negotiators indeed start negotiations under the "Fixed Sum Error", but it could also be shown that during negotiations, learning takes place and the extent of this error is reduced over time. Furthermore, this study, as well as several later studies (Thompson, 1991; Thompson & DeHarpport, 1994; Arunachalam & Dilla, 1995; Mumpower, Sheffield, Darling, & Miller, 2004) showed that in negotiation where the parties learn more about each other's preferences, the joint outcome is higher.

All these studies used very similar methods to measure knowledge about the opponent's preferences. Preferences were prescribed to experimental subjects in the form of a multi-attribute payoff table, which contained partial utility values for all possible outcomes in each issue being negotiated. After completion of the negotiation experiment, a blank table was given to negotiators and they were asked to fill in the values they believed were contained in their opponent's table. The differences between this table and the actual table given to the opponents was used as an indicator of the negotiators' understanding of their opponents' preferences.

Later studies extended the work of Thompson and Hastie (1990) mainly to analyze different factors which influence the amount of learning taking place during negotiations. For

example, Thompson (1991) studied whether it makes a difference if preference information is directly made available to negotiators by a third party, or if negotiators ask their opponents about their preferences. Thompson and DeHarpport (1994) studied how negotiators change their learning behavior about preferences when they have received different types of feedback in earlier negotiations. In a similar study, Nadler, Thompson, and Van Boven (2003) compared different methods how negotiators might learn how to find out about their opponent's preferences. In a different, but related line of research, Arunachalam and Dilla (1995) found that negotiators learn less about their opponent's preferences when they use a computer-based negotiation support system, than when they negotiate face-to-face.

The method used in these studies to measure knowledge about preferences was put into question by Mumpower et al. (2004), who distinguished between the "Payoff schedule estimation method" used in the earlier studies, and another method called "Holistic estimation method". In this method, subjects are directly asked to guess the utility value an alternative has to their opponent, rather than the different components of their opponent's utility function.

Both methods rely on information about the opponent's true utility function, which is used as a benchmark to measure the quality of a negotiator's preference information. This requirement is not problematic in experiments in which preferences are given to subjects in the form of a known payoff table. But this is not the situation which negotiators face in real life. Furthermore, the extent to which subjects actually follow such prescribed preferences in their negotiation behavior cannot be determined, this might influence empirical results obtained with this method.

From empirical research in decision analysis (Schoemaker & Waid, 1982), it is well known that different methods to elicit a multi-attribute utility function can lead to different results. It is therefore very likely that the direct specification of the opponent's utility function is also not a reliable method to elicit a negotiator's knowledge about the opponent's preferences, and that other methods could lead to other estimates.

The aim of this paper is to develop an alternative method to measure learning about a negotiator's preferences, which does not rely on pre-specified payoff tables or subjective guesses of the opponent's utility function. Rather, we propose a method which is based on an objective measure of the information about preferences which is available during a negotiation. Evidently, such a method can measure only the extent to which learning is possible given the information that could be obtained, and not whether such learning actually takes place. Negotiators might still ignore this information and learn nothing at all about their opponent's preferences. However, in sufficiently large empirical studies, we can expect individual differences in the ability to learn to cancel out, and thus such an objective measure can provide empirical insights into the relationships of preference information and the outcomes of negotiations.

Objective measures of the preference information available during a negotiation can also be useful for the development of negotiation support systems (NSS). An NSS can overcome possible cognitive limitations of negotiators, and thus can help negotiators to utilize all information about their opponent's preferences that becomes available during a negotiation. When an objective indicator about preference information is available, this might also be used to show to the parties that it is necessary to release more information in order to be able to find an efficient solution to the negotiation problem.

The present paper uses concepts from decision making under incomplete information. It thus builds on an analogy between negotiating with an opponent, whose preferences are (partly) unknown, and decision support for a decision maker who is not quite certain about his or her own preferences. While the perspective taken in these two cases is quite different, the fundamental problem of arriving at (or predicting) a decision without knowing the exact values of some parameters of a preference model is quite similar.

To study the external validity of our measurement approach, we perform an empirical study using existing negotiation data. The aim of this study is not to provide new insights into the already well documented relationship between preference information and negotiation outcomes. But by verifying that these relationships can also be established using our measure of preference information, we establish the external validity of our measure.

The remainder of the paper is structured as follows: In section two, we introduce our concept of measuring the preference information which can be learned during negotiations. Section three introduces the hypotheses of our empirical study, which is described in section four. Section five contains the results of this study, section six concludes the paper by discussing the results and gives an outlook on future research questions.

2 Measuring Preference Information

2.1 Overview

We consider a negotiation in which two parties negotiate about K issues. Each offer made by the parties, as well as the final compromise (if one is reached) can therefore be represented as a vector of issue values $\mathbf{x} = (x_1, \dots, x_K)$. The preferences of negotiators over alternatives \mathbf{x} can be represented by a utility function

$$u(\mathbf{x}, \mathbf{w}) \tag{1}$$

where $\mathbf{w} \in W$ is a parameter vector describing the individual preferences of a negotiator. The parameter vector \mathbf{w} can for example contain weights for the attributes, but also additional parameters describing the shape of the marginal utility functions. The set W

is the set of all possible parameter vectors that fulfill formal conditions required for the utility function, for example a scaling condition on attribute weights.

When one negotiator learns about the opponent's preferences, she can rule out certain parameter vectors \mathbf{w} , which would be inconsistent with observed behavior of the other negotiator. Thus the set of possible preference parameters, which the opponent might possess, is restricted to a set $S \subseteq W$. When perfect information about the opponent's preferences is obtained, set S is reduced to a singleton set and the negotiator exactly knows the opponent's utility function.

Thus, the size of set S describes the negotiator's information about the opponent's preferences. The smaller this set is, the better the information about preferences. To measure the extent to which learning about the opponent's preferences is possible in a negotiation, we have to measure the size of set S .

This problem is similar to the problem of sensitivity analysis in decision making. In sensitivity analysis, one studies the extent of a set Q of parameters for which the chosen solution would remain optimal. The larger this set, or the further away its boundaries are from the parameter vector which has originally been used to find the solution, the more robust is the solution.

Several methods have been developed in sensitivity analysis to measure the extent of such parameter sets Q . These methods can roughly be classified into three main categories:

- Single-dimensional distance approaches
- Multi-dimensional distance approaches
- Volume approaches

Single-dimensional distance approaches, e.g. (Mareschal, 1988) consider changes of one parameter at a time and calculate bounds within which a parameter can be changed without altering the optimal solution. This restricted view is not sufficient for our problem, since we need to consider learning about all attributes at the same time.

The method developed by Triantaphyllou and Sanchez (1997) marks the transition between single-dimensional and multi-dimensional approaches. They consider changes only in the weights of single attributes, but determine the most critical attribute, which is the attribute in which the smallest (relative) weight change will cause another alternative to become optimal.

In multi-dimensional distance approaches, one considers changes of several parameters at the same time and calculates how "far" in the sense of some distance measures one can move from the original parameter vector without altering the optimal solution. Different

distance measures have been proposed for this purpose, and also the direction in weight space in which this search is performed varies from method to method.

A very common distance measure is the Euclidean distance (l_2 -norm), which was proposed among others by Barron and Schmidt (1988). Ringuest (1997) extended this model to consider the l_1 and l_∞ norms, which both lead to linear programming models and thus can easily be calculated. Wolters and Mareschal (1995) developed a similar model, also based on the l_1 -norm, for the PROMETHEE outranking method.

These methods start from a given parameter vector, and locate the point on the boundary of set Q in parameter space which is located closest (in the sense of the distance used) to the original parameter vector. When the Euclidean norm is used as distance, this problem can be interpreted geometrically as finding the diameter of the largest hypersphere, which can be inscribed into set Q and which has the original parameter vector at its center. Based on this geometric interpretation, Evans (1984) proposed to use the diameter of the largest hypersphere centered around any point in Q , which can be inscribed into set Q , as a global measure of sensitivity. This measure does not require an initial parameter vector and is therefore more general, and also more closely related to our problem where no a priori information about the opponent's parameter values is available.

A different perspective is taken in the domain criterion introduced by Starr (1962). This criterion does not use a distance, but the volume of the region in parameter space in which an alternative remains optimal to indicate the sensitivity of a solution. The use of this criterion for multi-attribute decision problems was proposed by Charnetski and Soland (1978), and was further developed by Schneller and Sphicas (1983); Eiselt and Langley (1990); Erkut and Tarimcilar (1991); Eiselt and Laporte (1992).

The domain criterion seems to be particularly well suited for our problem. It takes into account the entire region in parameter space in which the chosen alternative remains optimal. Thus a measure based on the domain criterion reflects *exactly* the size of the region in parameter space that is compatible with the preference information available.

2.2 Measurement model

Our measurement model is based on an additive weighting model of preferences of the form

$$u(\mathbf{x}, \mathbf{w}) = \sum_{k=1}^K w_k v_k(x_k) \quad (2)$$

where w_k are attribute weights and $v_k(\cdot)$ is the marginal value function for attribute k . For our model, we use linear value functions of the form

$$v_k(x_k) = \frac{x_k - \underline{x}_k}{\overline{x}_k - \underline{x}_k} \quad (3)$$

where \overline{x}_k is the best and \underline{x}_k is the worst possible value in attribute k for the negotiators whose preferences are modeled. Without loss of generality, one can assume that the negotiator wants to maximize all attributes, so that $\overline{x}_k > \underline{x}_k$, however, this condition is not required for the following analysis. Furthermore, the weights w_k are standardized so that

$$\sum_{k=1}^K w_k = 1 \quad (4)$$

Linearity of the marginal value functions $v_k(\cdot)$ is a rather strong assumption. Conceptually, the approach can easily be extended to nonlinear functions as long as their shape can be described by few parameters. For example, the widely used negative exponential utility function

$$v_k(x_k) = 1 - e^{-\rho_k x_k} \quad (5)$$

requires one parameter ρ_k to specify the shape. However, for our purpose, the simpler function (3) is preferable for several reasons. First of all, the purpose of our model is not to support a negotiator in his or her decision making, but to obtain a measure of how much information about the negotiator's preferences is revealed during the negotiation. Thus it is not necessary to exactly capture the negotiator's preferences, a reasonably good approximation is sufficient. A more general function also requires more parameters. Since only a limited number of observations is available, the number of parameters should be kept as small as possible. This is a definite advantage of the linear form, which requires no additional parameters for the marginal utility functions. Finally, any nonlinear function also entails the risk of mis-specification. Thus it could be that even a function requiring more parameter than the simple linear form (3) would not provide a better approximation to the decision maker's true preferences.

Set W , which represents the state in which no information about preferences is available, is defined by the following conditions on the weights w_k :

$$\begin{aligned} 0 \leq w_k \leq 1 \quad k = (1, \dots, K) \\ \sum_{k=1}^K w_k = 1 \end{aligned} \quad (6)$$

Substituting for w_K , we obtain the following constraint set, which defines a polyhedron in

$K - 1$ -dimensional weight space:

$$\begin{aligned} 0 \leq w_k \leq 1 \quad k = (1, \dots, K - 1) \\ \sum_{k=1}^{K-1} w_k \leq 1 \end{aligned} \quad (7)$$

During the negotiations, preference information is revealed in the form of decisions made by a negotiator. Whenever a negotiator is observed to prefer an alternative (attribute vector) $\mathbf{x}^{(1)}$ over another alternative $\mathbf{x}^{(2)}$, a constraint of the form

$$\sum_{k=1}^K w_k v_k(x_k^{(1)}) \geq \sum_{k=1}^K w_k v_k(x_k^{(2)}) \quad (8)$$

can be generated. Substituting again for w_K , we obtain

$$\sum_{k=1}^{K-1} w_k \left(v_k(x_k^{(2)}) - v_k(x_k^{(1)}) - v_K(x_K^{(2)}) + v_K(x_K^{(1)}) \right) \leq v_K(x_K^{(1)}) - v_K(x_K^{(2)}) \quad (9)$$

During a negotiation, several decisions of a negotiator can be observed and encoded in conditions of the form (9). These decisions involve both the offers a negotiator himself or herself makes, and the acceptance or rejection of offers made by the opponent. In the following conditions, we denote offers made by a negotiator by $\mathbf{s}^{(i)}$, offers made by the opponent by $\mathbf{p}^{(j)}$ and the compromise, if one has been reached, by \mathbf{c} . To simplify the notation, we assume that vectors $\mathbf{s}^{(i)}$, $\mathbf{p}^{(j)}$, and \mathbf{c} are already transformed according to (3) and thus represent utility values rather than the original attribute values.

To formulate the relevant constraints, we distinguish between two cases, depending on whether an agreement has been reached or not. When an agreement has been reached, the compromise can be used as a reference value against which all other offers are compared. It is reasonable to assume that a negotiator will prefer all own offers made during the negotiation to the final compromise (which is the value accepted just at the end of the negotiation), and will prefer the compromise to all previous offers made by the opponent (otherwise it would be rational to return to an offer from the opponent which had already been on the table).

Adding these two types of constraints to (7), we obtain the following set of constraints:

$$\begin{aligned}
\sum_{k=1}^{K-1} w_k \left(s_k^{(i)} - c_k - s_K^{(i)} + c_K \right) &\leq c_K - s_K^{(i)} \quad \forall i \\
\sum_{k=1}^{K-1} w_k \left(c_k - p_k^{(j)} - c_K + p_K^{(j)} \right) &\leq p_K^{(j)} - c_K \quad \forall j \\
\sum_{k=1}^{K-1} w_k &\leq 1 \\
0 \leq w_k \leq 1 & \qquad \qquad \qquad k = (1, \dots, K-1)
\end{aligned} \tag{10}$$

(10) defines a polyhedron in $K - 1$ -dimensional space. We denote the volume of this polyhedron by V_1 . The smaller V_1 , the more information about preferences can be inferred from the decisions of the negotiator.

But V_1 can only be computed when a compromise has been reached during a negotiation. To analyze negotiations in which no compromise has been reached, a different model must be used. By transitivity, model (10) implies that a negotiator will prefer each of his or her own offers to each offer from the opponent. This condition can be directly formulated as a set of constraints, which leads to the following model:

$$\begin{aligned}
\sum_{k=1}^{K-1} w_k \left(s_k^{(i)} - p_k^{(j)} - s_K^{(i)} + p_K^{(j)} \right) &\leq p_K^{(j)} - s_K^{(i)} \quad \forall i, \forall j \\
\sum_{k=1}^{K-1} w_k &\leq 1 \\
0 \leq w_k \leq 1 & \qquad \qquad \qquad k = (1, \dots, K-1)
\end{aligned} \tag{11}$$

We denote the volume of the polyhedron defined by (11) by V_2 . Although model (11) contains more constraints than model (10), the constraints in (10) are tighter because the utility of the compromise has to lie between the utility of the worst own offer and the best offer of the opponent. Thus, V_1 will always be less than or equal to V_2 .

Both measures V_1 and V_2 can be defined for both parties in a negotiation. As long as we consider only linear partial utility functions of the form (3) and assume that the parties want to influence each attribute in the opposite direction (i.e. the negotiator wants to maximize all attributes and the opponent wants to minimize all attributes), the constraints (10) and (11) will be identical for both sides. Thus it is sufficient to consider them for one party in each negotiation.

3 Hypotheses

V_1 and V_2 measure the extent to which learning about the negotiators' preferences is possible during a negotiation. To test their external validity as measures, we analyze whether they empirically show the relationships to other constructs that would be predicted by negotiation theory. In particular, we study their relationships to other, subjective measures of learning as well as to outcome dimensions of negotiations.

One obvious alternative to the measures we have developed in section 2 are subjective measures of learning. Our first hypothesis tests whether V_1 and V_2 measure similar concepts as the negotiator's subjective statements:

Hypothesis H1: There is a positive relationship between measures V_1 and V_2 on one hand and subjective measures of the information learned about the opponent's preferences on the other hand.

In negotiation analysis, learning the preferences and goals of the other party is typically associated with an integrative, rather than distributive, bargaining style (Sebenius, 1992; Kersten, 2001). It is argued that integrative bargaining will make it more likely that an agreement is reached, and also improve the efficiency of agreements. In accordance with the empirical literature (Thompson & Hastie, 1990; Thompson, 1991; Thompson & DeHarpport, 1994; Arunachalam & Dilla, 1995), we formulate the following two hypothesis:

Hypothesis H2: There is a positive relationship between reaching an agreement and the extent to which information on preferences is available as measured by V_2 .

Hypothesis H2 refers only to V_2 , since V_1 can only be computed for negotiations which have reached an agreement.

Hypothesis H3: There is a positive relationship between the efficiency of agreements and the extent to which information on preferences is available as measured by V_1 and V_2 , respectively.

Both hypotheses are rather straightforward. The main aim of our empirical analysis is not to study the widely acknowledged benefits of an integrative bargaining style, but to analyze whether our measures can be considered a valid instrument to measure the learning of preferences that can take place in a negotiation. For this purpose, it is important to use well-established relationships to other constructs. If we fail to find these relationships, it is more likely that our instruments do not measure the appropriate concept than that the proposed relationships do not hold.

It should also be noted that our hypotheses do not specify a direction of causality for the presumed relationships. While it is obvious that the outcomes of a negotiation are the result of the preceding negotiation process, we do not claim that learning of preferences as

measured by our variables V_1 and V_2 is the direct cause of successful negotiations or efficient agreements. It might well be the case that both a high level of learning and an (efficient) agreement are both results of some other properties of an integrative bargaining process. But even if the relationship is only indirect, our measures, which are comparatively easy to obtain, could in future empirical studies serve as useful proxies for other, more complex and harder to measure variables.

4 Method and Measurement

The independent variable in all our hypotheses is the amount of information about preferences that is revealed during a negotiation as measured by V_1 and V_2 . As indicated in (10) and (11), this information depends on the offers made by both sides during the negotiation process. Thus, variables V_1 and V_2 cannot be directly controlled in an experimental setting.

Rather than manipulating V_1 and V_2 in a controlled experiment, our empirical study is based on an analysis of existing negotiation data. The data we are using was collected in electronic negotiations performed with the experimental Negotiation Support System Inspire (Kersten & Noronha, 1999) in the time 1996 to 2004. During this time, about 3,000 negotiations were carried out using Inspire. Inspire keeps a comprehensive log of all negotiations, including all offers made from both sides, the compromise if one was attained, and the utility functions of negotiators.

This large database allows us to perform an ex post analysis of the relationships specified in our hypotheses. While we are not able to directly control the independent variables in our analysis, the size of the database still makes it possible to obtain statistically significant results and to have a sufficient number of observations for all relevant conditions.

The analysis presented in this paper is based on 2,162 negotiation records. Inspire can be used with different cases, for this study we use only negotiations based on the "Cypress/Itex"-case. This case is a bilateral buyer-seller negotiation about the purchase of bicycle parts. The two parties negotiate about four issues: the price of the parts, delivery time, payment terms and the conditions under which defective parts may be returned for results. For each issue, the negotiators can choose from a menu of predefined values. The case specifies five values for price, four for delivery time and three each for the remaining two issues, for a total of 180 possible alternative contracts.

Most experiments using Inspire are set up as negotiations between students at various universities worldwide, involving courses like decision analysis, international negotiations or information systems, in which students take part in electronic negotiations are part of their course assignment. Students are typically credited for participating in the negoti-

ations, independent of the results they achieve. Negotiations are set up between classes in different universities, students negotiate anonymously and do not have direct contact with their opponents except through the systems. Negotiations last for three weeks, and participants are informed that they do not have to reach an agreement.

Inspire can be classified as an “Active facilitative-mediation” type of negotiation support system (Kersten, 2004). It elicits utility functions from users using a modified conjoint analysis method, and throughout the negotiations provides evaluations of offers and graphical representations of the negotiation history using this preference information. Apart from the negotiation logs and the utility functions elicited from negotiators, the Inspire database also contains the answers to two questionnaires, which are administered to subjects before and after the negotiations, respectively.

In total, 2,990 negotiations were set up using the “Cypress-Itex” case in the period under study. However, in some experiments, users did not perform utility elicitations, or at least one side remained inactive and did not make offers. Excluding experiments with this kind of data problems leaves 2,162 usable negotiation records, on which the present analysis is based. Among these 2,162 negotiations, 1,483 (68.6%) resulted in an agreement. 633 agreements (42.7%) were Pareto-optimal, the remaining agreements were dominated.

To calculate the variables V_1 and V_2 , all offers made by both sides, as well as the final compromise if one was reached were extracted from the Inspire database. All offer values were then transformed to the (0,1)-interval using the directions of improvement for each issue and negotiating party as specified in the case instructions. According to the case instructions, sellers should maximize the price, delivery time and the percentage of defective parts that would enable the buyer to return the parts for refund, and should minimize the payment terms. Buyers are supposed to influence each issue in the opposite direction as sellers. The resulting partial utility values in all offers were then entered into the constraint sets (10) and (11) respectively to define the polyhedra of feasible weight vectors.

The negotiation problem contains four issues. Since we substitute for the last weight, the volumes of three-dimensional polyhedra have to be computed. For these computations, a slightly modified version of the algorithm of Lasserre (1983) was used. Since there is only a comparatively small number of values in each attribute, in some instances identical constraints were generated. Although the algorithm of Lasserre is able to deal with redundant constraints in the usual sense (i.e. constraints which entirely lie outside of the feasible set), the exact duplication of constraints causes some faces of the polyhedron to be double-counted. Therefore, such constraints had to be eliminated. Since the volume of the three-dimensional unit simplex defined by (7) is $1/6$, the calculated volumes were multiplied by 6 to obtain a measure which is scaled between 0 (for full information) and 1 (for no information). As already indicated, the constraints are identical for buyers and

sellers, so only one value of V_1 and V_2 was calculated for each experiment.

Hypothesis H1 relates V_1 and V_2 to subjective measures on the information learned about the opponent’s preferences. Two questions from the post-negotiation questionnaire administered to Inspire users were used to obtain subjective measures: The first question asked whether the opponent was considered as informative (INFORM), the other question asked whether the negotiators felt they understood understanding their opponent’s priorities during the negotiation (UNDERST). Both variables were initially measured at the individual level. To obtain a value for the negotiation dyad, answers from both negotiators were added. Descriptive statistics about these aggregated variables are summarized in table 1.

Variable	Scale	Statistics
INFORM	Likert	N = 1592
	1 = informative	M = 4.0931
	5 = uninformative	SD = 1.8293
UNDERST	Likert	N = 1606
	1 = always	M = 3.8825
	5 = never	SD = 1.8293

Table 1: Subjective indicators of preference information used

Not all subjects did completely fill in the post negotiation questionnaire. Only negotiations in which answers from both parties were provided were used in the statistical analysis.

5 Results

5.1 Descriptive Statistics

Table 2 gives an overview of the descriptive statistics for the two volume measures. Since V_1 can be computed only for experiments in which a compromise was reached, we also show V_2 for this subset of experiments to enable a comparison of the two measures.

Figures 1 to 3 show the distributions of these values. In a considerable number of negotiations, ranging from almost 33.78% for V_1 to 49.26% for V_2 using all data, no preference information at all could be inferred from the behavior of the negotiators, and consequently, the value of V_1 or V_2 is equal to one.

For all measures there are some cases in which the volume is zero. A zero value can be obtained for two reasons: the corresponding set is a singleton set, or set S is empty, because the constraints are infeasible. While it might seem strange that these two cases




Figure 1: Distribution of V_1 , negotiations with compromise




Figure 2: Distribution of V_2 , all data

		V_1	V_2
All data	Mean		0.7366
	SD		0.3151
	N		2161
With compromise	Mean	0.5860	0.6844
	SD	0.3567	0.3282
	N	1483	1483

Table 2: Descriptive statistics for volume-based measures




Figure 3: Distribution of V_2 , negotiations with compromise

are considered as equal, this does not constitute a logical contradiction. In the case of a singleton set, perfect preference information has been obtained for the case of linear marginal utility functions. When set S is empty, no weight vector exists which would be compatible with observed behavior and the assumption of a linear utility function. This does not necessarily mean that the negotiator has behaved inconsistently. Observed behavior can still be compatible with a utility function with nonlinear marginal utilities. Thus in this case, additional information about preferences is available: the fact that at least one of the marginal utility functions is nonlinear. This view justifies placing these (few) cases at the same end of the spectrum as the case of full information about linear utilities. This point could be strengthened by assigning a negative value of V to these cases, but any such value would be arbitrary, therefore the value was kept at zero.

5.2 Test of Hypotheses

To test hypothesis H1, we calculated the correlation coefficients between the subjective measures of information and variables V_1 and V_2 . Results of this analysis are shown in Table 3.

		INFORM	UNDERST
V_1	ρ	** -0.0713	-0.0158
	p	0.0044	0.5261
V_2	ρ	-0.0166	** -0.0731
	p	0.5078	0.0034

Table 3: Correlation between volumes and subjective measures

Only two of these coefficients are significant. In both cases, the direction of the relationship contradicts H1: for the volumes as well as the subjective measures, lower values would indicate more information about preferences. Thus hypothesis H1 would predict a positive correlation, while in fact both significant coefficients are negative. Even the significant correlation coefficients have rather low absolute values, thus the relationship between the variables can at best be characterized as very weak. On the other hand, the correlations between the two volumes ($\rho = 0.3800, p < 0.001$) and the two subjective measures ($\rho = 0.6284, p < 0.001$) are both positive and highly significant. We therefore have to reject hypothesis H1 and conclude that the volumes V_1 and V_2 on one hand and the subjective indicators measure rather different concepts.

To test hypotheses H2 and H3, we compared the volume measures V_1 and V_2 for negotiations which reached an agreement vs. negotiations with no agreement, and for efficient and inefficient agreements, respectively. Since the variables are not normally distributed, a nonparametric Wilcoxon test was used to compare the two cases. The same test was also applied to the subjective measures of information, to verify whether they would offer a better explanation of the differences.

Hypothesis H2 involves negotiations in which no agreement is reached. For these negotiations, V_1 cannot be computed, we therefore can only use V_2 . Table 4 shows the results for this case.

In negotiations in which an agreement was reached, volume V_2 is significantly smaller than in negotiations without agreement, indicating that more information about preferences was available in successful negotiations. This confirms our hypothesis H2. In contrast to this result, negotiators in successful negotiations on average indicated significantly *less* understanding about their opponents' preferences than negotiators in failed negotiations. There is no significant difference in the perception of the opponents as being informative

		*** V_2	** UNDERST	INFORM
Agreement	Mean	0.6844	3.9139	4.0223
	SD	0.3282	1.6121	1.7460
	N	1483	1220	1209
No agreement	Mean	0.8506	3.6839	4.3029
	SD	0.2490	1.8624	2.0394
	N	679	386	383
p		< 0.0001	0.0011	0.0906

Table 4: Measures of information for negotiations which reached an agreement vs. negotiations without agreement

between successful and failed negotiations.

		* V_1	** V_2	*** UNDERST	* INFORM
Efficient	Mean	0.5631	0.6554	4.088	4.129
	SD	0.3540	0.3282	1.615	1.710
	N	633	633	616	614
Inefficient	Mean	0.6030	0.7061	3.806	3.939
	SD	0.3580	0.3267	1.648	1.797
	N	850	850	665	656
p		0.0439	0.0021	0.0006	0.0191

Table 5: Measures of information for negotiations with efficient and inefficient agreements

A similar picture emerges for efficient vs. inefficient agreements (Table 5): both volume-based measures indicate that more information about preferences was available in negotiations with an efficient agreement, although the difference is not as marked as for H2. Nevertheless, hypotheses H3 is confirmed by our analysis. On the other hand, both subjective measures indicate the opposite, users subjective felt better informed in negotiations in which the agreement was not efficient.

6 Discussion and Conclusions

Our empirical results indicate that the volume-based measure of learning about preferences, which we have introduced in this paper, exhibits a closer relationship to the fact that an agreement has been reached in a negotiation, and to the efficiency of the agreement, than purely subjective measures of learning. This is a positive sign for the external validity of our approach.

Since our empirical study was based on an ex-post analysis of existing data, we could not compare our approach to the more traditional methods to measure learning of preferences, like the payoff schedule estimation method or the holistic estimation method, since this data was not captured during the experiments. Further empirical studies will be required to compare these methods.

From a theoretical point of view, the two approaches could be seen as complements, rather than as alternatives. These methods can be classified according to two dimensions: The first dimension is the distinction between information about preferences that is available during the negotiation process, and the model of the opponent's preferences which a negotiator builds. Our method measures the information available, while methods like the payoff schedule estimation method refer to the model of the opponent's preferences created from this information.

The second dimension concerns the type of information. We consider preference information which is implicitly available through the decisions a negotiator makes. But preference information could also be provided explicitly, for example when one negotiator tells the other which criteria are more important to her. Figure 4 summarizes these two dimensions.

	Implicit information	Explicit information
Information available	Volume-based approach	Content analysis
Model of opponent	Payoff table estimation Holistic estimation	

Figure 4: Overview of measurement methods

Using methods which directly measure the negotiator's model of her opponent's preferences, it might not be possible to distinguish whether the information underlying the model was obtained in implicit or explicit form. Therefore, the lower row of Figure 4 is not split into the two parts.

Our method only deals with information which is made available implicitly, via offers. The preference information which is explicitly exchanged during a negotiation is part of the communication between negotiators. To analyze it, the content of the communication between negotiators needs to be analyzed. For e-negotiation systems, methods of text-based content analysis have already been successfully used to obtain insights into the structure of negotiations (Srnlka & Koeszegi, 2007), they could also be used for this purpose.

Simultaneous use of the different methods shown in Figure 4 could lead to interesting insights. By combining methods from the upper and lower row of this figure, one could study whether the information on preferences that is available during a negotiation is actually processed by negotiators. A difference between these two measures thus is an

indicator for the efficiency of the negotiator's learning process. Comparing the impact of implicit vs. explicit preference information could offer new insights into the efficiency of different bargaining strategies, in which different types of information are released. However, to perform such studies, controlled laboratory experiments which are explicitly designed for such comparisons are required.

The volume-based approach, which we have developed in this paper, can be used and extended in several directions. A negotiation support system advising one negotiator could use the information collected so far about the opponent's preferences to predict decisions of the opponent, like the acceptability of an offer that the negotiator considers to make. By a suitable partitioning of the remaining parameter set S , similar to the approach proposed in (Vetschera, 2000), one could obtain probabilities that the opponent prefers one offer over another, or the probability that a certain offer is optimal for the opponent. This information could also be used to optimize the offers of a negotiator so that more can be learned from the opponent's responses (Vetschera, 2004). The approach developed in this paper could thus serve as a starting point for further empirical as well as theoretical research.

References

- Arunachalam, V., & Dilla, W. N. (1995). Judgement accuracy and outcomes in negotiation: A causal modeling analysis of decision-aiding effects. *Organizational Behavior and Human Decision Processes*, 61(3), 289-304.
- Barron, H., & Schmidt, C. P. (1988). Sensitivity analysis of additive multiattribute value models. *Operations Research*, 36(1), 122-127.
- Blake, R., & Mouton, J. (1964). *The managerial grid*. Houston: Gulf.
- Carnevale, P. J., & Pruitt, D. G. (1992). Negotiation and mediation. *Annual Review of Psychology*, 43, 531-582.
- Charnetski, J. R., & Soland, R. M. (1978). Multiple-attribute decision making with partial information: The comparative hypervolume criterion. *Naval Research Logistics Quarterly*, 25, 279-288.
- Eiselt, H. A., & Langley, A. (1990). Some extensions of domain criteria in decision making under uncertainty. *Decision Sciences*, 21, 138-153.
- Eiselt, H. A., & Laporte, G. (1992). The use of domains in multicriteria decision making. *European Journal of Operational Research*, 61, 292-298.
- Erkut, E., & Tarimcilar, M. (1991). On sensitivity analysis in the analytic hierarchy process. *IMA Journal of Mathematics Applied in Business and Industry*, 3, 61-83.
- Evans, J. R. (1984). Sensitivity analysis in decision theory. *Decision Sciences*, 15(2), 239-247.

- Keeney, R. L., & Raiffa, H. (1991). Structuring and analyzing values for multiple-issue negotiations. In H. P. Young (Ed.), *Negotiation analysis*. Ann Arbor: University of Michigan Press.
- Kersten, G. E. (2001). Modeling distributive and integrative negotiations – review and revised characterization. *Group Decision and Negotiation*, 10(6), 493-514.
- Kersten, G. E. (2004, July 2004). *E-negotiation systems: Interaction of people and technologies to resolve conflicts* (Internege Research Report No. INR 08/04).
- Kersten, G. E., & Noronha, S. (1999). WWW-based negotiation support: Design, implementation, and use. *Decision Support Systems*, 25(2), 135-154.
- Lasserre, J. (1983). An analytical expression and an algorithm for the volume of a convex polyhedron in \mathbb{R}^n . *Journal of Optimization Theory and Application*, 39(3), 363-377.
- Mareschal, B. (1988). Weight stability intervals in multicriteria decision aid. *European Journal of Operational Research*, 33(1), 54 - 64.
- Mumpower, J. L., Sheffield, J., Darling, T. A., & Miller, R. G. (2004). The accuracy of post-negotiation estimates of the other negotiator's payoff. *Group Decision and Negotiation*, 13(3), 259-290.
- Nadler, J., Thompson, L., & Van Boven, L. (2003). Learning negotiation skills: Four models of knowledge creation and transfer. *Management Science*, 49(4), 529-540.
- Pruitt, D. G. (1983). Strategic choice in negotiation. *The American Behavioral Scientist*, 27(2), 167-194.
- Ringuest, J. L. (1997). Lp-metric sensitivity analysis for single and multi-attribute decision analysis. *European Journal of Operational Research*, 98, 563-570.
- Schneller, G., & Sphicas, G. (1983). Decision making under uncertainty: Starr's domain criterion. *Theory and Decision*, 15, 321-336.
- Schoemaker, P. J. H., & Waid, C. C. (1982). An experimental comparison of different approaches to determining weights in additive utility models. *Management Science*, 28, 182-196.
- Sebenius, J. K. (1992). Negotiation analysis: A characterization and review. *Management Science*, 38(1), 18-38.
- Srnka, K., & Koeszegi, S. (2007). From words to numbers - how to transform rich qualitative data into meaningful quantitative results: Guidelines and exemplary study. *Schmalenbach Business Review*, in print.
- Starr, M. (1962). *Product design and decision theory*. Englewood Cliffs: Prentice Hall.
- Thomas, K. W. (1992). Conflict and conflict management: Reflections and update. *Journal of Organizational Behavior*, 13, 265-274.
- Thompson, L. (1991). Information exchange in negotiation. *Journal of Experimental Social Psychology*, 27, 161-179.
- Thompson, L., & DeHarpport, T. (1994). Social judgement, feedback, and interpersonal learning in negotiation. *Organizational Behavior and Human Decision Processes*, 58, 327-345.

- Thompson, L., & Hastie, R. (1990). Social perception in negotiation. *Organizational Behavior and Human Decision Processes*, 47, 98-123.
- Triantaphyllou, E., & Sanchez, A. (1997). A sensitivity analysis approach for some deterministic multi-criteria decision making methods. *Decision Sciences*, 28(1), 151-194.
- Vetschera, R. (2000). A multi-criteria agency model with incomplete preference information. *European Journal of Operational Research*, 126(1), 152-165.
- Vetschera, R. (2004). Experimentation and learning in repeated cooperation. *Computational and Mathematical Organization Theory*, 9, 37-60.
- Wolters, W., & Mareschal, B. (1995). Novel types of sensitivity analysis for additive MCDM methods. *European Journal of Operational Research*, 81(2), 281-290.