Group Decision and Negotiation **15:** 111–125, 2006.
DOI: 10.1007/s10726-006-9023-0

Preference Structures of Negotiators and Negotiation Outcomes

RUDOLF VETSCHERA
Department of Business Studies, University of Vienna, Bruenner Strasse 72, A-1210 Vienna, Austria
(E-mail: rudolf.vetschera@univie.ac.at)

Abstract

Using a database collected during about 3,000 negotiation experiments, this paper analyzes how specific features of utility functions of negotiators, like attribute weights, monotonicity and convexity of marginal utility functions, are reflected in the outcomes of negotiations. We find that compromise values are to a considerable extent influenced by the utility functions. There are also significant impacts on the likelihood of achieving a compromise, but the model fit is considerably less than for compromise values.

1. Introduction

Many Negotiation Support Systems (NSS) like Inspire (Kersten & Noronha, 1999), Negotiation Assistant (Rangaswamy & Shell, 1997) or Joint Gains (Hämäläinen, 2003) incorporate formal models of the negotiators' preferences. These models are used to support users during negotiations and to identify a fair and efficient solution, which takes into account the users' preferences. This preference-based approach to negotiation support implicitly relies on a complex set of relationships, which encompasses several steps: The preference model used must reflect the true preferences of negotiators, negotiators must behave consistently with the model when conducting negotiations, and their behavior during the negotiation must influence the outcome. If this set of relationships breaks down in any step, the resulting outcome of the negotiation will no longer serve the true interests of the supported parties.

The question whether utility functions or other formal preference models are actually capable of capturing a decision maker's preferences has been studied quite extensively for individual decision makers, and a significant body of literature addresses this question (Winterfeldt & Edwards, 1986). This issue will therefore not be considered here. But the remaining steps, which link the preference model to negotiation results, have rarely been studied, despite of their importance. Mumpower (1991) distinguished several types of negotiation structures in terms of preferences of negotiators, and analyzed their impact on problem complexity and the efficiency of negotiation strategies. Northcraft et al. (1995) analyzed the relationship between non-linearity of marginal utility functions and negotiation outcomes first theoretically and later on (Northcraft et al. 1998) empirically. Stuhlmacher and Stevenson (1997) studied the relationship between negotiators' utility functions and negotiation processes in terms of sequence of offers made and their utility ratings.

In the present paper, we analyze the influence of characteristics of utility functions on negotiation outcomes in negotiations supported by an NSS. Although this is a considerable simplification of the very complex relationships, it provides important insights into the effectiveness of preference-based negotiation support. If negotiation outcomes fail to reflect properties of the preference model, this would cause serious doubt about the possibilities of NSS to influence negotiation outcomes at all, or at least to influence them in a way which is beneficial to their users.

This study is based on the analysis of data collected from experimental negotiations with the NSS Inspire (Kersten & Noronha, 1999) over several years. In these negotiation experiments, utility functionss were elicited from users and recorded along with the protocols of negotiations and their outcomes. This allows us to study the relationship between utility functions and outcomes ex post, in an exploratory way. Compared to the more common approach in negotiation experiments (Perkins et al. 1996; Northcraft et al. 1998; Foroughi et al. 2005), where utility functions are set up by the experimenters and provided to subjects, our approach corresponds more closely to the actual way in which NSS are used. The potential problem of this approach, that elicited utility functions lack the variety in the relevant characteristics required for statistical analysis, is overcome by the size of the database used.

In the following section, we will describe this database, the negotiation case used, and the utility measurement procedure. Section 3 introduces our research hypotheses and describes the measurement of variables. The empirical results are presented in section 4, and section 5 provides a critical discussion and conclusions.

2. Data

The analyses presented in this paper are based on negotiation experiments conducted with the negotiation support system Inspire (Kersten & Noronha, 1999). Using the classification of Kersten (2004), Inspire can be described as an active facilitative-mediation system. It provides not only communication features for the interaction between negotiators, but also analytical components to evaluate offers and counter-offers and to monitor the progress of the negotiations. In Inspire, this support is based on Multi-Attribute Utility Theory (Keeney & Raiffa, 1976).

Inspire has been available on the Internet since 1996. For this research, we use data collected between October 1996 and September 2004. Although Inspire can be used with different negotiation cases, all negotiations analyzed here are based on a single case, the "Cypress-Itex" case.

2.1. Negotiation case

The "Cypress-Itex" case is a buyer-supplier negotiation about the purchase of bicycle parts. The two parties negotiate about for attributes of a purchase contract: the price of the parts, delivery time, terms of payment and the conditions under which defective parts may be

Table 1. Attributes and possible values in the Cypress-Itex case.

Attribute	Possible values
Price	3.47$, 3.71$, 3.98$, 4.12$, 4.37$
Delivery time	20, 30, 45, 60 days
Payment terms	Payment 0, 30, 60 days after delivery
Returns	Items may be returned for refund when... Any part is defective 5% are defective 10% are defective

returned by the buyer for refund. The case specifies discrete values for each attribute, which are shown in Table 1.

Since there are five possible values for price, four for delivery time and three each for the two other attributes, the two parties have a total of 180 alternative contracts to choose from.

The case description states the direction of improvement for each attribute and each role. Buyers are instructed to minimize price, delivery time, and the percentage of defective parts which allows them to claim refund, and to maximize payment terms. Sellers are given instructions to influence the attributes in the opposite direction. Apart from these directions, the case description does not specify preferences. In particular, it does not contain information about the importance of the attributes (weights), nor does it specify how the particular values within each attribute should be rated. Thus, for example, a buyer is free to consider a payment term of 30 days to be almost as good as 60 days, or almost as bad as 0 days.

2.2. *Negotiation process and data collection*

Negotiations in Inspire follow a well-defined pattern: first, negotiators individually study the case description and perform an elicitation of their multi-attribute utility functions. Users are free to repeat the utility elicitation later on during the negotiation.

After completing the utility elicitation, users fill in a pre-negotiation questionnaire, in which demographic data, their expectations and reservation levels for the upcoming negotiation, as well as their experiences with the utility elicitation process are recorded. Users then negotiate with their partners by exchanging offers (in terms of the four attributes) and free text messages. There is a time limit of three weeks for the negotiations. Both parties are free to terminate the negotiations before that time expires and without reaching an agreement.

When the parties have reached an agreement, the system checks whether the agreement is Pareto optimal in utility space. When it is dominated, the system proposes other alternatives which dominate it and negotiations may be continued. Since we are mainly concerned with the direct impact of utilities on negotiation outcomes, this "post-settlement" phase is not included in our analysis. Finally, the subjects fill in a post-negotiation questionnaire, in which their satisfaction with the negotiation process and their attitudes towards the system are recorded.

The resulting database consist of four tables: One table each for the pre- and post-negotiation questionnaire, where one row is created for each user who fills in a questionnaire. A third table contains one row for each utility elicitation of a user and stores the partial utility values for all 15 attribute values shown in Table 1. In the last table, all offers sent (or accepted) and the corresponding attribute values are recorded. All table entries are time stamped and cross-linked via user and negotiation numbers, so that data referring to one experiment can be matched for statistical analysis.

In total, 2,990 negotiations using the Cypress-Itex case were set up in the time frame studied here. Almost all of the negotiators (5,977) filled in the pre-negotiation questionnaire. 2,808 negotiations actually started, i.e. at least one offer was sent. In 558 negotiations (19.87%), one side remained inactive. In the remaining 2,250 negotiations, which were used for this study, an agreement was reached in 1,626 experiments (72.3%).

2.3. Utility measurement in Inspire

Utility elicitation in Inspire is based on Conjoint Measurement (Green & Srinivasan, 1978). Each possible value in each of the four attributes is assigned a partial utility value, and an additive function is used to aggregate partial values across all criteria. A three-step approach is used to elicit the partial values (Kersten & Noronha, 1999):

In the first step, users are asked for a rating of the relative importance of the attributes (issues). The attribute weights obtained from this rating serve as upper bounds on the partial utility values in each attribute. In the second step, users specify partial values for the individual attribute values. In a third step, users are shown selected packages containing values for all attributes and the utility rating of these packages based on the data elicited so far. The system then asks the users to holistically evaluate the packages and modify the utility ratings. A least squares method is used to fit the partial utility values to the holistic ratings.

Although the conjoint measurement approach does not explicitly distinguish between attribute weights and marginal utility functions, it is equivalent to a more common additive multiattribute utility function of the form (Keeney & Raiffa, 1976)

$$u(X) = \sum_{i=1}^{4} w_i u_i(x_i) \tag{1}$$

where $X = (x_1, x_2, x_3, x_4)$ is the vector of attribute values of an offer in the four attributes, w_i is the weight and $u_i()$ is the marginal utility function of attribute i.

By conceptually separating the utility function into a weight and a marginal utility function for each attribute, we obtain a more detailed picture of the preference structure of a negotiator. Specifically, we consider the following characteristics of a negotiator's utility function:

– The *weight* of each issue. Weights represent the importance of an issue to the negotiator. Strictly speaking, a weight represents the utility increase when an attribute is changed from its worst to its best value. But since the attribute ranges were identical for all our subjects, we can safely state that a negotiator who has a higher weight in an attribute also considers that attribute more important.
– The *monotonicity* of the marginal utility functions. The case description specifies a direction of improvement for each attribute. Consequently, the marginal utility functions should be strictly monotonic. But the system does not force the elicited utility functions to conform to the case description. Thus non-monotonic utility functions or even functions which are monotonic in the "wrong" direction can occur.
– The *shape* of the marginal utility functions. Since the case does not involve risky outcomes, the marginal utility function should be concave to represent the usual economic assumption of decreasing marginal benefits. But this assumption is also not enforced by the system and marginal utility functions can be linear or convex.

These items are very similar to the classification of Mumpower (1991), who classified utility functions according to the weights, the functional form (convex, linear, or concave), and the organizing principle, which is the method by which evaluations in different attributes are aggregated. Since Inspire always uses an additive function, the last item in Mumpower's classification is not used here.

Data about utility functions is available from 5,143 users, who performed a total of 7,375 utility function elicitations. Thus, each user performed on average about 1.43 utility elicitations. However, as Table 2 shows, the distribution of utility elicitations is very uneven across users. While a large majority of users performed only one elicitation, some users performed as many as 49 utility elicitations.

One possible explanation for the repeated elicitation of utilities is that users detected a mistake and tried to correct it. If this is the case, we would expect repeated utility elicitations to take place at rather short intervals. This is not always the case. While about one third (34.41%) of subsequent utility elicitations took place within ten minutes, an even

Table 2. Number of utility elicitations performed.

Elicitations	*N*	Percent
1	3961	77,02
2	718	13,96
3	244	4,74
4	101	1,96
5	50	0,97
6	29	0,56
7	18	0,35
8	7	0,14
9	4	0,08
10+	11	0,21

larger number of users (42.13%) waited for one week or more before changing their utility functions. To avoid the possible effects of user errors, only the last utility function elicited from each user is used in the following analyses.

3. Hypotheses and Measurement

The main aim of this paper is to study the effects of negotiator preferences on negotiation outcomes. In the following hypotheses, we will argue how we expect utility functions to look like, and how we expect different types of utility functions to influence the outcomes of negotiations.

The outcome of a negotiation can be measured in different terms. One important characteristic is the fact that negotiators have achieved a consensus, or the negotiations have broken up. Different preference structures of negotiators could lead to a more friendly or more hostile behavior during the negotiations and thus affect the likelihood of agreement. This measure was used in several empirical studies on negotiations (Coursey 1982; Moore et al. 1999; Mosterd & Rutte, 2000; Neale & Bazerman, 1985).

Other possible measures of negotiation success relate to payoff levels and include the sum of the payoffs achieved by the two negotiators (Foroughi et al. 2005) or Pareto optimality of the compromise (Barkhi et al. 1999). Since we use characteristics of the individual utility functions as explanatory variables, we consider payoffs at the level of individual negotiators and attributes, rather than these aggregate measures.

3.1. Weights

The case description provided to subjects does not indicate any specific ranking of the negotiation issues. In decision analysis, there are also no typical assumptions about weights. Thus we do not expect any particular structure of the attribute weights.

Weights are therefore purely subjective values. In the context of our experiments, “subjective” refers not only to individual subjects, but also to the role subjects play in the experiment. There is considerable empirical evidence about value differences between buyers and sellers of a good (Kahneman et al. 1990). If subjects take the case seriously and adopt the roles of buyers or sellers, this difference should also be visible in their utility functions. If such differences exist, they can therefore be seen as an indicator that the case description made subjects really think of themselves as “Buyers” and “Sellers”. Therefore we formulate:

Hypothesis H1: Weights will be different for subjects assigned the roles of buyers and sellers.

The subjective importance of attributes, which is reflected in the attribute weights, should also be reflected in the negotiation outcomes. Therefore we formulate:

Hypothesis H2: Negotiators who have a higher weight in a particular attribute will obtain a better result in that attribute than users with lower weights.

Hypothesis H2 is based directly on the interpretation of weights as indicators of importance. A negotiator for whom an attribute is important will bargain harder with respect to that attribute, and consequently should achieve a better result in the attribute.

Weight vectors can also be used to measure conflict between negotiators, which in turn might influence their ability to reach a compromise. When marginal utility functions are linear and identically scaled, the weight vector is identical to the gradient of the utility function. Thus we can use the weight vector as an approximation of the gradient of a nonlinear utility function. When the gradients of the utility functions of the two negotiators are (close to) orthogonal, improving the utility of one side will leave the utility of the other side almost unchanged, thus there is a low level of conflict. On the other hand, conflict is strong when the gradients point into opposite directions. We can therefore use the scalar product of the weight vectors as a measure of conflict between parties. A similar argument was put forward by Mumpower (1991), who considered negotiations more difficult when negotiators have high weights for the same attributes. Since negotiations involving strong conflict are inherently more difficult, we propose:

Hypothesis H3: Negotiations in which the level of conflict, as measured by the scalar product of weight vectors, is low, will more often lead to an agreement than negotiations with a high level of conflict.

3.2. Monotonicity

The case description indicates the direction of improvement for each attribute given the role of buyer or seller. These directions should be reflected in the marginal utility functions for the attributes. Thus we formulate:

Hypothesis H4: Users will have strictly monotonic marginal utility functions in each attribute, corresponding to the direction of improvement indicated in the case description for their role.

Hypothesis H4 is based on two assumptions: that users will actually form preferences according to their role assigned in the case, and that users are able to correctly encode these preferences in the form of a utility function. Since we are not able to observe the "true" preferences of users separately from their encoding in the utility function, these two elements cannot be separated.

If users try to influence an issue in the "wrong" direction (e.g. a buyer wishing to maximize instead of minimizing price), we expect the outcome of such negotiations to favor their opponent. Thus we expect the following relationship between monotonicity and negotiation outcomes:

Hypothesis H5: A negotiator who does not exhibit a monotonic marginal utility function in the correct (according to the case description) direction for some attribute, will have a worse outcome in that attribute (in terms of the correct direction).

Consequently, when "wrong" preferences of one user lead negotiators to move in the same direction, the level of conflict decreases and an agreement should become more likely. Therefore, we formulate:

Hypothesis H6: Negotiations will more often lead to an agreement when the direction of preference is the same for both negotiators in some attributes.

To test these hypotheses, marginal utility functions were classified as strictly or non-strictly decreasing or increasing, or non-monotonic. These features were encoded on a five-point scale with end points -2 for strictly monotonically decreasing and $+2$ for strictly monotonically increasing.

3.3. Shape

The standard economic assumption of decreasing marginal utility would predict a concave shape of marginal utility functions. However, in a multi-attribute context, it is often argued that preferences need not follow that assumption (Keeney & Raiffa, 1976). To test whether this is really the case for our data, we formulate:

Hypothesis H7: Users have concave marginal utility functions in all attributes.

The shape of the utility function has important consequences for negotiations. When utility functions are convex, negotiators will lose considerably in terms of utility when they move away from their best value in an attribute, while for negotiators with a concave utility function, the same concession will lead to smaller losses. Thus it is often argued (Mumpower 1991; Northcraft et al. 1995, 1998) that the presence of negotiators with convex utility functions will make it more difficult to reach a compromise. Furthermore, since such negotiators are bound to negotiate toughly, we expect them to achieve better outcomes. Therefore, we formulate the following two hypotheses:

Hypothesis H8: Negotiators, who have a convex marginal utility function in an attribute, will achieve a better outcome in that attribute than negotiators who have linear or concave marginal utility functions.

Hypothesis H9: Negotiations will less often lead to an agreement when negotiators have convex marginal utility functions.

Convexity of a function is a local property. To obtain a measure of the shape of a marginal utility function across its entire domain, an approach similar to Pennings and Smidts (2003) was used: After scaling the attribute range to the (0,1) – interval, a negative exponential utility function of the form

$$g(x) = \frac{1 - e^{-\rho x}}{1 - e^{-\rho}} \tag{2}$$

was fitted to the utility values using a least-squares method. For parameter values $\rho < -0.01$, the function was classified as convex, for $\rho > 0.01$ as concave and for $0.01 \leq \rho \leq 0.01$ as linear.

4. Results

The hypotheses formulated above relate the structural properties of utility functions to the possibility of achieving an agreement and the compromise values of the attributes, if an agreement is reached. Since we analyze the effect of several factors on these outcome dimensions simultaneously, we discuss the results by outcome dimensions rather than by influence factors after providing a brief descriptive overview of the structural properties of the utility functions.

4.1. Structural properties

Table 3 compares the attribute weights of buyers and sellers for the four attributes. A nonparametric Wilcoxon test indicates that differences between the two roles are significant at $p < 0.0001$ for all four attributes, thus confirming our hypothesis H1.

The weights exhibit systematic patterns: For both roles, price is the most important attribute and receives a weight which is about twice as large as for the other attributes. Sellers assign higher weights than buyers to the attributes price and payment. Thus, buyers seem to be more interested in quality-related attributes, while sellers put a stronger focus on the financial aspects of the deal. A similar result was obtained by Stuhlmacher and Stevenson (1997), who found that for sellers, price was almost always the most important attribute, while for 21% of buyers, warranty, which is a quality-related attribute, was most important.

While the overwhelming majority of subjects exhibited monotonic (or strictly monotonic) utility functions corresponding to the case description, a considerable number of users deviate, as Table 4 shows.

There is again a significant difference between the roles of buyers and sellers. Sellers have significantly more monotonicity errors in the attributes price, delivery and returns, while buyers make more errors concerning the attribute payment. This coincides with the directions of improvement as specified in the case description. In all four attributes, the role

Table 3. Attribute weights for buyers and sellers.

Role		Price	Delivery	Payment	Returns
Buyer	Mean	0.3820	0.2311	0.1768	0.2102
	SD	(0.1298)	(0.0948)	(0.0861)	(0.1054)
Seller	Mean	0.4067	0.1705	0.2297	0.1931
	SD	(0.1301)	(0.0869)	(0.0905)	(0.0997)

Table 4. Fraction of users with incorrect monotonicity (in %).

Role	Price	Delivery	Payment	Returns
Buyer	10.94	9.18	19.10	8.08
Seller	23.83	34.09	12.77	16.31
χ^2	148.40	469.12	38.55	81.22
p	<0.0001	<0.0001	<0.0001	<0.0001

Table 5. Distribution of shapes (in %).

		Price	Delivery	Payment	Returns
Buyer	Convex	36.20	55.45	21.22	40.59
	Linear	0.43	1.14	24.20	21.22
	Concave	63.37	43.41	54.59	38.20
Seller	Convex	31.93	16.54	22.14	18.63
	Linear	0.27	0.04	21.44	23.41
	Concave	67.80	83.42	56.42	57.96
χ^2		11.66	893.29	5.55	317.16
p		0.0029	< 0.0001	0.0622	< 0.0001

who should maximize an attribute exhibits more errors than the role who should minimize the attribute.

Table 5 shows that there is a considerable amount of users with convex utility function, in one case, their share even exceeds 50%. Another interesting feature is the fraction of linear utility functions: in attributes in which only three values were available, a large fraction of users assigned a utility of 0.5 to the middle value, leading to a linear utility function. In attributes in which more values are available, the share of linear utility functions is negligible. Role differences are significant for all attributes except payment.

4.2. Compromise values

Hypotheses H2, H5 and H8 stated that negotiators will perform better in a given attribute when they have a high weight in that attribute, and when their marginal utility function is monotonic in the correct direction and convex. To test these hypotheses, we estimated linear regression equations on the compromise values in the four attributes. Since the compromise is jointly determined by both parties, the relevant properties of the utility functions of both parties were simultaneously used as explanatory variables. The results are shown in Table 6.

The regression analysis confirms most hypotheses. Higher attribute weights of buyers have a negative influence on the compromise values in the attributes price, delivery time and returns, which are to be minimized by buyers, and a positive influence on the attribute payment terms. For sellers, who have the opposite preferences, the influence of attribute weights is in the other direction. Compared to the other factors, attribute weights have the strongest impact on compromise values.

The impact of convexity of the marginal utility functions is less pronounced. While in most cases, hypothesis H8 is confirmed, there is no significant influence on outcomes in the attribute delivery for sellers, and the impact on attribute returns for sellers barely reaches the 5% significance level. The highly significant impact for buyers on attribute payment terms works in the opposite direction: buyers who have a convex utility function end up with a shorter payment period, which is a disadvantage for them.

Table 6. Influence factors on results.

Property	Role		Price	Delivery	Payment	Returns
Weight	Buyer	β	−0.4901	−33.8635	62.8115	−12.3701
		t	−14.0200	−14.2100	12.9200	−18.2300
		p	< .0001	< .0001	< .0001	< .0001
	Seller	β	0.3176	24.0434	−44.7200	9.3502
		t	9.1400	9.0900	−9.6300	12.6700
		p	< .0001	< .0001	< .0001	< .0001
Convexity	Buyer	β	−0.0471	−1.7474	−3.3958	−0.8971
		t	−4.7200	−3.7600	−3.2000	−5.9600
		p	< .0001	0.0002	0.0014	< .0001
	Seller	β	0.0286	−0.4058	−4.1724	0.3923
		t	2.8600	−0.6300	−4.0600	2.0200
		p	0.0043	0.5314	< .0001	0.0436
Monotonicity	Buyer	β	0.0290	1.6116	3.6320	0.6525
		t	4.6000	4.9600	10.5200	6.6000
		p	< .0001	< .0001	< .0001	< .0001
	Seller	β	0.0502	2.6286	4.5206	0.7040
		t	10.6400	13.1900	10.4700	10.3700
		p	< .0001	< .0001	< .0001	< .0001
		R^2	0.2439	0.2477	0.2521	0.3211

The parameters for monotonicity are all significant and have the expected signs: when parties want to maximize an attribute, the compromise value in that attribute increases. However, this effect might counteract the other effects. When a buyer assigns a high weight to attribute price, but wants to maximize price instead of minimizing it, the compromise value should be a higher. To correct for this problem, the same regressions (without monotonicity as explanatory variable) were also run on a subset of data containing only cases in which both sides had correct monotonicity. The results, which are not shown here because of limited space, were qualitatively identical to those of Table 6, except that the influence of convexity on the attribute delivery was insignificant also for buyers.

4.3. Agreement

According to hypotheses H3, H6 and H9, negotiations will more often end in an impasse when the level of conflict between negotiators, as measured by the scalar product of weight vectors, is high, when all negotiators have monotonic utility functions in the correct direction, and when at least one negotiator has a convex utility function in at least one attribute.

To test the hypotheses, a logistic regression model with impasse as the dependent variable was estimated. To test for interaction effects between convexity of utility functions of both negotiators, the number of attributes in which both negotiators have a convex marginal utility

Table 7. Factors influencing impasse (logistic regression).

Factor	Parameter	χ^2	p
Scalar product of weights	2.7444	7.4140	0.0065
Number same direction	−0.1724	5.1138	0.0237
N. Convex utilities buyer	0.1735	12.5093	0.0004
N. Convex utilities seller	0.1359	4.9038	0.0268
N. Convex utilities both	−0.00217	0.0005	0.9816
$r^2 = 0.0300$			

function was also included as an independent variable. The results are shown in Table 7. In interpreting this table, it should be noted that the the dependent variable represents the probability of not achieving an agreement.

The general fit of this model was rather low. Nevertheless, the hypotheses could be confirmed. The scalar product of weights significantly increases the probability of failure and has the strongest impact of all factors studied. Convex utilities also have a positive impact, which is highly significant for buyers and still significant (at p=0.0268) for sellers. However, it seems that it is sufficient to deteriorate the negotiations if one negotiator has a convex utility function, the interaction between negotiators has no significant effect. Surprisingly, the effect of both negotiators having preferences for an attribute in the same direction was only significant at $p = 0.0237$.

The reason for this weak impact is further analyzed in Table 8. An increase in the relative number of agreements occurs only when both parties want to influence three or four attribute in the same direction. But since such massive deviations from the case description occur only very infrequently, the whole effect, even if considered independently of the other factors, is barely significant ($\chi^2 = 9.4469$, $DF = 4$, $p = 0.0508$).

Table 8. Agreement when both parties have identical preferences.

		Agreement	
Attributes with same direction		No	Yes
0	*N*	520.00	1178.00
	%	30.62	69.38
1	*N*	130.00	289.00
	%	31.03	68.97
2	*N*	21.00	47.00
	%	30.88	69.12
3	*N*	1.00	20.00
	%	4.76	95.24
4	*N*	2.00	15.00
	%	11.76	88.24

Table 9. Summary of results.

Structural properties	Role effects, predictions	Effects on	
		Compromise	Agreement
Weights	H1 confirmed: Significant role differences	H2 confirmed: Better outcomes for higher weights	H3 confirmed: Conflict has negative impact
Monotonicity	H4 rejected: Some contradictions with role	H5 confirmed: Outcome reflects direction	H6 weakly confirmed: More agreements when same direction
Shape	H7 rejected: 20%–55% of utility functions convex	H8 confirmed: Better outcomes when utility function is convex	H9 weakly confirmed: Convex utilities lead to less agreements

5. Discussion and Conclusions

Table 9 summarizes the results of our analysis. The analysis of structural properties of the negotiators' utility functions has shown that utility functions are much more diverse than one would commonly assume.

The fact that a considerable number of subjects had utility functions which contradicted the direction of improvement for an attribute as specified in the case description causes some concern. Experiments using text-based cases are a common tool in negotiation research. When a significant number of subjects exhibit preferences which substantially deviate from their assigned role and the case description, this factor has to be taken into account when analyzing negotiation outcomes. If negotiation outcomes are analyzed only in terms of the case description as it is intended by the experimenters, wrong conclusions might be drawn.

One could argue that the problem lies in the encoding of preferences, and not in the actual preferences which negotiators were following. If subjects actually behaved as specified in the case description, but were not able to correctly represent their preferences in the form of utility functions, the problem would vanish. However, our results do not support this point of view: when subjects specified the "wrong" monotonicity in their utility functions, they also behaved according to their utility functions, and not according to the role they were assigned, as the results concerning hypotheses H5 and H6 clearly show.

While this result indicates potential problems in experimental negotiation research, it conveys a positive message for preference-based decision and negotiation support: even users who misinterpreted the case description were able to correctly encode their (incorrect) preferences in the form of utility functions. This indicates that the elicitation of utilities was not too difficult for subjects.

The consistent impact of preference structures on negotiation outcomes is also an encouraging result. Negotiator preferences captured in the form of a utility function do play a significant role in negotiations. Therefore an approach to negotiation support, which helps negotiators to specify their preferences even more precisely and apply them more consistently throughout the negotiation process, is not a purely theoretical, unrealistic concept, but builds on a solid basis already present in negotiations.

While the empirical evidence concerning the impact of preference structures on negotiation outcomes is very strong, results concerning agreements are less satisfactory. Although we were able to confirm all the hypotheses related to agreement, the overall fit of the model was rather low. This means that other factors not included in our model have a strong impact on whether an agreement can be reached. For example, factors like relationship building during the negotiation process (Koeszegi et al. 1904) could play a more important role for reaching an agreement than preferences, which relate only to the task-oriented outcome of the negotiation.

This fact points to one limitation of our study: by focusing on preferences encoded in multi-attribute utility functions, other characteristics of the negotiators, like their culture or personality, which could also influence the outcome of negotiations, are ignored. This study was also limited to one specific preference model, additive utility functions, and one type of utility elicitation. But even within these limitations, it clearly shows that formalized preference models can be an important factor in negotiations.

References

Barkhi, R., Jacob, V. S., and Pirkul, H. (1999). "An Experimental Analysis of Face to Face Versus Computer Mediated Communication Channels," *Group Decision and Negotiation* 8(4), 325–347.

Coursey, D. (1982). "Bilateral Bargaining, Pareto Optimality, and the Empirical Frequency of Impasse," *Journal of Economic Behavior and Organization* 3, 243–259.

Foroughi, A., Perkins, W. C, and Jessup, L. M. (2005). "A Comparison of Audio-Conferencing and Computer Conferencing in a Dispersed Negotiation Setting: Efficiency Matters!" *Journal of Organizational and End User Computing* 17(3), 1–26.

Green, P. E. and Srinivasan, V. (1978). "Conjoint Analysis in Consumer Research: Issues and Outlook," *Journal of Consumer Research* 5, 103–123.

Hämäläinen, R. P. (2003). "Decisionarium – Aiding Decisions, Negotiating and Collecting Opinions on the Web," *Journal of Multi-Criteria Decision Analysis* 12(2), 101–110.

Kersten, G. E. (2004). "E-Negotiation Systems: Interaction of People and Technologies to Resolve Conflicts," *Interneg Research Paper INR 08/04*.

Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1990). "Experimental Tests of the Endowment Effect and the Coase Theorem," *Journal of Political Economy* 98(6), 1325–1348.

Kersten, G. E. and Noronha, S. J. (1999). "WWW-Based Negotiation Support: Design, Implementation, and Use," *Decision Support Systems* 25(2), 135–154.

Keeney, R. L. and Raiffa, H. (1976). *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, J. Wiley & Sons, New York.

Koeszegi, S., Srnka, K., and Pesendorfer, E.-M. (2004). "Comparing Web-Based Negotiation Processes: A Combined Qualitative-Quantitative Approach," *Technical Report OP 2004-02*, University of Vienna, Dept. of Business Studies.

Moore, D. A., Kurtzberg, T. R., Thompson, L. L., and Morris, M. W. (1999). "Long and Short Routes to Success in Electronically Mediated Negotiations: Group Affiliations and Good Vibrations," *Organizational Behavior and Human Decision Processes* 77(1), 22–43.

Mosterd, I. and Rutte, C. G. (2000). "Effects of Time Pressure and Accountability to Constituents on Negotiation," *International Journal of Conflict Management* 11(3), 227–247.

Mumpower, J. L. (1991). "The Judgement Policies of Negotiators and the Structure of Negotiation Problems," *Management Science* 37(10), 1304–1324.

Neale, M. A. and Bazerman, M. H. (1985). "The Effects of Framing and Negotiator Overconfidence on Bargaining Behaviors and Outcomes," *Academy of Management Journal* 28(1), 34–49.

Northcraft, G. B., Brodt, S. E., and Neale, M. A. (1995). "Negotiating with Nonlinear Subjective Utilities: Why Some Concessions Are More Equal Than Others," *Organizational Behavior and Human Decision Processes* 63(3), 298–310.

Northcraft, G. B., Preston, J. N., Neale, M. A., Kim, P. H., and Thomas-Hunt, M. C. (1998). "Non-Linear Preference Functions and Negotiated Outcomes," *Organizational Behavior and Human Decision Processes* 73(1), 54–75.

Pennings, J. M. E. and Smidts, A. (2003). "The Shape of Utility Functions and Organizational Behavior," *Management Science* 49(9), 1251–1263.

Perkins, W. C., Hershauer, J. C., Foroughi, A., and Delaney, M. M. (1996). "Can a Negotiation Support System Help a Purchasing Manager," *Journal of Supply Chain Management* 32(2), 37–45.

Rangaswamy, A. and Shell, G. R. (1997). "Using Computers to Realize Joint Gains in Negotiations: Toward an "Electronic Bargaining Table"," *Management Science* 43(8), 1147–1163.

Stuhlmacher, A. F. and Stevenson, M. K. (1997). "Using Policy Modeling to Describe the Negotiation Exchange," *Group Decision and Negotiation* 6, 317–337.

Winterfeldt, D. and Edwards, W. (1986). *Decision Analysis and Behavioral Research*, Cambridge University Press, Cambridge.