

Bidding Support by the Pocket Negotiator Improves Negotiation Outcomes

Aydođan, Reyhan; Jonker, Catholijn M.

DOI

[10.1007/978-981-99-0561-4_4](https://doi.org/10.1007/978-981-99-0561-4_4)

Publication date

2023

Document Version

Final published version

Published in

Recent Advances in Agent-Based Negotiation

Citation (APA)

Aydođan, R., & Jonker, C. M. (2023). Bidding Support by the Pocket Negotiator Improves Negotiation Outcomes. In R. Hadfi, T. Ito, R. Arisaka, & R. Aydođan (Eds.), *Recent Advances in Agent-Based Negotiation: Applications and Competition Challenges* (pp. 52-83). (Studies in Computational Intelligence; Vol. 1092). Springer. https://doi.org/10.1007/978-981-99-0561-4_4

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository


'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Bidding Support by the Pocket Negotiator Improves Negotiation Outcomes

Reyhan Aydoğan^{1,2}  and Catholijn M. Jonker^{1,3} 

¹ Interactive Intelligence, Delft University of Technology, Delft, The Netherlands

C.M.Jonker@tudelft.nl

² Computer Science, Özyeğin University, Istanbul, Turkey

reyhan.aydogan@ozyegin.edu.tr

³ LIACS, Leiden University, Leiden, The Netherlands

Abstract. This paper presents the negotiation support mechanisms provided by the Pocket Negotiator (PN) and an elaborate empirical evaluation of the economic decision support (EDS) mechanisms during the bidding phase of negotiations as provided by the PN. Some of these support mechanisms are offered actively, some passively. With passive support we mean that the user only gets that support by clicking a button, whereas active support is provided without prompting. Our results show, that PN improves negotiation outcomes, counters cognitive depletion, and encourages exploration of potential outcomes. We found that the active mechanisms were used more effectively than the passive ones and, overall, the various mechanisms were not used optimally, which opens up new avenues for research. As expected, the participants with higher negotiation skills outperformed the other groups, but still they benefited from PN support. Our experimental results show that people with enough technical skills and with some basic negotiation knowledge will benefit most from PN support. Our results also show that the cognitive depletion effect is reduced by Pocket Negotiator support. The questionnaire taken after the experiment shows that overall the participants found Pocket Negotiator easy to interact with, that it made them negotiate more quickly and that it improves their outcome. Based on our findings, we recommend to 1) provide active support mechanisms (push) to nudge users to be more effective, and 2) provide support mechanisms that shield the user from mathematical complexities.

Keywords: Negotiation support · Bidding support · Experimental performance evaluation · User experience analysis

1 Introduction

Negotiation is a way to solve conflicts of interest among stakeholders [5, 19, 22, 25]. The negotiation outcome highly depends on the negotiation skills of the

The authors are alphabetically sorted. They put the same effort.

involved parties. Reaching optimal outcomes can be difficult, which explains why negotiation sometimes ends with suboptimal outcomes [23, 25]. Human negotiators can improve their negotiation outcomes by training before or being supported during the negotiation. Artificial Intelligence applications have been developed for both purposes. For example, Conflict Resolution Agent (CRA) is a virtual agent used to let humans train their negotiation skills [10, 11, 20, 21], and another example can be found in [3].

The research into providing computer support for negotiation dates back to the 1960s, see, e.g., [6]. Regarding computer support for people in their negotiations, the most frequently used term is Negotiation Support Systems (NSS). The definition of NSS, following [8], is “software which implements models and procedures, has communication and coordination facilities, and is designed to support two or more parties and/or a third party in their negotiation activities.” Note that the “third” party refers to an independent party, having no stake in the outcome of the negotiation. Gettinger *et al.* [9] suggest differentiating between behavior decision support and economic decision support, which touches more on the human-to-human relationship and interactions, versus the more mathematical analysis of the negotiation. In this work, we entirely focus on economic decision support mechanisms and their effectiveness in the process and outcomes.

In their 2015 paper [7], Foroughi and co-authors report that the decision support component of an NSS enables higher joint outcomes and more balanced contracts, and the communication and coordination facilities positively influence the negotiator attitudes. Various support mechanisms can be integrated into a negotiation system, each contributing in different ways to their user’s negotiations [4, 24, 26].

The Pocket Negotiator (PN) is a negotiation support system [13]) that aims at helping human negotiators improve their negotiation outcomes by guiding the negotiation process, and with a specialization in bidding support. It provides a list of support mechanisms such as analytical support mechanisms (e.g., utility estimation, graphical outcome space capturing the negotiation history, estimated Pareto Optimal Frontier) and strategic advice mechanisms (e.g., bidding advice). In this paper, we empirically investigate the effect of the PN support on negotiation outcomes and negotiation behavior during bidding. In particular, we aim to get more insight into the interaction between the system and human negotiators. For this purpose, we set up a balanced within-group experiment for supported and unsupported negotiations in which we measured which support options were clicked by the participants; we measured the outcome utility of the negotiations. Finally, we conducted a questionnaire on the user experience.

The structure of this paper is as follows. Section 2 briefly introduces the bidding support of PN. In Sect. 3, we also formulate the hypotheses for the experimental design. The experimental setup is described in Sect. 4. The negotiation results are presented in Sect. 5 and discussed in Sect. 6. The paper ends with conclusions and an outline for future research in Sect. 7.

2 Economic Decision Support for Negotiation by PN

The Pocket Negotiator [15] is developed to provide support in all negotiation phases; from domain and profile elicitation, through bidding, to closing. The research in this paper focuses on the effectiveness of the bidding support by PN. Therefore, we only describe the type of support provided during the bidding, which is summarized in Fig. 1.

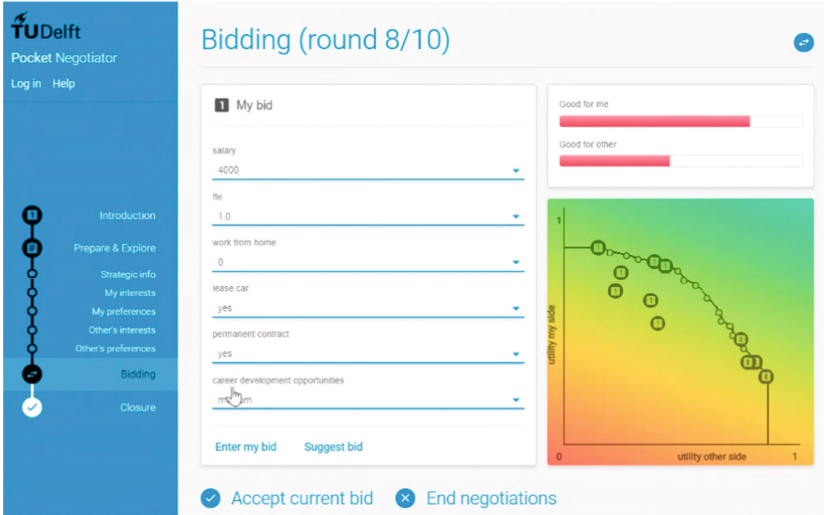


Fig. 1. Pocket negotiator bidding Interface.

In PN, the user can create a bid in three ways: by selecting a value for each issue (middle pane), by clicking on the Estimated Pareto Optimal Frontier (in the graph on the right), or by asking for a suggestion (“suggest bid”) underneath the middle pane. In all cases, the constructed bid’s content is listed in the middle pane. The Estimated Pareto Optimal Frontier (EPOF) is constructed on the basis of the profile elicitation phase covered by another part of PN, see [15]. Whenever the user has created a possible bid, the red bars above the graph on the right side indicate how good that offer is from the perspectives of the user and the opponent. Note that also this is done using the estimated opponent profile.

Furthermore, each offer made by one of the parties is plotted in the graph on the right (i.e., Negotiation History Display). This enables the user to analyse the progress of the negotiation. The user can accept a bid from the opponent or walk away without an agreement by way of the buttons at the bottom of the bidding interface. Either option ends the negotiation and takes the user to the closing phase which is covered by another part of PN not discussed here.

Finally, PN provides *Stopping Advice* to the negotiator. In particular, it advises accepting an offer when its calculations show that the chances of a better offer coming in the future are too low (which is the *Accept Offer* mechanism and suggests to *End the Negotiation without an Agreement* if it feels that the gap between the negotiators is too big, and the opponent appears unwilling to concede enough to get to the Zone of Agreement. The implementations of the *Bidding Advice* and *Stopping Advice* mechanisms depend on the chosen support agent, selected by the user by way of the *Bidding Strategy Selection* mechanism [15]. In the end, it is of course the user who decides whether to accept an offer from the opponent or to walk away without an agreement.

The Pocket Negotiator offers a range of agents that provide bidding suggestions. Well-known bidding strategies from the automated negotiating agents literature have been adapted for this purpose. Behaviour-based strategies, see e.g., [16], adapt their behaviour in response to the bids made by the opponent are used by the Simple Tit for Tat Agent, and Deniz Agent. Variants of concessions are used by e.g., Conceder, Bayesian Agent [12], and Tough Negotiator. Bayesian Agent uses Bayesian models to learn the opponent's preferences. Tough Negotiator only concedes near the deadline and is based on the HardHeaded agent presented in [17]. For our experiments we chose the Deniz agent as it is not easy to characterize by humans and is a negotiator returning Pareto Optimal bids, see [14].

3 Research Hypotheses

Our review of existing systems showed that all systems offer active analytical support for domain elicitation and preference profiling, and then use the information gained by these mechanisms to offer (actively, or passively) some bidding advice, e.g., in critiquing offers, providing the best k offers to make, bidding strategy selection, and on whether or not to accept an offer. Implicit is the choice to only present the best k options from a ranked list of potential offers as is done in FPJ [6] and eAgora [4]. In this manner, this gives an *implicit advice* not to offer lower than these options, or the choice to make only Pareto Optimal offers clickable as is done in PN [15] as this implicitly encourages people to make offers on the EPOF. Note that *explicit* support is visible in the negotiation and relational concepts that the system uses to present information or discuss negotiation aspects with the user.

From a design perspective, it is essential to look for the design deliberations on whether the support mechanisms were integrated into the systems as provide passive or active support. We define a support mechanism to provide *active support* if it pro-actively pushes advice or information to the user in a timely manner. Similarly, we define a support mechanism to provide *passive support* if the support is available upon user request. Overall we see that PN provides some additional forms of implicit support on the efficiency of the bids (EPOF), Graphical Outcome Space, Negotiation History Display), and passive mechanisms to support offer construction (Graphical Offer Selection, and k Best

Offers). We hypothesize that PN support for Bidding Advice would increase the efficiency of negotiation outcomes and that PN is beneficial for the negotiation experience of its users.

- **H1:** Better agreements will be reached by negotiators provided with PN support than when they do not have PN support.

Chen *et al.* [4] deliberately provide active support mechanisms and argue this to be more natural for potential users as it does not require advanced technical or decision analytical skills, suggesting that passive support mechanisms may require some technical and decision analytical skills. Thus, it might be the case that the mechanisms of PN for *Bidding Advice* will be used differently by users with different backgrounds, and that this effect might be more prominent for the passive mechanisms. In general, we would expect that effective usage of the Bidding advice mechanisms of PN requires a threshold level of technical skills as well as some negotiation knowledge and expertise. Our working hypotheses are as follows.

- **H2:** The utility gain for participants provided with Bidding Advice support mechanisms depends on their background.
- **H3:** The usage of passive Bidding Advice mechanisms by the participants is sub-optimal.
- **H4:** Bidding advice mechanisms that implicitly make use of negotiation knowledge nudge the user towards more effective negotiations.

We test these hypotheses only for PN and in our analysis of the data, we consider differences in the background of the participants. To research our hypotheses, we chose the experimental research approach elaborated in the next section.

4 User Experiments

In order to evaluate the bidding support of the Pocket Negotiator, we created a particular version of the system in which all support options are inaccessible for the user. This allows us to set up a balanced experiment in which participants get the opportunity to experience the system both with and without support. This section explains our experimental setup, the changes to PN to carry out the experiment, and the evaluation metrics used.

4.1 Experimental Setup

To properly test the system, we looked for a domain of negotiation that people are familiar with, so that they can easily engage in the negotiation. We hypothesize that negotiation support becomes more important for people with the increase of the complexity of the domain. So, when facing the choice of a negotiation domain, we were looking for a domain with several issues, but that is still relatively simple. The job negotiation domain we settled on satisfies these

criteria. In the job negotiation, we ask people to identify themselves with the role of the applicant that has to negotiate on the following issues: salary (range 2000 till 5000 euro), full-time equivalent (ranging from 0.6 to 1.0), work from home (0, 1 or 2 days), lease car (yes/no), permanent contract (yes/no) and career development opportunities (low, medium, high). We provided the participants with a complete specification of their own preferences, and with what they were told is an **estimate** of the other party's preferences, see Appendix A for the whole story.

In real life, proper preparation for a negotiation entails doing your best to acquire that estimate. The agents underlying the framework have their own approaches to model the opponent's preferences during the negotiation (and more work is currently being done in the research community to learn these preferences also from previous negotiations). To test our hypothesis, we did not want to potentially confuse our results with participants having incorrect or incomplete information on their opponent.

To research our hypotheses, we chose an experimental research approach with three groups with different backgrounds: Computer Science (CS) students, Industrial Engineering (IE) students, and Business Administration (BA) students. CS students have high analytical and technical skills and were given only one lecture on negotiation and ENS systems. IE students are similar to CS students, but we did not provide them with the lecture we gave to the CS students. BA students have less training in technical skills but high negotiation knowledge/skills, as they already had attended several lectures on negotiation, and we gave them the same lecture on negotiation and ENS systems as we gave to the CS students.

The Computer Science students (Group 1) were from Özyeğin University (Turkey), that we motivated to participate by a promise of a bonus point to their overall grade for a course in Collective Decision Making in Multi-Agent Systems. Group 2 consists of Industrial Engineering students from Özyeğin University (Turkey). They were just asked to volunteer, and there was no other connection to the researchers. Group 3 consists of business administration students from Erasmus University (The Netherlands). The Erasmus students were asked to participate in the experiment as a way to get some insight into negotiation tools. There was no relation to their participation and their grade, nor did the researchers have any other connection to the students than just for presenting the negotiation tools and conducting the experiment.

4.2 The Adaptation of PN for the Support and No-Support Conditions

For our experiments, we had to adapt PN to accommodate the No-support condition. In the unsupported version of PN, the users only have the middle section of Fig. 1 available, where they can enter their offers. The red bars, the graph, and the button to ask for a suggestion are not available.

4.3 Evaluation Metrics

The evaluation metric for the performance of the participants was individual utility scored at the end of a negotiation. The utility of outcomes was automatically computed by the Pocket Negotiator system on the basis of the profile information given in Appendix A. Furthermore, we evaluate the effectiveness of the support in terms of the number of bids made on the Pareto Optimal Frontier [23]. Finally, we measure to what extent the participants use the bidding advice mechanisms provided by PN by counting their clicks on the *k Best Offers* ($k = 1$) mechanism indicated by the button “Suggest Bid” and the Graphical Bid Selection mechanism. We run statistical tests to study their usage and effect.

In addition to the objective evaluation metrics, we consider subjective metrics as well. After their negotiation experience, we asked the participants about the usability of the PN, how PN influenced their negotiation process in terms of efficiency of the outcome, speed of the process, whether it distracted them, the overview they had of the process, their satisfaction level at the end of each negotiation, and what they think about their opponent in the negotiation. The evaluation of the subjective evaluation metrics was done using a questionnaire on which we ran an ANOVA analysis; see Appendix B for details.

5 Results

In this section, we present the results of the user experiments, first focusing on negotiation outcomes (Sect. 5.1). In the results, we provide the negotiation performance of the participants per group and the overall performance in each condition: *Support* and *No-Support*. Then, we study the usage of strategic bidding support mechanisms in Pocket Negotiator namely Graphical Offer Selection and Bidding Advice (Sect. 5.2). Finally, the results of the participant questionnaire are presented (Sect. 5.3).

5.1 Negotiation Outcome Results

We first present the outcome analysis to study the effect of bidding support mechanisms in PN per group, and then the overall analysis.

Results of Group 1 - Computer Science Students: The first group consisted of 24 Computer Science students (master and bachelor) at Özyeğin University participated in our experiments. The students received one lecture on the main challenges in negotiation and in building automated negotiating agents before participating in the experiment. Figure 2 shows the utilities gained by the participants against the same opponent at the end of their negotiations in both settings. The participants numbered 1 through 12 started with the *Support* condition, and the rest of the participants started with the *No-Support* condition. It can be seen that most participants received more utility when they negotiated with support (blue lines in the graph) than when they negotiated without

support (orange lines). Three participants ended with the same bid (thus same utility). Two participants (nr 3, and nr 21) did better without support, while the rest did better with support.)




Fig. 2. Utility distribution of Group 1 (CS)

Table 1 shows the average of the utilities received by the users in the *Support* and *No-Support* conditions. According to the Kolmogorov-Smirnov test, the data is not normally distributed; therefore, we applied a non-parametric test, namely Wilcoxon Signed-Rank test. With 0.95 confidence level, the utility received by the participants in the *Support* and *No-Support* versions are statistically significantly different ($z = -3.7191$ and $p = 0.002 < 0.05$), and on average, the users gained higher utilities when supported by PN. That is, the first hypothesis H1 holds for Group 1. In order to see whether the learning effect between sessions plays an important role on the negotiation results, we tested the performance of the participants by grouping them according to their start condition (S-NS denoting *Support* first and then *No-Support*, and NS-S denoting *No-Support* first and then *Support*). The non-parametric statistical test namely Mann-Whitney U test shows no statistically significant difference with a 0.95 confidence level ($p > 0.05$).

Results of Group 2 - Industrial Engineering Students: The second group consisted of 22 Industrial Engineering students at Özyeğin University who did not attend any negotiation lecture prior to the experiment. Figure 3 shows the utilities gained by the participants in both conditions. The participants numbered 1 through 11 started with the PN *Support* condition (S-NS), and the rest of the participants started with the *No-Support* condition (NS-S). Two participants failed to find an agreement in both sessions, and two participants ended with the same offer (thus same utility). Ten participants received higher utility with support, whereas eight participants got higher utility without support.

Table 1. Utility: means and standard deviations for Group 1 (CS). Higher mean values are presented in bold face.

Condition	Order	Mean	Std. deviation	N
No support	NS-S	0.470	0.041	12
Support	NS-S	0.532	0.071	12
No support	S-NS	0.407	0.156	12
Support	S-NS	0.497	0.044	12
No support	All	0.438	0.116	24
Support	All	0.515	0.060	24


**Fig. 3.** Utility distribution of Group 2 (IE)

Table 2 shows the average of the utilities received by the participants in the *Support*- and *No-Support* conditions. Since the data for *No-Support* condition was not normally distributed, we applied a non-parametric test, namely Wilcoxon Signed-Rank test. According to this result, there is no statistical significant difference ($z = -1.0017$ and $p\text{-value} = 0.317 > 0.05$). The first hypothesis does not hold for Group 2. Similarly, there is no significant difference between the NS-S and S-NS conditions.

Table 2. Utility: means and standard deviations for Group 2 (IE)

Condition	Order	Mean	Std. deviation	N
No support	NS-S	0.457	0.120	11
Support	NS-S	0.506	0.145	11
No support	S-NS	0.385	0.206	11
Support	S-NS	0.375	0.204	11
No support	All	0.421	0.169	22
Support	All	0.440	0.185	22

Results of Group 3 - Business Administration Students: This group consisted of 34 Business students at Erasmus University in The Netherlands. These students have a strong background in negotiation, having followed a course of several lectures on negotiation before participating in the experiment. The experiment took place during one of the last lectures of that course. Figure 3 shows the utilities gained by the participants in both conditions. The participants numbered 1 through 19 started with the *Support* condition (S-NS), and the rest with the *No-Support* condition (NS-S). The results show that 16 out of the 34 participants received a higher utility when they negotiated with support from the PN (blue lines) than when they negotiated without support (orange lines), 8 did better without support, and eight performed equally well with and without support.




Fig. 4. Utility distribution of Group 3 (BA)

Table 3 shows the average of the utilities received by the participants in the *Support*- and the *No-Support* conditions. According to the Kolmogorov-Smirnov test, the data is normally distributed; therefore, we applied the paired-samples t-test. With 0.95 confidence level, the utility received by the participants in the *Support* and *No-Support* versions are not statistically significantly different ($t = -1.741$, $p = 0.09 > 0.05$). Hypothesis H1 does not hold for this group.

Overall Results: In this section, we compose the result for the different groups to allow comparisons between the groups, and we merge the groups to analyze the effects of PN support for the merged group. By merging the groups, we obtain a merged group of size 80, with mean utility results as presented in Table 4. For the merged group, the average utility gained in the PN support condition is statistically significantly different than the one in the no-support condition ($p <$

Table 3. Utility: means and standard deviations of Group 3 (BA)

Condition	Order	Mean	Std. deviation	N
No Support	NS-S	0.564	0.075	15
Support	NS-S	0.582	0.086	15
No Support	S-NS	0.528	0.066	19
Support	S-NS	0.552	0.076	19
No Support	All	0.544	0.072	34
Support	All	0.565	0.081	34

0.05) irrespective of their starting condition (i.e., S-NS or NS-S). The same holds when we consider the ordering condition. In particular, for the S-NS condition we obtained $z = -2.0195$ and $p\text{-value} = 0.043$ and for the NS-S condition $z = -3.1159$ and $p\text{-value} = 0.002$. Thus, hypothesis **H1** holds for the merged group.

Table 4. Utility: means and standard deviations for the merged group. The higher mean values are presented in bold face. The difference between the mean utilities in the two conditions is significant according to the Wilcoxon Signed-Rank Test ($z = -3.6011$ and $p = 0.0003 < 0.05$).

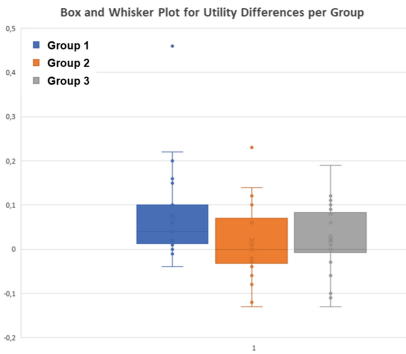
Condition	Order	Mean	Std. deviation	N
No support	NS-S	0.503	0.095	38
Support	NS-S	0.544	0.105	38
No support	S-NS	0.456	0.153	42
Support	S-NS	0.490	0.136	42
No support	All	0.478	0.130	80
Support	All	0.516	0.125	80

Even though we can conclude that for the merged group PN support improves the utility gained by the participants, when considering the data in Figs. 2, 3, and 4, clearly this was not true for all participants. In Group 3, about half of them (16 out of 34) did better with support from PN than without, 10 out of 34 did just as well with as without support, and 8 participants did better without support, see Table 5 for more details.

Table 5. The times that negotiation outcomes improved with support (S)

	S>NS	S<NS	S = NS	No agreement	Totals
Group 1 (CS)	19	2	3	0	24
Group 2 (IE)	10	8	2	2	22
Group 3 (BA)	16	8	10	0	34
Totals	44	18	15	2	80

Finally, we checked the utility difference Δ gained by each participant when they negotiated with and without PN support. We define $\Delta = U_S - U_{NS}$, where U_S and U_{NS} denote the utility obtained in the S and NS condition respectively for the same individual. Figure 5 presents the Box and Whisker plot for this data. There is a statistically significant difference between the utility difference in Group 1 and Group 3, as seen in Table 6. These results support for PN hypothesis **H2**: “The utility gain for participants provided with Bidding Advice support mechanisms depends on their background.”.



Group	Q1 Δ	Q2 Δ	Q3 Δ
Group 1	0.018	0.040	0.100
Group 2	-0.030	0.000	0.060
Group 3	0.000	0.000	0.075

Fig. 5. Utility gain: group quartiles, where $\Delta = U_S - U_{NS}$

Table 6. Utility gain for Bidding Advice support by PN, according to the Mann-Whitney U test. Note that, regarding Group 1 versus Group 2, the one-tailed test did indicate a significant difference ($z = \text{score} = 1.77019$ and $p = .03836 < .05$).

Group A	Mean Δ_A	Group B	Mean Δ_B	z-score	p-value
Group 1 (CS)	0.08	Group 2 (IE)	0.02	-1.786	0.074 > 0.05
Group 1 (CS)	0.08	Group 3 (BA)	0.02	-1.973	0.049 < 0.05
Group 2 (IE)	0.02	Group 3 (BA)	0.02	-0.313	0.754 > 0.05

5.2 Usage of Bidding Support Mechanisms

The Pocket Negotiator was developed with the aim to help humans optimize their negotiation results and at least to avoid leaving money on the table by making Pareto sub-optimal bids. Our experiments show that indeed the number of sub-optimal offers is reduced when people use PN. We measured this by logging how many offers made by the participant were Pareto Optimal. The averages per group and condition are presented in Table 7, along with the results of the Wilcoxon signed-Rank test for dependent variables¹. The difference between the *Support*- and the *No-Support* condition are significant at the level of $p < 0.05$ interval at the group level, even when checking for the starting conditions. The exceptions are Group 2 and Group 3; for the sub-group that started with support (NS-S condition) the difference was not statistically significant.

Table 7. Pareto efficiency: average number of participant offers on the POF per group, per starting-condition, and per condition, and their statistical significance according to the Wilcoxon test.

<i>Wilcoxon Test</i>	S-Pareto	NS-Pareto	z value	p-value	Sig. at p .05
Group 1	7.0	2.3	-3.9199	.00008	✓
Group 2	3.6	1.8	-2.5854	.0096	✓
Group 3	6.4	2.9	-3.7231	.0002	✓
Group 1					
NS-S	6.4	2	-2.6656		✓
S-NS	7.7	2.6	-2.9341	.00338	✓
Group 2					
NS-S	4.6	2.3	-1.6803		×
S-NS	2.6	1.4	-2.1704		✓
Group 3					
NS-S	6.3	2.8	-1.956	.05	×
S-NS	6.6	3	-3.1953	.00138	✓
All Groups					
NS-S	5.84	2.39	-3.7573	.00016	✓
S-NS	5.86	2.45	-4.8467	.00001	✓

The increase in Pareto efficiency cannot be attributed to an ordering effect. Running a non-parametric statistical test namely the Mann-Whitney U test for independent means on the average number of Pareto Optimal offers in the *Support* condition shows that the difference over the (NS-S) versus the (S-NS) condition is not statistically significant at $p < .05$, see Table 8.

¹ Note that some p-values are not specified due to the fact that the test is not able to give the actual p values because of a low number of samples. For those results, the W value is lower than the W -critical value.

Table 8. Comparing pareto efficiency over the (NS-S) and (S-NS) conditions. Presented are the average number of Pareto optimal offers in the merged group. The differences are not statistically significant according to the Mann-Whitney U test.

All groups	NS-S	S-NS	z-score	p-value	Significance at p. 05
S-Pareto	5.84	5.86	-0.10598	.912	×
NS-Pareto	2.39	2.45	1.27176	.204	×

Furthermore, we ran Spearman Rho’s tests on the correlation between the number of Pareto Optimal bids and the number of times that the participant used the Graphical OfferSelection. According to the results ($r_s = 0.612$ and p-value = 0), the correlation with the number of Graphical OfferSelection and the number of Pareto Optimal bids made by the participants is significant. Note that the number of times that participants used the Graphical OfferSelection mechanism is on average *higher* than the number of Pareto Optimal offers made by the participant, see Table 7. This suggests that the participants use the Graphical OfferSelection mechanism and the “Suggest Bid” (k Best Offer) mechanism as a way to explore their options.

Table 9. Marginal means per group and for the merged group. Pareto optimality versus explicit use of clicks to obtain that optimality.

Group	Bids	Pareto optimal	Graphical offer selection	Bid suggestion
Group 1	10	7	8.5	1
Group 2	9.6	3.6	4.9	1.1
Group 3	11.1	6.4	6.7	3.8
All	10.3	5.9	6.7	2.2

In light of hypotheses **H3** and **H4**, we analyzed the difference in how people from the different groups used the passive and implicit Bidding Support mechanisms, see Table 9. For the merged group Fig. 6a and Fig. 6b show how many participants made how much use of the implicit Graphical Offer Selection mechanism by clicking on the EPOF and how often they asked for bidding suggestions (k Best Offers) respectively. These figures show that the clickable mechanisms were not used optimally (some participants never used them at all, and a good portion of the participants only used one of both mechanisms). On the other hand, we also see that some people used these mechanisms more than 10 times, which is more than the number of rounds would account for². The logical follow-up question is, did the people that use these clickable mechanisms get better negotiation results?

² The default setting for the negotiation was 10 rounds.




Fig. 6. The number of participants that clicked/asked for a specific number of times on the specific support mechanism

In other words, can we find out which Bidding Advice mechanism has the most considerable effect on the increase in utility? Table 10 shows the average utility of those participants that made no use of the clickable bidding support mechanisms at all (None), those that only clicked on the Graphical Offer Selection - which corresponds to the offers on the Estimated Pareto Optimal Frontier - (Graphical only), those that only clicked on “Suggest bid” (Suggestions only), those that clicked on both (Both), and the accumulated set of participants that clicked on any of the clickable Bidding Advice mechanisms (Any). The checked for statistically significant differences between the utilities of these sets, see Table 11. Note that the data is normally distributed according to the Kolmogorov-Smirnov test; thus, we applied a t-test for two independent means. From these two tables, one can see that being able to click on the Estimated Pareto Optimal Frontier by way of the Graphical Offer Selection mechanism had the biggest impact on the utility. However, there was no significant difference between that mechanism and the “Suggest bid” (k Best Offers).

Table 10. Statistical information on the average utility and the usage of bidding support mechanisms. “Both” refers to using both the Graphical Offer Selection mechanism and the “Suggest bid” button (k Best Offers mechanism), while “Any” refers to using either of these.

	None	Graphical only	Suggestion only	Both	Any
# of Participants	9	21	18	32	71
Average Utility:	0.42	0.54	0.51	0.53	0.53
Standard Dev.:	0.21	0.07	0.10	0.13	0.11

Table 11. Statistical analysis of the effect of using (combinations of) Bidding Advice mechanisms. Results that are significant at $p < 0.05$ are bold.

Support mechanisms compared	t-value	p-value
None versus anything	-2.62009	.010561
None versus graphical offer selection & Suggest bid	-2.0418	.047976
Graphical offer selection only versus suggest bid only	0.85413	.398528
None versus graphical offer selection only	-2.35534	.025744
None versus suggest bid only	-1.65183	.111076

5.3 Results of the Participant Questionnaire

After completing the negotiations, all participants were asked to fill in an [online questionnaire](#). The questions are renumbered and available in Appendix B, along with the full ANOVA analysis of its results. In this section, we discuss the results of three clusters of questions. The first cluster concerns the impact of the PN bidding support mechanisms on the behavior of the participants. The second cluster focuses on the experience of negotiating with PN support. The third cluster researches the usability of the bidding support of PN. The results of the remaining questions are discussed in [14], which shows that the participants found the Deniz agent to be competitive and that it did not seem human-like to them.

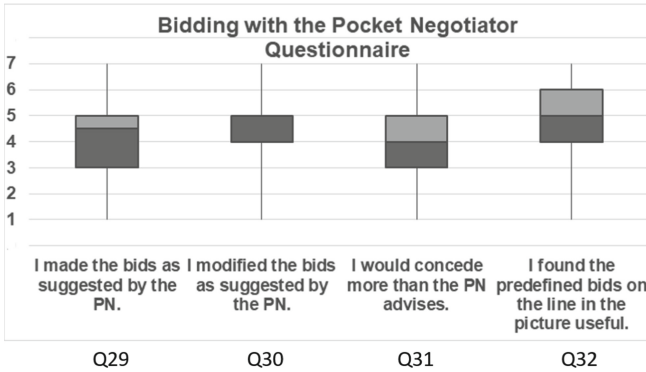


Fig. 7. Participant’s bidding behaviour with the PN where value of 1 indicates total disagreement, value of 4 is neutral, and value of 7 is total agreement.

Figure 7 presents the results of questions Q29, Q30, Q31, and Q32. The graphs show that the participants sometimes made bids as suggested by PN (Q29) and sometimes modified the suggested bids (Q30). So, apparently, some of the participants who asked for a recommendation, felt the need to modify these bids. This provides a possible explanation of the results, namely, that the effectiveness of the bidding advice by PN was reduced due to the participants’ modifications of those suggestions. This hypothesis is supported by the responses

to Q31, namely that the participants indicate that sometimes they concede more than the PN advice. Of course, their thinking about these modifications in itself might already help the participants to get better negotiation results.

The results of Q32 are also of interest to this discussion. Namely, Fig. 7 shows that the participants found the predefined bids of the Graphical Offer Selection mechanism on the EPOF, useful. This result corresponds well with the positive Pearson Correlation between the number of clicks on the EPOF through the Graphical Offer Selection mechanism with the utility that the participants scored in the negotiation. Combining all these results, we hypothesize that the participants felt that the recommendations made by the PN were sometimes too hardheaded (causing them to modify the bids), and that the Graphical Offer Selection mechanism makes it easier for them to make a concession that fits with their own bidding strategy. Note that even if the participants do not use either the Graphical Offer Selection mechanism or the *k* Best Offers mechanism, they still see where their bid is in bidding space by means of the Graphical Outcome Space mechanism, and how good that bid would be for them and for their opponent in the Fig. 1 because of the Critiquing Offers mechanism implemented in the form of the red bars. The gain that PN support provides to these participants is that it helps them to Pareto optimize their offers, and thus protects them from "leaving money on the table".

The responses on question Q14 ("The Pocket Negotiator improves my negotiation outcomes") in Fig. 9 are consistent with the utilities scored by the different groups. The groups return a statistically significant difference in responses. The average for Group 1 is 4.5, for Group 2: 4.9, and for Group 3: 3.8. Group 3 was the group that had taken part in a course of several lectures on negotiation prior to participating in the experiment, and indeed, for them on average, the benefit of the PN, measured as the difference in outcome was less than that for the other groups (NS: 0.544, S: 0.565, see also Table 3).




Fig. 8. Participants' negotiation experience with the PN

We also asked the participants about the quality of their experience with PN, see Fig. 8. We conclude that the participants found that PN made it easier for them to negotiate, and that they gained a better understanding of what is going on in the negotiation. Note that the following questions were phrased negatively. Therefore, average low-scored responses indicate that in fact they found the PN useful:

- Q21: Neutral - average 4.1, median 4
- Q22: Neutral - average 3.9, median 4
- Q23: Positive - average 3.5, median 3

The answers to the positively phrased Q24 “I felt more confident negotiating with the help of the PN” (average 4.3, median 4), shows that PN did not raise their confidence in the negotiation sessions. Overall, the participants were somewhat satisfied with the final result (Q28) (average of 4.6 and a median of 5). Note that this last question might also reflect on their satisfaction with their own results and not necessarily on the support by PN.

The variable *group* had a statistically significant effect on the responses to Q22, which on average had a neutral answer. To understand this, we looked at the average responses per group: Group 1 (average 3.1), Group 2 (average 4.2), Group 3 (average 4.3). We see again that Group 1 appreciated PN better than Group 3. Furthermore, the larger group size of Group 3 influences the median and average of the merged group. This effect also applies to Q23, where the averages per group are: Group 1: 2.7, Group 2: 3.8, Group 3: 3.8.

When asked about the statement “The PN improves my negotiation outcomes” (Q14), the participants’ responses are on average somewhat positive (average 4.3, median 4.9), see Fig. 9, and in this again, the group differences are statistically significant: the averages for these groups are: Group 1: 4.5, Group 2: 4.9, Group 3: 3.8.

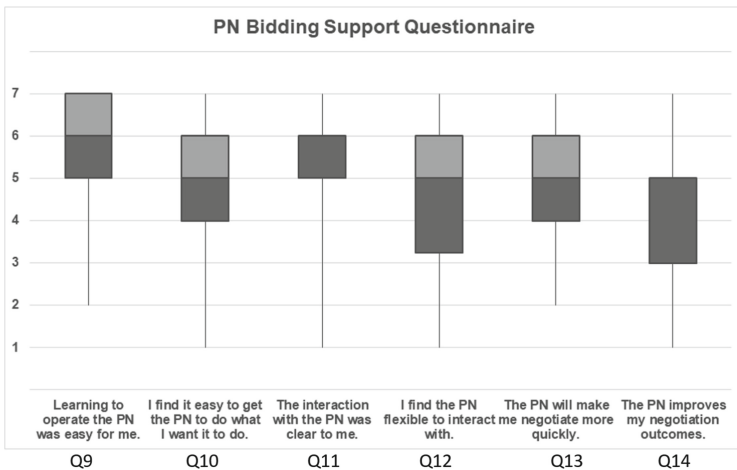


Fig. 9. Participants thoughts about the PN’s bidding support

The responses of the participants on statements regarding usability and effectiveness of PN were rather positive, see Fig. 9. They found it easy to learn to operate PN, and to make it do what they want, they found the interaction clear and flexible, and thought that it makes them negotiate more quickly. The statement about PN improving their negotiation outcomes is also answered somewhat positively (see discussion above), which is in line with the actual outcomes of their negotiations, see Tables 1, and 2, 3. More details can be found in Appendix B.2.



Fig. 10. Participants thoughts about Deniz agent as an opponent

6 Discussion

In this section, we discuss the results reported in the previous section in light of the hypothesis formulated in Sect. 3. First, in Sect. 6.1, we discuss hypotheses **H1** and **H2** for which we graphically summarized the utility data from the three groups. In the subsequent section, we examine the impact of PN on cognitive depletion (Sect. 6.2), the impact of a learning effect (Sect. 6.3), as well as the effect of training (Sect. 6.4). Finally, we discuss the effect of design choices (Sect. 6.5) and the limitations of our research (Sect. 6.6).

6.1 Utility Gain with PN

The results of Table 4 show that hypothesis **H1** is supported with a statistically significant difference, even though for Groups 2 and 3, there is no statistically significant difference: overall better agreements are reached with PN support. The self-reported gain by the participants, see question Q14 (“The Pocket Negotiator improves my negotiation outcomes”) in Fig. 9 are in line with our findings.

We found statistically significant evidence for Hypothesis **H2**: “The utility gain for participants provided with Bidding Advice support mechanisms depends on their background.” With respect to gaining utilities by getting support, we see in Table 6 and Fig. 5 that Group 1 statistically is significantly different from Group 3.




Fig. 11. Means of the utilities grouped by the start condition

Another way to study H2 is to plot the results per negotiation session. Figure 11 summarizes what happens with the marginal means of the different groups in the two conditions over two negotiation sessions, the precise data can be found in Tables 1, and 2, 3. The question is how to interpret this data. We consider two well-known effects in experimental research: cognitive depletion [1], and learning effect [18], and the influence of training and background knowledge in negotiation.

6.2 The Impact of PN on Cognitive Depletion

Cognitive depletion is the notion that performing an exhausting task can deplete a person’s cognitive resources. Negotiation is a complex task as motivated in the introduction, so arguably, after completing the first negotiation session, the participants’ cognitive resources are somewhat depleted so that they would not perform as well in the second negotiation.

If the cognitive depletion effect indeed holds, then what is the effect of supporting the participants in one of their negotiation sessions? If the participants benefit from the ENS, then the support should somewhat counter the cognitive depletion effect. If the participants do not benefit from the support then the cognitive depletion should be the same or worse than in the unsupported sessions. The blue lines, corresponding to the S-NS conditions, go down in Groups 1 and 3, and marginally go up in Group 2. This is in line with the cognitive depletion effect. Figure 11 shows that the red lines (corresponding to the NS-S conditions)

go up in all groups. Therefore, the cognitive depletion effect is reduced by Pocket Negotiator support.

6.3 Learning Effect

The learning effect is the notion that participants learn about the task when encountering the same task multiple times. Since in our experiment the participants are asked to perform exactly the same negotiation twice, arguably, a learning effect would occur, and the participants would perform better in the second negotiation session. As this learning effect is a well-known effect in experimental research, the experimental set-up balances the order of the two conditions.

To what extent can the learning effect also explain the results? Well, it can explain the fact that the red lines go up. However, if that would explain everything, then also the blue lines should go up, and they don't. The learning effect is clearly not enough to counter the cognitive depletion effect by itself, and that getting PN support positively influences the marginal mean of the utility reached by negotiators.

6.4 The Effect of Training

The three groups differ in the background in two ways; the participants study in different studies, and they received different training and education in negotiation. In brief, Group 1 - Computer Science, with one lecture on Negotiation and ENS to prepare them, Group 2 - Industrial Engineering, with no lecture to prepare them, and Group 3 - Business Administration, with a whole course on negotiation. Both groups 1 and 3 did have only one lecture in which they were introduced to PN, which entailed showing them in 10 min how an experienced person would use it, and then having a 15 min training session within a different negotiation domain than the sessions that we base our data on. Group 2 did not get any such introduction to PN, neither on negotiation.

Based on their background, one would assume that in the No-Support conditions, the means of the utilities of the groups would show the ranking: Group 3 > Group 1 > Group 2, which is confirmed by the data, see Table 12, which takes the data from Tables 1, 2, and 3. In fact, regarding the overall means of the two sessions, the utility of Group 2 is statistically significantly lower than that of Groups 1 and 3, and similarly, the utility of Group 1 is statistically significantly lower than that of Group 3, see Table 12. One might conclude that the proficiency of Group 3 stems from their strong negotiation background (yes, education helps!).

What is the effect of PN support on this? We make the following observations: in the first negotiations, Group 1 participants do slightly better with S than without, whereas Group 2 and Group 3 participants do better without support. Definitely, Group 2 struggles with PN in the support mode. One can see an effect here between the two sessions, in that for all groups considering the participants that get support, the means of the utilities in the second sessions is higher than that in the first session (see Fig. 11). In fact, here we see that the difference

Table 12. No-support utility: means for the no-support condition for all groups. Note that there is a statistically significant difference on the overall data between Group 3 and other groups according to Mann-Whitney U test ($p < 0.5$) while there is no significant difference between Group 1 and Group 2.

Group	1 st session	2 nd session	Difference	Overall	Knowledge
Group 1 (CS)	0.470	0.407	-0.063	0.438	1 lecture
Group 2 (IE)	0.457	0.385	-0.072	0.421	no lecture
Group 3 (BA)	0.564	0.528	-0.036	0.544	1 course

between the utilities achieved in those sessions is bigger for Group 2 than for the other groups: Group 2 > Group 1 > Group 3, see Table 13. A potential explanation for this is that the Group 2 participants who did a negotiation without support in the first session are less confused by the support PN has to offer than the other half of Group 1 who had to master both the negotiation concept and the support given by PN in their first session. As the difference in means for Groups 1 and 3 are also positive, to some extent, the same might hold for them. The subgroups of the groups are not big enough to say anything about the statistical significance. So more research is needed to find out to what extent the learning effect of two negotiation sessions impacts the ease with which participants can use PN for the first time.

Table 13. Utility: means in the support condition for all groups. Note that there is a statistically significant difference on the overall data between Group 3 and other groups according to Mann-Whitney U test ($p < 0.5$) while there is no significant difference between Group 1 and Group 2.

Group	1 st session	2 nd session	Difference	Overall	Knowledge
Group 1 (CS)	0.497	0.532	0.035	0.515	1 lecture
Group 2 (IE)	0.375	0.506	0.131	0.440	no lecture
Group 3 (BA)	0.552	0.582	0.030	0.565	1 course

6.5 The Impact of Design Choices

Our literature survey revealed that not only the support mechanisms by themselves but also how they are integrated in the system plays a key role in their effectiveness. Some literature studies showed empirically that the negotiation support tool they used improves the negotiation outcome, see e.g., [2, 7], while others do not, see [24]. In Sect. 3, we formulated two hypotheses on this. **H3:** The usage of passive Bidding Advice mechanisms is sub-optimal. Moreover, **H4:** Bidding Advice Mechanisms that implicitly make use of negotiation knowledge nudge the user towards more effective negotiations.

The results reported in Tables 10 and 11 provide insights regarding the usage of the passive bidding support mechanisms in PN, showing indeed that these mechanisms are not optimally used by the participants, but also that those participants that did use them had higher outcomes. Another argument for the conclusion that the support mechanisms were not optimally used is presented in Appendix C. If the participants would have clicked on the “Suggest bid” (k Best Offers mechanism with $k = 1$) in every round, and would offer that to their opponent in that round, on average they would have increased their outcome. However, in every group there were some students that achieved that utility or even more; some even managed that in the No-support version, see Table 14. None of them followed fully the Deniz-agent’s strategy.

Having this result on the clickable bidding support mechanisms (k Best Offers - “Suggest bid”, and Graphical Offer Selection) still leaves open the question of the impact of the active non-clickable bidding support mechanisms, i.e., the Graphical Outcome Space, Negotiation History Display, and the Critiquing Offers mechanism (the red bars for the user and for the opponent). Unfortunately, with the current version of PN we cannot trace that effect directly; we could try to indirectly measure it by looking at the results of people that use none of the clickable mechanisms and compare that with a group in the No-support conditions. Given the few people (9) who did not use any of clickable mechanisms, we leave that for future work. So, for now we conclude, that the Graphical Offer Selection mechanism had the most impact (and statistically significant positive impact at that) on the negotiation outcomes.

To what extent are the usage patterns of passive Bidding Advice mechanisms influenced by the participants’ main topic of study? The BA students (Group 3) have learned about making trade-offs, about reservation values, and that careful consideration of trade-offs might mean that you can get the opponent to concede more than you initially thought possible. These insights were not present in Group 1. Both groups understand the concept of Pareto Optimality and utility. Based on these considerations, one might argue that the BA students are more carefully considering what trade-offs are made in the different bids, and the CS students might be more easily satisfied by having an outcome on the Pareto Optimal Frontier that seems more or less equal to both negotiators. For the CS students, therefore, walking the Pareto Optimal Frontier using Graphical Bid Selection might be more attractive than for BA students, see Table 7 while BA students may prefer to ask for bidding suggestions more. These observations are in support of hypothesis H2, which says that the utility gain depends on the background of the user.

The results of Table 5, and the participant survey responses, show that some participants followed the bid suggestions while others modified them. We observed that the EPOF accompanied with the Graphical Bid Selection mechanism is useful. It implicitly leads human negotiators to avoid suboptimal offers. Moreover, we hypothesize that the participants felt that the recommendations made by PN were sometimes too hardheaded (causing them to modify the bids), and that the Graphical Bid Selection on the EPOF makes it easier for them to

make a concession that fits with their own bidding strategy. Based on those findings and observations, it is clear that there is room for improvement for which we propose the following two guidelines for the design of ENS systems for human negotiators:

- **Guideline-1:** Provide support mechanisms that actively push support, as those are more effective than passive support mechanisms.
- **Guideline-2:** Provide implicit support mechanisms that shield the user from mathematical complexities.

6.6 Limitations

As Gettinger *et al.* [9] pointed out, focusing attention primarily on the support tools may serve to distance negotiators from each other. With that in mind, one should realize that the experiments done in their work, our work, that of [4] and others that have run experiments with (parts of) ENS support mechanisms all were some sort of lab experiment; the participants were presented with a relatively simple negotiation task of our choosing in a setting that would allow us to research the potential effect of the support mechanism studied. They were not done in real negotiations, nor was the objective to find out in what ways these support mechanisms could be best deployed. This also becomes clear in how our results showed how our participants, that only had half an hour to get to know the bidding support of PN only made limited use of the various support mechanisms it offers to them. Given the insights of [9] one might speculate whether the significant improvement of the negotiation outcomes might be due to the fact that the participants were not allowed to freely interact with their counterpart, but only through the interfaces of the system. Given the complementary nature of the support mechanisms offered by the existing systems, ideally a flexible architecture for ENS would be set up in which these mechanisms can be toggled on and off and serve to provide passive or pro-active support to users in their negotiation roles, and to make this adaptive to the user's expertise and way of working.

7 Conclusions and Future Work

This paper researches the effect on utility and the extent of usage of the set of economic decision support mechanisms that can be found in the Pocket Negotiator [15]. We found evidence that the use of implicit knowledge can be beneficial for the user as it can shield the user from mathematical complexities. Finally, we found indications in the literature that the usage and effectiveness of the ENS system might well depend on the background of the user, in terms of their technical skills and negotiation skills and knowledge. The hypotheses we used in our research reflect these ideas.

To study our hypotheses, we conducted user experiments with the so-far untested bidding support mechanisms of the Pocket Negotiator (PN). The participants negotiated in two versions of PN, one with (S) and one without support

(NS) in a balanced setup over two negotiation sessions. A post-experiment questionnaire provided insight into the user experience of PN. To get insight into the impact of the background of the users, we formed three groups with different backgrounds in terms of technical skills and negotiation knowledge.

As our discussion in Sect. 6 shows, that, in general, the Pocket Negotiator improves negotiation outcomes, and that the way the passive mechanisms are used depends on the negotiation background of the participants. We found that cognitive depletion effects are countered by PN support. The implicit k Best Offers mechanism was used most effectively by the group with the most negotiation skills, but in general, was not used optimally. The implicit Graphical Offer Selection mechanism, which allows users to click on offers on the EPOF, nudges participants to make Pareto Optimal offers, which by itself already reduces the probability of “leaving money on the table”. Both of these mechanisms were used in an exploratory way, and not immediately followed up in making an offer to the opponent. As expected, the participants with higher negotiation skills outperformed the other groups, but still they benefited from the support. Our experimental results show that people with enough technical skills and with some basic negotiation knowledge will benefit more from PN support than others, in that their gain in utility is higher than that of other groups.

The subjective results, based on the questionnaire, show that participants found Pocket Negotiator easy to interact with. They reported that with Pocket Negotiator they could negotiate more quickly and reach better outcomes, which was in line with their negotiation results. Furthermore, they considered the opponent to be a competitive negotiator. Based on our findings and discussion, we conclude that the current form of Pocket Negotiator is an effective tool that increases the performance of the human negotiator, but that it is a tool that users need to familiarize themselves with before they use it in practice as they did not make the most effective use of the bidding support mechanisms provided to them.

Based on our findings, we formulated two guidelines for the design of ENS systems, i.e., to make mechanisms actively push their advice to the user, and to provide implicit mechanisms to shield the user from mathematical complexities.

In terms of objective measures, we found that

- Using the Pocket Negotiator as negotiation support system increases the outcome utility of negotiation in general. This effect is statistically significant according to Wilcoxon Signed-Rank Test ($z = -3.6011$ and $p = 0.0003 < 0.05$) as shown in Table 4;
- Using the Pocket Negotiator statistically significantly increases the average number of participants bids on the EPOF as shown in Tables 7 and 9;
- When compared to the k Best Offers mechanism (“Suggest bid” in PN), the Graphical Bid Selection mechanism had the higher positive impact on the negotiation outcomes: its impact was statistically significantly different as shown in Table 10 and Table 11.

- The Bidding Advice mechanisms were not used at its most effective, as following that advice to the full would significantly improve their outcomes (see Sect. 6.5).

As usual, our research results open up new avenues for further research. For example, in the experiments for this paper, the participants received correct and complete information on the opponent’s preferences, although we did not explicitly tell them so. For future research it would be interesting to study how participants deal with potentially incorrect and incomplete preference profiles and whether or not the ENS system could help them detect this. Furthermore, we realize that also the opponent bidding strategy might have an impact on our experimental results. It would be interesting to experiment with several different opponent bidding strategies, such as hardheaded, Tit-for-Tat, and random.

In general, to boost the research on support mechanisms, we should be able to measure how well people use the different support mechanisms of ENS systems. For this purpose, a method and framework has to be invented in which these mechanisms can be toggled to be included or not, toggled between passive and active, and toggle between implicit and explicit variants of the mechanism. Currently, systems are tested as a whole, which makes it difficult to assess the impact of the individual mechanisms. Such a framework would also make it possible to adapt the ENS to the needs of the user. Furthermore, in particular, bidding advice strategies should be tailored with respect to the negotiation attitude and personality of the users.

Acknowledgments. This research was (partly) funded by the <https://www.hybrid-intelligence-centre.nl/>, Hybrid Intelligence Center, a 10-year programme funded the Dutch Ministry of Education, Culture and Science through the Netherlands Organisation for Scientific Research, grant number 024.004.022 and by EU H2020 ICT48 project “Humane AI Net” under contract # 952026. Moreover, this work has been partially supported by the CHIST-ERA grant CHIST-ERA-19-XAI-005, and by (i) the Swiss National Science Foundation (G.A. 20CH21_195530), (ii) the Italian Ministry for Universities and Research, (iii) the Luxembourg National Research Fund (G.A. INTER/CHIST/19/14589586), (iv) the Scientific and Research Council of Turkey (TÜBİTAK, G.A. 120N680).

A The Pocket Negotiator - Job Negotiation Description

The information on the negotiation scenario given to the participants.

Suppose you have recently conducted a successful job interview, and you are now scheduled for a contract negotiation with your potential boss. In this experiment, you will negotiate twice with your future employer, once with negotiation support, and once without.

The following issues are at stake: salary, fte (full time equivalent), work from home, lease car, permanent contract, and career development opportunities.

There are four things that drive you in general: family life, wealth, status and ambition, and team spirit. You describe yourself as follows:

I feel more comfortable if I have some job security. I have a seven-year old daughter and a new-born son. I don't need to live in wealth, but to meet their expenses, it would be great if I can agree on a high salary. Furthermore, since my partner is working on Mondays and Fridays, I need to take care of our new-born son on those days. I am an ambitious person, and I consider myself a team-player; therefore, I prefer to work full time. I also prefer a job that provides some career development opportunities, such as being able to participate in several personal development workshops, although I could do without. The office is quite some distance from my home, and I would like to make some family trips. Therefore, it would be great if the company could provide me with a lease car.

Therefore, your preferences would be like:

- Permanent contract: Yes \succ No
- Work from home : 2 days \succ 1day \succ None
- Career development Opportunities: High \succ Medium \succ Low
- Lease car: Yes \succ No
- Salary: 4000 \succ 3500 \succ 3000 \succ 2500 \succ 2000
- FTE: 1.0 \succ 0.8 \succ 0.6

where (value 1 \succ value 2) means that you prefer value 1 over value 2. The ordering of the importance of those issues would be: FTE \succ Salary \succ Work from home \succ Permanent Contract \succ Lease car \succ Career Development Opportunities where (issue 1 \succ issue 2) means that issue 1 is more important than issue 2 for you. More details about this preference ordering can be found in Appendix 1 "My preference profile"³.

In your first exploratory meeting with your boss, you already got to know each other a little bit. During this exploration phase, you made the following notes about your boss:

My boss owns a small company and has only a limited budget, so the main issues for him are the salary and the contract duration. I am sure he would prefer to give less salary to me if he can. He probably will not be inclined to lease a car since it would be an extra cost for him. He mentions he had some bad experiences with his former employees that he hired in the past. Although he didn't like their performance, he couldn't fire them because of their permanent contract. Furthermore, he hinted that he likes to work with small, effective teams. I asked him about career development opportunities, but he was rather vague about it.

From this information, you extract that your boss preferences would be like:

- Permanent contract: No \succ Yes
- Work from home : None \succ 1 day \succ 2 days
- Career development Opportunities: Low \succ Medium \succ High
- Lease car: No \succ Yes
- Salary: 2000 \succ 2500 \succ 3000 \succ 3500 \succ 4000
- FTE: 1.0 \succ 0.8 \succ 0.6

³ The appendix mentioned here is left out for reasons of brevity and can be obtained from the authors.

where (value 1 \succ value 2) means that he prefers value 1 over value 2. The ordering of the importance of those issues would be: Work from home \succ Permanent Contract \succ Salary \succ FTE \succ Career Development Opportunities \succ Lease car, where (issue 1 \succ issue 2) means that issue 1 is more important than issue 2 for him. More details about the preference profile of your boss can be found in Appendix 2 “Boss profile”⁴.

B Online Questionnaire

After completing the negotiation, all participants were asked to fill in an [online questionnaire](#). The questions are renumbered and listed in Sect. B.1. We applied ANOVA all-between analysis on the responses we received from the participants, the results are presented in Sect. B.2.

B.1 List of Questions

1. Timestamp
2. What is your gender?
3. What is your age?
4. What is the highest level of education you have completed?
5. Group Type
6. Did you see the demonstration of the Pocket Negotiator before doing this experiment?
7. I am confident about my negotiation skills
8. I consider myself to be a strong negotiator.
9. Learning to operate the Pocket Negotiator was easy for me.
10. I find it easy to get the Pocket Negotiator to do what I want it to do.
11. The interaction with the Pocket Negotiator was clear to me.
12. I find the Pocket Negotiator flexible to interact with.
13. The Pocket Negotiator will make me negotiate more quickly.
14. The Pocket Negotiator improves my negotiation outcomes.
15. I could negotiate better when I had the PN support.
16. It was easier for me to negotiate when I had the PN support.
17. Without the PN support, I can concentrate better.
18. The negotiation outcome was better when I had the PN support.
19. Without the PN support, we found an agreement sooner.
20. I had a better understanding of what is going on during the negotiation when I had the PN support.
21. Using the Pocket Negotiator made it harder to find good agreements.
22. I could gain a better agreement without the PN support.
23. I found the PN distracted me from the negotiation.
24. I felt more confident negotiating with the help of the PN
25. The bids I made using the PN were more self-serving.

⁴ Also this appendix is left out for reasons of brevity and can be obtained from the authors.

26. The PN made it harder for me to pursue my own interests.
27. In the exercises without the PN, I frequently made complete bids.
28. I am satisfied with the end result.
29. I made the bids as suggested by the PN.
30. I modified the bids as suggested by the PN.
31. I would concede more than the PN advises.
32. I found the predefined bids on the line in the picture useful.
33. I would like to negotiate again with this opponent sometime in the future.
34. I took my opponent's preferences into account during the negotiation.
35. I took my own preferences into account during the negotiation.
36. I took my opponent's strategy into account while deciding my next move.
37. I adopted a collaborative negotiation strategy.
38. I was competitive during the negotiation.
39. My opponent was competitive during the negotiation.
40. My opponent used a collaborative strategy.
41. My opponent was a human.

B.2 ANOVA analysis of the Questionnaire Results

- "The Pocket Negotiator improves my negotiation outcomes.": Group has a significant difference (p value = 0.014; Fratio = 4.7 and dF = 2);
- "It was easier for me to negotiate when I had the PN support.": Group has a significant difference (p value = 0.002; Fratio = 6.8 and dF = 2) and age, group and sawdemo (p value = 0.036; Fratio = 4.6 and dF = 1).
- "I could gain a better agreement without the PN support.": Group has a significant difference (p value = 0.012; Fratio = 4.8 and dF = 2).
- "I found the PN distracted me from the negotiation.": Group has a significant difference (p value = 0.032; Fratio = 3.7 and dF = 2).
- "I am satisfied with the end results.": age and education have a significant difference (p value = 0.008; Fratio = 7.5 and dF = 1) and education and group (p value = 0.018; Fratio = 4.5 and dF = 2).
- "I would concede more than the PN advises.": Group has a significant difference (p value = 0.021; Fratio = 4.2 and dF = 2).
- "I found the predefined bids on the line in the picture useful.": Group has a significant difference (p value = 0.007; Fratio = 5.4 and dF = 2).
- "I would like to negotiate again with this opponent sometime in the future.": Age has a significant difference (p value = 0.021; Fratio = 5.6 and dF = 1) and group (p value = 0.017; Fratio = 4.4 and dF = 2).
- "I took my opponent's preferences into account during the negotiation.": Education has a significant difference (p value = 0.029; Fratio = 3.8 and dF = 2).
- "I took my own preferences into account during the negotiation.": Group has a significant difference (p value = 0.042; Fratio = 3.4 and dF = 2) and age and sawdemo (p value = 0.038; Fratio = 4.5 and dF = 1).
- "I took my opponent's strategy into account while deciding my next move). Group and sawdemo have a significant difference (p value=0.014; Fratio=4.7 and dF=2).

- “I adopted a collaborative negotiation strategy.”: Age has a significant difference (p value = 0.001; F ratio =11.4 and $dF=1$) and education and group (p value=0.018; f Ratio = 4.3 $dF = 2$);
- “My opponent was using a collaborative strategy”: Group has a significant difference (p value = 0.006; F ratio = 5.6 and $dF = 2$).

C Self-Play in PN and Individual Outcomes

The subjects in our experiments were negotiating against a software agent, called Deniz [14], and that a copy of Deniz supported the participants. This means that if the subjects would follow the recommendation exactly, the Deniz agent would negotiate with itself. However, none of our human participants that did that.

Table 14 shows the number of participants who achieved an outcome with a utility that was at least as high as Deniz agent would have achieved when playing against itself (Utility = 0.58).

Table 14. The number of participants who received at least as high utility as Deniz agent received when it plays against itself

Group	S in session 1	NS in session 1	S in session 2	NS in session 2
Group 1	–	–	2 (0.61;0.69)	–
Group 2	–	1 (0.70)	1 (0.76)	1 (0.59)
Group 3	4 (0.59–0.72)	7 (0.58–0.72)	9 (0.58–0.72)	3 (0.61–0.67)

In our experiments, the Bidding Advice mechanism is actually hardly asked for, and the advice is not always follow-up by the participants, as we reported as one of our findings. So there is no self-play in the experiment. Furthermore, the Deniz agent does not know what strategy the other is playing (neither as supporting agent, nor as opponent agent) and has no mechanisms for manipulating such foreknowledge, see [14], where the results of Deniz’s self-play are reported.

We could have let the participants play against multiple opponents however, we think one should first make sure that participants use support mechanisms of the agent more often and more effectively. After that, of course, more elaborate experiments are in order, with more agents, and many more negotiation scenarios (from 1 to many negotiation issues, with issue inter-dependencies and without).

References

1. Baumeister, R.F., Vohs, K.D., Tice, D.M.: The strength model of self-control. *Curr. Dir. Psychol. Sci.* **16**(6), 351–355 (2007)
2. Bosse, T., Jonker, C.M.: Human vs. computer behavior in multi-issue negotiation. In: *Rational, Robust, and Secure Negotiation Mechanisms in Multi-Agent Systems (RRS’05)*, pp. 11–24. IEEE (2005)

3. Broekens, J., Harbers, M., Brinkman, W.-P., Jonker, C.M., Van den Bosch, K., Meyer, J.-J.: Virtual reality negotiation training increases negotiation knowledge and skill. In: Nakano, Y., Neff, M., Paiva, A., Walker, M. (eds.) IVA 2012. LNCS (LNAI), vol. 7502, pp. 218–230. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33197-8_23
4. Chen, E., Vahidov, R., Kersten, G.E.: Agent-supported negotiations in the e-marketplace. *Int. J. Electron. Bus.* **3**(1), 28–49 (2005)
5. Fisher, R., William, L.: *Getting to Yes*. Penguin Group, New York (1981)
6. Foroughi, A., Perkins, W., Jelassi, T.: An empirical study of an interactive, session-oriented computerized negotiation support system (NSS). *Group Decis. Negotiat.* **4**, 485–512 (1995). [10.1007/BF01409712](https://doi.org/10.1007/BF01409712)
7. Foroughi, A., Perkins, W.C., Hershauer, J.C.: A study of asymmetrical decision support in computerized negotiation support systems (NSS). *Glob. Bus. Fin. Rev. (GBFR)* **20**(1), 25–42 (2015)
8. G., K., Lai, H.: Negotiation support and e-negotiation systems: an overview. *Group Decis. Negot.* **16**, 553–586 (2007)
9. Gettinger, J., et al.: Impact of and interaction between behavioral and economic decision support in electronic negotiations. In: Hernández, J.E., Zarate, P., Dargam, F., Delibašić, B., Liu, S., Ribeiro, R. (eds.) EWG-DSS 2011. LNBIP, vol. 121, pp. 151–165. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32191-7_11
10. Gratch, J., DeVault, D., Lucas, G.: The benefits of virtual humans for teaching negotiation. In: Traum, D., Swartout, W., Khooshabeh, P., Kopp, S., Scherer, S., Leuski, A. (eds.) IVA 2016. LNCS (LNAI), vol. 10011, pp. 283–294. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-47665-0_25
11. Gratch, J., Nazari, Z., Johnson, E.: The misrepresentation game: how to win at negotiation while seeming like a nice guy. In: Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, pp. 728–737. IFAAMAS (2016)
12. Hindriks, K., Tykhonov, D.: Opponent modelling in automated multi-issue negotiation using Bayesian learning. In: Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems-vol. 1. pp. 331–338. International Foundation for Autonomous Agents and Multiagent Systems (2008)
13. Jonker, C.M., et al.: An introduction to the pocket negotiator: a general purpose negotiation support system. In: Criado Pacheco, N., Carrascosa, C., Osman, N., Julián Inglada, V. (eds.) EUMAS/AT -2016. LNCS (LNAI), vol. 10207, pp. 13–27. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59294-7_2
14. Jonker, C.M., Aydoğan, R.: Deniz: a robust bidding strategy for negotiation support systems. In: Ito, T., Zhang, M., Aydoğan, R. (eds.) ACAN 2018. SCI, vol. 905, pp. 29–44. Springer, Singapore (2021). https://doi.org/10.1007/978-981-15-5869-6_3
15. Jonker, C.M., et al.: An introduction to the pocket negotiator: a general purpose negotiation support system. In: Criado Pacheco, N., Carrascosa, C., Osman, N., Julián Inglada, V. (eds.) EUMAS/AT -2016. LNCS (LNAI), vol. 10207, pp. 13–27. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59294-7_2
16. Jonker, C.M., Aydoğan, R., Baarslag, T., Fujita, K., Ito, T., Hindriks, K.V.: Automated negotiating agents competition (ANAC). In: AAAI, pp. 5070–5072 (2017)
17. van Krimpen, T., Looije, D., Hajizadeh, S.: Hardheaded. In: Ito, T., Zhang, M., Robu, V., Matsuo, T. (eds.) *Complex Automated Negotiations: Theories, Models, and Software Competitions*, vol. 435, pp. 223–227. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-30737-9_17

18. Lazar, J., Feng, L.H., Hochheiser, H.: *Research Methods in Human-Computer Interaction*. Willey, Chichester (2010)
19. Lewicki, R.J., Saunders, D.M., Barry, B., Minton, J.W.: *Essentials of Negotiation*. McGraw-Hill, Boston (2003)
20. Mell, J., Gratch, J., Baarslag, T., Aydođan, R., Jonker, C.M.: Results of the first annual human-agent league of the automated negotiating agents competition. In: *Proceedings of the 18th International Conference on Intelligent Virtual Agents*, pp. 23–28. IVA '18, ACM, New York, NY, USA (2018). 10.1145/3267851.3267907, <http://doi.acm.org/10.1145/3267851.3267907>
21. de Melo, C.M., Marsella, S., Gratch, J.: Do as I say, not as I do: challenges in delegating decisions to automated agents. In: *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pp. 949–956. IFAAMAS (2016)
22. Raiffa, H.: *The Art and Science of Negotiation*. Belknap Press, reprint edn. (2005), first version from 1982
23. Raiffa, H., Richardson, J., Metcalfe, D.: *Negotiation Analysis: The Science and Art of Collaborative Decision Making* (2002)
24. Schoop, M., van Amelsvoort, M., Gettinger, J., Koerner, M., Koeszegi, S.T., van der Wijst, P.: The interplay of communication and decisions in electronic negotiations: communicative decisions or decisive communication? *Group Decis. Negot.* **23**(2), 167–192 (2014)
25. Thompson, L.: *The Mind and Heart of the Negotiator*, 3rd edn. Prentice Hall Press, Upper Saddle River (2000)
26. Yuasa, M., Yasumura, Y., Nitta, K.: A negotiation support tool using emotional factors. In: *Proceedings Joint 9th IFSA World Congress and 20th NAFIPS International Conference* (Cat. No. 01TH8569), pp. 2906–2911. IEEE (2001)