Decision Biases in Preference Acquisition

Alexander Felfernig[1] and **Martin Stettinger**[1] and **Ralph Samer**[1]

Abstract. Decision support systems in many cases are based on user interfaces used to collect preferences and requirements of users. For example, configurators in the automotive domain ask users to provide preference information regarding the car color and car engine. Stakeholders in release planning scenarios provide feedback on software requirements in terms of importance evaluations of different interest dimensions. In such scenarios, decision biases can trigger situations where users take suboptimal decisions. In this paper, we provide a short overview of example decision biases and report the results of an empirical study that show the existence of such biases in the context of release planning (configuration) decision making.

1 Introduction

When interacting with decision support systems such as recommender systems [3] or configuration systems [5], users do not know their preferences beforehand but construct and frequently adapt these within the scope of a decision process [2]. In most of the cases, users also do not try to optimize decisions but apply heuristics to take a decision. For example, "elimination by aspects" (EBA) is based on the simple idea of an attribute-wise comparison of different decision alternatives where only those alternatives remain in a consideration set which satisfy a pre-defined set of preferences. This strategy can lead to suboptimal outcomes since alternatives that could become more preferable in the future have already been eliminated in the past.

Recently, decision support for groups became increasingly popular due to the fact that in many scenarios user groups are engaged in decision processes [3]. Examples thereof are (1) release planning where a group of stakeholders is in charge of prioritizing a given set of requirements and (2) open innovation scenarios where customer groups are contributing when deciding about the features of a new product. In many cases, the underlying decision scenario can be regarded as a group decision problem [3]. In this paper, we provide an overview of example decision biases that can occur in preference acquisition scenarios for individual users as well as groups. In this context, we discuss the results of an empirical study conducted in the context of preference acquisition for release planning.

The remainder of this paper is organized as follows. In Section 2, we exemplify decision biases on the basis of examples from the domain of software requirements engineering. In this context, we discuss the results of a related empirical study. In Section 3, we discuss issues for future work. We conclude this paper with Section 4.

2 Decision Biases in Preference Acquisition

The major goal of our study was to analyze biases in decision scenarios. The study focused on an analysis of the decision behavior of computer science students (N=222) working in groups of 6-8 persons in a software project. A structured questionnaire with A/B testing was used to analyze the decision behavior of students. The average time needed to complete the questionnaire was 4 minutes. In order to simulate decision scenarios, scenario descriptions where integrated in questions where needed. In the following, we provide an overview of the study results.

Framing. The way a decision alternative is presented can influence a user's decision behavior. One example of framing is that users prefer meat that is characterized with "80 percent lean" over meat that is "20 percent fat". Another example is the framing of prices: when comparing the offers of company x and y, the offer "pellets fore 24.50 per 100kg with a discount of 2.50 if the customer pays with cash" from x appears to be the more attractive one compared to the offer "pellets for 22.00 per 100kg with a 2.50 surcharge if the customer pays with credit card" from company y. The increased attractiveness of x's offer can be explained by prospect theory [6] which points out that alternatives are evaluated with regard to both, gains and losses, and losses (in our example, *fat meat* and *surcharge*) have a higher negative impact on a decision compared to equal gains.

Framing: Study Results. In our study, we described a scenario where stakeholders had to estimate the acceptability of a given probability of successful project completion. In one setting, the probability was specified as "probability of success", in the other setting, the probability was expressed as "failure probability". In the first setting, study participants evaluated the acceptability on an average with 86 out of 100 points (1: not acceptable, 100: definitely acceptable). In the second setting, study participants evaluated the acceptability on an average with 77 points (out of 100).

Anchoring. It is known that preference visibility has various negative impacts on the quality of a group decision. An example thereof is *anchoring* where indicated reference values (e.g., item evaluations of a group member) can have an influence on the evaluation behavior of other group members. An example thereof is the visualization of the average rating given to an item by a community: increasing the shown average item rating results in a increased rating by individual community members [1].

Anchoring: Study Results. In our study, the participants were asked whether it is important for them to have as soon as possible knowledge about the preferences (which requirements should be implemented when?) of other stakeholders. Nearly 70 percent of the study participants agreed on the fact that the mentioned preference visibility is important. These persons are vulnerable to limited information exchange which can result in suboptimal decisions [3].

Decoy Effects. Decisions are taken depending on the context in which the alternatives are presented. It can be the case that completely inferior decision alternatives added to a set of alternatives can change the selection behavior of users. Such alternatives are denoted as *decoy items* since they manage to draw the attention of users towards specific alternatives. An example of a decoy effect is *asymmet-*

[1] Graz University of Technology, Austria, email: {alexander.felfernig, martin.stettinger, ralph.samer}@ist.tugraz.at

ric dominance which denotes a situation where a decoy alternative is dominated by an item T in all dimensions. Dominance is evaluated in terms of a pairwise comparison of attribute values characterizing the alternatives. An example of asymmetric dominance is shown in Table 1. Alternative c can be regarded as a decoy item since it is outperformed by alternative a in both dimensions (higher project returns and lower project efforts) and thus makes alternative a even more attractive compared to alternative b.

release	project returns	project efforts
a	30.000	15.000
b	50.000	35.000
c	28.000	16.000

Table 1. An example of an asymmetric dominance effect.

Decoy Effects: Study Results. The study participants were asked to select one out of two alternative software release plans (characterized by the corresponding estimated returns and efforts). Release alternative c is completely dominated by release alternative a which has been selected in 86 percent of the cases (only 9 percent of the participants selected alternative b).

Table 2 includes a variant of the previous setting where alternative c is arranged near to alternative b. Compared to the setting shown in Table 1, the share of participants selected this alternative was only 77 percent wheres 22 percent of the participants selected alternative b. Consequently, the inclusion of inferior alternatives can trigger a shift in the selection behavior of stakeholders. One way to counteract such situations is to point out inferior alternatives or to simply delete these from the set of available options.

release	project returns	project efforts
a	30.000	15.000
b	50.000	35.000
c	52.000	40.000

Table 2. Another example of an asymmetric dominance effect.

Decision Strategies of Study Participants. In addition to the above mentioned biases, the study participants were asked a couple of questions regarding their practices in group decision making. *First*, early knowledge about the preferences of other stakeholders was considered as a positive element that helps to improve the quality of requirements prioritization (84% of the study participants supported this statement). However, as indicated in the literature, early knowledge about the preferences of other stakeholders can have a negative impact on decision quality since focusing on preferences triggers less efforts related to the exchange of decision-relevant information [7]. *Second*, participants regarded consensus as a positive aspect at the beginning of a decision process (80% support for this statement), i.e., consensus at the beginning is regarded as a precondition for high-quality prioritization. However, the contrary is the case: consensus at the very beginning contributes to the avoidance of knowledge interchange between stakeholders [8]. *Third*, study participants were asked regarding their opinion on the impact of preference visibility on the probability of decision manipulation. In this context, the majority of study participants (64% support) agreed that preference visibility increases the probability of manipulation. However, 36% still think that this is not the case.

Summarizing, biases in preference acquisition exist and can have a negative impact on the outcome of the decision process. As a result of our user study, it could be observed that study participants (in our case Computer Science students) were often not aware of this and thus vulnerable to such biases.

3 Future Work

There are a couple of issues that are within the scope of our future research. *First*, the majority of researchers still focuses on the identification of new biases and the analysis of biases in specific decision scenarios. A major goal of our ongoing and future work is to focus on approaches to automatically identify potential sources of suboptimal decisions and to adapt the underlying decision support. For example, decoy effects can be predicted on the basis of a formal model [4] - our focus for future work in this context is to figure out interactions between different decoy effects and to find ways to counteract such biases. *Second*, we will investigate how explanations can help to counteract biases and what kind of explanations are useful in which context. For example, in release planning, stakeholders could be informed about the fact that some of the candidate requirements should be analyzed in more detail. *Third*, we will extent the scope of our user studies to industrial settings.

4 Conclusions

In this paper, we discussed the results of a user study related to the existence of decision biases in preference acquisition. The results were discussed on the basis of an empirical study that was conducted with computer science students within the scope of a software engineering course. The outcomes of this study clearly indicate the existence of decision biases and suboptimal decision practices that can lead to suboptimal outcomes in group decisions. Our future work will include a.o. an analysis to which extent explanations can help to counteract decision biases. Furthermore, we will extend the scope of our user studies to industrial scenarios.

ACKNOWLEDGEMENTS

The work presented in this paper has been conducted within the scope of the Horizon 2020 project OpenReq (openreq.eu).

REFERENCES

[1] G. Adomavicius, J. Bockstedt, S. Curley, and J. Zhang, 'Recommender systems,consumer preferences, and anchoring effects', in *Decisions@RecSys11*, pp. 35–42, Chicago, IL, USA, (2011).

[2] A. Felfernig, 'Biases in decision making', in *Proceedings of the International Workshop on Decision Making and Recommender Systems 2014*, number 1278 in CEUR Proceedings, pp. 32–37, Bolzano, Italy, (2014).

[3] A. Felfernig, L. Boratto, M. Stettinger, and M. Tkalcic, *Group Recommender Systems – An Introduction*, Springer, 2018.

[4] A. Felfernig, B. Gula, G. Leitner, M. Maier, R. Melcher, S. Schippel, and E. Teppan, 'A dominance model for the calculation of decoy products in recommendation environments', in *AISB Symposium on Persuasive Technologies*, pp. 43–50, Aberdeen, Scotland, (2008).

[5] A. Felfernig, L. Hotz, C. Bagley, and J. Tiihonen, *Knowledge-based Configuration: From Research to Business Cases*, Morgan Kaufmann Publishers, 1st edn., 2014.

[6] D. Kahneman and A. Tversky, 'Prospect theory: An analysis of decision under risk', *Econometrica*, **47**(2), 263–291, (1979).

[7] A. Mojzisch and S. Schulz-Hardt, 'Knowing others' preferences degrades the quality of group decisions', *Journal of Personality and Social Psychology*, **98**(5), 794–808, (2010).

[8] G. Ninaus, A. Felfernig, and F. Reinfrank, 'Anonymous preference elicitation for requirements prioritization', in *ISMIS'12*, volume 7661 of *LNCS*, pp. 349–356, Macau, China, (2012).