The Impact of Simplicity on Financial Decision-Making

Marco Monti (monti@mpib-berlin.mpg.de)
Max Planck Institute for Human Development, Department for Adaptive Behavior and Cognition,
Königin-Luise-Straße 5
14195 Berlin Germany

Laura Martignon(martignon @ph-ludwigsburg.de)
University of Education, Ludwigsburg,
Reutealee 46
71634 Ludwigsburg Germany

Gerd Gigerenzer(sekgigerenzer@mpib-berlin.mpg.de)
Max Planck Institute for Human Development, Department for Adaptive Behavior and Cognition,
Königin-Luise-Straße 5
14195 Berlin Germany

Nathan Berg, (nberg@utdallas.edu)
School of Economic, Political, and Policy Sciences (EPPS) - University of Texas-Dallas
Mail Station GR31 - 800 West Campbell Road
Richardson, TX 75080-3021

Abstract

This paper reports new experimental and survey data collected from bank customers at several Italian banks. These data aim to uncover the decision processes used by investors, including their investment goals, the information sets they consider, and the number of factors that actually influence high-stakes financial decisions. Most subjects use a strict subset of the information available to them, ignoring variables that standard economic models typically assume drive investors' behavior. Rather than random trembling which would predict that omitted variables are dropped at random, fast and information-frugal heuristics appear to explain the information search and decision behavior of many subjects observed in this study, reflecting a lexicographic hierarchy of risk, time horizon and cost, in that order. A simple combination of a fast and frugal tree and a tallying rule predicts about 80% of investors' decisions.

Keywords: behavioral finance, decision making, heuristics.

1. Introduction

Bank customers are not experts, and yet they make high-stakes decisions that can change their welfare for better or worse. This raises the question of how non-experts actually go about making financial decisions and the processes that provide good empirical descriptions of their purposeful investment behavior. We report new evidence from customers at Italian mutual banks about the simplicity of non-experts' judgments, including how many pieces of information they typically consider, and the heuristic rules that map information in their consideration sets into actual decisions.

A growing literature in economics and psychology documents that decision makers typically do not incorporate all available information into their decisions, even when that information is statistically valid, non-redundant (i.e., non-collinear with other predictors), and costless to acquire (Chewning & Harrell, 1990; Lee & Lee, 2004; Berg & Hoffrage, 2008). If search is limited rather than exhaustive, how are the pieces of information considered mapped into actual behavior? Gigerenzer, Todd and the ABC Group (1999) put forward a positive theory regarding simple and information-frugal decision rules that have surprisingly attractive theoretical properties (e.g., accuracy in prediction tasks, as in Gigerenzer & Brighton, 2009; Martignon & Laskey, 1999, and Gigerenzer & Goldstein, 1996) and solid empirical support based on lab experiments.

Fast and frugal heuristics have been shown to perform well at prediction in a variety of domains (Gigerenzer, Todd and the ABC Group, 1999); they make inferences with very little knowledge and computational effort, largely by using only a small subset of the available information. They make no trade-offs and proceed lexicographically through each factor associated with the pair of objects being compared. Under certain circumstances, such heuristics can be as accurate as weighted linear models, falling only slightly behind the Bayesian approach (Martignon & Laskey, 1999 in Gigerenzer, Todd and the ABC Group, 1999) that takes all relevant factor correlations or conditional dependencies into consideration.

One simple decision heuristic that can be applied to binary decisions, such as whether to invest in stocks or bonds, is the take-the-best (TTB) rule, which ignores all correlations among predictors (i.e., the features of different investments being considered by an investor) and uses them, one by one, in a lexicographic decision tree that requires no weighting or averaging. TTB is a rule that maps two lists of features (or "cues" in the jargon of psychology, or "predictors" in the jargon of economics), one list for each of the two alternatives being considered, into a binary inference or decision. A key part of TTB's success in out-of-sample prediction is its noncompensatory, or lexicographic, structure according to which each feature is considered one at a time following a fixed ordering. If the first feature points in the same direction for both of the alternatives, then TTB makes no decision on the basis of that feature and the next feature is considered. As soon as one characteristic or feature points clearly in favor of one of the two alternatives, TTB makes the decision on the basis of that feature, ignoring all other features that could have been compared. Thus, TTB is fast and frugal, in the sense of depending only on a strict subset of the available information regarding the features of investments.

In contrast to TTB, the neoclassical model used in most analyses of financial decision-making assumes that decision makers exhaustively search the elements in their feasible sets, weigh costs and benefits of all features associated with each element of this set and, after weighting all relevant information, select the investment that is the global maximizer. TTB and the neoclassical model make predictions about the process that investors use when making a high-stakes decision. It implies that all information receives some weight (aside from the trivial cases of perfectly correlated pieces of information, or those that do not correlate at all with payoffs). Therefore, process tracing of bank consumers' decisions should, according to the neoclassical model, reveal that all investment features are looked up, and that all of them are integrated systematically into observed choice behavior. The neoclassical model implies that no relevant information is discarded or ignored. In terms of investors' goals, the neoclassical model implies that investor's behavior proceeds as if it is maximizing something, so that (at least for an interior optimum) goals should be described in terms of solving something akin to first-order equations, where marginal benefits just offset marginal cost.

The TTB model makes different predictions than the neoclassical model. Using experimental and survey data collected from Italian bank customers, this paper contributes to reveal that there is evidence that clearly rejects the neoclassical model and fits a lexicographic fast and frugal decision tree akin to TTB. Furthermore, normative assessment of the performance of real bank customers' decision processes (relative to the neoclassical benchmark) indicates that heuristic strategies appear to serve investors reasonably well. Whereas the biases and heuristics literature frequently assigns an automatic negative normative value to any decision procedure that deviates from the neoclassical ideal, we identify attractive normative properties of the fast and frugal heuristic approach to investment decision-making. TTB and similar lexicographic decision-tree heuristics consider the features of investments sequentially in a ranking determined by some measure of the goodness of each feature rather than considering all their inter-correlatedness. This helps reduce the cognitive processing required to execute the strategy and can improve robustness and accuracy of predictions (Gigerenzer & Brighton, 2009).

The experimental tasks that our subjects faced required them to search freely for information, unlike most experimental economics treatments of financial decisions, in which subjects are provided with a complete set of summary statistics such as expected values, variances and covariances, required by standard models such as Capital Asset Pricing Model.

In our case, participants could decide how much information they wanted to look at. Most chose not to look up all the information that was freely available to them and a significant proportion explored only a small set and overlapping of investment features (repeatedly looking up the same piece of information). Subjects exhibited remarkably similar information search behavior across trials, reporting that they spent on average only a short time on all portfolio decisions. Moreover, the bank customers indicated that, although they were handling meaningful amounts of money, they made investment decisions with relatively little cognitive effort. By far, the most important features of investments in the eyes of our subjects were risk, time horizon, and costs (brokerage fees), in that order. Information exploration was characterized by frugality and simplicity. Finally, investors' decisions can be described by a fast and frugal heuristic that has a very simple representation.

2. Methodology

The research project as a whole was developed in two steps. We first interviewed 20 professional financial advisors and 80 bank customers of the Italian mutual bank[1]. These interviews provided information on both the perspective of the advisors on the customers and that of the customers on the advisors concerning investment strategies. Data collected from these interviews were used to design the test taken by 15 customers (from the sample of 80 interviews). In this paper we focus on the results of the test and their analysis. Data collected from the test were analyzed both, at the aggregate level, and, at the individual level (within subject approach), by considering each single subject (across the 15 customers that took the test).

2.1 The instruments

The data treated here[2] consist of test results which track information look-ups and decisions in a hypothetical investment task. The hypothetical investment task was not incentivized by tying subjects' payments to the task. We discuss the rationale for this approach below and why we are confident that this design produces useful insights about inter-subject variability in terms of information usage in investment decision-making.

Fifteen customers were tested with a sequence of four experimental tests, based on Java Language. Our intention

[1] Cassa Rurale Giudicarie Valsabbia Paganella

[2] Data collected trough the questionnaires will be presented in another paper dealing with the customer and advisor relationship and advice taking strategy.

was to analyze subjects taking decisions in a naturalistic environment with maximally realistic investment tasks. Therefore, the information setup in portions of the investment task we asked our bank customer subjects to perform, relied in part on information about the investment features mentioned by them and advisors in the questionnaires and calibrated with realistic and typical values of expected return, risk, etc.

2.2 Experimental Design

We investigated decision strategies by considering three factors: the overall amount and the type of information that subjects need for taking their decisions, and the approach they follow in the information-search process.

Participants in our research were customers of an Italian mutual bank. A mutual bank is a nonprofit institution whose aim is to support the economic well-being of people living in a specific region. We selected this type of bank because its financial advisors are neither under the pressure of budget goals nor conditioned by other economic incentives that may distort their presentation of financial products.

Participants were randomly extracted from the bank database, which contains data on all active customers. The only requirement participants had to fulfill was to have deposits of at least 40,000 euros.

The computer-administered investment tasks were performed at different branch locations of the bank in Trento, Italy. The interviewer read the instructions to each participant and also explained the aim of the test. Each experimental session lasted approximately 75 min. (60 min. for the questionnaire interview and 15 min. for the investment tasks). Investors were not remunerated. They voluntarily participated in the tests and showed great enthusiasm, viewing their participation as a contribution to the quality of their mutual bank.

Tests were conducted on a touch-screen-based interface programmed in Java language. We chose touch-screen technology to facilitate the interaction of elderly investors, with dynamic information provided by the computer.

Each subject was placed in front of the touch-screen and trained on how to manage each single task. A personal computer ran a Java Virtual Machine, which recorded all the experimental data and, thus, all investors' decisions.

Each test was composed of four different phases that gave subjects a chance to implement a neoclassical strategy of exhaustive search of investment features while measuring the subjects' actual usage of information about features.

The task began by asking customers to choose between two investments, later extending the number of possible choices to six. Subsequently, customers were asked to design their preferred investment portfolio, basing their choices on investment labels and features. In the last phase, they were asked to repeat their assets allocation, based only on investment features and no labels. This facilitated a within-person test of the effect of labels on individuals' portfolio choice.

Test Phase 1: Pair-Wise Investment Choice When asked to choose between two investments, subjects were invited to explore a 6 x 2 matrix displaying in each of two rows the two alternative investments (Investment 1, Investment 2) and in each column six investment features: risk, time horizon, cost, liquidity, capital loss, returns. There were no constraints on how customers should look up feature information even if there was a constraint on the number of possible features looked up. Of the 12 features they could look up only 6[3]. The test began with a black matrix on the screen initially hiding all the information content.

Information popped up in a "flipping cards" fashion when the subject touched the display. Each subject was asked to explore those features that they considered helpful for identifying their preferred investment (see Figure 1). Each subject performed, on average, around four different trials at this test phase.

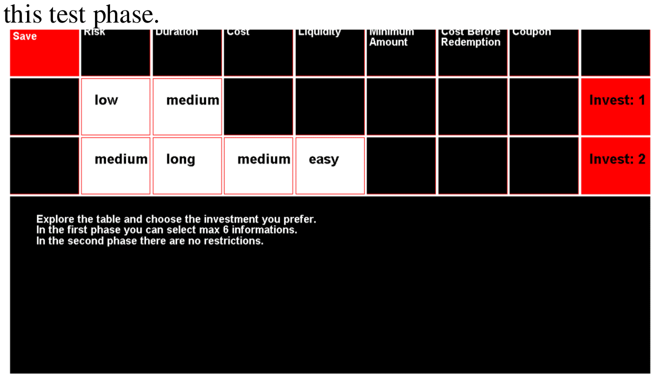

Figure 1: Graphic user interface for Test Phase 1

Test Phase 2: Extended Information Search - Financial Market Exploration The information provided was now arranged in a 7 x 6 matrix, displaying the same feature profile of an investment in each row for six different investments typically available in banks, namely, bank accounts, bonds issued by the mutual bank, by the government, by insurance companies, and balanced mutual funds (with a roughly 50-50 portfolio in corporate bonds and blue chip stock equity) and stocks.

Here again, the test is similar to the test Phase 2. Customers performed one exploration trial and were subsequently invited to continue the test by selecting their favorite investments portfolio (i.e., a set of weights on categories that add to 1) within the presented investment categories (see Test Phase 3).

Test Phase 3: Investment Portfolio – Categorization and Selection Unlike in Test Phase 2, in Phase 3 investors were now provided with the full information matrix uncovered from the very beginning of the test and they were asked to form a portfolio by allocating 100 units.

[3] The advisors' questionnaires revealed that there are limits, due to time constraints and to constraints in customers' understanding of financial information, to the exchange of the information on the investments. The upper limit usually considered is 6 pieces of information.

Figure 3 illustrates the matrix: The first column reported the name (label) of the investment, the set of white boxes revealed the investment features and the last column collected the investors' allocating decisions; with each touch of the box the allocated amount for that investment was increased by 5%.

Test Phase 4: Investment Portfolio 2 - "Blind" Categorization and Selection This phase was identical to the preceding one with the only difference that now the first column of the investments labels was hidden (see Figure 4); no changes were introduced for the other information.

3. Test Results

The tests results are presented in two blocks, one concerning the information search and the other concerning the decision strategies adopted by participants.

3.1 Part I: Information Search

In examining the approach followed by subjects in exploring financial information, we started by considering how much information an investor needed in order to arrive at his or her financial decision. We investigated the information search processes occurring both in Test Phase 1 (pair-wise investments comparison) and in Test Phase 2 (extended information search - financial market exploration).

Amount and Type of Explored Information In Test Phase 1, 86% of customers looked at all six pieces of information. In Test Phase 2, customers considered, on average, less than half of the available information (45%), revealing a clear preference for smaller information sets to act upon (Table 1); subjects probably focused on those subsets of financial products that mostly captured their interest.

Feature	Exploration (in %)
Risk	76.2
Time horizon	48.8
Costs	47.6
Liquidity	41.7
Coupon	39.3
Minimum Amount	38.1
Cost Before Redemption	26.2
Mean	45.4
Standard deviation	15.5

Table 1: Average Amount of Information Investigated from Each Single Investment Feature in Test Phase 2.

Information Search over Time In Test Phase 1, participants sequentially explored at most 6 different pieces of information dealing with the investment features in 65 trials; therefore, we analyzed data according to the 6 exploration steps denoted by $t_1 \dots t_6$. This sequential analysis reveals results consistent with Table 1; information concerning risk, time horizon and costs are looked up first

At time 1, risk was looked up in 89.2% of the cases. At Time 2, risk still was looked up in 41.5% and time horizon in 40% of the cases. At Time 3, cost was looked up in 35.4% and time horizon in 26.2% of the cases, and so on. This analysis reveals that, within the first two times, risk and time horizon are the most explored investment features. From Time 4 onwards, there appears to be no strong preference for any of the remaining features. During Times 1, 2 and 3 the preferred exploration path was risk -> time horizon -> cost. Figure 2 shows the aggregate view looking at the total number of cue look-ups in the grand pool of all customers over the six-step time path.

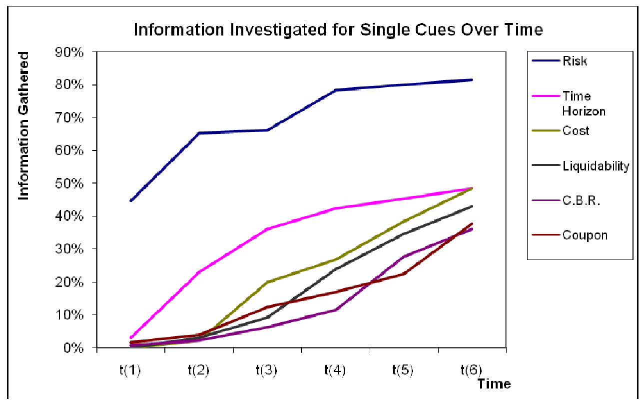

Figure 2: Information Gathered for Each Cue Over Tim in Test Phase 1.

We also estimated a **Markov transition matrix** with empirical probabilities of moving from one investment feature to another in the 6-step information search process. At the beginning, the feature most likely to be explored is risk (89%). The next feature after risk most likely to be explored is either time horizon (35%) or risk once again (23%). The next feature most likely to be explored after time horizon is either cost (46%) or time horizon once again (16%).

Payne's Analysis of Information Exploration Following the approach to information search analysis proposed by Payne, Bettman, and Johnson (2004), we looked at two types of exploration paths: feature-wise and investment-wise. A feature-wise path corresponds to an investor focusing on just one feature and exploring it across investments. An investment-wise path corresponds to an investor exploring features belonging to just one investment at a time.

Data collected in Test Phase 2 show that 8 out of the 14 participants (57%) adopted an investment-wise path; they focused their attention on information pertaining to a single investment at a time. The simultaneous protocol analysis revealed that most of those customers began their explorations from the investments they had already experienced in real life. The other customers explored the available information by adopting mixed strategies; some of them completely explored the information dealing with risk by adopting a cue-wise approach, while others gathered information across all the investments without revealing a predominant approach.

3.2 Decisions

In order to understand to what extent a decision tree (Figure 3) is able to capture a single investor's choice rule we introduce some definitions. Define the *cue profile* of an investment as a binary vector of 1's and 0's, according to whether cue values are "positive" or not and ordered by the sequence: risk, time horizon, liquidity, costs (intermediary fees), other costs[4], and returns. Based on the fast and frugal heuristic model (Gigerenzer, Goldstein, 2006) investment features are all transformed to binary values to simplify their comparison. The convention for assigning the values 1 or 0 to a cue reflects the preferences revealed by customers in their interviews. If risk, say, is medium or low, it is assigned the value 1. Similarly, if time horizon is medium or short, its value is 1, and if cost and liquidity are medium or low, they are also assigned the value 1. If the investment has no cost before redemption date this cue is assigned a 1, otherwise 0, and if there are "returns during the holding time" then this cue is assigned a 1, otherwise 0.

The heuristic that best modeled our data, lexicographically examines only the one cue that was explored most, namely risk, and processes all the remaining cues by means of a tallying rule (Figure 3); tallying is a heuristic that can be described by a linear model with weights equal to one for each investment feature. In this context, tallying means counting the number of 1's for both investments and choosing the investment with a higher score. For instance, if Investment A has a cue profile, (011111) and B has a cue profile (100000), then B is preferred because the first cue is treated lexicographically. As an example, if A is an investment with a cue profile (100101) and B is an investment with a cue profile (100100), investment A is chosen over investment B because its profile contains more 1's after the first entry.

How well does the previously described heuristic predict the choices observed in the Test Phase 2 task? Predictions for each subject are presented in Table 2.

By Fast and Frugal Tree we mean a tree that has at least one exit at each level; it is "minimal" among trees using all cues, because it has a minimal number of nodes (Martignon, Katsikopoulos & Woike, 2008).The tree in Figure 3 predicts 78.3% of the observed investment decisions in the experimental investment task of Phase 2. One of its key features is that, for most investors, there is no compensating trade-off for high risk investments. High risk investments are eliminated from consideration in the lexicographic formulation depicted in Figure 7. The second key feature is

Subject Id.	Scores
1	90%
2	80%
3	100%
4	50%
5	38%
6	90%
7	100%
8	88%
9	100%
10	88%
11	70%
12	60%
13	80%
14	40%
15	100%
Mean	78.3%
Standard Deviation	21.8%

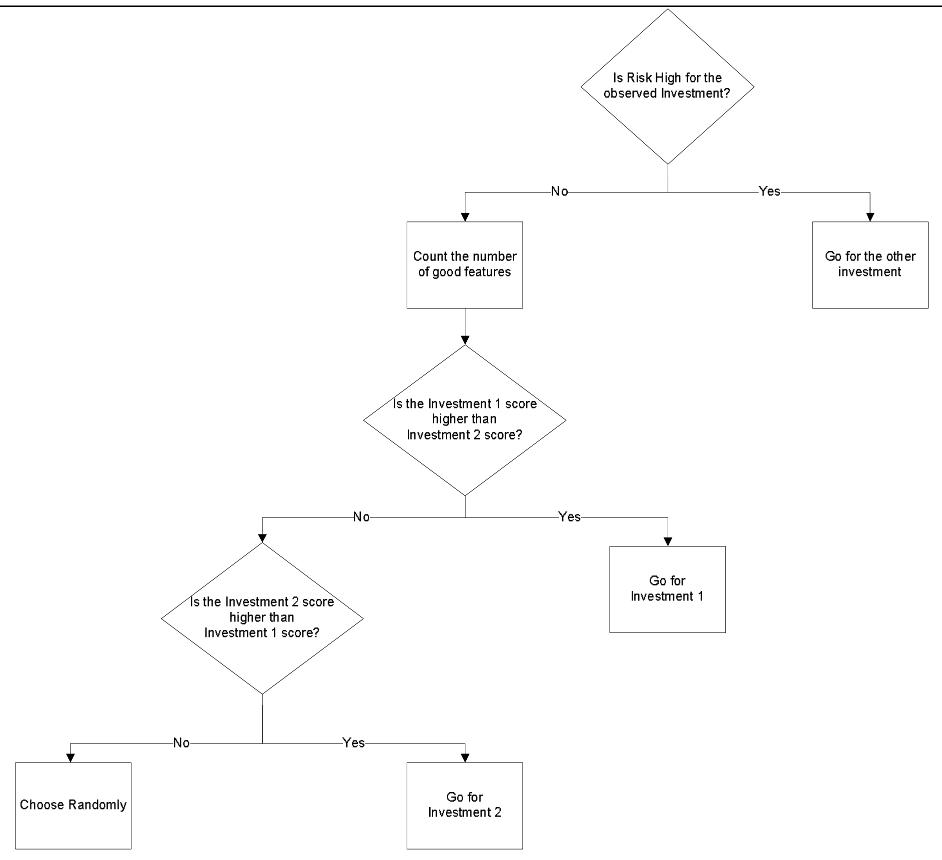

Table 2 (left): Decision Heuristic Predictions.
Figure 3 (right): Fast and Frugal Tree for Heuristic on Risk, Time Horizon and Costs and in Addition Tallying.

[4] For instance, costs for selling it before the redemption date.

that, beyond this lexicographic step to avoid high-risk investments, customers adopt a simple tallying rule that counts 1 for each cue value that matches their system of preferences, otherwise 0, and choose the investment with the higher score. In other words, rather than weighting different features differentially, the model suggests that investors simply count the number of features over which one investment dominates another one.

The Importance of Recognition for Portfolio Design We now present the effects of investment labels on subjects' classifications of investments. In Test Phase 4 all the 15 participants performed one task dealing with what we called "blind categorization". We collected data on participants' performances in reproducing the same investments allocation task they had already performed in Test Phase 3 but now without providing them with the investment names or labels, but just with their features. The idea was to test how consistent their choices remained when provided with just the investment features and not with their names.

The first result we obtained was that participants did not excel in the blind classification task. Given that participants paid most attention to risk, they should have split investments in two different categories, namely, high-risk investments versus medium- and low-risk investments, even when investment labels were absent. The empirical evidence shows us that 9 out of 15 participants (60%) made important errors, that is, they invested in much riskier portfolios than before and with asset allocations that diverged from the original ones, on average, by 69% (calculated on the amount of the originally invested money). These results give us a perspective on how people perceive, represent, and act upon financial information and reveal a delicate aspect for potential manipulation of decisions. The results confirm the power of the Recognition Heuristic: when it is not applicable, customers perform poorly[5].

4. Summary and Conclusions

The aim of this research was to investigate how average investors make financial decisions. We focused on two components of their decision processes: information search and decision. We interviewed 80 customers of an Italian mutual bank, of whom 15 were also tested in an interactive lab experiment. For this experiment, we designed naturalistic environments based on the information we collected from the customers and advisors' interviews: We preserved the characteristics of investment choices usually offered to those same bank customers. The experiment consisted of 4 tests. In Test Phase 1 we let customer choose between 2 investments constrained to looking up 6 possible feature variables at most. In Test Phase 2, by contrast, they had to choose between 6 types of investment and could still exploring the same financial features. Subjects in Test Phase 2 consulted less than half of the information at their disposal. The exploration of information is characterized by frugality and simplicity. We modeled customers' choices by means of a fast and frugal tree trough which we predicted namely 80% of investors' decisions. They consider only 1 feature lexicographically, as non compensated by others, and treats the remaining ones with a tallying rule. For our customers to tally features, when they do not know them well is simpler than establishing a ranking. The second experimental task was portfolio choice, designed to check whether providing participants with all the investment features but no labels/names can affect their behavior. We discovered that, when labels are missing, customers tend to select different and riskier investments than the original ones (namely when labels available). That is, the way financial products are presented to average investors plays a significant role in their investment decisions.

Acknowledgements

We thank Davide Donati, general director of the Cassa Rurale Giudicarie Valsabbia Paganella and all the board of executives for their full support for conducting this research. We also thank Marcel Jentsch for his valuable help in programming and designing interactive interfaces.

References

Chewning, E., & Harrell, A. (1990). The effect of information overload on decision makers' cue utilization levels and decision quality in financial distress decision task. Accounting, Organizations and Society, 15(6), 527–542.

Gigerenzer, G. (2004). Fast and frugal heuristics: The tools of bounded rationality. In D. Koehler & N. Harvey (Eds.), *Handbook of judgment and decision making.* (pp. 62–89). Oxford, UK: Blackwell.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review, 103(4),* 650–669.

Katsikopoulos, K., & Martignon, L., (2006). Naive Heuristics for Paired Comparision: Some Results on Their Relative Accuracy, *Journal of Mathematical Psychology, 50(5),* 488-494.

Lee, B. K., & Lee, W. N. (2004). The effect of information overload on consumer choice quality in an on-line environment. *Psychology & Marketing, 21,* 159–183.

Martignon, L., & Hoffrage, U. (2002). Fast, frugal, and fit: Simple heuristics for paired comparison. *Theory and Decision, 52(1),* 29–71.

Payne, J., Bettman, R., Johnson, E. G. (2004). Fast and frugal heuristics: The tools of bounded rationality. In D. Koehler & N. Harvey (Eds.), *Handbook of judgment and decision making.* (pp. 62–89). Oxford, UK: Blackwell.

[5] By Recognition Heuristic we mean a simple strategy that allows individuals to infer, for example, which of the two objects has a higher value on some criterion based on the fact the recognize one and not the other (e.g. which investment has a higher expected return). The recognition heuristic for such tasks is simply stated: If one of the two objects is recognized and the other is not, then infer that the recognized object has the higher value (Goldstein & Gigerenzer, 2002).