Investment decisions: New descriptive models of bank customers' behaviour

Monti, M.,* Martignon, L., ** Gigerenzer, G.,* Berg, N.***

* Max Planck Institute for Human Development

** University of Education, Ludwigsburg

*** University of Texas at Dallas

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' behaviour. 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 behaviour 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 more than 80% of investors' decisions.

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 behaviour. 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 (Balzer, Doherty, & O'Connor, 1989; 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 behaviour? 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 Brighton & Gigerenzer, 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. Empirical support for the use of TTB in human populations comes from a variety of sources: Bröder & Schiffer, 2003b; Bergert & Nosofsky, 2007; Bröder & Gaissmaier, 2007; Nosofsky & Bergert, 2007; Rieskamp & Otto, 2006). Theoretical research provides characterizations of the statistical features of decision-making environments in which TTB can be successfully applied (Gigerenzer & Brighton, 2009; Hogarth & Karelaia, 2005a, 2005b; Baucells, Carrasco, & Hogarth, 2008, Martignon & Hoffrage, 2002).

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 behaviour. The neoclassical model implies that no relevant information is discarded or ignored. In terms of investors' goals, the model implies that investor's behaviour 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, we report evidence that clearly rejects the neoclassical model and fits a lexicographic fast and frugal tree akin to TTB. Subjects in our study rely on simple decision trees. 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 (which accords with the normative assessments in Magni, in press); TTB and similar lexicographic decision-tree heuristics consider the features of investments sequentially in an 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).

Although it remains an open question as to what extent bank customers are able to judge the quality of their decisions in terms of outcomes and of processes, we noticed that our subjects clearly tended to consider specific combinations of investment features, and use these combinations according to predictable rules, even though they lacked theories on causal links between features and decision criteria.

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 CAPM.

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 behaviour 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 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 in the questionnaires. We used information from financial advisors' interviews to calibrate the features of the investment choices offered to subjects with realistic or typical values of expected return, risk, etc.

[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.

2.2 Experimental Design

We investigated decision strategies by considering three factors: the overall amount of information that subjects need for taking their decisions, the type of information (features or cues) that they consider before choosing their investments, and the approach they follow in the information-search process.

2.2.1 Participants

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.

2.2.2 Test description

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.

2.2.3 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[3]. 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[4]. 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.

[3] This terms are translations of the Italian terms (rischio, durata, costo, liquidabilita', perdita in conto capitale, interessi) used by our Bank without further explanations assuming the customers have familiarity with them.
[4] 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.

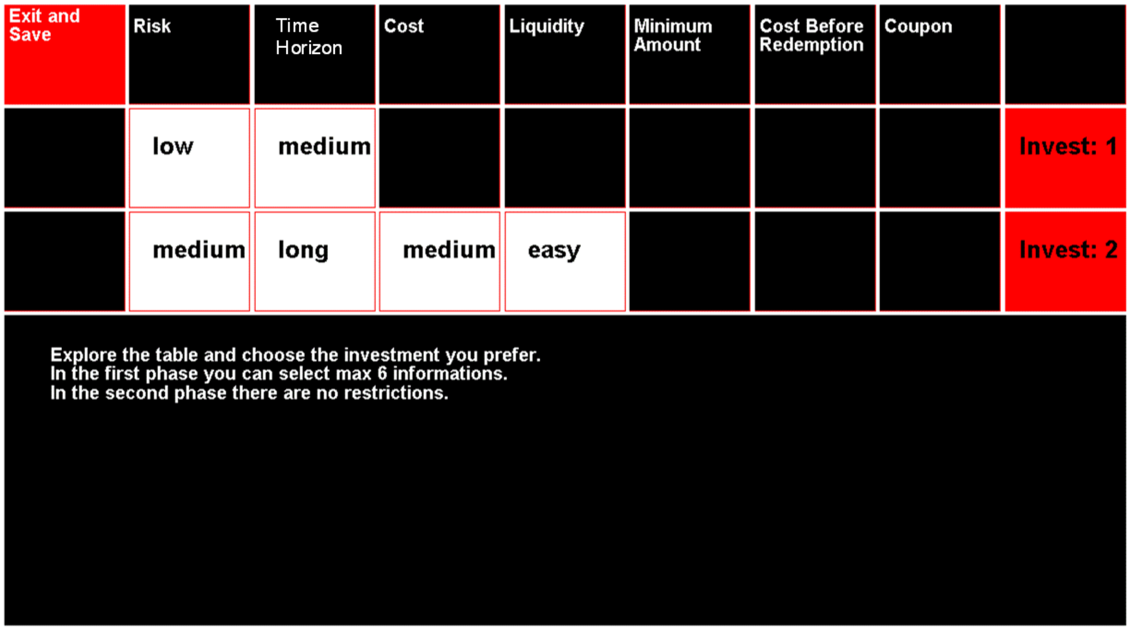

Figure 1: Test Phase 1 - Pair-Wise Investments Comparison.

2.2.4 Test Phase 2: Extended Information Search - Financial Market Exploration

Customers were asked to explore financial information they considered necessary for designing an investment portfolio.

The information provided was arranged in a 7 x 6 matrix, displaying the feature profile of an investment in each row, namely, its values on risk, time horizon, cost, liquidity, minimum amount[6], cost for anticipated demission and returns for six different investments typically available in banks, namely, bank accounts, bonds issued by the mutual bank, bonds issued by the government, bonds issued 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 began with a black matrix on the screen, hiding all information content. Pertinent information about a hypothetical "financial market" popped up when

[6] This is a new investment feature introduced at this phase of the test in order to complete the possible information set.

subjects touched the display. Participants were instructed to uncover those entries that they considered necessary for making their choices; no restrictions were imposed.

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).

Exit and Save	Risk	Time Horizon	Cost	Liquidity	Minimum Amount	Cost Before Redemption	Coupon
Bank Account							
Mutual Bank Bounds	low	short	low	easy		yes	yes
Governmet Bonds	low	medium	low			yes	yes
Insurance with Guaranteed Capital		medium					
Balanced Mutual Funds	medium						
Stocks	high					no	

Figure 2: Test Phase 2: Extended Information Search - Financial Market Exploration.

2.2.5 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.

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%.

Next Trial or Exit and Save	Liquidity	Coupon	Cost Before Redemption	Risk	Time Horizon	Cost	Split the Pie 70/100
Insurance with Guaranteed Capital	difficult	no	yes	low	medium	high	5%
Mutual Bank Bounds	easy	yes	yes	low	short	low	10%
Bank Account	easy	yes	no	low	short	low	15%
Governmet Bonds	easy	yes	yes	low	medium	low	20%
Balanced Mutual Funds	easy	no	no	medium	long	medium	20%
Stocks	easy	no	no	high	long	medium	0%

Figure 3: Test Phase 3: Investment Portfolio Selection – Categorization and Selection

2.2.6 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.

Next Trial or Exit and Save	Cost	Risk	Cost Before Redemption	Liquidity	Time Horizon	Coupon	Split the Pie 50/100
	medium	medium	no	easy	long	no	0%
	low	low	yes	easy	medium	yes	10%
	medium	high	no	easy	long	no	15%
	low	low	no	easy	short	yes	25%
	low	low	yes	easy	short	yes	0%
	high	low	yes	difficult	medium	no	0%

Figure 4: Test Phase 4: Investment Portfolio Selection 2 - “Blind” Categorization and Selection

The investors were asked to state their decision processes aloud while the experimenter kept a written protocol; we digitally recorded the descriptions subjects reported while performing the tests and then we summarized them into summary schemes.

3. Results

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

3.1 Part I: Information Search

In examining the approach followed by customers in exploring financial information, we started by considering how much information an investor needed in order to arrive at an investment 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).

3.1.1 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 (Figure 5); subjects probably focused on those subsets of financial products that mostly captured their interest.

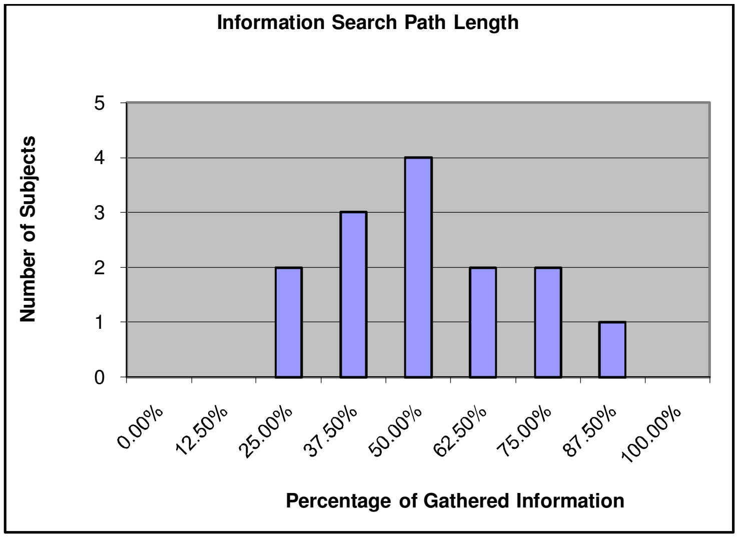

Feature	Exploration (in %)
Risk	76.2
Time horizon	48.8
Costs	47.6
Liquidity	41.7
Coupon	39.3
Minimum Amount	38.1
Cost Before Redemption	26.2
Mean	45.4
Standard deviation	15.5

Figure 5 (left): Differences in the Amount of Information Gathered in Test Phase 2.

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

Table 1 presents the most investigated cues across investments at the aggregate level. On average, investors focused mostly on information dealing with risk, time horizon and costs. For example, 76.2% of the information presented by risk was explored.

3.1.2 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… t_6. This sequential analysis reveals results consistent with Table 1; information concerning risk, time horizon and costs are looked up first (see the bold figures in Table 2). In table 2 we present the averages computed per trial.

Feature	t_1 (in %)	t_2 (in %)	t_3 (in %)	t_4 (in %)	t_5 (in %)	t_6 (in %)
Risk	**89.2**	**41.5**	1.5	**24.6**	3.1	3.1
Time horizon	6.2	**40.0**	**26.2**	12.3	6.2	6.2
Cost	0.0	4.6	**35.4**	13.8	**23.1**	20.0
Liquidity	0.0	6.2	12.3	**29.2**	**21.5**	16.9
Cost Before Redemption	1.5	3.1	7.7	10.8	**32.3**	16.9
Coupon	3.1	4.6	16.9	9.2	10.8	**30.8**

Table 2: Information Exploration Over Time in Test Phase 1 (in %).

Table 2 reveals that 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 information search 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 preferences for any of the remaining features. During Times 1, 2 and 3 the preferred exploration path was risk -> time horizon -> cost. Figure 6 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.

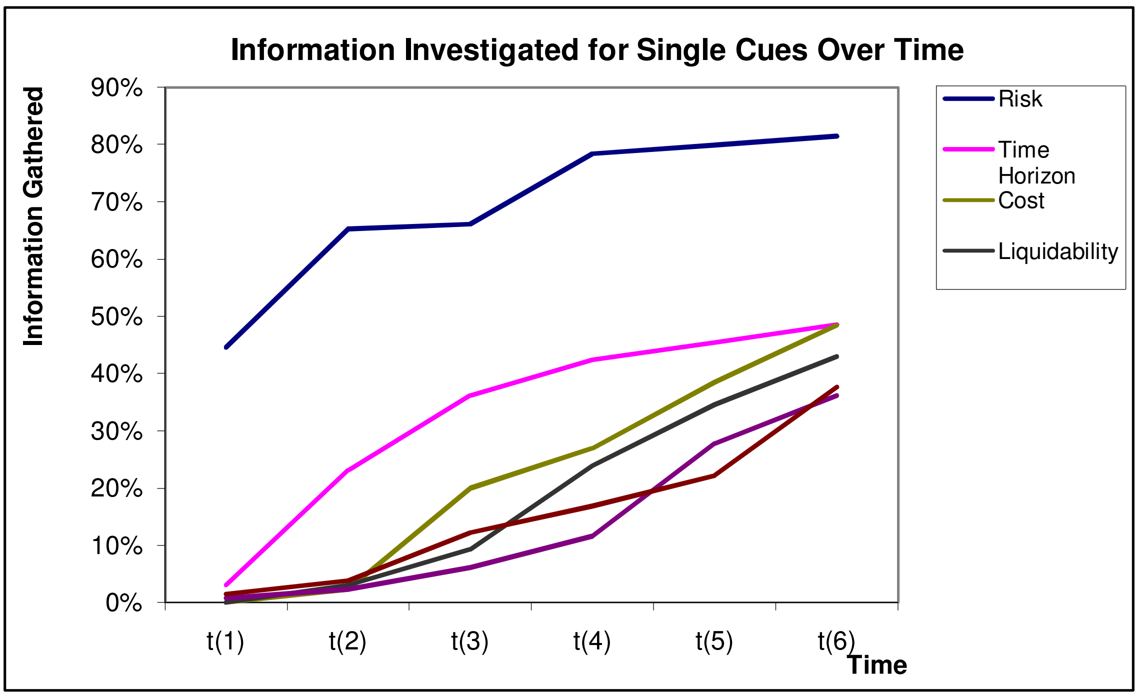

Figure 6: Information Gathered for Each Cue Over Time in Test Phase 1.

We estimated a Markov transition matrix with empirical probabilities of moving from one investment feature to another in the 6-step information search process, presented in Table 3. At the beginning (start position), 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%) – (see the bold figures in Table 3).

From Feature to Feature	Start (in %)	Risk (in %)	Time horizon (in %)	Cost (in %)	Liquidity (in %)	Cost Before Redemption (in %)	Coupon (in %)	End (in %)
Start	0	**89**	6	0	0	2	3	0
Risk	0	**23**	**35**	11	14	4	11	2
Time horizon	0	14	**16**	**46**	6	10	2	6
Cost	0	5	5	**14**	**33**	11	10	22
Liquidity	0	5	7	7	11	**32**	**16**	21
Cost Before Redemption	0	6	6	11	4	**17**	**30**	**26**
Coupon	0	12	4	8	**16**	6	10	**43**
End	0	0	0	0	0	0	0	0

Table 3: Transition Probabilities between Features Observed in Test Phase 1 (in %).

3.1.3 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 (e.g., government bond for subject 9, mutual bank bonds for subject 7, etc.). 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).

3.1.4 Overlapping Information Index and Order Preservation Index

We investigated information search to answer the following questions: Did customers look at identical information for both investments? Did individuals explore investment features by following a well-established common order? We adopted a within-subject approach.

We introduced two indices characterizing customers' information search. We noticed that in our experiments, customers did not necessarily collect overlapping information on different investments before making choices. In other words, when considering investments A and B (Test Phase 2), customers did not check the same features for A and B before choosing. Such measures of systematic search, as the 2 indices we introduced, are not commonly mentioned in the experimental psychology literature, although they are certainly relevant for describing information usage and decision process, which will provide an indication of whether neoclassical or heuristic models better describe how they choose investments.

These 2 specific measures are: The Overlapping Information Index (OII) and the Order Preservation Index (OPI). We defined the OII as of the percentage of identical features looked up for both investments across customers. OPI indicates the percentage of overlapping features explored in identical order[7]. Let us look at 2 examples: Suppose that a subject explores features of different investment options in the following order: risk, liquidity, then cost, for investment 1, and cost, risk, then liquidity, for investment 2. Another example: risk, cost, liquidity form investment 1, and risk, liquidity cost, for investment 2. For both examples the OII is 100%, since three out of the three features are looked up for both investments and for both examples the OPI is equal to 66.6%, since two out of the three overlapping features are looked up in the same order.

We decided to classify customers according to their OII; we chose a threshold of 50% and distinguished two groups:

- High Overlapping Information Index Group (HOI); customers show an OII greater or equal than 50%;
- Low Overlapping Information Index Group (LOI); customers show an OII lower than 50%.

[7] See Appendix 1 for details on the construction of these indices.

OPI and OII are highly correlated (corr. = 0.98), which means that customers belonging to the HOI group preserve their exploration order across explorations and focus their attention on a smaller set of financial features. This fact suggests that the selection of the considered information set is connected with the exploration approach and with the investment representation space and, thus, with the decision mechanism. From Table 4, we deduce that a part of our sample looked for coincident information for both investments and follow the same sequential order across cues.

Group	Subject Id.	Overlapping Information Index (in %)	Number of Explored Cues	Order Preservation Index (in %)
HOI	2	100.00	3.00	100.00
HOI	7	100.00	3.00	100.00
HOI	3	91.67	3.25	91.67
HOI	8	75.00	3.75	75.00
HOI	1	73.33	3.80	66.67
HOI	5	58.33	4.25	50.00
HOI	10	50.00	4.25	41.67
HOI	15	50.00	4.50	50.00
LOI	13	46.67	4.60	40.00
LOI	9	41.67	4.75	41.67
LOI	6	33.33	5.00	6.67
LOI	14	33.33	5.00	26.67
LOI	12	20.00	5.40	13.33
LOI	11	13.33	4.80	6.67
LOI	4	8.33	5.75	0.00
HOI Group Mean		74.79	3.72	71.87
LOI Group Mean		28.09	5.04	19.28
Sample Mean		53.00	4.34	47.33

Table 4: Overlapping Information Index and Order Preservation Index Calculated for Every Subject in Test Phase 1.

3.2 Decisions at the Aggregate Level

The following aggregate data (see Table 5) represent relationships between the customers' consideration of specific financial features and their decisions in the investment task. The data reveal that, across all the subjects' decisions, risk was selected in 41 cases for both investments and was discriminating (as different values) in 26 cases. In 20 of these cases (76.9 %), the bank customers preferred the less risky investment. Time until maturity was selected for both investments in 15 cases and was discriminating in 14 of these. In 10 of these cases (71.4%), the bank customers preferred the shorter investment. Cost was selected for both investments in 13 cases and was discriminating in 10 cases. In 8 of these cases (80%), the bank customers preferred the cheaper investment.

Investment Feature	N. of Lookups	Smaller Value Chosen	Larger Value Chosen	Frequency of Choice for the Lower Value (in %)
Risk	41	20	6	76.9
Time horizon	15	10	4	71.4
Costs	13	8	2	80
Liquidity	13	6	0	100
Cost Before Redemption	8	3	3	50
Return	11	0	8	0

Table 5: Relationships between Information Search and Decision in Test Phase 1.

3.2.1 How Do Customers Make Decisions? A Within-Subject Analysis

In order to understand to what extent a decision tree (Figure 8) 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[8], 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

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

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 7); 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 6.

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%

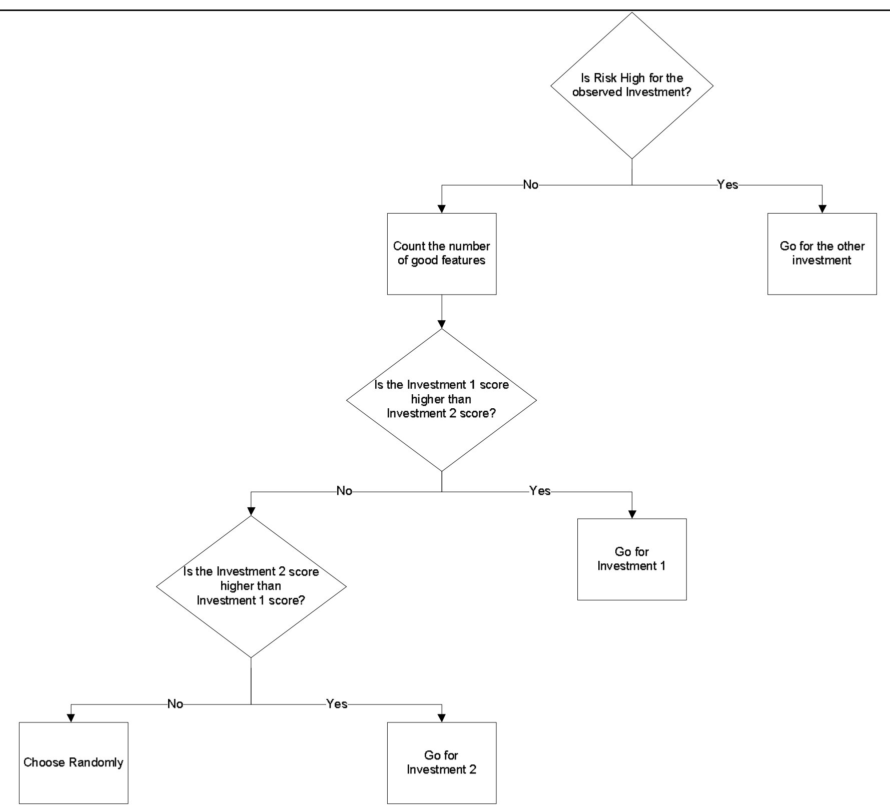

Table 6 (left): Decision Heuristic Predictions.

Figure 7 (right): Fast and Frugal Tree for Lexicographic Heuristic on Risk and on the score of the investment (which can be seen as a new cue of the investment).

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 7 predicts 80% 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 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.

Since we noticed that much of the observed information search focused on risk, time horizon and cost, we investigated the question: For which portion of participants' decisions was this sequential information search reflected in a lexicographic decision rule for those same three cues? Figure 8 presents a model in which customers' decisions follow a fast and frugal tree that lexicographically processes all those three cues. For investments with identical values for all three of these features, investors are modeled to proceed by means of tallying rule on the remaining ones; in case of a tie, they flip a coin.

Subject Id.	Scores (in %)
1	20
2	10
3	25
4	50
5	25
6	60
7	0
8	75
9	88
10	50
11	60
12	40
13	80
14	40
15	50
Mean	44.8
Standard Deviation	25.6

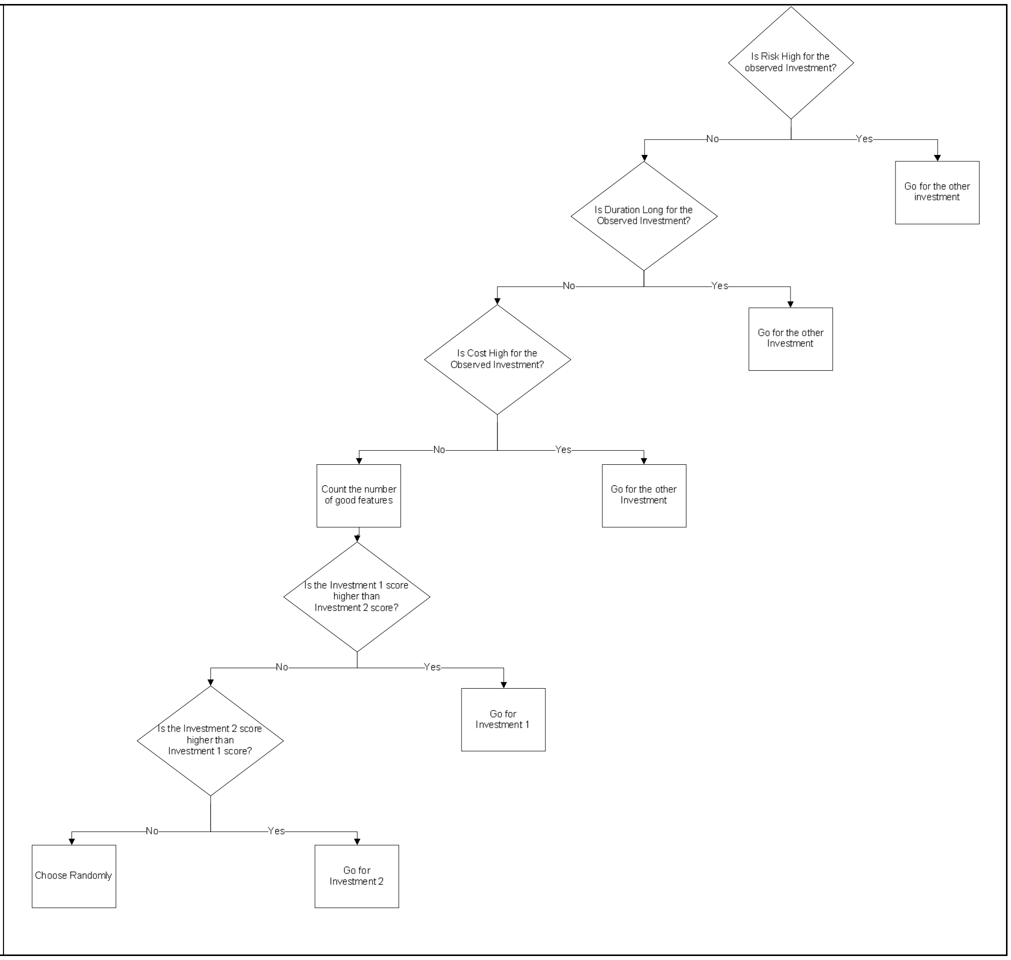

Table 7 (left): Decision Heuristic Predictions.

Figure 8 (right): Fast and Frugal Tree for Heuristic on Risk, Time Horizon and Costs and in Addition Tallying.

The tree described by Figure 8 operates as follows: An investment with a cue profile (100000) is preferred to any other investment having a 0 as first entry. If investments A and B have the same value on the first entry, and A has a 1 and B has a 0 as second entry, A is chosen. The same procedure holds for the third entry, when first and second entries coincide. The strategy changes after the first three cues are looked up. If both investments have the same values in the first three entries, the number of 1's after the third entry are counted in both cue profiles and the investment with the higher score is chosen. If there is a tie, one investment is chosen randomly. The policy thus described combines two prototypes of simple strategies, namely a lexicographic rule with one that integrates equally weighted cues (Dawes, 1974).

Data show that the first tree (Figure 7) captures most of the customers' policies (78.3%) while only in about half of the cases (44.8%) did customers decide lexicographically on all the first three mostly explored cues (Figure 8).

Both the lexicographic and the tallying strategies have been extensively examined by the ABC group and both are considered simple in their categories, namely trees and linear models (Martignon & Hoffrage, 2002; Katsikopoulos & Martignon, 2006), although the lexicographic rule is more frugal in terms of requiring few resources. The tallying rule is less "finely tuned" than the lexicographic one, because it does not discriminate between profiles having the same number of 1's, independent of where in the cue profile the 1's are. Furthermore, the tallying rule quite often makes the same choices as the lexicographic one.

3.2.2 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; the only difference we introduced at this phase is that we did not provide them with the investment names or labels, but just with their features. Therefore, they were still exposed to the same 6 investments types displayed in Test Phase 3; namely, stocks, mutual balanced bonds, government bonds, bank account, mutual bank bonds, insurances with guaranteed capital.

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[9].

The following table reveals investors' portfolio choices: data in bold represent customers who made mistakes in the "blind categorization" task.

Subject Nr.	With Labels (in %)	Without Labels (in %)	|delta| (in %)
13	**5**	**20**	**300**
12	**25**	**45**	**80**
1	**15**	**25**	**66.7**
3	**20**	**30**	**50**
2	**15**	**20**	**33.3**
16	**15**	**20**	**33.3**
11	**50**	**65**	**30**
10	**25**	**30**	**20**
4	**60**	**65**	**8.3**
5	40	40	0
6	0	0	0
7	0	0	0
8	40	40	0
14	5	5	0
15	0	0	0
Mean	21.0	27.0	41.4
Standard Deviation	19.0	21.4	6.6

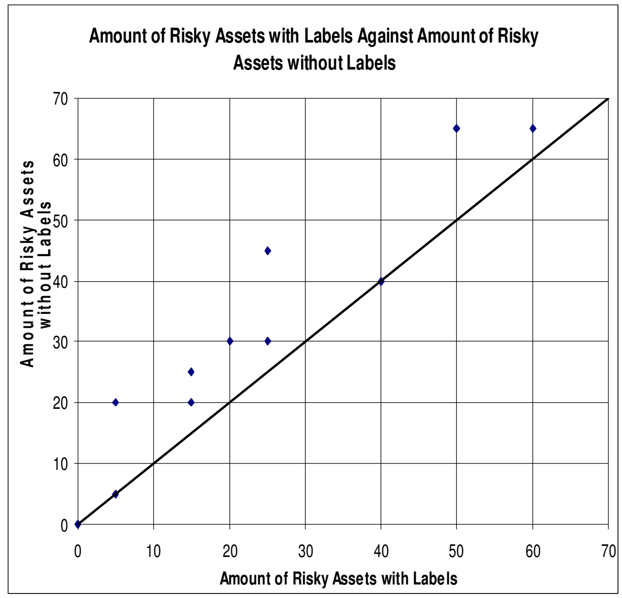

Table 8: Comparison Between Allocated Money in the 2 Test Conditions.

Figure 9: Comparison Between the Amount of Money Invested in Risky Assets When Labels are Provided (Test Phase 3) Versus Amount of Money Invested in Risky Assets When Labels are not Provided (Test Phase 4).

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

[9] **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).

decision (that is the mapping from observed information into an investment choice). 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' interviews: We preserved the characteristics of investment choices usually offered to those same bank customers. The experiment consisted of 4 tests focused on the building blocks of a decision process, namely, information search and decision rule. 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 consult all 36 pieces of information available. We observed that in Test Phase 2 they consulted less than half of the information at their disposal. In particular, a significant proportion of investors (57%) explored just a small set of non-overlapping features for pairs of investment alternatives. This provides new empirical insight on investors' search and investors' use of information when making financial decisions. The evidence suggests that investors do not necessarily search for similar sets of information, and in a similar sequence for different investments: The exploration of information is characterized by frugality and simplicity.

We modeled customers' choices by means of a fast and frugal tree where risk, time horizon, and costs are considered in a lexicographic manner and the remaining cues according to a tallying rule. This decision tree characterizes the approach followed by customers in 44% of their decisions.

A tree that covers more of investors' decisions, namely 80%, considers 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 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 can affect customer behaviuor. 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.

The recent tendency of financial markets to add huge numbers of new labels and to change their set of features, adding for instance ratings of expert agencies, may lead investors to make riskier choices. Observing how average investors tend to use simple heuristics for their portfolio selection suggests the conclusion that rapidly and, perhaps, strategically changing financial environments, can lead these investors to make undesirably riskier choices.

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

Balzer, W. K., Doherty, M. E., & O'Connor, R. (1989). Effects of cognitive feedback on performance. Psychological Bulletin, 106(3), 410–433.

Baucells, M., Carrasco, J. A., & Hogarth, R. M. (2008). Cumulative dominance and heuristic performance in binary multi-attribute choice. Operations Research Volume: 56(5), 1289-1304 .

Berg, N., Hoffrage, U. (2008). Rational ignoring with unbounded cognitive capacity. Journal of Economic Psychology, 29(6), 792-809.

Bröder, A. (2003). Decision-making with the "adaptive toolbox": Influence of environmental structure, intelligence, and working memory load. Journal of Experimental Psychology: Learning Memory and Cognition, 29(4), 611–625.

Bröder, A., & Schiffer, S. (2003b). Take the Best versus simultaneous feature matching: Probabilistic inferences from memory and effects of representation format. Journal of Experimental Psychology: General, 132(2), 277–293.

Chater, N., Oaksford, M., Nakisa, R., & Redington, M. (2003). Fast, frugal, and rational: How rational norms explain behaviour. Organizational Behaviour and Human Decision Processes, 90(1), 63–86.

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, Organization, and Society, 15(6), 527–542.

Dawes, R. M. (1979). The robust beauty of improper linear models in decision-making. American Psychologist, 34(7), 571–582.

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

Gigerenzer, G. & Todd, P. (1999). Fast and frugal heuristics: The adaptive toolbox. In Gigerenzer G., Todd P., M., Simple Heuristics that Make Us Smart, (pp.3—34)., New York, USA: Oxford University Press.

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. & Brighton, H. (2009). Homo Heuristicus: Why Biased Minds Make Better Inferences. Topics in Cognitive Sciences, 1, 107-143

Goldstein, D. G. & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. Psychological Review, 109, 75-90.

Hogarth, R. M., & Karelaia, N. (2005a). Ignoring information in binary choice with continuous variables: When is less "more"? Journal of Mathematical Psychology, 49(2), 115–124.

Hogarth, R. M., & Karelaia, N. (2005b). Simple models for multi-attribute choice with many alternatives: When it does and does not pay to face trade-offs with binary attributes. Management Science, 51(12), 1860–1872.

Katsikopoulos, K. V., & Martignon, L., (2006). Naïve 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.

Magni, C. A. (in press). Investment Decisions, Net Present Value and Bounded Rationality.

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

Martignon, L., Katsikopoulos, K. V., Woike, J. K., (2008). Categorization with limited resources: A family of simple heuristics. Journal of Mathematical Psychology, 52, 352–361.

Needleman, S. B., & Wunsch, C. D., (1970). A general method applicable to the search for similarities in the amino acid sequence of two proteins. J. Mol. Biol., 48, 443–453

Newell, B. R., Weston, N. J., & Shanks, D. R. (2003). Empirical tests of a fast and frugal heuristic: Not everyone "takes-the-best". Organizational Behaviour and Human Decision Processes, 91(1), 82–96.

Newell, B. R. (2005). Re-visions of rationality? Trends in Cognitive Sciences, 9(1), 11–15.

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.

Waldmann, M. R., Holyoak, K. J., & Fratianne, A. (1995). Causal models and the acquisition of category structure. Journal of Experimental Psychology: General, 124,181-206.

Weenig, M. W. H., & Maarleveld, M. (2002). The impact of time constraint on information search strategies in complex choice tasks. Journal of Economic Psychology, 23(6), 689–702.

Appendix 1

Overlapping Information Index

-Set of visible information for investment 1 and -Set of visible information for investment 2

$$OII = |A \cap B|$$

Order Preservation Index

-Sequence of revealed information for investment 1, meaning that was revealed before

-Sequence of revealed information for investment 1, meaning that was revealed before

F – the computation matrix of

<u>Inititialization of the Matrix:</u>

<u>Recurrence (</u>Needleman-Wunsch-algorithm)

There

<u>Result</u>

<u>Example (see the example in the text)</u>

		R	L	C
	0	0	0	0
C	0	0	0	1
R	0	1	1	1
L	0	1	2	2