Financial Advisors and Their Clients: Information Search and Portfolio Choice Among Bank Customers

Marco Monti[1], Nathan Berg[2], Gerd Gigerenzer[3], and Laura Martignon[4]

[1,3]Max Planck Institute for Human Development-Berlin

[2]School of Economic, Political and Policy Sciences, University of Texas-Dallas

[4]University of Education, Ludwigsburg

Corresponding author: Nathan Berg: Associate Professor of Economics; School of Economic, Political and Policy Sciences; University of Texas at Dallas; 800 W. Campbell Rd., GR31; Richardson, TX 75083-3021; nberg@utdallas.edu

Keywords: Lexicographic, Non-compensatory, Heuristic, Bank Customers, Behavioral Economics, Ecological Rationality

Acknowledgements: We thank John Payne for very helpful comments and Michela Balconi for her ideas regarding data analysis. We also thank Davide Donati, general director of the Cassa Rurale Giudicarie Valsabbia Paganella, and all the Board of Executives, for their full support in conducting on-site research at their banks.   We likewise thank Marcel Jentsch for valuable help programming and designing interactive interfaces.

Financial Advisors and Their Clients: Information Search and Portfolio Choice Among Bank Customers

Abstract: Bank customers are not financial experts, and yet they make high-stakes decisions that can substantively affect personal wealth. This raises questions about how non-experts actually make financial decisions. How do investors search for information? How does the information they look up map into decisions? Using data collected from financial advisors and bank customers at several branches of an Italian mutual bank, this paper aims to reveal new insights about the decision processes of average investors: their investment goals, the information sets they consider, and the factors that ultimately influence decisions about investment products. We designed four portfolio choice tasks based on data collected directly from financial advisors and their clients. Our subjects were actual bank customers with investment accounts at those banks containing 40,000 euro or more. Most subjects used a limited set of information, ignoring factors that conventional economic models usually assume drive investor behavior. We suggest that non-compensatory decision-tree models, which make no trade-offs among investment features, be considered as parsimonious descriptions of investor behavior useful in marketing applications and policy contexts alike.

Financial Advisors and Their Clients: Information Search and Portfolio Choice Among Bank Customers

1 Introduction

Given rapid innovation in financial markets and a virtually constant flow of new investment possibilities, investors' choice sets and the information upon which investors condition future expectations (at least according to standard portfolio choice models) are far from static. As exciting and useful as this innovation is, however, the rapidly changing investment environment hardly seems to match the small worlds inhabited by agents found in standard portfolio choice models upon which the overwhelming majority of investment advice and financial literacy campaigns are derived. For example, Nobel Laureate Harry Markowitz , inventor of modern portfolio theory from which CAPM is derived, does not follow the prescriptions of his own widely applied model of portfolio choice when investing his retirement funds (Gigerenzer, 2008).[1] The reason is that, although standard models of portfolio choice theory deliver *a priori* perfection via expected utility maximization in their respective small worlds, they require enormous amounts of information and assume the world never changes. If one were to choose portfolio weights according to the prescriptions of Markowitz's model or CAPM, one would have to proceed just as the omniscient agents in those models do—exhaustive search through the entire universe of risky returns to consider expected values and covariances of each financial security, which are all assumed to be

[1] Markowitz describes his investment behavior as not following the prescriptions of his model (Gigerenzer, 2008). Instead, Markowitz relies on a simple 1 over N heuristic, sometimes referred to as naïve diversification, that divides investable funds evenly over all fund categories offered in his retirement plan. The 1 over N heuristic performs surprisingly well by many risk/return metrics, resulting in statistical properties that are in many cases superior to those achieved by standard finance models derived from expected utility maximization (De Miguel et al., 2009). As financial information proliferates and the real-world complexity of choice increases, there is little or no evidence (that we are aware of) demonstrating gains in financial literacy. See, for example, Lusardi and Mitchell (2006) documenting very low rates of financial literacy in the U.S. according to standard definitions of that concept.

unchanging and instantaneously known parameters of the stochastic but nevertheless static returns-generating process. Investors may indeed be smart not to apply prescriptions from this small world to real-world environments characterized by profound uncertainty—where last period's correlation matrix (of risky returns and everything else one might think is relevant) is quite possibly irrelevant for today's investment environment. Financial advice derived from models in which the statistical distribution of risky returns is static, the list of investments is exhaustively known, and all means and correlations are known, brings to mind statistician John Tukey's (1962) concept of a type-3 error, which is the right answer (in a small world made artificially precise) to the wrong question.

This study seeks to describe in greater detail the decision-making processes of bank consumers who make portfolio choice decisions in a non-small world, in consultation with financial advisors employed by banks. Qualitative and quantitative data were collected using in-depth interviews with both financial advisors and their clients concerning the typical range of investment alternatives presented to consumer investors. Interviews and survey data also provided quantitative descriptions of the typical number of items in investors' consideration sets and the investment features (i.e., information) that investors typically ask for and are supplied with by investment advisors. We used these stylized facts based on interview and survey data to construct a laboratory-type experiment consisting of four portfolio choice tasks designed to match the investment environment that actual bank customers face.

Thanks to an unusually cooperative experimental partner, a chain of regional banks in Italy in the suburbs of Trento, we were able to conduct these investment task experiments using random samples of actual bank customers and conduct them on-site inside the banks at which these same customers have investment accounts. We think this provides an unusually high degree of match between the subject population and laboratory environment, on the one hand, and the real-world investment behavior of primary scientific interest on the other.

There are five primary components of our investigation that represent the main questions we seek to address. What information do common investors want when making high stakes investment decisions? Given a large field of information about different investment alternatives and their features, how do common investors search for information? Even when the size of the choice set is modest, can we describe patterns of information search and information usage to reveal how consumers eliminate investments from consideration? Is there confirmatory evidence of pairwise comparison of common investment features consistent with the widespread assumption that consumers rank items in their choice sets by product features? If not, can we specify a descriptive model informed by empirical data that better reveals patterns of information search and the decision processes actually used by investors?

Yee, Dahan, Hauser and Orlin (2007) show that consumer choice is sometimes better described and predicted by lexicographic, or non-compensatory, decision tree models rather than linear models (conjoint analysis) which are standard in marketing research and economics. We use this insight as motivation to fit non-compensatory decision trees and heuristic tallying rules to portfolio choice data, comparing predictive accuracy with larger linear models that simultaneously apply weights to all investment features to predict investor behavior.[2] We also borrow the technique of process tracing from psychology to filter the data in greater detail in hopes of revealing more about investors' decision process.

[2] Linear models of investor behavior, such as those derived from expected utility maximization with a mean-variance expected utility function, assume that all features of each investment alternative are weighted and summed (and possibly transformed by a monotonic function). In contrast, non-compensatory models have a fixed hierarchy of investment features. One investment feature (e.g., high risk) is enough to entirely discard investments from the investor's consideration set, without the possibility of compensation (e.g., standard models typically assume that, no matter how high the risk, this negative can be compensated with sufficiently large positives such as high expected return). Just as Yee, Dahan, Hauser and Olin (2007) found that consumers can effectively manage large choice sets by quickly shrinking the items in consideration using a single product feature, we consider models which allow data to reveal this kind of fast reduction of choice sets by lexicographically ruling out investments based on one undesirable feature.

One motivation for the approach taken in this study is that improved veridical descriptions of decision-making process which go beyond the as-if methodological goal of predicting decision *outcomes* would be a substantive advance for descriptive consumer science. Perhaps just as compelling, especially in light of the institutional failures surrounding recent financial market crises whose consequences continue to play out as we write, these descriptions of decision-making process should facilitate the design of improved institutions that are better matched to the decision-making processes commonly found among bank consumers. Because firms may have incentives to offer smaller ranges of choices to retail (versus online) customers (Lynch and Ariley, 2000), this study's rich data, drawing on retail branch bank customer decisions and information collected from those bank employees who dispense investment advice, provides a unique view into retail banking as experienced by consumers.

We would contend that mistaken models of the investor's mind underlie a wide variety of real-world problems, from faulty risk models used by ratings agencies on mortgage backed securities to misguided marketing campaigns that fail to successfully launch new investment products to consumers. Being well informed about decision process—not merely decision outcomes—matters greatly. We try to demonstrate a technique, implemented in the case study presented here of Italian bank customers (who are typical in that they possess no expert credentials or specialized academic training in finance), to reveal how average investors search for information and map that information into portfolio decisions.

A growing literature in economics and psychology shows that decision makers often do not incorporate all available information into their decisions, even when 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). Loewenstein (2006) illustrates the more general point about mismatch between models and underlying behavioral processes that, paradoxically, can get worse when modelers try to introduce new aspects of realism to their models. He writes (Loewenstein, 2006, p. 705) about the particular case of Stigler introducing information costs to an otherwise neoclassical model of choice in which information typically comes for free: "the new assumptions are as patently unrealistic as the original assumption… ."

If information search is limited rather than exhaustive, how are pieces of information that are considered mapped into actual decisions? Gigerenzer, Todd and the ABC Group (1999) put forward a positive theory regarding simple and information-frugal decision rules that have attractive 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 can make accurate predictions in a variety of domains (Gigerenzer, Todd and the ABC Group, 1999). Inferences from information-frugal models classify and infer with only limited knowledge and computational effort, using a strict subset of the available information. Some of the heuristics we consider, as in Yee, Dahan, Hauser and Olin (2007), make no trade-offs between investment features. Instead, they process sequentially and lexicographically, one investment feature at a time. Under certain circumstances, such heuristics can be as accurate as models based on standard definitions of "rational choice" in economics (e.g., weighted linear models and Bayesian approaches that take all features and their entire correlation structure into account, assuming that decision makers somehow have access to all the relevant parameters—see Martignon & Laskey, 1999, and Gigerenzer, Todd and the ABC Group, 1999).

The paper proceeds as follows. Section 2 describes the experimental design, the sample from which the data are collected, and details of the four experimental tasks tailored to match Italian bank customers' investment environment. Section 3 presents the empirical results, describing a number of patterns in subjects' search for information and fitted prediction models for consumers' ranking of investments and portfolio choice. Section 3 fits Markov transition matrices for information search, response-time data, process tracing, non-compensatory heuristic models of investment choice, and standard linear models that attempt prediction of consumer's choice of investment products actually offered to them by banks. Finally, Section 4 concludes with a discussion and interpretation of the empirical results.

2. Method

The research project has been developed in two phases. We first collected interviews from 20 professional financial advisors and then from 80 bank customers of an Italian mutual bank. The interviews focused on interactions between advisors and their clients with active investment accounts. An important focus was the investors' experiences communicating with advisors about risk, and the information they valued (or not) when allocating investment funds to different investment products. Fifteen customers from the sample were also invited to participate in four investment choice tasks. We analyze this experimental data at the aggregate level and then go on to investigate within-subject treatment effects across the four investment tasks, which systematically introduce or remove one real-world aspect of advisor-client communication with each treatment/task.

2.1 The Instruments

We designed four hypothetical investment tasks to keep track of financial consumers' information search and actual decision making in the field. These financial consumers who are our subjects were not incentivized by receiving rewards from completing the task. We discuss the rationale for this approach below and our reasons for thinking that this design provides useful insights about inter-subject variability in terms of information use in investment decision-making.

Fifteen subjects participated in a sequence of four experimental computer-interactive tasks, implemented using an original computer interface. To make the tasks as close to real investment environment as possible, the information setup of the investment tasks relied partially on the investment features mentioned by financial advisors in surveys and interviews we conducted as features that they commonly present to the customers they advise. We used information collected from the financial advisors to calibrate investment features presented in the experiments (e.g., typical values of expected return and risk, quantified as standard deviation of annual return) that realistically match the features of investment options commonly presented to customers.

2.2 Experimental Design

We investigated subjects' decision strategies according to three main characteristics: (i) the amount of information subjects use when making financial decisions, (ii) the type of information (features or cues) they consider before choosing their investments, and (iii) the decision process they follow when searching for information.

2.2.1 Subjects

Subjects were customers of an Italian mutual bank. A mutual bank is a nonprofit financial institution whose aim is to support the economic well-being of the people living in a specific region. We selected this type of bank because its financial advisors do not earn commissions and face little or no pressure to sell investment products. Therefore, it is commonly thought that their incentives are much more closely aligned with investors' interests than is common in other institutional arrangements. Therefore, the mutual bank offers what is perhaps the most conducive environment one might hope for in terms of clear and effective communication between financial advisors and their clients. This environment, we hope, provides a best-case scenario for learning about the information consumers want and use, and which elements of financial communication work and do not work from the consumer perspective, without the confounding influence of possibly adversarial incentives.

Subjects were randomly selected from the bank database, which contains all active customers. When randomly drawing customer/subjects, we imposed a minimum deposit threshold of 40,000 Euros to help ensure that investment decisions we observed were relatively high stakes and not merely peanuts. Therefore, the sample was drawn randomly among the database of bank customers with 40,000 Euros or more. We refer to "subjects" and "bank customers" interchangeably.

2.2.2 Overview of Experiment and Four Investment Tasks

Computer-administered investment tasks were carried out at different bank branches in and around Trento, Italy. An experimenter read the instructions to each participant out loud. Each session was run with a single bank customer. The average session lasted 60 minutes for the experimental tasks and 15 minutes for a post-experiment interview. Investors were not remunerated. They voluntarily participated in the experiments. Most subjects showed great enthusiasm and gave statements indicating that they viewed their participation as a contribution improving the quality of service at their mutual bank.

Investment tasks were conducted on a touch-screen-based interface programmed in Java. We designed this computer interface specifically to facilitate interaction of a wide range of consumers with different demographic characteristics. In particular, we wanted an experimental interface with large, easy-to-read, icons that would work well for older investors. Each subject was placed in front of the touch-screen and trained extensively in a one-on-one teaching format on how to operate the computer and the details of all tasks. All experimental decisions were recorded automatically by a PC running a Java Virtual Machine, and interviews were recorded as digital audio files that were later transcribed.

Each experiment was composed of four different tasks. In Task 1, subjects were asked to choose between two investments based on expected returns, standard deviations of returns, investment time horizons, and management fees. In Task 2, the number of investments was expanded to six, and we observed the information subjects chose to look up, their sequence of information search, and their most preferred investment portfolio. In Task 3, instead of choosing from a stated menu of investment choices, subjects were asked to design their preferred investment portfolio based on investment labels indicating general categories of asset classes (e.g., equity, treasury bills, other fixed income, real estate). Labels such as "treasury bills" are typically used by investment advisors as benchmarks when discussing risk, serving as the closest real-world analog to the risk-free returns that appear in many theoretical models of portfolio choice such as CAPM. We used these labels to better understand their effects, given that they are commonly used by investment advisors in consultations with bank customers. In Task 4, individuals were asked to repeat the asset allocation decision from Task 3, this time using only individual investment features but no labels for asset categories. This design allows for within-person analysis of the effect of labels on portfolio choice. In all experiments, we randomized the order in which investment features were presented. The goal was to not influence information search by presenting information in a fixed order.

2.2.3 Task 1: Binary Investment Choice

When asked to choose between two investments, subjects were invited to search for information presented in a 2x6 matrix. Each of the two rows represented one of the two investment alternatives. Each of the six columns contained information about one of the following investment features—risk, time horizon, management cost, liquidity, a binary indicator showing whether a capital loss is possible or not, and expected return (labeled "coupon" in the heading of the final column of Figure 1). The investment features, about which each column provided information, were labeled, but the information in those information matrix cells was hidden, allowing us to measure which information cells were "looked up" by subjects and in which order.[3] There are no monetary costs of information. The experimental design does, however, impose a limit of six as the maximum number of investment features that can be looked up (out of 12 that could be looked up in the absence of any limit) to generate meaningful opportunity costs associated with each investment feature that is looked up.[4] Risk was presented and implicitly considered as the probability of a capital loss.

Each task began by presenting a black matrix on the screen, initially hiding all information about investment features. Each subject was asked to explore those features that they considered helpful for identifying their preferred investment. Figure 1 shows a typical screen for Task 1 after six pieces of information about four investment features have been looked up. A decision is about to be made.

[3] These terms are translations of the Italian terms used by the banks' financial advisors. Their imprecision would likely make any student of financial economics blush or, better still, demand more specificity. What, for example, does "risk" mean? Nevertheless, our goal was to calibrate all experimental details to the actual investment environment faced by the consumers who served as subjects. Thus, the terms are regrettably vague: *rischio, durata, costo, liquidabilita, perdita in conto capitale* and *interessi.*

[4] The maximum number of investment features that can be looked up of six actually comes from information collected from investment advisors. They told us that, due to time constraints and limitations of customers' understanding of and capacity to absorb investment information, they commonly consider the upper limit to be six pieces of information.

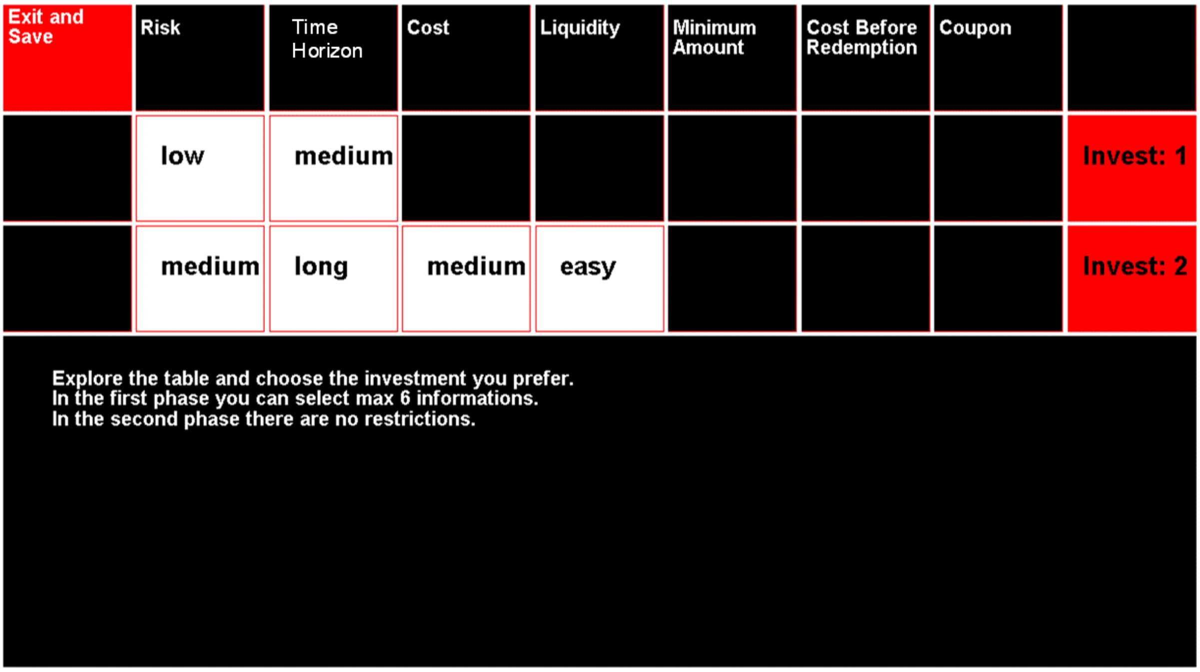

{Caption → Figure 1: Task 1, Choosing Between Two Investments}

2.2.4 Task 2: Extended Information Search

Task 2 forces the same subjects to confront an investment environment with a larger choice set and considerably more information that could be looked up. The investor's choice set now contains six alternatives. These were chosen based on different data collected about representative menus facing bank customers at the banks we studied. All six investment alternatives are those that are typically available in banks: bank accounts, bonds issued by the mutual bank, bonds issued by the government (i.e., the Italian equivalent of US treasury bills), bonds issued by insurance companies, balanced mutual funds (with a roughly 50-50 portfolio in corporate bonds and blue chip stock equities) and stocks (a value-weighted index of Italy's largest 40 publically traded companies trading in markets with euro dominated shares).

Figure 2 depicts this larger information field, consisting of six investment alternatives. Figure 2 also shows seven columns presenting feature-specific information. The newly added seventh investment feature is the "minimum required" feature, a binary indicator showing investments that require a minimum cash investment. The order in which investment features was presented to subjects was randomized in each trial. Therefore, the columns of Figure 1 were different for each task and each subject.

The screenshot in Figure 2 shows a subject who has already looked up 15 out of the total of 42 pieces of information. Subjects were asked to look up only the information they considered necessary for choosing their most preferred investment portfolio. Once again, the task began by presenting an information matrix in which all investment features were initially blacked out. No restrictions on the number of investment features that could be looked up were imposed.

Customers performed one information search per trial. Following each information search, subjects were asked to write down a set of portfolio weights on the six investment categories (summing to 1) representing their most preferred investment portfolio.

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	

{Caption → Figure 2: Task 2, Extended Information Search}

2.2.5 Task 3: Portfolio Choice with Asset Class Labels

Unlike Tasks 1 and 2, subjects in Task 3 were provided with the full information matrix revealing all investment features from the beginning. Their task was to choose an investment portfolio by writing down portfolio weights (i.e., a list of six numbers summing to 100 that represent the percentages of the investor's wealth allocated to the six investment categories in the rows of Figure 2). Financial advisors told us that none of their customers considered or invested in portfolios that take short positions. Therefore, the experimental user interface did not allow negative portfolio weights.

Figure 3 illustrates the screen that elicits subjects' portfolio weights. The first column contains the name (i.e., label) of each investment. The white boxes show the investment features. The last column contains the investors' portfolio weights (sometimes referred to as allocation decisions). Buttons on the boxes where subjects entered portfolio weights adjusted in 5 percentage-point increments. The heading of the last column dynamically displays the percentages of investor wealth allocated to each row, which corresponds to an investment or investment class serving as one element in the investor's portfolio. The screen shot shown in Figure 3 depicts an investor's portfolio choice in mid process. Only 70% of the investor's wealth has been allocated so far. The investor then continues adjusting the portfolio weights until a full 100% has been allocated across the six rows show in the final column.

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%

{Caption→ Figure 3: Task 3, Investment Portfolio Selection}

2.2.6 Task 4, Zero Information Portfolio Choice

Task 4 was identical to Task 3 except for the first column. The first column displayed asset class labels in Task 3. In Task 4, those labels are hidden, as shown in Figure 4. No other changes were introduced apart from a random reshuffle in columns to re-order investment features. 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 tasks and then summarized them into summary schemes.

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%

{Caption→ Figure 4: Task 4, “Blind” Portfolio Choice}

3. Results

The results are presented in two sections, one concerning the descriptive analysis of subjects’ information search behavior and the other concerning the modeling of decision strategies.

3.1 Information Search

In examining the approach followed by customers in searching for 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 Task 1 (pair-wise investments comparison) and in Task 2 (extended information search with an expanded choice set and list of investment features). In Task 1, 86% of customers looked at all six pieces of information. In Task 2, customers considered, on average, less than half of the available information (45%), revealing a preference for partial rather than full information to act upon.

3.1.1 Quantity of Lookups and Types of Information Searched

Table 1 shows the look-up rate per search for each investment feature. It shows the percentage of all searches, each consisting of selection of six pieces of information about two investments that could have been looked up, in which a feature was looked up at least once. A subject could look up all six features for one investment and no features for the other. Or a subject could look up three features for both investments. Thus, the sum of all look-up percentages in Table 1, must fall in the range of 300 to 600, representing the fewest versus maximum variety of feature look-ups that was possible. The realized sum was 318, which indicates that most subjects expended their information budget of six look-ups by looking up three features for both investments rather than four, five or six look-ups of features for just one of the pair of investments.

As shown in Table 1, the three main investment features that investors focused on were risk, time horizon and management costs. Subjects' demand for information about the risk feature dominated all other investment features by a sizable margin. In 65 trials, subjects looked up information about risk on more than 75% of searches, whereas the lookup rate for all other investment features was less than 50%.

Feature	Look-up Rate
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

{Caption→ Table 1: Rates of Look-up Among All Search Trials, by Investment Feature}

3.1.2 Temporal Analysis of Information Search

In Task 1, subjects sequentially searched for at most 6 different pieces of information dealing with the investment features in 65 trials. Table 2 presents the empirical distribution of information look-ups at each of the six search steps, denoted $t_1,\ldots, t_6$. This sequential analysis yields results consistent with Table 1 while providing more details about the sequence of search. Risk, time horizon and management costs tend to be looked up the earliest.

Feature	Time t_1	t_2	t_3	t_4	t_5	t_6
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**

{Caption→ Table 2: Percentage of Total Look-Ups Allocated to Various Investment Features in Task 1}

Table 2 reveals that at the first opportunity to look up an investment feature (t_1), risk was looked up in 89.2% of all trials. At t_2, risk was looked up in 41.5% of all cases, and time horizon 40%. At t_3, cost was looked up in 35.4% of all cases and time horizon 26.2%.

These per-period empirical distributions of information search reveal that, at t_1 and t_2, risk and time horizon are the most frequently looked up investment features. From t_4 onwards, the empirical distribution of look-ups is more evenly distributed over features. It is especially interesting that expected return (i.e., coupon) rather than appearing as the highest priority investment feature, is the most frequently looked up feature only in the very last time period. During the first three time periods, the average subject takes the following search path: risk, time horizon, and then cost.

Table 3 presents an estimated Markov transition matrix comprised of the empirical probabilities of moving from one investment feature to another in the six-step information search process. At the beginning (start position), the feature most likely to be looked up first is risk (89% of all first lookups). The feature most frequently looked up following risk is either time horizon (35% of look-ups following risk) or risk once again (23% of look-ups following risk). The most frequent look-up after looking up time horizon is either cost (46% of look-ups after looking up time horizon) or time horizon (16%).

Transition Probabilities From Feature to Feature	Start	Risk	Time horizon	Cost	Liquidity	Cost Before Redemption	Coupon	End
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
Redeem Early Penalty	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

{Caption→ Table 3: Transition Probabilities Between Investment Features in Individual Information Search in Task 1}

3.1.3 Process Tracing of Information Search

Following the approach to describing patterns in information search proposed by Payne, Bettman, and Johnson (2004), we looked at two types of search paths: feature-wise and investment-wise. A feature-wise path corresponds to an investor focusing on just one feature and looking up that same feature across investments. An investment-wise path corresponds to an investor looking up features belonging to just one investment at a time.

Data collected in Task 2 show that 8 out of the 14 bank customers (57%) adopted an investment-wise approach calculated according to Payne's measure. Subjects tended to focus their attention on information pertaining to a single investment at a time. The simultaneous protocol analysis revealed that most customers began searching for information about the investments they already owned in real life (e.g., government bonds [treasury bills] for subject 9, mutual bank bonds for subject 7, etc.). The other customers searched the available information by adopting mixed strategies; some of them exhaustively searched for information dealing with risk by adopting a cue-wise approach, while others gathered information across all investments without revealing a predominant approach.

3.1.4 Overlapping Information Index and Order Preservation Index

We investigated subjects' information search to answer the following questions: Did customers look at identical information for both investments? Did they search for investment features by following a stable ordering of investment features?

We adopted a within-subject approach by introducing 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 (Task 2), customers did not check the same features for A and B respectively before choosing. Such measures of systematic search, as the two indices we introduced, are not commonly mentioned in the experimental psychology literature, although they are certainly relevant for describing information usage and decision processes which provide an indication of whether neoclassical or heuristic models better describe how customers choose investments.

The two specific measures used here are: The Overlapping Information Index (OII) and the Order Preservation Index (OPI). We defined the OII as the percentage of identical features looked up for both investments across customers. OPI indicates the percentage of overlapping features searched in identical order. Let us look at two 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. As an alternative example, another subject might search in the following order: risk, cost, and liquidity for investment 1, and risk, liquidity, and cost, for investment 2. In 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 66.6%, because two out of the three overlapping features are looked up in the same order.

Following from above, we classified customers according to OII. We chose a threshold of 50% to split the sample into low- and high-OII groups:

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

OPI and OII are highly correlated, which means that customers belonging to the HOI group preserve their search order across different searches and focus their attention on a smaller set of financial features. This fact suggests that the selection of the considered information set is strongly linked to individual search processes 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 followed the same sequential order across cues.

Group	Subject Id.	Overlapping Information Index (in %)	Mean Number of Cue Look-ups	Order Preservation Index (in %)
HOI	2	100	3	100
HOI	7	100	3	100
HOI	3	91.6	3.2	91.6
HOI	8	75.	3.7	75
HOI	1	73.3	3.8	66.6
HOI	5	58.3	4.2	50
HOI	10	50.	4.2	41.6
HOI	15	50.	4.5	50
LOI	13	46.6	4.6	40
LOI	9	41.6	4.7	41.6

LOI	6	33.3	5.0	6.6
LOI	14	33.3	5.0	26.6
LOI	12	20.	5.4	13.3
LOI	11	13.3	4.8	6.6
LOI	4	8.3	5.7	0
HOI Group Mean		74.7	3.7	71.8
LOI Group Mean		28.1	5.0	19.2
Sample Mean		53.	4.3	47.3

{Caption→ Table 4: Overlapping Information Index and Order Preservation Index Calculated for Every Subject in Task 1}

3.2.1 Modeling Consumer Rankings of Investments

(1) Heuristic Model

This section is devoted to the description of a simple heuristic that describes quite well how bank customers made their decisions. Let us begin by establishing certain notations and introducing some important concepts. We define the *cue profile* of an investment as a binary vector of 1s and 0s according to whether cue values are “positive” or not, ordered according to the sequence: risk, time horizon, liquidity, costs (intermediary fees), other costs, and coupon (which is a synonym here for expected return). Based on the fast and frugal heuristic model (Gigerenzer & Goldstein, 1996), 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 for instance, risk 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.

Based on interviews with financial advisors combined with experimental data on customers’ information search, we identified a heuristic that is a hybrid between lexicographic and tallying rules. This hybrid heuristic first examines the one cue that was searched most, namely risk, which lexicographically over-rules all other investment features. For a pair of investments considered to have similar risk, the heuristic ranks the investments by means of a tallying rule depicted in Figure 5.

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 heuristic predict the choices observed in Task 1? Predictions for each subject are presented in Table 5. 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).

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%

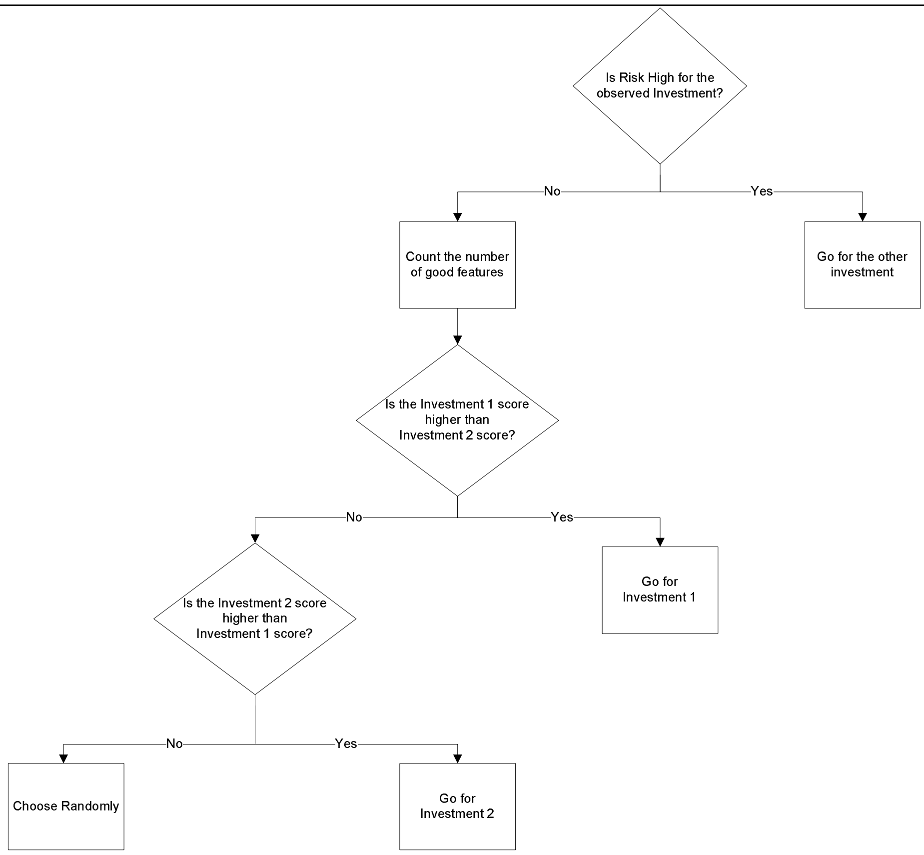

{Caption→Table 5 (left): Decision Heuristic Predictions}

{Caption→ Figure 5 (right): Fast and Frugal Tree for Investment Decisions}

The tree in Figure 5 predicts about 78% of the observed investment decisions in Task 2. The model predicts 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. Beyond this lexicographic step that avoids high-risk investments, the model describes investors as simply counting the number of their favorite features on which one investment dominates the other.

(2) Logit Model Prediction

The multinomial logit model linearly integrates effects of all investment features by taking a weighted average and then transforming (non-linearly) to estimated probabilities. We fit the logit model as a representative of standard linear prediction models based on rational choice (e.g., Train, 2003) as a benchmark against which to compare the predictive accuracy of the heuristic model.

The logit model is widely known and used in countless applications. It assumes that individuals can assign random utility scores to different objects according to the equation $u_i = v_i + \varepsilon_i$, composed of an observable component of utility v_i and an unobservable (random) utility component ε_i. The assumed decision mechanism is that the individual chooses the object with the highest utility. Hence, under the assumption that ε_i is identically and independently distributed, drawn from the Gumbel distribution, the probability of choosing an object is specified by the following functional form: $p_i = \exp(v_i) / \sum_j \exp(v_j)$, where *j* indexes all the objects in the choice set (which, in its present application, only includes two objects), and *i* represents one object of special interest to the modeler in the choice set. As is common practice, we also assume that observable utility is the summation of weighted cue utilities: that is, $v_i = \sum_k \beta_k x_{ik}$. The symbol x_{ik} represents the *k*th cue of object *i*, and β_k is the weight parameter that we external observers seek to estimate.

We discretized all cue values in order to normalize the utility value of investment features that were not looked up, even for free, to zero. For example, because the risk feature takes on three levels (low, medium, high), the utility specification coded this information as three dichotomous variables are used, r_l, r_m, and r_h. When the risk feature is not looked up, all these variables take on the value zero. When the risk feature is looked up and its value is revealed to be low, then the three risk variables are coded as $r_l = 1$, $r_m = 0$, and $r_h = 0$. Similarly for r_m or r_h. When risk information is looked up, only one of the three risk variables turns on to indicate a value of 1. As a consequence of this flexible coding that allows for any pattern of marginal utility effects conditional on low, medium, high risk, or no risk information at all, there are three parameters for risk, β_{rl}, β_{rm}, and β_{rh}.

In order to avoid as much as possible the problem of non-representative sample due to small sample size, we used a cross validation method. We divided the dataset by randomly drawing a 75% of the sample into the sub-sample for data fitting and a 25% sub-sample for validation. This procedure was repeated 100 times. Each time, the model was estimated on the fitting sample and validated on the validation sample.

Table 6 presents average realized values of the estimated coefficients in the linear utility model and counts on the number (out of these 100 prediction trials) in which a particular estimated coefficient was statistically significant. The results provide some evidence about which kinds of information are more important than others, at least in linear prediction. In Table 6, the four most robust cues seem to be high risk, high liquidity, unavailability of cost before redemption, and presence of regular coupon payments. The average parameter values appear to jibe with one's intuition about the direction in which various kinds of information would influence the investor's willingness to invest. High risk seems to have the largest magnitude impact on utility, which is negative. This is also, of course, consistent with the heuristic model introduced above. The positive sign on the other coefficients suggest that, all else equal, investments with high liquidity, no cost before redemption, and regular coupon payments are preferred.

The average percentage of correct hits for the fitting data is 84.5%; it is 80.4% for the validation data. Compared with the heuristic model, the number is 71.5% for the fitting data and 71.7% for the validation data. This suggests the full linear model could have over-fitted the data, pointing to another advantage of simpler heuristic models in out-of-sample prediction that has been commented on extensively elsewhere.

Parameter	Number of significant in 100 trials	Average coefficient value in the linear utility model
β_{rl} : low risk	20	1.3
β_{rm} : medium risk	8	2.7
β_{rh} : high risk	72	-3.3
β_{ds} : short time horizon	16	2.7
β_{dm} : medium time horizon	4	2.9
β_{dl} : long time horizon	0	n/a
β_{cl} : low cost	28	1.8
β_{cm} : medium cost	0	n/a
β_{ch} : high cost	0	n/a
β_{ll} : low liquidity	0	n/a
β_{lm} : medium liquidity	0	n/a
β_{lh} : high liquidity	100	2.1
β_{bu} : cost before redemption unavailable	48	1.7
β_{ba} : cost before redemption available	0	n/a
β_{iu} : coupon unavailable	4	-2.2
β_{ia} : coupon available	54	2.4

Mean hit rate for fitting: 84.5%
Mean hit rate in out-of-sample prediction: 80.4%

{Caption→ Table 6: Estimation results of the rational model}

3.2.2 Role of Recognition in Portfolio Choice

We now present the role of the investment labels on subjects' choice of portfolio weights. In Task 4, all 15 subjects performed one task dealing with what we called "blind categorization". We collected data on subjects' performance in reproducing the same investments allocation task they had already performed in Task 3; the only difference introduced in Task 4 is that we did not provide them with the investment names or labels—just the technical features of the investments. Therefore, financial consumers were still exposed to the same 6 types of investments they encountered in Task 3; namely, stocks, mutual balanced funds, government bonds (treasury bills), savings accounts, mutual bank bonds and insurance 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 was that bank customers did not excel in the blind classification task. Given that they paid most attention to risk in the exploration phase, 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 subjects (60%) made significant inconsistencies, that is, they invested in much riskier portfolios than before and with asset allocations that deviated 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 inconsistency may have to do with some customers relying on the recognition heuristic in

Task 3, while this was not possible in Task 4. The results confirm the role of the recognition heuristic: when it is not applicable, customers perform differently[5].

Table 7 and Figure 6 reveal investors' portfolio choices: data in bold represent customers who made "mistakes" in the blind categorization task by choosing different investment mixes in different treatments.

Subject Number	Weight on Risky Assets Chosen With Labels	Weight on Risky Assets Chosen Without Labels	Within-Subject Percentage Change
1	**15**	**25**	**66.7**
2	**15**	**20**	**33.3**
3	**20**	**30**	**50**
4	**60**	**65**	**8.3**
5	40	40	0
6	0	0	0
7	0	0	0
8	40	40	0
9	**25**	**30**	**20**
10	**50**	**65**	**30**
11	**25**	**45**	**80**
12	**5**	**20**	**300**
13	5	5	0
14	0	0	0
15	**15**	**20**	**33.3**
Mean	21.0	27.0	41.4
Std Dev	19.0	21.4	6.6

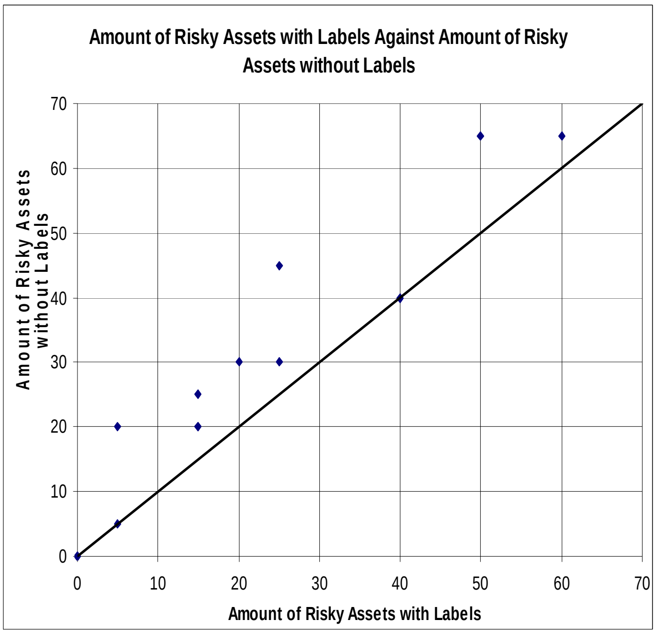

[5] By Recognition Heuristic, we mean a simple strategy that allows individuals to infer, for example, which of two objects has a higher value on some numerical criterion, based on the fact that one is recognized and the other is not (e.g., predicting which investment has a higher expected return based on whether one recognizes the name of one of the shares). 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).

{Caption→ Table 7: Comparison Between Portfolio Choice Across Treatments. Bold Figures identify subjects who unintentionally modified their asset allocations}

{Caption→ Figure 6: Comparison Between the Amount of Money Invested in Risky Assets When Labels are Provided (Task 3) versus Amount of Money Invested in Risky Assets When Labels are not Provided (Task 4)}

4. Summary Discussion and Interpretation

The aim of this research was to investigate how average investors make financial decisions. We focused on two components of investor behavior: information search and decision process. We interviewed 80 customers of an Italian mutual bank, and we invited 15 of them to participate in an interactive experiment with four investment tasks. We designed the experimental setup to reflect as realistically as possible the actual decisions that take place in transactions between financial advisors and their customers, based on information collected from financial advisors and bank customers in the earlier interview stage. To a large extent, we succeeded in presenting information to customer/subjects that matched the information they typically encounter when making real financial decisions.

In Task 1, we asked bank customers to choose between two investments while facing a strict information budget that limited information look-ups to a maximum of six out of 12 investment features that could have been looked up. In Task 2, subjects were free to search an extended financial information matrix without any limit on the maximum number of look-ups. We observed in Task 2 that subjects looked up less than half of the available investment features even though they were free to have more information and could have accessed all of it within a few additional seconds of search. More than half the subjects (57%) looked up only a small set of non-overlapping features (i.e., they did not look up the same feature from the two investments) even when they were free to access the entire information set and it was feasible to do so with very modest time costs.

Although the sample size is small (in part, due to our bank partner's indication that our data collection had continued long enough), the rich combination of qualitative interview data from 80 respondents, information search data, and investment decision data reveal new empirical insight into information search and investment decision processes among real investors. This evidence suggests that investors usually consider only a strict subset of available information about the set of investment alternatives from which they choose, even when the universe of investments and related information sets is relatively small and search costs minimal. In this sense, bank customers appear to follow a "less is more" principle that one finds as a key feature in a number of investment heuristics recently appearing in the behavioral finance literature (De Miguel et al., 2009; Gigerenzer, 2008).

Of particular interest was how infrequently bank customers undertook pairwise comparisons that we commonly assume take place in order to completely rank one's choice set (e.g., the expected return of investment A and the expected return of investment B). Pairwise comparison of product features is the essential behavioral assumption underlying the rational choice model in economics and the preponderance of consumer choice models used in marketing research. And yet our evidence directly contradicts this fundamental assumption about consumers' decision process.

We modeled binary investment choice with a lexicographic decision tree that predicts 78% of choices correctly. The decision tree orders risk first among various investment features. That means that risk is more important that all other features combined. The decision tree model then uses a tallying rule on the remaining investment features.

Tallying is cognitively less demanding than weighted additive models, because it depends only on counting up pros and cons rather than computing weighted sums. We compared the predictive performance of these models against a competing linear model based on full-information utility maximization. The performance of the rational model predicts the choice data slightly better than simpler models, with 81% versus 71% predictive accuracy. But in terms of model complexity, it costs many more parameters, which implies a large risk of overfitting, which would be true in any sample, but especially in a small sample like ours.

The second task focused on portfolio choice. Task 2 was designed to check whether providing subjects with all available information (i.e., all investment features) but no labels that name the different investment alternatives would affect investor behavior. We discovered that when labels are missing, individuals tend to select a riskier mix of investments than was selected when investment labels were available. This suggests that seemingly superficial differences in the naming of two financial products with identical mean and variance could play a significant role in investment decisions. The recognition heuristic theory helps explain how label- versus no-label tasks generated differences in observed choice behavior.

Financial products are presented to potential investors in the real world with a rapidly proliferating array of such labels. The names of things in the real-world investing environment very likely influence investor behavior, even though it should not according to standard portfolio choice theory. According to that theory, so long as the mean and standard deviation of all risky returns (including correlations, and possibly all higher moments for risk preference specifications more complicated than mean-variance preferences) are presented to the investor, his or her decisions should not be influenced by the order of presentation or names and labels applied to investment features. Behavior should be especially invariant to labeling changes that do not affect the first and second moments of random returns.

Using experimental and survey data collected from Italian bank customers, we tested a lexicographic fast and frugal tree against the multinomial logit model, which is a transformed but nevertheless linear model. The logit assumes that individuals, when choosing among objects, compute weighted additive utility scores for each object *en route* to choosing the object with the highest utility. We used a cross-validation method to test the out-of-sample predictive accuracy of both models.

The heuristic model suggests that subjects in our study rely on simple decision trees in which risk is, by far, the most important investment feature. Furthermore, normative assessment of the performance of real bank customers' decision processes (relative to the neoclassical benchmark) indicates 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 heuristic approach (similar to the normative assessments in Magni, 2009). Similar lexicographic decision-tree heuristics such as Gigerenzer and Goldstein's (1996) Take-The-Best Heuristic consider the features of investments sequentially in a ranking determined by some measure of pairwise correlation or univariate predictive accuracy (rather than considering all their inter-correlatedness with other regressors in the model as, for example, partial correlations do). This helps reduce the cognitive processing required to execute heuristics 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 tasks that subjects faced required them to, in some treatments, search freely for information, unlike most experimental economics tasks of financial decisions in which subjects are provided with a complete set of summary statistics (namely, expected values, variances, and covariances), which are required by standard models such as CAPM.

The subjects in this study could decide how much information they wanted to look at. Most chose to look up some quantity of information less than the maximum of all that was freely available. And a significant proportion explored only a very small subset of overlapping investment features (i.e., looking up the same piece of information more than once). Subjects exhibited remarkably similar information search behavior across trials, and their subjective reports about their search behavior coincided remarkably in that nearly all reported spending very little time in deliberation (and never calculation) in choosing portfolio weights. Moreover, subjects indicated that, although they were handling meaningful amounts of money, they nevertheless made investment decisions with relatively little cognitive effort in both lab experiments and in actual investment decision making. By far, the most important features of investments in the eyes of the subjects were risk, time horizon, and costs (brokerage fees), in that order. Information search was characterized by frugality and simplicity.

Given the crisis we are living through at the time of writing, and the accumulating evidence of mismatch between models of omniscient investors and the badly designed institutions that result from those models, we believe that marketing research can play an important and positive role in improving risk communication and deepening the relationship between firms selling investment advice and their customers. We hope that improvements in descriptive models of investors' use of information and decision processes will, in turn, improve the investment environment that average financial consumers face. Insofar as researchers discover better descriptive information about consumers' information needs and decisions processes, advancement toward consumers' satisfaction with the investment experience and with firms' and policy makers' ability to predict consumer behavior becomes more likely. Far from a zero-sum game, there appears to be room for substantial improvements, from the vantage points of all parties, in the processes by which consumer-level financial transactions take place.

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, 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 behavior. Organizational Behavior 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.

DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy. Review of Financial Studies, 22, 1915–1953.

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. (2008). Why heuristics work. Perspectives on Psychological Science, 3, 20-29.

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 Comparison: 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.

Loewenstein, G. (2006). The pleasures and pains of information. Science, 312, 704-706.

Lusardi, A. & Mitchell, O. S., (2006), Financial literacy and planning: implications for retirement wellbeing. Working Paper, The Wharton School.

Lynch, J. G. Jr., & Ariely, D. (2000), Wine online: Search costs and competition on price, quality, and distribution, Marketing Science, 19, 83-103.

Magni, C. A. (2009), Investment decisions, net present value and bounded rationality. Quantitative Finance, Volume 9, Issue 8 December 2009, pages 967-979

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 Behavior 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.

Train, K. E. (2003). Discrete Choice Methods with Simulation. Cambridge University Press, Cambridge, UK.

Tukey, J. (1962). The future of data analysis. Annals of Mathematical Statistics, 33, 1-67.

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.

Yee, M., Dahan, E., Hauser, J. R., & Orlin, J. B. (2007). Greedoid-based non-compensatory two-stage consideration-then-choice inference. Marketing Science, 26, 532–549.

Appendix A:

Overlapping Information Index

Give a set of information for investment 1, denoted A, and a set of information for investment 2, denoted B, the OII measure of overlapping information (introduced in the body of this paper) is simply the cardinality of the intersection of these two sets:

Eq. A.1 $$OII = |A \cap B| .$$

Order Preservation Index

The Order Preserving Index takes not of the order in which the elements of the information sets A and B were generated by subjects' search behavior. We write the elements of the set A as a vector that records this sequence: $a = [a_1, a_2, \ldots, a_n]$, where n denotes the number of items searched for or looked up by a particular subject. The index of subjects' identities is suppressed here for expositional clarity in describing the information indexes used in the paper's section on process tracing. Thus, a_i was looked up before a_j if and only if $i<j$. Similarly, a vectorized version of the set B that records the order in which information was looked up is denoted as $b = [b_1, b_2, \ldots, b_n]$.

The (n+1)x(m+1) matrix referred to as the so-called F matrix (see Payne, Bettman and Johnson, 2004) plays a key role in process tracing. The first column and first row of the matrix F are initialized as 0. In MATLAB command language, this is expressed as F(1,1:m+1) = 0 and F(1:n+1,1) = 0.

The rest of the matrix is computed according to a recurrence relation, the Needleman-Wunsch algorithm. The order of the steps follow from the sequence given by vectors a and b.

Eq. A.2 $$F(i,j) = max\begin{cases} F(i-1,j-1) + s(a_{i-1}, b_{j-1}) \\ F(i-1,j) \\ F(i,j-1) \end{cases},$$

where

Eq. A.3 $$s(a_i, b_j) = \begin{cases} 1, if\ a_i = b_j \\ 0, otherwise \end{cases}.$$

Finally, the OPI is computed as:

Eq. A.4 $$OPI = \left(100\left(\frac{F(n+1,m+1)}{\max(n,m)}\right)\right)\%.$$

The expression above is either normalized to a 0-100 percentage-point scale or interpreted as decimal representations of percentages that lie in the unit interval. An example that visually illustrates the recursion relation described above symbolically follows. This visual corresponds to the example discussed in Section 3.1.4.

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