

Search and Satisficing*

Andrew Caplin, Mark Dean, and Daniel Martin[†]

August 12, 2009

Abstract

Decision makers often have imperfect information. We develop a choice theoretic experiment to explore choice mistakes that result from incomplete search. Our choice process methodology generates data on how choices change with contemplation time, thereby illuminating the search process. We demonstrate that most subjects behave in line with a reservation-based model of sequential search, altering their reservation utilities in response to the size of the choice set and the complexity of the environment. These findings support Simon's model of satisficing behavior and suggest simple measures of contextual effects on the quality of decisions.

Key Words: Revealed preference, search, incomplete information, bounded rationality, stochastic choice, decision time

1 Introduction

When faced with large or complicated choice sets, it is unsurprising that people make significant mistakes, in the sense of failing to choose the best possible alternative. Understanding the nature and prevalence of such mistakes is an important theoretical and practical challenge. In practical terms, policy makers are looking to develop decision making protocols and rules that reduce consumer confusion. In parallel, economic theorists have begun to model the behavior resulting

*We thank Isabelle Brocas, Juan Carrillo, Vince Crawford, Stefano DellaVigna, Drew Fudenberg, Daniel Gottlieb, David Laibson, Alessandro Lizzeri, Ulrike Malmendier, Rosemarie Nagel, Efe Ok, Antonio Rangel, Ariel Rubinstein, and Yuval Salant for insightful suggestions.

[†]Center for Experimental Social Science and Department of Economics, New York University.

from choices being made from subjective “consideration sets” that are strictly smaller than the objectively available set of choices.¹

A key question is whether the fact that people make mistakes when they have incomplete information can be reconciled with the concept of revealed preference.² We introduce a novel choice theoretic experiment for purposes of reconciliation and use it to show that many apparent mistakes can indeed be rationalized by a model that incorporates search into the choice procedure. While subjects regularly violate standard rationality conditions, their behavior is well described by a simple model of sequential search with a reservation stopping rule. This is the satisficing model of boundedly rational behavior proposed by Simon [1955]. Moreover, both estimated reservation values and the order of search respond systematically to changes in the choice environment. In combination, these factors have strong explanatory power in identifying choice environments in which people make large mistakes, and those in which they do not. Thus, the process of information search provides a natural framework for understanding how environmental factors affect the quality of decisions that people make.

In order to explore the process of information search, we present a choice-based experiment that makes visible aspects of search that are not revealed in standard choice data. Our design elicits “choice process” data that records not only the final choices that subjects make, but also how choices change with contemplation time (see Campbell [1978] and Caplin and Dean [2009]).³ We obtain such data using an experimental design in which subjects’ choices are recorded at a random point in time unknown to them, incentivizing them to always report their currently preferred alternative. This represents a choice-based experiment constructed precisely to enrich our understanding of search behavior and imperfect information.

In order to pin down the effect of search on choice quality, we use choice objects which allow for an intuitive notion of mistakes. The objects of choice in our baseline experiments are simple mone-

¹See Manzi and Mariotti [2007] and Masatlioglu and Nakajima [2009] for examples of decision theoretic models with consideration sets. See also Eliaz and Spiegler [2008]. Rubinstein and Salant [2006] present a model of choice from lists, in which a decision maker searches through the available options in a particular order. Ok [2002] considers the case of a decision maker who is unable to compare all the available alternatives in the choice set.

²The appropriateness of categorizing particular decisions as mistakes is taken up by Bernheim and Rangel [2008], Gul and Pesendorfer [2008] and Koszegi and Rabin [2008].

³Compared to other novel data used to understand information search, such as those based on eye tracking or Mouselab (Payne, Bettman and Johnson [1993], Gabaix et al. [2006], Reutsaja et al. [2009]), choice process data is more closely tied to standard choice data and revealed preference methodology.

tary prizes, making preferences trivial and universal. However, this dollar value is not immediately clear to the subject, as it is expressed in the form of a sum, or sequence of addition and subtraction operations. The act of choice is non-trivial because mental effort is needed to understand the value of each prize on offer. Subjects regularly make choice mistakes in the sense of failing to select the object with the highest dollar value. We show that the size of these mistakes is affected both by the number of available alternatives and by the complexity of each alternative, as measured by the number of mathematical operations that make up each option.

We use choice process data to test whether models of information search can explain this pattern of mistakes. We first show that search behavior is well described by sequential “alternative-based” search: subjects behave as if they are searching through alternatives one by one, always selecting the best of the alternatives that they have come across. This contrasts with other boundedly rational models of search, such as those that are “attribute-based”, in which different attributes of the goods are examined in sequence. Caplin and Dean [2009] provide a general characterization of the implications of alternative based search (ABS) for choice process data, and in section 4 we show that the vast majority of data satisfy this characterization. In fact allowance for this simple form of search removes almost all evidence of mistakes for most subjects. While apparent violations of rationality in final choice are large and context dependent, appropriately measured improvements during the process of search are not.

More striking still is the apparent applicability of the simple satisficing model of Simon [1955] to describing choice process data. Most of our experimental subjects appear to engage in sequential search that stops once a satisfactory, or reservation, level of utility is achieved. Mistakes are found to be large in environments associated with low levels of reservation utility. We show that such behavior is optimal for a decision maker (DM) facing fixed per-alternative psychic search costs. The optimal reservation level is decreasing in the complexity of each object, but is unaffected by size of the choice set.

Given the applicability of the satisficing framework, we investigate experimentally how changes in the decision making environment impact mistakes by estimating the corresponding changes in reservation utility. We find that reservation levels do indeed decrease as the complexity of choice objects increase, in line with the optimal model. However, we also find that reservation levels increase with the size of the choice set, suggesting that subjects search relatively too hard in larger choice sets, as compared to optimal behavior.

In order to elicit choice process data we use an experimental design in which final choices may not in fact be implemented. In section 7 we explore the impact of this design feature on our results. To this end we use data from the pure choice experiments in which only final choices were implemented. As with the choice process data, these pure choice experiments allowed individuals to click to new options during the pre-decision period. The key difference is that these switches were payoff irrelevant in the pure choice case. Despite this difference, the results in section 8 show that the key features of the analysis hold when we examine these unincentivized changes in choice. It appears that the incentive structure of the choice process experiment has little impact on the nature of the search process.

In addition to providing information on reservation values, choice process data can also shed light on the order in which people search through the choice set. In the final section of the paper, we study search order in settings in which options vary both in the order on the screen and in their complexity. We identify some individuals whose search order is governed by screen position, and others whose search order is governed by complexity. We show that individual differences impact the mistakes that subjects make: those who search in screen order miss good objects at the base of the list, while those who search by complexity miss good objects if they are complex.

2 Measuring Mistakes

2.1 Experimental Design

In our first experiment (experiment 1), we use a standard choice task to identify choice environments in which people make mistakes, in the sense of failing to select the best possible objects. In order to make such mistakes obvious, we use choice objects that have a clear underlying value, but whose value takes effort to uncover. Each object is displayed as an arithmetic expression, a sequence of addition and subtraction operations, with the value of the object equal to the value of the sum in dollars.⁴ As we demonstrate below, choice among these objects produces evidence of significant and environmentally sensitive mistakes.

⁴Given that the subjects (NYU students) were unusually numerate and made negligible mistakes when purely numerical options were presented, we wrote out the arithmetic expressions in word form rather than in symbolic form.

Experiment 1 consisted of six treatments, differing in the complexity of choice object (3 or 7 addition and subtraction operations for each object) and the total number of objects (10, 20 or 40 alternatives) in the choice set. Figure 1 shows a 10 option choice set with objects of complexity 3.

FIGURE 1 ABOUT HERE

Each round began with the topmost option on the screen selected, which had a value of \$0, and so was worse than any other option. While only the final choice was recorded, subjects could select whichever option they wanted at any time by clicking on the radio button next to that option. The alternative that the subject currently selected would then be displayed at the top of the screen. Once they had finalized their selection, they could proceed by clicking on the submit button at the bottom of the screen. The changes that were made over the entire pre-decision period were recorded and their properties are explored in section 8. However it was only the final choices that were payoff relevant. There was no constraint on decision time.

The value of each alternative was drawn from an exponential distribution with $\lambda = 0.25$, truncated at \$35 (a graph of the distribution was shown in the experimental instructions - see appendix A).⁵ Once the value of each object was determined, the operations used to construct the object were drawn at random.

Subjects for experiment 1 took part in a single experimental session consisting of 2 practice rounds and between 27 and 36 regular rounds, drawn from all 6 treatments. At the end of the session, two regular rounds were drawn at random, and the subject received the value of the selected object in each round, in addition to a \$10 show up fee. Each session took about an hour, for which subjects earned an average \$32. In total we observed 22 subjects making 657 choices.

2.2 Mistakes

Table 1 presents information on the extent to which mistakes were made in each treatment. We report three measures of error. The first row reports “failure rate” - the proportion of rounds in which the subject did not choose the best option (i.e. the option with the highest dollar value). The second row reports average absolute loss - the difference in dollar value between the chosen

⁵For each of the three choice set sizes we generated 12 sets of values, which were used to generate the choice objects at both the low and the high complexity levels.

item and the highest value item in the choice set. The third row reports average percentage loss - the absolute loss expressed as a percentage of the highest value in the choice set.

TABLE 1 ABOUT HERE

Our experimental design creates an environment in which subjects make suboptimal choices. Averaging across all treatments, subjects fail to select the best option 38% of the time. These failures of rationality are also significant in terms of dollar amounts. On average, subjects leave \$3.12, or 17% of the available money, on the table in each round.⁶

The degree to which subjects make mistakes varies significantly and systematically across treatments. All measures reported in table 1 increase both with the size and the complexity of the choice set. Failure rates vary from 7% for the size 10, low complexity (3 operations) treatment to 65% for size 40, high complexity (7 operations) treatment. Average losses range from \$0.41 (3.44%) in the size 10, low complexity treatment to \$7.12 (33.25%) in the size 40, high complexity treatment. Regression analysis shows that the difference in losses between treatments is significant.⁷

There is also some evidence that the effect of complexity is higher in larger choice sets - the difference in loss between low and high complexity objects in size 10 choice sets is \$1.29 (10.2%) and not significant at the 10% level. For size 40 choice sets, the difference is \$4.83 (22.8%) and significant at the 1% level.⁸

⁶There is no evidence for any effect of learning or fatigue on mistakes. The order in which choice rounds were presented was reversed for half the subjects, and the order of presentation did not have a significant effect on performance. This may in part be because our experimental design is structured to remove learning effects. The decision making context, including the distribution of prizes, is known to the decision maker at the start of each experimental round.

⁷Absolute differences in value were regressed on dummies for choice set size, complexity and interactions, with standard errors calculated controlling for clustering at the subject level. Losses were significantly higher at the 1% level for complexity 7 vs. complexity 3 for size 20 and 40 choice sets, though not for size 10 choice sets. Losses were also significantly higher at the 1% level for size 40 vs. size 10 choice sets at both levels of complexity.

⁸While not the primary subject of study in the current paper, there are significant individual differences in mistakes. Estimates obtained from a regression of absolute loss on individual specific dummies, controlling for treatment effects, indicate that the 25th percentile subject does on average \$1.10 better than the median subject, while the 75th percentile subject does \$1.23 worse, averaging across all rounds.

3 The Choice Process

3.1 Ideal Data

While the mistakes identified in section 2 are unsurprising, standard choice theory has little to say about them - either the process by which such mistakes come about, or the relationship between factors in the choice environment and the likelihood and the significance of mistakes.

In order to explore these issues, we introduce choice process data, which is designed to shed light on search-based causes of mistakes. Rather than recording only the final alternative that is chosen by the DM, choice process data tracks how the choices that people make evolve with contemplation time. As such, choice process data come in the form of sequences of observed choices. For each non-empty set of alternatives A , choice process data specify not just the final choice $C(A) \subset A$, but rather a sequence of choices, representing the DM's choices after considering the problem for different discrete lengths of time.

We introduce now the formal version of the choice process data set from Caplin and Dean [2009].⁹ Let X be a nonempty finite set of elements representing possible alternatives, with \mathcal{X} denoting non-empty subsets of X . Let \mathcal{Z} be the set of all infinite sequences from \mathcal{X} with generic element $Z = \{Z_t\}_1^\infty$ with $Z_t \in \mathcal{X}$ all $t \geq 1$. For $A \in \mathcal{X}$, define $Z \in \mathcal{Z}_A \subset \mathcal{Z}$ iff $Z_t \in A$ all $t \geq 1$.

Definition 1 *A (deterministic) choice process* (X, C) *comprises a finite set* X *and a function,* $C : \mathcal{X} \rightarrow \mathcal{Z}$ *such that* $C(A) \in \mathcal{Z}_A \forall A \in \mathcal{X}$ *and* $|Z_t| = 1 \forall t$.

Given $A \in \mathcal{X}$, choice process data assign not just final choices, but a sequence of such choices, representing the DM's choices after considering the problem for different lengths of time. We let $C_A(t)$ refer to the object chosen after contemplating A for t periods.

Choice process data represent a relatively small departure from standard choice data, in the sense that all observations represent choices, albeit indexed by time. We therefore see this approach as complementary to other attempts to use novel data to understand information search, such as those based on eye tracking or Mouselab (Payne, Bettman and Johnson [1993], Gabaix et al. [2006], Reutsaja et al. [2009]). These approaches make aspects of the search process observable, yet do

⁹Caplin and Dean [2009] consider the generalized case with set-valued choice functions.

not connect these intermediate acts of search with their implications for choice. On the other hand, choice process data misses out on potentially relevant cues to search behavior, but captures the moment at which search changes a DM’s assessment of the best option thus far encountered.

3.2 Experimental Design

For each set of alternatives presented to an experimental subject, our aim is to generate a time series of observations that records their preferred alternative from the choice set at each moment in time. Our design has two key features. First, subjects were allowed to select any alternative in the choice set at any time, changing their selected alternative whenever they wished. Second, actualized choice was recorded at a random point in time unknown to the experimental subject. At the end of each choice round, a random time was generated, and whatever the subject had selected at that time was recorded as their choice. This incentivized subjects to always have selected their current best option in the choice set. We therefore interpret the sequence of selections as choice process data.¹⁰

Appendix A reproduces the experimental instructions. As in the standard choice experiment, each round began with the topmost and worst option of \$0 selected, subjects could at any time select any of the alternatives on the screen either by clicking on the alternative itself or the radio button next to it, with the currently selected object being displayed at the top of the screen. Unlike in the standard choice experiment, there was a time constraint, with subjects having up to 120 seconds to complete the choice task (though, as we shall see below, this time constraint is rarely binding).¹¹ Subjects were instructed that at the end of the round, a random time would be picked from distribution between 1 and 120 seconds according to a truncated beta distribution with parameters $\alpha = 2$ and $\beta = 5$, and the selected alternative at this time would be recorded as the choice for that round.¹² A subject who finished in less than 120 seconds could press a submit

¹⁰In support of this interpretation, 58 of 76 subjects in a post-experiment survey responded directly that they always had their most preferred option selected, while others gave more indirect responses that suggest similar behavior (e.g. having undertaken a re-calculation before selecting a seemingly superior alternative).

¹¹In experiment 1, which had no time limit, 56% of rounds were completed inside 2 minutes. This difference in time usage may explain why final choices were slightly worse in the choice process treatment, as we discuss below.

¹²A graph of this distribution was shown in the experimental instructions, which are reproduced in appendix A. The beta distribution was chosen in order to “front load” the probability of a time being selected in the first minute of the choice round, as most subjects made their choices inside 120 seconds.

button, which completed the round as if they had kept the same selection for the remaining time. Typically, a subject took part in a single session consisting of 2 practice rounds and 40 regular rounds, and two recorded choices were actualized for payment, which was added to a \$10 show up fee.

The choice process experiment (experiment 2) made use of exactly the same treatments as the standard choice experiments of experiment 1: choice sets contained 10, 20 or 40 alternatives, with the complexity of each alternative being either 3 or 7 operations. Moreover, exactly the same choice sets were used in the choice process and standard choice experiments.¹³

3.3 Basic Properties of Choice Process Data

Before using the choice process apparatus to estimate models of search, we establish two properties that are important for its usefulness. First, we show that final choices made under the choice process regime are similar to those made under standard choice conditions. This suggests that there is some similarity in the choice making procedure used in the choice process and standard choice experiments. Second, people do indeed change their selection with consideration time. This is a necessary condition for choice process data to contain more information than standard choice data alone.

For this analysis we discard observations from rounds in which the subject does not press the submit button before the allotted 120 seconds. In such rounds, we assume that subjects have not finished their choice process, so we cannot assume that we are observing their final choice. In doing so, we lose 94 rounds, or 8% of our total observations.

3.3.1 Impact on Final Choices

Table 2 compares failure rates and average absolute loss by treatment for choice process and non-choice process data. It also shows the number of observations per treatment for the choice process data.

¹³We also conducted experimental sessions in which the grand set of objects was 29 lotteries of the form P% chance of \$X and 1-P% chance of \$Y. In each round, 11 of the 29 lotteries were presented in a list on the screen, and choices were recorded at randomly selected times distributed uniformly between 1 and 60 seconds.

TABLE 2 ABOUT HERE

The comparative statics of loss and failures of optimality are very similar for the choice process experiment and the standard choice experiment. In both cases, subjects fail to optimize more frequently and lose more money in larger and more complicated choice sets. While it appears that choice process data leads to somewhat higher losses (and less optimal selection) on average, regression analysis suggests the effect is insignificant for percentage loss, and on the border of significance for failure rate.¹⁴

To the extent that there is a difference in the quality of final choices, it goes in the expected direction. The incentive to continue searching is higher in the standard choice experiment, since it is certain that any identified improvements will be implemented. The corresponding probability is less than one in the choice process experiment, and falls toward zero as the 2 minutes come to an end. In this light, it is noteworthy how limited was the impact of the incentive changes induced by the choice process interface. When we compare the distribution of final choices in each choice set from choice process and non-choice process sessions using Fisher's exact test, we find that 12 (20%) of the 60 choice sets have distributions that are significantly different at the 5% level. More tellingly, the analysis of section 7 shows that all of the results that follow concerning the nature of the search and decision process from the choice process experiments are closely mirrored using data from the pre-decision period in the pure choice experiment.

3.3.2 Number of Switches

Choice process data provide significantly more information than standard choice data in the form of switches in the period prior to finalization. Figure 2 shows histograms of the number of choice switches per round for each treatment. We define a choice switch as an occasion in which the subject changes selection from one alternative to another, excluding the initial change away from the \$0 option. Across all trials, 67% of rounds contain at least one switch and 37% contain at least two switches, indicating that people do use the choice process technology to update their choices as they contemplate the problem.

¹⁴To test this hypothesis, we repeat the regression analysis of section 2.2 on the combined standard choice and choice process data set with an additional dummy for whether or not choice process was implemented. The estimated coefficient is 0.637 (p-value of 0.260) for absolute loss and 9.01 (p-value of 0.052) for failure rate.

FIGURE 2 ABOUT HERE

One feature of this data is that there are many instances in which there are zero switches - subjects switch away from the initial zero option then stop. This suggests the possibility that there may be changes of mind that are not recorded in the choice process data, possibly due to perceived transactions costs of making the switch. It is important to note that none of our analyses are impacted by this possibility: our results are all consistent with behavior in which there is a private threshold of significance that has to be crossed before a change is recorded. In other words, our analysis is robust to the possibility that we do not observe all changes in preferences.

4 ABS and Mistakes

4.1 ABS

We now introduce a model of information search to shed light on apparent mistakes. The model we consider is ABS, the process of sequential search with recall, in which the DM evaluates over time an ever-expanding set of objects, choosing at all times the best object thus far identified.¹⁵ ABS is a common feature of classic models of search within economics [McCall, 1970; Stigler, 1961]) and of many boundedly rational models such as that of Simon [1955].

As defined by Caplin and Dean [2009], choice process data has an ABS representation if there exists a fixed utility function and a non-decreasing search correspondence for each choice set such that what is chosen at any time is utility maximizing in the corresponding searched set.

Definition 2 *Choice process (X, C) has an **ABS** representation (u, S) if there exists a utility function $u : X \rightarrow \mathbb{R}$ and a search correspondence $S : \mathcal{X} \rightarrow \mathcal{Z}^{ND}$, with $S_A \in \mathcal{Z}_A$ all $A \in \mathcal{X}$, such that,*

$$C_A(t) = \arg \max_{x \in S_A(t)} u(x)$$

where $\mathcal{Z}^{ND} \subset \mathcal{Z}$ comprises non-decreasing sequences of sets in \mathcal{X} , such that $Z_t \subset Z_{t+1}$ all $t \geq 1$.

¹⁵This contrasts with other more intricate forms of search involving partial understanding of all options (e.g. those based on exploring attributes and/or continuously learning about multiple options).

Caplin and Dean [2009] provide a general method of identifying whether or not choice process data has an ABS representation. The key to this representation is understanding what type of behavior implies a revealed preference in the context of the ABS model. It is not the case that final choice of x over y necessarily indicates that x is preferred to y , as the decision maker may simply be unaware of y . However, if we see a subject at some point choose y and then replace it with x then under the ABS model they must be interpreted as preferring x to y . The fact that y has previously been chosen indicates that the subject is aware of it. However, the subject has later rejected y in favor of x , indicating that the latter must be preferred.

In general, choice process data will have an ABS representation if and only if this revealed preference information is consistent with some underlying linear order - in other words, it must be acyclic. However, in our experiments, we have an externally observable ranking over the objects of choice, given by their underlying dollar value. The corresponding result is therefore trivial: an ABS representation exists for our data if and only if all switches are to higher value alternatives. This result is noted in remark 1:

Remark 1 *Let $v : X \rightarrow \mathbb{R}$ be the externally observable value of a set of choice objects. A choice process model (X, C) permits an ABS representation (v, S) if and only if $v(C_A(t)) \leq v(C_A(t + s))$ for all $A \in \mathcal{X}$ and $t, s \geq 1$ (**Condition 1**).*

4.2 Testing ABS

In order to measure how close our data is to satisfying condition 1, we use a measure of consistency proposed by Houtman and Maks [1985]. The Houtman-Maks (HM) index is based on calculating the largest number of observations that are consistent with a particular condition, which can be determined by finding the minimum number of observations that have to be removed before the condition is satisfied. The underlying idea is that a data set that requires fewer such removals is “closer” to satisfying condition 1 than one that requires more removals. In this case, we specifically ask how many selections have to be removed from a subject’s data set before condition 1 is satisfied. The resulting HM Index is normalized by dividing through by the total number of observations, so that the HM Index takes a value between 0 and 1, which can be interpreted as the largest fraction of a data set that satisfies condition 1

To give a concrete example, consider that for one subject we observe that they initially select

an option worth 7, then one worth 6, then 8 then 9. Such a subject would not be consistent with condition 1, as their initial switch would be to a lower value. However, if we removed their second selection, their choice process data would show them switching from value 7 to value 8 to value 9 - in line with condition 1. Thus this subject would have an HM index of 0.75, as 1 of their 4 observations would have to be removed to make their data consistent with condition 1. This subject we consider closer to satisfying condition 1 than one who switched from value 7 to value 6 to value 9 to value 8. We would have to remove two observations from this subject's data to make them consistent with condition 1, giving them an HM index on 0.50

To determine the relative consistency of subjects in the choice process experiment, we compare their selections to a benchmark of random choice, as proposed by Bronars [1987]. For each subject, a benchmark choice process data set is constructed by replacing each selection with a random selection from the corresponding choice set, so that the resulting random choice process data has the same number of selections in each round as the original data.

Figure 3 shows the results of the benchmarking. The top histogram shows the distribution of HM Index scores for all 76 subjects using their actual selections, and the bottom histogram shows the distribution of HM Index scores for 1,000 simulations of random data for each subject in the way described above, which gives a total of 76,000 simulated scores. A two-sample Kolmogorov-Smirnov test indicates that the distributions are significantly different ($p < .001$).

FIGURE 3 HERE

4.3 Identifying ABS Types

Figure 3 suggests that, for the population as a whole, ABS does a good job of describing search behavior. We can also ask whether the behavior of a particular subject is well described by the ABS model - if so, we describe this subject as an ABS type.¹⁶

¹⁶While the choice process data for this experimental setting can be modeled well with ABS, it remains to be shown that ABS is appropriate for other choice objects. Therefore, we ran an additional treatment of 20 rounds with 21 subjects using the lotteries. The grand set of objects was 29 lotteries of the form P% chance of \$X and 1-P% chance of \$Y. In each round, 11 of the 29 lotteries were presented in a list on the screen. Despite the complexity and novelty of the choice objects, many of these 21 subjects can be modeled well with ABS. Because preferences are not immediate for these objects, we performed a more general test for acyclicity. For 16 subjects, 90% or more selections

To identify ABS types, we compare each subject's HM Index with the median HM Index of the 1,000 simulations of random data for that subject, which have exactly the same number of observations in each round. Only 1 subject (727) has an HM Index below the median HM Index of the corresponding random choice process data, and only 4 subjects (638, 680, 727, and 826) have an HM Index lower than the 75th percentile. For the remainder of the paper we focus on the 72 out of 76 subjects we classify as ABS types.¹⁷

4.4 ABS and Mistakes

Under the standard model of decision making, preferences are revealed through final choice: one object is revealed preferred to another if it is chosen when the other was available. Because our experiment makes use of objects with externally observable values, we have defined a mistake relative to the standard model as a case when revealed preference is not in line with the external valuation - in other words, when one object is chosen though a more valuable object was available.

The ABS model incorporates a different notion of revealed preference: preference is revealed not through final choice, but by switching from one alternative to another. We can therefore define the concept of a mistake relative to the ABS model as a case when this definition of revealed preference is not in line with the external value. Viewed this way, the HM index calculated above counts the proportion of observations which are consistent with an absence of mistakes.

Figure 4 compares the proportion of mistakes according to the ABS model and according to the standard model for each treatment. Two key facts stand out in this figure. First, the level of irrationality as measured by the standard definition of revealed preference is far higher than that with the ABS measure. Second, while there is strong evidence of increasing irrationality in larger and more complex choice sets according to the standard measure, such effects are minimal according to the ABS measure - using the latter, there is no effect of set size, and only a small effect of complexity on mistakes.

FIGURE 4 ABOUT HERE

are consistent with acyclicity

¹⁷Using a cutoff of the 95th percentile would lead to the loss of 3 more subjects, and would not change any of the following results.

Figure 4 suggests that simple search theoretic explanations can help make sense of the mistakes that we observe. In large choice sets, people still recognize preferred objects and choose them when they come across them. However, their final choices may not be maximal because they do not search through all available alternatives. Thus, the ABS model can resurrect the concept of revealed preference in environments in which decision makers must search for information on the available alternatives.

5 Satisficing

In his pioneering model of bounded rationality, Simon [1955] suggested that decision makers do not optimize, but rather search through a decision set until they achieve a “satisfactory” (or reservation) level of utility. One factor that has held back research on satisficing behavior is that the model has typically been interpreted in terms of its implications for final choices alone. The problem in this regard is that the simplest form of satisficing cannot be separated from utility maximization on the basis of choice alone: both are characterized by final choices that obey the weak axiom of revealed preference.¹⁸

In this section we use choice process data to shed new light on satisficing behavior. The essential advantage that choice process data provides is that it opens up observation of both unsatisfactory as well as satisfactory choices, in that we directly observe occasions when a subject continues to search having uncovered an unsatisfactory object. This allows us to estimate reservation values for our different treatments.

The bottom line is that a simple model of satisficing behavior in which the satiation level is dependent on ex ante known features of the decision making context has great explanatory power in our data set. Moreover, reservation levels depend in a predictable way on our two treatment variables, choice set size and complexity.

¹⁸This is true in the version of the satisficing model in which decision makers always search through choice objects in the same order, and the set of satisficing objects is fixed. If the order of search can change over time, then the satisficing model has no implication for final choice.

5.1 Satisficing and Reservation Utility

In search theoretic terms, satisficing behavior corresponds to ABS behavior coupled with a reservation level of value (or utility): a subject searches through the choice set item by item, stopping if and only if this reservation level is achieved. This connection to sequential search based on a simple stopping rule link satisficing with our experimental data.

The first indication that our subjects exhibit satisficing behavior is shown in figure 5. This shows how the value of the selected object changes with order of selection for each of our six treatments. Each graph has one isolated point and three lines. The isolated point shows the average object value for those who stop at the first object chosen.¹⁹ The first line shows the average value of each selection from rounds in which one switch was made. The next line shows the average value of each selection in rounds where 2 switches were made, and the final line for rounds in which 3 switches were made.

FIGURE 5 ABOUT HERE

Figure 5 is strongly suggestive of satisficing behavior. First, as we would expect from the proceeding section, in aggregate people behave in line with ABS: in all but one case, the average value of selections is increasing. Second, we can find reservation values for each treatment such that aggregate behavior is in line with satisficing according to these values. The horizontal lines drawn on each graph show candidate reservation levels, estimated using a technique we describe below. In every case, the aggregate data show search continuing for values below the reservation level, and stopping for values above the reservation level, as with satisficing behavior.

5.2 The Estimator

In order to estimate the reservation utility for each treatment, we assume a stochastic generalization of the reservation strategy. We assume that all individuals in a given choice environment have the same constant reservation value \bar{v} and experience variability ε in this value each time they decide whether or not to continue search. Further, we assume this stochastically enters additively and is drawn independently and identically from the standard normal distribution. Let v be the value

¹⁹Following the initial switch away from the zero value option.

of the item that has just been evaluated, and so the DM uses the following strategy to determine whether to continue searching through the choice set:

$$\begin{aligned} \text{search stops if } v &> \bar{v} + \varepsilon : \\ \text{search continues if } v &\leq \bar{v} + \varepsilon; \end{aligned}$$

where $\varepsilon \sim N(0, 1)$.

We can recast this procedure as a binary choice model. Let k be a decision node, v_k be the value of the object uncovered and x_k be the choice made at that decision node, with $x_k = 1$ if search stops and $x_k = 0$ if search continues. Then

$$x_k = 1(v_k - \bar{v} - \varepsilon_k > 0),$$

where $1(\cdot)$ is the indicator function.

An individual will stop searching if $\varepsilon_k < v_k - \bar{v}$, so the probability of stopping is search is $\Phi(v_k - \bar{v})$, where Φ is the cumulative density function of the standard normal distribution. Similarly, search will continue if $\varepsilon_k > v_k - \bar{v}$, so the probability of search continuing is given by $1 - \Phi(v_k - \bar{v}) = \Phi(\bar{v} - v_k)$.

Thus, to estimate the parameter \bar{v} with maximum likelihood estimation, we use the log likelihood function

$$\ln \mathcal{L} = \sum_{k=1}^K [x_k \ln(\Phi(v_k - \bar{v})) + (1 - x_k) \ln(\Phi(\bar{v} - v_k))]$$

and find the value of \bar{v} maximizes $\ln \mathcal{L}$.

To employ this procedure using our data, we consider each selection made by a subject as a decision node. We then need to identify occasions when we observe that search has stopped, and when we observe that it has continued. The latter is simple: search continues if a subject switches to another alternative after the current selection. Identifying stopped search is slightly more complicated. If we observe that a subject does not make any more selections after the current one, then there are three possibilities. First, they could have continued to search, but run out of time before they found a better object. Second, they could have continued to search, but already have selected the best option. Third, they could have stopped searching. We therefore consider a subject to have stopped searching at a decision node only if they made no further selections, pressed the submit button, and the object they had selected was not the highest value object in the choice set.

Choice process data is clearly vital for the estimation of reservation values. If we ignore data on the choice process and instead consider only standard choice data, we cannot use the same estimation strategy because it requires observations of subjects continuing to search as well as observations in which they stop searching. Choice data is composed entirely of the latter, so it only indicates when search has stopped, not when it continues.

5.3 Estimated Reservation Levels

Because we assume that all individuals have the same distribution of reservation values in a given environment, we pool together all selections within each treatment. We estimate reservation levels for the 72 participants whose choice data is best modeled with ABS. Table 3 shows the estimated reservation levels for each treatment, with standard errors in parentheses.

TABLE 3 ABOUT HERE

Table 3 reveals two robust patterns in the estimated reservation levels. First, reservation levels decrease with complexity: using a likelihood ratio test, estimated reservation levels are significantly lower for high complexity treatments than for low complexity treatments for all set sizes ($p < 0.001$). Second, reservation levels increase monotonically with set size (significantly different across set sizes for both complexity levels with $p < 0.001$).

One question that this estimation strategy does not answer is how well RBS behavior explains our experimental data. In order to shed light on this question, we calculate the equivalent of the HM index for the RBS model with the estimated reservation levels of table 3. For each treatment, we calculate the fraction of observations which obey the reservation strategy (i.e. subjects continue to search when they hold values below the reservation level and stop when they have values above the reservation level).

TABLE 4 ABOUT HERE

The results, aggregated across all subjects, are shown in table 4. The estimated RBS model describes about 85% of observations for treatments with simple objects and about 80% for complicated objects. Both of these figures are significantly higher than the random benchmark of 50%

(where people arbitrarily stop or continue at each decision node) at the 1% level.

As with the ABS model, there is significant heterogeneity across individuals with respect to how well they are described by the RBS model. While the majority of subjects have an HM index above 75%, some have extremely low scores and are clearly poorly described by the RBS model with the given estimated reservation levels. In order to ensure these individuals are not affecting our estimates in table 3, we repeat the estimation of reservation strategies while dropping subjects who have an HM index below 50%. These results are in table 3 under the rows for “RBS” types. The estimated reservation levels are essentially the same as those for the whole sample.

5.4 Reservation Utility or Reservation Time?

A natural question is whether our data is consistent with other stopping rules. One obvious candidate is a stopping rule based on a reservation time, in which subjects search in a given environment for a fixed time, selecting the best option found subject to this time constraint. In order to test this possibility, we redraw the graphs of figure 5, but showing the average time of each switch, rather than the average value. If subjects are using a fixed stopping time strategy then we expect the graphs to look like those in figure 5 - on average, subjects stop searching when time is over the stopping time, and continuing when it is less than the stopping time.

FIGURE 6 ABOUT HERE

The results of the above analysis are shown in figure 6. The figures are completely destructive of the reservation time alternative. Unlike in figure 5, there is generally no “reservation time” such that subjects continue to search for times below this level and stop for times above that level. It appears that those who identified a high value object with their first selection stopped quickest, while those who made the most switches along the way took far longer. This is precisely as the reservation utility model would suggest, and runs completely counter to the predictions of the reservation time model.

6 Optimal Stopping

In this section we explore the connection between the reservation stopping rules that we identify in the experiment and optimal stopping rules. We establish that reservation stopping rules of the kind that we uncover are optimal in the context of our experimental design. We consider a standard model of sequential search with a search cost specified in utility terms, as in Gabaix et al. [2006]. The DM is an expected utility maximizer with a specific final utility function $u : X \rightarrow \mathbb{R}$ that represents object values. The agent’s search strategy from any non-empty finite subset $A \subset X$ is based only on the size M of the set of available objects in A , not the identities of these objects. Each available option is assumed ex ante to have a utility level that is independently drawn from some distribution $F(z)$. Note that this is explicitly true in our experiment.

We endow the searcher with information on one available option. At each subsequent time $t \geq 1$, the decision maker faces the option of selecting one of the options already searched, or examining an extra option and paying the additional psychic search cost $\kappa > 0$. Once search stops, the agent must choose one of the uncovered objects.²⁰ There is no discounting. In this environment, we establish that the optimal search strategy is based on a fixed reservation level of utility.

Theorem 1 *Given that search costs satisfy $0 < \kappa < \int_0^\infty z dF(z)$, define reservation utility R as the unique solution to the equation*

$$\int_R^\infty (z - R) dF(z) = \kappa.$$

The expected utility maximizing strategy is to continue search until and unless an option is uncovered with utility strictly above the cutoff level R , with immediate selection of any such object.

Proof. We prove the result inductively on n , the number of remaining unsearched elements in a set of initial cardinality $N \geq n$. Supposing that search continues until there is only one element left unsearched, let x_1 be the highest utility object encountered in prior search. The optimal strategy is either to stop immediately and take this option, or to continue. The continuation results in net expected utility gain $G_1(x_1)$ as the result of one additional search, comprising the possible surplus above x_1 if the final object uncovered has such a utility balanced against the additional search

²⁰This method of modeling makes the process of uncovering an option equivalent to the process of “locating” it as feasible. The strategy is more intricate if we allow unexplored options to be selected.

costs,

$$G_1(x_1) = \int_{x_1}^{\infty} (z - x_1) dF(z) - \kappa.$$

The upper bound we have imposed on the search costs imply that $G_1(0) > 0$, implying that continued search is worthwhile. In addition note that $G_1(x_1)$ is strictly decreasing in x_1 , that $G_1(R) = 0$, and that $\lim_{x_1 \rightarrow \infty} G_1(x_1) = -\kappa$. Hence it is uniquely optimal to search the final object if $x_1 < R$, strictly optimal to stop if $x_1 > R$, with indifference between searching and stopping if $x_1 = R$.

Assume now that this precise search strategy is optimal if search continues until there are some $n \geq 1$ elements left unsearched: defining x_n as the maximum value object encountered in the prior search, assume that it is uniquely optimal to search the final object if $x_n < R$, strictly optimal to stop if $x_n > R$, with indifference between searching and stopping if $x_n = R$. Now consider the optimal strategy with $n + 1$ elements left unsearched, defining x_{n+1} as the maximum value object encountered in the prior search.

- If $x_{n+1} > R$, any optimal search strategy involves searching at most one more time by the inductive hypothesis. Hence the net gain from continued search is precisely as identified by the function G_1 introduced above, so that the strict optimality of immediately stopping follows from the fact that $G_1(x_{n+1}) < G_1(R) = 0$.
- If $x_{n+1} < R$, the expected utility gain from continued search is bounded below by $G_1(x_{n+1}) > 0$, which is the value of the strategy of searching for one more period and then stopping for sure. Hence the unique optimal strategy for $x_{n+1} \leq R$ is to so continue.
- If $x_{n+1} = R$, the expected utility gain from continued search is bounded below by $G_1(R) = 0$, so that continuation is an optimal strategy. On the other hand, by the inductive hypothesis it is also optimal to continue one more period and then stop for sure, which gives rise to an expected gain of precisely 0. Hence stopping immediately is also an optimal strategy.

■

Theorem 1 also tells us how the optimal reservation level varies across our experimental treatments. First, the optimal reservation level falls as the per unit search cost rises. Thus, assuming that search costs are higher for the 7 than for the 3 complexity objects, this implies that optimal

reservation levels are lower in the higher complexity environment. Second, optimal reservation levels are independent of the size of the choice set: there is no increase in the optimal reservation level as the size of the choice set increases.

Thus, the comparative statics properties of our estimated stopping rules do not align perfectly with those of the optimal stopping rule. While we do find that subjects reduce their reservation level in response to higher search costs, they also tend to *increase* their reservation level as the size of the choice set increases.

There are two possible reasons for this discrepancy between optimal behavior and this observation. The first is that subjects are behaving optimally with respect to a different maximization problem. For example, theorem 1 assumes that no learning takes place with respect the distribution of values of the objects in the choice set. While our subjects are explicitly told the distribution from which values are drawn, it may be that they in fact try to learn this distribution for every new choice set. In such a case, estimated reservation levels would in many cases be greater in larger choice sets.

A second possibility is that subjects are acting sub-optimally by increasing their reservation levels in larger choice sets: they are searching “too much” in larger choice sets relative to smaller ones. This result may relate to findings from the psychology and experimental economics literature that show that people have preferences for smaller choice sets [Iyengar and Leper, 2000; Seunanz-Salgado, 2006]. One factor that potentially links these two findings is the concept of regret. Zeelenberg and Pieters [2007] show that decision makers experience more regret in larger choice sets, and suggest that can lead them to search for more information.

7 Choice Process vs Non-Choice Process Data

An important question is how the elicitation of choice process data impacts the decision making process. Clearly, our results are of more interest if the experimental techniques used to elicit choice process data are not having a marked impact on the way in which people make decisions. In order to explore this issue, we re-run the above analyses of choice process data using data from the standard choice experiment described in section 2. Recall that, in this experiment, subjects could select options prior to their final choice just as they could in the choice process experiment. The

only difference was that in the standard choice experiment, there was no incentive for them to do so. However, as figure 7 shows, subjects still did record switches of their own volition. We can therefore treat these switches as choice process data and test whether the results derived above survive.

FIGURE 7 ABOUT HERE

It turns out that the results from switches recorded in the standard choice treatment are remarkably similar to those from the choice process data. Firstly, subjects still exhibit ABS behavior: figure 8 compares the estimated distribution of HM indices for the choice process and standard choice experiments, which shows that, if anything, standard choice data are more in line with ABS than the choice process data. Figure 9 repeats the analysis of figure 4 for the standard choice data, comparing the proportion of “mistakes” by treatment, as measured by the ABS model and standard revealed preference. Again, we see little effect of treatment on mistakes for the ABS based measure.

FIGURE 8 AND 9 ABOUT HERE

The data from the standard choice experiments is also in line with RBS behavior. Figure 10 recreates the analysis of figure 5, and suggests that a reservation stopping rule broadly describes the aggregate data. Table 5 shows the estimated reservation levels for the standard choice data exhibit the same comparative statics as do those for the choice process data, while table 6 shows that the estimated HM indices for these reservation levels are only slightly lower than for the choice process data.

FIGURE 10 AND TABLE 5 AND 6 ABOUT HERE

8 Search Order and Choice

In this section we show that choice process data provide insight into the order in which people search through available objects, and that this information can help predict when subjects will

do badly in particular choice sets. We are interested in two particular factors that can determine search order: screen position and object complexity. In order to explore both factors, we ran an additional experimental treatment which contained objects of varying complexity. This treatment contained choice sets of size 20, and the objects in each set varied in complexity from between one and nine operations. We ran the new treatment on 20 subjects for a total of 206 observed choice sets.

8.1 Aggregate Search Order

Figure 11 shows how average screen position and complexity of selection change with selection order. As with figure 5 separate lines show the average screen position and complexity for rounds in which 0, 1, 2, 3 and 4 switches were made. Screen position is encoded from top to bottom (i.e. the top object on the screen has position 1, the second has position 2 and so on), while complexity is encoded as the number of arithmetic operations needed to evaluate each prize.

FIGURE 11 ABOUT HERE

These figures suggest that average search behavior has systematic patterns. The first graph shows that, on average, subjects search the screen from top to bottom: screen position is higher for later selections. The second panel shows that subjects also tend to search from simple to complex objects. As complexity is uncorrelated with value, this is in line with optimal strategy. While neither relationship is completely monotonic, regression analysis confirms that the both are significant.²¹

8.2 Individual Search Order

We can augment our analysis of aggregate search behavior by looking at the search patterns of individual subjects. We look for subjects whose behavior is consistent with “Top-Bottom” (TB) search, and those whose behavior is consistent with “Simple-Complex” (SC) search. The former are subjects whose search order takes them from the top to the bottom of the screen, while the latter are subjects whose search takes them from simple to complex objects.

²¹Regressing selection number on the screen position and complexity of the selection gives coefficients of 0.034 and 0.132 respectively, both significant at the 1% level (allowing for clustering at the subject level).

We categorize subjects by calculating their HM indices assuming each of these two search orders. We first assume that the subject is searching top to bottom and calculate the fraction of observations that are consistent with this search order. We then repeat the procedure assuming that the subject searches from simple to complex. A subject is categorized as being a TB or SC searcher if their HM Index for that search order is in the 95th percentile compared to a benchmark distribution constructed using random search orders. Note that the majority of subjects are well described by either TB or SC search or both. As table 7 shows, only 2 subjects fall into neither category.

TABLE 7 ABOUT HERE

In the experimental treatments described in section 3, we would categorize 51% of subjects as TB searchers using the same metric.

8.3 Search Order and Choice

We provide two simple examples that illustrate how knowledge of a subject's search order helps predict those choice sets in which they will make large mistakes and those in which they will not. Example 1 is from a round in which the highest valued item is very short and occurs at the end of the list (see figure 12 panel A- best option highlighted in green). We would expect this to be a choice set in which TB searchers would do badly and SC searchers would do well. This turns out to be the case. Pure top-bottom searchers find the best option least often (66% of the time), those that search top-bottom but also simple-complex find it more often (83%) and pure simple-complex searchers find it most often (100%).²² Unfortunately, due to the small sample size, these numbers are not significant at the standard levels of significance (difference between pure simple-complex searchers and other subjects has a p-value of 0.12)

FIGURE 12 ABOUT HERE

Example 2 is from a round in which the highest valued item is very long and occurs very early in the list (figure 12 panel B). In this case, we would expect TB searchers to do well and SC searchers

²²In order to avoid potential circularity in our argument, we re-estimate subjects' search types excluding these two example rounds. The results are unchanged.

to do badly. Pure top-bottom searchers find the best option most often (80% of the time), those that search top-bottom but also simple-complex find it less often (71%) and pure simple-complex searchers find it even less (66%).

Note that in categorizing people as TB searchers, we are simply saying that options which are selected later occur further down the list. We have little to say about how their first selection is made - when they quickly move away from the option that gives them \$0 for sure. One could plausibly model this initial selection as a stochastic process, after which the TB searcher searches downward from that starting point. Our findings suggest that this is the case, as TB searchers are much less likely to find the best option if their initial choice is above the best option than below it. In example 2, those TB searchers whose initial selection comes after the best item in the list are much less likely to find it (70% of the time against 100% of the time for those whose first click is before the best option).

Note that above provides only the most rudimentary indication of the insights into search order that choice process data can provide. We see this as a very important subject of continued investigation, and an area in which complementing choice process data with other data on the search process, particularly eye-tracking data, will be of particular value.

9 Concluding Remarks

An important challenge for researchers is to unite revealed preference theory and the theory of search. We introduce a choice-based experiment that answers this challenge. We have used it to classify search behaviors in different decision making contexts. Our central finding concerns the prevalence of satisficing behavior. Models of sequential search based on achievement of context-dependent reservation utility closely describe our experimental data. More broadly, we believe that the search theoretic lens will be of significant value in systematizing our understanding of boundedly rational behavior.

10 Bibliography

References

- [1] Bernheim, B. Douglas, and Antonio Rangel. 2008. "Choice-Theoretic Foundations for Behavioral Welfare Economics." *The Foundations of Positive and Normative Economics*, ed. Andrew Schotter and Andrew Caplin.
- [2] Bronars, Stephen. 1987. "The Power of Nonparametric Tests of Preference Maximization." *Econometrica*, 55(3): 693-698.
- [3] Campbell, Donald. 1978. "Realization of Choice Functions." *Econometrica*, 46(1): 171-180.
- [4] Caplin, Andrew, and Mark Dean. 2009. "Choice Anomalies, Search, and Revealed Preference." Unpublished.
- [5] Eliaz, Kfir, and Ran Spiegler. 2008. "Consideration Sets and Competitive Marketing." Unpublished.
- [6] Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg. 2006. "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model." *American Economic Review*, 96(4): 1043-1068.
- [7] Gul, Faruk, and Wolfgang Pesendorfer. 2008. "The Case for Mindless Economics." *The Foundations of Positive and Normative Economics*, ed. Andrew Schotter and Andrew Caplin.
- [8] Houtman, Martijn, and J. A. H. Maks. 1985. "Determining all Maximal Data Subsets Consistent with Revealed Preference." *Kwantitatieve Methoden*, 19: 89-104.
- [9] Iyengar, S. S., and M. R. Lepper. 2000. "When Choice is Demotivating: Can One Desire Too Much of a Good Thing?" *Journal of Personality and Social Psychology*, 79: 995-1006.
- [10] Koszegi, Botond, and Matthew Rabin. 2008. "Choice-Theoretic Foundations for Behavioral Welfare Economics." *The Foundations of Positive and Normative Economics*, ed. Andrew Schotter and Andrew Caplin.
- [11] Manzini, Paola, and Marco Mariotti. 2007. "Sequentially Rationalizable Choice." *American Economic Review*, 97(5): 1824-1840.

- [12] Masatlioglu, Yusufcan, and Daisuke Nakajima. 2008. "Choice by Iterative Search." Unpublished.
- [13] McCall, John J. 1970. "Economics of Information and Job Search." *Quarterly Journal of Economics*, 84(1): 113-126.
- [14] Ok, Efe. 2002. "Utility Representation of an Incomplete Preference Relation." *Journal of Economic Theory*, 104: 429-449.
- [15] Payne, J. W. , J. R. Bettman, and E. J. Johnson. 1993. *The Adaptive Decision Maker*. Cambridge, MA: Cambridge University Press.
- [16] Reutskaja, E., J. Pulst-Korenberg, R. Nagel, C. Camerer, and A. Rangel. 2008. "Economic Decision-Making under Conditions of Extreme Time Pressure and Option Overload: An Eye-tracking Study." Unpublished.
- [17] Rubinstein, Ariel, and Yuval Salant. 2006. "A Model of Choice from Lists." *Theoretical Economics*, 1(1): 3-17.
- [18] Seuanez-Salgado, Maria. 2006. "Choosing to have Less Choice." FEEM Working Paper No. 37. Available at SSRN: <http://ssrn.com/abstract=889311>.
- [19] Simon, Herbert. 1955. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics*, 69(1): 99-118.
- [20] Stigler, George. 1961. "The Economics of Information." *Journal of Political Economy*, 69(3): 213-225.

Figure 1: A typical choice round

Round 2 of 30	Current selection: four plus eight minus four
Choose one:	
<input type="radio"/>	zero
<input type="radio"/>	three plus five minus seven
<input type="radio"/>	four plus two plus zero
<input type="radio"/>	four plus three minus six
<input checked="" type="radio"/>	four plus eight minus four
<input type="radio"/>	three minus three plus one
<input type="radio"/>	five plus one minus one
<input type="radio"/>	eight plus two minus five
<input type="radio"/>	three plus six minus five
<input type="radio"/>	four minus two minus one
<input type="radio"/>	five plus five minus one
<input type="button" value="Finished"/>	

Figure 2: Number of Switches per Choice Round, Experiment 2




Figure 3: Distribution of HM Indices for Experiment 1 (Actual vs. Random Data)

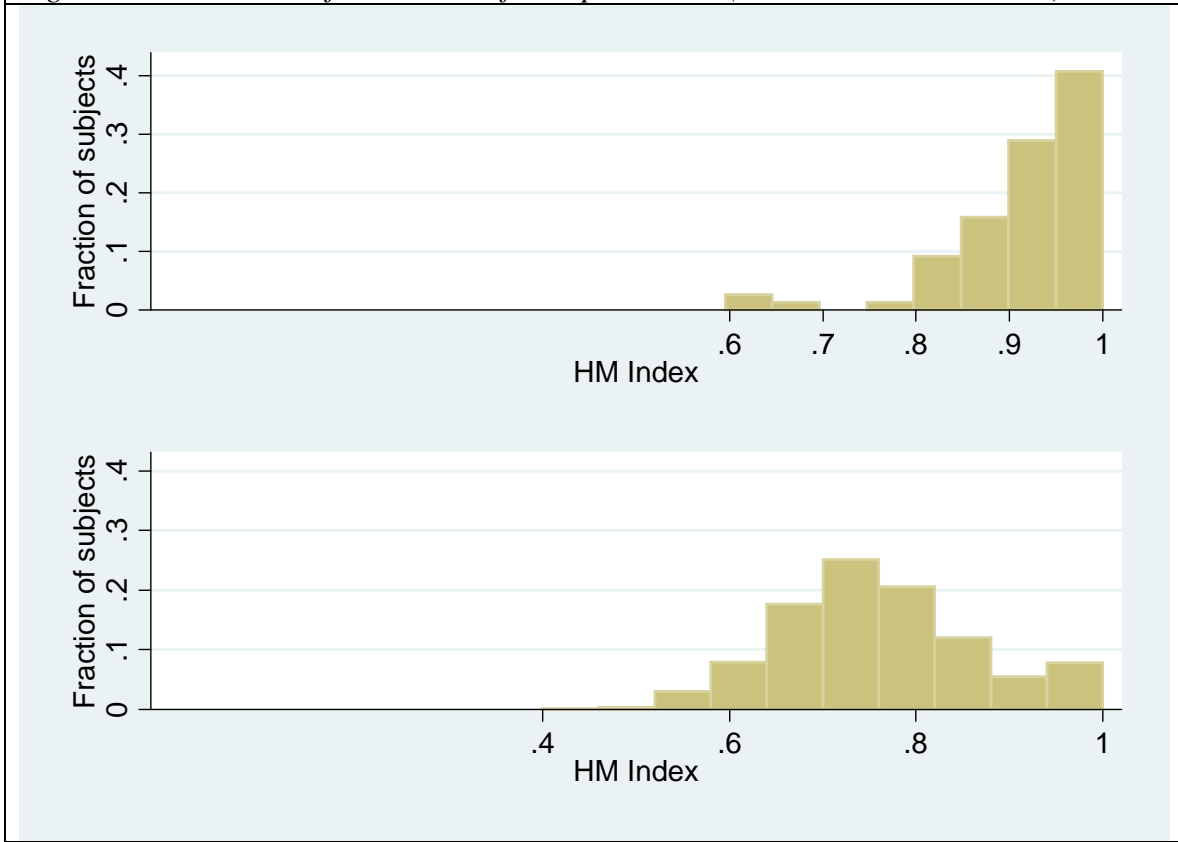


Figure 4 Proportion of Mistakes According to the Standard Model and ABS




Figure 5: Average Value by Switch

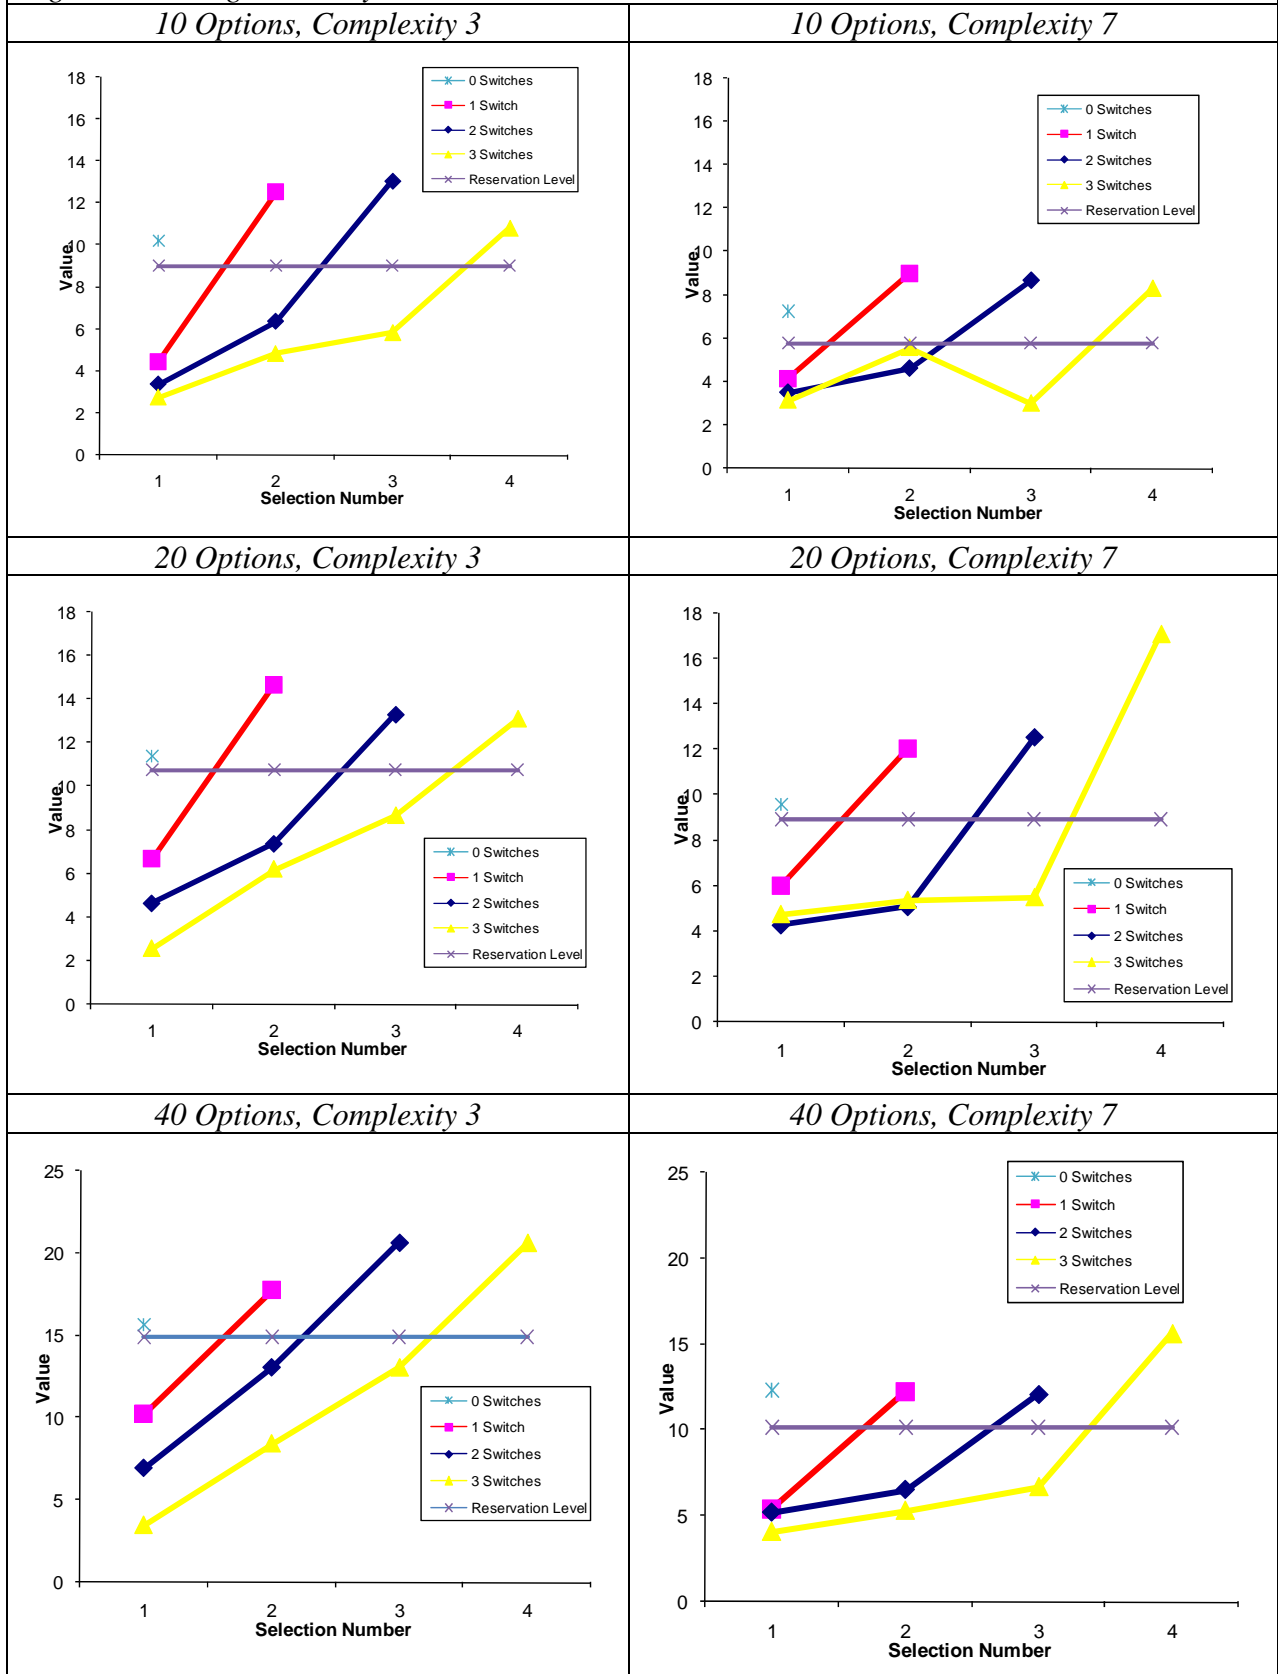


Figure 6: Average Time by Switch




Figure 7: Switches in Choice Process and Standard Choice Experiments

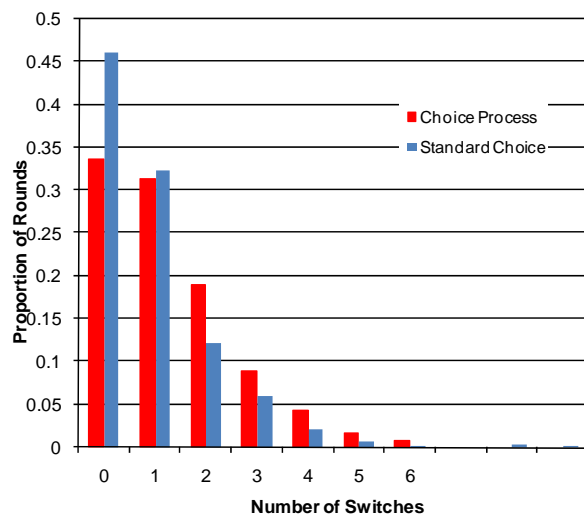


Figure 8: HM indices for Choice Tracking and Standard Choice Experiments

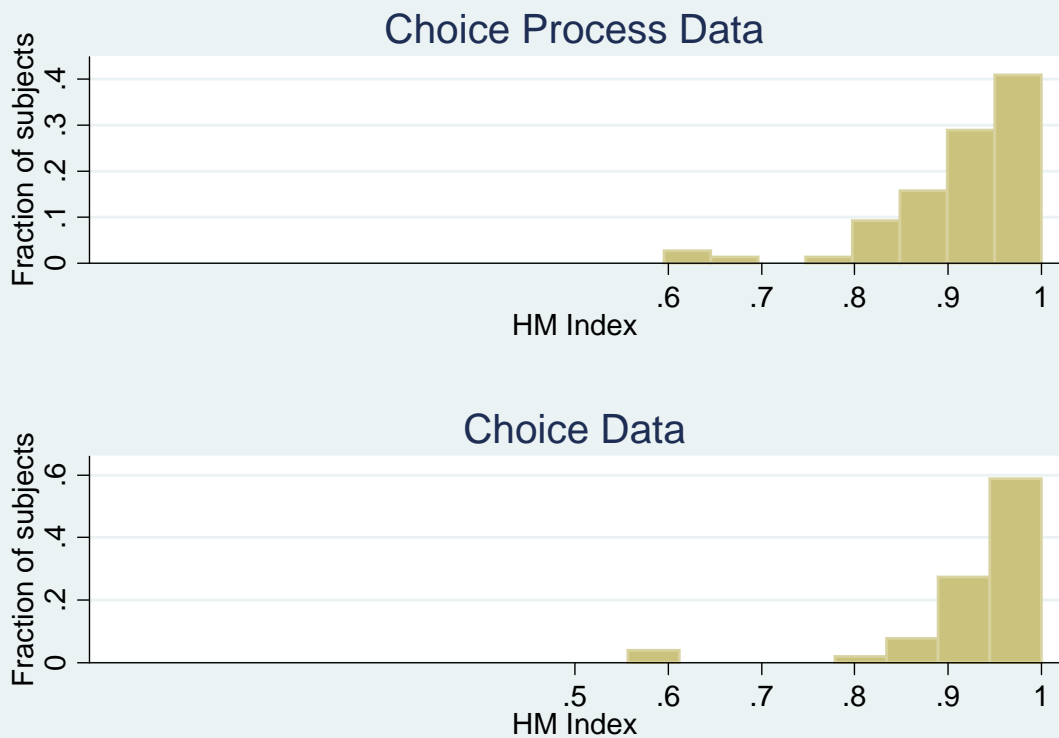



Figure 9: Proportion of Mistakes According to the Standard Model and ABS – Standard Choice Data

Panel A: Standard Model



Panel B: ABS Model

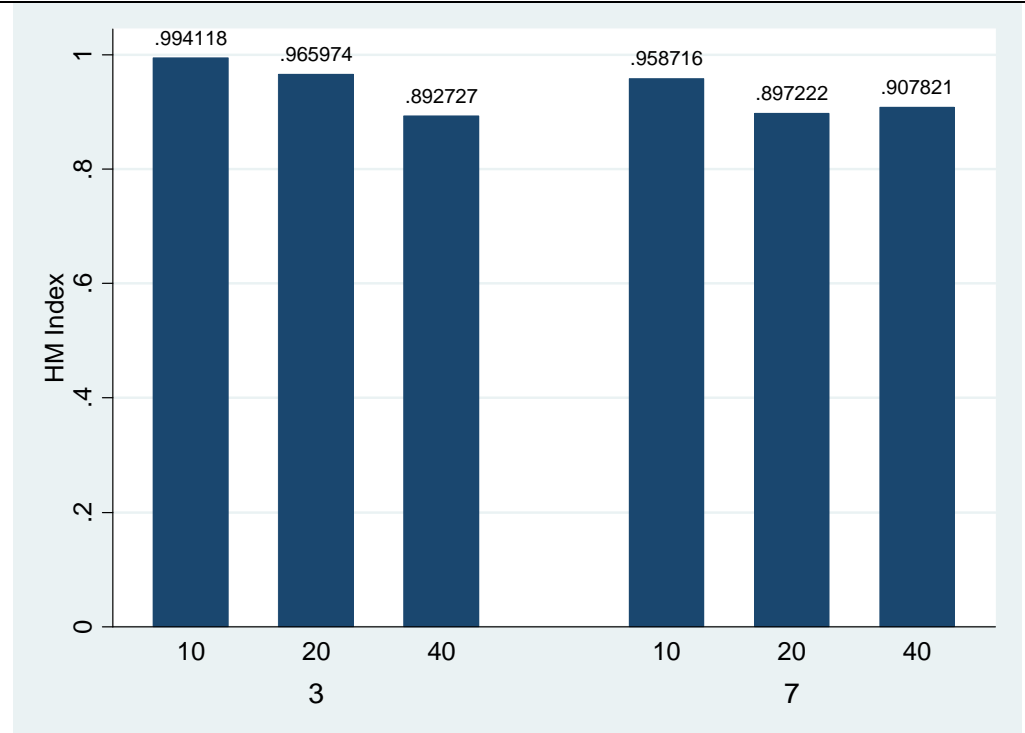


Figure 10: Average Value by Switch – Standard Choice Data

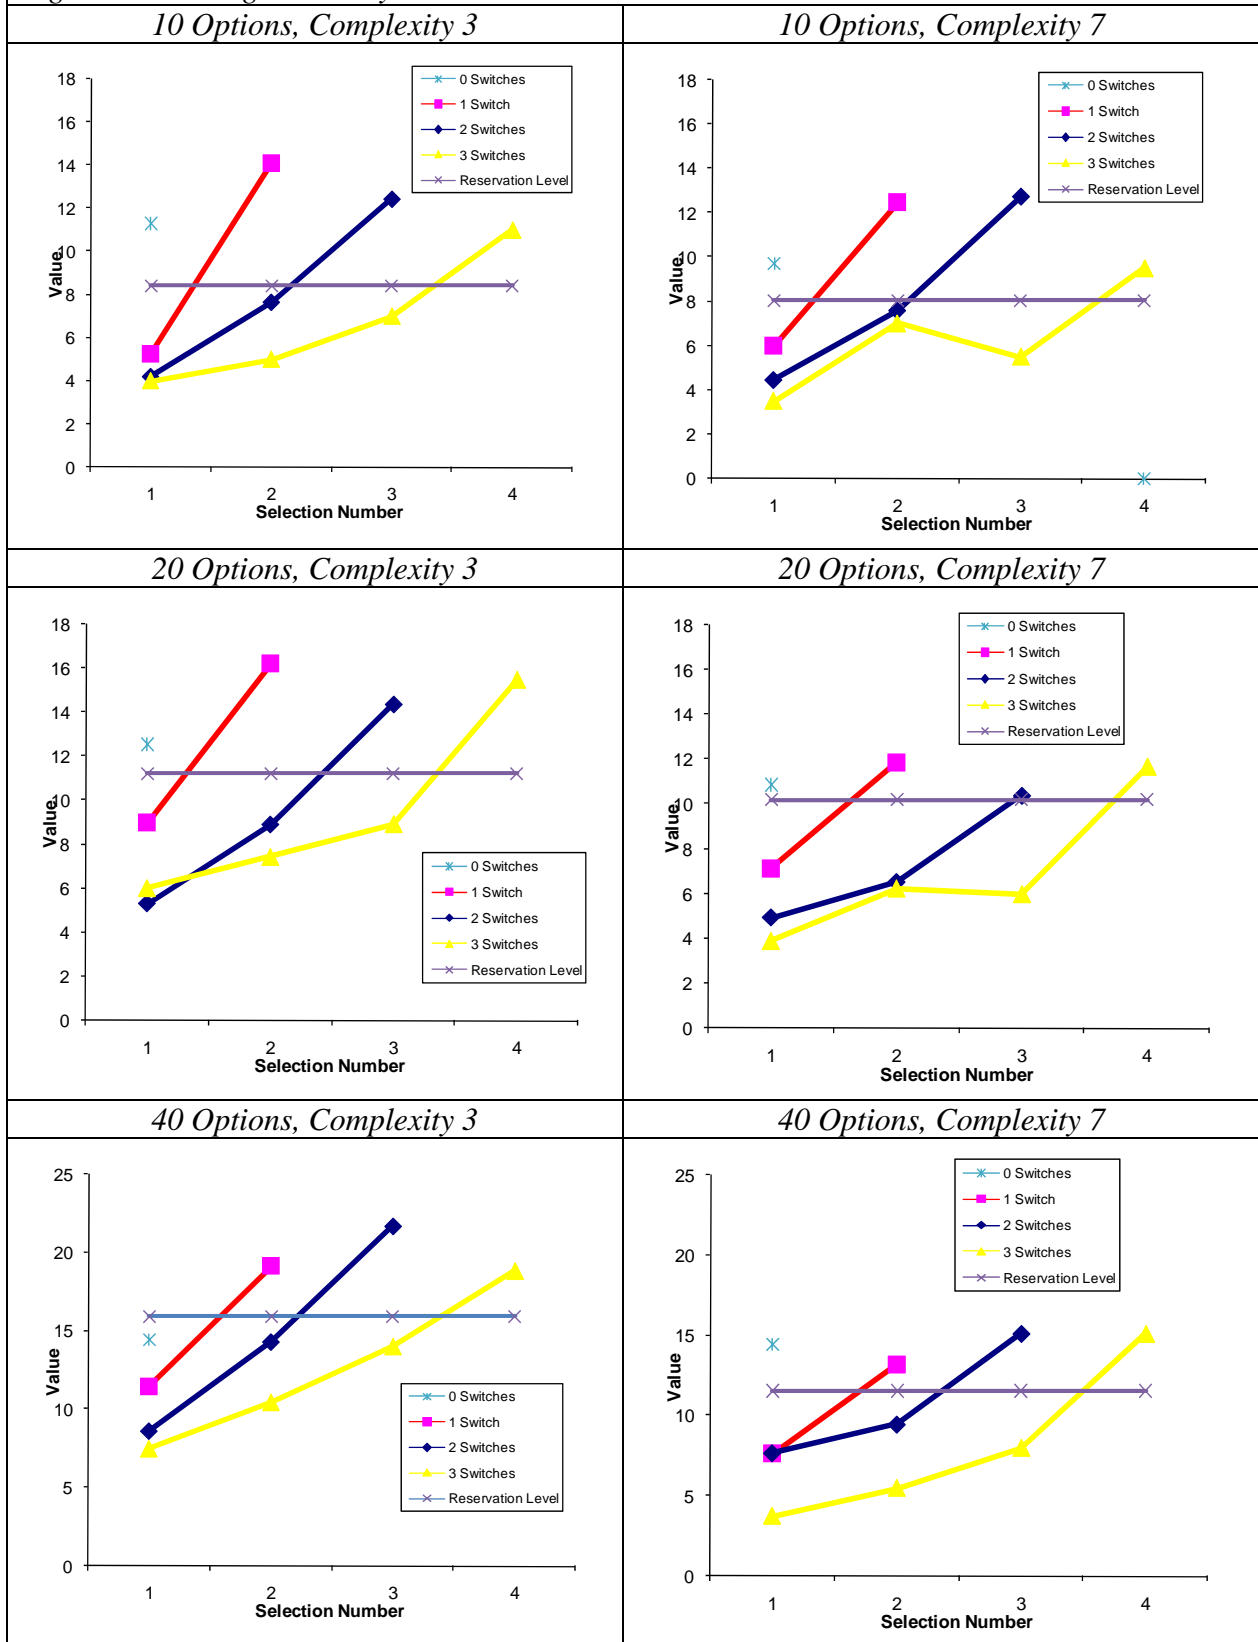


Figure 11: Screen Position and Complexity by Switch

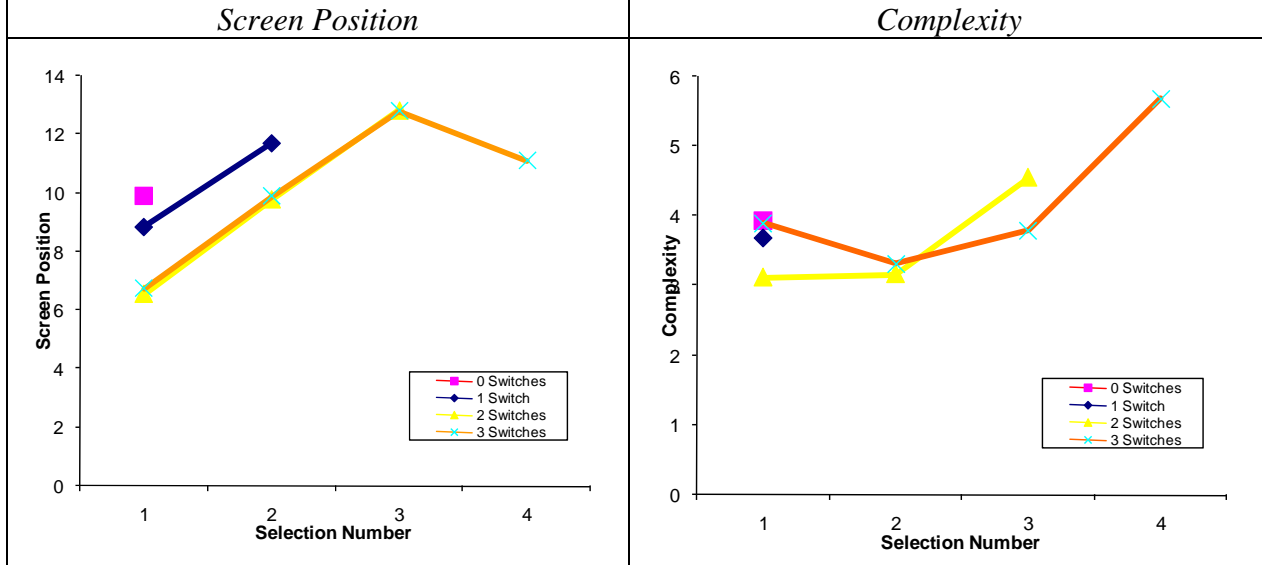


Figure 12

Panel A

Round 9 of 30 Current selection: zero

Choose one:

- zero
- four minus four plus five
- three plus two
- seven plus three plus two
- five plus nine minus four minus four
- three plus two minus eight minus four plus five plus six minus nine plus eight plus seven
- four plus zero plus two plus one minus two
- four minus ten plus zero minus one plus two plus zero plus five plus two
- six minus two minus two minus four plus four
- four minus one
- five plus four minus six plus one
- eight plus four minus three minus two plus one minus three plus three
- five plus six minus seven minus nine plus two plus five plus three minus one
- seven plus zero minus eight minus one plus five plus six minus one minus four minus two
- four plus zero plus three plus two minus two minus nine plus six
- three minus one
- four minus four minus two plus four minus ten plus seven plus three plus three plus one
- five plus zero minus four minus two plus five plus three minus five
- two
- four plus five minus four minus one minus one
- four plus one plus ten

Finished

Panel B

Round 21 of 30 Current selection: zero

Choose one:

- zero
- seven minus one
- two minus six plus seven plus three plus seven minus three minus one
- three plus eight plus one minus ten plus two
- three minus ten plus two plus five plus three plus one
- five minus one minus eight plus six plus eight minus nine plus six minus four
- eight
- four plus three minus seven plus one
- three minus four plus three
- seven minus two plus zero minus two plus two minus nine plus six plus four minus one
- three plus three plus three plus five minus five minus three plus six minus nine minus one
- eight plus one minus four minus six plus three
- eight minus one minus three minus one minus three plus four plus three
- six plus three
- five minus three plus six plus one plus one minus three minus three plus one
- five plus one minus one plus zero plus six minus five
- three plus zero plus two minus two minus three minus three plus five
- seven plus five minus eight
- seven minus four plus three minus one minus four
- four minus two minus two plus five
- five minus three plus zero

Finished

<i>Table 1: Magnitude of Mistakes (Experiment 1)</i>				
Set Size		Complexity		Total
		3	7	
10	Failure Rate (%)	6.78	23.61	16.03
	Average Loss (\$)	0.41	1.69	1.11
	Average Loss (%)	3.44	13.66	9.05
	<i>Observations</i>	59	72	131
20	Failure Rate (%)	21.97	56.06	39.02
	Average Loss (\$)	1.10	4.00	2.55
	Average Loss (%)	7.07	24.70	15.89
	<i>Observations</i>	132	132	264
40	Failure Rate (%)	28.79	65.38	46.95
	Average Loss (\$)	2.30	7.12	4.69
	Average Loss (%)	10.49	33.25	21.79
	<i>Observations</i>	132	130	262
Total	Failure Rate (%)	21.98	52.69	37.60
	Average Loss (\$)	1.46	4.72	3.12
	Average Loss (%)	7.81	25.65	16.88
	<i>Observations</i>	323	334	657

<i>Table 2: Choice Process vs. Normal Choice Data</i>				
Failure Rate				
Set Size		Complexity		Total
		3	7	
10	Choice Process	11.38	46.53	27.23
	Normal Choice	6.78	23.61	16.03
20	Choice Process	26.67	58.72	40.55
	Normal Choice	21.97	56.06	39.02
40	Choice Process	37.95	80.86	57.42
	Normal Choice	28.79	65.38	46.95
Total	Choice Process	27.26	64.14	43.66
	Normal Choice	21.98	52.69	37.60
Absolute Loss				
Set Size		Complexity		Total
		3	7	
10	Choice Process	0.42	3.69	1.90
	Normal Choice	0.41	1.69	1.11
20	Choice Process	1.63	4.51	2.88
	Normal Choice	1.10	4.00	2.55
40	Choice Process	2.26	8.30	5.00
	Normal Choice	2.30	7.12	4.69
Total	Choice Process	1.58	5.73	3.43
	Normal Choice	1.46	4.72	3.12
Number of Observations - Choice Process				
Set Size		Complexity		Total
		3	7	
10		123	101	224
20		225	172	397
40		195	162	357
Total		543	435	978

Table 3: Estimated Reservation Levels

Set Size		Complexity	
		3	7
10	ABS Types	9.03 (.19)	5.78 (.12)
	RBS Types	9.56 (.21)	5.78 (.12)
20	ABS Types	10.76 (.10)	8.85 (.09)
	RBS Types	11.22 (.11)	9.45 (.09)
40	ABS Types	14.91 (.09)	10.16 (.09)
	RBS Types	15.32 (.10)	10.57 (.09)

Table 4: Aggregate HM Indices

Set Size	Complexity	
	3	7
10	0.91	0.82
20	0.79	0.77
40	0.75	0.78

Table 5: Reservation Levels - Standard Choice Data

Set Size		Complexity	
		3	7
10	ABS Types	8.41	9.04
20	ABS Types	11.22	10.02
40	ABS Types	15.92	11.54

Table 6: HM Indices - Standard Choice Data

Set Size		Complexity	
		3	7
10	ABS Types	0.80	0.79
20	ABS Types	0.81	0.70
40	ABS Types	0.82	0.70

Table 7: Search Types

		TB Search	
		Yes	No
SC Search	Yes	7	4
	No	7	2



Instructions

In this session of this experiment, you will take part in 2 practice rounds and then 24 regular rounds.

In each round, you will be shown a group of options from which you will be asked to make a selection. Here is an example of an option:


nine minus two plus four minus three

And here is another example of an option:

five minus eleven plus thirteen minus two plus three

Options are valued at the total of the numbers shown. In this example, both options are valued at \$8 because $9 - 2 + 4 - 3 = 8$ and $5 - 11 + 13 - 2 + 3 = 8$. Options can contain up to 9 addition and subtraction operations. Feel free to use scratch paper as you do calculations.

The value of each object is independent of the value of other objects, and the value of each object is determined by drawing a random integer number between 1 and 35 from an particular distribution, which is pictured below. The possible values are on the bottom, and the probability of each value is on the left. For example, there is around a 5% chance of an object having a value of exactly \$7.



Note that the value of an option is not related to the length of an object or the number of addition or subtraction operations.

When a round begins, 4, 10, 20, or 40 options will be presented to you on the computer screen.

Each round will last up to 2 minutes. At any time, you can end a round by clicking on the 'Finished' button, but there is no penalty for using the entire 2 minutes. Clicking the 'Finished' button is the same as letting the remaining time run out with your current selection still selected. If you click the 'Finished' button, you cannot make any more changes in that round.


When the first round starts, the option that is located on the top of the list will have been selected for you. This option will give you \$0. This is the 'worst' option, in that all other options will give you more than \$0. You can change which option is selected by clicking on the button to the left of the option you want or by clicking anywhere on the option itself. You are free to change which option is selected at any time and as many times as you like.

You can highlight an option without selecting it by moving the mouse cursor over that option. The option will remain highlighted as long as the cursor stays anywhere on top of it.

After 2 minutes or whenever you click on the 'Finished' button, the first round will come to an end. At that point, a time will be picked at random from between 1 and 120 seconds (as described below). THE OPTION WHICH WAS SELECTED AT THAT TIME WILL BE RECORDED AS YOUR CHOICE FOR THAT ROUND. Therefore, **if at any point you prefer a different option to the one you have selected, you should change your selection** as this will reduce the chance of the less preferred option being recorded as your choice.

Remember, the option that is recorded as your choice is not necessarily the one selected when you click the finish button, but rather the one selected at the randomly determined time.

The time that the selected option is recorded as your choice is determined by drawing a random integer number between 1 and 120 from an particular distribution, which is pictured below. The possible times are on the bottom, and the probability of each time is on the left. For example, there is a 7% chance of your selected choice being recorded at exactly 16 seconds and around a 3% chance of your selected choice being recorded at exactly 60 seconds.



After the round has ended, you will be told which option was recorded as your choice. After a brief

pause, you will be given the opportunity to either review the instructions again on the computer screen or proceed to the second round. The second round and all following rounds will proceed exactly like the first round. This will continue until you have completed 2 practice rounds and 24 regular rounds, for a total of 26 rounds. After these rounds are complete, you will proceed to a final section of the experiment, which will have different directions.

At the end of the experiment, one of the 24 regular rounds will be picked at random, and you will be paid the value of your choice in that round.

REMEMBER:

Selecting a option

- You can select an option by clicking on the empty circle to the left of that option or by clicking anywhere on the option.
- Only one option can be selected at a time.
- Initially, the top option will be selected, which gives \$0 for sure.
- You are free to change your selected option to any other option at any time, whether or not you have picked that option previously.
- You can change the selected option as many times as you would like.
- After clicking on the 'Finished' button, the round will end, and it will be as if the selected option remained selected for the remainder of the round.

How your choice is recorded

- You should select an option as soon as you know that it is better than your currently selected option.
- After each round, a time between 1 and 120 seconds will be picked at random.
- The option that was selected at that time will be recorded as your choice.
- The option that is recorded as your choice is not necessarily the one selected when you click the finish button, but rather the one selected at the randomly determined time.
- At the end of each round, you will be told which option was recorded as your choice for that round.

How you are paid

- At the end of the experiment, we will pick one round at random, and you will be paid the value of your choice in that round. That money will be paid in addition the \$5 show-up fee and any payments from the last section of the experiment.
- At the beginning of each round, the option selected pays \$0. If this option is still selected at the random time allotted for that round, it will be recorded as your choice for that round. If this round is the one randomly picked at the end of the experiment, you will receive no money for this selection. Thus, you should move off of the \$0 option as quickly as possible.
- Choices recorded during the practice rounds will not be picked at the end of the experiment to play for money.

Next