# Decisions from Experience: Modeling Choices due to Variation in Search Strategies

**Neha Sharma (neha724@gmail.com) & Varun Dutt (varun@iitmandi.ac.in)**
Applied Cognitive Sciences Laboratory, Indian Institute of Technology Mandi, Kamand, Himachal Pradesh, India – 175005

**Abstract**

Decisions from Experience (DFE) research involves a paradigm (called, sampling paradigm), where decision-makers search for information before making a final consequential choice. Although DFE research involving the sampling paradigm has focused on accounting for information search and final choices using computational cognitive models. However, little attention has been paid to how computational models could account for final choices of participants with different information-search strategies. In this paper, we perform an individual-differences analysis and test the ability of computational models to explain final choices of participants with different search strategies. More specifically, we take an Instance-Based Learning (IBL) model, which relies on recency processes, and we calibrate this model to final choices of participants exhibiting more-switching (piecewise strategy) or less-switching (comprehensive strategy) between options in different problems. Our results indicate more reliance on recency of information among participants exhibiting piecewise strategy compared to comprehensive strategy. Overall, the IBL model calibrated to individual participants using a single set of parameters could account for both piecewise and comprehensive strategies. We highlight the implications of our results for DFE research involving information search before consequential decisions.

**Keywords:** information search; experience; search strategy; computational cognitive models; Instance-Based Learning Theory; multi-arm bandit problems.

## Introduction

In words of famous philosopher Plato, a good decision is based on knowledge and not numbers (Stutman & Kevin, 2015). Knowledge can be obtained by searching the environment for information before making consequential decisions. For example, investment decisions are likely to be affected by an investor's previous knowledge of a company's stocks (Subramanyam, 2016). An investor could invest in a wide range of companies in the stock market. To ensure a good decision, one must gather information about various returns offered by different stocks before making a consequential choice for a company's stocks. While gathering information, some people may explore the prices of a company's stock repeatedly before switching to a different company's stock (comprehensive strategy). However, some people may explore prices of a company's stock once and then switch to exploring the stock prices of a different company (piecewise strategy). In both cases, it is important to investigate how influential computational cognitive models account for consequential choices among both kinds of search strategies. This investigation is the main goal of this paper.

The act of making choices based on information search is a common exercise involving people in different facets of their daily life (choosing smartphones, choosing TV channels etc.). In fact, information search before a choice constitutes an integral part of Decisions from Experience (DFE) research, where the focus is on explaining human maximizing decisions based upon one's experience with sampled information (Hertwig & Erev, 2009). To study people's search and choice behaviors in the laboratory, DFE research has proposed a "sampling paradigm" (Hertwig & Erev, 2009).

In the sampling paradigm, people are presented with two or more options to choose between. These options are represented as blank buttons on a computer screen. People are first asked to sample as many outcomes as they wish and in any order they desire from different button options (information search). This sampling of information among different options is costless. Once people are satisfied with their sampling of options, they decide from which option to make a single final (consequential) choice for real.

Hills and Hertwig (2010) have analyzed the search strategies of people asked to make choices in the sampling paradigm. Hills and Hertwig (2010) report two search strategies prevalent among participants: comprehensive and piecewise. In the comprehensive strategy, people search one option repeatedly before switching to the other option. In contrast, in the piecewise strategy, people search for one option once and then switch to the other option. They sample the other option once and again switch back to the first option, searching for information in a zigzag manner.

Computational cognitive models of human choice behavior have thus far predicted choices at an aggregate level in the sampling paradigm, i.e., when people's final choices are averaged over several participants (Busemeyer & Wang, 2000; Gonzalez & Dutt, 2012; Lejarraga, Dutt, & Gonzalez, 2012). For example, the Instance-Based Learning (IBL) model is a popular DFE algorithm for explaining aggregate choices (Erev et al., 2010; Gonzalez & Dutt, 2011; Lejarraga, Dutt, & Gonzalez, 2012; Hertwig, 2012). The IBL model (Gonzalez & Dutt, 2011) consists of experiences (called instances) stored in memory. Each instance's activation is used to calculate the blended values for each option, thereby helping the model make a final choice. The IBL model relies on ACT-R framework for its functioning (Anderson & Lebiere, 1998).

Prior DFE research has shown that, at the aggregate level, the IBL model exhibits superior performance compared to other computational models in the sampling paradigm (Erev et al., 2010; Gonzalez & Dutt, 2011).

Although computational cognitive models have been evaluated at the aggregate level; yet, less attention has been paid to the evaluation of models in their ability to account for individual differences, especially in terms of search strategies. Given that people exhibit two specific search strategies (Hills and Hertwig, 2010), comprehensive and piecewise, it would be interesting to see how computational cognitive models with a set of parameters account for consequential choices for participants exhibiting these strategies.

In this paper, our main goal is to evaluate how computational cognitive models, which explain choice behavior at the aggregate level (e.g., IBL model), perform in capturing consequential decisions of participants exhibiting different search strategies with single set of parameters. For this purpose, we use risky problems involving two options and outcomes with different probabilities (rare events and common events). We calibrate an IBL model, which was evaluated in prior research at the aggregate level, to preferences of participants showing different search strategies. In what follows, we detail the problems used and the working of the IBL model. Then, we discuss the methodology of calibrating the IBL model to consequential decisions in different problems. Next, we present the results of model evaluation and the role of recency and frequency mechanisms in accounting for consequential decisions involving different search strategies. We close the paper by discussing the implications of our results for DFE research in the sampling paradigm.

## Problem Dataset

Eighty students at Indian Institute of Technology Mandi, India, participated in a study where the objective was to evaluate participant preferences for options after information search. The study involved the sampling paradigm, where participants searched for information and then decided an option they preferred across two between-subjects problem conditions: Rare-Event (RE; N = 40) and Common-Event (CE; N = 40). In the CE problem, a variable option had a high probability (0.8) value associated with a high (H) outcome (1.18 return on the allocated amount); whereas, in the RE problem, the variable option had a low probability (0.1) associated with the H outcome (3.28 return on the allocated amount). Across both problems (CE and RE), the low (L) outcome (0.88) in the variable option always occurred with a complementary chance. An alternative with a fixed return on investment (1.1 return on the invested amount with certainty) was present in both RE and CE conditions as second option. Thus, in each problem, participants were presented with two options: an option with a fixed return on allocation (non-maximizing option); and, an option with a variable return on allocation (maximizing option). The maximization was defined based upon the expected value of options in problems. The nature of outcomes and probabilities in different CE and RE problems were like those described in Hertwig et al. (2004).

In each problem, participants were first asked to sample options (presented as blank buttons; sampling phase). During the sampling phase, every time an option was chosen in a problem, participants could see an outcome based upon the associated probability in the option. Sampling of options was costless in the sampling phase and participants were free to sample options in any order and as many times as they desired. At any time during the sampling phase, participants could click the “Make a Final Decision” button. Clicking this button terminated the sampling phase and moved participants to the final-decision phase. In the final-decision phase, participants were asked to make a final choice for one of the options for real.

To understand the effect of different sampling strategies, we calculated the switch ratio, which was defined as the total number of switches made by a participant between options divided by the total number of switches possible (= number of samples – 1). Like done by Hills and Hertwig (2010), we calculated the median value of switch ratio by pooling participants across both CE and RE problems. Participants possessing switch-ratios less than median were classified as following comprehensive search strategy (called LM) and participants possessing switch-ratios greater than or equal to median were classified as following piecewise strategy (called GM). By pooling the CE and RE problems, there were N = 40 participants in the LM group and N = 40 participants in the GM group.

## Human Results

Figure 1 shows proportion of final choices by human participants in the GM and LM condition. As seen in the Figure, the pattern of preferences across problems in the LM and GM conditions was similar in human data: higher allocation to the fixed option compared to the variable option. Next, we consider whether an IBL model can account for these effects via its cognitive mechanisms (model results will be described in a future section).

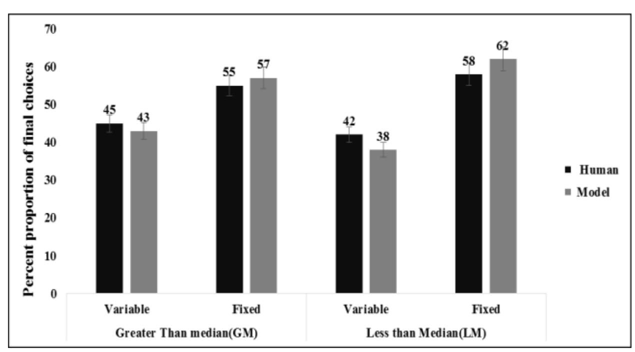


Figure 1: Percentage Proportion of final choices in each option by human and model for both GM and LM condition.

## The Model

In this section, we detail the working of a model based upon Instance-Based Learning Theory (IBL model; Gonzalez & Dutt, 2011; 2012), which was calibrated to LM and GM search strategy groups separately.

## Instance-Based Learning (IBL) Model

The IBL model (Dutt & Gonzalez, 2012; Gonzalez & Dutt, 2011; 2012; Lejarraga, Dutt, & Gonzalez, 2012) is based upon the ACT-R cognitive framework (Anderson & Lebiere, 1998). In this model, every occurrence of an outcome of an option is stored in the form of an instance in memory. An instance is made up of the following structure: SDU, here S is the current situation (two blank option buttons on a computer screen), D is the decision made in the current situation (choice for one of the option buttons), and U is the goodness (utility) of the decision made (the outcome obtained upon making a choice for an option). When a decision choice needs to be made, instances belonging to each option are retrieved from memory and blended together. Blended value of an option is a function of activation of instances corresponding to outcomes observed on the option. Activation of an instance is a function of the frequency and recency of observed outcomes that occur on choosing options during sampling. The blended value of option *j* at any trial *t* is defined as (Lebiere, 1999):

$$V_{j,t} = \sum_{i=1}^{n} p_{i,j,t} x_{i,j,t} \qquad (1)$$

where $x_{i,j,t}$ is the value of the U (outcome) part of an instance *i* on option *j* at trial *t*. The $p_{i,j,t}$ is the probability of retrieval of instance *i* on option *j* from memory at trial *t*. Because $x_{i,j,t}$ is value of the U part of an instance *i* on option *j* at trial *t*, the number of terms in the summation changes when new outcomes are observed within an option *j* (and new instances corresponding to observed outcomes are created in memory). Thus, $n = 1$ if *j* is an option with one possible outcome. If *j* is an option with two possible outcomes, then $n = 1$ when one of the outcomes has been observed on an option (i.e., one instance is created in memory) and $n = 2$ when both outcomes have been observed (i.e., two instances are created in memory).

At any trial *t*, the probability of retrieval of an instance *i* on option *j* at trial *t* is a function of the activation of that instance relative to the activation of all instances (1, 2, … *n*) created within the option *j*, given by

$$p_{i,j,t} = \frac{e^{(A_{i,j,t})/\tau}}{\sum_{i=1}^{n} e^{(A_{i,j,t})/\tau}} \qquad (2)$$

Where τ, is random noise defined as $\sigma.\sqrt{2}$ and σ is a free noise parameter. Noise captures the imprecision of recalling past experiences from memory. The activation of an instance *i* corresponding to an observed outcome on an option *j* in each trial *t* is a function of the frequency of the outcome's past occurrences and the recency of the outcome's past occurrences (as done in ACT-R). At each trial *t*, activation $A_{i,j,t}$ of an instance *i* on option *j* is

$$A_{i,j,t} = \sigma * \ln\left(\frac{1-\gamma_{i,j,t}}{\gamma_{i,j,t}}\right) + \ln \sum_{t_p \varepsilon \{1..,t-1\}} \left(t - t_p\right)^{-d} \qquad (3)$$

where *d* is a free decay parameter; $\gamma_{i,j,t}$ is a random draw from a uniform distribution bounded between 0 and 1 for instance *i* on option *j* in trial *t*; and $t_p$ is each of the previous trials in which the outcome corresponding to instance *i* was observed in the binary-choice task. The IBL model has two free parameters that need to be calibrated: *d* and σ. The *d* parameter controls the reliance on recent or distant sampled information. Thus, when *d* is large (> 1.0), then the model gives more weight to recently observed outcomes in computing instance activations compared to when *d* is small (< 1.0). The σ parameter helps to account for the sample-to-sample variability in an instance's activation. In the IBL model, we feed the sampling of individual human participants to generate instance activations and blended values. Every time a choice is made and outcome is observed, the instance associated with it is activated and thereafter blended values are computed for options faced by an individual participant.

In one version of the IBL model, we use parameters suggested by Lejarraga, Dutt, and Gonzalez (2012) to test the model's ability in capturing final choices for different search strategy groups, LM and GM. In a second version of the model, we found values for the *d* and σ parameters by calibrating these parameters to final choices from human participants separately in the two strategy groups. For this calibration, we determine the model's likelihood for making the same choice as made by each human participant given a set of model parameters.

For each model participant, the model applied the following softmax function across both options in a problem (Bishop, 2006; Sutton & Barto, 1998):

$$Prob(Option\ X) = \frac{e^{S_{Mean}X}}{e^{S_{Mean}X} + e^{S_{Mean}Y}} \qquad (4)$$

Where, $S_{Mean}X$ and $S_{Mean}Y$ are the blended values calculated for the two options and *Prob (Option X)* is the probability of choosing Option X given a set of model parameters (also, called the "likelihood"). If *Option X* was chosen by a human participant in a problem, then the *Prob (Option X)* is used to calculate the likelihood value of making the same choice from the IBL model given its set of parameters. The log-likelihood *L* is defined as:

$$L = \sum_{i=1}^{i=N} \log(Prob(Option\ X_i)) \qquad (5)$$

Where, *i* refers to the *i*th model participant playing a problem and *N* is the total number of human participants in the LM and GM groups (the model was calibrated separately to each of the two switching groups). The *log* in equation 5 is the natural logarithm and we calibrated the IBL model by minimizing the negative of the log-likelihood value (- *L*).

Furthermore, to derive a choice from the IBL model, we use the following rule: If the human chose Option X and the value of *Prob (Option X)* is greater than or equal to 0.5, then the model makes a choice like the human choice; else, the model chooses the option that is opposite of what human participant chose. We calculated the error proportion by comparing the model participant's choice to the human participant's choice.

## Method

### Dependent Variables

In this paper, we account for the final choices of participants with different search strategies. For this purpose, a choice made by a model participant is evaluated against a choice made by a corresponding human participant in either of the LM and GM groups, separately.

A choice in a problem is classified as maximizing if the chosen option's expected value is greater than the expected value of the non-chosen option. Those cases for which this criterion failed were termed as having non-maximizing choice. The expected value of an option was calculated by multiplying the probability of occurrence of outcomes with the outcomes and summing the multiplications together. For a model, the error proportion was calculated in a problem as:

$$ErrorProportion = (M_H N_M + N_H M_M)/(M_H N_M + N_H M_M + N_H N_M + M_H M_M) \quad \ldots (6)$$

Where, $M_H N_M$ was the number of cases where the human participant made a maximizing choice but the model predicted a non-maximizing choice. $N_H M_M$ was the number of cases where the human participant made a non-maximizing choice but the model predicted a maximizing choice. Similarly, the $M_H M_M$ and $N_H N_M$ were the number of cases, where the human participant made the same choice (maximizing or non-maximizing) as predicted by the model. Smaller the value of the error proportion, the more accurate the model is in accounting for maximizing individual choices of human participants.

### Model Calibration

The IBL model described here had two free parameters *d* and *σ*. The model was calibrated on final choices for both groups, GM and LM, using a genetic algorithm program. A single set of parameters were used to calibrate the model by minimizing the negative of the Log-Likelihood value. The genetic algorithm has features that help prevent the algorithm getting trapped in local minima. The genetic algorithm repeatedly modifies a population of individual parameter tuples to find the tuple that minimizes -*L*. In each generation, the genetic algorithm selects individual parameter tuples randomly from a population to become parents and uses these parents to select children for the next generation. Over successive generations, the population evolves toward an optimal solution. The population size used here was a set of 20 randomly-selected parameter tuples in a generation (each parameter tuple was a value of *d* and σ parameters). The mutation and crossover fractions were set at 0.1 and 0.8, respectively, for an optimization over 150 generations. The model was calibrated separately in the LM and GM groups. Within each group, for each parameter tuple, the model was run 10-times across participants in a problem and the average –*L* value across 10-runs was minimized. The 10-runs ensured that the run-to-run variability in the –*L* value was small and the 10 value was derived after trying different integer values between 1 and 20 runs.

## Model Results

First, we evaluated the IBL model's ability to account for final choices in the GM group. The best calibrated values of *d* and σ parameters in the IBL model were found to be 15.05 and 0.29, respectively (see Table 1). The large *d* value exhibited extreme reliance on recency during sampling. Also, the smaller σ value exhibited lower sample-to-sample variability in instance activations. The lowest value of log-likelihood obtained during calibration was -25.19.

Table 1: Parameters and Likelihood Values

| Condition | Parameters | Log-Likelihood |
|---|---|---|
| **GM** | d=15.05 | -25.19 |
| | σ=0.29 | |
| **LM** | d=8.82 | -29.03 |
| | σ=0.73 | |
| **GM-LDG** | d=5.0 | -127.07 |
| | σ=1.5 | |
| **LM-LDG** | d=5.0 | -106.33 |
| | σ=1.5 | |

The parameters obtained from the IBL model for the LM group were *d* = 8.82 and σ = 0.73. The value of *d* in the LM group again made participants rely on recency of information during sampling; however, this reliance on recency processes was less than that for the GM group. Furthermore, the noise parameters value represented lesser variability in activations across samples. Overall, the calibrated likelihood value was -29.03, which was slightly lesser than that in the GM group. Furthermore, the calibration of IBL model to both LM and GM groups resulted in improved likelihoods compared to the parameters suggested by Lejarraga, Dutt, and Gonzalez (2012) (*d* = 5; σ =1.5). The model parameters fitted using log-likelihoods by us in this paper are for individual participant choices in the two groups, LM and GM. However, the model parameters fitted by Lejarraga, Dutt, and Gonzalez (2012) were for choices aggregated across several participants. Given the high values of *d* parameter in our results, it seems that the recency and frequency processes are stronger among individual participants compared to the average across several participants.

Figure 1 shows proportion of final choices by model participants compared to human data in the GM and LM conditions. In both conditions, the IBL model performed like human participants: The model showed greater preferences for the fixed option compared to the variable option in both GM and LM conditions. The model's preference for fixed (variable) option was slightly higher (lower) compared to those for human participants. Due to recency effect, the model's account for human preferences was better for those who switched more and followed the piecewise search strategy compared to participants who switched less and followed the comprehensive strategy. Thus, perhaps, recency processes were more prevalent among the piecewise strategy group compared to the comprehensive strategy group.

Lastly, we analyzed the IBL model's performance in accounting for individual decisions. According to error proportion criterion, more number of $N_HN_M$ and $M_HM_M$ combinations help minimize the error proportion (which is desirable), while higher number of $M_HN_M$ and $N_HM_M$ combinations increase the error proportion. Table 2 shows the individual-level results from different LM and GM groups. As seen in Table 1, the calibrated IBL model for GM group produced 55% of $N_HN_M$ combinations and 30% of $M_HM_M$ combinations, respectively. In contrast, the erroneous $N_HM_M$ and $M_HN_M$ combinations were 12 % and 3%, respectively, from the model. Due to comparatively higher values for the $N_HN_M$ and $M_HM_M$ combinations in the GM group compared to the LM group, the IBL model possessed smaller error proportion in the GM group compared to the LM group. Overall, the IBL model showed superior performance for GM group compared to the LM group (15% error proportion < 32% error proportion). Finally, the error proportions from models fitted in this paper at the individual participant level were comparatively less compared to the error proportions from models fitted to the aggregate data by Lejarraga, Dutt, and Gonzalez (2012) (LDG model in GM and LM groups). Thus, it seems that fitting models using individual choices makes such models perform better compared to when the same models are fitted using aggregate choices.

Table 2: The error proportions from IBL model in the LM and GM groups

| Human and Model data combination H/M | GM | LM | GM (LDG) | LM (LDG) |
|---|---|---|---|---|
| No. of Observations | 40 | 40 | 40 | 40 |
| $N_HN_M$ | 55 | 45 | 22 | 33 |
| $M_HM_M$ | 30 | 23 | 13 | 20 |
| $N_HM_M$ | 12 | 15 | 45 | 27 |
| $M_HN_M$ | 03 | 18 | 20 | 20 |
| Error Proportion | 0.15 | 0.32 | 0.65 | 0.47 |

## Discussion & Conclusions

So far, models in decisions from experience (DFE) paradigms had been evaluated to aggregate human choices (Gonzalez & Dutt, 2011; 2012). In such comparisons, the average risk-taking from the model was compared to the average risk-taking from human data. However, in this paper, we compared a model's performance by calibrating the model to individual human choices. More specifically, we calibrated an Instance-Based Learning (IBL) model to individual preferences with different information-search strategies. Overall, the IBL model showed superior performance when calibrated to both search-strategy groups, piecewise and comprehensive. The high value of decay parameter showed stronger reliance on recency processes among individual participants. In fact, the recency effect was stronger among participants who switched more and followed the piecewise search strategy compared to participants who switched less and followed the comprehensive strategy.

One likely reason for differing recency effect among different search strategy is that when participants use the piecewise strategy, they tend to compare the most recent outcome on one option with the most recent outcome on the other option. For this comparison to work, participant needs to rely on recent information. Furthermore, this comparison is less prevalent in the comprehensive strategy, where participants tend to search one option repeatedly before moving to investigate the other option.

In fact, the observation about high *d* parameter value for the piecewise strategy also helps us explain why the error proportion from the model was much less for the piecewise strategy compared to the comprehensive strategy. That is because recency is more suited to piecewise strategy compared to comprehensive strategy.

In this paper, we took one model of experiential choice; however, as part of future research, we plan to extend this investigation to a larger set of models and application areas. Also, it would be interesting to investigate how recency effects explain choices among different search strategies in environments where the outcomes and probabilities are non-stationary and change overtime. Some of these ideas and others form the immediate next steps for us to pursue in the near future.


## Acknowledgments

The authors are grateful to Indian Institute of Technology, Mandi for providing computational resources for this paper. Also, the authors would like to thank Tata Consultancy Services' research fellowship to Neha Sharma for supporting this project.


## References


Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Hillsdale, NJ: Lawrence Erlbaum

Associates. ISBN 0-8058-2817-6
Busemeyer, J. R., & Wang, Y. (2000).Model comparisons and model selections based on the generalization criterion methodology. *Journal of Mathematical Psychology*, 44, 171–189.
Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.
Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., & Hau, R.(2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Mak*ing, 23, 15–47.
Gonzalez, C., & Dutt, V. (2011). Instance-Based Learning: Integrating Sampling and Repeated Decisions From Experience. *Psychological Review*,118(4), 523-551.
Gonzalez, C., & Dutt, V. (2012). Refuting data aggregation arguments and how the instance-based learning model stands criticism: A reply to Hills and Hertwig *Psychological Review*, *Vol* 119(4),893-898.
Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, 13, 517-523.
Hills, T. T., & Hertwig, R. (2010). Information Search in Decisions From Experience Do Our Patterns of Sampling Foreshadow Our Decisions?. *Psychological Science*.
Hertwig, R. (2012). The psychology and rationality of decisions from experience. *Synthese*, 187, 269-292.
Lebiere, C. (1999). Blending: An ACT–R mechanism for Aggregate retrievals. *Paper presented at the 6th Annual ACT–R Workshop at George Mason University*. Fairfax County, VA.
Lejarraga, T. & Dutt, V. & Gonzalez, C. (2012). Instance-Based Learning: A general model of repeated binary choice. *Journal of Behavioral Decision Making*,25: 143-153.
Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1, No. 1). Cambridge: MIT press.
Stutman, M., & Conklin, K. (2015). *The Ultimate Book of Inspiring Quotes For Kids.* CreateSpace Independent Publishing Platform.
Subramanyam, P.V. (2016, Nov 21). 8 biases that may lead to bad investment decisions. *Economic Times.* Retrieved from http://economictimes.indiatimes.com/wealth/invest/8-biases-that-may-lead-to-bad-investment-decisions/articleshow/55508378.cms