

American Economic Association

A Boundedly Rational Decision Algorithm

Author(s): Xavier Gabaix and David Laibson

Source: *The American Economic Review*, Vol. 90, No. 2, Papers and Proceedings of the One Hundred Twelfth Annual Meeting of the American Economic Association (May, 2000), pp. 433-438

Published by: [American Economic Association](#)

Stable URL: <http://www.jstor.org/stable/117264>

Accessed: 03/10/2011 12:16

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



American Economic Association is collaborating with JSTOR to digitize, preserve and extend access to *The American Economic Review*.

<http://www.jstor.org>

A Boundedly Rational Decision Algorithm

By XAVIER GABAIX AND DAVID LAIBSON*

Cognition requires scarce inputs, including time and concentration. Since cognition is costly, sophisticated decision-makers should use mental short-cuts, or heuristics, to reduce cognitive burdens. We propose and test a model motivated by these principles.

We base our model on algorithms that simplify decision trees. When cognition is costly, the actor will not calculate perfectly, relying instead on a simplification of the tree. The algorithms are parameterized to include perfect rationality (i.e., zero cognition costs) as a special case. We focus on one algorithm that simplifies trees by removing branches with low probability. This reduces to a psychologically natural forward simulation rule.

We believe that this model achieves four goals. First, the model makes quantitative behavioral predictions and, hence, provides a precise alternative to the rational-actor hypothesis. Second, the model is psychologically plausible because it is based on the actual decisions algorithms that subjects claim to use. Third, the model is empirically testable; such a test is provided in this paper. The data overwhelmingly reject the rational model in favor of our bounded-rationality alternative. Fourth, the model is broadly applicable, because it can be used to analyze decision problems that can be represented in tree form.

Our approach combines four strands in the

existing literature on bounded rationality. First, our conceptual approach extends the satisficing literature which was pioneered by Herbert Simon (1955) and subsequently formalized with models of deliberation cost.¹ Second, our approach extends the heuristics literature started by Daniel Kahneman and Amos Tversky (1974). Third, our approach is based on work by Colin Camerer et al. (1994) which shows that decision-makers often solve problems by looking forward, rather than using backward induction. Fourth, our emphasis on empirical evaluation is motivated by the work of Ido Erev and Alvin Roth (1998) and Camerer and Teck-Hua Ho (1999). Our approach differs from theirs since we study first-period play in non-strategic settings instead of learning dynamics in repeated games.

I. Decision Trees and an Algorithm

We propose an algorithm that mimics the way that decision-makers evaluate trees. We base our analysis on decision trees since trees can be used to represent a wide class of problems.

Consider the tree in Figure 1. Each starting box in the left-hand column leads to boxes in the second column. The numbers inside the boxes are flow payoffs. Probabilistic branches connect boxes. For example, the first box in row five contains a flow payoff of 3, and five branches exit from this box, with probabilities 0.55, 0.1, 0.15, 0.1, and 0.1. Starting from row five, there exist 95 possible outcome paths. The path that follows the highest-probability branch at each node is $\langle 3, -1, 5, 4, -5 \rangle$. Integrating over all 95 paths, the expected payoff of starting from row five is 5.04.

Imagine that a decision-maker must choose one of the boxes in the first column of Figure 1. The decision-maker will be paid the expected value associated with his chosen box. What box

* Gabaix: Department of Economics, Massachusetts Institute of Technology, Cambridge, MA 02142 (e-mail: xgabaix@mit.edu); Laibson: Department of Economics, Harvard University, Cambridge, MA 02138, and NBER (e-mail: dlaibson@harvard.edu). We are grateful for the advice of Gray Becker, Ed Glaeser, Lars Hansen, Larry Katz, George Loewenstein, Sendhil Mullainathan, Matthew Rabin, Al Roth, Andrei Shleifer, Richard Thaler, and Richard Zeckhauser. Numerous research assistants helped implement and analyze the experiment; we are particularly indebted to Meghana Bhatt, David Hwang, Guillermo Moloche, and Stephen Weinberg. Gabaix acknowledges the hospitality of the University of Chicago. Laibson acknowledges financial support from the National Science Foundation, the National Institute of Health, and the MacArthur Foundation.

¹ See John Conlisk (1996) for a review.




FIGURE 1. DECISION TREE

would you choose?² We propose a heuristic, dubbed “follow the leaders” (FTL), for evaluating options in such trees. FTL has one parameter, a cutoff probability p :

From any starting box b follow all branches that have a probability greater than or equal to p . Continue in this way, moving from left to right across the columns in the decision tree. If a branch has a probability less than p , consider the box to which the branch leads but do not advance beyond that box. Weight all boxes that are considered with their associated cumulative probabilities, and calculate the weighted sum of the boxes.

² The four most frequent picks are rows 10 (42 percent), 5 (20 percent), 3 (16 percent), and 4 (12 percent). The respective expected values are \$5.52, \$5.04, \$3.20, and \$5.54.

This heuristic describes the FTL value of a starting box b . In our implementation we set $p = 0.25$; perfect rationality with zero cognition cost corresponds to $p = 0$.

To make FTL precise, call a *path* a sequence of boxes and associated branch/transition probabilities. Abusing notation, let $\langle 3, (0.15, 2), (0.5, -1), (0.15, -1), (1.0, 4) \rangle$ represent a path originating in the fifth row of Figure 1. A p -constrained path $\langle u_0, (p_1, u_1), \dots, (p_N, u_N) \rangle$ has the following properties: (i) all probabilities, except possibly the last one, are greater than or equal to p , and (ii) the path cannot be lengthened without violating (i). So, if $p = 0.25$, then the path described above is not a p -constrained path, but $\langle 3, (0.15, 2) \rangle$ is a p -constrained path. Since $0.15 < p$ the p -constrained path does not extend past the branch with probability 0.15. For any path, define a cumulative probability $\Pi = \prod_{t=1}^N p_t$, and cumulative payoff $U = \sum_{t=0}^N u_t$.

Finally, to get the FTL value associated with a starting box b ,³ identify the set of p -constrained paths starting with box b , the associated set of cumulative probabilities, $\{\Pi_k\}_{k=1}^K$, and cumulative payoffs, $\{U_k\}_{k=1}^K$. The FTL value of this box b is: $V_b^{FTL} = \sum_{k=1}^K \Pi_k U_k$. Ten p -constrained paths originate in row five of Figure 1.⁴ The FTL value of row five is 4.91, which is not far from the true expected value of 5.04. In this example, FTL dramatically simplifies the decision tree. Ninety-five paths originate in the fifth row, but application of FTL reduces consideration to only ten paths, five of which are truncated.

We chose FTL because it corresponds to our intuition about how decision-makers “see” decision trees that do not contain outlier payoffs. Subjects follow paths that have high-probability branches at every node. In the process of determining which branches to follow, subjects see low-probability branches that branch off of high-probability paths. Subjects encode such

³ For simplicity we describe the algorithm for payoffs with mean 0. The generalization to nonzero means is straightforward.

⁴ Namely, $\langle 3, (0.1, -3) \rangle$, $\langle 3, (0.15, 2) \rangle$, $\langle 3, (0.1, -3) \rangle$, $\langle 3, (0.1, -2) \rangle$, $\langle 3, (0.55, -1), (0.55, 5), (0.15, 4) \rangle$, $\langle 3, (0.55, -1), (0.55, 5), (0.85, 4), (0.55, -5) \rangle$, $\langle 3, (0.55, -1), (0.55, 5), (0.85, 4), (0.45, 2) \rangle$, $\langle 3, (0.55, -1), (0.45, -1), (1.0, 4), (0.33, -5) \rangle$, $\langle 3, (0.55, -1), (0.45, -1), (1.0, 4), (0.57, 2) \rangle$, $\langle 3, (0.55, -1), (0.45, -1), (1.0, 4), (0.1, 4) \rangle$.

low-probability branches, but subjects do not follow them deeper into the tree. FTL is only one among many sensible algorithms that subjects use to evaluate decision trees. There may be close variants that better predict subject choices. We analyze FTL to demonstrate that bounded-rationality models can be successfully empirically implemented.

In the next section, we empirically compare FTL to the rational-actor model and two other choice models, which we call the “column-cutoff model” and the “discounting model.” The rational-actor model with zero cognition cost is the standard assumption in almost all economic models. This corresponds to the limit case $p = 0$. For presentational simplicity, we refer to this benchmark as the *rational-actor model*, but we emphasize that this standard model assumes both rationality and zero cognition costs. The *column-cutoff model* assumes that decision-makers calculate perfectly rationally but only pay attention to the first Q columns of the tree, completely ignoring the remaining columns. The *discounting model* assumes that decision-makers follow all paths, but exponentially discount payoffs, using discount factor δ , according to the column in which those payoffs arise.

Anticipating these tests, we need to close the model by providing a theory that translates rational and quasi-rational payoff evaluations into choices. Assume that there are B possible choices with evaluations $\{V_1, V_2, \dots, V_B\}$ given by a candidate model. Then the probability of picking box b is

$$(1) \quad P_b = \frac{\exp(\beta V_b)}{\sum_{b'=1}^B \exp(\beta V_{b'})}$$

We estimate the parameter β in our econometric analysis.

II. Experimental Design

We tested the model using data from 259 Harvard undergraduates. Subjects were guaranteed a payment of \$7 and earned more if they performed well; the average payment was \$20.08.

A. Decision Trees

The experiment was based around 12 randomly generated trees, one of which is reproduced in Figure 1. We chose large trees, because we wanted undergraduates to use heuristics to “solve” these problems. Half of the trees have 10 rows and five columns of payoff boxes; the other half of the trees are 10×10 . The branching structure, probabilities, and payoffs are independently randomly generated.

Experimental instructions described the concept of expected value to the subjects. Subjects were told to choose one starting row from each of the 12 trees. Subjects were told that one of the trees would be randomly selected and that they would be paid the true expected value for the starting box that they chose in that tree. Subjects were given a maximal time of 40 minutes to make their choices on all 12 trees.

B. Debriefing Form and Expected-Value Calculations

After passing in their completed tree forms, subjects filled out a debriefing form which asked them to describe their choice algorithm and their background information. Subjects were then asked to solve 14 simple expected-value problems (e.g., calculate the expected value of the starting boxes in a 2×3 tree). To eliminate subjects who did not understand the concept of expected value, our econometric analysis excludes subjects who solved fewer than half of these simple problems.⁵ Out of the 251 subjects who provided answers for all of the 12 decision trees, 230, or 92 percent, answered correctly at least half of the simple expected-value questions. Out of this subpopulation, the median score was 12 out of 14 correct.

III. Results and Analysis

Had subjects chosen randomly, they would have chosen starting boxes with an average payoff of \$1.30.⁶ Had the subjects chosen the starting boxes with the highest payoff, the

⁵ This restriction did not affect our results.

⁶ In theory this would have been zero, since our games are randomly drawn from a mean-zero distribution. Our realized value is within the two-standard-error bands.

average chosen payoff would have been \$9.74. In fact, subjects chose starting boxes with an average payoff of \$6.72.

We are interested in comparing four different classes of models: rationality, FTL, column-cutoff, and discounting. All of these models have a nuisance parameter, β , which captures the tendency to pick starting boxes with the highest evaluations [see equation (1)]. For each model, we estimate a different β parameter. FTL, column-cutoff, and discounting, also require an additional parameter.

FTL requires the probability cutoff value p . We exogenously set $p = 0.25$. We did not experiment with other values. We choose 0.25 because it corresponds to our intuition about which paths could feasibly be considered given the 40-minute time constraint on the decision process. Because we do not believe that either column-cutoff or discounting are psychologically plausible models, we do not impose a single value on either Q or δ . Instead, we evaluate the models for the full range of Q and δ values. This analysis is provided for comparison to our preferred model, FTL.

Formal Model Testing

There are $G = 12$ games, $B = 10$ starting boxes per game, and $I = 230$ subjects who satisfy our inclusion criterion. Call $\hat{P} \in \mathbb{R}^{BG}$ the empirical distribution of choices (i.e., \hat{P}_{bg} = the fraction of players that played box b in game g). A model m (rationality, FTL, etc.) gives the value V_{bg}^m to box b in game g . Recall that we parameterize the link between this value and the probability that a player will choose b in game g by

$$P_{bg}^m(\beta) = \frac{\exp(\beta V_{bg}^m)}{\sum_{b'=1}^B \exp(\beta V_{b'g}^m)}$$

for a given β . To determine β , we minimize the Euclidean distance (in \mathbb{R}^{BG}) between the empirical distribution \hat{P} and the distribution predicted by the model, given β . Specifically, we take

$$\begin{aligned} L^m(\hat{P}, \beta) &= \|P^m(\beta) - \hat{P}\|^2 \\ &= \sum_{b,g} [P_{bg}^m(\beta) - \hat{P}_{bg}]^2 \\ \hat{L}^m &= \text{Min}_{\beta \in \mathbb{R}^+} L^m(\hat{P}, \beta). \end{aligned}$$

Hence, \hat{L}^m is the empirical (quadratic) distance between the model and the data. In a companion paper (Gabaix and Laibson, 1999) we derive asymptotic standard errors for \hat{L}^m and $\hat{L}^m - \hat{L}^{m'}$, where $m \neq m'$.

Table 1 reports test statistics for each of our models. The rows of the table represent the different models, $m \in \{\text{rational, FTL, } Q\text{-column-cutoff, } \delta\text{-discounting}\}$. Column (i) reports distance metric \hat{L}^m , which measures the (square of the) Euclidean distance between the estimated model and the empirical data. Column (ii) reports the difference $\hat{L}^{\text{rational}} - \hat{L}^m$. When $\hat{L}^{\text{rational}} - \hat{L}^m > 0$ the rational model performs relatively poorly, since the distance of the rational model from the empirical data ($\hat{L}^{\text{rational}}$) is greater than the distance of the competitor model from the empirical data (\hat{L}^m). Column (iii) reports the difference $\hat{L}^{\text{FTL}} - \hat{L}^m$.

Three findings stand out. First, several models statistically significantly outperform rationality: FTL, $Q = 8$ column-cutoff, $Q = 9$ column-cutoff, $\delta = 0.8$ discounting, $\delta = 0.9$ discounting, and $\delta = 0.91$ discounting. We report $\delta = 0.91$ discounting, since this value of δ minimizes $\hat{L}^{\text{rational}} - \hat{L}^{\delta \text{ discounting}}$. Second, FTL outperforms the rational model by the biggest margin: $\hat{L}^{\text{rational}} - \hat{L}^{\text{FTL}} = 0.195$. We reject $\hat{L}^{\text{rational}} - \hat{L}^{\text{FTL}} = 0$ with a t statistic of 8.75. The next closest contender is $Q = 8$ column-cutoff: $\hat{L}^{\text{rational}} - \hat{L}^{Q=8 \text{ column-cutoff}} = 0.104$. Third, we reject all models in favor of FTL. For the closest competitor, we reject $\hat{L}^{\text{FTL}} - \hat{L}^{Q=8 \text{ column-cutoff}} = 0$ with a t statistic of 6.06. Hence, the experimental evidence overwhelmingly supports FTL.

IV. Extensions

FTL exemplifies the general cognitive tendency to simulate the future by identifying typical or representative scenarios. For example, if numerous future outcomes are equally likely, a

TABLE 1—MODEL STATISTICS

<i>m</i> = Model	(i) <i>L</i> (<i>m</i>)	(ii) <i>L</i> (Rational) − <i>L</i> (<i>m</i>)	(iii) <i>L</i> (FTL) − <i>L</i> (<i>m</i>)
Rational	0.789 (0.063)		−0.195 (0.022)
FTL	0.594 (0.051)	0.195 (0.022)	
Column cutoff 1	1.903 (0.085)	−1.114 (0.081)	−1.309 (0.083)
Column cutoff 2	1.253 (0.076)	−0.464 (0.074)	−0.659 (0.070)
Column cutoff 3	0.819 (0.059)	−0.031 (0.055)	−0.225 (0.047)
Column cutoff 4	0.806 (0.061)	−0.017 (0.032)	−0.212 (0.026)
Column cutoff 5	0.852 (0.062)	−0.064 (0.033)	−0.258 (0.028)
Column cutoff 6	0.831 (0.060)	−0.042 (0.030)	−0.237 (0.027)
Column cutoff 7	0.824 (0.064)	−0.035 (0.023)	−0.229 (0.025)
Column cutoff 8	0.685 (0.059)	0.104 (0.013)	−0.091 (0.015)
Column cutoff 9	0.728 (0.062)	0.061 (0.008)	−0.133 (0.019)
δ = 0.5 discounting	1.111 (0.086)	−0.322 (0.064)	−0.516 (0.060)
δ = 0.6 discounting	0.978 (0.087)	−0.189 (0.056)	−0.384 (0.052)
δ = 0.7 discounting	0.864 (0.089)	−0.075 (0.047)	−0.270 (0.042)
δ = 0.8 discounting	0.769 (0.092)	0.019 (0.036)	−0.175 (0.031)
δ = 0.9 discounting	0.704 (0.096)	0.085 (0.020)	−0.110 (0.019)
δ = 0.91 discounting	0.704 (0.096)	0.085 (0.018)	−0.109 (0.019)

Notes: *L*(*m*) is the squared Euclidean distance between the predictions of model *m* and the empirical data. Asymptotic standard errors appear in parentheses.

sensible forecast may focus on the median or the mean outcome. Such simplifications are generally effective strategies for planning in complex environments. Likewise, incomplete contracts arise because it is too costly to develop contingencies for every state of the world. Strategic planning is simplified when one assumes that other parties and nature will behave in typical ways.

Occasionally, however, these simplifications prove to be counterproductive. For example, in adverse selection (“winner’s curse”) problems, simplifying uncertainty leads generally to important deviations from rational decision-mak-

ing. Consider the well-studied example of a bidder who values a widget 50-percent more than a seller. The bidder does not know the underlying value of the widget to the seller. Only the seller knows this value, which is distributed uniformly from 0 to 100. Replacing this distribution of outcomes with the typical outcome (a valuation of 50), leads to a bid of 50 + ε, far from the optimal bid of 0.⁷

Endogenous Algorithms

FTL successfully reproduces the decisions of subjects in our experiment. This result is likely to be robust, since our trees are randomly generated. However, it is easy to construct special examples, say, with extreme outlier payoffs, in which FTL may not do well. We view FTL as a special case of a richer model that combines several decision algorithms. For decision trees with extreme outlier payoffs, consider the following algorithm:

Identify boxes with outlier payoffs ($|\text{payoff}| > \mu$). Start at these outlier payoffs and apply FTL backward, moving from right to left.

This algorithm highlights certain paths, some of which intersect with the forward simulation paths generated by the standard FTL algorithm. In general, decision-makers evaluate a given starting box by integrating over the union of forward-FTL paths and the intersecting backward-FTL paths. This combined algorithm is parameterized by the original probability threshold, *p*, and μ , a threshold value that defines the minimal magnitude that a payoff must have to attract attention as an extreme outlier.

Further research will extend this framework in two ways. First, new research will identify a parsimonious set of parameterized algorithms. Second, new research will provide a theory that describes how the parameters adjust across problems. Natural adjustment candidates include reinforcement learning and expected-payoff

⁷ See Richard Thaler (1994) for a discussion of the “winner’s curse.”

maximization subject to constraints on calculation and memory. We are currently exploring these research directions.

REFERENCES

- Camerer, Colin and Ho, Teck-Hua.** "Experience Weighted Attraction Learning in Normal-Form Games." *Econometrica*, July 1999, 67(4), pp. 827–74.
- Camerer, Colin; Johnson, E. J.; Rymon, T. and Sen, S.** "Cognition and Framing in Sequential Bargaining for Gains and Losses," in Ken Binmore, Alan Kirman, and P. Tani, eds., *Frontiers of game theory*. Cambridge, MA: MIT Press, 1994, pp. 27–47.
- Conlisk, John.** "Why Bounded Rationality?" *Journal of Economic Literature*, June 1996, 34(2), pp. 669–700.
- Erev, Ido and Roth, Alvin.** "Predicting How People Play Games." *American Economic Review*, September 1998, 88(2), pp. 848–81.
- Gabaix, Xavier and Laibson, David.** "A Positive Theory of Bounded Rationality." Mimeo, Massachusetts Institute of Technology, 1999.
- Kahneman, Daniel and Tversky, Amos.** "Judgment Under Uncertainty: Heuristics and Biases." *Science*, 27 September 1974, 185(4157), pp. 1124–31.
- Simon, Herbert.** "A behavioral model of rational choice." *Quarterly Journal of Economics*, February 1955, 69(1), pp. 99–118.
- Thaler, Richard.** *The winner's curse*. New York: Russell Sage, 1994.