THE ALLOCATION OF ATTENTION: THEORY AND EVIDENCE

Xavier Gabaix*
MIT and NBER

David Laibson
Harvard University and NBER

Guillermo Moloche
NBER

Stephen Weinberg
Harvard University

Current Draft: October 2002

Abstract. The combination of limited mental processing speed and scarce time bounds the quality of decision-making and defines a fundamental resource allocation problem: decision makers need to continuously decide what to think about. In this paper, we report the first empirical test of an economic model of endogenous attention allocation. We report the results of an experiment in which decision time is a scarce resource and in which attention allocation is continuously measured using Mouselab. We compare measured attention allocation choices to the choices predicted by a tractable attention allocation model based on option value principles. Subject behavior correponds well to the quantitative predictions of the model.

JEL classification: C7, C9, D8.

Keywords: bounded rationality, satisficing, attention, Mouselab.

*We thank Miguel Costa-Gomes, Vince Crawford, Colin Camerer and seminar participants at Harvard University, Massachusetts Institute of Technology, Stanford University, University of California Berkeley, and University of Montréal. Numerous research assistants helped us run the experiment. We are particularly indebted to Rebecca Thornton and Natalia Tsvetkova. Gabaix and Laibson acknowledge financial support from the NSF (SES-0099025). Weinberg acknowledges support from an NSF Graduate Research Fellowship, and Laibson acknowledges support from the Olin Foundation. Xavier Gabaix: MIT, 50 Memorial Drive, Cambridge, MA 02142. David Laibson and Stephen Weinberg: Harvard University, Department of Economics, Cambridge, MA, 02138. Guillermo Moloche: NBER, 1050 Massachusetts Avenue, Cambridge, MA 02138.

1. Introduction

Most economic models contain the useful assumption that all thinking is instantaneous. However, real people have limited processing speeds and consequently make most decisions under time pressure. Like players in a chess tournament or students taking a test, we do not have the luxury of taking our time to make decisions. We necessarily stop thinking about most problems before we have a perfect solution.[1]

The combination of limited processing speed and scarce time bounds the quality of our decision-making and defines a fundamental resource allocation problem: decision makers need to continuously decide what to think about. This attention allocation process has numerous intuitive implications. Decisions with important consequences — Which car should I buy? — receive more attention than unimportant decisions — Which car should I rent for the weekend? Likewise, "close" decisions — Should I buy a Toyota or a Honda? — receive more attention than "no-brainers" — Should I buy a Toyota or a Lada?

Economics is often defined as the study of the allocation of scarce resources. So the allocation of attention, and its consequences for decision-making, seems like a natural topic for economic research. But, little empirical work has been done on this problem. Two hurdles have slowed down this research. First, attention allocation is difficult to observe empirically. Second, models of attention allocation face several conceptual problems. Thinking about thinking generates an infinite regress of thinking levels (Simon 1955, Conlisk 1996).

In this paper, we address these problems and report the first empirical test of an economic model of endogenous attention allocation. We report the results of an experiment in which decision time is a scarce resource and in which attention allocation is continuously measured. We compare measured attention allocation

[1] See Kahneman and Tversky (1974) and Gigerenzer et al (1999) for a description of some of the simple heuristics that decision-makers use to solve complex problems.

choices to the choices predicted by a tractable attention allocation model based on cost-benefit analysis.

In our experiment we make time scarce in two different ways. In one part of the experiment we give subjects an exogenous amount of time to complete a choice problem in which subjects choose one good from a set of N goods. We call this an *N-good game with an exogenous time budget*. Here we measure how subjects allocate their time within the game, as they think about each good in the set of N goods.

In another part of the experiment we give subjects an open-ended sequence of N-good problems. In this design each N-good problem is followed by another N-good problem. Subjects keep facing different N-good problems until their total budget of time runs out. Although the total time budget is fixed, the amount of time chosen for each game is a choice variable. Hence, we refer to each game as an *N-good game with an endogenous time budget*. Since payoffs are cumulative and each N-good problem has a positive expected value, subjects have an incentive to move through the N-good problems quickly. But, moving too quickly reduces the quality of their decisions. As a result, the subjects trade off quality and quantity. Spending more time on any given N-good problem raises the quality of the choice in that problem but reduces the time available to participate in future problems with accumulating payoffs.

Throughout the experiment we measure how our subjects endogenously allocate their attention. We measure attention allocation within each N-good problem (i.e., among the N goods). In addition, in the open-ended sequence design, we also measure attention allocation across N-good problems.

Following the lead of other economists (Camerer *et al.* 1993 and Costa-Gomes, Crawford, and Broseta 2001), we use the "Mouselab" programming language to measure subject's attention allocation.[2] Mouselab tracks subjects' information search during an experiment. Information is hidden "behind" boxes on a computer screen. Subjects use the computer mouse to open the boxes. Mouselab records what information subjects access and when they access it. Since only one screen box can be open at any point in time, the Mouselab software enables an experimenter to pinpoint exactly what information the subject is contemplating on a second-by-second basis throughout the experiment.[3]

Gabaix and Laibson (2002) propose a bounded rationality model based on the microfoundation of limited processing speed. In this model, time-pressured agents decide which information to analyze, allocating their attention according to the expected economic value of marginal analysis. We use this model to evaluate the choices that our subjects make.

Subject behavior correponds well to the predictions of the model. Intuitively, subjects allocate thinking time when the marginal value of such thinking is high, either because competing goods are close in value or because there is a "large" amount of remaining information to be revealed. We demonstrate this in numerous ways. We evaluate the pattern of information acquisition within N-good games. In the open-ended sequence design, we evaluate the relationship between economic incentives and subjects' decisions about when to stop working on one N-good game so they can move on to the next N-good game. We find that the economic model of attention allocation outperforms myopic information acquisition models. The economic value

[2] Bettman, Payne, and Johnson (1993), developed the Mouselab language in the 1970's. Mouselab is one of many "process tracing" methods. For example, Payne, Braunstein, and Carrol (1978) elicit mental processes by asking subjects to "think aloud." Russo (1978) records eye movements.

[3] Mouselab has the drawback that it uses an artificial decision environment, but several studies have shown that the Mouselab environment does not significantly distort final choices over goods/actions (e.g., Costa-Gomes, Crawford, and Broseta 2001 and Costa-Gomes and Crawford 2002). However, Mouselab's interface does generate an "upper-left" search bias, which we discuss below.

of information turns out to be the most important predictor of attention allocation.

However, the experiment also reveals a deviation from the economic model. Subject choices are partially predicted by a "boxes heuristic." Specifically, subjects become more and more likely to end analysis of a problem the more boxes they open, holding *fixed* the economic benefits of additional analysis. In this sense, subjects display a partial insensitivity to the particular circumstances of each problem that they face. This finding supports research on the "$1/N$ heuristic" (Bernatzi and Thaler 2001) and on "system neglect" (Camerer and Lovallo 1999, Massey and Wu 2002). In the context of our experiment, we interpret system neglect to be another reflection of bounded rationality. Subjects generally look for the information with the highest economic value, but imperfectly identify/locate this useful information. We also test for other evidence of heuristic decision-making, but we find little evidence for commonly studied heuristics, including satisficing and elimination by aspects.

Section 1 describes our experimental set-up. Section 2 summarizes an implementable one-parameter attention allocation model (Gabaix and Laibson 2002). Section 3 summarizes the results of our experiment and compares those results to the predictions of our model. Section 4 concludes.

2. Experimental Set-up

Our experimental design facilitates second-by-second measurement of attention allocation in a decision problem. We begin this section by describing the basic building block of the experiment: a choice among N goods.

2.1. An N-good game. An N-good game is an N-row by M-column matrix of boxes (Figure 1). Each box contains a random payoff (in units of cents) generated with normal density and zero mean. After analyzing an N-good game, the subject "consumes" a single row from that game. The subject is paid the sum of the boxes in the consumed row.

Consuming a row is an abstract representation of a very wide class of choice problems. We call this problem an N-good game, since the N rows conceptually represent N goods. The subject is asked to consume one of these N goods. The columns represent different attributes of the N goods.

The importance or variability of these attributes declines as the columns move from left to right. In particular, the variance decrements across columns equal one tenth of the variance in column one. For example, if the variance used to generate column one is 1000 (squared cents), then the variance for column 2 is 900, and so on, ending with a variance of 100 for column 10.

So far the game sounds simple: "Consume the best row." We make the task harder by masking the contents of boxes in columns 2 through M. Subjects are shown only the box values in column 1. However, a subject can mouse click on a masked box in columns 2 through M to unmask the value of that box (Figure 2). Only one box from columns 2 through M can be unmasked at a time. This procedure enables us to record exactly what information the subject is analyzing at every point in time. Note that unmasking a box does not imply that the subject consumes that box. Only after the subject has unmasked some boxes does the subject decide which row to consume. Note too that if a row is picked for consumption, then all boxes in that row are consumed, whether or not they have been previously unmasked.

We introduce time pressure, so that subjects will not be able to unmask — or will not choose to unmask — all of the boxes in the game. Mouselab enables us to record which of the $N(M-1)$ masked boxes the subject chooses to unmask. Of course, we also record which row the subject ultimately consumes.

2.2. Games with exogenous and endogenous time budgets.. In our experiment, subjects played two different types of N-good games: games with exogenous time budgets and games with endogenous time budgets, which we will respectively refer to as "exogenous" and "endogenous" games.

For each "exogenous" game a game-specific time budget is generated from the uniform distribution over the interval [10 seconds, 49 seconds]. A clock shows the subject the amount of time remaining for each isolated game (see clock in Figure 2).

In endogenous games, subjects have a fixed budget of time — 25 minutes — in which to play as many different N-good games as they choose. In this design, each N-good games is separated by a 20 second buffer screen, which counts toward the total budget of 25 minutes. Subjects are free to spend as little or as much time as they want on each game and they are free to play as many games as they can before the 25 minute time budget expires.

2.3. Experimental logistics. Subjects receive written instructions explaining the structure of an N-good game and the logistics for the isolated and open-ended designs. Computer instructions explain the Mouselab interface. All instructions are available in appendix A.

At the beginning of the experiment, subjects play three test games which do not count toward their payoffs. Then subjects play 12 games with seperate exogenous time budgets as well as a set of endogenous games with a joint 25-minute time budget.

For half of the subjects we reverse the order of the exogenous and endogenous games. At the end of the experiment, subjects answer demographic and debriefing questions.

Subjects are paid the cumulative sum of all rows that they consume. Subjects are given feedback after every game. The feedback reports the running cumulative value of the consumed rows.

3. Directed cognition model

The previous section describes a complex attention allocation problem. In isolated games, subjects must decide which boxes to unmask before their time runs out. In open-ended games, subjects must jointly decide which boxes to unmask and when to move on to the next game.

We want to determine whether economic principles guide subjects' attention allocation choices. We compare our experimental data to the prediction of an attention allocation model proposed by Gabaix and Laibson (2001a). This 'directed cognition' model approximates the economic value of attention allocation using two simple economic ideas that can be deduced from option value analysis.

First, when many different choices are being compared and a particular choice gains a large edge over the available alternatives, the option value of continued analysis declines. Second, when cognitive analysis yields little new information (i.e., the standard deviation of marginal information is low), the economic value of continued analysis declines. The directed cognition model quantifies these two effects and integrates them in a formal framework that makes sharp quantitative predictions about boundedly rational attention allocation choices.

We first introduce our notation and then formally derive an equation that measures these option value effects. Since our games all have eight rows (goods), we label the rows, A, $B, ..., H$. We will use lower case letters — $a, b, ..., h$ — to track a subject's expectations of the values of the respective rows.

At the beginning of the game, the row expectations will equal the value of the payoff in the left-most cell of each row. Recall that the subject knows the value of all boxes in column 1 when the game begins. For example, if row A has a 22 in its first cell, then at time zero $a = 22$. This value will be updated as the subject unmasks information in row A. If the subject unmasks the second and third cells in row A, revealing cell values of 10 and 5, then a would be updated to $37 = 22 + 10 + 5$.

Our application of the directed cognition model is built on a basic and natural cognitive operation: unmasking boxes in the respective rows of the matrix. Such operations enable the decision-maker to improve her forecast of the expected value of any given row. For example, O_A^Γ represent the operation: "open Γ additional boxes in row A."

We assume that an operator that opens Γ boxes has cost $\Gamma \cdot \kappa$, where κ is the cost of unmasking a single box. We take this cost to include many components, including the time involved in opening a box with a mouse, reading the contents of the box, and updating one's expectations. Such updating includes an addition operation as well as two memory operations: recalling the prior expectation and rememorizing the updated expectation.

Operator O_A^Γ leads a subject to update his expectations of the value of consuming row A. Formally, we write

$$a' = a + \varepsilon$$

where ε is the sum of the values in the Γ newly unmasked boxes in row A.

We can now derive an expression that serves as a proxy for the economic value of cognitive operator O. Examine the case in which the decision-maker is considering analyzing row A. Assume that row B would be the leading row if row A were eliminated, so row B is the next best alternative to row A.

The agent is thinking about learning more about row A by executing a cognitive operater O_A^Γ. Executing the cognitive operator will enable the agent to update the expected payoff of row A from a to $a' = a + \varepsilon$. If the agent doesn't execute the cognitive operator, her expected payoff will be

$$\max(a, b).$$

If the agent executes the cognitive operator, her expected payoff will be

$$E[\max(a', b)].$$

This expected payoff is based on the assumption of partial myopia. In particular, the agent only considers the immediate value of the information revealed by the cognitive operator and ignores the option value of executing subsequent cognitive operators. As discussed below, this assumption is made to preserve computational tractability.

The value of executing the cognitive operator is the difference between the previous two expressions:

$$E[\max(a', b)] - \max(a, b). \tag{1}$$

This value can be represented with a simple expression. Let σ represent the standard deviation of the updated estimate resulting from application of cognitive operator:

$$\sigma^2 = E(a' - a)^2.$$

Appendix B shows that the value of the cognitive operator is[4]

[4]This result assumes Gaussian innovations, which is the density used to generate the games in

$$E\left[\max\left(a',b\right)\right]-\max\left(a,b\right)=w\left(a-b,\sigma\right), \tag{2}$$

with

$$w(x,\sigma)\equiv\sigma\phi\left(\frac{x}{\sigma}\right)-\left|x\right|\Phi\left(-\frac{\left|x\right|}{\sigma}\right), \tag{3}$$

where ϕ represents the standard normal density function and Φ represents the associated cumulative distribution function. The $w\left(x,\sigma\right)$ function captures the expected value of implementing a cognitive operator O, where x is the estimated value gap between the active row and its next best alternative and σ is the standard deviation of the payoff information revealed by O. Figure 3 plots $w\left(x,1\right)$. The general case can be deduced from the fact that $w\left(x,\sigma\right)=\sigma w\left(x/\sigma,1\right)$

Two fundamental comparative statics are captured in this option value framework. First, the value of a path exploration decreases the larger the gap between the active row and the next best row: $w\left(x,\sigma\right)$ is decreasing in $|x|$. Second, the value of a path exploration increases with the variability of the information that will be obtained: $w\left(x,\sigma\right)$ is increasing in σ. In other words, the more information that is likely to be revealed by a path exploration, the more valuable such a path exploration becomes.

Now that we have described how to calculate the economic value of a cognitive operator, we can describe how we model subjects' choices among cognitive operators. We assume that at every point in time the subject implements the cognitive operator with the highest ratio of the expected benefit to the cost. Recall that the expected benefit of an operator is given by the $w(x,\sigma)$ function and that the implementation cost of an operator is the number of boxes that it unmasks. The subject executes

our experiment.

the cognitive operator with the greatest benefit/cost ratio,

$$G \equiv \max_{O} w(x_O, \sigma_O)/\Gamma_O.$$

Hence, the subject executes cognitive operator

$$O^* = \arg\max_{O} w(x_O, \sigma_O)/\Gamma_O.$$

Appendix C contains an example of such a calculation.

In games with an exogneous time budget, the subject keeps selecting cognitive operators in this way until time runs out. In games with an endogenous time budget, the subject keeps executing cognitive operators until G falls below the marginal value of time.

We calibrate the marginal value of time in two ways. First, we estimate the value of time as perceived by our subjects. Advertisements for the experiment implied that subjects would be paid about \$20 for their participation. Since they were promised \$5 of guaranteed payoffs, this implied that they would receive \$15 for the rest of the experiment. Dividing this in half implies an expectation of about \$7.50 for each half of the experiment: the exogenous games and the endogenous games. Since the opened-ended games were budgeted to take 25 minutes, this implies a price per second of

$$(750 \text{ cents})/\left(25 \text{ minutes} \cdot 60 \text{ seconds/minute}\right) = 0.50 \text{ cents/second}.$$

Since subjects took on average 0.98 seconds to open each box, we end up with an effective price per box of

$$(0.50 \text{ cents/second})(0.98 \text{ seconds/box}) = 0.49 \text{ cents/box}.$$

We also explored a 1-parameter version of the directed cognition model, in which the cost of cognition — κ — was chosen to make the model partially fit the data. Calibrating the model so it matches the average number of boxes explored in the open-ended games implies $\kappa = 0.18$ cents/box. Here κ is chosen only to match the average amount of search per open-ended game, not to match the order of search or the distribution of search across games.

The model described in this section has the virtue that it is easy to analyze and computationally tractable. This simplicity arises because of two special assumptions. First, the directed cognition model assumes that the agent only calculates the immediate expected gains from executing each cognitive operator, and does not anticipate the option value of continued search. This assumption resembles the approach taken by Jehiel (1995), who assumes a constrained planning horizon in a game-theory context. Second, the directed cognition model avoids the infinite regress problem (i.e., thinking about thinking about thinking, etc...), by implicitly assuming that solving for O^* is costless. Without some version of these two simplifying assumptions the model would not be practically implementable. Without some myopia (i.e., a limited evaluation horizon), multi-good and multi-attribute option value problems cannot be solved either analytically or even computationally.[5] In addition, without eliminating cognition costs at some basic stage of reasoning, maximization models are not well-defined.[6]

[5] See Gittins (1979) and Weitzman (1979) for restricted cases in which the solution to an option value problem is computationally tractable.

[6] See Conlisk (1996) for a description of the infinite regress problem and an explanation of why it plagues all decision cost models. We follow Conlisk in advocating exogenous truncation of the

3.1. Other decision algorithms . In this subsection, we discuss several other models that we compare to the directed cognition model.

The *column model* unmasks boxes column by column, stopping either when time runs out (in games with exogenous time budgets) or stopping according to a satisficing rule (in games with endogenous time budgets). Specifically the column model unmasks all the boxes in column 2 (top to bottom), then in column 3,..., etc. In exogenous games, this column-by-column unmasking continues until the simulation has explored the same number of boxes as a yoked subject. In endogenous games, the unmasking continues until a row has been revealed with estimated value greater than or equal to $A^{\text{Column model}}$, an aspiration level estimated so that the simulations generate an average number of simulated box openings that matches the average number of empirical box openings (26 boxes per game).

The *row model* unmasks boxes row by row, starting with the "best" row and moving to the "worst" row, stopping either when time runs out (in games with exogenous time budgets) or stopping according to a satisficing rule (in games with endogenous time budgets). Specifically the row model ranks the rows according their values in column 1. Then the row model unmasks all the boxes in the best row, then the second best row, etc. In exogenous games, this row-by-row unmasking continues until the simulation has explored the same number of boxes as a yoked subject. In endogenous games, the unmasking continues until a row has been revealed with estimated value greater than or equal to $A^{\text{Row model}}$, an aspiration level estimated so that the simulations generate an average number of simulated box openings that matches the average number of empirical box openings.

A choice algorithm called *Elimination by Aspects* (EBA) has been widely studied in the psychology literature (see work by Tversky 1972 and Payne, Bettman, and

infinite regress of thinking.

Johnson 1993). We apply this algorithm to our decision framework to analyze games with endogenous time budgets. We exploit the interpretation that each row is a good with 10 different attributes or "aspects" represented by the ten different boxes of the row. Our EBA application assumes that the agent proceeds aspect by aspect (i.e., column by column) from left to right , eliminating goods (i.e. rows) with an aspect that falls below some threshold value A^{EBA}. This elimination continues, stopping at the point where the *next* elimination will eliminate all remaining rows. At this stopping point, we pick the remaining row with the highest estimated value. As above, we estimate A^{EBA} so that the simulations generate an average number of simulated box openings that matches the average number of empirical box openings.

4. Results

4.1. Subject Pool. Our subject pool is comprised of 388 Harvard undergraduates. Two-thirds of the subjects are male (66%). The subjects are distributed relatively evenly over concentrations: 11% math or statistics; 21% natural sciences; 20% humanities; 29% economics; and 20% other social sciences. In a debriefing survey we asked our subjects to report their statistical background. Only 15% report taking an advanced statistics class; 40% report only an introductory level class; 45% report never having taken a statistics course.

Subjects received a mean total payoff of \$29.23, with standard deviation of \$5.49. Payoffs range from \$13.07 to \$46.69. All subjects played 12 games with exogenous times. On average subjects played 28.7 games with endogenous time, with standard deviation of 7.9. The number of games played range from 21 to 65.

Payoffs do *not* systematically vary across concentrations, class years, age, and statistics coursework and self-described ability in statistics. Table A1 reports average pay by concentration and by statistics coursework. For example, in the three concentrations for which we have at least 50 observations, the average payouts per subject are respectively \$28.9, \$29.0 and \$29.1. Average payoffs correlate with the number of statistics courses taken by the subject (Table A1), but this is actually a gender effect (Table A2). Men earn \$3.1 more than women, and males tend to take relatively more statistics courses. Of the six demographic variables that we measured, only gender has significant predictive power in a regression framework. In light of the lack of much predictable variation across groups, we have chosen to adopt the useful approximation that all subjects have identical strategies. Relaxing this assumption is beyond the scope of the current paper, but we regard such a relaxation as an important future research goal.

4.2. Games with exogenous time budgets. We begin by analyzing attention allocation patterns in the games with exogenous time budgets. Our analysis focuses on the pattern of box openings across columns and rows. We compare the empirical patterns of box openings to the patterns of box openings predicted by the directed cognition model.

Figure 4 reports the average number of boxes opened in columns 2-10. We report the average number of boxes unmasked, column by column, for both the subject data and the model predictions.

The empirical profile is calculated by averaging together subject responses on all of the exogenous games that were played. Specifically, each of our 388 subjects played 12 isolated games, yielding a total of $388 \times 12 = 4656$ exogenous games. Of these 4656 games, 160 are unique, though no subject played the same game twice. To calculate the empirical column profile (and all of the profiles that we analyze) we only count the first unmasking of each box. So if a subject unmasks the same box twice, this only counts as one opening. Approximately 90% of box openings are first-time unmaskings.[7]

[7] We do not analyze repeat unmaskings because they are relatively rare in our data and because

Figure 4 also plots the theoretical predictions generated by "yoked" simulations of our model. Specifically, these predictions are calculated by simulating the directed cognition model on the exact set of 4656 games played by the subjects. We simulate the model on each game from this set of 4656 games and instruct the computer to unmask the same number of boxes that were unmasked by the subject who played each respective game. The analysis compares the particular boxes opened by the subject to the particular boxes opened by the yoked simulation of the model.

Figure 4 also reports an $\widetilde{R}^2$ measure, which captures the extent to which the empirical data matches the theoretical predictions. This measure is simply the $\widetilde{R}^2$ statistic[8] from the following constrained regression:[9]

$$Boxes(col) = \text{constant} + \widehat{Boxes}(col) + \varepsilon(col).$$

Here $Boxes(col)$ represents the empirical average number of boxes unmasked in column col and $\widehat{Boxes}(col)$ represents the simulated average number of boxes opened in column col. Note that col varies from 2 to 10, since the boxes in column 1 are always unmasked. This $\widetilde{R}^2$ statistic is bounded below by $-\infty$ (since the coefficient on $\widehat{Boxes}(col)$ is constrained equal to unity) and bounded above by 1 (a perfect fit). Intuitively, the $\widetilde{R}^2$ statistic represents the fraction of squared deviations around the

we are interested in building as simple and parsimonious model as possible. Incorporating memory constraints would improve the fit of the model but would also partially undermine our goal of parsimony.

[8]In other words, the $\widetilde{R}^2$ is

$$\widetilde{R}^2 = 1 - \frac{\sum_i \left(Boxes(col) - \langle Boxes \rangle - \widehat{Boxes}(col) + \left\langle \widehat{Boxes} \right\rangle \right)^2}{\sum_{col} \left(\widehat{Boxes}(col) - \left\langle \widehat{Boxes} \right\rangle \right)^2}$$

where $\langle \cdot \rangle$ represents means.

[9]In this section of the paper, the constant is redundant, since the dependent variable has the same mean as the independent variable. However, in the next subsection we will consider cases in which this equivalence does not hold, necessitating the presence of the constant.

mean explained by the model. For the column predictions, the $\widetilde{R}^2$ statistic is 86.6%, implying a very close match between the data and the predictions of the model.

Figure 5 reports analogous calculations by row. Figure 5 reports the number of boxes opened on average by row, with the rows ranked by their value in column one. (Recall that column one is never masked.) We report the number of boxes opened on average by row for both the subject data and the model predictions. As described above, the model predictions are calculated using yoked simulations.

Figure 5 also reports an $\widetilde{R}^2$ measure analogous to the one described above. The only difference is that now the variable of interest is $Boxes(row)$, the empirical average number of boxes opened in row row. For our row predictions our $\widetilde{R}^2$ measure is -14.4%, implying a poor match between the data and the predictions of the model. The model simulations predict too many unmaskings on the top ranked rows and far too few unmaskings on the bottom ranked rows. The subjects are much less selective than the model. The $\widetilde{R}^2$ is negative because we constrain the coefficient on simulated boxes to be unity. It turns out that this is the only bad prediction that the model makes.

Figure 6 reports similar calculations using an alternative way of ordering rows. Figure 6 reports the number of boxes opened on average by row, with the rows ranked by their value at the *end* of each game. Figure 6 also reports the associated $\widetilde{R}^2$ measure. For our row predictions the $\widetilde{R}^2$ statistic is 87.2%.

4.3. Endogenous games. We repeat the analysis above for the endogenous games. As discussed in section 3, we consider two variants of the Directed Cognition Model when analyzing the endogenous games. One variant is calibrated by exogenously setting κ to match the subjects' expected earnings per unit time in the endogenous games: $\kappa = 0.49$ cents/box opened (see calibration discussion in section 3). With this calibration, subjects are predicted to open 15.72 boxes per game. In the data, however, subjects open 26.06 boxes per game. To match this frequency of box opening, we consider a second calibration with $\kappa = 0.18$. With this lower level of κ, the model opens the empirically "right" number of boxes.

Figure 7 reports the average number of boxes unmasked in columns 2-10 in the endogenous games. We report the average number of boxes unmasked by column for the subject data and for the model predictions with $\kappa = 0.49$ and $\kappa = 0.18$.

The empirical data is calculated by averaging together subject responses on all of the endogenous games that were played. The theoretical predictions were generated by "yoked" simulations of our model. Specifically, we use the directed cognition model to simulate play of the 10,931 endogenous games that the subjects played. The model generates its own stopping rule (so we no longer yoke the number of box openings).

Figure 7 also reports the associated $\widetilde{R}^2$ statistic for these column comparisons in the endogenous games. For these endogenous games, the column $\widetilde{R}^2$ statistic is 99.0% for $\kappa = 0.49$ and 85.5% for $\kappa = 0.18$.

Figure 8 reports analogous calculations by row for the endogenous games. Figure 8 reports the number of boxes opened on average by row, with the rows ranked by their value in column one. We report the number of boxes opened on average by row for both the subject data and the model predictions. Figure 8 also reports the associated $\widetilde{R}^2$ statistic. For these endogenous games, the row $\widetilde{R}^2$ statistics are 87.6% for $\kappa = 0.49$ and 70.5% for $\kappa = 0.18$.

Figure 9 reports similar calculations using the alternative way of ordering rows. Figure 9 reports the number of boxes opened on average by row, with the rows ranked by their value at the *end* of each game. Figure 9 also reports the associated $\widetilde{R}^2$ statistics. For these endogenous games, the alternative row $\widetilde{R}^2$ statistics are 96.3% for $\kappa = 0.49$ and 91.8% for $\kappa = 0.18$.

These figures show that the model explains a very large fraction of the variation in attention across rows and columns. However, some of the results in this section are partially confounded by a bias that Costa-Gomes, Crawford, and Broseta (2001) have found in their analysis. In particular, subjects who use the Mouselab interface tend to have a bias toward selecting cells in the upper left corner of the screen. The up-down component of this bias does not affect our results. But the left-right bias does, since information with greater economic relevance is located toward the left-hand side of the screen. One way to evaluate this bias would be to flip the presentation of the problem during the experiment so that the rows are labeled on the right and variance declines from right to left. Unfortunately, the internal constraints of Mouselab make this relabeling impossible. Future work should use a more flexible programming language that facilitates such spatial rearrangements.

4.4. Stopping Decisions in Endogenous Games:. Almost all of the analysis above reports *within*-game variation in attention allocation. The analysis above shows that subjects allocate most of their attention to economically relevant columns and rows within a game, matching the patterns predicted by the directed cognition model. Our experimental design also enables us to evaluate how subjects allocate their attention *between* games. In this subsection we focus on several measures of such between-game variation.

First, we compare the empirical variation in boxes opened per game to the predicted variation in boxes opened per game. Most importantly, we ask whether the model can correctly predict which games receive the most attention from our subjects. Let $Boxes(g)$ represent the average number of boxes opened by subjects who played game g. Let $\widehat{Boxes}(g)$ represent the average number of boxes opened by the model when playing game g. In this subsection we analyze the first and second moments of the empirical sample $\{Boxes(g)\}_{g=1}^{160}$ and the simulated sample $\{\widehat{Boxes}(g)\}_{g=1}^{160}$. Note that these respective vectors each have 160 elements, since we are analyzing game-specific averages. Our experiment utilized 160 unique games, though no single subject played all 160 games.

We begin by comparing first moments. The empirical (equally-weighted) mean of $Boxes(g)$ is 26. By contrast the 0-parameter version of our model (with $\kappa = 0.49$) generates a predicted mean of 16. Hence, unless we pick κ to match the empirical mean (i.e., $\kappa = 0.18$), our model only crudely approximates the average number of boxes opened per game.

We now turn to second moments. The empirical standard deviation of $Boxes(g)$ is 6.04, while the 0-parameter version of our model (with $\kappa = 0.49$) generates a predicted standard deviation of 11.61. Moreover, when we set $\kappa = 0.18$ to match the average boxes per game, the standard deviation rises to 15.91. The relatively high standard deviations predicted by the model, reflect the model's sophisticated search strategy. The model is highly attuned to instantaneous variation in the economic incentives that the subjects face, and the model continuously adjusts its search strategies accordingly. By contrast, the subjects are less sensitive to high frequency variation in economic incentives.

Despite these shortcomings, the model successfully predicts the pattern of empirical variation in the number of boxes opened in each game. The correlation between $Boxes(g)$ and $\widehat{Boxes}(g)$ is 0.66 when we set $\kappa = 0.49$. Similarly, the correlation is 0.64 when we set $\kappa = 0.18$. See Figure 10 for a plot of the individual (160) datapoints for the $\kappa = 0.18$ case. These high correlations imply that the model is doing a very good job predicting which games the subjects will analyze most thoroughly.

The model also does a good job predicting the relationship between economic incentives and depth of analysis. Figure 11 plots a non-parametric kernel estimate of the expected number of additional box openings in a game, conditional on the current value of the ratio of the benefit of marginal analysis to the cost of marginal analysis:

$$G \equiv \max_O w(x_O, \sigma_O)/\kappa_O.$$

The solid line represents the subject data. The crosses represent the relationship predicted by the model (with $\kappa = 0.18$). The dashed lines represent bootstrap estimates of the 95% confidence intervals. For most levels of G (the benefit-cost ratio), the model's predictions are close to the pattern in the subject data. Subjects do more analysis (i.e., open up more boxes) when the economic incentives to do so are high. Moreover, the functional form of this relationship roughly matches the form predicted by the theory.

Finally, we evaluate the directed cognition model by asking whether G predicts when subjects decide to stop working on the current game and move on to the next game. We ran a stopping logit (1 = stop, 0 = continue) with explanatory variables that include the measure of the economic value of continued search or G, the number of different boxes opened to date in the current game (*boxes*), the expected value in cents of the leading row in the current N-good game (*leader*), and subject fixed effects. Note that satisficing models predict that G and *boxes* should not have predictive power in this logit regression, but that a higher value of the leader variable will increase the stopping probability.

Each observation for this logit is a decision-node in our endogenous games, generating 330,873 observations.

$$\text{Probability of continuation } = \frac{\exp(x\beta)}{1 + \exp(x\beta)}.$$

We find that the variables with the greatest *economic* significance are G and *boxes,* with G playing the most important role. The coefficient on G is -0.3660 (with standard deviation 0.0081); the coefficient on *boxes* is .0404 (0.0008); the coefficient on *leader* is 0.0064 (0.0002). At the mean values of the explanatory variables, a one-standard deviation reduction in the value of G more than doubles the probability of stopping. A one-standard-deviation increase in the value of *boxes* has an effect roughly 1/2 as large and a one-standard-deviation increase in the value of *leader* has an effect less than 1/4 as large.

The logistic regression shows that the economic value of information — G — is by far the most important predictor of the decision to stop searching. However, our subjects are also using other information, as shown most importantly by the strong predictive power of the *boxes* variable. If subjects place partial weight on the "*boxes* heuristic" — i.e., increasing the propensity to stop analyzing the current N-good game as more and more boxes are opened — they will be less likely to spend a long time on any one game. For an unsophisticated player who can only imperfectly calculate G, the boxes heuristic is a useful additional decision input. We view our subjects' partial reliance on *boxes* as an example of sensible — perhaps constrained optimal — decision-making.

The predictive power of the *boxes* variables supports experimental research on "system neglect" by Camerer and Lovallo (1999) and Massey and Wu (2002). These authors find that subjects use sensible rules, but fail to adequately adjust those rules to the particular problem at hand. The boxes heuristic is a good general rule, but it is not the first-best rule since it neglects the idiosyncratic incentives generated by each specific N-good game. Such system neglect is a reflection of a type of bounded rationality that is important, but beyond the scope of the current paper.

4.5. Comparisons with Other Models. The analysis to date focuses on the predictions of the directed cognition model. In this subsection we consider alternative models and evaluate their performance. Table 1 reports $\widetilde{R}^2$ measures for all of the alternative models summarized in subsection 3.1.

It is immediately apparent that none of these models do nearly as well as the zero-parameter directed cognition model. For the exogenous games, the directed cognition model has an average $\widetilde{R}^2$ value of 53%. The Column and Row Models have respective averages of -18.5% and -137.0%.

For the endogenous games, the zero-parameter directed cognition model has an average $\widetilde{R}^2$ of 94.2%.[10] The Column Model with a Satisficing stopping rule has an average $\widetilde{R}^2$ of 44.5%. The Row Model with a Satisficing stopping rule has an average $\widetilde{R}^2$ of 38.2%. The Elimination by Aspects Model has an average $\widetilde{R}^2$ of 1.7%.

We also evaluate the different models' ability to forecast game-by-game variation in average time allocation.[11] Table 1 shows that all of the Satisficing Models yield effectively no correlation between the game-by-game average number of box openings predicted by these models and the game-by-game average number of box openings in the data. The Elimination by Aspects Model actually has a large negative correlation. By contrast, the zero-parameter directed cognition model generates predictions that are highly correlated (0.69) with the empirical average boxes opened by game.

5. Conclusion

Using Mouselab, we have experimentally evaluated a zero-parameter economic model of attention allocation, the directed cognition model. The model successfully predicts both the within-game and between-game allocation of attention. The model outperforms psychological models of attention allocation including satisficing models and elimination by aspects. The directed cognition model correctly predicts that subjects allocate more attention when the economic gains from such allocation are

[10] Other variants of the directed cognition model have average $\widetilde{R}^2$ values of 82.6% (one parameter model) and 59.1% (with satisficing stopping rule).

[11] Recall that this analysis uses the average number of boxes opened by game for the 160 unique games in the dataset. This data is compiled from the endogenous time segment of the experiment.

the highest.

The current paper adopts the useful assumption that all subjects have identical cognitive abilities and identical attention allocation strategies. Future work should test this assumption and identify the most important sources of differences in attention allocation across subjects. It is important to know whether all subjects allocate attention according to cost-benefit analysis, or whether such sensible rules only apply to a subset of the population. Generalizing beyond the population of Harvard students is an important step in this broader research program.

Appendix A: Instructions to Subjects

Paper Instructions

At the beginning of the experiment, subjects reviewed the following instructions on paper.

Instructions

Please take 5-6 minutes to read these instructions. These written instructions describe the decisions you will be making in this experiment. You will also see a computer demo on a laptop computer that will explain how the computer works. After the demo, you will spend approximately 35 minutes making decisions on the laptop.

In this experiment, we are studying how people make decisions under time limits. We will ask you to solve several problems where the choices you make can increase (or decrease) the amount of money that you will earn in this experiment. You will have a limited amount of time for each problem. Each problem will consist of eight rows of ten random numbers. Boxes will cover nine of the random numbers in each row; you can see these numbers by clicking on them with the mouse. You will always be able to see the first number in each row. For the sake of these instructions, we will show you pictures where the boxes have all been opened; the computer demo will show you how the boxes work. With all the boxes opened, two rows might look like this:

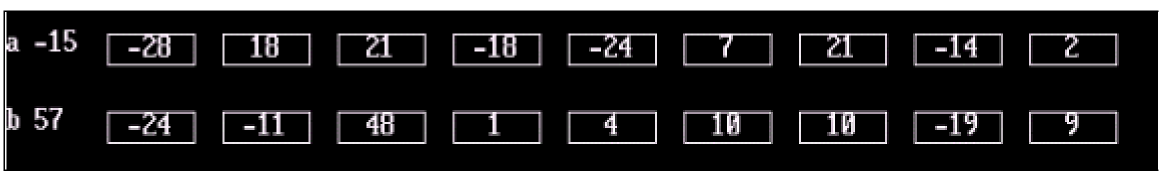

We will pay you the sum of the row you pick (in pennies), so your best choice will be the row with the highest sum. We will discuss payment in more detail at the end of these instructions.

I. Random Numbers

You may find it useful to understand how we generate the random numbers in the problems you will solve. We use bell curves to generate these numbers.

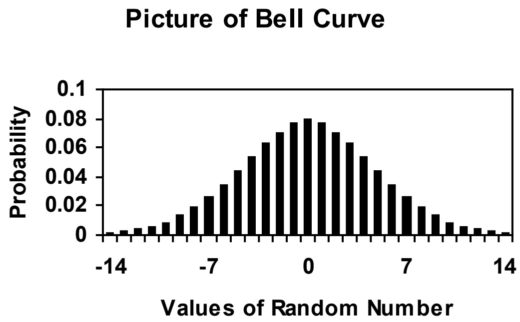

A bell curve, also called a "normal distribution," describes the way that random numbers are spread around some central value. In this experiment, the random numbers cluster around 0. A bell curve has three key properties:

- The higher the level of the bell curve for a particular number, the more likely we are to draw that number.
- There is an equal chance of getting a positive value or a negative value.
- The curve is highest at zero and then slopes downward on either side of zero, so the bell curve is more likely to generate numbers the closer they are to zero. For example, 1 should be more likely than 11.

How do we generate random numbers from bell curves? The bars show the percentage of the time that a random number takes a given integer value. There's only a small chance of getting any single number. For example, if we generated 100,000 random numbers from this bar graph, 7.6% of them should be –1, and 4.3% of them should be 5. For a smaller number of draws, these percentages may be only approximations.

Here is another look at the same bell curve.

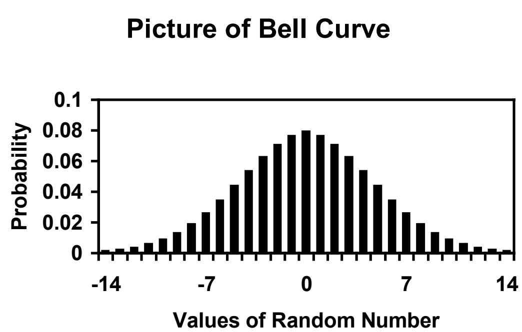

II. Ranges of Random Numbers

Because we are unlikely to generate any single number, it is easier to think about ranges of numbers. If we generate lots of numbers from this bell curve, 95% of the numbers should fall between –10 and 10.

For example, the following twelve numbers were generated from this bell curve:

-1
12
6
-9
1
-4
2
2
-5
0
5
-7

Notice that six of the numbers are positive; five of the numbers are negative; and one of the numbers is 0. Eleven of the numbers are between –10 and 10, the 95% range. The lone exception is the 12.

For each problem, we will generate random numbers from different bell curves. **We will describe these different bell curves to you by giving you the 95% range, the range within which we expect about 95% of the random numbers to fall.** Remember that while about 95% of the numbers will be within the 95% range, most of the numbers will actually be considerably closer to 0. In the few cases where a number falls outside the 95% range, it is most likely to fall very close to the range.

Let's recap the main points you need to remember:

1) The wider the 95% range is, the more likely you are to see very big numbers, either positive or negative.
2) The random numbers are just as likely to be positive as negative.

Please keep these facts in mind. We will now describe the problems you will encounter in this experiment.

III. The Problems

Each problem will consist of 8 rows of 10 random numbers. For example, a problem might look like the following, if we open all the boxes so you can see all the numbers:

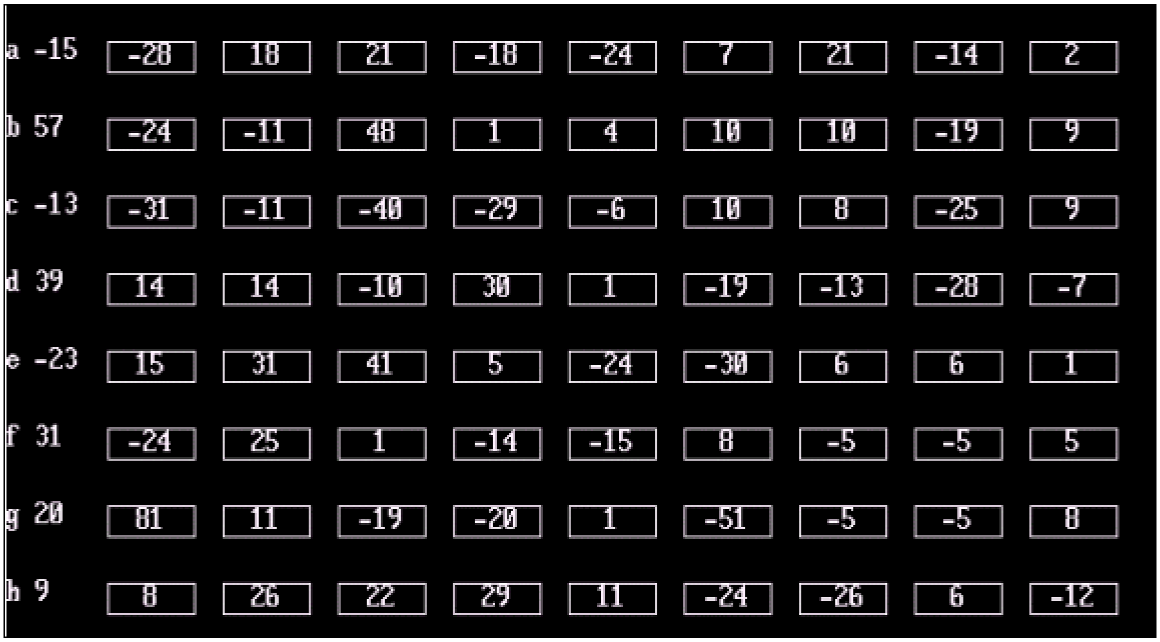

a	-15	-28	18	21	-18	-24	7	21	-14	2
b	57	-24	-11	48	1	4	10	10	-19	9
c	-13	-31	-11	-40	-29	-6	10	8	-25	9
d	39	14	14	-10	30	1	-19	-13	-28	-7
e	-23	15	31	41	5	-24	-30	6	6	1
f	31	-24	25	1	-14	-15	8	-5	-5	5
g	20	81	11	-19	-20	1	-51	-5	-5	8
h	9	8	26	22	29	11	-24	-26	6	-12

You want to choose the row you think has the highest sum.

There are 10 columns in each problem; in the example above, column 1 contains the numbers
–15, 57, -13, 39, -23, etc. The numbers in column 1 are not in boxes; numbers in the other columns are in boxes. In the actual problems, you will need to click on the boxes with your mouse to see those numbers; we will show you how the boxes work in the computer demo.

We generate each column from its own bell curve with its own 95% range. Within a single row, such as row (a), every number is generated from a different bell curve.

The leftmost column (-15, 57, -13, 39, etc) comes from a bell curve that has a 95% range of –90 to 90. The next column (-28, -24, -31, 14, etc) comes from another bell curve, which has a slightly narrower 95% range of -86 to 86. We thus expect numbers from the second column, on average, to be a little closer to 0. The 95% ranges narrow gradually from left to right, so the rightmost column (2, 9, 9, -7, etc) comes from a tenth bell curve, which has a 95% range of -25 to 25. **This means the left-most numbers are most likely to be far from 0, and as you move from left to right the numbers should tend to be closer to 0.**

To review, if you look down the columns of the problem, every number is generated from the same bell curve. As you look left-to-right along the rows of the problem, every number is generated from a different bell curve, going from widest to narrowest. **Before each problem, you will be told the 95% range for the bell curves generating the leftmost and rightmost random numbers in that problem.**

Feel free to skip the following paragraph, provided for those who prefer a more technical explanation:

All you need to understand is that the ranges narrow gradually and smoothly from left to right. In case you are interested and have taken a statistics class, we will tell you the actual process used to generate these ranges. The bounds on the 95% confidence intervals we report are 1.96 times the standard deviation used for the bell curve; the statistics are based on the squares of the standard deviations, or the variances. Moving from left to right the variances fall linearly in increments of 1/10ths. Therefore, the squares of the bounds fall linearly. For example, if the leftmost range is -10 to 10, then its bound is 10 and the square is 100. The square of the bound of the second range is 9/10 times 100, or 90. The actual range comes from the square root of this, or 9.5. Moving to the 3rd column, we take 8/10 of the bound of the first range 100, and then take the square root, for a bound of 8.9. (If this paragraph is confusing, don't worry about it.)

We have randomly generated all the problems as described above; furthermore, we have also randomized the order of the problems, and even the order of the different parts of the experiment. You do not need to spend any mental energy looking for patterns in the problems or other tricks. Economics has a professional standard against deceiving our subjects in any way. The editors of our journals require us to be completely straightforward with you.

IV. Payment

You will begin the experiment with a payment account of $5. **After each problem, we will add the sum of the row you selected to your account (in pennies).** Remember that numbers in the first column, though not covered by boxes, still count towards your payment. Suppose the example problem were the first problem in the experiment. At the beginning of the problem, you have your starting payment account of $5.00. In the example problem, row (a) has a sum of -30. If you selected row (a), your account would decrease from $5 to $4.70. You would then move on to the next problem with an account worth $4.70. Row (b) has a sum of 85. If you selected row (b) instead, your account would increase from $5 to $5.85, and you would start the next problem with an account worth $5.85. **At the end of the experiment we will record the value of your payment account and send you a check for that amount within 72 hours**. So suppose you played a total of 50 problems and in each problem you chose a row with a sum of 50 (worth $.50). Then your final account would be $.50 times 50, plus the initial $5, so you would receive a check for $30. To recap: Your payoff in each problem will be the sum of the boxes in the row you choose, and we will pay you $5 plus the sum of your payoffs from every problem you finish.

There is one friendly exception to this calculation. We guarantee you a payment of $5 if you finish the experiment, even if the total payoff from your choices would be less than $5. We believe if you try hard you can earn $10 to $30. You may choose to leave at any time.

You will now practice on a computer demo. Please pay close attention to the demo instructions, which explain the various twists and turns of our user interface. **Please raise your hand for an experimental assistant.**

Computer Instructions

Subjects received the following instructions on a laptop computer.

DECISION MAKING EXPERIMENT

You will now begin the computer portion of the experiment.

You should have just finished reading a written set of instructions that discusses bell-curves. If not, please raise your hand and notify an assistant.

Before beginning, you will first be guided through some screens
to help you get comfortable with this software.

[next screen]

Remember, you will be shown 8 rows of 10 numbers. Your goal is to choose the row that you think has the highest sum.

To record how you approach these problems, we are using a computer program that will cover the numbers with boxes. Only the left-most numbers in each row will not be covered and will always be visible. To look at the numbers covered by boxes, you must use your mouse to click on the box that you want to see.

Each box covers one number.

[next screen]

To open a box, move the cursor into the box and click the LEFT mouse button.
To close a box, click on the RIGHT mouse button. You do not need to have the cursor in a box to close it. Only one box can be open at a time.

After deciding which row you think has the highest sum, close all of the boxes by right clicking, and then make a choice by clicking on the corresponding button at the bottom of the screen. You MUST close all the boxes before making a choice. You will then be asked to confirm your selection. The buttons at the bottom of the screen are labeled with the letter and initial value for each row.

[next screen]

Please remember the following main points:

1. To open a box -> left click with your mouse
   To close a box -> right click with your mouse
      NOTE: your cursor can be anywhere on the screen to close a box

2. When you are ready to choose a row with the highest sum, first close all of the open boxes by right clicking. Then click on the starting value of the row you believe to have the highest sum and confirm your choice.

3. You will be asked to confirm your choice. To do this, click on the long horizontal button on the bottom of the screen.

4. You cannot make a choice unless all of the boxes are closed.
   Remember to right click to close any open box.

The next screen is a sample for you to practice opening and closing the boxes.

THE NEXT SCREEN IS PRACTICE #1

[The next screen is a demo.]

[Note: this text applies to subjects who play the Exogenous Time segment first, followed by the Endogenous Time segment. This ordering was randomly determined for each subject.]

In Part I of the experiment, you will have a fixed amount of time for each problem. You will have a different amount of time for each game. There will be a round clock in the top right corner of the screen that will show you how much time you have left for each game.

When the clock runs out, all the boxes will close and you will hear a beep indicating that your time has expired. After you hear the beep and are prompted to make a selection, choose the row you think has the highest sum.

If you make a choice before the time has expired, you will be returned to the screen for the remaining amount of time.

WARNING: If there are any boxes open, the clock will temporarily pause
            ...but the time will continue to countdown!!!

[next screen]

The next screen will let you practice making a selection with a time restriction.

You will have 1 1/2 minutes to practice. (This is more time than you will have in the actual problems)

You must wait for the entire 1 1/2 minutes before making a selection. If you make a choice before the time has expired, you will be returned to the screen for the remaining amount of time.

THE NEXT SCREEN IS PRACTICE #2

[The next screen is a demo.]

There will be a 20-second delay between each problem. During these delays, you will see a "buffer screen" instead of a math problem.

During the buffer screen you will be told:
1. The 95% ranges of the left-most columns and right-most columns.
2. The time limit of the next screen
3. How much money you have earned

You must wait for 20 seconds during the buffer screen. After the time has expired and you hear a beep, you will be able to advance to the next problem.

[next screen]

The next practice will involve a 20 second buffer screen, followed by a 45 second fixed time problem.

Remember, you must wait for the time to expire before making your choice.

THE NEXT SCREEN IS PRACTICE #3

[The next screen is a demo.]

You will now begin Part I of the experiment.

You will have 12 problems with randomly selected time limits. For these problems, you must use the entire time before making your selection.

Remember:
- Look on the buffer screens to know how much time you will have for the next game
- You must wait for the time to expire before making your choice

- The higher the sum of the rows you choose, the more money you will earn in this experiment.

GOOD LUCK!!!!

THE NEXT PROBLEMS WILL ALL COUNT TOWARDS YOUR PAYOFF!!!

[The subjects would then begin the experiment. At the end, they take the following survey.]

You have now completed the experiment!
Thank you for your time.

Please take a few minutes to answer a few questions.

[next screen]

What would you say best describes your current concentration or expected concentration at Harvard?

a. Economics
b. History or History and Literature
c. Humanities
d. Language
e. Math, Statistics, or Applied Math
f. Psychology
g. Other Natural Sciences
h. Other Social Sciences

[next screen]

Which of the following years do you expect to graduate from Harvard?
(If you are advanced standing, please choose the year you expect to finish.)

a. 2000
b. 2001
c. 2002
d. 2003
e. 2004
f. 2005

g. Other

[next screen]

Which of the following best describes your class?

a. First Year
b. Sophomore
c. Junior
d. Senior

[next screen]

How old are you?

a. under 18
b. 18
c. 19
d. 20
e. 21
f. 22
g. over 22

[next screen]

How would you best describe your gender?

a. Female
b. Male

[next screen]

We would like to know about your statistical background. Please give us your best assessment of your previous coursework

Pick the HIGHEST number that applies to you.

a. I have never taken a statistics course.

b. I took a statistics class in high school.

c. I took an introductory statistics course such as statistics 100 or statistics 104, or I took the AP statistics test in high school.

d. I have taken a higher-level statistics course such as statistics 110 or higher, economics 1123, or government 1000.

e. I have taken more than one advanced level statistics courses.

[next screen]

Please give your best estimate of your statistical ability, compared to all Harvard undergraduates.

a. I have below average skills in statistics for a Harvard undergraduate.

b. I have about average skills in statistics for a Harvard undergraduate.

c. I have above average skills in statistics for a Harvard undergraduate

[next screen]

Please raise your hand for assistance.

Thank you.

[An experimental assistant must enter a special code to access the next screen.]

Please open the sealed envelope given to you at the beginning of the experiment. Please sit at a desk without a computer and take a few moments to answer the questions found inside the envelope. Your answers will be helpful for us to study how you approached the problems in this experiment.

Final score:

[The subject's final score (i.e. total payment) would now be displayed.]

Questionnaire

Before leaving, subjects filled out the following written questionnaire.

Subject #:

Questionnaire

1. Please describe how you approached the problem of finding the best rows in the problems you just analyzed. Your detailed response to this question will be very helpful to us. Please feel free to use the back of this form if needed.

2. Has any part of this experiment been confusing? Please explain.

7. Appendix B: Derivation of equation (2)

To gain intuition for equation (2), begin by assuming that $b \geq a$. In this case,

$$\max(a', b) - \max(a, b) = \begin{cases} 0 & \text{if} \quad b > a + \varepsilon \\ a + \varepsilon - b & \text{if} \quad b \leq a + \varepsilon \end{cases}$$

As ε is drawn from a Normal$(0, \sigma^2)$ distribution, integrating over the relevant states yields the right-hand-side of equation (2):

$$\int_{b-a}^{\infty} (a + \varepsilon - b)\, \phi\left(\frac{\varepsilon}{\sigma}\right) \frac{d\varepsilon}{\sigma} = \int_{|a-b|}^{\infty} (\varepsilon - |a-b|)\, \phi\left(\frac{\varepsilon}{\sigma}\right) \frac{d\varepsilon}{\sigma}$$

We can reexpress the option value formula by noting that $x\phi(x) = -\phi'(x)$, so that:

$$\begin{aligned} \int_{|a-b|}^{\infty} (\varepsilon - |a-b|)\, \phi\left(\frac{\varepsilon}{\sigma}\right) \frac{d\varepsilon}{\sigma} &= \int_{|a-b|}^{\infty} \varepsilon\phi\left(\frac{\varepsilon}{\sigma}\right) \frac{d\varepsilon}{\sigma} - |a-b| \left(1 - \Phi\left(\frac{|a-b|}{\sigma}\right)\right) \\ &= \sigma\phi\left(\frac{|a-b|}{\sigma}\right) - |a-b|\, \Phi\left(-\frac{|a-b|}{\sigma}\right) \\ &= w\left(|a-b|, \sigma\right). \end{aligned}$$

Likewise, when $b \leq a$,

$$\max(a', b) - \max(a, b) = \begin{cases} b - a & \text{if} \quad b > a + \varepsilon \\ \varepsilon & \text{if} \quad b \leq a + \varepsilon \end{cases}$$

Integrating over the relevant states again yields the right-hand-side of equation (2):

$$\begin{aligned} (b-a)\Phi\left(\frac{b-a}{\sigma}\right) + \int_{b-a}^{\infty} \varepsilon\phi\left(\frac{\varepsilon}{\sigma}\right) \frac{d\varepsilon}{\sigma} &= (b-a)\Phi\left(\frac{b-a}{\sigma}\right) + \sigma\phi\left(\frac{b-a}{\sigma}\right) \\ &= w\left(|b-a|, \sigma\right). \end{aligned}$$

8. Appendix C: Example calculation of G

Consider a game where the payoffs in column 1 are drawn from a $N(0, \sigma^2)$ distribution. The variance of payoffs in column m is then $\sigma_m^2 = \sigma^2 (11 - m)/10$, for $m \in \{1...10\}$. Calling b_{nm} the payoff in row n, column m, the current expected value of row n is

$$a_n = \sum_{\text{opened boxes } m \text{ in row } n} b_{nm}$$

For instance, suppose that at the current time in row n a subject has already opened boxes $\{1, 2, 4, 7, 8\}$ and that boxes $\{3, 5, 9, 10\}$ remain unopened. The current value of row n is

$$a_n = b_{n1} + b_{n2} + b_{n4} + b_{n7} + b_{n8}$$

Call a_{-nt} the value of the best competitor of n. In most instances a_{-nt} is the value of the leading row, except if n is the leading row itself, in which case a_{-nt} is the value of the second leading row.

The relevant cognitive operators for exploring row n are 'open box 3,' 'open boxes 3 and 5,' 'open boxes 3, 5, and 9' and 'open boxes 3, 5, 9, and 10." Their costs are respectively 1, 2, 3 and 4. So, if $x = a_{nt} - a_{-nt}$, then the value of the benefit/cost ratio G_n of row n is, $G_n =$

$$\max\left\{\frac{w(x,\sigma_3)}{1}, \frac{w\left(x, \sqrt{\sigma_3^2+\sigma_5^2}\right)}{2}, \frac{w\left(x, \sqrt{\sigma_3^2+\sigma_5^2+\sigma_9^2}\right)}{3}, \frac{w\left(x, \sqrt{\sigma_3^2+\sigma_5^2+\sigma_9^2+\sigma_{10}^2}\right)}{4}\right\}$$

and the value of G at the current time is simply the highest possible benefit/cost ratio across all rows:

$$G = \max_n G_n.$$

9. References

Bernatzi, Shlomo and Richard Thaler. "Naive Diversification in Defined Contribution Savings Plans", *Amer. Econ. Rev.*, March 2001, Vol. 91(1), pp. 79-98.

Camerer, Colin, Eric J. Johnson, Talia Rymon, and Sankar Sen, (1993) "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.

—- and Daniel Lovallo. "Overconfidence and Excess Entry: an Experimental Approach." *American Economics Review* (1999), March, 89, pp306-318.

Camerer, Colin and Ho, Teck-Hua (1999) "Experience weighted attraction learning in normal-form games." *Econometrica* 67 (1999), pp. 837-874.

Conlisk, John. "Why bounded rationality?" *Journal of Economics Literature*, June 1996, *34* (2), pp.669-700.

Costa-Gomes, Miguel, Vincent Crawford, and Bruno Broseta. "Cognition and Behavior in Normal-Form Games: An Experimental Study." *Econometrica*, 2001, 69(5): 1193-1235.

— and —. "Cognition and Behavior in Two Person Guessing Games: an Experimental Study." July 2002. University of California, San Diego mimeo.

Erev, Ido and Roth, Alvin. "Predicting how people play games." *American Economic Review*, Sept. 1998, 88 (2), pp.848-81.

Gabaix, Xavier and David Laibson. "Bounded Rationality and Directed Cognition," mimeo, 2002.

Gigerenzer, Gerd, Peter Todd, and the ABC Research Group (1999). "Simple Heuristics that Make Us Smart," Oxford University Press.

Gittins, J.C. (1979) "Bandit processes and dynamic allocation indices", *Journal of the Royal Statistical Society B*, 41, 2, p.148-77.

Jehiel, Philippe. (1995) "Limited horizon forecast in repeated alternate games," *Journal of Economic Theory*, 67: 497-519.

Kahneman, Daniel and Tversky, Amos. "Judgment Under Uncertainty: Heuristics and Biases". *Science*, 1974.

Massey, B. Cade and George Wu. "Detecting Regime Shifts: The Psychology of Over- and Under-reaction," mimeo, University of Chicago Graduate School of Business, 2002.

Payne, John W., James R. Bettman, and Eric J. Johnson, *The Adaptive Decision Maker*, 1993, Cambridge: Cambridge University Press.

—, M.L. Braunstein, and J.S. Carroll. "Exploring Predecisional Behavior: an Alternative Approach to Decision Research." *Organizational Behavior and Human Performance* (1978), 22, pp17-44.

Russo, J. Edward. "Eye Fixations can Save the World: a Critical Evaluation and a Comparison between Eye Fixations and other Information Processing Methodologies." *Memory and Cognition* (1978), 3(3), pp267-276.

Simon, Herbert. "A behavioral model of rational choice," *Quarterly Journal of Economics*, Feb 1955, 69 (1).

Tversky, Amos. "Elimination by aspects: A theory of choice." *Psychological Review, 1972,* 79:281-299.

Weitzman, Martin (1979) "Optimal search for the best alternative", *Econometrica* 47, 3, p.641-54.

Table 1: Evaluation of Alternative Models

	Games with Exogenous Time Budgets			Games with Endogenous Time Budgets					
	Directed Cognition	Column Model	Row Model	**Directed Cognition**	Directed Cognition	Column Model Satisficing	Row Model Satisficing	Directed Cognition Satisficing	Elimination by Aspects
Free parameters?	**No**	No	No	**No**	$\kappa = 0.18$	S = 42	S = 43	S = 45	A = -6
R'^2 for column profile	**86.6%**	-106.6%	39.3%	**99.0%**	85.5%	67.8%	-16.3%	43.0%	-40.4%
R'^2 for row profile	**-14.6%**	26.2%	-539.4%	**87.4%**	70.5%	36.7%	82.5%	75.4%	4.0%
R'^2 for alt. row profile	**87.0%**	25.0%	89.2%	**96.3%**	91.9%	29.2%	48.4%	58.7%	41.5%
Average R'^2	**53.0%**	-18.5%	-137.0%	**94.2%**	82.6%	44.5%	38.2%	59.1%	1.7%
Correlation between Predicted and Empirical Number of Boxes per Game	**---**	---	---	**0.69**	0.64	0.00	0.02	0.03	-0.38

R'^2 refers to the R^2 statistic from a regression of the empirical number of boxes opened by subjects in each row or column on a constant plus the number of boxes predicted by the models, where the co-efficient on the predictions is fixed to equal 1.

Correlation refers to the correlation between the number of boxes opened by subjects in each of the 160 games and the number of box openings predicted by the models.

Figure 1: Sample Game with All Values Unmasked

In the actual experiment, the subjects see only the value of one box at a time (see Figure 2). Values are all drawn independently from normal distributions. Values in each column are drawn from a distribution with the same variance, with variances declining linearly from left to right. In this sample, the left-most column is generated with a standard deviation of 30.6 cents, which is explained to subjects as a 95% confidence interval of –60 to 60 cents.

Figure 2: Sample Game with Values Concealed

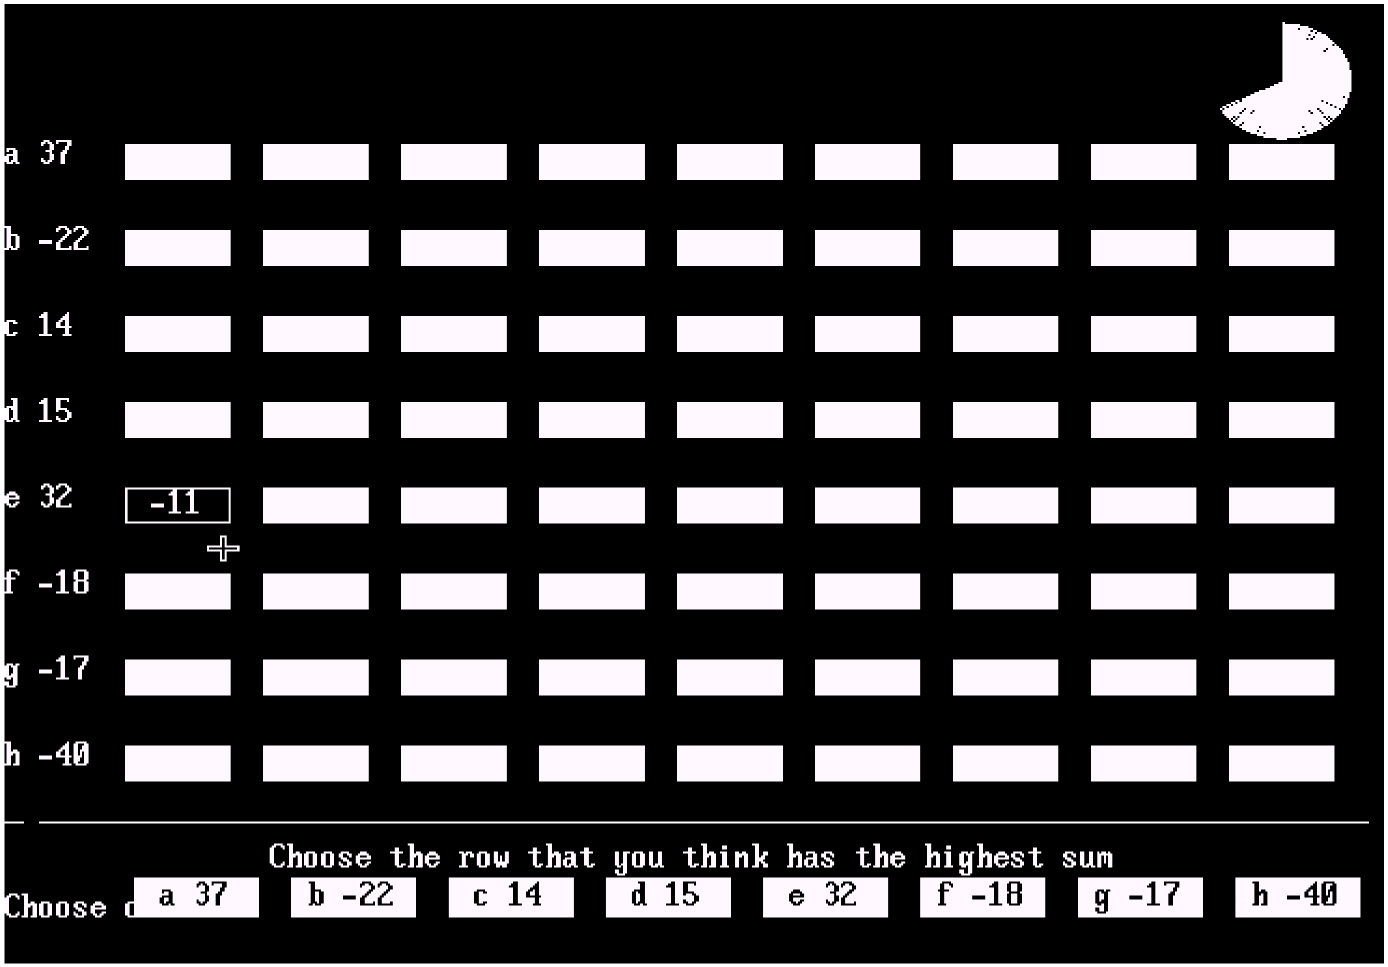

This is how a sample game would appear to subjects, with values concealed by boxes. Subjects can use the mouse to open one box at a time. In this game the subject faces a set time limit; the clock in the upper right corner reveals the fraction of time remaining. The preceding screen will have told subjects the time limit and the distribution structure for this game.

Figure 3: w(x,σ) for σ=1

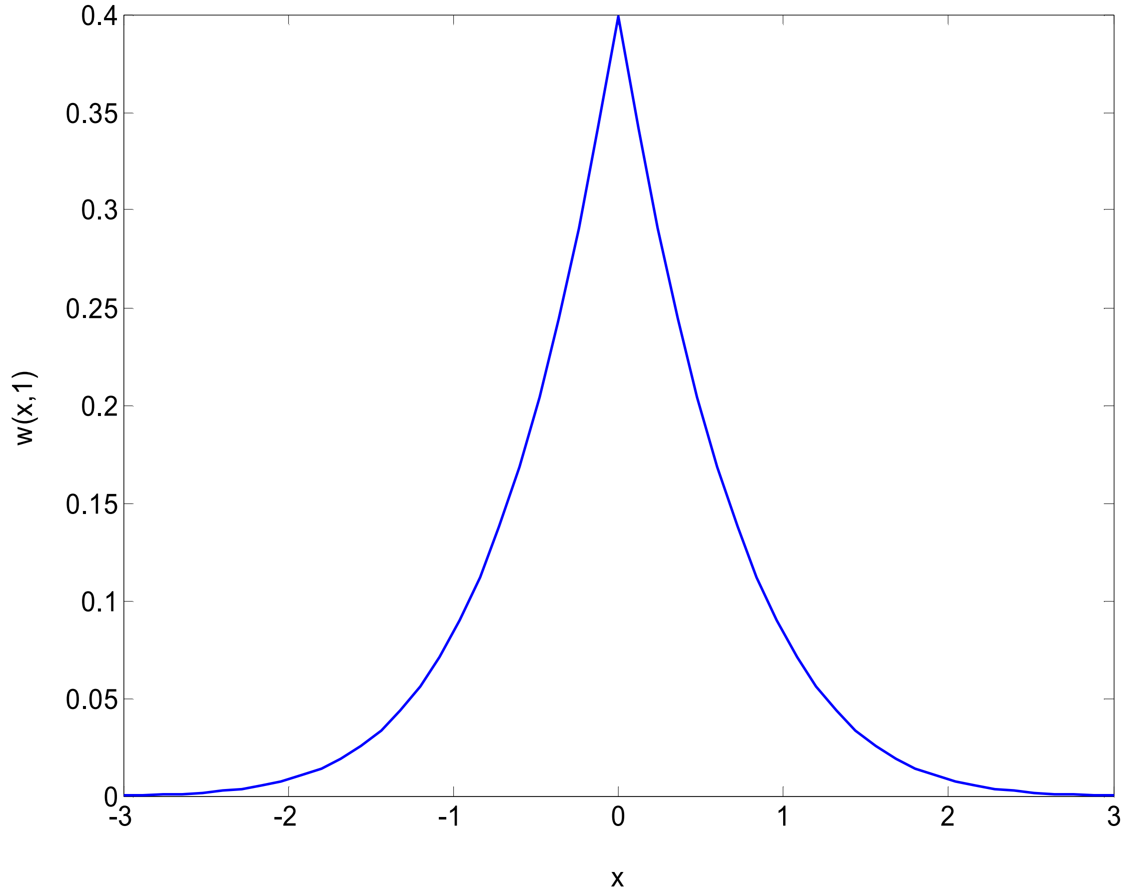

Figure 3 plots the expected benefit from continued search on an alternative if the difference between the value of the searched alternative and its best alternative is x, and the standard deviation of the information gained is σ=1, as defined in Equation 3. This w function is homogeneous of degree one.

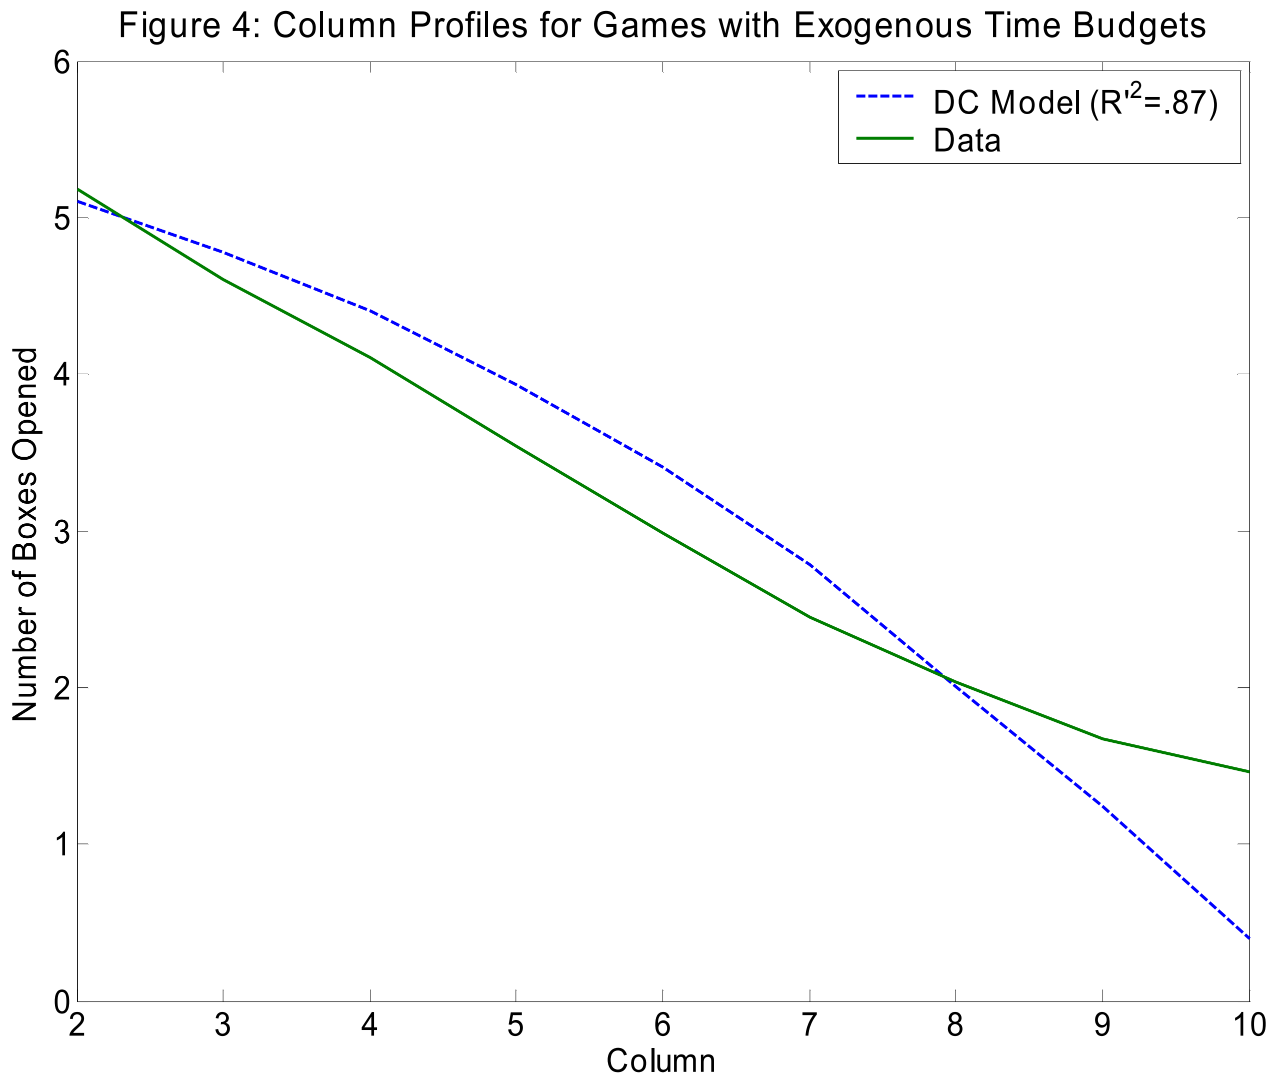

Figure 4 plots the mean number of boxes opened by subjects in each column from 2 to 10 and the predicted number of boxes opened by the Directed Cognition algorithm. The algorithm is simulated on the same 4656 games played by the subjects.

Figure 5: Row Profiles for Games with Exogenous Time Budgets

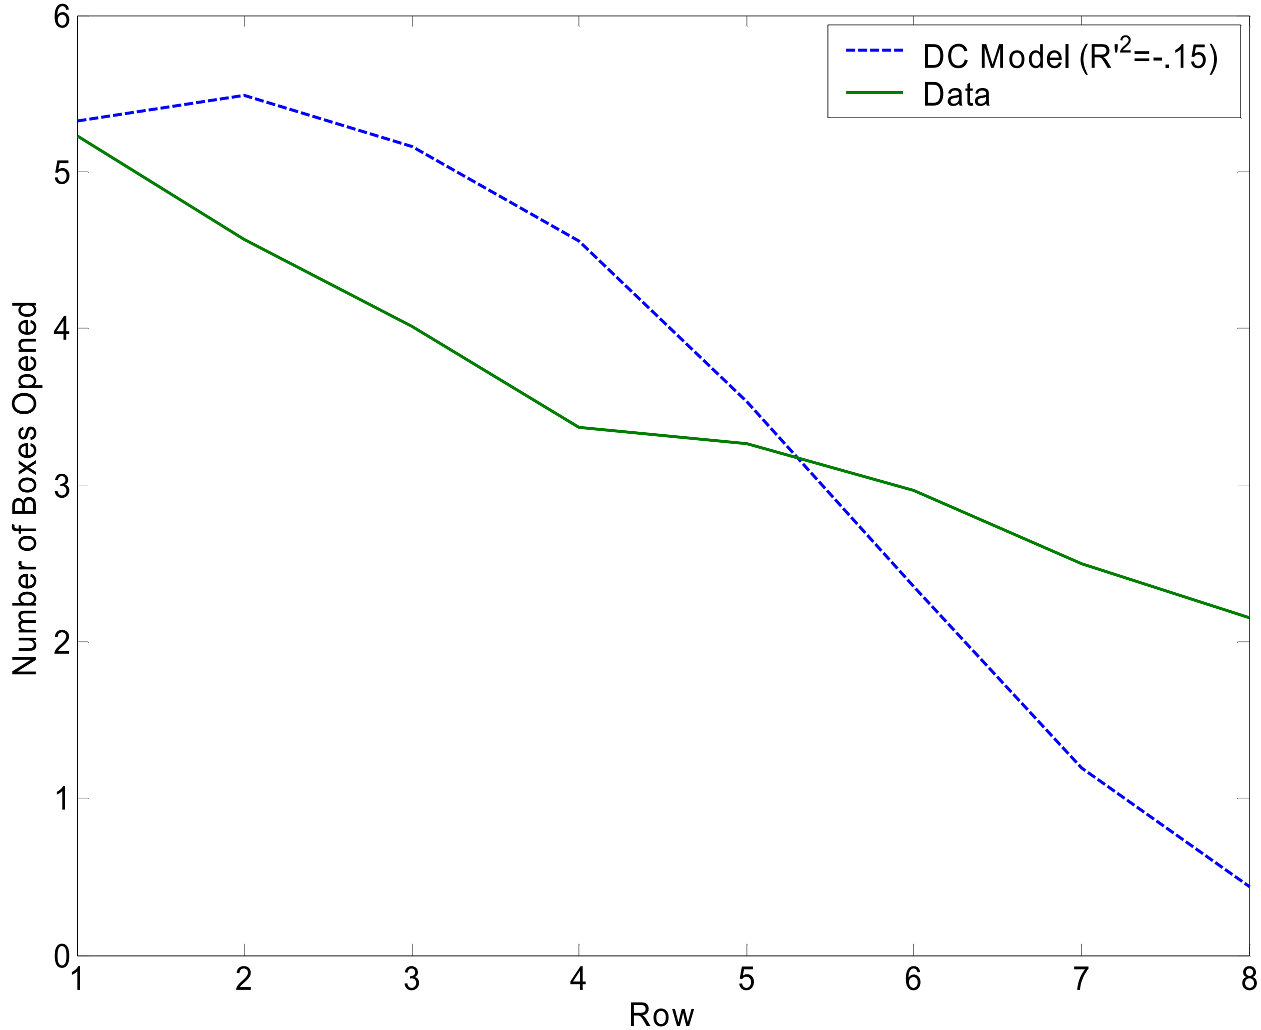

Figure 5 plots the mean number of boxes opened by subjects in each row and the predicted number of boxes opened by the Directed Cognition algorithm, simulated on the same games played by the subjects. The rows are ordered according to the values in the initial column.

Figure 6: Row Profiles (alternative ordering) for Games with Exogenous Time Budgets

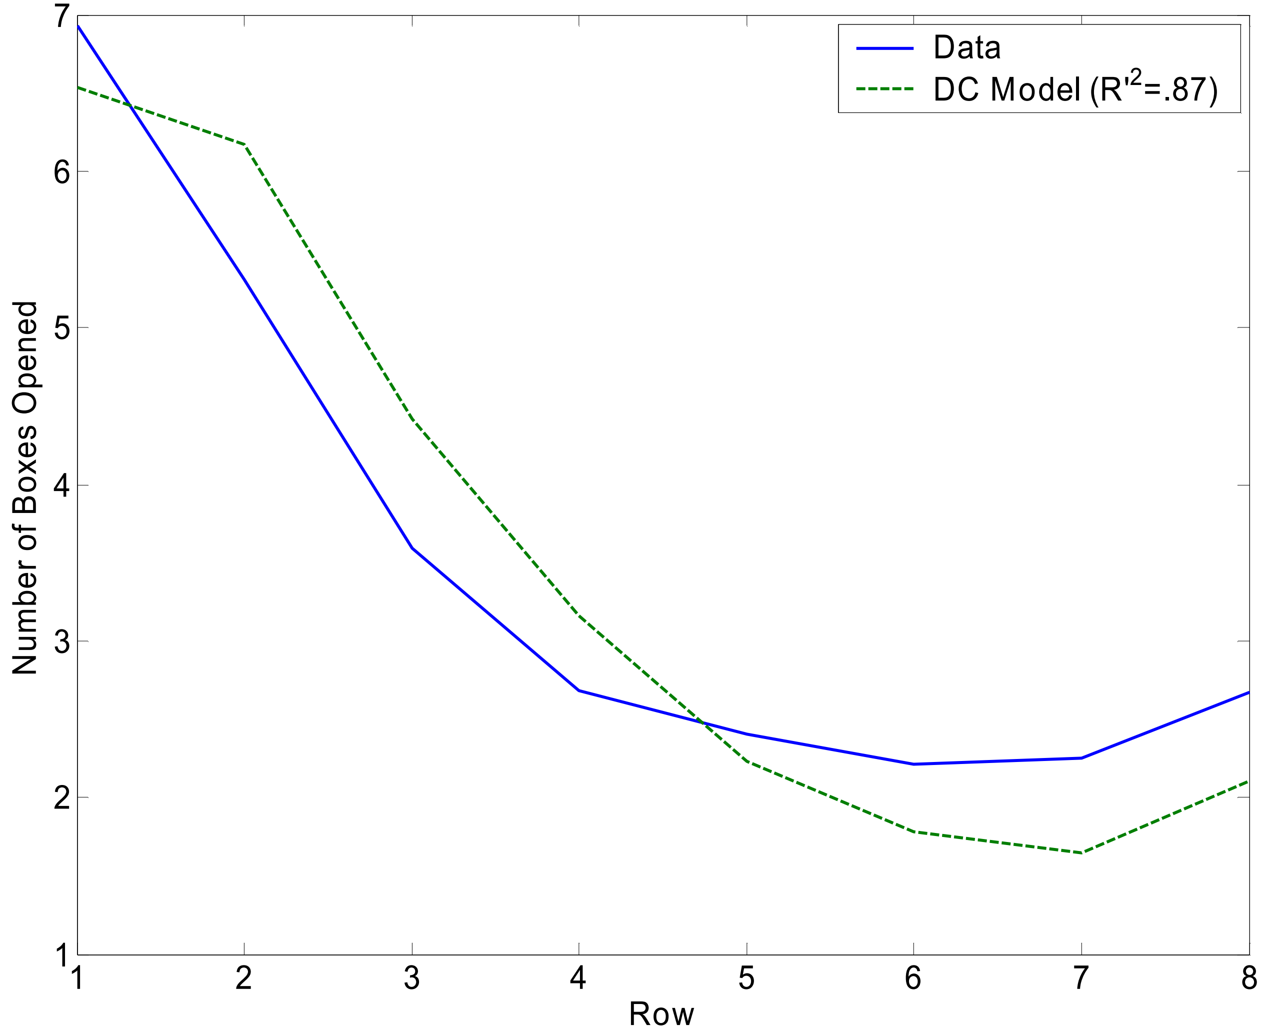

Figure 6 plots the mean number of boxes opened by subjects in each row and the predicted number of boxes opened by the Directed Cognition algorithm. In this figure, the rows are ordered according to all the information available to subjects after they have concluded their search, just prior to making a choice.

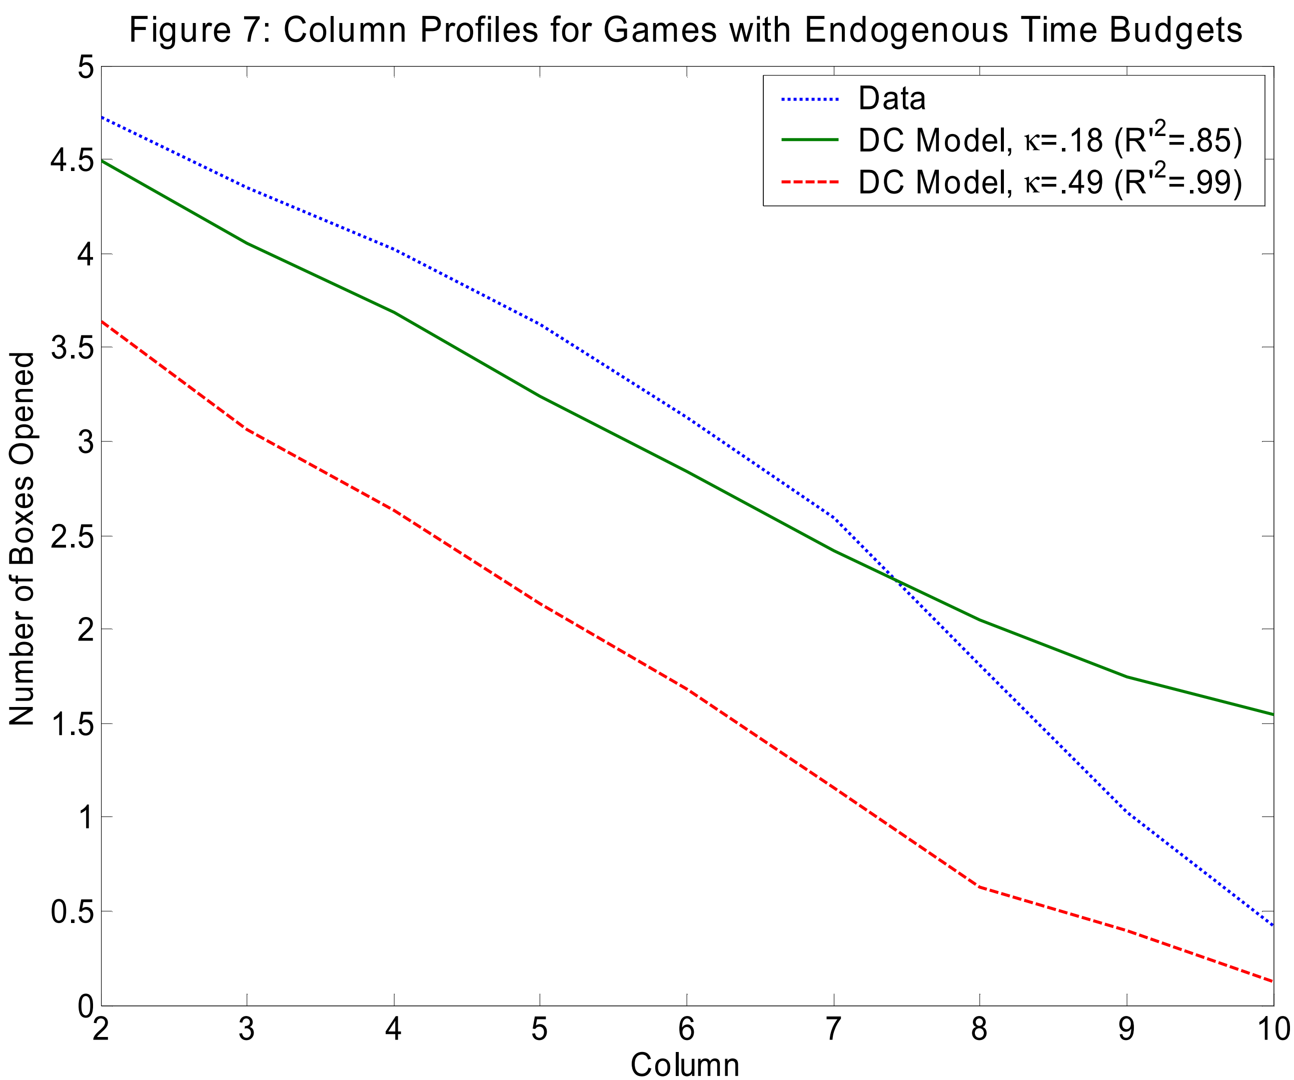

Figure 7 plots the mean number of boxes opened by subjects in each column from 2 to 10 and the predicted number of boxes opened by the Directed Cognition algorithm, for games in which agents choose how much time to allocate from a fixed multiple-game time budget.

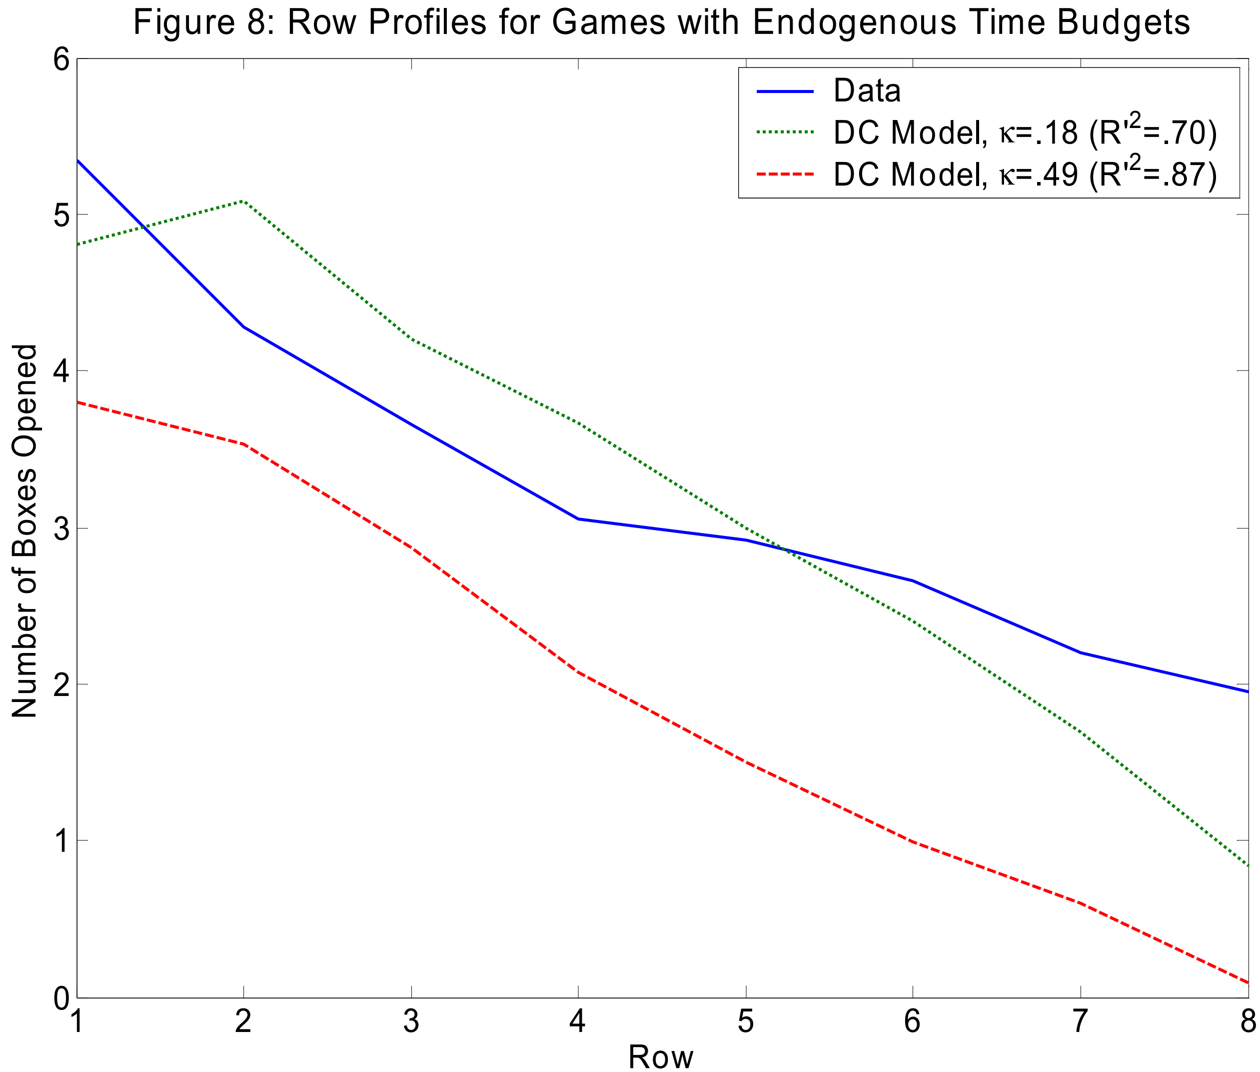

Figure 8 plots the mean number of boxes opened in each row by subjects and by both calibrations of the Directed Cognition model, for games in which agents choose how much time to allocate from a fixed multiple-game time budget. The rows are ordered according to the values in the initial column.

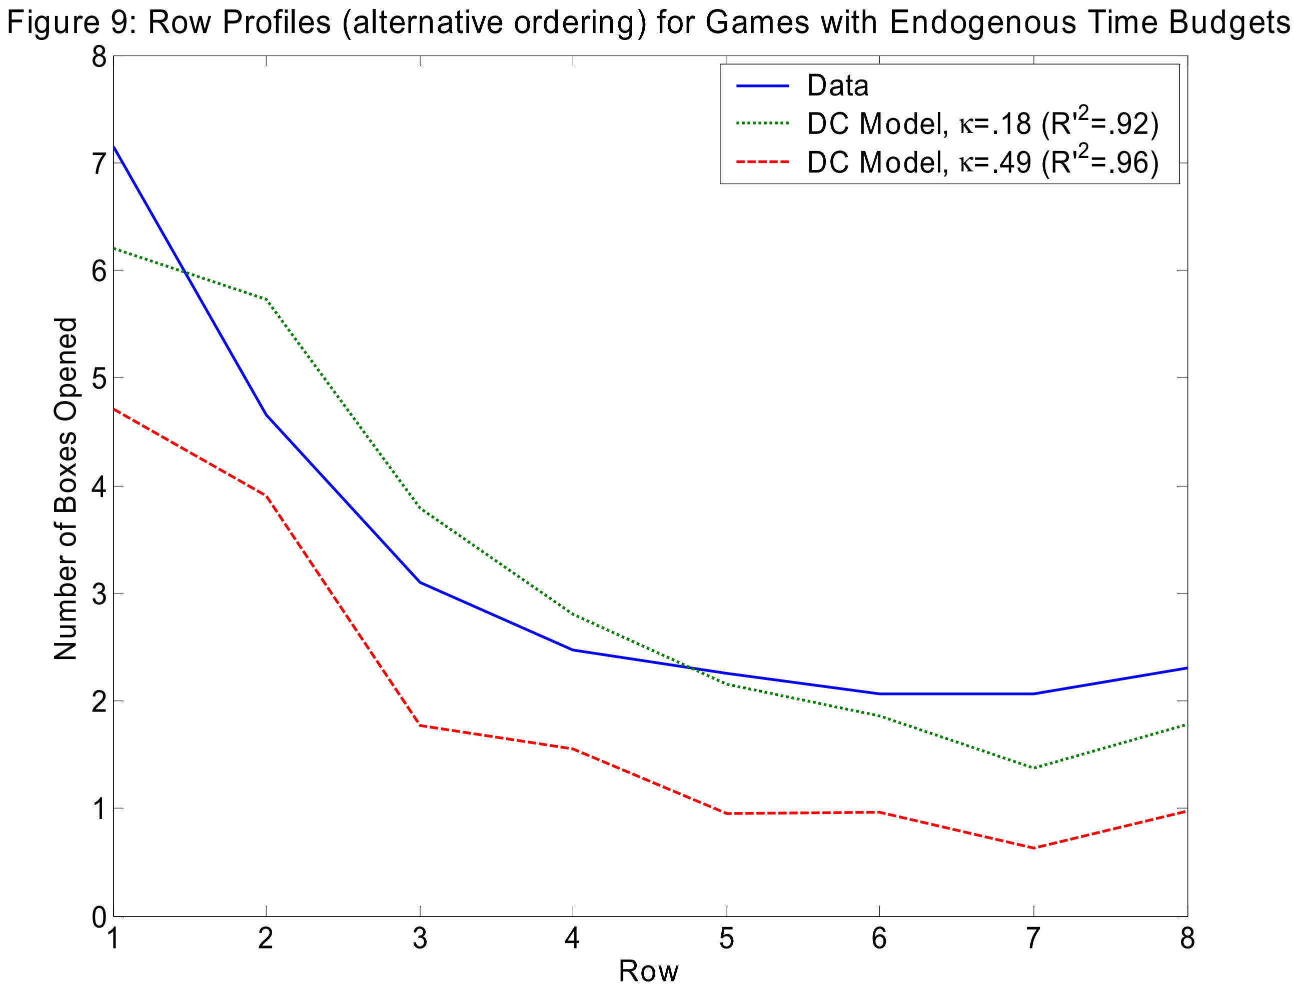

Figure 9 plots the mean number of boxes opened in each row by subjects and by both calibrations of the Directed Cognition model, for games in which agents choose how much time to allocate from a fixed multiple-game time budget. In this figure, the rows are ordered according to all the information available to subjects after they have concluded their search, just prior to making a choice.

Figure 10: Boxes Opened by Game

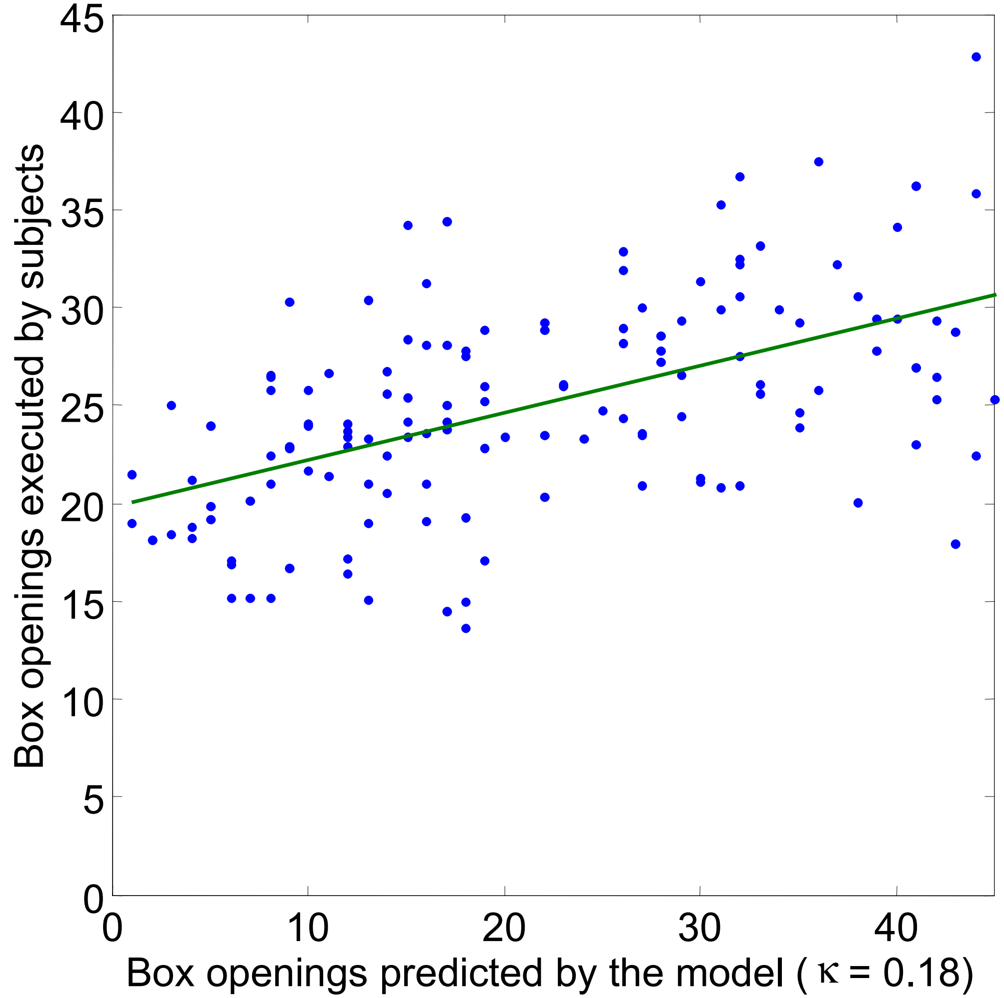

Figure 10 takes each of the 160 game-types played by subjects and compares the number of boxes opened by subjects to the number of box openings predicted by the model, for the games in which agents choose their cross-game time allocations. The correlation between the model and the data is .64.

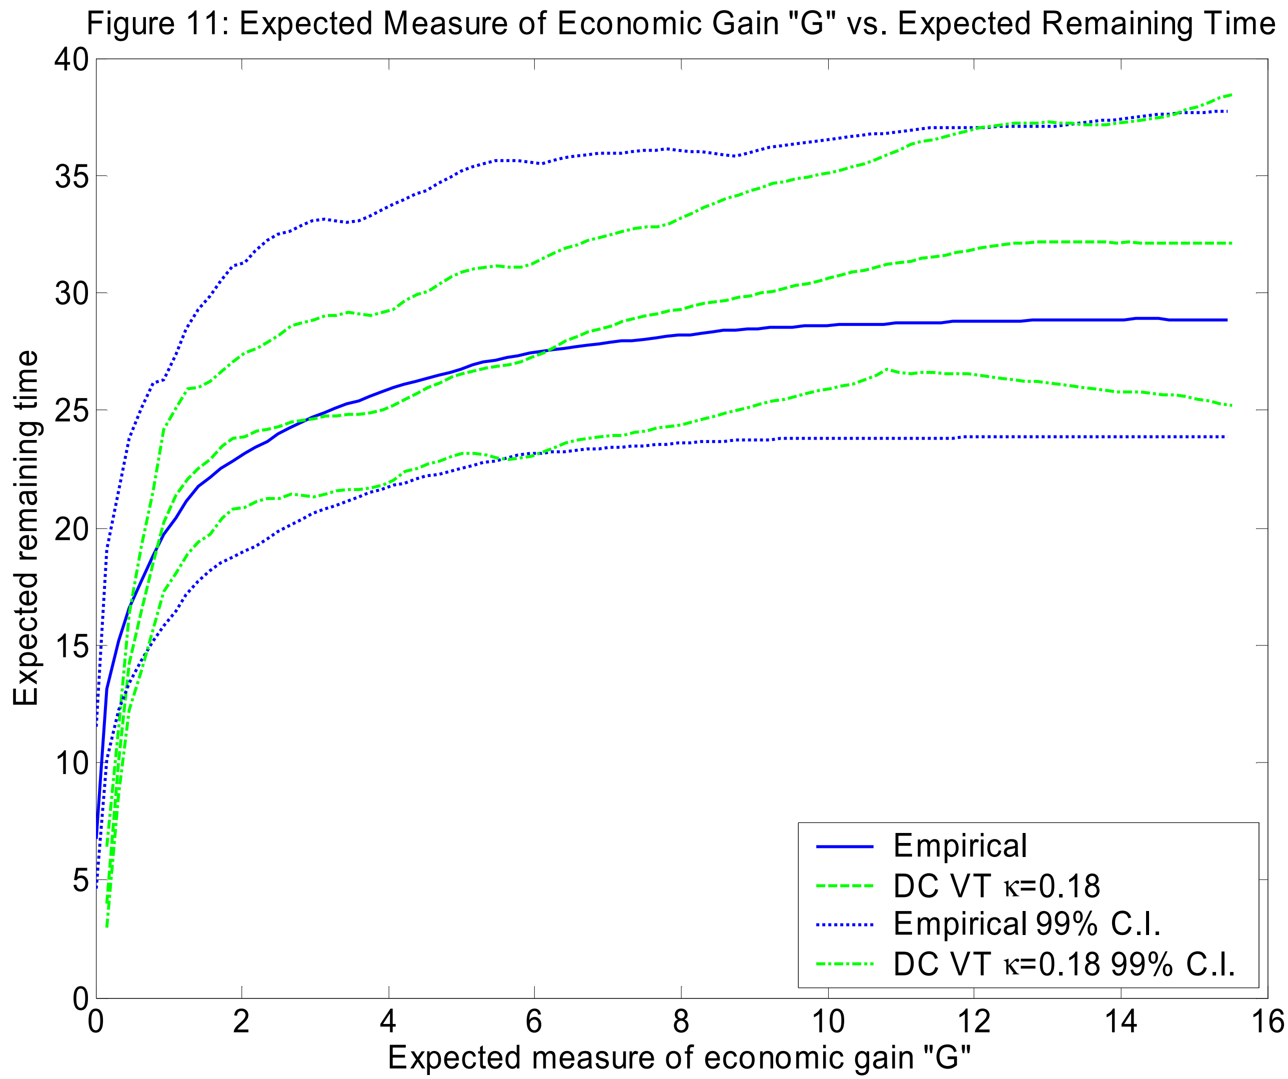

Figure 11 plots non-parametric (kernel) estimates of the expected number of additional box openings in a game, conditional on the current "G" value. "G" is the benefit-to-cost ratio of marginal analysis in the Directed Cognition model, as defined in equation (xx).

Table A1: Total Payoff by Major and Statistics Coursework

	Coursework in Statistics					
	A	B	C	D	E	All
Economics	30.4 (n=20)	30.1 (n=4)	27.8 (n=53)	30.5 (n=30)	28.6 (n=2)	29.1 (n=109)
History, History & Literature	29.3 (n=23)	30.7 (n=4)	29.1 (n=7)	-- (n=0)	-- (n=0)	29.5 (n=34)
Humanities	28.9 (n=31)	31.5 (n=4)	29.0 (n=6)	-- (n=0)	-- (n=0)	29.2 (n=41)
Language	33.9 (n=1)	-- (n=0)	28.8 (n=1)	-- (n=0)	-- (n=0)	31.4 (n=2)
Math, Statistics, Applied Math	28.6 (n=15)	35.9 (n=2)	27.6 (n=5)	33.1 (n=12)	32.4 (n=7)	30.8 (n=41)
Psychology	27.4 (n=8)	30.6 (n=1)	29.3 (n=11)	24.6 (n=2)	-- (n=0)	28.2 (n=22)
Other Natural Sciences	28.3 (n=52)	29.9 (n=7)	29.9 (n=20)	30.7 (n=2)	30.2 (n=1)	28.9 (n=82)
Other Social Sciences	29.1 (n=25)	27.3 (n=9)	29.6 (n=20)	-- (n=0)	-- (n=0)	29.0 (n=54)
All	28.9 (n=175)	29.9 (n=31)	28.7 (n=123)	30.9 (n=46)	31.4 (n=10)	

A=none
B=high school statistics
C=an introductory college-level course in statistics
D=an advanced college-level course in statistics
E=more than one advanced college level course in statistics

Table A2: Total Payoff by Gender and Statistics Coursework

	Coursework in Statistics					
	A	B	C	D	E	All
Female	26.5 (n=67)	29.1 (n=14)	27.5 (n=38)	27.7 (n=10)	27.4 (n=1)	27.2 (n=130)
Male	30.4 (n=108)	30.5 (n=17)	29.3 (n=85)	31.8 (n=36)	31.9 (n=9)	30.3 (n=255)
All	28.9 (n=175)	29.9 (n=31)	28.7 (n=123)	30.9 (n=46)	31.4 (n=10)	

A=none
B=high school statistics
C=an introductory college-level course in statistics
D=an advanced college-level course in statistics
E=more than one advanced college level course in statistics