Published in final edited form as:

J Eur Econ Assoc. 2009 ; 7(2): 628–637.

The allocation of time in decision-making*

Christopher F. Chabris,
Union College

David Laibson,
Harvard University

Carrie L. Morris,
Washington University in St. Louis

Jonathon P. Schuldt, and
University of Michigan

Dmitry Taubinsky
Harvard University

Abstract

We study the allocation of time across decision problems. If a decision-maker (1) has noisy estimates of value, (2) improves those estimates the longer he or she analyzes a choice problem, and (3) allocates time optimally, then the decision-maker should spend less time choosing when the difference in value between two options is relatively large. To test this prediction we ask subjects to make 27 binary incentive-compatible intertemporal choices, and measure response time for each decision. Our time allocation model explains 54% of the variance in average decision time. These results support the view that decision-making is a cognitively costly activity that uses time as an input allocated according to cost-benefit principles.

1. Introduction

Agents should allocate time according to cost-benefit tradeoffs. Such tradeoffs should also influence the amount of time that agents allocate to *decision-making itself*. It is natural to assume that agents have expected utility estimates that are noisy and that these noisy estimates can be improved by thinking longer about a choice problem. In this case, agents should allocate more decision time to choices between options of similar expected utility than to choices between options of dissimilar expected utility.

Intuitively, if an agent's estimate of the value of the best option is close to her estimate of the next best alternative, then it is optimal to take more time to refine her estimates since this extra time has a good chance of leading her to change her choice. In contrast, if an agent has a noisy estimate of the value of an option but the option almost surely dominates its alternatives, then it is not worth taking the time to refine her estimates, because additional thinking is unlikely to change her choice. These intuitions generalize, predicting an inverse relationship between

*We thank Kirill Babikov, Lee Chung, Alison H. Delargy, Margaret E. Gerbasi, J. Richard Hackman, Stephen M. Kosslyn, Melissa A. Liebert, Sarah Murphy, Jacob Sattelmair, Aerfen Whittle, and Anita W. Woolley for their assistance with and support for this research. We thank seminar participants at USC and the EEA for helpful comments.

Address correspondence to: Christopher F. Chabris, Department of Psychology, Union College, 807 Union Street, Schenectady, NY 12308; chabris@gmail.com.

average decision time and the difference in expected value between rewards in a binary choice set (Gabaix and Laibson, 2005; Gabaix et al., 2006). The current paper tests these predictions.

Psychologists have studied the amount of time that subjects take to make comparative judgments since the work of Henmon (1906). In perception, Johnson (1939) had subjects decide which of two simultaneously presented lines was longer, and found that response time (RT) was longer the more similar the lines: RT was a negative linear function of the logarithm of the difference in length. This pattern has been found with other perceptual quantities like area and luminance, with the size of visualized objects, and even with the magnitudes of abstract numbers (Moyer and Landauer, 1967). Brain imaging studies suggest that all of these comparison tasks invoke a common cognitive process localized in the parietal lobes (e.g., Pinel et al., 2004).

In the experiments reported in the current paper, we find that response times are longer the more similar the discounted values of the rewards in a binary choice set. We ask laboratory subjects to make 27 binary intertemporal choices. We measure intertemporal preferences using data from these choices. We find that the average RT for each of the choices is predicted by the implied difference in discounted expected value. Choices with relatively large differences in discounted value have low average RTs. Differences in discounted value account for 54.1% of the variance in average RT. When we estimate different discount rates for different reward size categories, differences in discounted value account for 69.9% of the RT variance.

2. Data description

We created a computerized version of a 27-question discounting task developed by Kirby et al. (1999). Each question asks the subject to choose between a smaller, immediate reward (SIR) and a larger, delayed reward (LDR), both denominated in U.S. dollars. For example, the first question asks “Would you prefer \$54 today, or \$55 in 117 days?” Rewards range from \$11 (the smallest SIR) to \$85 (the largest LDR). Delays range from 7 to 186 days. Nine trials involve “small” LDRs (\$25–35), nine involve “medium” LDRs (\$50–60), and nine involve “large” LDRs (\$75–85). Following Kirby et al. (1999), we refer to these sets of nine trials as reward size categories. We administered the questions as described by Kirby et al. (1999), except that we also recorded the time that each subject took to answer each question (with a keypress), starting from the time it was displayed on the screen.

This task was included in three separate studies that yielded data from a total of 712 subjects: (1) The *Weight* study examined associations between body mass index (BMI) and discounting, as well as other measures of reward-related behavior. (2) The *Cognition* study examined individual differences in cognitive abilities, decision-making, and personality. (3) The *Web* study examined differences in cognition and personality that may be associated with differences in academic disciplines and career choices.¹

In all three studies, each subject had a 1-in-6 chance of having one randomly-selected question played out for real stakes. In the *Weight* and *Cognition* studies, the subject rolled a six-sided die at the end of the testing session. If a 6 was rolled, the subject blindly drew from a box containing cards labeled 1–27, corresponding to the 27 trials of the discounting task. If the subject chose the SIR on the randomly-drawn trial, the amount of the SIR was added to the

¹In total 751 subjects were tested, but 18 subjects were excluded from the *Weight* study and 16 were excluded from the *Cognition* study for the following reasons: reported brain injury; reported mental illness; reported drug use; had difficulty understanding directions; was unable to complete the protocol; was previously tested. Five subjects from the *Web* study were excluded because they had three or more responses under 200 milliseconds on the discounting task. The average ages of remaining participants in the *Weight*, *Cognition*, and *Web* studies were, respectively, 32, 27, and 31 years; the average numbers of years of education were 15.0, 14.3, and 18.4; the percentages of female participants were 50%, 52%, and 67%.

subject's show-up fee for participating, and a check request for the total was submitted to the research administration office within one business day. If the subject chose the LDR, a separate check request for the LDR amount was made after the specified delay. In the Web study, a spreadsheet was used to generate the necessary random numbers, and “winning” subjects received payment through an Amazon.com gift certificate, which was e-mailed by the next business day (SIRs) or after the specified delay (LDRs).

In the Weight and Cognition studies, subjects were tested in the laboratory; for the Web study, they completed the task independently via the internet using their own computers.² Because the studies varied in recruitment strategies, participant characteristics, payment methods, and apparatus used to measure response time, we analyze them separately in this paper.

3. Models and estimation techniques

3.1. Discount function

We use a generalized hyperbola to model time preferences (Mazur 1987). The discounted value of a one unit reward delayed by τ days is

$$D(\tau) = \frac{1}{1 + \alpha\tau}.$$

The discount rate—the rate of decline of the discount function—is decreasing with τ . To simplify notation, we refer to parameter α as the “discount rate,” though it is actually the instantaneous discount rate at horizon $\tau = 0$.

3.2. Choice data model

We assume that time preferences are homogeneous within each of our three datasets. A subject chooses delayed reward Y in τ days over an immediate reward X if and only if the net present value of Y exceeds X :

$$\frac{Y}{1 + \alpha\tau} - X \geq 0.$$

We assume that subjects experience preference shocks with the logistic distribution, so subjects choose the delayed reward with probability

$$F_{\text{Logit}} \left(\frac{Y}{1 + \alpha\tau} - X \right) = e^{\frac{\omega Y}{1 + \alpha\tau}} / \left(e^{\omega X} + e^{\frac{\omega Y}{1 + \alpha\tau}} \right), \quad (1)$$

where $1/\omega$ is the variance of the logistic distribution. The parameters, $\hat{\alpha}$ and $\hat{\omega}$, are estimated in each dataset by maximizing the likelihood function

²Complete details of the study procedures, including software used, can be found in Chabris et al. (2008a). Examination of the data suggested that response time was measured with similar accuracy in all three studies.

$$L(\alpha|d) = \prod_{i=1}^N \prod_{t=1}^T \left[F_{\text{Logit}} \left(\frac{Y_t}{1+\alpha\tau_t} - X_t \right) \right]^{d_{it}} \left[1 - F_{\text{Logit}} \left(\frac{Y_t}{1+\alpha\tau_t} - X_t \right) \right]^{1-d_{it}} \quad (2)$$

Where T is the number of trials, N is the number of subjects, and $d = \{d_{it}\}_{1 \leq i \leq N, 1 \leq t \leq T}$ is the set of the NT binary decisions made by the N subjects. In trial t , dummy variable $d_{it} = 1$ if subject i chose the delayed reward Y_t in τ_t days over the immediate reward X_t (otherwise $d_{it} = 0$).

3.3. Response time model

Gabaix and Laibson (2005) and Gabaix et al. (2006) predict that decisions should take the longest when the two options are similar in (discounted) value. Formally, subjects should take the most time on the trials for which the difference in the discounted values of the choices,

$$\Delta_t \equiv \left| \frac{Y_t}{1+\alpha\tau_t} - X_t \right|, \quad (3)$$

is closest to 0. The proxy for decision difficulty, $\Gamma(\Delta_t) \in [0,1]$, that we use is given by the transformation

$$\Gamma(\Delta_t) \equiv 2 / (1 + e^{\omega\Delta_t}). \quad (4)$$

This transformation has three desirable properties. First, Γ is convex and decreasing, which means that as the difference between the discounted values of the options grows larger, Γ exhibits less sensitivity to changes in Δ . Second, this transformation is based on the logit distribution function (1); another interpretation of Γ is that subjects take longest to decide when the probability of choosing one or the other reward is close to $1/2$. Third, the precision parameter ω incorporates the effect of the variance of taste shocks on response time—greater variance in taste shocks corresponds to a smaller precision parameter ω , which corresponds to larger response times.

The response time (RT) model is

$$RT_{it} = b_0 + b_1 t + b_2 \Gamma(\Delta_t) + \varepsilon_{it} \quad (5)$$

where t is trial number, RT_{it} is the response time of subject i on trial t , and ε_{it} is the subject and trial specific noise term (section 4.1 motivates the $b_1 t$ term). We make no assumptions about the errors except that the random vectors $\varepsilon_i = (\varepsilon_{i1}, \dots, \varepsilon_{iT})$ are mutually independent. Our approach to estimating the model is simply to minimize the sum of squared residuals; in other words, we seek the vector $(\hat{b}_0, \hat{b}_1, \hat{b}_2)$ that minimizes

$$\sum_{it} [RT_{it} - (\hat{b}_0 + \hat{b}_1 t + \hat{b}_2 \Gamma(\Delta_t))]^2, \quad (6)$$

where

$$\hat{\Delta}_t = \frac{2}{1 + e^{\hat{\omega} \left| \frac{y_t}{1 + \hat{\sigma} \tau_t} - x_t \right|}}$$

Note that $\hat{\alpha}$ and $\hat{\omega}$ are taken from the revealed preference estimates obtained from (2). However, in some cases (see subsection 4.3) we estimate the complete vector $(\hat{b}_0, \hat{b}_1, \hat{b}_2, \hat{\alpha}, \hat{\omega})$ in the response time model (i.e., we use the RT data to estimate the preference parameters directly).

4. Results

4.1. Using choice data to estimate discount rates

Table 1 reports estimates $\hat{\alpha}$ obtained by using choice data to estimate the discount function. Specifically, we maximize the choice-based likelihood function in equation (2) for each of our three studies. We also estimate separately the discount rates for each of the three different reward size categories. Consistent with others (e.g., Kirby et al., 1999; Jaroni et al., 2004) we find that subjects' discount rates decrease with reward size. Moreover, we find that subjects in the Weight study have higher discount rates than subjects in the Cognition study, and that subjects in the Cognition study have slightly higher discount rates than subjects in the Web study. This is probably at least partially due to demographic differences between the samples.³

4.2 Using discount rates (estimated from choice data) to predict response time

We use the discount rate estimates from choice data to predict response times. Specifically, for each study we use the estimated parameters $\hat{\alpha}$ and $\hat{\omega}$ to define $\hat{\Delta}_t$, as in equation (3), apply Γ as in equation (4) to get our proxy for decision difficulty, $\Gamma(\hat{\Delta}_t)$, and then show that $\Gamma(\hat{\Delta}_t)$ predicts response times.

We also control for the effect of experience—as subjects practice the choice task, they answer questions more quickly. Subjects take a substantial amount of time on the first two trials and thereafter the effect of experience is linear. Therefore, we drop the first two trials from the RT analysis and control for the linear learning effect in the remaining 25 trials. Finally, we trim our dataset by removing observations with extremely high RTs—we demean RTs at the subject level and drop observations that are in the top 2% of the distribution of demeaned RTs.⁴ We apply these exclusions to all subsequent analyses.

The columns labeled *choice data* in Tables 2 and 3 report the estimated discounting parameters when the model is estimated with choice data that are taken from the trimmed sample.⁵ We form the decision difficulty measure, $\Gamma(\hat{\Delta}_t)$, using these parameters. To test our hypothesis that $\Gamma(\hat{\Delta}_t)$ predicts response time, we estimate the linear regression model

$$RT_{it} = b_0 + b_1 t + b_2 \Gamma(\hat{\Delta}_t) + \varepsilon_{it}.$$

³For example, the Weight study recruited subjects with high BMI and the Web study recruited college graduates. The differences in the group-level discount rates are reduced when we exclude overweight and obese subjects ($BMI \geq 25$) from the estimation samples. Chabris et al. (2008) estimate discount rates at the individual level and show statistically significant relationships between the estimated discount rates and behavioral variables such as BMI and smoking.

⁴Response time is a very noisy variable. An example of an outlier is a 59-second response in the Weight study (compared to a mean of 3.4 and SD of 2.8 seconds). In other analyses, varying the cutoff point from 2% yielded similar results to those reported here.

⁵These estimates are nearly identical to those in Table 1.

The effect of $\Gamma(\hat{\Delta}_t)$ is highly significant ($p < 0.001$) and positive in the regression. Figure 1 shows, for each of the three datasets, a plot of $RT_{it} - b_1 t$ (averaged by trial) against $\Gamma(\hat{\Delta}_t)$, and the corresponding regression lines. When we average by trial and control for the learning effect, $\Gamma(\hat{\Delta}_t)$ accounts for 51.2%, 49.8%, and 61.3% of the variance, respectively, in the Weight, Cognition, and Web studies. When we estimate a different α for each of the different reward categories, the results are even more statistically significant—the corresponding R^2 values are 74.8%, 64.1%, and 70.8%.

4.3. Using response times to predict preferences

We also invert the relationships reported in the previous subsection: we use response time data to estimate $\hat{\alpha}$. Specifically, we estimate equation (5) by minimizing the sum of squared residuals (6). Table 2 reports $\hat{\alpha}$ and $\hat{\omega}$ for each of the three datasets, with the *choice data* columns showing estimates obtained by estimating the choice data model (2) and the *RT data* columns showing estimates obtained by estimating the RT model (5). The correlation between the $\hat{\alpha}$'s estimated from choice data and the $\hat{\alpha}$'s estimated from response time data is 0.97.⁶ This does not imply that one should estimate discounting parameters from RT data. Rather it implies that there is an underlying structural relationship between revealed preferences and RT data.

Additionally, we estimate α and ω for each of the three reward categories. We estimate the choice model separately for each reward size category. The response time model (5) now includes a distinct discount rate, $\hat{\alpha}_k$, and precision parameter, $\hat{\omega}_k$, for each of the three size categories. Table 3 shows the estimates of $\hat{\alpha}$ and $\hat{\omega}$ for each of the three size categories in each of the studies. Again, the estimates based on choice data covary closely with the estimates based on response times. The correlation between the $\hat{\alpha}$'s estimated from choice data and the $\hat{\alpha}$'s estimated from response time data is 0.82.⁷

In an earlier version of this paper (Chabris et al., 2008b), we also compare the predictive power of the discounting parameters estimated from choice data and the discounting parameters estimated (exclusively) from RT data. The discounting parameters estimated using RT data predict out-of-sample choices almost as well as the parameters estimated using choice data.

5. Conclusion

Optimization theory predicts that agents will allocate more decision time to choices between options of similar expected utility than to choices between options of dissimilar expected utility (Gabaix and Laibson, 2005; Gabaix et al., 2006). Our experiments confirm this prediction. The value-gap variable explains 54% of the variance in average decision time across 27 binary choices. These results support the view that decision-making is a cognitively costly activity that allocates time according to cost-benefit principles.

Acknowledgments

We acknowledge financial support by a NARSAD Young Investigator Award and DCI Postdoctoral Fellowship awarded to Christopher F. Chabris, an NSF ROLE grant to J. Richard Hackman and Stephen M. Kosslyn, and NIA (P01 AG005842, R01 AG021650) and NSF (0527516) grants to David Laibson.

⁶The $\hat{\omega}$ estimates do not move together due to a discrepancy in the Weight study. The correlation between the $\hat{\omega}$'s estimated from choice data and from response time data is -0.66 .

⁷The correlation between the $\hat{\omega}$'s estimated from choice data and from RT data is 0.94.

References

- Chabris, Christopher F.; Laibson, David; Morris, Carrie L.; Schuldt, Jonathon P.; Taubinsky, Dmitry. Individual Laboratory-Measured Discount Rates Predict Field Behavior. *Journal of Risk and Uncertainty*. 2008aForthcoming
- Chabris, Christopher F.; Laibson, David; Morris, Carrie L.; Schuldt, Jonathon P.; Taubinsky, Dmitry. NBER Working Paper 2008b:14353.
- Gabaix, Xavier; Laibson, David. Bounded Rationality and Directed Cognition. Harvard University; 2005.
- Gabaix, Xavier; Laibson, David; Moloche, Guillermo; Weinberg, Stephen. Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model. *American Economic Review* 2006;96:1043–1068.
- Henmon, Vivian AC. The Time of Perception as a Measure of Differences in Sensations. *Archives of Philosophy, Psychology, and Scientific Methods* 1906;8
- Jaroni, Jodie L.; Wright, Suzanne M.; Lerman, Caryn; Epstein, Leonard H. Relationship between education and delay discounting in smokers. *Addictive Behaviors* 2004;29:1171–1175. [PubMed: 15236819]
- Johnson DM. Confidence and speed in the two-category judgment. *Archives of Psychology* 1930;241:1–52.
- Kirby, Kris N.; Petry, Nancy M.; Bickel, Warren K. Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls. *Journal of Experimental Psychology: General* 1999;128:78–87. [PubMed: 10100392]
- Mazur, James E. An adjusting procedure for studying delayed reinforcement. In: Commons, ML.; Mazur, JE.; Nevin, JA.; Rachlin, H., editors. *Quantitative Analysis of Behavior: The Effects of Delay and Intervening Events on Reinforcement Value*. Vol. 5. Hillsdale, NJ: Erlbaum; 1987. p. 55-73.
- Moyer, Robert S.; Landauer, Thomas K. Time Required for Judgements of Numerical Inequality. *Nature* 1967;215:1519–1520. [PubMed: 6052760]
- Pinel, Philippe; Piazza, Manuela; Bihan, Denis Le; Dehaene, Stanislas. Distributed and Overlapping Cerebral Representations of Number, Size, and Luminance During Comparative Judgments. *Neuron* 2004;41:983–93. [PubMed: 15046729]

Figure 1.

Plots of $RT - b_1t$ versus Γ for each of the three studies

Table 1

Discount rates estimated using choice data

	Weight	Cognition	Web
Alpha	0.0090 [0.0005]	0.0060 [0.0003]	0.0050 [0.0001]
Alpha small	0.0248 [0.0021]	0.0143 [0.0011]	0.0104 [0.0005]
Alpha medium	0.0129 [0.0010]	0.0063 [0.0005]	0.0056 [0.0002]
Alpha large	0.0065 [0.0006]	0.0045 [0.0003]	0.0037 [0.0001]

Notes: Standard errors in brackets.

Table 2

Estimated discounting parameters

	Weight		Cognition		Web	
	Choice	RT	Choice	RT	Choice	RT
Alpha	0.0094 [0.0005]	0.0112 [0.0013]	0.0060 [0.0003]	0.0043 [0.0003]	0.0050 [0.0001]	0.0051 [0.0003]
Omega	0.0837 [0.0028]	0.1545 [0.0339]	0.1111 [0.0039]	0.0845 [0.0158]	0.1199 [0.0024]	0.12457 [0.0121]

Notes: Standard errors in brackets. For the RT model, observations are clustered by subject and robust standard errors are reported.

Table 3

Estimated discounting parameters for each reward category

	Weight		Cognition		Web	
	Choice	RT	Choice	RT	Choice	RT
Alpha small	0.0207 [0.0018]	0.0116 [0.0017]	0.0147 [0.0012]	0.0089 [0.0012]	0.0106 [0.0005]	0.0065 [0.0006]
Alpha medium	0.0102 [0.0009]	0.0106 [0.0020]	0.0061 [0.0005]	0.0042 [0.0006]	0.0057 [0.0002]	0.0066 [0.0006]
Alpha large	0.0061 [0.0006]	0.0073 [0.0013]	0.0043 [0.0004]	0.0040 [0.0005]	0.0035 [0.0002]	0.0055 [0.0005]
Omega small	0.1509 [0.0079]	0.0945 [0.0368]	0.1917 [0.0106]	0.1345 [0.0262]	0.1898 [0.0059]	0.1144 [0.0176]
Omega medium	0.0804 [0.0046]	0.0435 [0.0156]	0.1125 [0.0066]	0.0773 [0.0152]	0.1195 [0.0040]	0.0547 [0.0082]
Omega large	0.0699 [0.0038]	0.0395 [0.0154]	0.0892 [0.0053]	0.0682 [0.0140]	0.0979 [0.0034]	0.0471 [0.0076]

Notes: Standard errors in brackets. For the RT model, observations are clustered by subject and robust standard errors are reported.