FICTIVE LEARNING IN CHOICE UNDER UNCERTAINTY: A LOGISTIC REGRESSION MODEL

By

Donald Brown, Oliver Bunn, and Caterina Calsamiglia

March 2013
Revised March 2014

COWLES FOUNDATION DISCUSSION PAPER NO. 1890R

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.econ.yale.edu/

Fictive Learning in Choice under Uncertainty: A Logistic Regression Model

Donald Brown, Oliver Bunn, and Caterina Calsamiglia

March 21, 2014

Abstract

This paper is an exposition of an experiment on revealed preferences, where we posite a novel discrete binary choice model. To estimate this model, we use general estimating equations or GEE. This is a methodology originating in biostatistics for estimating regression models with correlated data. In this paper, we focus on the motivation for our approach, the logic and intuition underlying our analysis and a summary of our findings. The missing technical details, including proofs, are in the working paper by Bunn, et al (2013).

The experimental data is available from the corresponding author: donald.brown@yale.edu. The recruiting poster and the informed consent form are included as appendices to the section on Experimental Protocols.

JEL Classification: C23, C35, C91, D03

Keywords: Counterfactual outcomes, Odds ratios, Alternating logistic regression

1 Introduction

We propose and analyze a revealed preference experiment on fictive learning in choice under uncertainty, where subjects are offered a sequence of binary choices between risky and ambiguous binary lotteries. Subjects know the relative frequencies of risky payoffs, but are ignorant of the relative frequencies of ambiguous payoffs. Inspired by Ellsberg's well known two-urn paradox (1961), where if the risky and ambiguous urn have the same payoffs then optimistic subjects choose the ambiguous urn and pessimistic subjects choose the risky urn, Bracha and Brown (2012) introduced affective utilities as representations of a subject's optimism or pessimism in making binary choices under uncertainty. We assume that subjects are endowed with random utility functions, where they evaluate risky lotteries with expected utility and ambiguous lotteries with affective utility. Subjects chose the risky lottery if its expected utility exceeds the affective utility of the ambiguous lottery by some random threshold.

Each subject's sequence of binary choices is divided into three phases: pre-learning, learning and post-learning. In the learning phase, the payoffs of actual and counterfactual choices are revealed to subjects. No payoffs are revealed to subjects in the pre-learning and the post-learning phases. The subjects in our experiment are

Yale undergraduates, randomly assigned to a control group or a treatment group. In the treatment group, subjects are exposed to the factual and counterfactual payoffs of lotteries in the learning phase, allowing subjects to estimate the relative frequencies of ambiguous payoffs. In the control group, subjects are exposed to noisy factual and counterfactual payoffs of ambiguous lotteries in the learning phase, where they cannot estimate the relative frequencies of ambiguous payoffs.

Conditioning current choices under uncertainty on counterfactual payoffs of previous choices, i.e., fictive learning, is well documented in $fMRI$ studies of gambling behavior in humans — see Lohrenz et al. (2007) and decision-making under uncertainty in monkeys — see Hayden et al. (2009). Recently Boorman et al. (2011) identified neural circuits for counterfactual payoffs and fictive learning. A common practice in experimental studies of decision-making under uncertainty, such as the $fMRI$ studies in Huettel et al. (2006) and Levy et al. (2009) is to posit a cross-sectional model for the experimental data. Unfortunately, a cross-sectional analysis ignores that each subject's repeated binary responses are correlated. In fact, this is the generic property of most revealed preference experiments in neuroeconomics.

Recently, Li et al. (2008) proposed the longitudinal analysis of neuroimaging data with general estimating equations (GEE), due to Liang and Zeger (1986) and Zeger and Liang (1986). Li et al. argue that the existing statistical methods for analyzing neuroimaging data are primarily developed for cross-sectional neuroimaging studies and not for panel neuroimaging data. We find this critique of the current practice in neuroimaging studies equally compelling as a critique of the current statistical practice in neuroeconomic studies of revealed preferences for risk and ambiguity. To this end, we propose a marginal longitudinal model of revealed preferences for risk and ambiguity, where for each subject the covariates in each trial are the payoffs of the ambiguous lottery and the payoffs and probabilities of payoffs of the risky lottery.

Two of the frequently used models for discrete repeated measurements of experimental outcomes are: mixed effects models, used extensively in econometrics to estimate individual specific parameters, and marginal models, where the regression coefficients are population parameters of subgroups. For a lucid discussion of the relative merits and limitations of the two approaches we recommend the lecture notes of Fitzmaurice published in power point on the internet under the title: **Applied Longitudinal Analysis: Contrasting Marginal and Mixed Effects Models.** General estimating equations or GEE is a widely used methodology in biostatistics for estimating the population-specific parameters in a marginal model. The GEE approach has a number of appealing properties for estimation of the regression coefficients in marginal models. First, we need only make assumptions on the first two moments of the distribution of the vector of responses. The GEE estimates of the regression coefficients are consistent and asymptotically normal, where the covariance matrix is consistently estimated using a sandwich estimator, even if the within subject associations among the repeated measurements have been misspecified. In many cases, GEE is almost as efficient as maximum likelihood estimation. We interpret the parameters in the marginal model as population averages in a given group.

There is an important difference between the application of longitudinal analysis to neuroimaging data, where the within-subject association of responses is considered a nuisance and our application of longitudinal analysis. In our experiment, subjects make a sequence of binary choices between risky and ambiguous binary lotteries. For paired binary data, Lipsitz et al. (1991) introduced odds ratios as a measure of the within-subject association of binary responses. We use alternating logistic regression (ALR), as proposed by Carey et al. (1993), with constant log odds ratios (LOR) as the within-subject association of responses in each phase of the sequence of binary choices to estimate the regression equations for both the first and second moments of the marginal model. In ALR the within-subject association of responses is not a nuisance for our model, but an essential part of the longitudinal analysis. It is the within-subject association of responses as log odds that allows us to test for fictive learning in the revealed preferences derived from the dependent, clustered responses of subjects.

We define fictive learning for each group as statistically significant changes in the responses of subjects before and after exposure to in the learning phase of the experiment. In each group, we estimate a constant (LOR) of the odds of choosing the risky lottery in a trial in the post-learning phase, conditional on the choice in a trial in the pre-learning phase. Our null hypotheses is the absence of fictive learning in the learning phase for each group. For the treatment group, we reject the null hypothesis that the LOR, is zero, i.e., there is fictive learning in the learning phase. This finding is significant at the 1% level. The significant fictive learning in the treatment group, given the sample data in the learning phase is as expected. The surprising finding in our experiment is that we also reject the null hypothesis that the LOR,between trials in the pre and post learning phases in the control group, is also zero, i.e., there is fictive learning in the learning phase for the control group. We expected no fictive learning for the control group. Again the finding is significant at the 1% level. A possible but problematic explanation of the choice behavior of the control group is apophenia: "seeing meaningful patterns in meaningless or random data." For an evolutionary rational of this type of behavior, see Shermer's article **"Patternicity: Finding Meaningful Patterns in Meaningless Noise"** in *Scientific American* (2008).

Whatever subjects in the two groups "learn" in the treatment phase, we can ask if the effects of the treatments are significantly different between groups. To compare the effects of the treatment phase in each group, we use the log-odds-ratio test proposed in chapter 10 of Fleiss et al. (2004). The analysis begins with the calculation of whatever the subjects in the two groups "learn" from the treatment phase, then we ask if the effects of the two treatments are significantly different. The null hypothesis is that the LOR, between trials in the pre and post learning phases, in the treatment group and the LOR, between trials in the pre and post learning phases, in the control group are equal. We reject the null at the 1% level. That is, the fictive learning in the two groups produced significantly different choices in the post learning phase relative to the choices in the pre-learning phase. See the working paper for details. A more dramatic illustration of the different effects of fictive learning in the two groups are the box plots in figure 1. That is, if we plot the amount of time where the ambiguous lottery is chosen in each phase of the experiment then the curve is a piece-wise linear concave function for the treatment group and a piece-wise linear convex function for the control group.

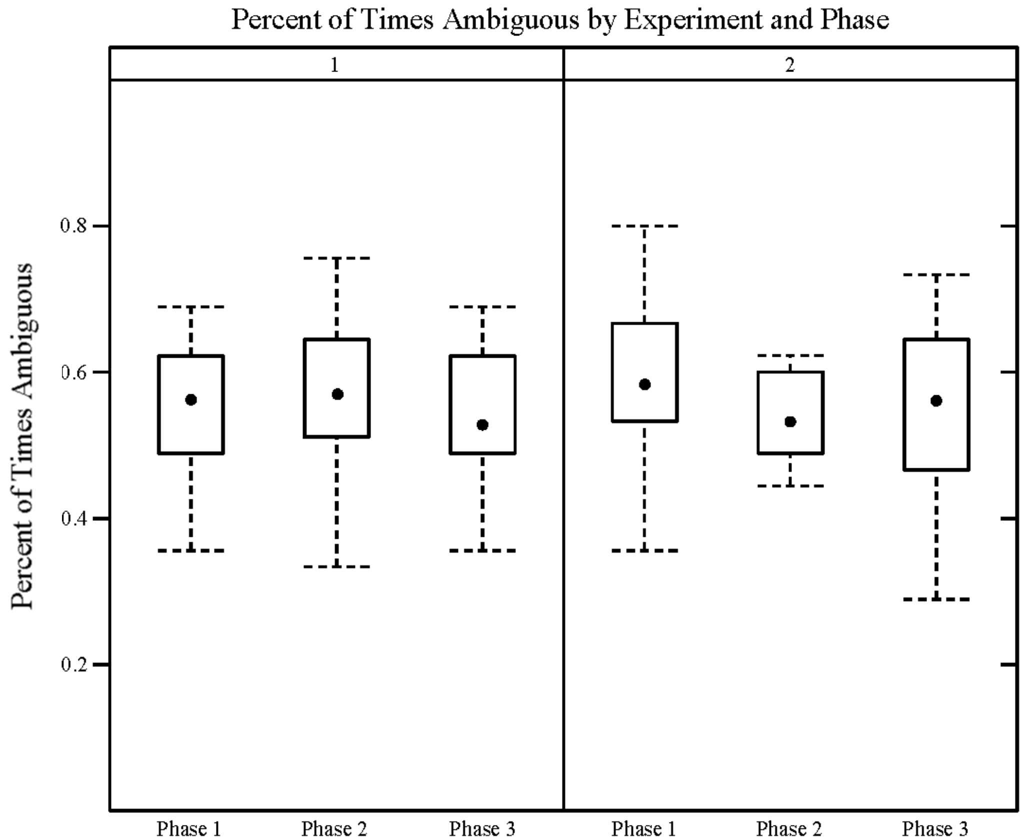

Figure 1

2 Experimental Protocols

To test for the presence of fictive learning in revealed preferences for risk and ambiguity, we propose an experiment on revealed preferences for choices under uncertainty, consisting of 36 Yale undergraduates as subjects randomly chosen from the 2011 Yale Fall term. Each subject makes a sequence of 100 binary choices between risky and ambiguous lotteries. Risky lotteries are defined as lotteries where the relative frequencies of outcomes are known. Ambiguous lotteries are lotteries where the relative frequencies of outcomes are unknown.

The experiment is divided into three phases. Subjects face the same sequence of 30 binary choices between risky and ambiguous lotteries in the first and third phase of the experiment. That is, the trials in phase 1 and phase 3 are clustered matched-pairs, but the lotteries in phase 1 and 3 for the two groups are independent. To test for fictive learning, we reveal to each subject the outcomes of her 40 actual and counterfactual choices in phase 2. In the treatment group, the relative frequencies of counterfactual ambiguous outcomes in phase 2 are relatively easy to learn, using sample averages of the outcomes of the ambiguous lotteries. In the control group, the relative frequencies of counterfactual ambiguous outcomes in phase 2 are quite difficult, if not impossible, to learn, since the sample averages do not converge. See the working paper for a formal proof. The binary choices in phase 2 are the same in both groups and independent of the binary choices in phases 1 and 3. Subjects are unaware that they will be exposed to counterfactual outcomes in phase 2 before they are presented binary choices in phase 3. In particular, subjects do not know if the relative frequencies of outcomes of ambiguous lotteries in phase 2 is a sample average of the probabilities of ambiguous outcomes in phase 1 and 3. In fact, they are in the treatment group, but not in the control group.

No outcomes are revealed to subjects in the first and third phase of the experiments. The lotteries are displayed as pie graphs on each subject's computer screen. Probabilities for the risky lotteries are displayed. The probabilities determining the payoffs of ambiguous lotteries are constant in phase 1 and phase 3 of both experiments, but never revealed to the subjects. We randomly vary the placement and colors of the lotteries on the computer screen to control for positional bias. We randomly choose one group of 17 students from the 36 students as the treatment group. At the end of the experiment, a trial is randomly chosen for each subject and the subject is given the payoff of her choice.

We define fictive learning for each group as statistically significant changes in the responses of subjects before and after exposure to outcomes in phase 2 of the experiment. In each group, we estimate a constant log odds ratio (LOR) of the odds of choosing the risky lottery in a trial in phase r, conditional on the choice in a trial in phase s. We use $GENMOD$ in SAS with the $LOGOR$ option to estimate the regression equations for both the first and second moments of the marginal model. We assume the LOR is constant in phase 1; phase 2; phase 3; between phases 1 and 2; between phases 1 and 3 and between phases 2 and 5. In ALR, the odds ratio for each pair of trials is

$$\frac{Pr(Yij=1;Yik=1)Pr(Yij=0;Yik=0)}{Pr(Yij=1;Yik=0)Pr(Yij=0;Yik=1)}=\frac{\frac{Pr(Yij=1|Yik=1)}{Pr(Yij=0|Yik=1)}}{\frac{Pr(Yij=1|Yik=0)}{Pr(Yij=0|Yik=0)}}$$

where i is the subject index and j and k are the indices of the trials.$Yij=1$ means subject i choose the risky lottery in trial j.

The recruiting poster and informed consent form are attached as appendices.

3 A Marginal Analysis of Fictive Learning

In Ellsberg's well-known two-color paradox (1961), where the risky and the ambiguous urn have the same payoffs, optimistic subjects choose the ambiguous urn, where the relative proportions of the black and white balls are unknown, and pessimistic subjects choose the risky urn, where the relative proportions of the black and white balls are known. Recently, Bracha and Brown (2012) introduced affective utilities where optimistic subjects are endowed with a convex affective utility function and pessimistic subjects are endowed with a concave affective utility function. In their model β is a proxy for risk-aversion and α is a proxy for ambiguity-aversion. The concavity (convexity) of the utility functions in this class of non-expected utility functions depends on the ratio of α and β. In our model we restrict attention to the parametric class of linear-quadratic concave (convex) utility functions introduced by Rockafellar (1987).

This is the technical section of the working paper and we limit our exposition to a discussion of the results and refer the reader to the working paper for technical details, such as proofs. We denote risky lotteries as X and ambiguous lotteries as Y, where

$$X \equiv (x_1, x_2; \pi_1, \pi_2) \text{ and } Y \equiv (y_1, y_2).$$

Subjects evaluate risky lotteries, X, using expected utility:

$$E_{\beta,K}(U(X)) \equiv \pi_1 u_{\beta,K_R}(x_1) + \pi_2 u_{\beta,K_R}(x_2).$$

where the Bernoulli utility function.

$$u_{\beta,K_R}(w) \equiv K_R w + \frac{\beta}{2} w^2 \text{ and } \beta \text{ is the proxy for risk-aversion.}$$

If

$$u_{\beta,K_R}(x_1) = K_R x_1 + \frac{\beta}{2} x_1^2 \text{ and } u_{\beta,K}(x_2) = K_R x_2 + \frac{\beta}{2} x_2^2$$

then

$$E_{\beta,K_R}(U(X)) = K_R(\pi_1 x_1 + \pi_2 x_2) + \frac{\beta}{2}[\pi_1 x_1^2 + \pi_2 x_2^2].$$

Subjects evaluate ambiguous lotteries, Y, using affective utility:

$$J_{\alpha,\beta,K_A}(Y) \equiv K_A(y_1 + y_2) + \frac{[\alpha - \beta]}{2}[y_1^2 + y_2^2].$$

In the parametric specification of $J_{\alpha,\beta,K_A}(Y)$, the affective utility of the ambiguous lottery Y, β is the proxy for risk-aversion and α is the proxy for ambiguity-aversion

In each binary choice between a risky and an ambiguous lottery, we assume that subjects choose the lottery that maximizes random utility, where the parametric nonrandom components of the random utility of a risky and an ambiguous lottery are $E_{\beta,K_R}(U(X))$ and $J_{\alpha,\beta,K_A}(Y)$. These are linear-quadratic multivariate functions. The important technical aspect of the linear-quadratic formulation is that for any pair of risky and ambiguous lotteries, the difference in the expected utility of the risky lottery and the affective utility of the ambiguous lottery is linear in the parameters. The binary discrete choice model is a generalized linear model where the link function is a cdf. In this paper, the link function is the logistic cdf. The argument of the link function is the difference of the parametric nonrandom components of the random utility of a risky and an ambiguous lottery. If the nonrandom component of the random utility function is linear in the parameters, then the log-likelihood is strictly concave in the parameters defining the nonrandom components of the random utility function.

$\Phi(\alpha, \beta, K_R, K_A)$, the argument of the logistic cdf, is the difference of the expected utility of the risky lottery $X = (x_1, x_2; \pi_1, \pi_2)$ and the affective utility of the ambiguous lottery $Y = (y_1, y_2)$. Hence the choice probability for X, $p_X(\alpha, \beta, K_R, K_A)$, is implicitly defined by the logistic cdf

$$\Lambda[\Phi] \equiv \frac{\exp \Phi}{1 + \exp \Phi}$$

where

$$\Phi(\alpha, \beta, K_R, K_A) \equiv \log \frac{p_X(\alpha, \beta, K_R, K_A)}{1 - p_X(\alpha, \beta, K_R, K_A)} = [E_{\beta, K_R}(U(X)) - J_{\alpha.\beta, K_A}(U(Y))]$$

is the log-odds of choosing X

$$\begin{aligned} & [E_{\beta, K_R}(U(X)) - J_{\alpha, \beta, K_A}(U(Y))] \\ = \; & K_R(\pi_1 x_1 + \pi_2 x_2) + \frac{\beta}{2}[\pi_1 x_1^2 + \pi_2 x_2^2] - \left\{ K_A(y_1 + y_2) + \frac{[\alpha - \beta]}{2}[y_1^2 + y_2^2] \right\} \end{aligned}$$

$$p_X(\alpha, \beta, K_R, K_A) = \frac{\exp[E_{\beta, K_R} U(X) - J_{\alpha, \beta, K_A}(U(Y))]}{(1 + \exp[E_{\beta, K_R} U(X) - J_{\alpha, \beta,}(U(Y))])}$$

$p_X(\alpha, \beta, K_R, K_A)$ is the explicit probability of choosing the risky lottery X in the pair-wise comparison between the risky lottery X and the ambiguous lottery Y. In each experiment, let $\theta_{i,j} = 1$ if the risky lottery is chosen by subject i in trial j and 0 otherwise, then the density of $\theta_{i,j}$ is

$$[p_X(\alpha, \beta, K_R, K_A)]^{\theta_{i,j}} [1 - p_X(\alpha, \beta, K_R, K_A)]^{1 - \theta_{i,j}}$$

We estimate the regression parameters for each phase of the experiment:

$$\pi^k \equiv (\alpha_{1k}, \alpha_{2k}, \alpha_{3k}; \beta_{1k}, \beta_{2k}, \beta_{3k}; K_{A1k}, K_{A2k}, K_{A3k}; K_{R1k}, K_{R2k}, K_{R3k}) \in R^{12}$$

in the marginal model, using generalized estimating equations (GEE). Our primary reference is the monograph on applied longitudinal analysis by Fitzmaurice et al. (2011). See the working paper for details.

Table 1. Analysis of GEE Parameter Estimates: Treatment Group						
	Estimates	Standard error	95% confidence limits		Z	$\Pr > \|Z\|$
$kr1$	2.1074	0.3118	1.4863	2.7186	6.76	< 0.0001
$kr2$	1.4113	0.2397	0.9416	1.8810	5.89	< 0.0001
$kr3$	1.7255	0.2616	1.2127	2.2383	6.59	< 0.0001
$\beta 1$	-0.0955	0.0172	-0.1292	-0.0617	-5.54	< 0.0001
$\beta 2$	-0.0507	0.0125	-0.0752	-0.0261	-4.04	< 0.0001
$\beta 3$	-0.0697	0.0150	-0.0992	-0.0402	-4.64	< 0.0001
$ka1$	0.8134	0.1418	0.5355	1.0914	5.74	< 0.0001
$ka2$	0.6000	0.1217	0.3615	0.8385	4.93	< 0.0001
$ka3$	0.6250	0.1303	0.3695	0.8804	4.80	< 0.0001
$\alpha 1$	-0.0206	0.0072	-0.0347	-0.0065	-2.87	< 0.0041
$\alpha 2$	-0.0103	0.0068	-0.0236	0.0030	-1.51	< 0.1300
$\alpha 3$	-0.0090	0.0071	-0.0229	0.0048	-1.28	< 0.2016

Table 2. Analysis of GEE Parameter Estimates: Control Group						
	Estimates	Standard error	95% confidence limits		Z	$\Pr > \|Z\|$
$kr1$	1.6809	0.3058	1.0815	2.2803	5.50	< 0.0001
$kr2$	1.4511	0.2580	0.9454	1.9568	5.62	< 0.0001
$kr3$	1.4122	0.3367	0.7513	2.0711	4.19	< 0.0001
$\beta 1$	-0.0661	0.0156	-0.0966	-0.0355	-4.24	< 0.0001
$\beta 2$	-0.0525	0.0147	-0.0813	-0.0238	-3.58	< 0.0003
$\beta 3$	-0.0526	0.0172	-0.0863	-0.0189	-3.06	< 0.0022
$ka1$	0.7642	0.1503	0.4697	1.0587	5.09	< 0.0001
$ka2$	0.6261	0.1408	0.3501	0.9022	4.45	< 0.0001
$ka3$	0.0587	0.1821	0.2117	0.9256	3.12	< 0.0018
$\alpha 1$	-0.0195	0.0080	-0.0351	-0.0039	-2.44	< 0.0145
$\alpha 2$	-0.0134	0.0083	-0.0296	0.0028	-1.62	< 0.1055
$\alpha 3$	-0.0094	0.0099	-0.0287	0.0100	-0.95	< 0.3442

In the working paper, we show that the estimated parameter values are consistent with the concavity and monotonicity of the Bernoulli utilities of wealth and consistent with the convexity and monotonicity of the affective utility for each group. That is, both the treatment group and the control group are (on average) risk-averse, but both groups are (on average) optimistic.

The odds ratio in ALR is

$$\frac{Pr(Yij=1;Yik=1)Pr(Yij=0;Yik=0)}{Pr(Yij=1;Yik=0)Pr(Yij=0;Yik=1)} = \frac{\frac{Pr(Yij=1|Yik=1)}{Pr(Yij=0|Yik=1)}}{\frac{Pr(Yij=1|Yik=0)}{Pr(Yij=0|Yik=0)}}$$

$$\left[\frac{Pr(Yij=1;Yik=1)Pr(Yij=0;Yik=0)}{Pr(Yij=1;Yik=0)Pr(Yij=0;Yik=1)}\right]^{-1} = \frac{\frac{Pr(Yij=1|Yik=0)}{Pr(Yij=0|Yik=0)}}{\frac{Pr(Yij=1|Yik=1)}{Pr(Yij=0|Yik=1)}}$$

where i is the subject index and j and k are the indices of the trials.$Yij = 1$ means the subject choose the risky lottery in trial j. We assume the LOR is constant in phase 1; phase 2; phase 3; between phases 1 and 2; between phases 1 and 3 and between phases 2 and 5. In the following tables for the treatment and control groups, the estimated constant LOR are denoted Alpha Q for $Q = 1, 2, ..., 5$. We test the null hypothesis H_0: the log odds ratio is equal to zero, against the alternative hypothesis H_A: the log odds ratio is unequal to zero. Here are the estimates for the log-odds ratios.

Table 3. Analysis of GEE Parameter Estimates: LOR for the Treatment Group

	Estimates	Standard error	95% confidence limits		Z	Pr > $\|Z\|$
$\alpha1$ (Phase 1)	−0.0096	0.0468	−0.1013	0.0822	−0.20	0.8381
$\alpha2$ (Phase 2)	0.1124	0.0641	−0.0132	0.2380	1.75	0.0796
$\alpha3$ (Phase 3)	0.0334	0.0379	−0.1077	0.0408	−0.88	0.3778
$\alpha4$ (Phase 1 & 2)	0.0208	0.0376	−0.0529	0.0944	0.55	0.5803
$\alpha5$ (Phase 1 & 3)	0.1057	0.0389	0.0293	0.1820	2.71	0.0067
$\alpha6$ (Phase 2 & 3)	0.0112	0.0454	−0.0778	0.1003	0.25	0.8046

Table 4. Analysis of GEE Parameter Estimates: LOR for the Control Group

	Estimates	Standard error	95% confidence limits		Z	Pr > $\|Z\|$
$\alpha1$ (Phase 1)	0.1217	0.0744	−0.0242	0.2675	1.64	0.1020
$\alpha2$ (Phase 2)	−0.0463	0.0222	−0.0898	−0.0027	−2.08	0.0374
$\alpha3$ (Phase 3)	0.1349	0.0832	−0.0282	0.2980	1.62	0.1050
$\alpha4$ (Phase 1 & 2)	−0.0061	0.0304	−0.0656	0.0535	−0.20	0.8420
$\alpha5$ (Phase 1 & 3)	0.2051	0.0779	0.0525	0.3577	2.63	0.0084
$\alpha6$ (Phase 2 & 3)	−0.0239	0.0303	−0.0833	0.0356	−0.79	0.4315

Alpha 2 and Alpha 5 are the only significant statistics in each group. The LOR in phase 2 of the control group

$$\frac{Pr(Yij = 1 \mid Yik = 0)}{Pr(Yij = 0 \mid Yik = 0)} > \frac{Pr(Yij = 1 \mid Yik = 1)}{Pr(Yij = 0 \mid Yik = 1)}$$

and the LOR in phase 2 of the treatment group is

$$\frac{Pr(Yij = 1 \mid Yik = 1)}{Pr(Yij = 0 \mid Yik = 1)} > \frac{Pr(Yij = 1 \mid Yik = 0)}{Pr(Yij = 0 \mid Yik = 0)}.$$

For Alpha 5, the LOR between phase 1 and phase 3 in each experiment is

$$\frac{Pr(Yij = 1 \mid Yik = 1)}{Pr(Yij = 0 \mid Yik = 1)} > \frac{Pr(Yij = 1 \mid Yik = 0)}{Pr(Yij = 0 \mid Yik = 0)}.$$

4 Acknowledgements

The WhiteBox Foundation and SOM provided generous financial support for our experiments for several summers. Le-Minh Ho was a truly exceptional RA. Finally we wish to thank our colleagues and graduate students for their ongoing interest in this project. This paper is a revision of CFDP 1890.

References

[1] Boorman, Erie D., Timothy E. Behrens, and Matthew F. Rushworth (2011). "Counterfactual Choice and Learning in a Neural Network Centered on Human Lateral Frontopolar Cortex."

[2] Bracha, A., and D. J. Brown (2013). "Keynesian Utilities: Bulls and Bears," CFDP 1898.

[3] Bunn, O., C. Calsamiglia, and D.J. Brown (2013). "Testing for Fictive Learning in Decision-Making under Uncertainty," CFDP 1890.

[4] Carey, V., S.L. Zeger and P. Diggle (1993). "Modelling Multivariate Binary Data with Alternating Logistic Regressions," *Biometrika*, 517–526.

[5] Ellsberg, D., (1961). "Risk, Ambiguity and the Savage Axioms," *Quarterly Journal of Economics*, 643–669.

[6] Fitzmaurice, Garrett M., Nan M. Laird, James H. Ware (2011). *Applied Longitudinal Analysis* (Wiley Series in Probability and Statistics).

[7] Fleiss, Joseph L., Bruce Levin, and Myunghee Cho Paik (2003). *Statistical Methods for Rates and Proportions*, Third Edition, John Wiley & Sons, Inc.

[8] Hayden, Benjamin Y., John M. Pearson, and Michael L. Platt (2009). "Fictive Reward Signals in Anterior Cingulate Cortex," *Science,* 948–950.

[9] Huettel, Scott A., C. Jill Stowe, Evan M. Gordon, Brent T. Warner, and Michael L. Plat (2006). "Neural Signatures of Economic Preferences for Risk and Ambiguity," *Neuron* 49, 765–775.

[10] Keynes, J.M. (1937). "The General Theory of Employment," *Quarterly Journal of Economics*, 209–223.

[11] Levy I„ J. Snell, A.J. Nelson, A. Rustichini, and P.W. Glimcher (2010). "Neural Representation of Subjective Value under Risk and Ambiguity," *Journal of Neurophysiology*, 1036–1047.

[12] Li, Yimei, Hongtu Zhua Lia, Hongtu, Yasheng Chenb, Hongyu, John Gilmorec, Weili Lin and Dinggang Shenb (2009) "STGEE: Longitudinal Analysis of Neuroimaging Data"Proceeding," Proceeding of the SPIE.

[13] Liang, K.-Y. and S.L. Zeger (1986). "Longitudinal Data Analysis Using Generalized Linear Models," *Biometrika*, 13–22.

[14] Lipsitz, Stuart R., Nan M. Lair and David P. Harrington (1991). "Generalized Estimating Equations for Correlated Binary Data: Using the Odds Ratio as a Measure of Association," *Biometrika*, 153–160.

[15] Lohrenz, Terry, Kevin McCabe, Colin F. Camerer, and P. Read Montague (2006). "Neural Signature of Fictive Learning Signals in a Sequential Investment Task," PNAS, 9493-9498.

[16] Platt, M., and Hayden,B. (2011). "Learning: Not Just the Facts, Ma'am, but the Counterfactuals as Well," PLoS Biol 9(6): e1001092..

[17] Rockafellar, R.T. (1987). "Linear-Quadratic Programming and Optimal Control," *SIAM J. Control and Opt.*, 781–814.

If you want to try your luck, don't go to Foxwoods. Join us in a 75-minute research experiment on decision-making under uncertainty, i.e. gambling.

Our games are rigged in your favor: You will receive **$8** for participating. In addition, you have a chance to win between **$2** and **$15** during the experiment.

You can participate in the experiment at one of the following five times:

1. Tuesday, October 25, at 3pm,
2. Tuesday, October 25, at 4.30pm,
3. Tuesday, October 25, at 7pm,
4. Wednesday, October 26, at 2pm,
5. Wednesday, October 26, at 7pm.

In order to be **eligible** for participation, you must:

- be at least 19 years old,
- be enrolled at Yale in the fall term 2011,
- NOT be taking a course from Professor Donald Brown or Professor Laurie Santos in the fall term 2011.

If you meet the above requirements and would like to participate, please contact Donald Brown at donald.brown@yale.edu or Oliver Bunn at oliver.bunn@yale.edu and specify your preference with regard to the time of the experiment (choices 1-5 above).

The principal investigator for this study is Professor Donald Brown, HSC #1104008396.

Decision Making Under Uncertainty

Experimental Study Performed at Yale University

Principal Investigator: Professor Donald J. Brown

IRB Protocol Number: 1104008396

Research Informed Consent Form

This is a research experiment on decision-making under uncertainty. Although this study will not benefit you personally, we hope that our results will add to the knowledge about sequential decision-making under risk and ambiguity.

This is a 75-minute experiment where you will be asked to make 100 pair-wise choices between lotteries with two outcomes. In some lotteries you will be told the probabilities of the outcomes and for other lotteries you will not be told the probabilities of the outcomes. The experiment is divided into 3 blocks. At the end of each block, you can take a break at your own discretion.

You may leave the experiment at anytime and we will pay you $8. At the end of the experiment you will you will earn additional dollars from your choices. That is, if you complete the experiment, then we will randomly select one of your 100 choices and you will also receive the payoff of that choice. The payoffs range from $2 to $15.

All of your responses will be confidential. Your responses will be numbered and any information linking your number with your name will be destroyed at the end of the experiment, after you are paid. This anonymous data will be available upon request to other social scientists.

To participate in this experiment you must be at least 19 years old. The risks of this study include fatigue or mild stress. These risks are no greater than those found in taking a 75 minute midterm exam.

If you have any questions about this study, you may contact the investigator, Professor Donald Brown at donald.brown@yale.edu. If you would like to talk to someone other than the researchers to discuss problems or concerns, to discuss situations in the event that a member of the research team is not available, or to discuss your rights as a research participant, you may contact the Yale University Human Subjects Committee, Box 208010, New Haven, CT 06520-8010, 203-785-4688, human.subjects@yale.edu.

If you wish to participate in this study, please sign the consent form.

______________________________ ______________________

Signature of Study Participant **Date of the Study**

PERSONAL COPY OF STUDY PARTICIPANT