

On the relative importance of the hot stove effect and the tendency to rely on small samples

Takemi Fujikawa*

Centre for Policy Research and International Studies

Universiti Sains Malaysia

Abstract

Experiments have suggested that decisions from *experience* differ from decisions from *description*. In experience-based decisions, the decision makers often fail to maximise their payoffs. Previous authors have ascribed the effect of underweighting of rare outcomes to this deviation from maximisation. In this paper, I re-examine and provide further analysis on the effect with an experiment that involves a series of simple binary choice gambles. In the current experiment, decisions that bear small consequences are repeated hundreds of times, feedback on the consequence of each decision is provided immediately, and decision outcomes are accumulated. The participants have to learn about the outcome distributions through sampling, as they are not explicitly provided with prior information on the payoff structure. The current results suggest that the “hot stove effect” is stronger than suggested by previous research and is as important as the payoff variability effect and the effect of underweighting of rare outcomes in analysing decisions from experience in which the features of gambles must be learned through a sampling process.

Keywords: decisions from experience, payoff variability, rare events, uncertainty, undersampling.

1 Introduction

Much attention has been given to the distinction between decisions from *description* and decisions from *experience*. In experience-based decisions, people experience difficulty in estimating and understanding uncertainty. Erev and Barron (2005) hypothesised that two main behavioural tendencies determine the effect of rare events on repeated decisions from experience. The first is a tendency to rely on small samples of past experiences (also proposed by Fox & Hadar, 2006). This tendency leads to underweighting of rare events, as most samples are not likely to include the rare events. The second is a tendency to rely on recent experiences. When the information available to the decision makers (DMs) is limited to the obtained payoffs, this tendency leads to the “hot stove effect”, which implies overweighting of the worst outcomes. The hot stove effect was first introduced by Mark Twain with his observation that if a cat jumped on a hot stove, then she would never jump on a hot stove again. However, the cat would never jump even on a cold stove. Coutu (2006) states that the hot stove effect is a funda-

mental problem of learning that reduces the DMs’ likelihood of repeating decisions that got them in trouble. The hot stove effect implies a bias against a risky alternative in binary experience-based decisions (Denrell & March, 2001). The bias is a product of the tendency to reproduce actions that have been successful and avoid recent actions that have led to poor outcomes.

Previous research on experience-based decisions has led to mixed conclusions with regard to the descriptive value of the hot stove effect. Whereas some studies (e.g., Denrell & March, 2001) demonstrate its importance, other studies (e.g., Barron & Erev, 2003; Erev & Barron, 2005) suggest that this effect is weak. In this paper, I try to clarify this picture by focusing on choice problems in Barron and Erev (2003) and Erev and Barron (2005). The authors conducted experiments in which three choice problems (Problem 1, 2 and 3) were performed by the participants, each involving 400-fold binary choice between H (an alternative with higher expected value) and L (an alternative with lower expected value). Table 1 shows the payoff structure of each problem. For example, one selection of H in Problem 1 made the participants earn four points with probability 0.8 and zero point otherwise. The participants in their study were told that the experiments included many trials, and their goal in each trial t ($t = 1, \dots, 400$) was to select (click on) one of the two unmarked buttons that appeared on the

*I thank Hidenori Oda for his helpful comments and valuable research support. I gratefully acknowledge the valuable suggestions and comments of Jon Baron, Greg Barron, Ido Erev, Nick Feltovich and two anonymous reviewers. All errors remain my own. Address: Takemi Fujikawa, Centre for Policy Research and International Studies, Universiti Sains Malaysia, 11800 Penang, Malaysia. Email: takemi@usm.my.

Table 1: Choice problems in Barron and Erev (2003) and Erev and Barron (2005). P_H (P_L) is % of H (L) choices over 400 rounds.

Problem	H	P_H	L	P_L
1 ($N=48$)	4, 0.8	63%	3, 1	37%
2 ($N=48$)	4, 0.2	51%	3, 0.25	49%
3 ($N=48$)	32, 0.1	28%	3, 1	72%

computer screen. Each click resulted with an immediate payoff (random draw from the payoff distribution associated with the selected button). Thus, the prior information was minimalistic, and the participants had to base their decisions on experience. The participants deviated from maximisation. Table 1 shows the maximisation rate (the overall proportion of H choices) in each problem: for example, the overall proportion of H choices was 0.63 in Problem 1.

In the data considered by Erev and Barron (2005), the tendency to rely on small samples appeared to be stronger than the hot stove effect. The clearest support for this conclusion came from Problem 1, which used the clicking paradigm, where: (1) the participants were asked to select between unlabelled buttons on the computer screen; (2) each selection/click led to a random draw from the payoff distribution associated with the different buttons; and (3) in choosing among possible options, the participants had to rely on the immediate feedback obtained in similar situations in the past.

Notice that in Problem 1 the worst outcome (0 from H) is also the rare outcome (probability of 0.2). In Problem 1, reliance on small samples and the hot stove effect lead to contradicting predictions. Reliance on small samples implies that the rare outcome (0 from H) will be *underweighted*: this prediction implies that H will be preferred. The hot stove effect predicts the participants' learning that reduces their likelihood of repeating decisions, with which they have done poorly (i.e., getting burned on a hot stove in Twain's example, and thus referring to earning the worst outcome from H). Thus, the hot stove effect implies that the worst outcome (0 from H) will be *overweighted*: this prediction implies that L will be preferred. Barron and Erev (2003) and Erev and Barron (2005) reported that the observed proportion of H choices (over 400 trials) was 0.63. Their results suggest that the tendency to rely on small samples is stronger than the hot stove effect.

Follow-up studies demonstrated the descriptive value

of the assumed tendency to rely on small samples, and of the hot stove effect. For example, all the leading models in a recent choice prediction competition that focused on repeated decisions from experience can be described as alternative quantifications of these assumptions (see Erev et al., 2009). However, some of the recent results appear to question Erev and Barron's (2003) conclusions with regard to the relative magnitude of the two effects. Review of Erev and Barron (2003) suggests that the clearest indications for underweighting of rare events come from studies that examine decisions from experience with complete feedback (e.g., Ert & Erev, 2007). This design controls the hot stove effect with the provision of complete feedback.

A different picture is, however, shown in studies that focus on decisions from experience with limited feedback (e.g., Fujikawa, 2007; Fujikawa & Oda, 2007): the feedback is limited to the obtained payoff, and the foregone payoff (the payoff from the unselected option) is not presented. These studies reveal strong underweighting of attractive rare events (when reliance on small samples and the hot stove effect lead to the same predictions) but no clear indication of underweighting of unattractive rare events (when the two tendencies lead to contradicting predictions). This verbal summary of the results is consistent with the predictions of the leading models in the choice prediction competition. For example, the best baseline model (explorative sampler with recency in Erev et al. (2009)) predicts a H-rate of only 0.54 in Problem 1.

The main goal of the current paper is to clarify this picture: a picture that the hot stove effect is stronger than suggested by Barron and Erev (2003) and Erev and Barron (2005). In order to achieve this goal I implemented Problem 1, 2 and 3. Note again that the hot stove effect implies a bias toward L (the low variability option) in Problem 1 and 3.

2 Experiment

The current experiment was conducted at the Kyoto Experimental Economics Laboratory (KEEL) in Japan with 42 paid subjects — undergraduates from various faculties at Kyoto Sangyo University. On their arrival at the KEEL, each participant was assigned a workstation that displayed an experimental screen, and distributed a written instruction of the experiment. (The instruction and experimental screen are available in Appendix.) The instruction was read aloud and the participants were given an opportunity to ask questions individually. The participants engaged in Problem 1, 2, and 3 in order. They were instructed to operate a "computerised money machine" and to choose one of two unmarked buttons shown in Figure 1 which corresponded to H and L for 400 times

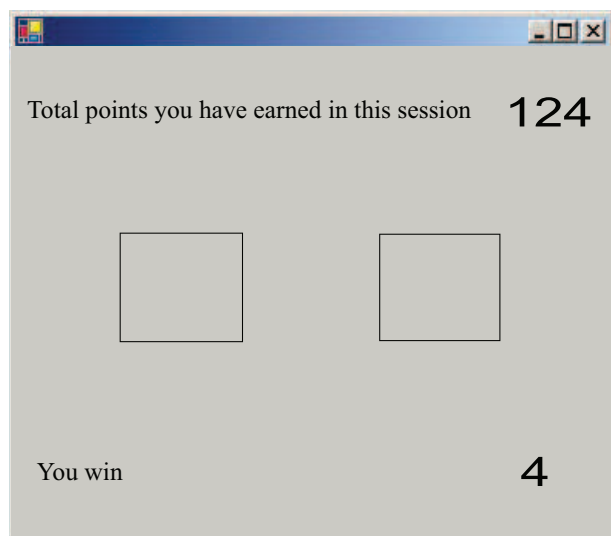


Figure 1: Computerised money machine

in each of the three problems. They made a choice between the two unmarked buttons on a computer screen to which each participant was assigned. In each trial t ($t = 1, 2, \dots, 400$), the participants were asked to click on one of the two buttons. Each click led to a random draw from the outcome distribution associated with the selected button. The participants were disclosed *neither* prior information on possible outcomes and probabilities, *nor* the exact length of the experiment.¹ They could see the drawn value (the obtained payoffs) after each trial on their computer screens. That is, the information available to the participants was limited to feedback concerning the outcomes of their previous decisions. The money machine provided the participants with binary types of feedback immediately following each choice: (1) the payoff for the choice that appeared on the screen for the duration of one second; and (2) an update of an accumulating payoff counter, which was constantly displayed.

The protocol of the experiment was as follows. At first, the participants played Problem 1, 2 and 3; that is, they were played 1200 trials in the experiment (400 trials for each problem). As noted above, they were not informed that they were to play exactly three choice problems, in each of which the participants were presented with a 400-fold repetition of a binary choice. Hence, the participants were not aware that they had 1200 trials to play in the experiment. Instead, they were aware that they faced several choice problems in the experiment. The participants started with Problem 1 and made 400 selections in Problem 1. Then, the participants were prompted to move to Problem 2 by the automatically-generated message on the

¹The participants were informed at the time of recruitment that an estimated duration of the whole experimental procedure was two hours.

screen on their completion of Problem 1. (The message is presented in the instruction that is available in Appendix.) Hence, they were aware when a change from Problem 1 to Problem 2 was generated; that is, on their completion of Problem 1, they were advised that Problem 1 had been completed and they moved on Problem 2. The same procedure applied to when a change from Problem 2 to Problem 3 was generated. At the conclusion of the experiment, the participants were paid individually and privately at a conversion rate of one point to 0.3 Yen (about 0.25 US cent at the time of the experiment), and received no initial (showing up) fee.

3 Results and discussion

The overall maximisation rate (*choiceH*) is 0.48, 0.55 and 0.22 for Problem 1, 2 and 3 respectively. It follows that H, for example, was chosen on average 192 out of 400 times in Problem 1. Figure 2 illustrates *choiceH* for each problem. The individual *choiceH* is presented in Table 2.

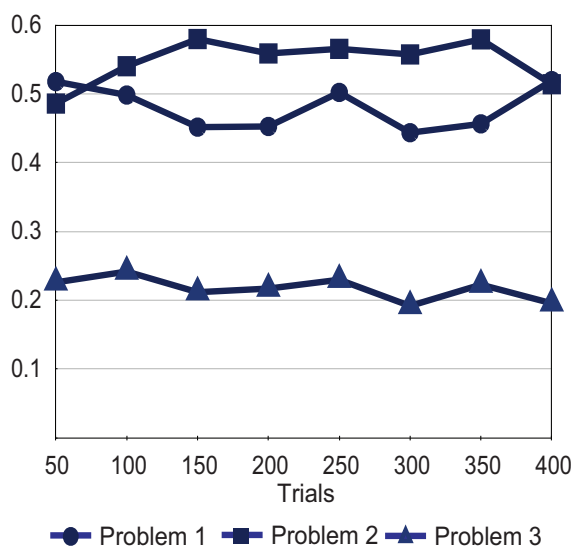


Figure 2: *choiceH* in Problem 1, 2 and 3

Here, I should like to raise a question as to what extent finding of the deviation from maximisation in Barron and Erev (2003) and Erev and Barron (2005) — also in my experiment — can be attributed to the hot stove effect, which appears to be as important as the payoff variability effect and the effect of reliance on small samples (and underweighting of rare outcomes). The payoff variability effect is a change of preference between two alternatives in experience-based binary decisions, associated with a change in the payoff variability of the alternatives. In the current choice problems, the payoff variability effect is what makes the DMs move toward random

Table 2: The individual overall *choiceH* in each problem, grouped by prevalent patterns of responding. An absolute difference of 5.5% or more from 50% is significant ($p < 0.05$ after Bonferroni correction). In Problem 1, 16 participants exhibited the stronger hot stove effect (i.e., their *choiceH* is less than 0.5). In Problem 2, 17 participants behaved toward random choice (i.e., their *choiceH* is not significantly deviated from 0.5). In Problem 3, more than half of the participants (26 of them) exhibited the stronger effect of underweighting of rare outcomes and the stronger hot stove effect (i.e., their *choiceH* is less than 0.22).

Problem 1	Problem 2	Problem 3
0.0000	0.6225	1.0000
0.6250	0.7100	0.4900
0.3350	0.6675	0.4850
0.6875	0.5200	0.5725
0.0950	0.5675	0.6000
0.5525	0.3350	0.6425
0.7175	0.5575	0.4925
0.5600	0.6075	0.7125
0.5350	0.5100	0.5800
0.7200	0.8775	0.4050
0.8250	0.4900	0.4800
0.5175	0.5900	0.4825
0.6625	0.3900	0.0275
0.6575	0.5200	0.0200
0.9050	0.6600	0.0075
0.7150	0.6550	0.2250
1.0000	0.4750	0.0175
0.5250	0.6725	0.0225
0.7000	0.5075	0.1925
0.7475	0.5200	0.0675
0.6025	0.4100	0.0050
0.7725	0.5425	0.2025
0.9925	0.4075	0.0075
0.9300	0.6625	0.0300
0.8425	0.5250	0.0125
0.6800	0.5800	0.0250
0.9750	0.5525	0.1400

	Problem 1	Problem 2	Problem 3
	0.0175	0.5125	0.0025
	0.2200	0.3600	0.0225
	0.0025	0.3475	0.0125
	0.0025	0.4600	0.0100
	0.0350	0.4825	0.0150
	0.0025	0.8775	0.0050
t	0.3400	0.7275	0.1575
	0.3675	0.4925	0.0375
	0.2275	0.5775	0.0725
	0.0000	0.4700	0.0025
	0.3825	0.6175	0.0025
	0.2400	0.4550	0.3575
	0.2675	0.6175	0.1375
	0.1925	0.3450	0.3400
	0.0050	0.5325	0.0000

choice between an alternative with higher expected value and an alternative with lower expected value when the payoff variability is associated with the alternative with lower expected value in experience-based binary decisions. Specifically, when the payoff variability of an *attractive* alternative (an alternative with higher expected value) increases, choice of the alternative decreases. On the other hand, when the payoff variability of an *unattractive* alternative (an alternative with lower expected value) increases, the DMs are sensitive to a bias toward random choice between both alternatives, rather than being sensitive to expected values. Erev and Barron (2005) describe the payoff variability effect as an obvious class of failures to maximisation. As said above, when the higher payoff variability is associated with an attractive alternative, the DMs would feel that it is less attractive. They then behave worse in terms of maximising expected value (by choosing an unattractive alternative often). When the higher payoff variability is associated with an unattractive alternative, they would be indifferent between an attractive and unattractive alternative so as to move toward random choice between both alternatives.

Denrell and March (2001) document that the hot stove effect implies a bias against a risky alternative in experience-based decisions, and the bias is a product of the tendency to reproduce actions that have been successful and avoid actions that led to loss. Thus, the hot stove effect implies a bias toward L (the low variability option) in Problem 1, 2 and 3. Here is the explanation that low payoffs from H reduce the probability of additional H

choices, and for that reason their effect of the estimated value from H is large. In an extreme case, a sequence of two “0” outcomes in Problem 3 can eliminate additional H choices and keep the participants’ estimate that H yields only “0” outcomes.

In Problem 1, reliance on small samples (and thus the effect of underweighting of rare outcomes) leads to the prediction, implying more H choices. Central to this prediction is the supposition that H provides better payoff (4 vs. 3) in most of trials (80%). The results of Barron and Erev (2003) and Erev and Barron (2005) suggest that the hot stove effect is *not* very strong: they observed almost 60% H choices in Problem 1. On the contrary, I observed almost 50% H choices in Problem 1 in the current experiment. The current results can be summarised with the assertion that they reflect a stronger hot stove effect that implies more L choices. It seems that the two effects (the hot stove effect and the effect of underweighting of rare outcomes) cancel each other and *choice*H is close to 50% in the current experiment. The hot stove effect appears to be as important as reliance on small samples (the effect of underweighting of rare outcomes), though Barron and Erev (2003) and Erev and Barron (2005) seem to have paid little attention to the hot stove effect in analysing behavioural tendencies in Problem 1.

In Problem 3, reliance on small samples — that causes underweighting of the attractive rare outcome (32 from H) — leads to the same prediction as the hot stove effect: it implies more L choices. Central to this prediction is the supposition that L provides better payoff (0 vs. 3) in most of trials (90%). Results in Barron and Erev (2003) and Erev and Barron (2005) reveal strong reliance on small samples (underweighting of the attractive rare outcome) when both reliance on small samples and the hot stove effect lead to the same predictions, though the authors do not further discuss the hot stove effect. They observed almost 30% H choices in Problem 3. On the contrary, I observed almost 20% H choices in Problem 3. The current results suggest that the participants’ less selection of H is the consequence of *both* the effect of underweighting of rare outcomes *and* the hot stove effect, as the two effects lead to the same prediction in Problem 3 (a bias toward L).

I suggest that *both* the payoff variability effect *and* the hot stove effect can account for behavioural tendencies in Problem 2, though much attention to the latter is not given by Barron and Erev (2003) and Erev and Barron (2005). They observed almost 50% H choices in Problem 2. Their results suggest that choice behaviour moves toward random choice in Problem 2, where the payoff variability is associated with an alternative with lower expected value — an alternative L.² Note that the payoff

²We can measure the payoff variability for an alternative that has two outcomes as follows: Variance of H in Problem 1, $s_{H1}^2 = 0.8(4 -$

variability is not associated with L in Problem 1 and 3, as it only yields a sure payoff of three points in the two problems. When the payoff variability of an unattractive alternative is maximal, the payoff variability effect implies a bias toward random choice: in Problem 2, both alternatives yield worst outcomes (“0”) for most of the rounds (i.e., 80% from H and 75% from L). Hence, the participants might have been considered to be indifferent between H and L. This might have caused a bias toward random choice. However, I observed 55% H choices in Problem 2 in the current experiment. This phenomenon seems to be caused by existence of the two effects — the payoff variability effect and the *attenuated* hot stove effect. The payoff structure in Problem 2 is more complicated than that in Problem 1 and 3, as both H and L involve uncertain prospect in Problem 2. Hence, the hot stove effect is attenuated in Problem 2, as there is more decay of the participants’ memory of past experience in Problem 2 than in Problem 1 and 3. Thus, both two effects can account for 55% H choices in Problem 2 in the current experiment: (1) the payoff variability effect, implying a bias toward random choice; and (2) the attenuated hot stove effect, implying a bias toward less L choices.

4 Concluding remarks

This paper has revisited the roles of mechanisms of individual decision making in experience-based decisions to complement a work of Barron and Erev (2003) and Erev and Barron (2005). They showed that the participants deviated from maximisation. They argued that the participants’ choice mainly reflected the payoff variability effect and the effect of underweighting of rare outcomes.

In this paper, I replicated the choice problems in Barron and Erev (2003) and Erev and Barron (2005) to re-examine their results. Consistent with their results, the participants in the current experiment deviated from maximisation. The current results suggested that choices were consistent with the prediction of the hot stove effect in addition to the payoff variability effect and the effect of underweighting of rare outcomes. Although the hot stove effect was not further discussed in Barron and Erev (2003) and Erev and Barron (2005), I found that the effect appears to be as important as the payoff variability effect and the effect of underweighting of rare outcomes in examining the behavioural tendencies in experience-based decisions. These conclusions are consistent with the fact that most of the clearest demonstrations of underweighting of rare events were observed in environments that

$3.2)^2 + 0.2(0 - 3.2)^2 = 2.56$, Variance of H in Problem 2, $s_{H2}^2 = 0.2(4 - 0.8)^2 + 0.8(0 - 0.8)^2 = 2.56$, and Variance of L in Problem 2, $s_{L2}^2 = 0.25(3 - 0.75)^2 + 0.75(0 - 0.75)^2 = 1.6875$. Variance of L in Problem 1 is zero.

control the hot stove effect by the addition of information concerning the forgone payoffs, with free sampling, or with forced sampling.

References

- Barron, G. & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, 16, 215–233.
- Coutu, D. (2006). Ideas as art. *Harvard Business Review*, 84, 82–89.
- Denrell, J. & March, J. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, 2, 523–538.
- Erev, I. & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, 112, 912–931.
- Erev, I., Ert, E., Roth, A., Haruvy, E., Herzog, S., Hau, R., Hertwig, R., Stewart, T., West, R. & Lebiere, C. (2009). A choice prediction competition, for choices from experience and from description. Manuscript submitted for publication.
- Ert, E. & Erev, I. (2007). Loss aversion in decisions under risk and the value of a symmetric simplification of prospect theory. Technion, Working Paper.
- Fox, C. R., & Hadar, L. (2006). “Decisions from experience” = sampling error + prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment and Decision Making*, 1, 159–161.
- Fujikawa, T. (2007). Perfect bayesian vs. imperfect bayesian in small decision making problems. *Behaviormetrika*, 34, 27–44.
- Fujikawa, T. & Oda, S. H. (2007). Judgement in small decision-making problems. In S. H. Oda (Eds.), *Developments on Experimental Economics*, pp. 149–154. Germany: Springer Verlag.

Appendix

Instruction (translation from Japanese)

Introduction

Thank you very much for joining our economics experiment. In this experiment you are asked to play easy games. Your goal is to complete the experiment with as many points as possible. The more points you earn, the more cash you can receive. The procedure of this experiment is explained along this instruction.

Distributions Please confirm whether you have received the following four items:

- Instruction (This leaflet)

- Questionnaire form
- Receipt form
- Subject NO card

Receipt Please write in the form your name, ID number, address, and the date today in advance. Keep the amount blank.

Notice

- You may NOT create a disturbance.
- You may NOT leave the laboratory during the experiment.
- You may keep switching your portable phone off during the experiment.
- You must leave all items distributed by personnel in the laboratory.
- You may NOT touch a keyboard.
- Do NOT click on right.
- You may NOT attempt to tamper with a computer.

Failure to comply with administrator’s directions can result in points you earned being cancelled and no money will be paid.

If you need an administrator If at any time during the experiment you believe you have a problem with your computer or need an administrator for any reason, raise your hand.

Payment At the conclusion of the experiment, points will be converted to monetary payoff according to the exchange rate: 100points = 30yen. The amount below 10 yen is rounded up.

Procedure

Registration Check that Figure 1 is displayed on your screen. (If it is not, raise your hand.) Click on a triangular button on your screen in order to equalise the number appeared on the screen with your subject number then press “Correct”. Assuming that your subject number is 19, press “Correct” in Figure 2.

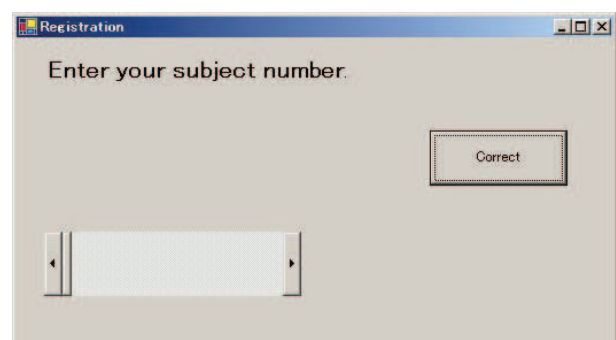


Figure 1

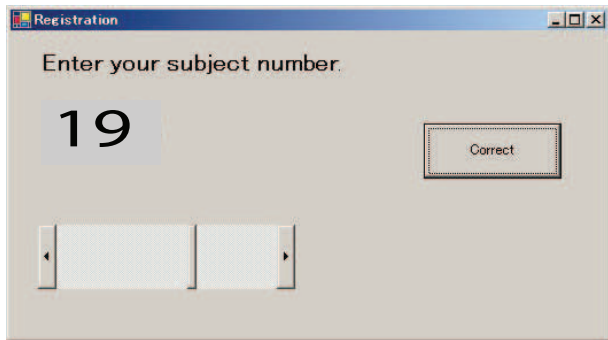


Figure 2

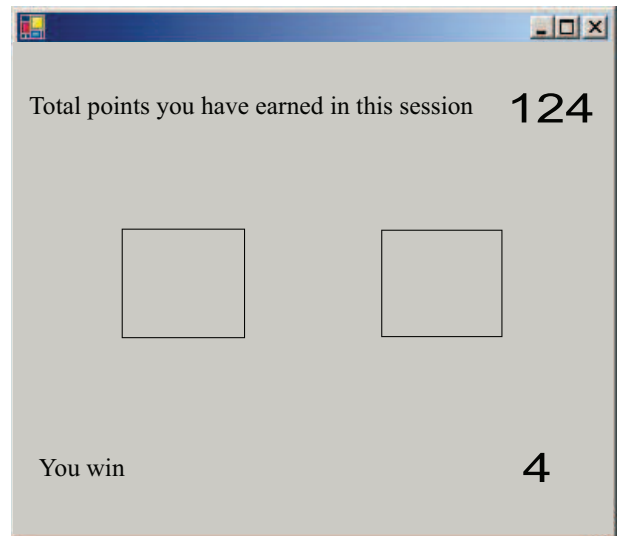


Figure 4

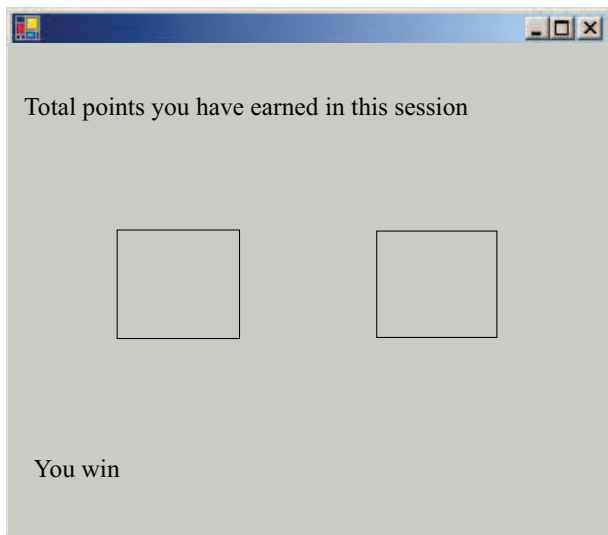


Figure 3




Figure 5



Figure 6

You are asked to play along this procedure for specific times. Points are contingent upon the button chosen. The different session has the different structure of the experiment. Your score is not affected by other's behaviour. An update of an accumulating score is constantly displayed on the right side of "Total points you have earned in this session". After completing each session, Figure 5 appears. Then Figure 6 appears after pressing "OK" in the Figure 5.

How to operate? The experiment consists of several sessions. Each session consists of several rounds. You are asked to choose either the right or the left button in each round as seen in Figure 3. The points corresponding to the selected button appear on the right side of "You win" (see Figure 4 as an example) and you can get it at that round. Your income is calculated by the computer.