The hot stove effect in repeated-play decision making under ambiguity

Fujikawa, Takemi

Universiti Sains Malaysia

3 October 2009

The hot stove effect in repeated-play decision making under ambiguity

Takemi Fujikawa
Universiti Sains Malaysia

Abstract

The "hot stove effect" has been studied for repeated-play decision making under uncertainty (also referred to as experience-based decision making) in which the decision makers repeatedly face the Allais-type binary choice problems, and have to learn about the outcome distributions through sampling as the decision makers are not explicitly provided with prior information on the payoff structure. The previous studies have found mixed evidence: some studies have found that the hot stove effect is strong in repeated-play decision making under uncertainty, while other studies have found that the effect is weak. Thus, the evidence is inconsistent. This paper reports an experimental investigation of the hot stove effect in repeated-play decision making under *ambiguity*. The current experiment involves an ambiguity treatment in which (1) the participants perform two binary repeated-play choice problems, each involving 400-fold choice between a risky option and a riskless option; and (2) in each problem, there are two states of nature available: a favourable state and an unfavourable state, but only one of them obtains on any given trial. The realisation of the actual state is not disclosed to the participants, thus they would be expected to discover the actual state through sampling with immediate feedback. The current results suggest that the magnitude of the hot stove effect is significantly different between repeated-play decision making under uncertainty and repeated-play decision making under ambiguity. I shall show that the hot stove effect is *attenuated* in repeated-play decision making under ambiguity.

We are living in an uncertain and ambiguous world, and facing decision making under uncertainty and ambiguity. Many popular definitions of "ambiguity" have been introduced (e.g., Einhorn & Hogarth, 1986; Klibanoff, Marinacci, & Mukerji, 2005; Mukerji, 1998; Mukerji & Tallon, 2001, 2003, 2004). This paper obeys the definition employed in Einhorn and Hogarth (1986) and Mukerji (1998): Ambiguity is an intermediate state between uncertainty (i.e., people receive no information to rule out any probability distribution possibilities) and risk (i.e., people receive one defined probability distribution).

I thank Hidenori Oda for his helpful comments and valuable research support. All errors remain my own. Comments should be sent to Takemi Fujikawa at takemifujikawa@gmail.com.

Ambiguity or "uncertainty about uncertainties" is a pervasive element of much real world decision making (Einhorn & Hogarth, 1986). An example of our real world decision making includes investment behaviour. Traders and investors often face the complexity of the situation and ambiguity in asset markets. Traders and investors do not often have perfect, prior information about asset markets, but ambiguous information about them. An example is concerned with investors who face a choice between a safe asset (e.g., cash) and a risky asset with an unknown return distribution. The distribution is unknown, but the investors know that the distribution obtains either a favourable distribution or an unfavourable distribution. If they consider the favourable distribution to be possible, then they may choose to invest in the risky asset. If, on the other hand, they consider the unfavourable distribution to be possible, then they consider not to invest in the risky asset (Easley & O'Hara, 2005).

This paper reports the "hot stove effect" in repeated-play decision making under ambiguity. The hot stove effect may push the traders and investors into trouble if they generalise from experience. An application of the hot stove effect to behavioural finance implies that repeated experience often leads the traders and investors to deviate from maximisation. 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) stated that the hot stove effect is a fundamental problem of learning that reduces the decision makers' (DMs') likelihood of repeating decisions that got them in trouble. Munichor, Erev, and Lotem (2006) provided the explanation on the hot stove effect that bad experience decreases the tendency to select the same alternative, even though it is a preferable alternative (i.e., an alternative with higher expected value). 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.

This paper aims to examine to what extent the hot stove effect persists in repeated-play decision making under ambiguity. For this aim, I implemented the current experiment that involved an ambiguity treatment in which (1) the participants performed two binary repeated-play choice problems, each involving 400-fold choice between a risky option and a riskless option; and (2) in each problem, there were two states of nature available: a favourable state and an unfavourable state, but only one of them obtained on any given trial. The current results suggest that the hot stove effect is *attenuated* in repeated-play decision making under ambiguity.

Previous research on repeated-play decision making under uncertainty, also referred to as "experience-based" decision making, has led to mixed conclusions with regard to the descriptive value of the hot stove effect. Whereas some studies (e.g., Denrell & March, 2001) demonstrated its importance, other studies (e.g., Barron & Erev, 2003; Erev & Barron, 2005) suggested that this effect is weak. Recently (Fujikawa, 2009), I accounted for the existence of the hot stove effect in repeated-play decision making under uncertainty. Below, we revisit the experiment in Fujikawa (2009).

Table 1: Choice problems in Fujikawa (2009). P_A (P_B) is % of A (B) choices over 400 rounds.

Problem	A	P_A	B	P_B
1a (N=42)	4, 0.8	48%	3, 1	52%
2a (N=42)	32, 0.1	22%	3, 1	78%

The hot stove effect in repeated-play decision making under uncertainty

In repeated-play decision making under uncertainty, the DMs have to learn about the outcome distributions through sampling, as they are not explicitly provided with prior information on the payoff structure. Previous and recent studies (e.g., Barron & Erev, 2003; Erev & Barron, 2005; Fox & Hadar, 2006; Hau, Pleskac, Kiefer, & Hertwig, 2008; Fujikawa, 2009) reported behavioural tendencies in repeated-play decision making under uncertainty.

Fujikawa (2009) presented an experiment that involved two choice problems (Problem 1a and 2a). Table 1 shows the payoff structure of each problem. For example, one selection of A in Problem 1a made the participants earn four points with probability 0.8 and zero point otherwise. Each problem included 400-fold binary choice between A (a risky alternative) and B (a riskless alternative). The participants' goal in each trial t ($t = 1, \ldots, 400$) was to select (click on) one of the two unmarked buttons that appeared on the 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. Table 1 shows the overall maximisation rate (the overall proportion of A choices) in each problem: for example, the overall maximisation rate was 0.48 in Problem 1a.

Notice that in Problem 1a the worst outcome (0 from A) is also the rare outcome (probability of 0.2). In Problem 1a, reliance on small samples and the hot stove effect lead to contradicting predictions. Reliance on small samples implies that the rare outcome (0 from A) will be *underweighted*: this prediction implies that A 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 A). Therefore, the hot stove effect implies that the worst outcome (0 from A) will be *overweighted*: this prediction implies that B will be preferred.

Denrell and March (2001) documented that the hot stove effect implies a bias against a risky alternative in experience-based decisions (i.e., repeated-play decision making under uncertainty), 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 B (the safe and low variability option) in Problem 1a and 2a. One explains that low payoffs from A reduce the probability of additional A choices, and for that reason their effect of the estimated value from A is large. In an extreme case, a sequence of two "0" outcomes in Problem 2a can eliminate additional A choices and keep the participants' estimate that A yields only "0" outcomes.

In Fujikawa (2009), I implemented Problem 1a and 2a to examine to what extent the hot stove effect could mediate behavioural tendencies in repeated-play decision making under uncertainty. In Problem 1a, the hot stove effect implies more B choices, while reliance on small samples (and thus the effect of underweighting of rare outcomes) leads to the prediction that implies more A choices. Central to this prediction is the supposition that A provides better payoff (4 vs. 3) in most of trials (80%). I observed almost 50% A choices in Problem 1a: it seems that the two effects (the hot stove effect and the effect of underweighting of rare outcomes) cancelled each other, and the maximisation rate was close to 50%. In Problem 2a, reliance on small samples — that causes underweighting of the attractive rare outcome (32 from A) — leads to the same prediction as the hot stove effect: it implies more B choices. Almost 20% A choices were observed in Problem 2a. This observation suggests the significant consequence of the hot stove effect in repeated-play decision making under uncertainty.

The hot stove effect in repeated-play decision making under ambiguity

Repeated-play decision making under uncertainty has been discussed in previous studies on experience-based decisions, where the participants made repeated choices under uncertainty (i.e., choices with *unknown* outcome distributions). That is, the participants received no prior information to rule out any outcome distribution possibilities. On the contrary, this paper presents the current experiment in which the participants made repeated choices under ambiguity.

Experiment. The current experiment was conducted at the Kyoto Experimental Economics Laboratory (KEEL) in Japan with 33 paid participants — 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 was read aloud and the participants were given an opportunity to ask questions individually.

The participants engaged in the following Problem 1b and 2b:

Problem 1b
State A: Choose between L: (4, 0.8) and R: (3, 1).
State B: Choose between L: (4, 0.7) and R: (3, 1).

Problem 2b
State A: Choose between L: (32, 0.1) and R: (3, 1).
State B: Choose between L: (32, 0.05) and R: (3, 1).

Each problem consisted of 400 rounds. The participants engaged in Problem 1b and 2b in order. At the beginning of each problem, they were presented with two equally likely states of the world: a priori relatively high state and a priori relatively low state. Undisclosed to the participants, in both problems, State A (a priori relatively high state) was set to be an actual state of the world, and State B (a priori relatively low state) was set to be a dummy state. They were not disclosed which of the two states of the world was being realised in each problem, but were disclosed that the same state of the world was being

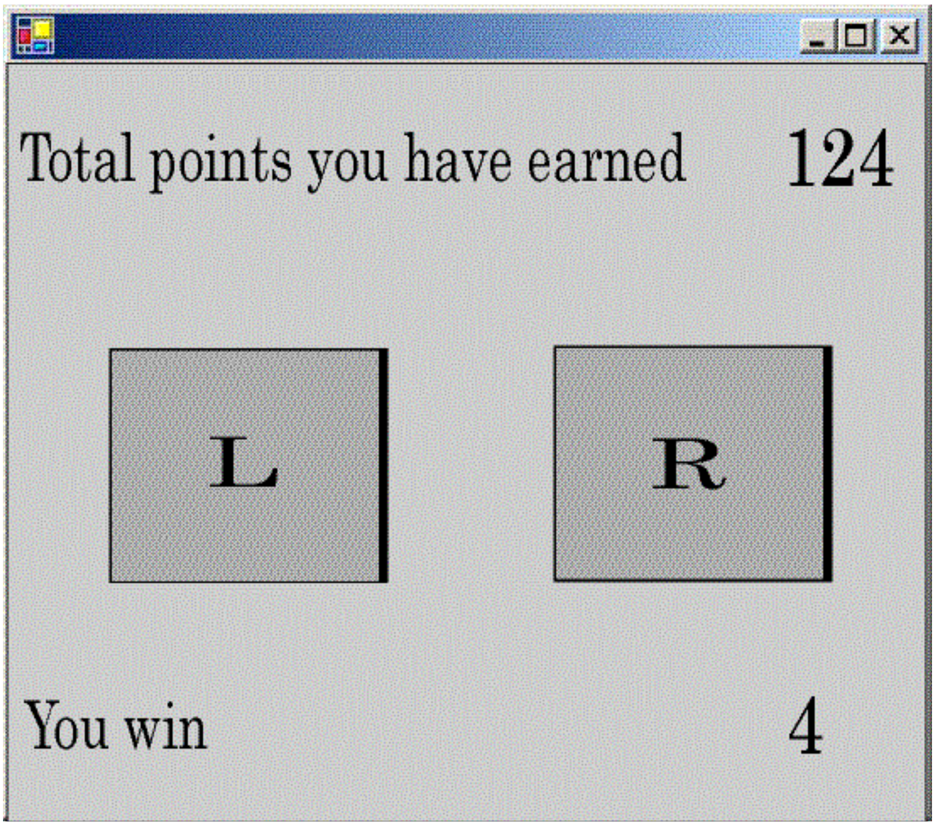

Figure 1. A computerised money machine

realised across the problem. The participants would be expected, through sampling with feedback, to discover which of the two states of the world was actually being realised in each problem.

In each problem, the participants were instructed to operate a "computerised money machine" shown in Figure 1. Their basic task at each round t ($t = 1, \ldots, 400$) was a binary choice between L and R in each problem. They made a choice between the two buttons on a computer screen to which each participant was assigned. In each trial, 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. They started with Problem 1b and made 400 selections in Problem 1b. Then, they were prompted to move to Problem 2b by the automatically-generated message on the computer screen on their completion of Problem 1b. Hence, they were aware when a change from Problem 1b to Problem 2b was generated; that is, on their completion of Problem 1b, they were advised that Problem 1b had been completed and they moved on Problem 2b. The money machine provided the participants with binary types of feedback immediately following each choice: (1) the payoff for the button chosen that appeared on the screen for the duration of one second and; (2) an update of an accumulating payoff counter, which was constantly displayed. 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.

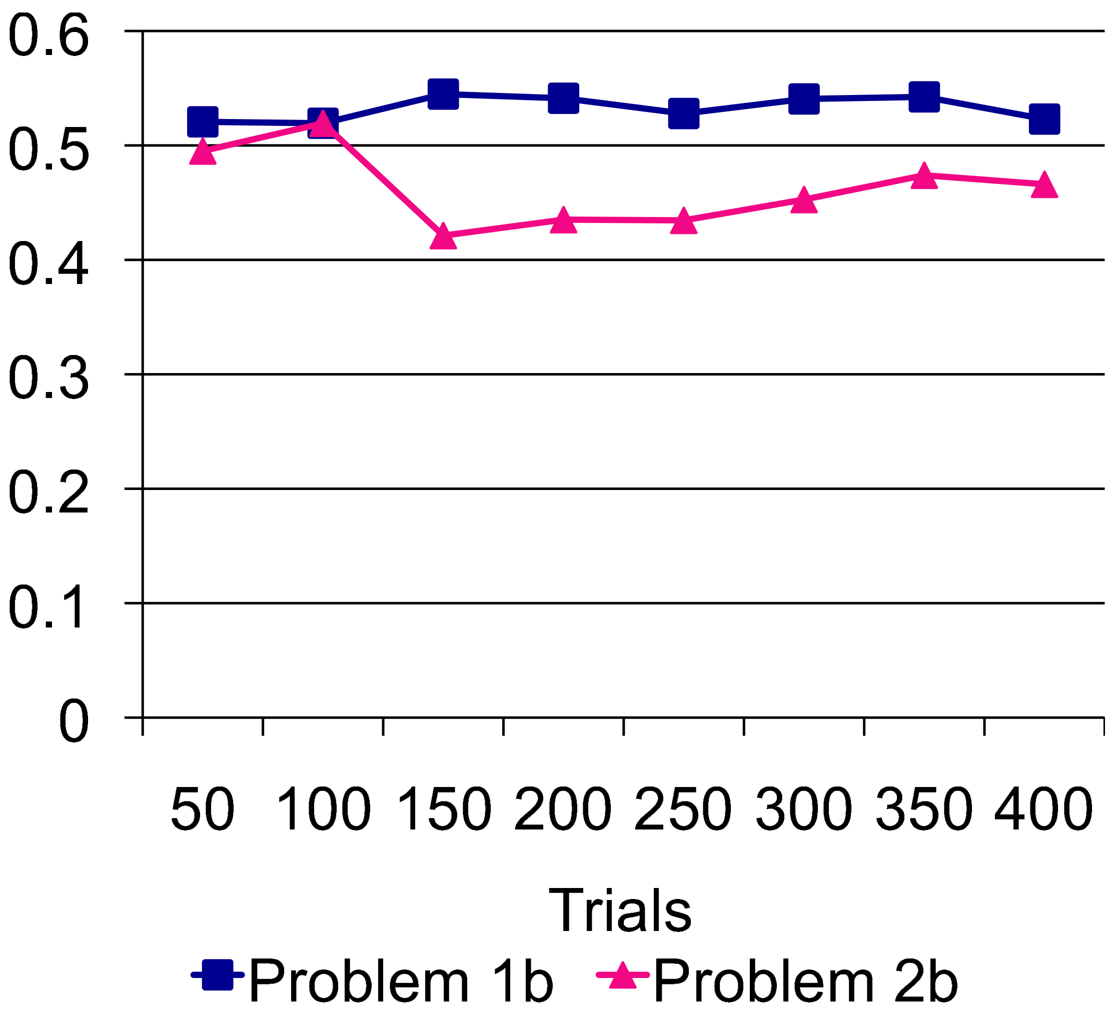

Figure 2. The maximisation rate in blocks of 50 rounds.

Results and Discussion

The overall maximisation rate (i.e., the proportion of L choices) is 0.53 and 0.46 in Problem 1b and 2b, respectively. It follows that L, for example, was chosen on average 212 out of 400 times in Problem 1b. The maximisation rate in Problem 1b is significantly larger ($p = 0.0000$) than in Problem 2b. Figure 2 illustrates the maximisation rate in blocks of 50 rounds. Table 2 shows a comparison between the maximisation rate observed in Fujikawa (2009) and that in the current experiment. The maximisation rate in the current experiment (i.e., the maximisation rate in Problem 1b and 2b) is significantly larger ($p = 0.0000$) than the maximisation rate in Fujikawa (2009) (i.e., the maximisation rate in Problem 1a and 2b).

Table 2: The overall maximisation rate in the current and previous experiments. The left-hand (right-hand) column presents the maximisation rate in Problem 1a (2a) in Fujikawa (2009) and in Problem 1b (2b) in the current experiment.

	Risky: (4, 0.8) Safe: (3, 1)	Risky: (32, 0.1) Safe: (3, 1)
Fujikawa (2009) (N=42)	48%	22%
Current experiment (N=33)	53%	46%

It can be seen that the *attenuated* hot stove effect and strong reliance on small samples can well explain behavioural tendencies in repeated-play decision making under ambiguity. This section presents the two-fold finding: (1) the hot stove effect gets attenuated in repeated-play decision making under ambiguity (the current experiment); and (2) reliance on small samples is strong in repeated-play decision making under ambiguity (the current experiment). To show the finding, I performed the comparison of the results of Problem 1a and 1b, and the comparison of the results of Problem 2a and 2b.

The hot stove effect implies more B and R choices in Problem 1a and 1b, respectively; while reliance on small samples (and thus the effect of underweighting of rare outcomes) leads to the prediction, implying more A and L choices, as A and L provide better payoff (4 vs. 3) in most of trials (80%). Fujikawa's (2009) results indicated less than 50% A choices in Problem 1a, implying the strong hot stove effect in repeated-play decision making under uncertainty. On the other hand, the current results indicate more than 50% L choices in Problem 1b, implying the *attenuated* hot stove effect and strong reliance on small samples in repeated-play decision making under ambiguity.

Both reliance on small samples and the hot stove effect lead to the same predictions, implying more B and R choices in Problem 2a and 2b, respectively. Fujikawa (2009) revealed strong reliance on small samples (underweighting of the rare outcome) and the hot stove effect in Problem 2a. On the contrary, the current results of Problem 2b suggest that, when the DMs engage in repeated-play decision making under ambiguity, they exhibit strong reliance on small samples and the *attenuated* hot stove effect, resulting the higher choice rate of the risky alternative than when the DMs engage in repeated-play decision making under uncertainty.

Concluding remarks

Previous studies on repeated-play decision making under uncertainty, also referred to as experience-based decision making, have led to mixed conclusions with regards to the descriptive value of the hot stove effect. Some (Denrell & March, 2001; Fujikawa, 2009) demonstrated its importance. Others (Barron & Erev, 2003; Erev & Barron, 2005) suggested that the effect is weak.

This paper aimed at examining to what extent the hot stove effect would persist in repeated-play decision making under ambiguity, implementing the current experiment that involved the Allais-type binary choice problems. For this aim, I reviewed my recent work (Fujikawa, 2009) on repeated-play decision making under uncertainty. The comparison of its results to the current results reveals that the hot stove effect is attenuated in repeated-play decision making under ambiguity.

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. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, *12*, 523-538.

Easley, D., & O'Hara, M. (2005). Regulation and return: The role of ambiguity. *Working Paper, Cornell University*.

Einhorn, H. J., & Hogarth, R. M. (1986). Decision making under ambiguity. *Journal of Business*, *59*, S225-S250.

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, *112*, 912-931.

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. (2009). On the relative importance of the hot stove effect and the tendency to rely on small samples. *Judgment and Decision Making*, *4*, 429-435.

Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description-experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, *21*, 493-518.

Klibanoff, P., Marinacci, M., & Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, *73*, 1849-1892.

Mukerji, S. (1998). Ambiguity aversion and the incompleteness of contractual form. *American Economic Review*, *88*, 1207-1231.

Mukerji, S., & Tallon, J. M. (2001). Ambiguity aversion and incompleteness of financial markets. *Review of Economic Studies*, *68*, 883-904.

Mukerji, S., & Tallon, J. M. (2003). Ellsberg 2-color experiment, portfolio inertia and ambiguity. *Journal of Mathematical Economics*, *39*, 299-316.

Mukerji, S., & Tallon, J. M. (2004). Ambiguity aversion and the absence of wage indexation. *Journal of Monetary Economics*, *51*, 653-670.

Munichor, N., Erev, I., & Lotem, A. (2006). Risk attitude in small timesaving decisions. *Journal of Experimental Psychology: Applied*, *12*, 129-141.