Economics Letters 156 (2017) 32–35

Contents lists available at ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

Framing effects and the reinforcement heuristic

Carlos Alós-Ferrer *, Sabine Hügelschäfer, Jiahui Li

University of Cologne, Department of Economics, Albertus-Magnus Platz, 50923 Cologne, Germany

HIGHLIGHTS

- We examine the effect of win/loss framing on a reinforcement heuristic.
- The experiment used an incentivized probability-updating task.
- A loss frame strengthens the basic impulse to shift away after losing.
- The effect occurs only if the losing option was chosen freely.

ARTICLE INFO

Article history:
Received 16 November 2016
Received in revised form
17 March 2017
Accepted 31 March 2017
Available online 12 April 2017

JEL classification:
C91
D03
D83

Keywords:
Framing
Reinforcement heuristic
Bayesian updating

ABSTRACT

We examine the effect of win/loss framing on individuals' use of a reinforcement heuristic in an incentivized probability-updating task. A loss frame strengthens the basic impulse to shift away from an unsuccessful option, but only if this option was chosen freely before.

1. Introduction

Framing effects describe the phenomenon that different but logically equivalent descriptions of the same decision problem lead to astonishingly different choices. For instance, people tend to avoid risks that are described in terms of benefits, but tend to take the same risks when described in terms of losses (Tversky and Kahneman, 1981). Such effects motivated the development of prospect theory (Kahneman and Tversky, 1979), which incorporates both loss aversion ("losses loom larger than gains") and different risk attitudes in the gain vs. the loss domain.

Optimal decision making under risk requires integrating all available information, which calls for the use of Bayes' rule. This is particularly true if the outcomes of previous decisions deliver information on underlying uncertain events. However, if those outcomes also provide feedback in a success/failure format (e.g., in the form of absolute or relative performance, profits and losses, etc.), human beings have a tendency to focus on past performance only. Previously successful decisions are repeated, and those that led to failure are revised, creating a simple "win-stay, lose-shift" decision rule. This "reinforcement heuristic," which might be an effective shortcut in simple settings, can conflict with normative behavior in more complex settings. Charness and Levin (2005) and Achtziger and Alós-Ferrer (2014) showed that, under such a conflict, individuals frequently rely on the faulty heuristic, hence committing many decision errors.

The present study investigates whether such errors are affected by framing, and specifically by whether a failure is presented as a loss or as the absence of a gain. There is some previous evidence (in simple settings) that loss frames are more motivating than gain frames, which is consistent with a general negativity bias as described in psychology (see, e.g., the reviews by Baumeister et al., 2001; Rozin and Royzman, 2001). Research on the behavioral effects of different reinforcement schedules (e.g., Rasmussen and

* Corresponding author.
E-mail addresses: carlos.alos-ferrer@uni-koeln.de (C. Alós-Ferrer), sabine.huegelschaefer@uni-koeln.de (S. Hügelschäfer), psylijiahui@gmail.com (J. Li).

http://dx.doi.org/10.1016/j.econlet.2017.03.038

Newland, 2008) has shown that monetary reinforcement and punishment are not symmetric, but rather that punishment is weighted more heavily. Accordingly, the results of several early studies involving children demonstrated that learning was better promoted by removing tokens for errors than by providing tokens for successes (e.g., Costantini and Hoving, 1973). In a student sample, Bereby-Meyer and Erev (1998) observed faster learning towards the optimal choice in a probability-learning task under negative compared to positive payoffs, in line with a model assuming that reinforcements are evaluated relative to an adjustable reference point. Similarly, Niznikiewicz and Delgado (2011) found enhanced learning effects and a stronger modulation of activity in the brain's reward circuitry when participants learned to avoid monetary losses, compared to when they learned to earn monetary rewards. However, whether a similar effect also holds for decision making in complex economic settings is still an open question.

Following the literature on framing and loss aversion, we concentrated on errors of the "lose-shift" type. We hypothesized that the impulse to shift away from an unsuccessful option would be strengthened when failure feedback is presented under a loss frame compared to a gain frame. We manipulated the framing of the goal of a decision task (win vs. loss frame) and tested the effect of this manipulation on participants' decision behavior. We hypothesized that a loss frame would result in more lose-shift errors compared to a gain frame.

2. Bayesian-updating task

To test our hypothesis, we relied on a probability-updating paradigm (Charness and Levin, 2005; Achtziger and Alós-Ferrer, 2014) in which participants are repeatedly confronted with situations where rational reasoning (Bayesian updating) conflicts with the reinforcement heuristic. Hence, participants can commit both win-stay errors and lose-shift errors. There are two urns, the Left Urn and the Right Urn, each containing 6 black and white balls. The urns are presented on the computer screen, with masked colors for the balls. Participants are asked to choose which urn a ball should be extracted from (with replacement) by pressing one of two keys, and are paid for drawing balls of a predefined color, say black (the winning ball color is counterbalanced across participants). After observing the result of the first draw, participants are asked to choose an urn a second time, and are paid again if the extracted ball is of the appropriate color. The payment per winning ball was 18 Euro cents.

The urn composition varies according to a "state of the world," Up or Down, which is not revealed to participants (see Table 1). The (known) prior probability of both states is 1/2. The state of the world is constant within the two-draw decision, but is randomized according to the prior for each new round. Hence, after the first draw, by observing the first ball's color, the decision maker can update his/her beliefs on the state of the world. For the second draw, an optimizer should choose the urn with the highest expected payoff, given the posterior probability of the state of the world updated through Bayes' rule.

The important decisions for our analysis are the second-draw decisions after the first draw was made from the Left Urn. Given the posterior probability updated through Bayes' rule, elementary computations show that in this case an optimizer should stay after a loss and switch after a win (win-shift, lose-stay), which is opposed to the prescriptions of reinforcement (see Charness and Levin, 2005; Achtziger and Alós-Ferrer, 2014, for details). Those situations correspond to a conflict among decision processes. In contrast, after a first draw from Right, optimizing behavior is fully aligned with behavior prescribed by a reinforcement heuristic,

Table 1
Urn composition in the Bayesian-updating task.

State (Prob)	Left Urn	Right Urn
Up (1/2)	● ● ● ● ○ ○	● ● ● ● ● ●
Down (1/2)	● ● ○ ○ ○ ○	○ ○ ○ ○ ○ ○

creating an alignment of processes which typically results in rather low error rates.

Participants repeated the two-draw decision 60 times. Following Achtziger and Alós-Ferrer (2014) and Charness and Levin (2005), we included both forced first draws (where the choice is dictated to the participant) and free first draws. Participants made forced draws and free draws alternately. Forced draws are included in the design to ensure enough first draws from Left, because, given the urn composition, choosing the Right Urn in the first draw reveals the state of the world and the decision for the second draw is straightforward. A simple computation shows that a sophisticated decision maker should always start with the Right Urn, as this maximizes the expected payoff for the two draws.

We hypothesized that a loss frame would strengthen the impulse to avoid unsuccessful decisions, and accordingly lead to more lose-shift errors after receiving negative feedback (i.e. drawing an unpaid or a losing ball in the first draw) from the Left Urn.

3. Experimental study

3.1. Methods

Participants. The study was conducted at the Cologne Laboratory for Economic Research (CLER) using z-Tree (Fischbacher, 2007). 64 participants were recruited via ORSEE (Greiner, 2015) and were randomly assigned to two framing conditions (win vs. loss frame) and two counterbalance conditions (winning ball color). In exchange for participation, they received a payment based on the outcomes in the decision task plus a show-up fee of 7.5 Euros. Two participants were excluded from data analysis due to failure to properly understand the instructions. Thus 62 participants (36 female, age range 19–35, $M = 24.6$, $SD = 3.70$) were considered for data analysis, 31 in each condition. Average earnings were 19.41 Euros ($SD = 0.83$), including the show-up fee.

Procedure. The manipulation was implemented by means of the initial experimental instructions regarding the goal of the decision task and the payment mechanism. In the win-frame condition, the goal was described as drawing as many winning balls as possible to earn as much money as possible. Participants were told that they would earn 18 Euro cents for each winning ball, and nothing for losing balls. In the loss-frame condition, the goal was described as drawing as few losing balls as possible to lose as little money as possible. Participants were told that they would be given an endowment of 36 Euro cents for each round (i.e., for each two-draw decision). From this endowment, 18 Euro cents would be deducted for each losing ball, and nothing for winning balls. Half of participants received win-framed instructions, and the other half received loss-framed instructions. The rest of the instructions was identical. Participants answered several control questions to ensure they understood the rules of the task. After that, they proceeded with the decision task, which lasted around 10 min.

3.2. Results

For all tests below, the unit of analysis is the individual-level error rate. That is, for each participant and each relevant class of errors, we compute the participant's percentage of errors and treat it as one observation.