Cognition 166 (2017) 272–276

Contents lists available at ScienceDirect

Cognition

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

Brief article

Conditioning task switching behavior

Senne Braem *

Department of Experimental Psychology, Ghent University, Belgium

ARTICLE INFO

Article history:
Received 4 March 2017
Revised 24 May 2017
Accepted 29 May 2017
Available online 5 June 2017

Keywords:
Cognitive control
Reward
Task switching
Reinforcement learning

ABSTRACT

Most reward studies focus on the reinforcement of simple tasks or stimulus-response rules. However, recent theories (re)emphasized that cognitive control representations should adhere to the same reinforcement learning principles as do more basic stimulus and response representations. This study focused on the act of switching between different tasks, and investigated the effects of disproportionally rewarding task alternations or repetitions in a cued task switching paradigm on subsequent voluntary task switching behavior (i.e., when participants could choose which task to perform). The results show that subjects who were more rewarded for task alternations (relative to those more rewarded for repetitions) showed more task switching behavior. Moreover, this increased task switching behavior also came with a cost, with participants more rewarded for task repetitions showing a better task focus (i.e., smaller task-rule congruency effects). These results demonstrate that reward can reinforce more abstract control representations, beyond low-level stimulus or response representations.

1. Introduction

Imagine cooking dinner while keeping an eye on the kids, and suddenly the phone rings. A situation like this presents a daily challenge for our cognitive control abilities: we have to adapt our behavior in response to novel stimuli in the environment (i.e., the phone suddenly ringing), switch back and forth between multiple tasks (i.e., cooking, watching the kids, and picking up the phone), and regulate our habits and impulses appropriately. In distinguishing cognitive control from other functions of the brain, it is often pitted against more automatic forms of low-level learning, such as associative learning or reinforcement learning. According to this distinction, cognitive control is not learned, but, instead, originates from a general supervisory system that corrects low-level learned behavior. This way, cognitive control and (reinforcement) learning are often portrayed as opposite forces in driving our behavior – reminiscent of the century-old dichotomy between the works of Tolman (1925) and Thorndike (1911). However, inspired by computational models that effectively modelled control functions in terms of simple reinforcement learning algorithms (Botvinick, Braver, Barch, Carter, & Cohen, 2001; Holroyd & Coles, 2002; Verguts & Notebaert, 2008), recent theories started to argue against this dichotomy, claiming that cognitive control adheres to, and might even be embedded in, the same (reinforcement) learning principles that drive more "low-level" forms of learning (Abrahamse, Braem, Notebaert, & Verguts, 2016; Egner, 2014; Verbruggen, McLaren, & Chambers, 2014).

Central to reinforcement learning is the century-old *law of effect* which posits that learned behavior is more likely to re-occur when it is followed by a reward (Thorndike, 1911). A growing number of studies already started to investigate the impact of reward on cognitive control (for a recent review, see Botvinick & Braver, 2015). However, these studies mostly focused on the anticipation of reward using explicit block-wise manipulations (e.g., Locke & Braver, 2008), cues (e.g., Padmala & Pessoa, 2011), or task features (e.g., Krebs, Boehler, & Woldorff, 2010; Umemoto & Holroyd, 2015) that predict reward. Although these experiments on reward anticipation are highly interesting for other reasons, they do not directly investigate the reinforcing effect of reward signals. To investigate how we learn from reward history (e.g., can reward strengthen the behavior that led to it?), we need studies that test the impact of reward delivery on subsequent behavior (for a similar distinction, see Notebaert & Braem, 2016).

This study zoomed in on the process of task switching, which is considered an important cognitive control function that allows us to flexibly switch attention between different tasks (Kiesel et al., 2010; Vandierendonck, Liefooghe, & Verbruggen, 2010). Specifically, an important open question is whether not only task sets themselves (Schiffer, Muller, Yeung, & Waszak, 2014; Umemoto & Holroyd, 2015), but also the overarching cognitive control process of task switching can be subject to reward modulations. No study thus far examined under closely controlled lab conditions

* Corresponding author at: Department of Experimental Psychology, Henri Dunantlaan 2, B – 9000 Ghent, Belgium.

E-mail address: Senne.Braem@ugent.be

http://dx.doi.org/10.1016/j.cognition.2017.05.037
0010-0277/

whether people can be conditioned into more, or less, switching between tasks. To this end, this study will disproportionally reward task alternations versus task repetitions, and test whether this can impact task switching behavior in a subsequent voluntary task switching phase (see method). The anticipated results would be inconsistent with the idea of cognitive control as a strategic supervisory system that is insensitive to more basic forms of reinforcement learning, and contradictory to the notion that reward promotes exploitative behavior only (e.g., Thorndike, 1911); but concordant with the hypothesis that cognitive control is sensitive to reinforcement by reward, just as any other type of behavior (Abrahamse et al., 2016; Verbruggen et al., 2014).

2. Method

Subjects performed a typical task switching study where their task was to categorize words on either their animacy or size (for a review, see Kiesel et al., 2010). They were indicated which task to perform during the first half of each block (cued task phase), and were free to choose a task during each second half (free task phase). Crucially, half of the subjects was rewarded more on task alternations, whereas the other half was rewarded more on task repetitions, during the cued task phase. We studied their task switching behavior during the free task phase, when no more reinforcements were distributed. By using a stimulus set to which both tasks could be applied we could also examine the task-rule congruency effect (see below), which can inform us about whether the ability to maintain task focus (and its shielding from the other task set) was also affected.

Participants. 49 students (range = 18–33 years, 28 women, 12 left-handed) took part in return for 10 €. The study was approved by the Ghent University Psychology and Educational Sciences Ethical Committee. The sample size was determined using sequential Bayes hypothesis testing by increasing the sample until a decisive Bayes factor smaller than 1/6 or larger than 6 was obtained (p. 12, Schönbrodt, Wagenmakers, Zehetleitner, & Perugini, 2017).

Stimuli and material. The task stimuli consisted of 320 high frequent words, which could all be evaluated on whether they were living or non-living, and whether they were smaller or larger than a basketball. Both dimensions were crossed orthogonally resulting in four lists of 80 words (i.e., living/small, living/large, non-living/small, and non-living/large; for more information, see Supplementary Material). Importantly, these 320 words were presented in a random order and ensured the presentation of a novel stimulus on each and every trial. The task cues consisted of the vowels A, E, I, O, or U, or the consonants V, F, L, Q, or C, to indicate one task or the other. A task cue would never reappear within three consecutive trials. During the second half of each block, task cues would be replaced by the symbol # to indicate a free choice of tasks. Task cues were presented slightly above, while the stimulus words were presented slightly below, the centre of the screen (5% of the distance between the centre and the upper or lower boundary of the screen, respectively). Feedback was centrally presented and consisted of the "+10" symbols for high reward, and "+01" for low reward. The Dutch word for false (i.e., "FOUT!") was used as error feedback in the training block. All instructions and stimuli were presented in a regular Arial 24-point font, in white on a black background. The left and right response keys were the letters S and D for the task assigned to the left hand, and the letters K and L for the task assigned to the right hand, on a standard QWERTY keyboard.

Procedure. The task was to categorize the word according to its animacy or size, depending on whether the above-presented letter was a vowel or a consonant. Specifically, participants had to press left (right) when the word was smaller (larger) than a basketball, with one hand, and to press left (right) when the word was non-living (living), with the other. This letter-type-to-task and task-to-response-hand assignment was counterbalanced across subjects, for each reward condition separately (see below). Participants were explained that living could refer to every living organism, including animals, trees, plants, fruits, or vegetables. This way, half of the words required a left or right button press on both tasks (i.e., both hands), whereas the other half required a different side response per task. These words were categorized as congruent and incongruent trials, respectively. Last, the participant was instructed that on some trials they would see the # symbol, instead of a vowel or consonant, upon which they should choose a task randomly, as if decided by flipping a coin (i.e., free choice trials).

The reward instructions indicated that they could win a (multimedia) store coupon of 50€ if they won the most points out of all participants. They were explained that they could win points on every trial for correct responses within the response deadline, but that the amount of points would be randomly determined (which was untrue, see below): sometimes this could be one point, sometimes ten points. Last, subjects were instructed that they could not win points on free choice trials, but that accurate performance and an honest attempt to choose tasks randomly was still important to remain in the competition. In other words, a failure to respond accurately on most free choice trials, or a failure to choose tasks randomly (for example, by constantly repeating one task), would exclude them from potentially winning the store coupon.

Importantly, and unknown to the participant, half of the participants were assigned to the switching group, whereas the other half was assigned to the repetition group. Participants in the switching group could receive ten points on 80% (and one point on the remaining 20%) of the trials where the current task was different from the preceding task (i.e., a task alternation), while they could receive ten points on only 20% (and one point on the remaining 80%) of the trials where the current task was similar to the preceding task (i.e., a task repetition). This reward manipulation was reversed in the repetition group, where repetitions were more highly rewarded.

After a first run of twelve practice trials (with task cues, no rewards, and four words that were not part of the above-mentioned lists), participants had to inform the experimenter if the task was still unclear, after which they were shown the instructions once more, and presented with four blocks of 80 trials. Each block consisted of 40 trials using task cues followed by 40 trials using free choice cues. The exact randomization is explained in the Supplementary Material. The trial procedure is visualized in Fig. 1. Each trial started with the presentation of two vertically aligned central fixation crosses for 500 ms, after which the task cue replaced the upper fixation cross and stayed on screen for 1000 ms. Next, the target word replaced the lower fixation cross, and both target word and task cue remained on screen until response, or the response deadline of five seconds. Following response or the response deadline, the screen turned blank for 500 ms, followed by a feedback presentation of another 500 ms (high or low reward following correct responses, blank screen following incorrect responses or response omission), and an inter-trial interval of 1000 ms. Free choice trials were identical to forced choice trials, except for the task cue always being a #, and the absence of feedback.

Questionnaires. After the experiment, participants filled in the BIS/BAS questionnaire (Carver & White, 1994) to assess their reward responsiveness. Analyses with this scale are reported in the Supplementary Material. Finally, participants were asked either verbally, or by questionnaire (the last nine subjects), whether they noticed an imbalance in the reward schedule: "Did