

Higher incentives can impair performance: neural evidence on reinforcement and rationality

Anja Achtziger,¹ Carlos Alós-Ferrer,² Sabine Hügelschäfer,² and Marco Steinhauser³

¹Department of Social Sciences, Zeppelin University Friedrichshafen, Friedrichshafen, Germany, ²Department of Economics, University of Cologne, Cologne, Germany, and ³Department of Psychology, Catholic University of Eichstätt-Ingolstadt, Eichstätt, Germany

Standard economic thinking postulates that increased monetary incentives should increase performance. Human decision makers, however, frequently focus on past performance, a form of reinforcement learning occasionally at odds with rational decision making. We used an incentivized belief-updating task from economics to investigate this conflict through measurements of neural correlates of reward processing. We found that higher incentives fail to improve performance when immediate feedback on decision outcomes is provided. Subsequent analysis of the feedback-related negativity, an early event-related potential following feedback, revealed the mechanism behind this paradoxical effect. As incentives increase, the win/lose feedback becomes more prominent, leading to an increased reliance on reinforcement and more errors. This mechanism is relevant for economic decision making and the debate on performance-based payment.

Keywords: ERPs; FRN; Bayesian updating; reinforcement; incentives

INTRODUCTION

Economic practice rests on the assumption that higher monetary incentives induce higher effort and improve performance. This view underlies many employee compensation plans, from piece-rate pay to performance boni. Incentive-based interventions have also been proposed to improve school performance and to control harmful habits (Gneezy *et al.*, 2011). It can be safely said that the question of whether monetary incentives increase performance is of fundamental importance for our society.

Research in economics (Camerer and Hogarth, 1999) and psychology (Jenkins *et al.*, 1998) shows that increased incentives result in increased performance in simple tasks. However, in general settings, the link between incentives and performance is complex and might break down at different points. First, monetary incentives might fail to elicit higher effort, e.g. if extrinsic incentives crowd out intrinsic motivation (Gneezy and Rustichini, 2000; Ariely *et al.*, 2009). Second, increased effort might fail to translate into increased performance, e.g. due to ceiling effects (Camerer and Hogarth, 1999). It has also been shown that ‘choking under pressure’ (Baumeister, 1984; Baumeister and Showers, 1986) might lead to a performance decrease in spite of increased effort.

In the present research, we concentrate on decision making under uncertainty and aim to demonstrate the existence of an additional, surprising mechanism leading to counter-intuitive effects of incentives on performance. We argue that reinforcement learning, a basic component of human behavior (Thorndike, 1911; Sutton and Barto, 1998), might cause increased incentives to produce more errors in many real-life decision tasks. The reasoning is as follows. Previous research (see below) has shown that the basic human tendency to repeat successful actions and avoid those which led to failure can impair performance by

focusing attention on win/lose outcomes and away from the probabilities of the relevant uncertain events. The objective of the present study is to show that, whenever reinforcement processes conflict with rational behavior, increased incentives might fail to increase performance because they make win/lose feedback more salient, leading to a stronger reliance on faulty reinforcement.

THE TROUBLE WITH REINFORCEMENT

Optimal decision making under uncertainty requires the integration of all available information. Rational decision makers should integrate new information with previous beliefs through the use of Bayes’ rule (Bayesian updating). However, many decisions lead to feedback in a win/lose or success/failure format, e.g. whether a business strategy leads to profits or losses, whether a medical treatment results in the improvement of a patient’s condition, or even whether a household appliance of a certain manufacturer performs satisfactorily or not. The association of success or failure with the original decision creates a tendency to rely on past performance as an indicator of future outcomes. A given decision will tend to be repeated if it led to success in the past, and to be changed if it led to failure. This ‘win-stay, lose-shift’ associative heuristic is the simplest form of reinforcement learning, and might be a shortcut to optimal decisions in simple settings. However, in a complex world, the outcome of a decision also delivers information on underlying uncertain events. Previous research (Charness and Levin, 2005; Achtziger and Alós-Ferrer, 2014) has shown that human decision makers frequently make use of a faulty ‘reinforcement heuristic’ which causes errors whenever a win-stay, lose-shift decision scheme conflicts with optimal decision making.

The conflict between reinforcement learning and rational decision making can be interpreted in terms of dual-process theories from psychology, according to which human decision making is influenced by competing processes of two broad types, automatic and controlled (Kahneman, 2003; Strack and Deutsch, 2004; Evans, 2008). Automatic processes, e.g. reinforcement, are fast, immediate, and require few cognitive resources and no conscious intent (Schneider and Shiffrin, 1977; Strack and Deutsch, 2004). Rationality in an economic sense prescribes normatively optimal solutions for each decision task, often leading to complex patterns which are closer to the idea of controlled processes.

The tendency to focus on past performance for the evaluation of decisions is especially relevant for economic and managerial decision

Received 30 October 2013; Revised 3 July 2014; Accepted 25 March 2015

Advance Access publication 29 March 2015

The authors are grateful to Armin Falk, Alexander Jaudas, Fernando Vega-Redondo, Nick Yeung, two anonymous referees, and seminar participants at the Center for Economics and Neuroscience (University of Bonn), the European University Institute (Florence), the 12th Meeting of the Social Psychology Section of the German Psychological Society (Luxembourg), the IAREP/SABE/ICABEEP Conference in Cologne, the Third Conference on Non-Cognitive Skills in Berlin, and the 2014 NeuroPsychoEconomics Conference in Munich for helpful comments. This work was supported by the Center for PsychoEconomics at the University of Konstanz.

Authors in alphabetical order. Correspondence should be addressed to Carlos Alós-Ferrer, Department of Economics, University of Cologne, Albertus-Magnus-Platz, 50923 Cologne, Germany. Email: carlos.alos-ferrer@uni-koeln.de

making (e.g. Ater and Landsmann, 2013), but is also relevant in other domains as e.g. medical decision making (Diwas et al., 2013). It is closely related to the outcome bias (Baron and Hershey, 1988), by which the evaluation of decisions is often based on outcomes only, neglecting other available information, and which has been shown to have negative consequences for manager evaluation and learning in organizations (Dillon and Tinsley, 2008). In such settings, successes and failures are often associated with explicit monetary consequences, and a natural question is whether increased incentives shift the balance towards more rational decision processes, and hence increase performance by reducing errors caused by a reinforcement heuristic.

We argue that this might not be the case. Our hypothesis was that, whenever reinforcement processes conflict with rational behavior, increased incentives might fail to increase performance because they make win/lose feedback more salient, leading to a higher reliance on faulty reinforcement. To examine this hypothesis, we conducted two studies on the effects of monetary incentives in a task in which optimal decisions could be achieved by means of Bayes' rule while application of a 'win-stay, lose-shift' heuristic led to errors. In a first study, immediate feedback on decision outcomes was made available, hence naturally enabling reinforcement processes. In a second study, feedback on decision outcomes was not provided, hence avoiding the cues on which reinforcement processes might act.

To directly measure automatic reinforcement in the brain, we analyzed event-related potentials (ERPs) associated with the processing of decision outcomes. ERPs are brain potentials time-locked to specific events (e.g. feedback) and are manifestations of brain activity occurring in response to these events (Fabiani et al., 2000). Because of their high time resolution, they are viewed as real-time indicators of psychological processes underlying thinking, feeling, and behavior. ERPs have already been used to investigate processes of economic decision making, particularly the processing of incentives in simple gambling tasks (Nieuwenhuis et al., 2004; Hajcak et al., 2006) and the role of heuristics in Bayesian updating tasks (Achtziger et al., 2014).

We focused on the feedback-related negativity (FRN), an ERP elicited by negative feedback (e.g. loss of money) on decisions, occurring about 200–300 ms after feedback (Miltner et al., 1997; Holroyd and Coles, 2002; Holroyd et al., 2003; Yeung and Sanfey, 2004). The FRN is linked to neural mechanisms of reinforcement learning (Holroyd and Coles, 2002; Holroyd et al., 2003), involving the basal ganglia and the dopamine system (Schultz et al., 1997). The basal ganglia evaluate the outcomes of ongoing behaviors against one's expectations. If an outcome is better or worse than anticipated, this information is conveyed to the dorsal anterior cingulate cortex, which is involved in cognitive control and has been identified as the most likely generator of the FRN (Holroyd and Coles, 2002). Communication happens via a phasic increase or decrease in the activity of midbrain dopaminergic neurons (Montague et al., 2004). According to the Reinforcement Learning Theory of the FRN (Holroyd and Coles, 2002), the FRN reflects a negative reward prediction error—a large FRN is generated whenever feedback indicates that an outcome is worse than expected. Due to its early occurrence, the FRN is considered as an indicator of automatic reinforcement learning, with its amplitude reflecting the decision maker's involvement in this process.

MATERIALS AND METHODS

Participants

Ninety-nine healthy, right-handed participants with normal or corrected-to-normal vision were recruited from the student community at the University of Konstanz (Germany), excluding students majoring in economics. In exchange for participation, participants received a show-up fee plus a monetary bonus that depended upon the outcomes

of the computer task. The individual sessions lasted ~90 min. Both studies were conducted in accordance with the ethical guidelines of the *American Psychological Association* (APA) and the Declaration of Helsinki; all participants signed an informed consent document before the experiment started.

A total of seven participants were excluded from data analysis due to a lack of understanding of the rules of the decision task. Seven further participants had to be eliminated from data analysis due to an excessive number of EEG trials containing artifacts and consequently too few valid trials to ensure reliable FRN indices. Of the remaining participants, 40 were in Study 1 (17 women and 23 men, ranging in age from 19 to 29 years, $M=21.9$, $s.d.=2.42$) and 45 in Study 2 (27 women and 18 men, ranging in age from 19 to 28 years, $M=22.3$, $s.d.=2.26$).

Procedure

Participants took part in the experiment individually assisted by two female experimenters in the EEG laboratory. Each participant was seated in a soundproof experimental chamber in front of a desk with a computer monitor and a keyboard. The experiment was run on a personal computer using Presentation® software 12.2 (Neurobehavioral Systems, Albany, CA). Stimuli were shown on a 19" computer monitor (resolution: 1024 × 768 pixels) at a distance of about 50 cm. The two keys on the keyboard that had to be pressed in order to make a decision ('F' key and 'J' key) were marked with glue dots to make them easily recognizable. Stimuli were images of two urns containing six balls each (image size 53 × 32 mm) and images of colored balls (blue and green, image size 14 × 14 mm) shown on the screen of the computer monitor. In order to reduce eye movements, the distance between the two urns was kept to a minimum (21 mm).

After application of the electrodes, each participant was asked to read through detailed instructions explaining the experimental setup. By means of an experimental protocol, the experimenter checked that the central aspects had been comprehended and clarified any misconceptions that the participant had with the rules or mechanisms. In addition, participants were instructed to move as little as possible during the computer task, to keep their fingers above the corresponding keys of the keyboard, and to maintain their gaze focused at the fixation cross in the center of the screen.

Each participant, while under supervision of the experimenter, completed three practice trials in order to become accustomed to the computer program. These trials did not differ from the following experimental trials except for the fact that outcomes were not rewarded. After the practice trials, the Bayesian updating experiment was started, during which the EEG was recorded. For the duration of the computer task, the participant was alone in the soundproof experimental chamber. When the experimental trials were completed, the amount of money earned during the task was displayed on the screen. Depending on the time the participant took for his/her decisions, the experimental task lasted about 10–15 minutes. After the computer experiment, participants filled out a questionnaire comprising several questions about personality characteristics, skills, and demographic information. Finally, participants were thanked, paid, and debriefed.

Decision task

The decision task, previously used by Charness and Levin (2005) and Achtziger and Alós-Ferrer (2014), captures the conflict between reinforcement learning and rational decision making. Uncertainty is represented by two urns, left and right, each containing six balls of different colors (blue or green). Participants choose which urn a ball should be randomly extracted from (with replacement), and are only

paid for drawing balls of a fixed color (e.g. blue; winning colors were counterbalanced between participants). After the outcome of this first draw is revealed, the participant chooses an urn again, a new ball is extracted, and the participant is paid again if the ball is of the winning color. The distribution of balls in the urns is known but varies depending on an unknown state of the world (up or down; Figure 1A). The sequence of events for Study 1 is depicted in Figure 1B. The state of the world is constant across both draws, and participants know the prior probability of the states (1/2). The composition of the urns is such that expected payoff for the first draw according to the prior is identical for both urns. For the second draw, the rational decision involves choosing the urn with the highest expected payoff, given the posterior probability updated from the prior through Bayes' rule after observing the color of the first extracted ball.

Participants repeated the two-draw decision 60 times. Following previous studies (Charness and Levin, 2005; Achtziger and Alós-Ferrer, 2014), in the first 30 rounds the initial draw was forced, alternatingly from left or right, and for the remaining rounds it was free. In this paradigm, many participants deviate from rational behavior by following reinforcement in a 'win-stay, lose-shift' form, i.e. choose the same urn as in the first draw if they won and switch otherwise (Charness and Levin, 2005). However, after an initial draw from left, Bayes' rule prescribes to switch to the right urn if the draw was successful, and stay if not. In this case both processes are in conflict (prescribe different responses), and errors are called reinforcement errors. After an initial draw from right, the opposite pattern is prescribed. In this case both processes are aligned, and errors are usually rare. Research on response times (Achtziger and Alós-Ferrer, 2014) has shown that reinforcement errors are faster than correct responses,

which is evidence for the automaticity of reinforcement in this paradigm.

In Study 1, we varied the magnitude of incentives (high vs low) to test for the effect of monetary incentives on reinforcement and performance. We conjectured that increased incentives would lead to increased reliance on reinforcement learning and fail to improve performance. There were 20 participants in a high-incentive condition (18 cents for each winning ball) and 20 in a low-incentive condition (9 cents for each winning ball). On average, participants earned 8.66 € (11.60 €, resp. 5.72 € per condition), in addition to a 5 € show-up fee.

To confirm that the automatic reinforcement process counteracts the effect of incentives, we conducted Study 2, in which reinforcement learning was blocked. This was implemented by disentangling information from win/lose feedback, using a removal-of-valence manipulation which has been shown to reduce error rates in our paradigm (Charness and Levin, 2005; Achtziger and Alós-Ferrer, 2014). The changes are as follows. All first draws are forced left and unpaid; the color of a rewarded ball in the second draw varies randomly from trial to trial and is revealed only shortly before the second draw. Hence, feedback on the first draw (ball color) now provides non-affective information about the decision outcome so that reinforcement processes are not elicited automatically, but participants can still use the information as an indicator of the distribution of the balls in the urns. Since only every second draw was remunerated, in Study 2 the payment for a winning ball was set at 18 cents for the low-incentive condition and 36 cents for the high-incentives one, keeping total incentives comparable across studies. The low-incentive condition consisted of 22 participants and the high-incentives condition of 23

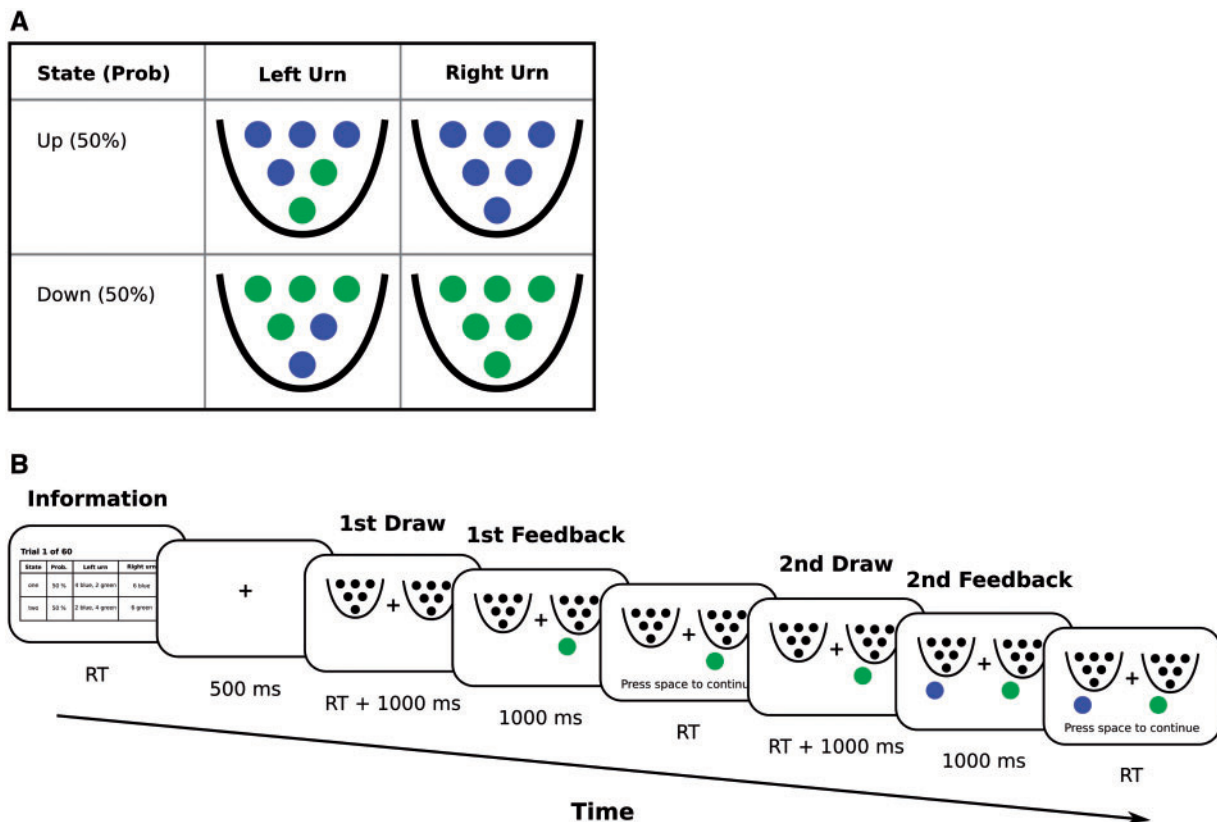


Fig. 1 A. Distribution of balls in the urns. In state 'up', the left urn contains four blue and two green balls, and the right urn contains only blue balls. In state 'down', the left urn contains two blue and four green balls, and the right urn contains only green balls. B. Sequence of stimuli and decisions in Study 1.

participants. On average, participants earned 8.55 € (11.38 €, resp. 5.59 € per condition), in addition to an 8 € show-up fee.

EEG Procedures

EEG acquisition

Data were acquired using Biosemi Active II system (BioSemi, Amsterdam, The Netherlands, www.biosemi.com) and analyzed using Brain Electrical Source Analysis (BESA) software (BESA GmbH, Gräfelfing, Germany, www.besa.de) and EEGLAB 5.03 (Delorme and Makeig, 2004). Data were acquired using 64 Ag–AgCl pin-type active electrodes mounted on an elastic cap (ECI), arranged according to the 10–20 system, and from two additional electrodes placed at the right and left mastoids. Eye movements and blinks were monitored by electro-oculogram (EOG) signals from two electrodes, one placed approximately 1 cm to the left side of the left eye and another one approximately 1 cm below the left eye (for later reduction of ocular artifact). As per BioSemi system design, the ground electrode during data acquisition was formed by the Common Mode Sense active electrode and the Driven Right Leg passive electrode. Both EEG and EOG were sampled at 256 Hz. All data were re-referenced off-line to an average mastoids reference and corrected for ocular artifacts with an averaged eye-movement correction algorithm implemented in BESA software.

FRN analysis

Stimulus-locked data were segmented into epochs from 500 ms before to 1000 ms after stimulus onset (presentation of the drawn ball); the prestimulus interval of 500 ms was used for baseline correction. In Study 1, epochs locked to first-draw positive and negative feedback were averaged separately, producing two average waveforms per participant. In Study 2, epochs were locked to first-draw pooled feedback. Epochs including an EEG or EOG voltage exceeding $\pm 100 \mu\text{V}$ were omitted from the averaging, in order to reject trials with excessive electromyogram (EMG) or other noise transients. Grand averages were derived by averaging these ERP waveforms across participants. Before subsequent analyses, the resulting ERP waveforms were filtered with a high-pass frequency of 0.1 Hz and a low-pass frequency of

30 Hz. On average 24% (Study 1) respectively 22% (Study 2) of trials were excluded due to artifacts.

To quantify the FRN in the averaged ERP waveforms for each participant, the peak-to-peak voltage difference between the most negative peak occurring 200–300 ms after feedback onset and the positive peak immediately preceding this negative peak was calculated (Yeung and Sanfey, 2004; Frank et al., 2005). This time window was chosen because previous research has found the feedback negativity to peak during this period (Miltner et al., 1997; Holroyd and Coles, 2002). In accordance with previous studies, the FRN amplitude was evaluated at channel FCz, where it is normally maximal (Miltner et al., 1997; Holroyd and Krigolson, 2007). To create topographies of the FRN, peak-to-peak differences were calculated at each electrode site using voltages at the peak latencies identified for channel FCz. Positive values (indicating the absence of a negative peak) were set to zero.

In the results section of the paper, we report on the FRN analysis for Study 1 only. In Study 2, no FRN on first-draw feedback was expected as no valence was attached to this draw. Indeed, the grand-average ERP waveform at FCz on first-draw feedback revealed no typical FRN in the time window 200–300 ms.

RESULTS

Behavioral Results

The average individual error rates for reinforcement errors were 58.3% in Study 1 and 25.1% in Study 2. Figure 2 illustrates the error rates disentangling *win-stay* and *lose-shift* errors and distinguishing participants with low and high incentives. The *win-stay* average error rate decreased from 62.8% (Study 1) to 21.9% (Study 2). *Lose-shift* error rates decreased from 52.4% (Study 1) to 28.4% (Study 2). Mann-Whitney-Wilcoxon tests (MWW, one-sided) on individual error rates confirmed a significant difference across studies for *win-stay* ($z = 4.68$, $P < 0.001$) and *lose-shift* errors ($z = 3.23$, $P < 0.001$). Please note that one-sided tests were used here because we had directional hypotheses on the effect of incentives and treatment effects. In Study 1, errors after a first draw from right, when both processes are aligned, were far less frequent (average error rate of 9.29%), as in Achtziger and Alós-Ferrer (2014). Error rates remained stable over time (see Supplementary Figure S1).

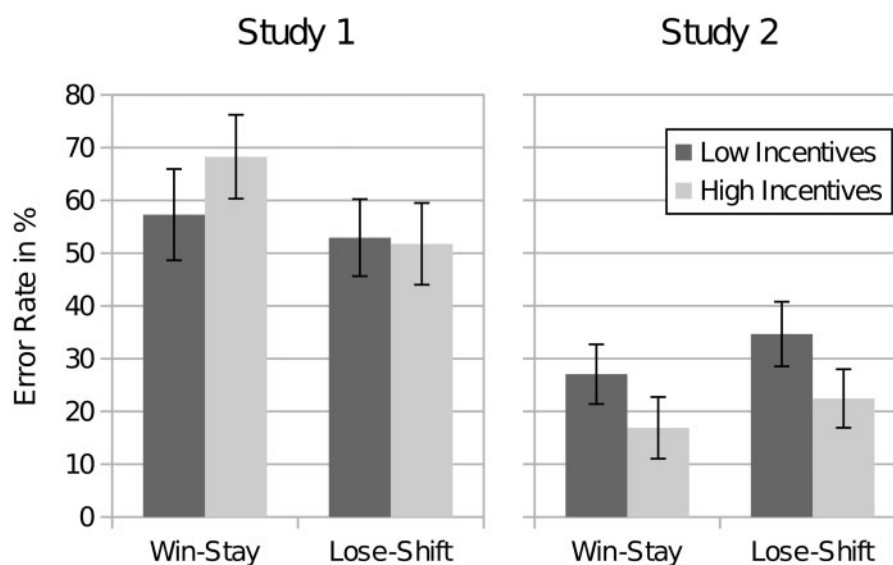


Fig. 2 Error rates in Studies 1 and 2 for reinforcement errors of the *win-stay* and *lose-shift* type, classified according to condition (high and low incentives). Error bars indicate one standard error from the mean.

Incentives did not significantly decrease error rates in Study 1 (MWW, one-sided; win-stay errors: low incentives 57.3%, high incentives 68.3%; $z = 0.78$, $P = 0.778$; lose-shift errors: low incentives 52.9%, high incentives 51.8%, $z = 0.04$, $P = 0.484$). In Study 2, however, error rates were lower in the high-incentive condition (MWW, one-sided; win-stay errors: low incentives 27.1% high incentives 16.9%, $z = 1.99$, $P < 0.05$; lose-shift errors: low incentives 34.7%, high incentives 22.4%, $z = 1.68$, $P < 0.05$). Hence, we find the expected link between higher incentives and increased performance in Study 2, where the reinforcement process was shut down, but this link is absent in Study 1, where the reinforcement process was active.

ERP Analysis

To analyze whether automatic reinforcement learning is responsible for the failure of incentives in Study 1, we analyzed the FRN following first-draw outcomes. A substantial FRN was observed, being more pronounced on negative (non-rewarded ball) than positive feedback (rewarded ball; Figure 3A).

To investigate the relation between the FRN and error rates in Study 1, participants were classified as having low or high reinforcement error rates using a median-split procedure, separately for each incentive condition. To test for differences in the FRN, we conducted a two-way analysis of variance (ANOVA) using the between-subjects factors incentive magnitude (low vs high) and reinforcement error rates. Because the FRN reflects a negative prediction error occurring on negative feedback only (Holroyd and Coles, 2002), we included only trials with negative feedback.

The ANOVA revealed a significant main effect of the reinforcement error rate, $F(1,36) = 4.44$, $P < 0.05$, $d^2 = 0.11$, i.e. the FRN was substantially more pronounced for participants with high reinforcement error rates (Figure 3B,C). However, this effect is due to the high-incentives condition. Indeed, the interaction between reinforcement error rates and incentives was significant, $F(1,36) = 7.05$, $P < 0.05$, $d^2 = 0.16$, supporting our main hypothesis. That is, the FRN amplitude was stronger for high incentives for high-error participants. Separate t -tests for the two incentive conditions confirmed this interpretation. For the high-incentives group (Figure 3H,I), low-error-rate participants had a significantly smaller FRN than those with high error rates (low rate: $M = 2.65$, s.d. = 1.98; high rate: $M = 8.80$, s.d. = 7.42; $t(7.671) = 2.29$, $P < 0.05$, one-sided). This effect was absent for the low-incentives condition (Figure 3E,F; low rate: $M = 3.96$, s.d. = 3.24; high rate: $M = 3.25$, s.d. = 2.69; $t < 1$). We remark for completeness that the main effect of incentives without separating participants according to error rates missed significance (Figure 3D,G; $F(1,36) = 2.70$, $P = 0.109$, $d^2 = 0.07$), but of course our hypothesis concerned the (significant) interaction.

The results indicate that, for high incentives, there is a link between reinforcement processes (revealed by the FRN) and error rates, which is absent under low incentives. To confirm this link beyond the median split, we analyzed the relation between reinforcement error rates and the FRN. Since individual error rates take values in the $[0,1]$ interval, no linear relation should be expected, hence we started the analysis by looking at Spearman's correlation. This revealed a strong positive correlation for the case of high incentives ($\rho = 0.449$, $P < 0.05$) but no relation for low incentives ($\rho = 0.047$, $P = 0.844$). To confirm this difference while further controlling for individual differences, we conducted a regression analysis. Again, since the dependent variable is the individual reinforcement error rate, which is bounded but is not a censored variable, linear or tobit regressions are not appropriate. We conducted a maximum-quasi-likelihood estimation for fractional dependent variables with binomial distribution, logit link function, and robust standard errors ('fractional logit'; Papke and Wooldridge,

1996). The regressions, which control for gender, age, whether the participant followed a statistics course, knowledge of probabilities, self-reported mastery of difficult problems, and self-reported effort confirm our results (Table 1; see [Supplementary material](#) for a discussion of the controls). As an illustration, separate regressions show a strong positive relation between FRN amplitude and reinforcement error rate for high incentives (first column of Table 1), but not in the case of low incentives (second column). Our main regression (third column in Table 1) pools all the data to confirm that the difference is significant, that is, the interaction of incentives and FRN is significant and positive, indicating a stronger correlation in case of high incentives.

DISCUSSION

Increased monetary incentives do not always result in increased performance. Research in economics and psychology has reported on a number of phenomena which greatly complicate the mappings from incentives to performance, including crowding out of intrinsic motivation (Gneezy and Rustichini, 2000; Ariely *et al.*, 2009), ceiling effects (Camerer and Hogarth, 1999), and choking under pressure (Baumeister, 1984). The present research identified an additional mechanism leading to counter-intuitive effects of incentives on performance. Strikingly, the basic human tendency to repeat successful actions and avoid those which led to failure (reinforcement learning; Thorndike, 1911; Sutton and Barto, 1998) can impair performance by focusing attention on win/lose outcomes and away from the probabilities of the relevant uncertain events.

We explored this phenomenon in two experimental studies. In Study 1, higher incentives failed to increase performance. Instead, higher incentives produced a positive correlation between FRN amplitudes and error rates. Since the FRN is a neural correlate of reinforcement learning (Holroyd and Coles, 2002; Holroyd *et al.*, 2003), this indicates that additional errors were created by an increased reliance on reinforcement processes. In the low-incentive condition, there was no correlation between FRN amplitudes and error rates. A natural explanation is that, while incentives do have a positive effect on performance, this effect is offset by automatic reinforcement learning, leading to the observed relationship between FRN and reinforcement error rates in the high-incentive condition. In Study 2, where the reinforcement process was blocked, the link between incentives and performance was restored. We conclude that increased incentives might induce errors, and the neural mechanism behind this link is a reinforcement learning process. Higher incentives make win/lose feedback more prominent and lead to a higher reliance on faulty reinforcement.

Some previous brain imaging studies have looked at paradoxical effects of large incentives. Mobbs *et al.* (2009) and Chib *et al.* (2012) examined performance in skilled tasks requiring eye-hand coordination and found that too high incentives led to performance decreases (choking under pressure). Mobbs *et al.* (2009) found that performance reduction for large incentives was correlated with midbrain activation as measured by the BOLD signal, and interpreted it as an overmotivation signal in the face of large incentives. Chib *et al.* (2012) showed that performance reductions correlated with ventral striatum deactivations for high incentives. Both were associated with a behavioral measure of loss aversion, which suggests that high incentives might be encoded as a potential loss arising from failure. Both studies are qualitatively aligned with ours in the sense that high incentives make the loss/gain feedback more prominent. However, a direct comparison is not possible because our paradigm uses a cognitive decision-making task, while Mobbs *et al.* (2009) and Chib *et al.* (2012) use manual ability tasks. Further, our Study 2, where higher incentives successfully

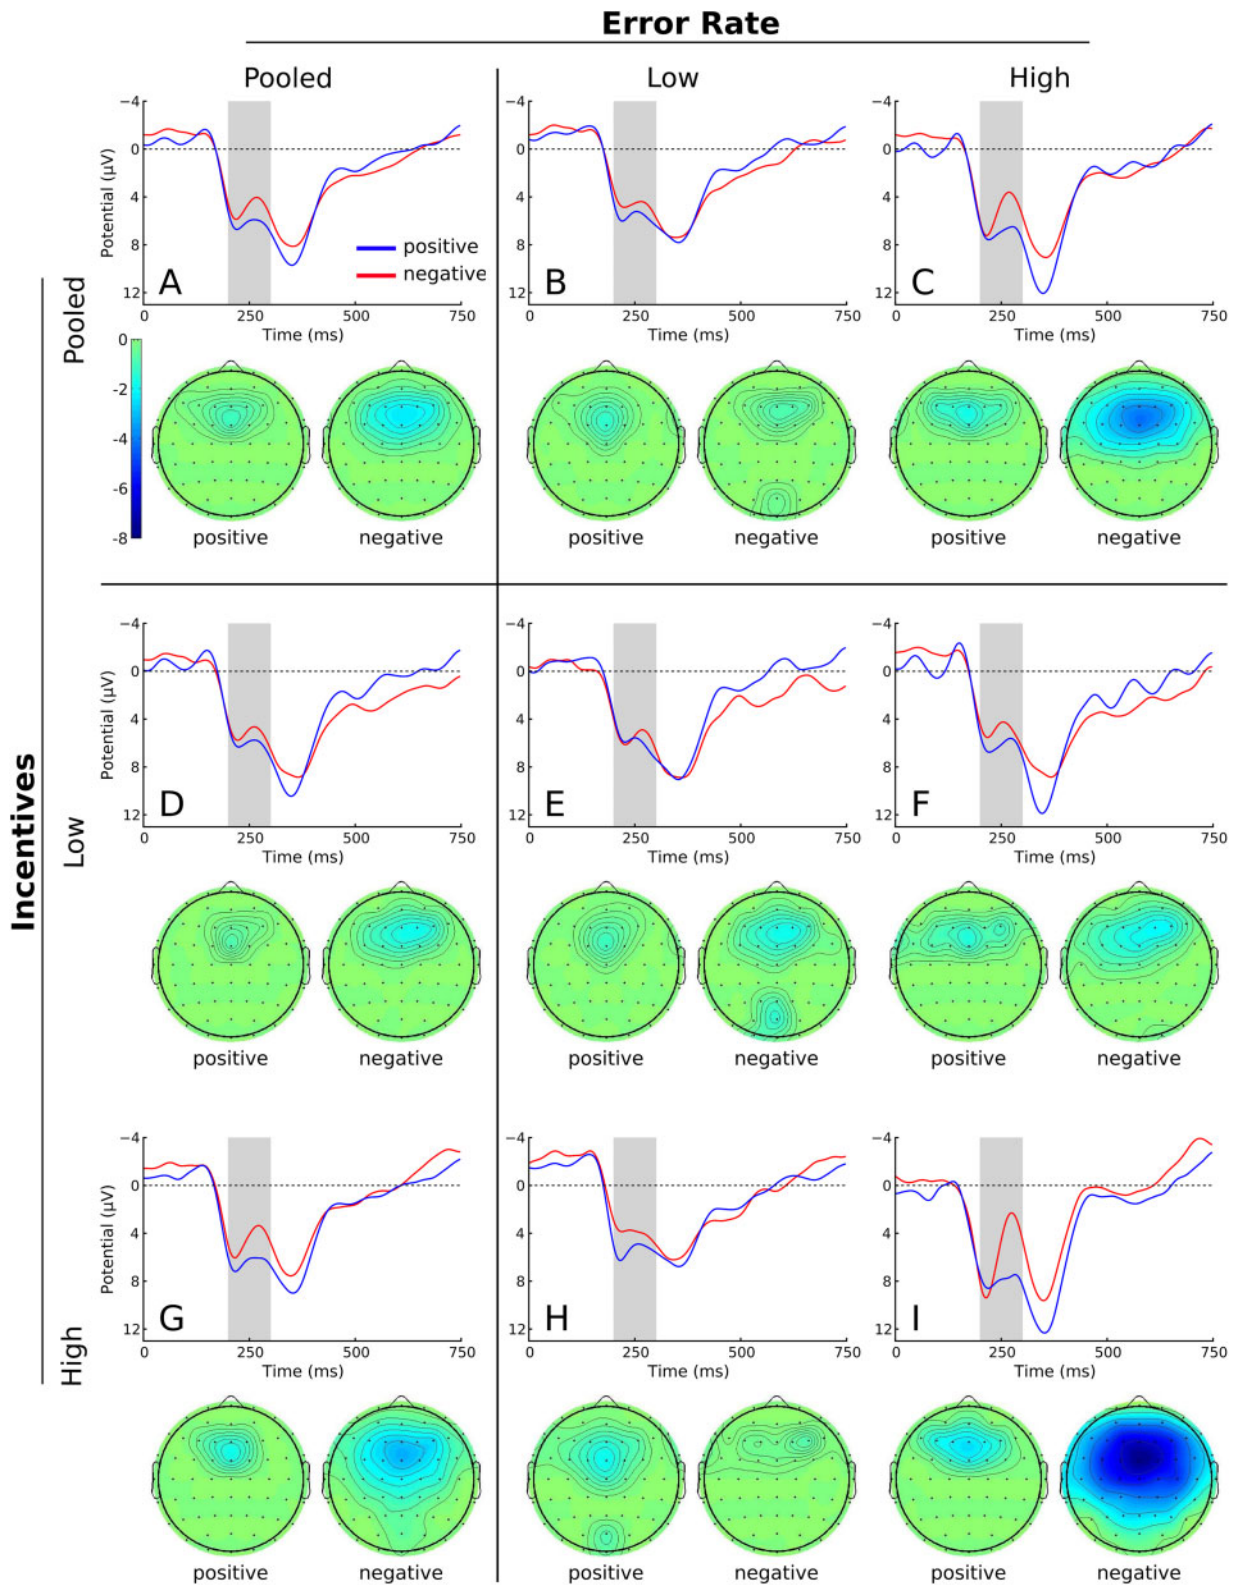


Fig. 3 Grand-average FRN as a function of feedback valence on the first draw, and corresponding topographies in Study 1. The nine panels correspond to the possible combinations of incentives (high, low, pooled) and reinforcement error rates following a median split (high, low, pooled). Waveforms were filtered (0.1/15 Hz) for presentation.

improved performance, shows that choking under pressure does not occur in our setting.

The insight that reinforcement processes might break the link from incentives to performance is highly relevant for economic decision

making and the debate on performance-based pay. Reinforcement processes are linked to extremely early brain responses and are very difficult to control. An excessive increase of performance-based monetary payments in e.g. managerial settings where similar decisions are

Table 1 Fractional logit estimations for reinforcement error rates

Variable	High incentives		Low incentives		Joint regression	
	β	SE	β	SE	β	SE
FRN (normalized)	0.531***	0.159	-0.346	0.493	-0.439	0.487
SVD	-0.898***	0.266	0.855***	0.310	0.283	0.228
Incentives (1=High)					1.139	0.951
Incentives X FRN					1.151**	0.519
Incentives X SVD					-0.509	0.313
Gender (1=Male)	-0.362	0.588	2.110**	0.916	-0.188	0.402
Age	-0.261***	0.098	-0.155***	0.055	-0.185***	0.067
Mastery	0.451**	0.229	1.208***	0.236	0.717***	0.150
Knowledge of Probabilities	-0.176	0.142	-0.282**	0.122	-0.383***	0.100
Statistics Course (1=Yes)	-1.016	0.674	3.604***	0.650	0.934*	0.485
Effort	0.186	0.272	-1.960***	0.401	-0.798***	0.215
Cb	0.831	0.730	-0.754*	0.421	-0.162	0.463
Constant	4.448*	2.708	7.935***	2.279	6.721***	2.057
pseudo-log likelihood	-7.002		-5.989		-14.731	

The individual FRN is the normalized FRN after negative feedback. Cb is counterbalance (1 if the winning color was blue). SVD is the subjective (self-reported) valence difference between the winning and the losing colors. Mastery and knowledge of probabilities (self-reported) were measured in a 0–10 scale. Statcourse is a dummy indicating that the participant reported having followed a statistics course. Effort was the self-reported effort invested in the task (0–10). $N=20$ for low incentives, $N=19$ for high incentives (1 subject dropped due to missing questionnaire data). * $P < 0.1$. ** $P < 0.05$. *** $P < 0.01$.

made frequently in a changing environment (lending and investment, supplier contracts, hiring, personnel allocation) might lead to an increased weight of reinforcement-based decisions (ex-post justified as ‘managerial gut feeling’) instead of increasing performance. Our research also points out the need for decision-making strategies removing the emotional attachment generated by win/lose outcomes, hence restoring the incentives-performance link.

SUPPLEMENTARY DATA

Supplementary data are available at SCAN online.

Conflict of Interest

None declared.

REFERENCES

Achtziger, A., Alós-Ferrer, C. (2014). Fast or rational? A response-times study of Bayesian updating. *Management Science*, 60, 923–38.

Achtziger, A., Alós-Ferrer, C., Hügelschäfer, S., Steinhauser, M. (2014). The neural basis of belief updating and rational decision making. *Social Cognitive and Affective Neuroscience*, 9, 55–62.

Ariely, D., Gneezy, U., Loewenstein, G., Mazar, N. (2009). Large stakes and big mistakes. *Review of Economic Studies*, 76, 451–69.

Ater, I., Landsmann, V. (2013). Do customers learn from experience? Evidence from retail banking. *Management Science*, 59, 2019–35.

Baron, J., Hershey, J.C. (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology*, 54, 569–79.

Baumeister, R.F. (1984). Choking under pressure: self-consciousness and paradoxical effects of incentives on skillful performance. *Journal of Personality and Social Psychology*, 46, 610–20.

Baumeister, R.F., Showers, C.J. (1986). A review of paradoxical performance effects: Choking under pressure in sports and mental tests. *European Journal of Social Psychology*, 16, 361–83.

Camerer, C.F., Hogarth, R.M. (1999). The effects of financial incentives in experiments: a review and capital-labor-production framework. *Journal of Risk and Uncertainty*, 19, 7–42.

Charness, G., Levin, D. (2005). When optimal choices feel wrong: a laboratory study of Bayesian updating, complexity, and affect. *American Economic Review*, 95, 1300–9.

Chib, V.S., De Martino, B., Shimojo, S., O’Doherty, J.P. (2012). Neural mechanisms underlying paradoxical performance for monetary incentives are driven by loss aversion. *Neuron*, 74, 582–94.

Delorme, A., Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, 134, 9–21.

Dillon, R.L., Tinsley, C.H. (2008). How near-misses influence decision making under risk: a missed opportunity for learning. *Management Science*, 54, 1425–40.

Diwas, K.C., Staats, B.R., Gino, F. (2013). Learning from my success and from others’ failure: Evidence from minimally invasive cardiac surgery. *Management Science*, 59, 2435–49.

Evans, J.S.B.T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255–78.

Fabiani, M., Gratton, G., Coles, M.G.H. (2000). Event-related brain potentials. In: Cacioppo, J.T., Tassinari, L.G., Berntson, G.G., editors. *Handbook of psychophysiology* 2nd edn. New York, NY: Cambridge University Press, pp. 53–84.

Frank, M.J., Worocho, B.S., Curran, T. (2005). Error-related negativity predicts reinforcement learning and conflict biases. *Neuron*, 47, 495–501.

Gneezy, U., Rustichini, A. (2000). Pay enough or don’t pay at all. *Quarterly Journal of Economics*, 115, 791–810.

Gneezy, U., Meier, S., Rey-Biel, P. (2011). When and why incentives (don’t) work to modify behavior. *Journal of Economic Perspectives*, 25, 191–210.

Hajcak, G., Moser, J.S., Holroyd, C.B., Simons, R.F. (2006). The feedback-related negativity reflects the binary evaluation of good versus bad outcomes. *Biological Psychology*, 71, 148–54.

Holroyd, C.B., Coles, M.G.H. (2002). The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, 109, 679–709.

Holroyd, C.B., Krigolson, O.E. (2007). Reward prediction error signals associated with a modified time estimation task. *Psychophysiology*, 44, 913–17.

Holroyd, C.B., Nieuwenhuis, S., Yeung, N., Cohen, J.D. (2003). Errors in reward prediction are reflected in the event-related brain potential. *NeuroReport*, 14, 2481–84.

Jenkins, G.D., Jr., Mitra, A., Gupta, N., Shaw, J.D. (1998). Are financial incentives related to performance? A meta-analytic review of empirical research. *Journal of Applied Psychology*, 83, 777–87.

Kahneman, D. (2003). Maps of bounded rationality: psychology for behavioral economics. *American Economic Review*, 93, 1449–75.

Miltner, W.H.R., Braun, C.H., Coles, M.G.H. (1997). Event-related brain potentials following incorrect feedback in a time-estimation task: evidence for a ‘generic’ neural system for error detection. *Journal of Cognitive Neuroscience*, 9, 788–98.

Mobbs, D., Hassabis, D., Seymour, B. et al. (2009). Choking on the money. Reward-based performance decrements are associated with midbrain activity. *Psychological Science*, 20, 955–62.

Montague, P.R., Hyman, S.E., Cohen, J.D. (2004). Computational roles for dopamine in behavioural control. *Nature*, 431, 760–7.

Nieuwenhuis, S., Yeung, N., Holroyd, C.B., Schurger, A., Cohen, J.D. (2004). Sensitivity of electrophysiological activity from medial frontal cortex to utilitarian and performance feedback. *Cerebral Cortex*, 14, 741–7.

Papke, L.E., Wooldridge, J.M. (1996). Econometric methods for fractional response variables with an application to 401(k) plan participation rates. *Journal of Applied Econometrics*, 11, 619–32.

Schneider, W., Shiffrin, R.M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, 84, 1–66.

Schultz, W., Dayan, P., Montague, P.R. (1997). A neural substrate of prediction and reward. *Science*, 275, 1593–9.

Strack, F., Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, 8, 220–47.

Sutton, R.S., Barto, A.G. (1998). *Reinforcement learning: an introduction*. Cambridge, MA: MIT Press.

Thorndike, E.L. (1911). *Animal intelligence. Experimental studies*. Oxford, England: Macmillan.

Yeung, N., Sanfey, A.G. (2004). Independent coding of reward magnitude and valence in the human brain. *The Journal of Neuroscience*, 24, 6258–64.