NeuroImage 125 (2016) 868–879

Contents lists available at ScienceDirect

NeuroImage

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

Randomised prior feedback modulates neural signals of outcome monitoring

Faisal Mushtaq [a,*], Richard M. Wilkie [a], Mark A. Mon-Williams [a], Alexandre Schaefer [b,**]

[a] School of Psychology, University of Leeds, Leeds, West Yorkshire, UK
[b] School of Business, Monash University, Sunway Campus, Selangor, Malaysia

A R T I C L E I N F O

Article history:
Received 13 June 2015
Accepted 18 October 2015
Available online 20 October 2015

Keywords:
Feedback-related negativity
Reward prediction error
Short-term expectations
Salience
Monitoring
Uncertainty
Affective priming

A B S T R A C T

Substantial evidence indicates that decision outcomes are typically evaluated relative to expectations learned from relatively long sequences of previous outcomes. This mechanism is thought to play a key role in general learning and adaptation processes but relatively little is known about the determinants of outcome evaluation when the capacity to learn from series of prior events is difficult or impossible. To investigate this issue, we examined how the feedback-related negativity (FRN) is modulated by information briefly presented before outcome evaluation. The FRN is a brain potential time-locked to the delivery of decision feedback and it is widely thought to be sensitive to prior expectations. We conducted a multi-trial gambling task in which outcomes at each trial were fully randomised to minimise the capacity to learn from long sequences of prior outcomes. Event-related potentials for outcomes (Win/Loss) in the current trial (Outcome$_t$) were separated according to the type of outcomes that occurred in the preceding two trials (Outcome$_{t-1}$ and Outcome$_{t-2}$). We found that FRN voltage was more positive during the processing of win feedback when it was preceded by wins at Outcome$_{t-1}$ compared to win feedback preceded by losses at Outcome$_{t-1}$. However, no influence of preceding outcomes was found on FRN activity relative to the processing of loss feedback. We also found no effects of Outcome$_{t-2}$ on FRN amplitude relative to current feedback. Additional analyses indicated that this effect was largest for trials in which participants selected a decision different to the gamble chosen in the previous trial. These findings are inconsistent with models that solely relate the FRN to prediction error computation. Instead, our results suggest that if stable predictions about future events are weak or non-existent, then outcome processing can be determined by affective systems. More specifically, our results indicate that the FRN is likely to reflect the activity of positive affective systems in these contexts. Importantly, our findings indicate that a multifactorial explanation of the nature of the FRN is necessary and such an account must incorporate affective and motivational factors in outcome processing.

Introduction

The ability to rapidly update information about reward probability is necessary for goal-directed behaviour. Monitoring and evaluating an outcome relative to prior expectations is essential to this process (Kerns et al., 2004; Schall et al., 2002; Sohn et al., 2007). A large body of research on outcome monitoring has focused on a scalp event-related potential known as the feedback-related negativity (FRN; Gehring and Willoughby, 2002; Miltner et al., 1997). The FRN is usually operationalised as a contrast between ERPs to negative and positive outcomes. It has a frontocentral topography and is characterised by a negative deflection maximal at ~250–350 ms after feedback onset that is larger for non-reward compared to reward outcomes (Ferdinand et al., 2012; Yeung and Sanfey, 2004). Substantial evidence indicates that the FRN is linked to activity in medial frontal areas including the ACC (e.g. Hauser et al., 2014). The FRN is influenced by outcome history and varies as a function of prior reward expectations and probability: In fact, the FRN produced in the majority of previous studies relates to information that relies on learning contexts established over multiple trials and blocks (Bellebaum et al., 2010a; Cohen and Ranganath, 2007; Donkers et al., 2005; Pfabigan et al., 2010; Pietschmann et al., 2011; Sailer et al., 2010; Santesso et al., 2008).

A few studies have shown that the influence of prior outcome history on the FRN can be observed on very brief time scales. Specifically, information presented in the trial immediately preceding a current trial (Outcome$_{t-1}$) can modulate the FRN related to the current trial (Outcome$_t$) (e.g. Gehring and Willoughby, 2002; Goyer et al., 2008; Holroyd and Coles, 2002). These findings suggest that the FRN is sensitive to factors beyond the learning of probabilistic relationships

* Correspondence to: F. Mushtaq, School of Psychology, University of Leeds, Leeds LS2 9JT, UK.
** Corresponding author.
E-mail addresses: f.mushtaq@leeds.ac.uk (F. Mushtaq), alexandre.schaefer@monash.edu (A. Schaefer).

http://dx.doi.org/10.1016/j.neuroimage.2015.10.046

between events and outcomes over a long period of time. These results are of importance to the field of decision neuroscience as they suggest that expectations can be formed rapidly and (ultimately) bias decision-making on a very short time scale. Whilst the determinants of FRN effects have been the focus of intense debates in recent years, the processes that modulate short-term effects on the FRN have received relatively little attention (we refer to these effects as the stFRN hereafter). Examining the stFRN promises not only to shed light on the debate about the determinants of the FRN, but also speaks to the broader issue of how the brain keeps track of changing expectancies in a rapidly-changing environment. Thus, the goal of this study is to test four contrasting explanations of stFRN encoding effects derived from existing FRN models.

First, the most prevalent account of the FRN has been provided by the reinforcement learning error-related negativity theory (RL-ERN; Holroyd and Coles, 2002). The original version of the RL-ERN theory proposed that the FRN indexes negative reward prediction error (-RPE) — a key component of Reinforcement Learning theories (e.g. Sutton and Barto, 1998). A -RPE occurs when an event (e.g. a decision outcome) violates a prediction learned from previous outcomes in such a way that the event constitutes an outcome that is "worse than expected". More specifically, the RL-ERN theory posited that dopaminergic systems in mesencephalic areas monitor and detect when learned predictions are violated. When a -RPE is detected, there is a decrease in dopaminergic firing rate. This change in dopamine activity produces a signal that is sent to the ACC, causing a disinhibition of ACC neurons and thus leads to a larger FRN (Holroyd and Coles, 2002). These prediction error signals are suggested to be signals that trigger the implementation of top-down cognitive control processes (Kerns et al., 2004; Mushtaq et al., 2011; Swainson et al., 2003).

Second, a number of studies have reported a greater FRN not only for -RPE, but also for positive reward prediction errors (+RPE, e.g. Ferdinand et al., 2012; Oliveira et al., 2007). These findings contradict the RL-ERN model and support current models that emphasise a valence-independent explanation of the FRN. For instance, these results are consistent with accounts such as the predicted response-outcome (PRO) model (Alexander and Brown, 2011). The PRO model posits that ACC neurons keep track of the history of previous positive and negative reinforcements to specific actions, and formulate predictions about the probabilities of future outcomes. When a predicted outcome does not occur — a surprise — then ACC activity is maximal. Importantly, according to the PRO model, surprising outcomes are processed by ACC neurons in response to both reward and non-reward. This leads to the prediction that the FRN should index prediction errors irrespective of the sign of the error — as observed by Ferdinand et al. (2012) and Oliveira et al. (2007). Similarly, Talmi et al. (2013) recently suggested that the FRN codes salience prediction errors irrespective of outcome valence.

Beyond the original RL-ERN and valence-independent accounts, an interesting development in FRN research comes from the growing evidence showing that the FRN seems to be driven mainly by sensitivity to positive outcomes. A recent review of the literature (Walsh and Anderson, 2012) reported that a number of FRN studies tend to show that the negativity of the component is attenuated for outcomes that are better than expected. This effect results in a greater positivity for +RPEs, whereas the FRN waveform related to negative outcomes often remains a clear negative peak that varies little (or not at all) as a function of -RPE. This contrast between a varying FRN positivity to +RPEs and a relatively stable FRN to -RPEs would, in many cases, be sufficient to account for the classic FRN component. Walsh and Anderson (2012) noted that this predominance of FRN positivity is present in a majority of FRN studies, whereas experiments reporting an increased negativity for -RPEs are less frequent. This trend in the literature has led to a re-formulation of the RL-ERN model by its original authors, who more recently proposed that the FRN observed on the scalp is the product of two distinct outcome-monitoring processes (Holroyd et al., 2008). Specifically, the revised account holds that both phasic increases and decreases modulate the FRN. A negative deflection (N2) is produced by low probability (i.e. unexpected) task relevant events, irrespective of valence. However, unexpected rewards also produce a positive deflection induced by a phasic increase in dopaminergic activity — referred to as the "reward positivity" (Baker and Holroyd, 2011; Holroyd and Yeung, 2012; Holroyd et al., 2011). This increase in dopamine firing rate inhibits ACC neurons, thus causing a reduction in the N2-like negativity typical of the FRN. This model fits with the majority of the data reviewed by Walsh and Anderson (2012), and has received further recent support (Holroyd and Yeung, 2012).

Nevertheless, data exist that do not seem to be explained by the updated RL-ERN model (we refer to this as reward positivity [RP] model from hereon in). Apart from evidence supporting valence-independent accounts, there are also studies reporting a more positive FRN amplitude when positive outcomes are expected rather than unexpected. For instance, in a gambling task, San Martín et al. (2010) found that the FRN was more positive for "win" outcomes when the probability of rewards was higher, compared to win outcomes in a low reward probability context. Similarly, Mushtaq et al. (2013), in a different decision task, found greater FRN positivity for "win" outcomes in a context of positive compared to negative expectations. In related findings, Yu and Zhang (2014) did not find a greater FRN positivity for rewards compared to non-rewards when losses were more likely- contradicting a key prediction of the reward positivity framework. In addition, they found a more positive FRN for reward (compared to non-reward) outcomes in the context of positive expectations.

These studies point towards a fourth account of the FRN. It is possible that, in the studies mentioned above, a positive context (e.g. a "gain" domain, or reward expectation) could have primed a positive affect system. In other words, a positive context may have pre-activated affective systems sensitive to appetitive stimuli, which in turn became more sensitive to the delivery of reward feedback. This possibility is consistent with results reported in the literature on affective priming (Fazio et al., 1986; Hermans et al., 2001, 2003; Musch and Klauer, 2003). This explanation implies that FRN positivity can reflect positive affect over and above prediction error computations. In line with this interpretation, it has previously been suggested that the FRN is sensitive to emotional variables (Hajcak and Foti, 2008; Yeung and Sanfey, 2004), and substantial evidence exists demonstrating a relationship between E/FRN amplitude and: (i) affective ratings (Holroyd et al., 2006; Moser and Simons, 2009; Yeung and Sanfey, 2004); (ii) affective traits in healthy participants (Hajcak and Simons, 2002; Hajcak et al., 2003; Luu et al., 2000; Wiswede et al., 2009); and (iii) affective traits in clinical populations (Gehring et al., 2000; Ruchsow et al., 2004; Weinberg et al., 2010).

In summary, four main theoretical models can be derived from the existing literature on the FRN component. First, the original RL-ERN — which suggests that the FRN is a signal of -RPE. Second, valence-independent accounts such as the PRO model, which posit that the FRN is an index of prediction error regardless of outcome valence. The third and fourth models are driven by data showing that the FRN is preferentially modulated by positive rather than negative outcomes in FRN effects. The reward positivity model (or updated RL-ERN) suggests that FRN positivity increases reflect a +RPE signal, whilst a positive affective priming account predicts that the FRN response to rewards should be more positive when positive affective systems have been previously activated.

The goal of the present study was to evaluate whether any of these models could explain the specific case of short-term effects on the FRN (stFRN), defined as the sensitivity of the component to information presented very briefly prior to the decision outcome time-locked to the FRN. In order to address this, we asked participants to perform a multi-trial gambling task where the outcome could be either monetary wins or losses relative to an initial endowment. We separated FRN activity for current trials (Outcome_t) according to their valence (Win vs. Loss) and according to the valence of the immediately preceding two outcomes (Outcome_{t-1} and Outcome_{t-2}). Crucially, the sequence of gains and losses was fully randomised in such a way that participants could