The dynamics of frustration

Bowen J Fung (bjfung@caltech.edu)
Dean Mobbs (dmobbs@caltech.edu)
Colin F Camerer (ccamerer@caltech.edu)
Division of the Humanities and Social Sciences, California Institute of Technology
Pasadena, California 91101, USA

Abstract

Frustration is a widely experienced emotional state that has been linked to a wide range of societal and individual issues. Early research developed a number of frameworks that exemplified frustration and frustrative non-reward in classical conditioning. However, contemporary approaches to reward learning - routinely used to precisely estimate psychological variables - have not been applied to frustration. These computational methods could help to identify under which conditions frustration arises, and whether the consequences of frustration play an adaptive role in behavior. Here we present an experimental approach that aims to measure and model the effect of frustrative non-reward on motor vigor, and to assess how frustration is represented and implemented in the brain.

Keywords: frustration; frustrative non-reward; reinforcement learning; Pavlovian instrumental transfer; fMRI

Introduction

Frustrative non-reward is a construct of the NIMH Research Domain Criteria (RDoC), here defined as "reactions elicited in response to withdrawal/prevention of reward, i.e., by the inability to obtain positive rewards following repeated or sustained efforts". As an emotion or mood, frustration is likely to play a critical role in aggression (Munyo & Rossi, 2013), domestic violence (Card & Dahl, 2011), gambling (Banks, Tata, Bennett, Sekuler, & Gruber, 2017), and education (Wilde, 2012), and thus is of high relevance to current societal problems.

Psychological research into frustrative non-reward has a long history, with early classical conditioning research characterizing a "frustration effect" (or "extinction burst") whereby the discontinuation of expected rewards resulted in a seemingly paradoxical subsequent increase in vigor (Amsel, 1958). While these initial studies relied on evidence in rodents, the canonical effect is also present in humans. For example, a human study by Otis and Ley (1993) investigated the frustration effect by establishing an instrumental contingency between the force of a lever press and monetary reinforcement. When this contingency was discontinued (i.e. extinction), the force of subjects' responses to an irrelevant "reset" lever press was significantly increased after the first non-rewarded press. This was accompanied by an increase in skin conductance, pointing to physiological correlates of the effect. Other recent research has turned to investigating the brain regions involved in frustration (Yu, Mobbs, Seymour, Rowe, & Calder, 2014). When action-reward contingencies were "blocked" in a multi-trial reward schedule, Yu et al. observed increased BOLD activity in amygdala, periaqueductal grey (PAG), insula, and dorsal anterior cingulate cortex (dACC), and prefrontal cortex (vmPFC) - areas previously associated with reactive aggression (Mobbs et al., 2007). This activity, as well as self-reported frustration, was modulated as a function of proximity and effort expenditure to the reward.

While these recent investigations have shed some light on the physiological mechanisms associated with frustration, it remains stubbornly resilient to a precise computational definition. Given that it is widely appreciated as an important phenomenon, this paper suggests a basic experimental approach that could help specify and ultimately understand the algorithmic and cognitive mechanisms associated with frustration.

Research questions and hypotheses

The primary research agenda outlined here aims to show, at least initially, that frustration, and its consequences, can be modeled as a psychological parameter using a basic reinforcement learning framework. We adopt the simple idea that frustration is the difference, if positive, between the amount of reward initially expected, and the amount received. We also assume that the increased motor vigor observed after a frustrating event (i.e. the frustration effect) can be taken as an implicit index of subjective frustration. It is also of interest to establish whether such an increase in motor vigor is able to impact instrumental performance - that is, whether frustration could be adaptive of maladaptive within the constraints of our paradigm.

Our hypotheses are straightforward. If frustration is a function of expected and received reward, then motor vigor following non-reward will be positively correlated with the initial reward expectation. Over multiple trials, non-reward after recent histories of high average rewarded will increase motor vigor more than a non-rewarded trial after recent histories of lower rewarded trials. For a task that relies on accurate motor responses, this will lead to a decrease in performance immediately subsequent to frustrating events.

Given previous imaging studies of frustration (Yu et al., 2014), and the canonical functions of brain regions identified in these studies, we hypothesize that specific regions should be associated with specific components of this simple model. For example, vmPFC and dACC will encode the expectation of reward (Stalnaker, Liu, Takahashi, & Schoenbaum, 2018), whereas amygdala and PAG will encode the prediction error that is correlated with the frustration effect (Tye, Cone,

Schairer, & Janak, 2010).

Experimental methods

To capture our proposed index of frustration, we will ask participants to complete a series of trials in which they must apply a target amount of force to a dynamometer in order to earn monetary reward (see Figure 1). Participants will also have to provide a non-contingent "reset" response between each trial which will give us an incentive-irrelevant force measure (Otis & Ley, 1993; Yu et al., 2014). This incentive-irrelevant measure will constitute our index of the frustration effect. By artificially blocking participants success on specific trials, we will investigate the dynamics and motor consequences of frustration on task performance and generalized motor output as a function of reward history. We will then test whether representations of the behavioral measures are observable in brain activity.

Participants

Power analysis of the effect sizes reported in (Otis & Ley, 1993)(which used a similar paradigm) varied from 0.04 to 1.38 (Cohens *d*) depending on the delay from the frustrating event (5 or 20 seconds). For an analysis with 80% power, detection of these effects at the 0.05 alpha level would require either 6 or 4042 individuals. Our paradigm will constrain the delays after a frustrating event to within 5 seconds, and will use an analysis that exploits the information in all trials (Otis and Ley used only a single trial per participant per condition). With this in mind, and given that the frustration effect is a reasonably robust phenomenom (Ditkoff & Ley, 1974; Dudley & Papini, 1997; Scull, 1973; Yu et al., 2014), we consider 40 individuals (by 128 trials) to be an ample sample size to elicit an observable frustration effect.

Paradigm

A paradigm schematic can be seen in Figure 1. Participants will first undergo a procedure to calibrate their grip strength, and to determine experimental parameters that will allow them to have high ($<$90%) performance. This means that trials that are not artificially "blocked" will generally be successfully completed. The sequence of predetermined win and "blocked" trials will be counterbalanced using a pseudo-random, high-order deBruijn sequence (Aguirre, Mattar, & Magis-Weinberg, 2011). This will maximize the variability in reward expectation over time with local sequences of both high and low reward, while maintaining an equal number of instances in which rewards follow non-rewards and vice versa.

Analysis

We will apply the model of Daly and Daly (1982) to generate (see Table 1) predictions for our index of frustration. In this model, three associative values are updated by a delta rule on each trial. The first value is the same as in the standard Rescorla-Wagner model, and represents the associative value of the stimulus (V_{AP}). The second is updated only on non-reward trials, and represents an "active" aversive association (V_{AV})[1]. It is this value that reflects primary frustration, and we hypothesize that this value, or its update size, will be correlated with response force.

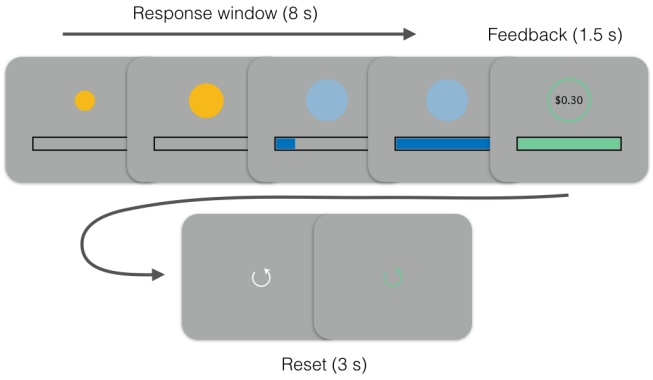

Figure 1: Paradigm schematic. In each trial of the task, the size of a circle stimulus will provide feedback about the current grip force. Participants will have to maintain a target grip force (the force at which the circle turns blue) for a target duration (3 s). The blue bar represents how long the target force has been maintained for. If participants fill this bar within the response window (8 s), and be presented with positive feedback and a monetary reward. In some predetermined trials ("blocked" trials), the circle stimulus will drift away from the target range, and the participant will be unable to successfully complete the trial.

Note that in accordance with our hypothesis, the observed frustration effects ought to vary depending on the previous reward history. Phasic "one-off" non-reward will elicit increases in motor vigor, whereas longer sequences of non-reward will diminish reward expectation, and therefore not elicit increases in motor vigor. The predetermined sequence of wins/blocked trials is designed to maximize our ability to detect these effects.

Importantly, provided that participants are able to successfully complete trials which are not "blocked", the predictions of this model are entirely independent from participants behavioral data. This avoids any potential circularity between behavior and model parameters, and in our perception constitutes a very strong test of the model. This also means we can present the model predictions for our predetermined win/blocked trial sequence (which repeats 3 times) in Figure 2.

Imaging

After standard preprocessing (fmriprep), we will run a whole-brain GLM on the fMRI data, including standard nuisance re-

[1] The associative (V_{AP}) and avoidance (V_{AV}) value components of the model have a theoretical correspondence with the *Go* (D1 striatonigral) and *NoGo* (D2 striatopallidal) striatal dopamine pathways.

Table 1: Update functions from Daly & Daly, 1982.

	Approach (V_{AP})	Avoidance (V_{AV})	Counterconditioning (V_{CC})
Reward	$\Delta V_{AP} = \alpha\beta_I(\lambda - \bar{V}_{AP})$	$\Delta V_{AV} = \alpha\beta_D(0 - \bar{V}_{AV})$	$\Delta V_{CC} = \alpha\beta_I[(0 - \bar{V}_{AV}) - \bar{V}_{CC}]$
No reward	$\Delta V_{AP} = \alpha\beta_D(0 - \bar{V}_{AP})$	$\Delta V_{AV} = \alpha\beta_I[2(0 - \bar{V}_{AP}) - \bar{V}_{AV}]$	$\Delta V_{CC} = \alpha\beta_D(0 - \bar{V}_{CC})$

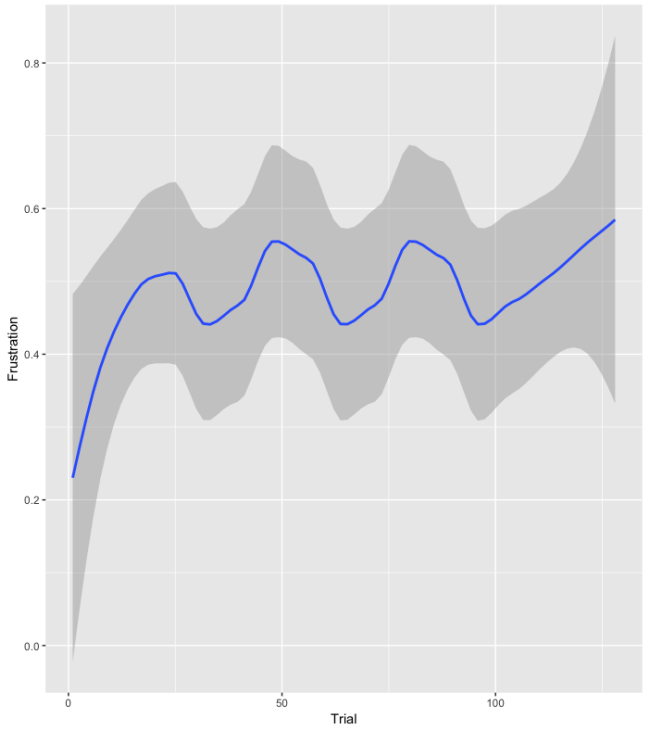

Figure 2: Model predictions from Daly & Daly, 1982. On the y-axis, "frustration" indicates the associated avoidance value of the trial. Note that the predicted values for the "frustration" parameter have periodic high and low values. Our hypothesis is that this quantity will be positively correlated with the force of motor responses throughout the experiment.

gressors, as well as the parameters from the computational model. We will constrain our statistical analysis to independent anatomical ROIs, and correct for multiple comparisons using small volume correction (SVC). These will be consistent with those used in the Anatomical Automatic Labelling Atlas from the WFU Pickatalas package. We will also perform a multivariate pattern analysis, by training an SVM classifier on brain activity after feedback and (separately) during trial reset, in order to see whether there is information that can distinguish normal trials from frustration trials (and error trials). Similarly, we will train a classifier on brain activity after feedback in order to predict the strength of motor response during trial reset. The remainder of the whole brain statistical maps will be used only for exploratory analysis.

Acknowledgments

This work is supported by a Conte Center grant to DM and CFC.

References

Aguirre, G. K., Mattar, M. G., & Magis-Weinberg, L. (2011, June). de Bruijn cycles for neural decoding. *NeuroImage*, *56*(3), 1293–1300.

Amsel, A. (1958, March). The role of frustrative nonreward in noncontinuous reward situations. *Psychological bulletin*, *55*(2), 102–119.

Banks, P. J., Tata, M. S., Bennett, P. J., Sekuler, A. B., & Gruber, A. J. (2017, July). Implicit Valuation of the Near-Miss is Dependent on Outcome Context. *Journal of Gambling Studies*, 1–17.

Card, D., & Dahl, G. B. (2011). Family violence and football: the effect of unexpected emotional cues on violent behavior. *The quarterly journal of economics*, *126*(1), 103–143.

Daly, H. B., & Daly, J. T. (1982). A mathematical model of reward and aversive nonreward: Its application in over 30 appetitive learning situations. *Journal of Experimental Psychology: General*, *111*(4), 441–480.

Ditkoff, G., & Ley, R. (1974). Effects of positive and negative force-contingent reinforcement on the frustration effect in humans. *Journal of Experimental Psychology*, *102*(5), 818–823.

Dudley, R. T., & Papini, M. R. (1997). Amsel's Frustration Effect: A Pavlovian Replication With Control for Frequency and Distribution of Rewards. *Physiology & Behavior*, *61*(4), 627–629.

Mobbs, D., Petrovic, P., Marchant, J. L., Hassabis, D., Weiskopf, N., Seymour, B., ... Frith, C. D. (2007). When fear is near: threat imminence elicits prefrontal-periaqueductal gray shifts in humans. *Science*, *317*(5841), 1079–1083.

Munyo, I., & Rossi, M. A. (2013, May). Frustration, euphoria, and violent crime. *Journal of Economic Behavior & Organization*, *89*, 136–142.

Otis, J., & Ley, R. (1993, November). The effects of frustration induced by discontinuation of reinforcement on force of re-

sponse and magnitude of the skin conductance response. *Bulletin of the Psychonomic Society*, *31*(2), 97–100.

Scull, J. W. (1973, June). The Amsel frustration effect: interpretations and research. *Psychological bulletin*, *79*(6), 352–361.

Stalnaker, T. A., Liu, T.-L., Takahashi, Y. K., & Schoenbaum, G. (2018, February). Orbitofrontal neurons signal reward predictions, not reward prediction errors. *Neurobiology of Learning and Memory*, 1–0.

Tye, K. M., Cone, J. J., Schairer, W. W., & Janak, P. H. (2010, January). Amygdala Neural Encoding of the Absence of Reward during Extinction. *Journal of Neuroscience*, *30*(1), 116–125.

Wilde, J. (2012, June). The Relationship between Frustration Intolerance and Academic Achievement in College. *International Journal of Higher Education*, *1*(2), 1–8.

Yu, R., Mobbs, D., Seymour, B., Rowe, J. B., & Calder, A. J. (2014, May). The neural signature of escalating frustration in humans. *Cortex*, *54*, 165–178.