Title

Incidental discrete emotions influence processes of evidence accumulation in reinforcement-learning

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

ISSN

1069-7977

Authors

LaFollette, Kyle
Kotynski, Anne Elizabeth
Merner, Amanda
et al.

Publication Date

2021

Peer reviewed

Exposure to incidental discrete emotions influence processes of evidence accumulation in reinforcement-learning

Kyle J. LaFollette (kjl113@case.edu), Anne E. Kotynski (aek79@case.edu), Amanda R. Merner (arm251@case.edu), Rock Lim (rxl471@case.edu), Huangqi Jiang (hxj244@case.edu), & Heath A. Demaree (had4@case.edu)

Department of Psychological Sciences, Case Western Reserve University, Cleveland, OH 44106 USA

Abstract

Discrete emotions are known to elicit changes in decision-making. Previous research has found that affect biases response times and the perception of evidence for choices, among other key factors of decision-making. However, little is known how affect influences the specific cognitive mechanisms that underlie decision-making. We investigated these mechanisms by fitting a hierarchical reinforcement-learning decision diffusion model to participant choice data. Following the collection of baseline decision-making data, participants took part in a writing exercise to generate neutral or discrete emotions. Following the writing exercise, participants made additional decisions. We found that exposure to discrete emotions modulates decision-making through several mechanisms including rates of learning and evidence accumulation, separation of decision thresholds, and sensitivity to noise. Furthermore, we found that exposure to each of the four discrete emotions modulated decision-making differently. These findings integrate learning and decision process models to expand on previous research and elucidate processes of affective decision-making.

Keywords: emotion; decision-making; reinforcement-learning; drift diffusion

Introduction

Emotions can direct us to make decisions to approach or withdraw from other people or situations. For example, fear can narrow our attention to thinking about possible escapes or ways to avoid punishment, whereas desire can drive us toward objects of our appetitive urges. Considerable extant literature exists supporting the “affect-as-information” hypothesis, which suggests that affect itself can support a particular action (much like expected utility and cost; Greifeneder et al., 2010; Schwarz, 2011). Emotion can also influence the speed at which actions are taken, such as in the case where negative valence strengthens post-error slowing (Inzlicht et al., 2015). However, different discrete emotions can have greatly disparate effects. For example, anger, while a negative emotion, can lead to approach behaviors (Carver & Harmon-Jones, 2009). This highlights a critical gap in the literature, namely that the influence of different discrete emotions on specific cognitive mechanisms underlying decision-making remains unclear.

Affect-as-information theory suggests that valence, intensity, and the weight of affect in evidence accumulation processes all contribute to decision-making (Hartley & Sokol-Hessner, 2018). It is possible that each of these dimensions are independently weighted in the decision-making processes, giving rise to the observed behavioral effects of different discrete emotions. However, little work has been done to test this hypothesis.

Rather, the variability in decision-making has traditionally been explained by the confluence of choice behavior, response time, and the integration of feedback by the decision-maker. The former two indices can be modeled as a process of evidence accumulation or drift (Ratcliff, 1978). Decision diffusion models (DDM) explain how choices and latencies arise from latent mechanisms of information processing, such as drift rate, response biases, and decision thresholds. The latter index can be modeled with reinforcement-learning (see O’Doherty et al., 2017), which typically formalize sequential decision-making as learning processes influenced by reward prediction errors and noise. Together, these two approaches to modeling decision-making can identify which latent processes are affected by specific discrete emotions.

Indeed, some prior work suggests that certain discrete emotions modulate decision diffusion processes, such as emotions induced with fear-related words (Mueller & Kuchinke, 2016) and aversive images (Warren et al., 2020). That said, these pursuits are sparse in the literature and it is unclear how different discrete emotions comparatively influence these processes. Furthermore, to the best of our knowledge no study to date has incorporated reinforcement-learning attributes with DDM to investigate their combined susceptibility to emotion. Considering these information gaps, we aimed to develop a hierarchical reinforcement-learning DDM with applications for explaining decision-making influenced by four incidental discrete emotions – happiness, sadness, anger, and desire – as compared to a neutral emotional state. We fit this model to data from the following experiment to explore relationships between specific discrete emotions, information processing, and learning. Inferences taken from this model are exploratory and reported for completeness and to inspire future modeling efforts. Replication is needed to validate the findings as reported. Future modeling work can consider these findings when defining prior expectations in emotion-related parameter weights.

Method

Participants

Thirty-one individuals (15 women; mean age = 19.61, SD = 1.02) participated in this experiment. Participants were recruited from a private Midwestern university and received partial course credit for their participation. All participants provided written informed consent in accordance with the university's institutional review board.

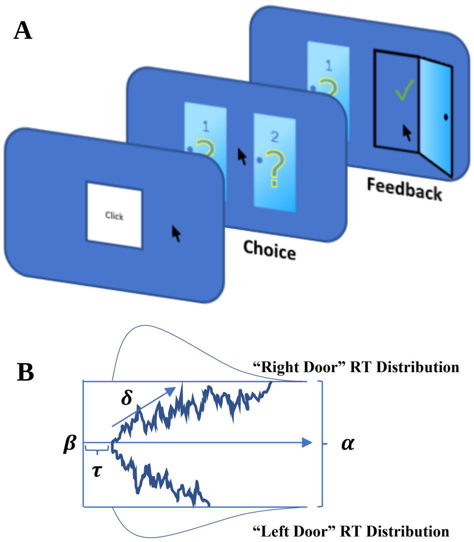

Figure 1: Task and decision diffusion model. (A) The 2AFC task. Participants were presented with images of two doors representing different reward probability distributions, and were asked to maximize their reward. (B) The decision diffusion model. In this model, decisions are made by continuously accumulating noisy evidence until a threshold is reached. Participants respond according to whether the threshold was met for either the upper or lower bound.

2AFC Task

All participants completed a novel 2-alternative forced choice task (2AFC; Figure 1a.). At the start of each trial, participants were presented with a centered white "click" box and were instructed to move their mouse cursor to and click on this box. Afterwards, the cursor was re-centered and the box was replaced with images of two doors. Participants were instructed to move their cursor and click on either of these two doors with the goal of receiving a reward. Participants were also made aware that each door represented a probability distribution of reward; one had a 75% chance of success (the exploitation door) whereas the other had only a 25% chance (the exploration door). Seven blocks of ten trials each were completed for a total of 70 trials.

Discrete Emotion Manipulation

The 2AFC task was completed in five between-subjects conditions differing in manipulated discrete emotion: Happiness, sadness, anger, desire, and neutral (control). After participants had completed the first five blocks, they were asked to reflect on and write about a real-life experience that made them feel the emotion congruent with their assigned condition. For example, a participant in the angry/sad/happy condition was asked to, "In detail, please write about the one situation that has made you the most angry/sad/happy you have been in your life, and describe it such that a person reading the description would become angry/sad/happy just from reading about the situation." Participants in the desire condition were asked to, "In detail, please write about the one situation that you most desire right now, and describe it such that a person reading the description would have great desire just from reading about the situation. Examples might include acing a major test, kissing a person you're attracted to, etc.". Participants in the neutral condition were instead asked to write about their daily routine and "describe it such that a person reading the description would have a clear understanding of your daily routine". After this manipulation, all participants continued with their remaining two blocks of the 2AFC task.

Participants also completed a computerized Discrete Emotions Questionnaire (DEQ; Harmon-Jones et al., 2016) both immediately prior to and after experimentation. The DEQ was used to measure four distinct state emotions: Happiness, sadness, anger, and desire. DEQ state emotion scores were considered as manipulation checks for change in subjective emotional experience.

Hierarchical Drift Diffusion Model

2AFC task performance was modeled by fitting a hierarchical DDM to participant choice and response time data. The DDM formalizes evidence accumulation by calculating the likelihood of response time for choice *x* with the Wiener first-passage time distribution:

$$RT(x) \sim WFPT[\alpha, \tau, z, \delta]$$

The DDM decomposes response time entirely into four parameters: boundary separation, non-decision time, initial bias z, and drift rate (Figure 1b). For this 2AFC task, it is assumed that participants noisily accumulate evidence at a drift rate for either door, and select a door once the evidence reaches its corresponding threshold. Thresholds are distanced by a boundary separation, with larger separations demanding the accrual of more evidence and reflecting an emphasis on accuracy over speed. Initial bias represents that starting point of the diffusion process, and non-decision time encompasses time spent on extraneous processes such as stimulus encoding and motor planning. We fit the model to participant data to estimate the joint probability of participant choices and response times from the values of these four parameters.

Reinforcement-Learning Parameters To account for the influence of learning on this task, we modified the DDM to update trial-by-trial expected values of reward using a delta learning rule (Rescorla & Wagner, 1972). Each door was associated with an expected value, calculated as the product of its previous value plus the product of learning-rate η and a reward-prediction error:

$$Q_{o,i} = Q_{o,i-1} + \eta(Reward_{o,i-1} - Q_{o,i-1})$$

Additionally, the choice rule for selecting reinforced options was formalized as a softmax logistic function (Luce, 1959), where the sensitivity β scales the probability of choosing option o as a function of how much larger its expected value is compared to other options n:

$$p_{o,i} = \frac{e^{(\beta x Q_{o,i})}}{\sum_{j=1}^{n} e^{(\beta x Q_{j,i})}}$$

Change in sensitivity is associated with the exploration-exploitation dilemma (Daw et al., 2006). As $\beta \rightarrow \infty$, choice becomes deterministic or exploitative. β=0 is purely stochastic or explorative. Together, these two processes capture learning in the model and introduce two additional free parameters.

To combine these reinforcement-learning and DDM processes, we further decomposed the drift rate parameter of the DDM to be the scaled difference between the expected value of reinforced options (Frank et al., 2015):

$$\delta_{i,\,t} = \, v_i(EV_U - EV_L)$$

Where i, t is the drift rate for participant i at trial t, and v is a free parameter describing the rate at which the chooser accumulates evidence for the expected value difference between the upper and lower bounds.

Model Specifications To account for individual differences in parameter estimates, the model was fit to individual participants' data using hierarchical Bayesian analysis and were partially pooled such that subject-level parameters were drawn from common normal group distributions. Posteriors were inferred with the Hamiltonian Monte Carlo No-U-Turn sampler, which is a specific Markov Chain Monte Carlo sampler available in the Stan probabilistic programming language (Carpenter et al., 2017). We collected 60,000 samples for each parameter across 6 chains run in parallel. The first 30,000 samples of each chain were discarded as warm-up. Chain convergence was diagnosed with traceplots and the Gelman-Rubin convergence diagnostic (Gelman & Rubin, 1992).

Priors for the group-level normal means and standard deviations were weakly informative and set to standard normal (normal (0, 1)) and half-Cauchy (half-Cauchy (0, 5)), respectively. We re-parameterized the model to be non-centered ("Matt trick") to optimize sampling and reduce autocorrelation between group-level parameters (Betancourt & Girolami, 2013).

Parameters were permitted to either vary freely between experimental conditions or be fixed across conditions (Vandekerckhove & Tuerlinckx, 2007). We fixed initial bias to a non-biased value of 0.5 and assumed that non-decision time would not vary with emotional state. As such, we considered differences in the posterior distributions of four free parameters (boundary separation, drift rate, learning rate, and sensitivity) across four condition contrasts (Happy-Neutral, Sad-Neutral, Angry-Neutral, and Desire-Neutral).

Data preparation scripts and model code are available at https://github.com/kjlafoll/emotionddm.

Equivalence Testing

After successful parameter estimation, we used Bayesian equivalence testing to determine whether a null value was among the credible values of the posterior distribution differences for the four contrasts. This is analogous to frequentist equivalence testing, but with highest density intervals instead of confidence intervals. For each parameter delta (difference in estimates between conditions) we established the 95% highest density interval (HDI) representing the 95% most credible difference values. We further established a region of practical equivalence (ROPE) around the null difference value, specified as half of Cohen's conventional small effect size (Cohen, 1988), the range of -0.1 to 0.1, scaled by the standard deviation of the dependent observation variable, as suggested by Kruschke (2018). Using a HDI+ROPE decision rule, we accepted that a parameter difference was equivalent to null if the 95% HDI fell completely inside the ROPE, and we rejected equivalence to the null if it fell completely outside the ROPE. All other cases where the 95% HDI was partially within the ROPE were considered weakly informative and left to interpretation.

Results

DEQ Manipulation Checks. A one-way ANOVA for DEQ happiness change scores revealed a significant change in self-reported happiness; $F(4, 26) = 3.13$, $p = 0.03$; Happy M = 2, Sad M = -0.625, Angry M = -2.6, Desire M = -0.37, Neutral M = 0. A one-way ANOVA for DEQ anger change scores also revealed a significant change in self-reported anger; $F(4, 26) = 6.07$, $p < 0.001$; Happy M = -3.33, Sad M = 0.125, Angry M = 3.2, Desire M = 1.125, Neutral M = -0.428. These results suggest that our task was successful in manipulating conscious subjective experiences of happiness and anger.

One-way ANOVAs for DEQ sadness and desire change scores failed to yield significant differences in sadness nor desire. DEQ sadness was not found to significantly change post manipulation; $F(4, 26) = 2.02$, $p = 0.12$, nor did DEQ desire post-manipulation; $F(4, 26) = 1.45$, $p = 0.24$. These results suggest that our task was not successful at manipulating conscious subjective experience of sadness nor desire.

Boundary Separation We found strong evidence supporting a negative effect of the happiness condition on boundary separation, such that the difference between the Happy and Neutral estimates had a probability greater than 99.9% of being negative (Mean = -0.11, 95% HDI [-0.220, 0.017]; Figure 2). This suggests that the happiness condition reduced the amount of evidence necessary to reach a decision threshold, and that happy individuals placed a greater emphasis on speed over accuracy.

We also found strong evidence for positive effects of the sadness, anger, and desire conditions on boundary separation.

The difference between the Sad and Neutral estimates had a probability greater than 99.9% of being positive (Mean = 0.043, 95% HDI [0.035, 0.053]), as was the difference between the Angry and Neutral estimates (Mean = 0.16, 95% HDI [0.150, 0.170]) and the difference between the Desire and Neutral estimates (Mean = 0.059, 95% HDI [0.049, 0.070]). This suggests that sadness, anger and desire conditions all increase the amount of evidence necessary to reach a decision threshold, shifting the emphasis on accuracy over speed.

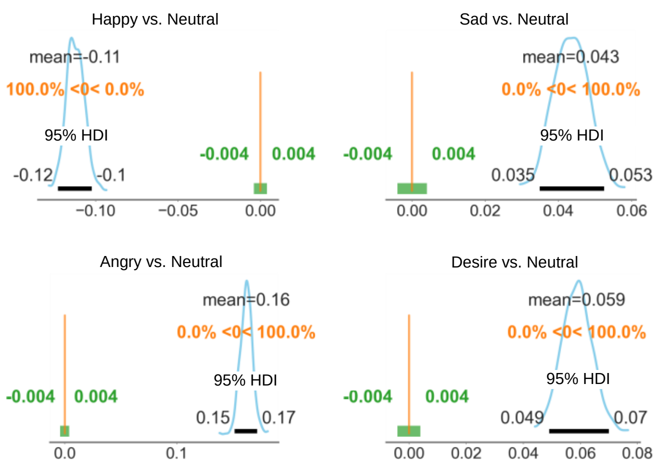

Figure 2: Posterior predictive density plots for effects of condition on boundary separation. Black bar is the 95% highest density interval, indicating the 95% most credible values for the mean difference between conditions. Green bar is the region of practical equivalence.

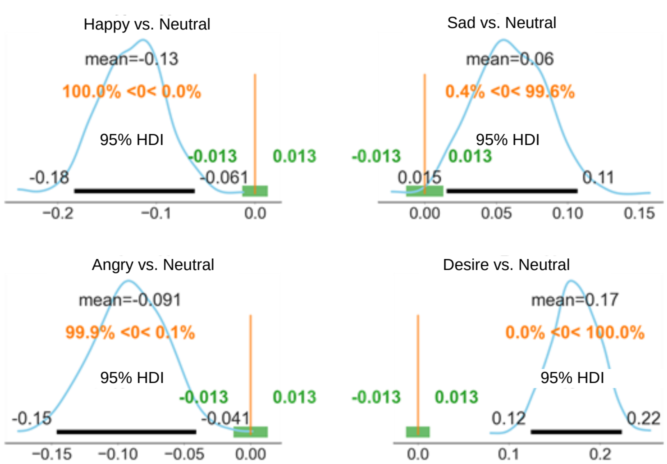

Figure 3: Posterior predictive density plots for effects of condition on the drift rate scaler *v*.

Drift Rate We found strong evidence supporting negative effects of the happiness and anger conditions on drift rate. The difference between the Happy and Neutral estimates had a probability greater than 99.9% of being negative (Mean = -0.13, 95% HDI [-0.180, -0.061]; Figure 3), as did the difference between the Angry and Neutral estimates (Mean = -0.091, 95% HDI [-0.150, -0.041]). This suggests that happiness and anger conditions both slow the rate of evidence accumulation.

We also found strong evidence for positive effects of the sadness and desire conditions on drift rate. The difference between the Sad and Neutral estimates had a probability of 99.6% of being positive (Mean = 0.06, 95% HDI [0.015, 0.110]). The difference between the Desire and Neutral estimates had a probability greater than 99.9% of being positive (Mean = 0.17, 95% HDI [0.120, 0.220]). This suggests that sadness and desire conditions both quicken the rate of evidence accumulation.

Learning Rate We found weak evidence for positive effects of the sadness and desire conditions on learning rate. The difference between the Sad and Neutral estimates had a probability of 96.6% of being positive (Mean = 0.018, 95% HDI [-0.001, 0.036]; Figure 4). The difference between the Desire and Neutral estimates had a probability of 75% of being positive (Mean = 0.0065, 95% HDI [-0.012, 0.024]). The HDI partially fell within the ROPE for both of these difference estimates (ROPE [-0.013, 0.013]). While these results neither support effects of condition nor equivalence, should they continue in the positive direction, it would suggest that sadness and desire both increase the weight of prediction errors in expected value updating. This suggests that sadness and desire conditions quicken learning but also increase susceptibility to larger post-error swings.

We also found weak evidence for equivalence between the happiness, anger and neutral conditions with respect to learning rate. The difference between the Happy and Neutral estimates had a probability of 58.7% of being negative (Mean = -0.003, 95% HDI [-0.031, 0.022]) and the difference between the Angry and Neutral estimates had a probability of 55.2% of being positive (Mean = 0.002, 95% HDI [-0.017, 0.021]). While the HDI only partially fell within the ROPE for both of these estimates, the probabilities of direction were near chance level, suggesting happiness and anger conditions had no effect on learning rate.

Sensitivity We found strong evidence supporting positive effects of the sadness and desire conditions on sensitivity. The difference between the Sad and Neutral estimates had a probability greater than 99.9% of being positive (Mean = 0.63, 95% HDI [0.260, 0.980]; Figure 5) as did the difference between the Desire and Neutral estimates (Mean = 0.77, 95% HDI [0.420, 1.100]). This suggests that sadness and desire conditions both increase how deterministic or exploitative individuals are in their decision-making.

We also found weak evidence for a negative effect of the happiness condition on sensitivity. The difference between the Happy and Neutral estimates had a probability of 95.7% of being negative (Mean = -0.38, 95% HDI [-0.760, 0.026]). Although the HDI did partially fall within the ROPE, this does suggest that the happiness condition served to increase stochasticity or exploration in decision-making (ROPE [-0.027, 0.027]).

Lastly, we found weak evidence for equivalence between the anger and neutral conditions with respect to sensitivity. The difference between the Angry and Neutral estimates had a probability of 69% of being negative (Mean = -0.094, 95% HDI [-0.460, 0.300]), indicating that while the HDI only partially fell within the ROPE, the probability of direction was near chance level. This suggests that the anger condition had no effect on sensitivity.

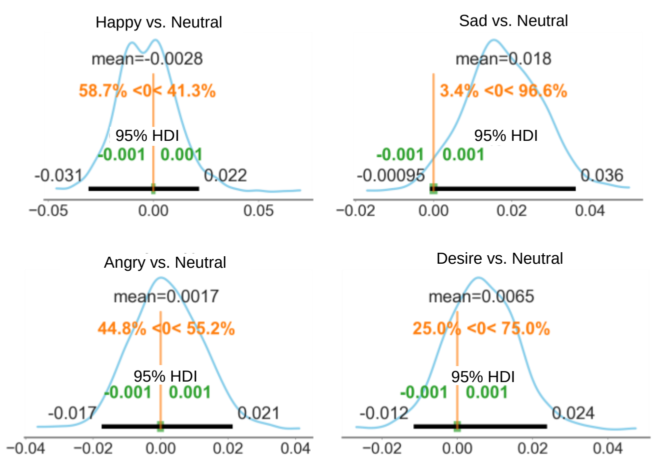

Figure 4: Posterior predictive density plots for effects of condition on learning rate.

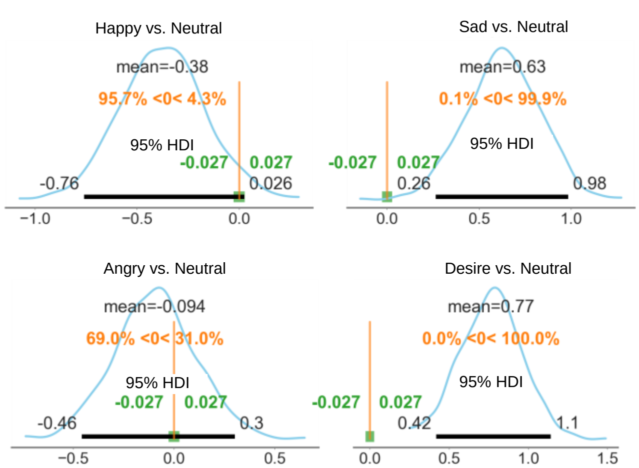

Figure 5: Posterior predictive density plots for effects of condition on sensitivity.

Discussion

To the best of our knowledge, this is the first analysis of reinforcement-learning and decision diffusion processes in the context of manipulated discrete emotion. This analysis demonstrates the utility of computational models in studying the influence of affect on decision-making.

With respect to the efficacy of our manipulation, we were successful in increasing self-reported happiness and anger in the Happy and Anger conditions, respectively. We were not successful in manipulating self-reported sadness and desire. When interpreting these results, it is important to remember that the intent of this work is to present a novel reinforcement learning DDM with applications for studying affective decision-making, and that therefore all parameter inference is de facto exploratory. We report these results for completeness and to guide future replication efforts to ensure that these effects are robust. Future research may consider why these effects are observed if not for differences in subjective emotional experience. For example, although our manipulation of sadness was not sustained, participants may have associated the explicitly recalled memory with implicit mood-congruent memories, biasing behavior while remaining unaware of unconscious emotions.

In the Happy-Neutral contrast we discovered that people in the Happy condition had a greater emphasis on speed over accuracy, a slower rate of evidence accumulation, and a greater propensity to explore. Each of these findings are in accordance with what theory would predict with the exception of the reduced drift rate. Contrary to our findings, Mueller and Kuchinke (2016) discovered that decisions related to words with happy connotations had greater drift rate than those related to neutral and fear-related words. According to Personality Systems Interaction theory (PSI; Kuhl, 2000), we should instead expect that negative affective states diminish the attentional resources that are needed for evidence accumulation, which in fact is supported by empirical work (Smallwood et al., 2009). One possible explanation for our observed smaller drift rate is that discrete emotions on our task were incidental rather than integral. Feelings-as-information theory would suggest that incidental emotions might not influence unrelated target decisions (Schwarz, 2011). Rather, target non-specific mood processes may better explain the difference in observed drift rates. Future studies should further investigate the role of incidental emotion on drift rate.

In the Sad-Neutral contrast we discovered that people in the Sad condition had a greater emphasis on accuracy over speed, faster rates of evidence accumulation, faster learning rates and a greater propensity to exploit. Both the emphasis on accuracy and tendency to exploit could be consequences of efforts to reduce current negative affect and maximize utility (e.g., Tice et al., 2001). Interestingly, sadness was the only discrete emotion associated with a greater learning rate, suggesting that sadness was aversive to such an extent that learning was accelerated at the expense of greater instability when faced with unexpected outcomes. A possible explanation for this is that sadness disproportionally weights positive and negative feedback. Indeed, differences in feedback valence have been related to differences in learning rate (Gershman, 2015) and affect has been found to differentially influence the perception of these outcomes (Eldar & Niv, 2015). To better understand the independent influences of positive and negative feedback on learning, future studies can incorporate separate learning rates for positive and negative prediction errors. This has been accomplished in previous reinforcement-learning DDMs (e.g., Pedersen et al., 2017), but outside the context of discrete emotion where it could be especially important.

In the Angry-Neutral contrast we discovered that people in the Angry condition had a greater emphasis on accuracy over speed and slower rates of evidence accumulation. Curiously, this combines the drive for accuracy that we associated with sadness with the impairment to evidence accumulation that we associated with happiness. With both far-reaching thresholds and slowed diffusion, anger particularly debilitates response time. It is possible that anger becomes less debilitating as separation narrows, but our model cannot evaluate this as we did not incorporate dynamic decision thresholds. Fontanesi and colleagues (2019) accomplished this by formalizing boundary separation as the natural exponential of the sum of a fixed separation ηi and the average expected value of a decision scaled by a threshold modulation parameter ηc:

$$\eta_{i,\,t} = \exp(\eta i_i + \eta c_i \left(\frac{EV_U + EV_L}{2}\right))$$

In the Desire-Neutral contrast we discovered that people in the Desire condition had a greater emphasis on accuracy over speed, faster rates of evidence accumulation, and a greater propensity to exploit. All three of these observations are in accordance with what theory would expect of a positively valenced, high intensity emotion such as desire. Furthermore, these findings support the notion that desire motivates a sense of urgency to approach, resulting in highly focused attention (Gable & Harmon-Jones, 2010).

While the current analysis demonstrates the utility of computational models at the junction of studying discrete emotion and decision-making, our model is not without limitations. Notably, our model did not incorporate effects of discrete emotion at the algorithmic level. Rather, we calculated differences in parameter estimates between conditions by comparing the condition-specific posterior predictive densities. In place of this rather binary approach to appraising the weight of discrete emotion, future studies should consider more directed concept-to-parameter mapping. For example, the weighted integration model of evidence accumulation posits that emotions varying in valence and intensity can be differentially weighted in evidence accumulation (Hartley & Sokol-Hessner, 2018). Specifying such weights directly within a model could provide a more precise measure of emotion's influence of specific decision-making processes.

Our findings are generally consistent with theories of emotion regulation and integration, with some exceptions. Unfortunately, little empirical work exists at the junction of computational modeling and discrete emotions theory to better interpret these exceptions. As a first step, this analysis provides some empirical clarity at this junction and demonstrates the advantages of modeling discrete emotion at the process-level. Furthermore, our findings motivate interesting theoretical questions about dimension-specific effects of emotion on processes of evidence accumulation and encourage future work to consider the advantages of sophisticated modeling for addressing those questions.

References

Betancourt, M., & Girolami, M. (2013). Hamiltonian Monte Carlo for hierarchical models. *arXiv*, 1312.0906.

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, *76*(1), 1-32.

Carver, C. S., & Harmon-Jones, E. (2009). Anger is an approach-related affect: Evidence and implications. *Psychological Bulletin*, *135*(2), 183-204.

Cohen, J. (1988). Statistical power analysis for the behavioral sciences. New York, NY: Routledge Academic.

Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, *441*, 876-879.

Eldar, E., & Niv, Y. (2015). Interaction between emotional state and learning underlies mood instability. *Nature Communications*, *6*, 1-9.

Fontanesi, L., Gluth, S., Spektor, M. S., & Rieskamp, J. (2019). A reinforcement learning diffusion decision model for value-based decisions. *Psychonomic Bulletin & Review*, *26*, 1099-1121.

Frank, M. J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T. V., Cavanagh, J. F., & Badre, D. fMRI and EEG predictors of dynamic decision parameters during human reinforcement learning. *Journal of Neuroscience*, *35*, 485-494.

Gable, P. A., & Harmon-Jones, E. (2010). The motivational dimensional model of affect: Implications for breadth of attention, memory, and cognitive categorization. *Cognition and Emotion*, *24*(2), 322-337.

Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, *7*, 457-511.

Gershman, S. J. (2015). Do learning rates adapt to the distribution of rewards? *Psychonomic Bulletin & Review*, *22*, 1320-1327.

Greifeneder, R., Bless, H., & Tuan Pham, M. (2010). When do people rely on affective and cognitive feelings in judgment? A review. *Personality and Social Psychology Review*, *15*, 107-141.

Harmon-Jones, C., Bastian, B., & Harmon-Jones, E. (2016). The discrete emotions questionnaire: A new tool for measuring state self-reported emotions. *PLOS One*, *11*(8): e0159915.

Hartley, C. A., & Sokol-Hessner, P. (2018). Affect is the foundation of value. In *The Nature of Emotion: Fundamental Questions* (Fox, A. S., *et al*., eds), pp. 348-351, Oxford University Press.

Inzlicht, M., Bartholow, B. D., & Hirsh, J. B. (2015). Emotional foundations of cognitive control. *Trends in Cognitive Sciences*, *19*(3), 126-132.

Kruschke, J. K. (2018). Rejecting or accepting parameter values in Bayesian estimation. *Advances in Methods and Practices in Psychological Science*, *1*(2), 270-280.

Kuhl, J. (2000). The volitional basis of Personality Systems Interaction Theory: Applications in learning and treatment contexts. *International Journal of Educational Research*, *33*, 665-703.

Luce, R. D. (1959). Individual choice behavior: A theoretical analysis. New York, NY: Wiley.

Mueller, C. J., & Kuchinke, L. (2016). Individual differences in emotion word processing: A diffusion model analysis. *Cognitive, Affective, & Behavioral Neuroscience*, *16*(3), 489-501.

Pedersen, M. L., Frank, M. J., & Biele, G. (2017). The drift diffusion model as the choice rule in reinforcement learning. *Psychonomic Bulletin & Review*, *24*(4), 1234-1251.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59-108.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning II: Current research and theory* (Black, A. H., & Prokasy, W. F., eds), pp. 62-99, New York, NY: Appleton-Century-Crofts.

Schwarz, N. (2011). Feelings-as-information theory. In *Handbook of Theories of Social Psychology* (Van Lange, P. A. M., *et al.*, eds), pp. 289-308, SAGE Publications.

Smallwood, J., Fitzgerald, A., Miles, L. K., & Phillips, L. H. (2009). Shifting moods, wandering minds: Negative moods lead the mind to wander. *Emotion*, *9*(2), 271-276.

Tice, D. M., Bratslavsky, E., & Baumeister, R. F. (2001). Emotional distress regulation takes precedence over impulse control: If you feel bad, do it! *Journal of Personality and Social Psychology*, *80*(1), 53-67.

Vandekerckhove, J., & Tuerlinckx, F. (2007). Fitting the Ratcliff diffusion model to experimental data. *Psychonomic Bulletin & Review*, *14*(6), 1011-1026.

Warren, S. L., Zhang, Y., Duberg, K., Mistry, P., Cai, W., Qin, S., Bostan, S.-N., Padmanabhan, A., Carrion, V. G., & Menon, V. (2020). Anxiety and stress alter decision-making dynamics and causal amygdala-dorsolateral prefrontal cortex circuits during emotion regulation in children. *Biological Psychiatry*, *88*(7), 576-586.