Cognitive, Affective, & Behavioral Neuroscience
https://doi.org/10.3758/s13415-023-01147-8

RESEARCH ARTICLE

Affective valence does not reflect progress prediction errors in perceptual decisions

Alan Voodla[1,2] · **Andero Uusberg**[1] · **Kobe Desender**[2]

Accepted: 10 December 2023

Abstract
Affective valence and intensity form the core of our emotional experiences. It has been proposed that affect reflects the prediction error between expected and actual states, such that better/worse-than-expected discrepancies result in positive/negative affect. However, whether the same principle applies to progress prediction errors remains unclear. We empirically and computationally evaluate the hypothesis that affect reflects the difference between expected and actual progress in forming a perceptual decision. We model affect within an evidence accumulation framework where actual progress is mapped onto the drift-rate parameter and expected progress onto an expected drift-rate parameter. Affect is computed as the difference between the expected and actual amount of accumulated evidence. We find that expected and actual progress both influence affect, but in an additive manner that does not align with a prediction error account. Our computational model reproduces both task behavior and affective ratings, suggesting that sequential sampling models provide a promising framework to model progress appraisals. These results show that although affect is sensitive to both expected and actual progress, it does not reflect the computation of a progress prediction error.

Keywords Affect generation · Perceptual decisions · Evidence accumulation

Introduction

Affect is the dimension of experiences that ranges from positive to negative valence, with some intensity (Russell, 2009). According to various cognitive theories of emotions, affective valence arises as the current situation gets appraised with respect to various motives (Moors et al., 2013; Uusberg et al., 2019). One of the key appraisal dimensions involves evaluating progress in obtaining valued outcomes, with more positive/negative affect reflecting higher/lower progress (Carver, 2015; Carver et al., 1996). However, it is currently unclear what role *progress* expectations play in this appraisal and whether and how expected and actual progress jointly impact our affective experiences.

✉ Alan Voodla
alan.voodla@ut.ee

1 Institute of Psychology, University of Tartu, Tartu, Estonia

2 Brain and Cognition, KU Leuven, Leuven, Belgium

Progress appraisals and affect

A core idea expressed in different accounts of affect generation is that affect is a regulatory signal (Carver, 2015; Proust, 2014; Velasco & Loev, 2022). Similar to a signal inside a thermostat regulating the heating in a room, affect functions as a monitoring signal, detecting discrepancies (e.g., conflicts, errors) and indicating the need to invest effort to overcome them. One influential account, Control-Process theory (CPT, Carver, 2015; Carver & Scheier, 1990) further specifies the nature of the monitoring signal and suggests that it reflects an evaluation of the discrepancy between the actual and desired progress in a task (see review of evidence: Carver & Scheier, 2013). For example, if a researcher wants to write 1000 words per day but only manages to write 500 words, this lower-than-desired progress is hypothesized to generate negative affect. In turn, this negative affect functions as a signal to increase effort. On the other hand, if the researcher manages to write 1300 words, this better-than-desired progress is hypothesized to generate positive affect.

Published online: 05 January 2024

Role of expectations in progress appraisals

Empirical evidence suggests that affective experiences are sensitive to expectations and specifically track the violations of these expectations, i.e., prediction errors (Emanuel & Eldar, 2022; Rutledge et al., 2014). For instance, affect has been shown to reflect the difference between expected and actual rewards. In a monetary gambling task, Rutledge et al. (2014) found that participants reported more positive affect when an actual reward exceeded the expected reward (i.e., a positive prediction error) and more negative affect when it was below expectation (i.e., a negative prediction error). This pattern also has been found outside the laboratory. For instance, in an experience sampling study, college students reported more positive/negative affect when their exam results were better/worse than expected (Villano et al., 2020). In addition to reward prediction errors, affect has been shown to reflect perceptual prediction errors, with higher levels of predictability leading to more positive affect (Chetverikov & Kristjánsson, 2016).

As prediction errors have been robustly associated with affect, this computation also may underlie prediction errors about progress in task. In line with this hypothesis, various Predictive Processing accounts of affect generation have proposed that affect is computed as a function of expected and/or actual rate of prediction error minimization (Joffily & Coricelli, 2013; Van de Cruys, 2017; Velasco & Loev, 2020). For example, affect has been proposed to reflect the rate of prediction error minimization (Joffily & Coricelli, 2013), the expected rate of prediction error minimization (Velasco & Loev, 2020) or the difference between expected and actual rate of prediction error minimization (Van de Cruys, 2017). These accounts suggest that before engaging in a task, people have expectations of progress that are compared with the actual progress. If the actual progress is better or worse than expected, it generates positive or negative affect, respectively.

Toward a computational account of progress prediction errors

Although affect has been shown to be sensitive to reward prediction errors (Emanuel & Eldar, 2022; Rutledge et al., 2014), to our knowledge no empirical studies have investigated whether and how affective experiences are sensitive to progress prediction errors. One reason for this lack of empirical work is that it is difficult to operationalize expected and actual progress. In the current work, we do so by building a perceptual-decision-making paradigm that enables computationally operationalizing expected and actual progress in a decision. This is useful, because in perceptual decisions, progress (of moving toward a correct answer) happens very fast within a trial, and previous work has established means to systematically manipulate this progress.

To understand how perceptual decisions can have high or low progress, and what it means to expect high or low progress in such decisions, we leverage the theoretical notion of evidence accumulation. Evidence accumulation refers to the process of gathering noisy evidence with respect to possible decision options, usually until a decision boundary is reached and a decision is made. This idea has been formalized in various sequential sampling models, of which the Drift Diffusion Model (DDM), a specific instance of the larger family of accumulation-to-bound models, has received much empirical support (Gold & Shadlen, 2007; Ratcliff et al., 2016). In the DDM, the evidence accumulation process is described by a small set of latent parameters that explain variability in decision accuracy and reaction times. The most important parameter for the current purpose is the drift rate, which characterizes the average rate of evidence accumulation. When there is a clear signal (e.g., high signal-to-noise ratio), the drift rate will be high and evidence will accumulate quickly toward the correct boundary (i.e., leading to fast and correct decisions). When the signal is unclear (e.g., low signal-to-noise ratio), the drift rate will be low and evidence will accumulate slowly (i.e., slow and error-prone decisions).

To model affect generation, we will use the drift rate of the DDM as a measure of actual progress. Because the DDM does not have expectations about the rate of progress (i.e., different from, e.g., expectations about the most likely response; Mulder et al., 2012), we assume a parallel process of expected evidence accumulation (controlled by an expected drift rate parameter), which reflects expected progress (for a similar approach in the field of decision confidence, see Drugowitsch & Pouget, 2012; Fleming & Daw, 2017; Khalvati et al., 2021; Van Marcke et al., 2022). From a theoretical perspective, one might expect that affect reflects the prediction error between the expected and actual amount of evidence in a decision. To formalize this insight, we implemented this model in a computational framework, for convenience referred to as affectDDM, which enables simultaneous estimation of expected and actual drift rate of evidence accumulation in a decision.

Current study

Does affect reflect progress prediction errors? To answer this empirical question, we systematically varied expected and actual progress in a perceptual decision-making paradigm and measured their influence on affect. Moreover, we extend the DDM framework with an expected drift rate to evaluate whether affect also reflects progress prediction errors on a computational level. In addition to explaining choices and reaction times, affectDDM makes predictions about trial-by-trial changes in affect as a function of expected and actual progress in the evidence accumulation process. As a consequence, by fitting the affectDDM to the empirical data, we can unravel the computational mechanisms of affect generation in perceptual decisions.

Methods

Participants

The study was conducted online on the *Cognition.run* platform, and participants were recruited using the Sona systems platform of the University of KU Leuven. Ethical approval was granted by the local Ethics Committee. Participants were recruited from the KU Leuven 1st year Bachelor Psychology students pool and received credits in return for participation. Expecting data loss and following Brysbaert & Stevens (2018) suggestion that we needed at least 1600 observations per analysis cell to obtain 80% statistical power, we initially collected data from 67 participants. Participants were excluded if they failed one of the following criteria: 1) overall accuracy below 60% (8 participants); and 2) context checks below 60% correct (participants were asked at random times “Is the next trial probably hard or easy?”) (20 participants). Note that behavioral findings remain qualitatively similar when including these participants in the analyses. In addition, on a trial-by-trial level, we excluded trials with reaction times below 100 ms (0.24% of the data). The final sample comprised of 39 participants (age range 18–30, mean = 18.85, standard deviation [SD] = 3.46, mostly female (32 F, 7 M). There was one left-handed participant and 38 right-handed participants.

Experimental task

The task (Fig. 1) consisted of 50 mini-blocks of six trials. On each trial, participants reported the direction of a patch of dynamically moving dots and afterwards indicated their decision-related affect. The main advantage of the random-dot-motion task is its simplicity and elegant manipulation of the actual drift rate (i.e., actual progress) by manipulating the proportion of coherently moving dots (Gold & Shadlen, 2007). To influence expected progress, we used a block level manipulation. At the beginning of each mini block of six trials, as well as before each trial, participants were informed whether they were currently in a hard or easy block. Easy blocks consisted of 75% easy coherence trials (and 25% hard coherence trials) and hard blocks of 75% hard trials (and 25% easy trials).

The experimental task was programmed using the jsPsych library (De Leeuw, 2015) and implemented online within the cognition.run (*www.cognition.run*) experiment-hosting environment. The random-dot stimuli were constructed by using a plugin within the jsPsych library (Rajananda et al., 2018). The code for the experimental task is available at https://osf.io/z85td/.

The experiment had a 2x2 within-subjects design, with expected difficulty manipulated on two levels (hard block vs. easy block) and actual difficulty manipulated on two levels (2 levels of motion coherence: 0.05 hard vs. 0.4 easy). These levels of coherences were established in a pilot experiment, so that the hard condition would produce on average ~60% and easy condition ~90% correct trials. The trials were divided into two kinds of blocks (hard and easy) that aimed to generate either low or high expected progress for the current trial. The random sequences of trials within blocks were generated using a fixed seed, so all participants proceeded through the same sequence of randomized trials. To validate the effect of the expectation manipulation, we conducted a pilot study, where we asked for prospective difficulty ratings before each trial, which confirmed that in easy blocks participants indeed expected easier and in difficulty blocks more difficulty trials.

On the decision level, we measured overall performance, i.e., percentage of correct responses and reaction times to establish that the experimental manipulations worked as expected. The main dependent variable, affect, was measured on a discrete 6-point scale. Participants responded by using specific finger positioning so that their ring, middle, and index fingers of both hands were placed on numbers 1, 2, 3 for the left hand and 7, 8, 9 for the right hand. Because most participants used their dominant (right) hand to report higher levels of affect and less dominant hand to report lower levels, this could introduce a bias. However, even if this bias is present, it should not influence our main conclusions about the expected and actual progress. We instructed participants to try to report even slight changes in their affective state on the scale. The scale started with the prompt (“This trial made me feel….”) and had the following discrete response options: ”Rather negative”; “Quite negative”; “Slightly negative”; “Slightly positive”; “Quite positive”; and “Rather positive.” We did not use extreme labels (e.g., very negative or very positive for the affect scale), because we expected participants not to experience large emotional fluctuations on a trial-by-trial and instead opted for a discrete scale to make the experience of going through 300 trials as smooth as possible.

Procedure

Participants provided informed consent before the experiment, after which they completed a computer screen-size calibration procedure (the credit card method, where the size of stimuli is calibrated by a participant dragging a

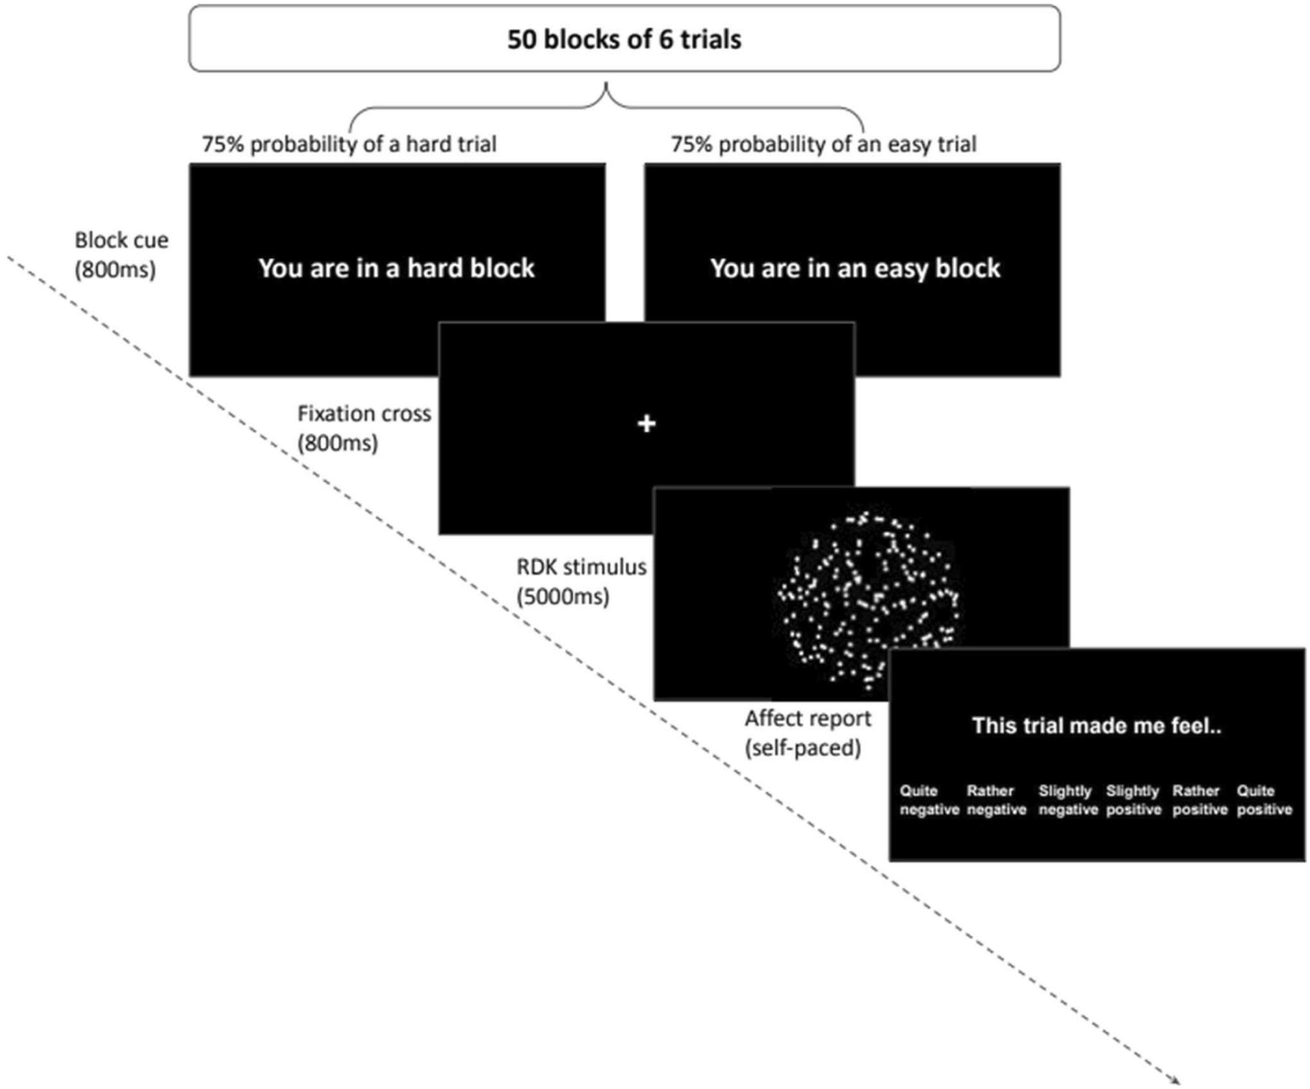

Fig. 1 Structure of the experimental task. For each trial, participants indicated the direction of the random-dot-motion stimulus. The coherence of the stimulus was manipulated on two levels: 0.05 (very low coherence - hard condition) or 0.4 (very high coherence – easy condition). For expected progress manipulation, the task was divided into two types of blocks, with block indicators before each trial. In hard blocks, the probability of getting a hard trial was 75%, and in easy blocks, the probability of getting an easy trial was 75%. Participants reported their affect on a six-level scale after every trial. Following the message on each trial, a fixation cross was shown for 800 ms followed by random-dot motion stimuli for 5000 ms. Participants had to respond within the 5000-ms time limit; otherwise, the message "Too late!" was displayed. Participants had to use keyboard buttons C and N with their thumbs to indicate leftward or rightward motion

rectangle on a screen to match a physical credit card). They were then shown a series of instructions about the task and how they should position their hands on the keyboard to give responses. Before moving to the experimental phase, all participants had to complete two training blocks. The first consisted of only easy (0.4) coherence trials with a threshold of 75% correct responses to move forward to the next block. If participants failed the criteria, they could attempt the training block again. The second training block had an equal amount of low and high coherence trials, with 60% correct responses as the threshold. In these training blocks, participants only reported the direction of the moving dots, without any affect reports. Before moving to the experimental phase, participants were introduced to the scales to be used and had to pass two test questions to validate that they had understood the keyboard to scale mapping (i.e., "Please indicate Rather negative"; "Please indicate Quite positive") before they could move on.

In separate phases of the experiment, participants were asked to report either affect or confidence, after having responded to the direction of the moving dots (counterbalanced across participants). In the current paper, we only analyze the affect phase of the experiment. In the experimental phase, participants got feedback about their performance after every 25 trials, during which they were also encouraged to take small breaks. After finishing the experimental phase, participants were required to fill out a short post-task questionnaire, in which their demographics (gender, age) and impressions about the experiment (subjective importance of the task, current mood, experienced boredom, used strategies) were queried. Completing the whole experiment took on average approximately 50 minutes.

Data analysis

Behavioral analyses were conducted in R with RStudio using linear mixed-effect models for affect and reactions times and generalized linear mixed effects models for analyzing accuracy relying on the R packages lme4 (Bates et al., 2009) and lmerTest (Kuznetsova et al., 2015). To estimate statistical significance of the factors in the models, we used Satterthwaite's approximation with Type III Anova if the interaction effects were included. To find the best models, we started from a model with the full fixed effect structure and the maximum random effect structure and compared this to models with the most random effects that would converge. If a model failed to converge, we did not include this in the comparison. We used BIC as criterion for model selection and picked the model with the lowest BIC value for final analysis. For all models, we checked that variation inflation factor (VIF) were below 5. The means reported in comparisons refer to averages computed from raw data.

Computational modeling

We model the decisions as an accumulation-to-bound process (Fig. 2), with the strength of the accumulation controlled by the drift rate v. In addition, the accumulation process is characterized by decision threshold a controlling the amount of evidence required to trigger a decision, starting point z controlling initial bias toward one choice option and nondecision time *ter* capturing the nondecision part (e.g., perceptual and motor) of the reaction time. To model expectations, we assume an additional accumulator racing to the same thresholds starting from the same starting point whose strength is not controlled by the stimuli presented on the screen but by the expectation that people have about the upcoming trial, v_e. Thus, the affectDDM consists of a stimulus-based accumulator, controlled by the drift-rate, and a parallel expectancy-driven accumulator, controlled by the expected drift-rate.

In line with the theoretical notion that affect reflects the prediction error between expected and actual progress, decision-related affect in our model was quantified as a function of the accumulated evidence in the stimulus-driven accumulator and the expectancy-driven accumulator. Note that we did not constrain the model to compute affect as the

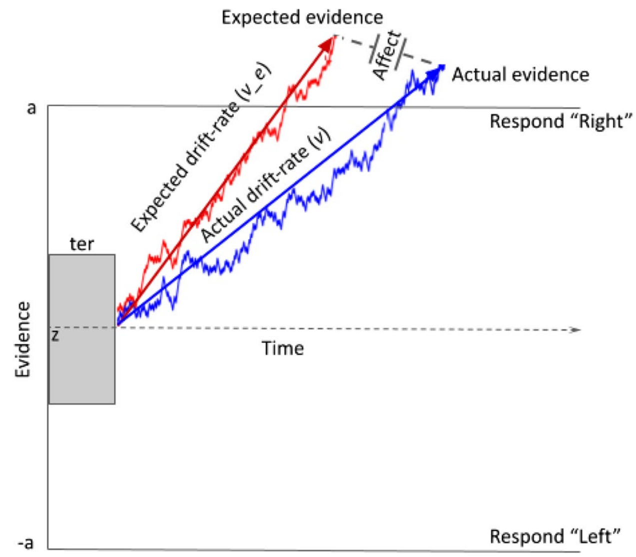

Fig. 2 Structure of the AffectDDM. Perceptual decisions are thought to be the result of a noisy process of evidence accumulation controlled by drift rate v, departing at starting point z, until one of two opposing decision boundaries a or $-a$ have been reached. In addition to the stimulus-driven accumulator (characterized by actual drift rate v) that is collecting evidence about the direction of dot motion, AffectDDM assumes an additional expectancy-driven accumulator (characterized by expected drift rate v_e), from which expected amount of evidence for a current trial is estimated. The difference between actual and expected amount of evidence underlies affect computation in the model (e.g., affect = actual evidence – expected evidence)

difference between these two quantities, such that during fitting the model also could land on affect as the sum of these quantities in case this would provide a better fit to the data. Following perceptual metacognition literature (Desender et al., 2021; Pleskac & Busemeyer, 2010), the model assumed that both the stimulus-based and expectancy-based accumulator processes continue after reaching the decision threshold for a fixed amount of time (controlled by a free parameter t). This is consistent with our trial structure where participants report affect only after they first reported their decision, such that their appraisal of progress can still be informed by post-decisional evidence accumulation.

We implemented the affectDDM using a random walk process starting from an unbiased starting point z, and updated evidence at each time step τ until reaching either a or $-a$ following Eq. (1):

$$\Delta e = v * \tau + \sigma * \sqrt{\tau} * \mathcal{N}(0, 1) \tag{1}$$

Where Δe is the change in evidence at each timestep, v is the drift rate, τ is the step size, σ is within-trial variability and $\mathcal{N}(0, 1)$ is Gaussian noise with a mean of 0 and standard deviation of 1. For our purposes, σ was fixed to 0.1, τ to 0.001 (s) and z to 0 in fitting the model. We used this equation to implement both the stimulus-driven accumulator and the novel expectancy-based accumulator. Specifically, Δe in Eq. 1 was computed by using the actual drift rate (v) for the stimulus-based accumulator and using the expected drift rate (v_e) for the expectancy-based accumulator. The actual drift rate was determined by the motion coherence of each trial and expected drift rate by the average motion coherence of each mini-block (i.e., easy or hard block).

Given that affect judgments were given post-decision, and during this time decision-makers could obtain new evidence, we allowed both accumulators to continue collecting evidence, after a decision boundary had been reached. The post-decisional evidence accumulation followed the same update rule as described in Eq. 1, using the same drift rates as the pre-decision phase. The post-decisional evidence accumulation terminated after a fixed amount of time, which was controlled by a free parameter ι . To produce an affect rating, the model computed the difference between the expected (generated by expectancy-driven accumulator) and actual (generated by stimulus-driven accumulator) amount of post-decisional accumulated evidence.

$$Affect\ raw_i = evidence\left(e_{t,i}\right) - Expected\ evidence\left(ee_{t,i}\right) \quad (2)$$

Affect on trial i is computed as the difference between actual accumulated evidence on trial i at time t and expected accumulated evidence on trial i on time t, where t reflects the sum of RT and l. This operationalization of affect is directly informed by the idea that affect reflects progress prediction error, where expected progress is compared to actual progress. To map raw affect predictions (which in theory range between -∞ and +∞) onto the affect scale used by participants, we used within participant standardization parameters (*bias_sd* and *bias_mean*), which scaled affect ratings according to idiosyncrasies of the scale use of each participant (Eq. 3, see Desender et al., 2022 for a similar approach). Finally, we obtained discrete values for predicted affect by dividing standardized predictions into equal ranges that determined the bins and mapping them onto a scale from 1 to 6 as in the empirical data.

$$Affect_i = affect\ raw_i / bias\ sd + bias\ mean \quad (3)$$

Fitting procedure

For each participant, we estimated two expected and two actual drift rates following our experimental design. Separate (actual) drift rates reflecting actual progress were estimated for trials with high and low motion coherence; separate expected drift rates were estimated for trials expected to be easy or difficult. The other parameters (*a*, *ter*, *l*, *bias_sd*, *bias_m*) were fixed within participants across conditions. To fit the parameters, we used quantile optimization where we defined a simultaneous cost function in terms of the discrepancy between model-predicted and observed reaction times and affect ratings. For reaction times, we separated the data into correct and error trials and computed the proportion of trials at five reaction time quantiles (.1, .3, .5, .7, .9) and compared these with model predicted quantiles. For affect ratings we computed the proportion of trials for each of the six response options and compared these with model predicted ratings. The cost function we used simultaneously minimized squared error for both quantities:

$$Cost = \sum \left(obsRT_{q,i} - predRT_{q,i}\right)^2 + \sum \left(obsAffect_q - predAffect_q\right)^2 \quad (4)$$

where obsRT refers to observed proportions of reaction times in quantile q separately for error and correct trials (i) and predRT to predicted proportions of reactions times for quantile q for error and correct trials (i). Similarly, obsAffect refers to observed affect ratings in for each discrete level (e.g., from 1 to 6) and predAffect to predicted affect ratings in the respective bins. To fit this model, we used the Differential evolution algorithm implemented in the DEoptim package (Mullen et al., 2011). We set the number of populations to 100 and set the fitting stopping criteria either to 1000 generations or if no updates in cost function have been obtained with the last 100 attempts.

Results

In the following, we first ask if the experimental task and manipulations worked as intended. We then compare predictions from the affectDDM to empirical data. Finally, we analyze whether fitted parameters from affectDDM adhere to the behavioral patterns to see if affect tracks progress prediction errors.

Reaction times and choice accuracy

To validate our experimental manipulations, we first analyzed choice reaction times and choice accuracy as a function of motion coherence (2 levels: high vs. low; reflecting actual progress), expected difficulty (2 levels: high vs. low; reflecting expected progress), and their interaction. For choice reaction times, we found a statistically significant main effect of motion coherence ($F(1, 38) = 34.82$, $p < .001$), but no effect of expected difficulty ($F(1, 38.5) = 0.008$, $p = 0.928$), nor an interaction between the two factors ($F(1, 11404.9) = 2.96$, $p = 0.086$). As illustrated in Fig. 3A, reaction times were faster in easy motion coherence trials (M = 0.707 s) than in hard motion coherence trials (M = 1.01 s), whereas expected difficulty had no effect on choice reaction times (M = 0.920 s, M = 0.807

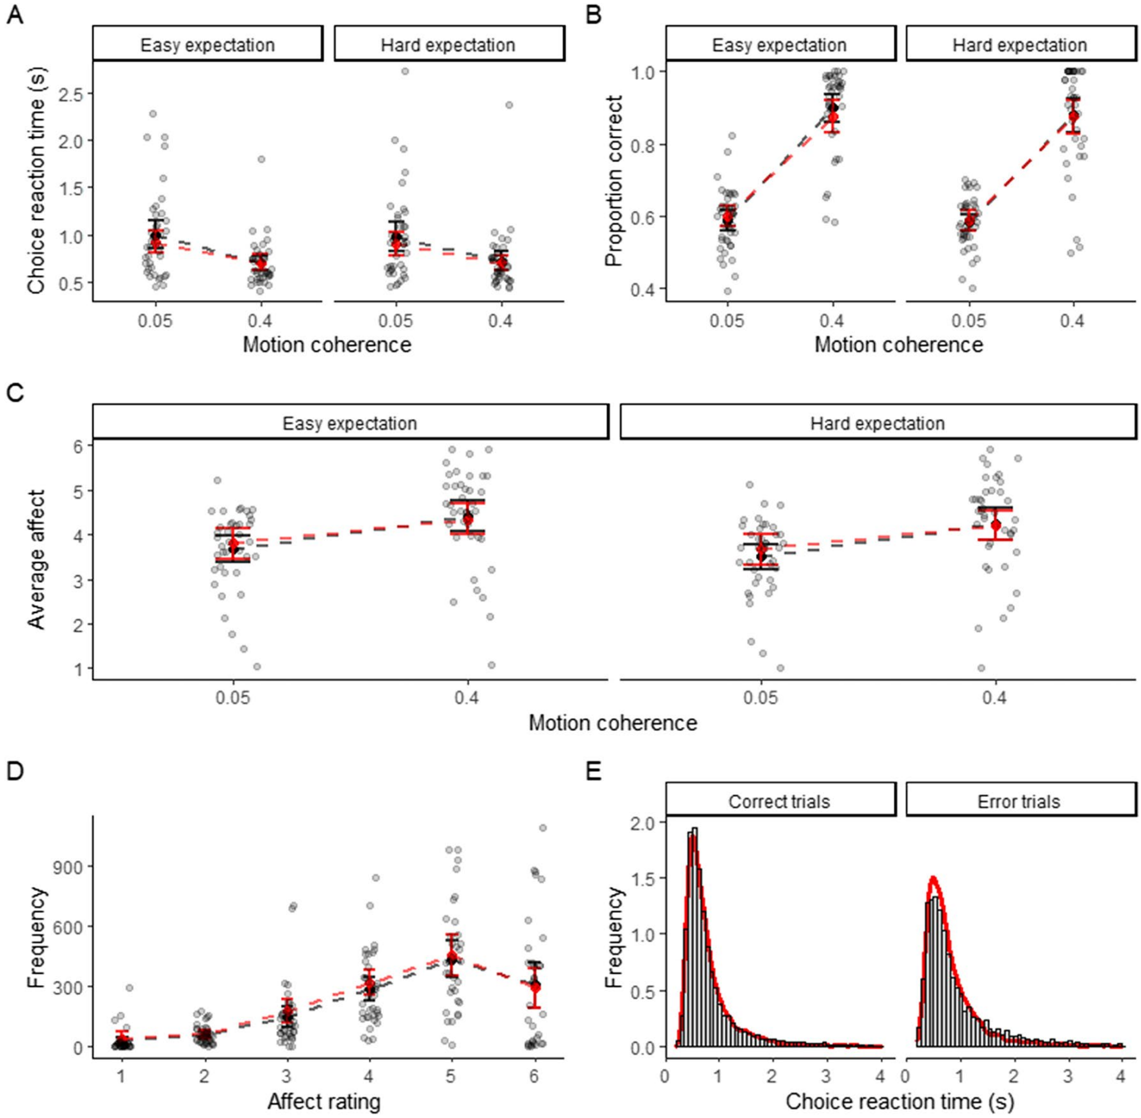

Fig. 3 Behavioral results and model fits. **A.** Choice RTs were shorter with high versus low motion coherence, but it was unaffected by expected difficulty. **B.** Choice accuracy was higher with high versus low motion coherence, but it was unaffected by expected difficulty. **C.** Affect ratings were influenced by both expected difficulty and motion coherence. **D.** The distribution of affect ratings across all conditions. The affectDDM closely captured the empirically observed distribution. **E.** The distribution of choice reaction times in error and correct trials. Again, the affectDDM closely matched the empirically observed distributions. Note: error bars indicate 95% confidence intervals; black and red data points refer to empirical data and affectDDM model fits, respectively.

s). For choice accuracy, we found a statistically significant effect of motion coherence (χ^2 (1) = 1113.056, p < .001), but not of expected difficulty (χ^2 (1) = 2.195, p = .139). The interaction effect was not statistically significant (χ^2 (1) = 2.251, p = .1335. As shown in Fig. 3B, choice accuracy was higher in trials with easy coherence (M = 0.894) compared with trials with hard coherence (M = 0.585), whereas accuracy did not differ depending on expected difficulty (M = 0.798 vs. M = 0.669). These results indicate that our actual progress manipulation worked as expected, whereas the expected progress manipulation did not influence behavioral performance.

Effect of expected and actual progress effects on affect

Our next analyses focused on the reported levels of affect, which were expected to be sensitive to both motion coherence and expected difficulty. Indeed, we found that both motion coherence (F(1, 38) = 44.07, p < .001) and expected difficulty (F(1,38) = 16.301, p < .001) had a statistically significant effect on affect. However, contrary to what can we expected from a prediction error account, the interaction effect was not significant (F(1, 11402) = 0.07, p = 0.792). Figure 3C shows that the reported affect was more positive in the easy motion coherence condition (M = 4.36) than in the hard motion coherence condition (M = 3.53). Similarly, affect was more positive in the easy expectation condition (M = 4.16) than in the hard expectation condition (M = 3.70). These results confirm that both actual and expected progress independently contributed to affect ratings. However, the specific pattern that we observed was not in line with the predictions from the progress prediction error account. According to this account, affect should be most positive when participants expected a difficult trial but the trial was easy, whereas the empirical data showed most positive affect when participants expected an easy trial and the trial was indeed easy. As an exploratory analysis, we ran a linear mixed-effects model with block index, expected, and actual difficulty as predictors of affect and found no main or interaction effects of the indexing variable, which suggests that the effect of expectation did not depend on their placement in the block.

Comparing AffectDDM model fits to empirical data

Having demonstrated that both actual- and expected difficulty contribute to affect ratings, whereas only actual difficulty influenced decision performance, we next turn to computational modeling to examine whether the affectDDM could account for these empirical findings. In Fig. 3, we added the model predictions (in red) on top of the behavioral data (black) in order to evaluate model fit. As can be seen, the model provided a good fit and accurately captured the distributional properties of both correct and error RT distributions (Fig. 3E). In addition, the model showed a close correspondence with the empirical affect ratings (Fig. 3D). To formally test whether the model reproduces the specific patterns seen in the empirical data, we analyzed the model predictions. Similar to the empirical pattern, the simulated data showed a significant main effect of motion coherence on reaction times ($F(1, 38) = 26.23$, $p < .001$) but no effect of expected difficulty ($F(1, 11620) = 0.076$, $p = .782$) nor a significant interaction effect ($F(1, 11620) = 0.096$, $p = .757$). Likewise, there was a significant main effect of motion coherence on accuracy ($\chi^2 (1) = 70.836$, $p < .001$) but no effect of expected difficulty ($\chi^2 (1) = 0.798$, $p = .371$). Similar to the empirical data, model accuracy was highest in high motion coherence trials (M = 0.875) and lowest on low motion coherence trials (M = 0.559).

Crucially, model-predicted affect showed similar effects of expected and actual progress as in the empirical data. Specifically, we found main effects of expected difficulty ($F(1, 38) = 11.862$, $p < .01$) and motion coherence ($F(1, 38) = 30.147$, $p < .001$) and no interaction effect on affect ($F(1, 11582) = 0.647$, $p = .421$). In the model predictions, affect was again most positive when expected and actual progress were both easy (M = 4.34) and the least positive when both were hard (M = 3.64). When the expected progress was hard but actual progress easy, affect was higher (M = 4.21) than when expected progress was easy and actual progress hard (M = 3.8).

Having established that our model closely captured the patterns seen in the empirical data, we next turned toward the estimated parameters (Table 1). In the model, we estimated two drift-rates per actual progress manipulation (v_easy and v_hard) and two drift rates per expected progress manipulation (v_e_easy and v_e_hard). In addition to standard DDM parameters (a - threhold, ter - nondecision time), we used additional free parameters to control the duration of post-decision evidence accumulation (l) and for adjusting idiosyncrasies of scale use (bias_m, bias_sd).

Parameters are the following: a- decision threshold, ter – non-decision time, v_easy – drift-rate in the easy coherence condition, v_hard – drift-rate in the hard coherence condition, v_e_easy – expected drift-rate in easy blocks, v_e_hard – expected drift-rate in the hard blocks, l - duration of the post-decisional evidence accumulation, bias_sd – standard deviation of individual level affect ratings, bias_m – mean of individual level affect ratings. Note, within-trial noise (sigma) was fixed to .1.

We expected higher drift-rates in easy coherence trials compared to hard coherence trials and higher expected drift-rates in easy difficulty blocks, compared with hard difficulty blocks. These predictions were confirmed for actual but not for expected drift rates. Actual drift rates were higher for high coherence trials (M = 0.187) than for low coherence trials (M = 0.027), $t(38) = 10.581$, $p < .001$. By contrast, the expected drift rates differed significantly between blocks, but in the opposite direction to what we expected. The estimated expected drift rate was lower in easy blocks (M = 0.187) compared with hard blocks (M = 0.209), $t(38) = -3.678$, $p < .001$). Importantly, these findings do not imply that participants expected easy trials to be more difficult. Instead, these findings are a direct consequence of our implementation that

Table 1 Estimated means and standard deviations of the model parameters

Estimate	a	ter	v_easy	v_hard	v_e_easy	v_e_hard	l	bias_sd	bias_m
Mean	0.07	0.36	0.19	0.03	0.19	0.21	1.70	0.37	0.81
Sd	0.02	0.09	0.10	0.02	0.11	0.11	1.22	0.28	0.86

affect reflects the *difference* between actual and expected difficulty. Given that the empirical data shows the opposite pattern (i.e., with affect reflecting the *sum* between both), the model likewise shows the opposite pattern. Importantly, if we would instead implement our model such that affect reflects the *sum* between actual and expected difficulty then the estimates would be flipped and would show higher expected drift rates for the easy condition. Put differently, the pattern in the expected drift rates can be interpreted as *reductio ad absurdum*, where the implementation of affect computation as progress-prediction error leads to empirically invalid predictions (e.g., higher/lower expected drift-rate in hard/easy blocks) which consolidates the previous behavioral finding that affect does not reflect progress prediction error but rather the sum of expected and actual progress.

Discussion

In the current work, we tested whether affect reflects progress prediction errors. This has been proposed by various theoretical accounts of affect generation (Joffily & Coricelli, 2013; Van de Cruys, 2017; Velasco & Loev, 2020), but direct evidence for this claim is mostly lacking. We designed a perceptual decision-making task where we manipulated both expected and actual progress in evidence accumulation and measured affective ratings after each trial. We also developed an extension of the drift diffusion model to estimate expected and actual drift rates of evidence accumulation, and model affect as a trial-by-trial function of their prediction error. Our results showed that both expected and actual progress influenced affect, but not following the hypothesized prediction error logic. Instead, both expected and actual progress influenced affect in the same direction with actual progress having a larger effect. Participants felt more positive when they had high expected progress and high actual progress and more negative when they had low expected progress and low actual progress. Results from the modelling effort are twofold: first, our analyses indicated that affectDDM, which generates affect ratings from a pair of evidence accumulators, is able to simultaneously and accurately account for choice reaction times, choice accuracy, and affect. Second, modeling results are consistent with the behavioral data and show that affect does not reflect a progress prediction error, but instead reflects the sum of the expected accumulated evidence and the actually accumulated evidence.

Our findings support the idea that progress appraisal plays a role in the computation of affect, by showing that higher actual progress (i.e., higher drift-rate) is associated with more positive affect (Carver, 2015). Our results also show that expected and actual progress both contribute to affective experience (as proposed by Van de Cruys, 2017) as opposed to affect reflecting only expected (Velasco & Love, 2020) or only actual progress (Joffily & Coricelli, 2013). However, although affect has been shown to track prediction error in cases of reward (Rutledge et al., 2014) or perception (Chetverikov & Kristjánsson, 2016), in our perceptual decision-making experiment we did not find evidence that affect reflects progress prediction errors. Instead, the results show that expected and actual progress contribute to affective experience in an additive manner, such that expecting higher/lower progress and actually experiencing higher/lower progress leads to more positive/negative affect.

Why does affect reflect both expected and actual progress?

First, the additive influence of expected and actual progress on reported affect can be explained by assuming that progress prediction errors do not elicit affect themselves. Instead, progress predictions errors might influence the mental construal of ongoing progress, which is then compared to a progress goal, which then generates affect. For example, on a difficult trial in an easy block, a participant with an initially high progress expectation (generated by the frequency of easy trials in the ongoing block) would gradually learn that actual progress on the ongoing trial is low. As outlined in Predictive Processing accounts (Clark, 2013; Van de Cruys, 2017; Velasco & Loev, 2020), the perceived progress on that trial would be a mental construal governed not only by sensory evidence but also by prior expectation. As a result, an optimistic expectation would bias perceived progress to be higher than experience evidence alone. Affect would then be generated by appraising this biased construal of ongoing progress in relation to a progress goal. Thus, more positive/negative progress prediction errors can lead to more positive/negative affect by introducing upward/downward biases to the construals that are used in progress appraisal.

An second explanation for the additive influence of expected and actual progress on affect builds on the concept of affective inertia (Kuppens et al., 2010). This concept refers to the degree to which affective states are resistant to change. In our experiment, participants were always informed about the difficulty of the upcoming six trials. Therefore, it is possible that prior to engaging in a trial, the expected progress manipulation already induced an affective state that was consistent with the expectation. As a consequence, this prior affective state could then influence the final affect rating in a direction consistent with the prior. A similar mechanism has been described by the mood-as-momentum account (Eldar et al., 2016) where rewards are perceived better when one is in a positive as opposed to negative mood. Note that the first and second explanation are not per se mutually exclusive. Instead, they identify two potential representations that might be updated based on expectations: in the first case it is the actual progress representation; in the second case it is the representation of affect itself.

A third explanation leverages the idea that affect reflects changes in expected value (Emanuel & Eldar, 2022). Given that expected value is a function of both costs and benefits, it is possible that participants experienced a higher cost in more difficult blocks, resulting in more negative affect. Expected value and higher costs have also been related to increased effort (Shenhav et al., 2013), which is generally negatively valenced (Inzlicht et al., 2018; Kurzban, 2016). Therefore, harder blocks, which require more effort, might be experienced as more negative. It is important to stress that all three explanations are constructed *post-hoc*, and therefore, future work should explicitly aim to dissociate between them and examine whether progress induced affect is influenced by progress discrepancy, previous affect (i.e., the affective inertia mechanism) or expected value.

Goals, progress expectations and epistemic goal pursuit

The conceptual analysis above also demonstrates that when conceptualizing progress appraisals, it is important to distinguish at least three features: expectations, goals, and construals (i.e., representations of current situations). Even though goals are a form of prediction as they usually represent states that do not yet exist (Smith et al., 2022), they have a very different function compared with other expectations with respect to affect generation (Klein, 2018; Moors et al., 2021; Van de Cruys et al., 2022). Goals function as reference values that can be used to appraise the construals of the current situation. For instance, a goal to write 1000 words per day can be compared to the construal of writing 1300 or 500, yielding positive/negative discrepancies that result in positive/negative affect. One implication of the present findings is that prediction errors may have an affective impact mainly as inputs to goals and construals, rather than as direct affect generators. As an example of expectations shaping construals, consider how an expectation to write a lot in a day can lift a persons' sense of how much they have written (before they count the words), thus biasing affect to be more positive. As an example of expectations shaping goals, consider a person who desires to write a lot but expects not to meet this goal and decides to set a more realistic goal. Thus, expectations and their violations can clearly be important indirect drivers of affect by updating both goals and construals, but they may lack a direct affective impact.

At first sight, these results might go against the hypothesis that people always prefer higher predictability, i.e., accuracy of their models of world (Clark, 2013). However, this apparent contradiction can be reconciled by considering accuracy as a (more or less salient) epistemic goal that, like all other goals, can be used to appraise (construals of) current situations. Under this view, it feels bad not to understand things not because of the prediction errors elicited in such situations but simply because low level of understanding diverges from desired level of understanding. This proposal therefore predicts that perceptual prediction errors do not generate affect themselves even as they are often correlated with affect, but because people have epistemic goals (e.g., desire to be accurate in their predictions). Consequencely, similarly to generic goal-pursuit, affect obtains the function of a regulatory signal, where more negative affect can induce more intense epistemic goal-pursuit, i.e., trying to be more accurate. This proposal also has preliminary evidence from the current findings, as the main goal in the current decision-making task is response accuracy, for which higher/lower progress resulted in more positive/negative affect but was uncorrelated to progress prediction error.

AffectDDM framework

We developed an extended Drift Diffusion model to test the idea that affect generation is based on how well one is progressing toward a goal (Carver, 2015). This model allowed us to measure expected and actual progress from the task and to see whether affect depends on both of these factors. The model results showed that affect indeed combines expected and actual progress, but not as a difference between them, rather as the sum. Future studies could use the affectDDM to predict how changing the parameters of evidence accumulation (i.e., response threshold; starting-point) would influence the reported affect, and then test these predictions empirically. Previous studies have used computational methods to study how affect is generated (Hesp et al., 2021; Loossens et al., 2020; Roberts & Hutcherson, 2019; Rutledge et al., 2014; Scherer, 2009), but they did not focus on how affect relates to decision-making and evidence accumulation. An exception to this is Givon et al. (2020), who used an evidence accumulation model to analyze affect reports. However, our approach is different, because we do not view affect as a decision outcome, but rather as a signal that arises in the context of the decision process itself.

Conclusions

We found computational and behavioral evidence that during perceptual decision-making, affect is positively associated with both actual and expected progress. This finding goes against the idea that affect reflects progress prediction errors. These conclusions were corroborated by an extension of the drift diffusion model, which was able to accurately fit affect ratings as reflecting the sum of expected and actual accumulated evidence.

Open practices statement The data, experiment and analysis code are available at https://osf.io/z85td/. The experiment was not preregistered.

Funding This research was supported by a CELSA grant from the KU Leuven (CELSA/21/010) and Estonian Research Council grant PSG525.

References

Bates, D., Maechler, M., Bolker, B., Walker, S., Christensen, R. H. B., Singmann, H., Dai, B., Scheipl, F., & Grothendieck, G. (2009). Package 'lme4'. *URL Http://Lme4. r-Forge. r-Project. Org*.

Brysbaert, M., & Stevens, M. (2018). Power analysis and effect size in mixed effects models: *A tutorial. Journal of cognition, 1*(1).

Carver, C. S. (2015). Control processes, priority management, and affective dynamics. *Emotion Review, 7*(4), 301–307.

Carver, C. S., Lawrence, J. W., & Scheier, M. F. (1996). *A control-process perspective on the origins of affect.*

Carver, C. S., & Scheier, M. F. (1990). Origins and functions of positive and negative affect: A control-process view. *Psychological Review, 97*(1), 19.

Carver, C. S., & Scheier, M. F. (2013). Goals and emotion. *Guilford Handbook of Cognition and Emotion*, 176–194.

Chetverikov, A., & Kristjánsson, Á. (2016). On the joys of perceiving: Affect as feedback for perceptual predictions. *Acta Psychologica, 169*, 1–10.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences, 36*(3), 181–204.

De Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a web browser. *Behavior Research Methods, 47*, 1–12.

Desender, K., Ridderinkhof, K. R., & Murphy, P. R. (2021). Understanding neural signals of post-decisional performance monitoring: An integrative review. *Elife, 10*, e67556.

Desender, K., Vermeylen, L., & Verguts, T. (2022). Dynamic influences on static measures of metacognition. *Nature Communications, 13*(1), 4208.

Drugowitsch, J., & Pouget, A. (2012). Probabilistic vs. non-probabilistic approaches to the neurobiology of perceptual decision-making. *Current Opinion in Neurobiology, 22*(6), 963–969.

Eldar, E., Rutledge, R. B., Dolan, R. J., & Niv, Y. (2016). Mood as representation of momentum. *Trends in Cognitive Sciences, 20*(1), 15–24.

Emanuel, A., & Eldar, E. (2022). Emotions as computations. *Neuroscience & Biobehavioral Reviews, 104977*.

Fleming, S. M., & Daw, N. D. (2017). Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. *Psychological Review, 124*(1), 91.

Givon, E., Itzhak-Raz, A., Karmon-Presser, A., Danieli, G., Meiran, N. (2020). How does the emotional experience evolve? Feeling generation as evidence accumulation. *Emotion 20*(2), 271.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience, 30*, 535–574.

Hesp, C., Smith, R., Parr, T., Allen, M., Friston, K. J., & Ramstead, M. J. (2021). Deeply felt affect: The emergence of valence in deep active inference. *Neural Computation, 33*(1), 1–49.

Inzlicht, M., Shenhav, A., & Olivola, C. Y. (2018). The effort paradox: Effort is both costly and valued. *Trends in Cognitive Sciences, 22*(4), 337–349.

Joffily, M., & Coricelli, G. (2013). Emotional valence and the free-energy principle. *PLoS Comput Biology, 9*(6), e1003094.

Khalvati, K., Kiani, R., & Rao, R. P. (2021). Bayesian inference with incomplete knowledge explains perceptual confidence and its deviations from accuracy. *Nature Communications, 12*(1), 5704.

Klein, C. (2018). What do predictive coders want? *Synthese, 195*(6), 2541–2557.

Kuppens, P., Allen, N. B., & Sheeber, L. B. (2010). Emotional inertia and psychological maladjustment. *Psychological Science, 21*(7), 984–991. https://doi.org/10.1177/0956797610372634

Kurzban, R. (2016). The sense of effort. *Current Opinion in Psychology, 7*, 67–70.

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2015). Package 'lmertest'. *R Package Version*, *2*(0), 734.

Loossens, T., Mestdagh, M., Dejonckheere, E., Kuppens, P., Tuerlinckx, F., & Verdonck, S. (2020). The Affective Ising Model: A computational account of human affect dynamics. *PLoS Computational Biology, 16*(5).

Moors, A., Ellsworth, P. C., Scherer, K. R., & Frijda, N. H. (2013). Appraisal theories of emotion: State of the art and future development. *Emotion Review, 5*(2), 119–124.

Moors, A., Van de Cruys, S., & Pourtois, G. (2021). Comparison of the determinants for positive and negative affect proposed by appraisal theories, goal-directed theories, and predictive processing theories. *Current Opinion in Behavioral Sciences, 39*, 147–152.

Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the brain: A diffusion model analysis of prior probability and potential payoff. *Journal of Neuroscience, 32*(7), 2335–2343.

Mullen, K., Ardia, D., Gil, D. L., Windover, D., & Cline, J. (2011). DEoptim: An R package for global optimization by differential evolution. *Journal of Statistical Software, 40*(6), 1–26.

Pleskac, T. J., & Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review, 117*(3), 864.

Proust, J. (2014). The representational structure of feelings. In *Open MIND. Open MIND*. MIND Group.

Rajananda, S., Lau, H., & Odegaard, B. (2018). A random-dot kinematogram for web-based vision research. *Journal of Open Research Software, 6*(1).

Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences, 20*(4), 260–281.

Russell, J. A. (2009). Emotion, core affect, and psychological construction. *Cognition and Emotion, 23*(7), 1259–1283.

Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2014). A computational and neural model of momentary subjective well-being. *Proceedings of the National Academy of Sciences, 111*(33), 12252–12257. https://doi.org/10.1073/pnas.1407535111

Scherer, K. R. (2009). The dynamic architecture of emotion: Evidence for the component process model. *Cognition and emotion, 23*(7), 1307–1351.

Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron, 79*(2), 217–240.

Smith, R., Ramstead, M. J., & Kiefer, A. (2022). Active inference models do not contradict folk psychology. *Synthese, 200*(2), 81.

Uusberg, A., Taxer, J. L., Yih, J., Uusberg, H., & Gross, J. J. (2019). Reappraising reappraisal. *Emotion Review, 11*(4), 267–282.

Van de Cruys, S. (2017). *Affective value in the predictive mind*. MIND Group.

Van de Cruys, S., Bervoets, J., & Moors, A. (2022). Preferences need inferences: Learning, valuation, and curiosity in aesthetic

experience. In *The Routledge international handbook of Neuroaesthetics* (pp. 475–506). Routledge.
Van Marcke, H., Denmat, P. L., Verguts, T., & Desender, K. (2022). Manipulating prior beliefs causally induces under-and overconfidence. *BioRxiv*, 2022–2003.
Velasco, P. F., & Loev, S. (2020). Affective experience in the predictive mind: A review and new integrative account. *Synthese*, 1–36.
Velasco, P. F., & Loev, S. (2022). *Cognitive feelings in the predictive mind: Emotion* Meta-cognition and Predictive Processing.
Villano, W. J., Otto, A. R., Ezie, C. E., Gillis, R., & Heller, A. S. (2020). Temporal dynamics of real-world emotion are more strongly linked to prediction error than outcome. *Journal of Experimental Psychology: General, 149*(9), 1755.