Affective Science
https://doi.org/10.1007/s42761-022-00118-5

RESEARCH ARTICLE

Stimulus-Driven Affective Change: Evaluating Computational Models of Affect Dynamics in Conjunction with Input

Niels Vanhasbroeck[1] · Tim Loossens[1] · Nil Anarat[2] · Sigert Ariens[1] · Wolf Vanpaemel[1] · Agnes Moors[1,3] · Francis Tuerlinckx[1]

Received: 15 October 2021 / Accepted: 10 April 2022

Abstract

The way in which emotional experiences change over time can be studied through the use of computational models. An important question with regard to such models is which characteristics of the data a model should account for in order to adequately describe these data. Recently, attention has been drawn on the potential importance of nonlinearity as a characteristic of affect dynamics. However, this conclusion was reached through the use of experience sampling data in which no information was available about the context in which affect was measured. However, affective stimuli may induce some or all of the observed nonlinearity. This raises the question of whether computational models of affect dynamics should account for nonlinearity, or whether they just need to account for the affective stimuli a person encounters. To investigate this question, we used a probabilistic reward task in which participants either won or lost money at each trial. A number of plausible ways in which the experimental stimuli played a role were considered and applied to the nonlinear Affective Ising Model (AIM) and the linear Bounded Ornstein-Uhlenbeck (BOU) model. In order to reach a conclusion, the relative and absolute performance of these models were assessed. Results suggest that some of the observed nonlinearity could indeed be attributed to the experimental stimuli. However, not all nonlinearity was accounted for by these stimuli, suggesting that nonlinearity may present an inherent feature of affect dynamics. As such, nonlinearity should ideally be accounted for in the computational models of affect dynamics.

Keywords Affect · Affect dynamics · Computational modeling · Context · Input

A central topic in the study of affect (i.e., the subjective component of emotions) is its dynamics, or how it changes over time. Affect dynamics can be studied by repeatedly assessing individuals' affective states, either in the laboratory or in daily life. Using this method, researchers were able to identify systematic individual differences in people's baseline affective state, the variability around this baseline, and the strength of the regulation towards it (Congard et al., 2011; Kalokerinos et al., 2020; Kuppens et al., 2010; Wendt et al., 2020). These individual differences have been linked to, among others, the development and maintenance of mood disorders (Ebner-Priemer et al., 2015; Kuppens et al., 2012; Sperry et al., 2020; Trull et al., 2015), albeit with varying degrees of success (Dejonckheere et al., 2019).

There are currently two dominant approaches to studying affect dynamics (see also Loossens et al., 2020). A first approach uses summary statistics to characterize an individual's affective time series (Dejonckheere et al., 2019; Wendt et al., 2020). This approach is especially useful for investigating specific characteristics of affective time series. For example, some studies have investigated group differences in person-specific affective means and standard deviations (e.g., Congard et al., 2011; Kalokerinos et al., 2019) while others have focused on more complicated affect dynamical characteristics, such as affective instability (Houben et al., 2021; Sperry et al., 2020). Despite this approach's predominance in the literature, it has one major limitation: it cannot provide a full account of affect dynamics. While summary statistics are useful for describing specific data patterns, they cannot

Handling editor: Jonathan Gratch

✉ Niels Vanhasbroeck
niels.vanhasbroeck@kuleuven.be

[1] Research Group of Quantitative Psychology and Individual Differences, KU Leuven, Leuven, Belgium

[2] Developmental Psychiatry, KU Leuven, Leuven, Belgium

[3] Center for Social and Cultural Psychology, KU Leuven, Leuven, Belgium

Published online: 14 July 2022

explain how these data come about, or what the data-generating mechanism looks like. Achieving these goals requires the use of a different approach, which consists of the development and application of computational models that formalize the fundamental dynamical features of an individual's affective life. A central question within this approach concerns the principles on which computational models should be built in order to properly describe affect dynamics.

The most commonly used class of computational models are the discrete-time (e.g., Adolf et al., 2017; Albers & Bringmann, 2020) and continuous-time (Boker & Nesselroade, 2002; Oravecz et al., 2011; Voelkle & Oud, 2013) autoregressive models, which postulate a linear relationship between current and past affective states. In spite of their popularity; however, they cannot accommodate several nonlinear features of affective data, including skew in affective measurements, an L-shaped relationship between positive and negative affect (PA and NA respectively; Diener & Iran-Nejad, 1986; Kuppens et al., 2013; Loossens et al., 2020; Norris et al., 2010; Schimmack, 2001), and multiple baselines to which affect can be regulated — a feature sometimes referred to as multistability (Hollenstein, 2015; Bonsall et al., 2012). These features, two of which are visualized in Fig. 1, have been found to describe fundamental characteristics of affective time series (Loossens et al., 2020), indicating that there might be a need to move from linear to nonlinear models in order to capture relevant dynamical features of affect.

Whereas the evidence of nonlinearity is starting to accumulate, the interpretation of this evidence is more contentious. While nonlinearity may be an inherent characteristic of affect, it can also be induced by external events, in which case the observed nonlinearity in affective time series is a reflection of the dynamics of the environment, and not an indicator of underlying nonlinearity in affect. Consider multistability as an example. It can arise as an inherent characteristic of one's affect dynamics (e.g., in bipolar disorder; Holmes et al., 2016; Bonsall et al., 2012; for a modeling perspective, see Steinacher & Wright, 2013), but could also arise due to the occurrence of positive and negative events (see Fig. 2). To avoid drawing the wrong conclusions about the operation of

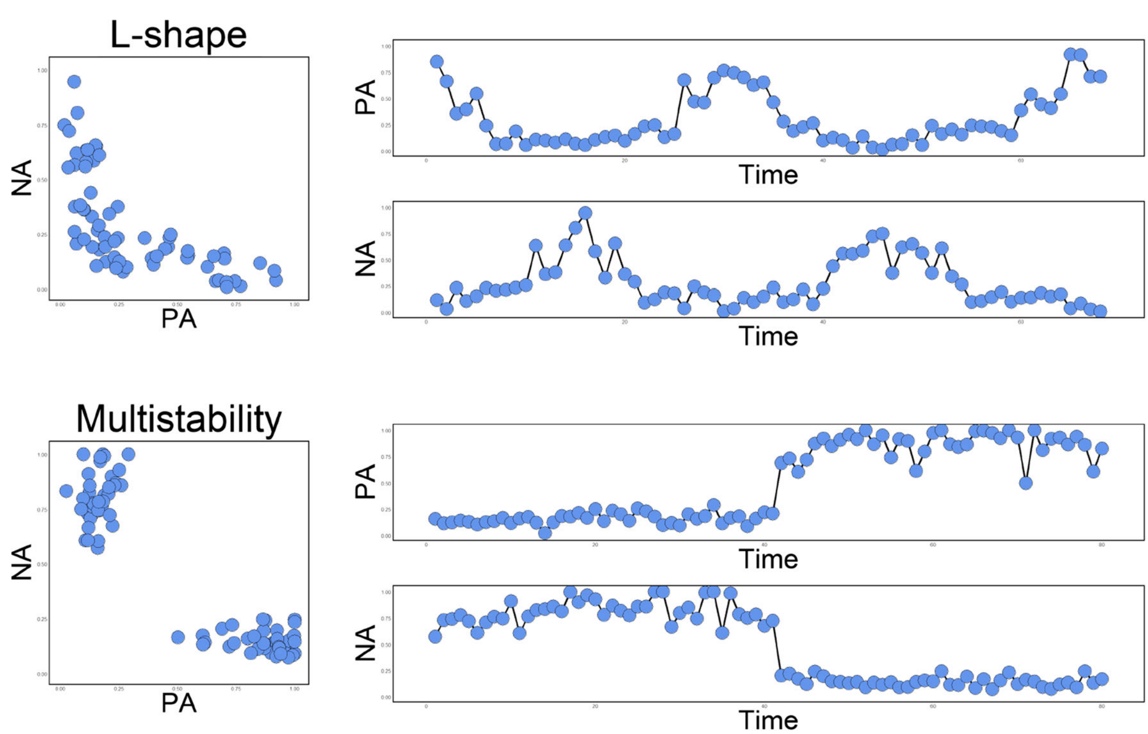

Fig. 1 Visualization of the two nonlinear features that will be central in this paper. The plots on the left show the distribution of PA and NA scores within the (bounded) PA-NA space. The plots on the right show (potential) time series that may accompany these kinds of distributions. The scatter plot and time series at the top visualize an L-shaped relationship between PA and NA. This means that when PA/NA is high, NA/PA will probably be low. Importantly, low affective states in both dimension may also occur, resulting in an L-shaped distribution within the PA-NA space. The second nonlinear characteristic, visualized in the lower plots, is multistability. Multistability implies that an individual experiences discrete-like affective states, such as overall positive (high PA, low NA) and overall negative (low PA, high NA) states, with little to no transitioning period between these states. This is most easily seen in the plotted time series, where one can see a sudden transitioning point in which the person switches from an overall negative to an overall positive state. Importantly, multistability may also involve faster transitions between discrete-like affective states than the ones plotted here

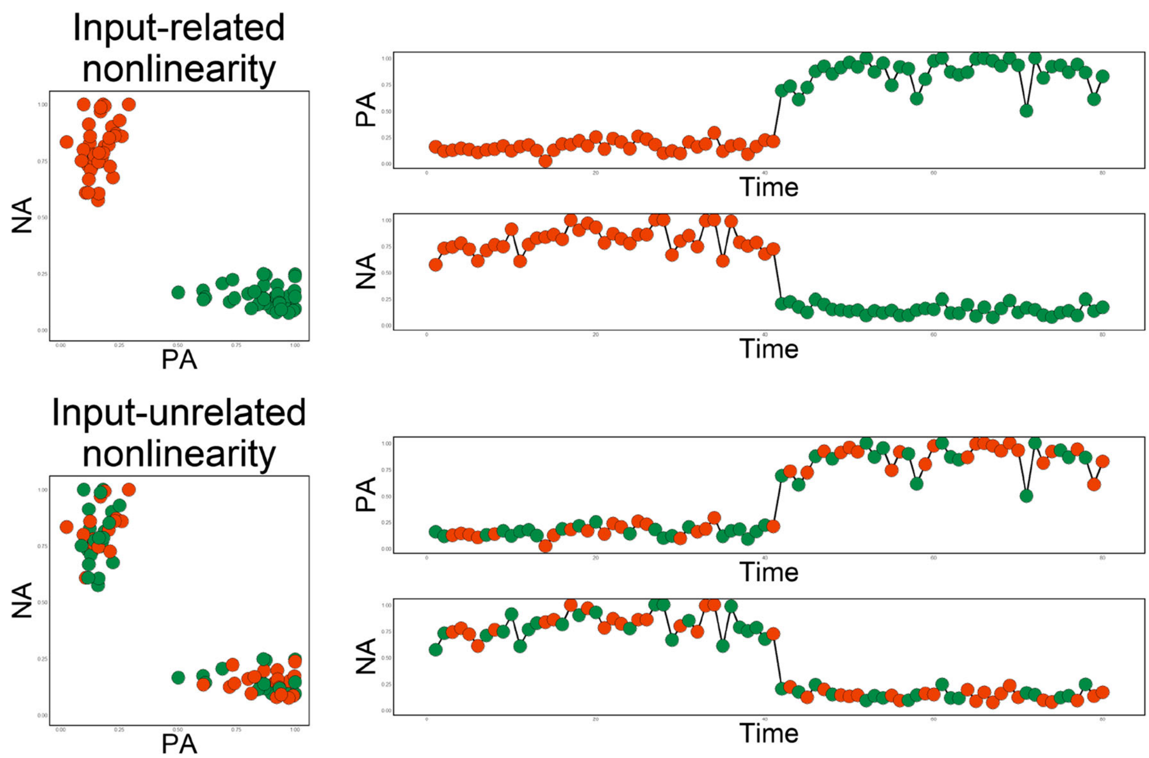

Fig. 2 Visualization of multistability that arises due to environmental stimuli (top) or due to environmentally unrelated nonlinearity of the affective system (bottom). Red dots denote the affective state during negative affective events, while green dots denote the affective state during positive affective events. If multistability is to be a feature of the affective system itself, then it should not be explained away by the environmental stimuli that a participant encounters, such as in the upper case

the affective system, it is thus important to study affect dynamics in relation to its immediate environment (Asutay et al., 2020, 2021; Kuppens et al., 2012; Kuppens & Verduyn, 2017; Lapate & Heller, 2020; Rutledge et al., 2014; Villano et al., 2020).

In the current study, we investigate whether nonlinearity is inherent to the affective system or whether it is solely a consequence of not accounting for affective stimuli. To this end, we compared the nonlinear Affective Ising Model (AIM) with a linear competitor, namely the Bounded Ornstein-Uhlenbeck model (BOU; see Loossens et al., 2020), a bounded version of the more often used Ornstein-Uhlenbeck model (OU; Driver & Voelkle, 2018; Oravecz et al., 2011; Uhlenbeck & Ornstein, 1930; Voelkle & Oud, 2013). Given that both models account for the experimental stimuli, we can answer our question by assessing the performance of the models though a variety of evaluation tools. More specifically, we will evaluate the models in a relative manner — through comparing model fit and predictive accuracy — and in an absolute manner — through assessing how well the models can reproduce specific characteristics of the data. If we find that the AIM outperforms the BOU, this is an indication that the experimental stimuli are not the sole source of nonlinearity in the affective data.

Method

Participants

A total of 178 individuals were recruited through Prolific, an online data collection platform (https://www.prolific.co/). The sample was well balanced with regards to gender (46% female) and educational background (31% high school, 38% bachelor, 25% master, 6% other). Participants were on average 27 years old (SD = 8, range = [18,59]).

As participants completed a probabilistic reward task, they were told that they would receive their total earnings as the reward for participation. Unbeknownst to them, this total was predetermined, so that everyone received £8 after successfully completing the experiment, irrespective of their behavior during the task. On average, participants spent 31min 42s to complete the task (SD = 11min 35s).

Materials

To assess participants' momentary affective states, we used a modified version of the Evaluative Space Grid (ESG; Larsen et al., 2008). The ESG is a two-dimensional grid whose axes are formed by PA and NA, allowing participants to report mixed feelings (i.e., affective states that are both positive and negative). For our experiment, we slightly modified the ESG to be continuous rather than discrete, to have the labels "Positive" and "Negative" on its axes, and to have four qualitative labels in the corners of the grid (clockwise starting at the lower left: "neutral," "bad," "mixed," and "good"). For our analyses, we scaled the participants' momentary affective states to take on values between 0 and 1, based on the boundaries of the grid.

Within the scope of two associated (unpublished) master's theses, we included several questionnaires which had to be filled out at the end of the experiment in a pseudo-random order, based on the time at which the participant started the study. Given that the research goals of these master's theses were different from the ones of this article, we will not discuss them further.

Fig. 3 Structure of a trial of the probabilistic reward task. Participants were presented with four doors, behind two of which hid a monetary reward (win), and behind the other two a monetary punishment (loss). At the top of the screen, participants saw their current total winnings, denoted as trial total. The modified version of the Evaluative Space Grid was shown at the bottom of the screen

Procedure

After filling out an informed consent form, participants were asked to provide some basic demographic information. Participants then received instructions and started the experiment.

The experiment was a probabilistic reward task. On each trial, participants were shown four doors that each hid either a monetary reward (win) or punishment (loss; see Fig. 3). Participants were told that two doors concealed a win, while the other two hid a loss, so that there was a presumed equal probability of winning or losing on each trail. Participants chose one of the four doors, after which their choice remained visible on screen for 500ms and the door opened, showing the trial outcome for 2s. One second after the trial outcome was known, the net total — which was always visible at the top of the screen — was updated.

At the end of each trial, participants were asked to report their affective state on the ESG at the bottom of the screen. The grid kept track of participants' temporary responses by placing a red dot at the previously clicked location. Participants confirmed their response by pressing the spacebar, and then moved on to the next trial after a 1-s delay. The red dot in the ESG disappeared with the start of each new trial to avoid imposing dependence between two subsequent trials. At no point during the study were participants limited by a time constraint; the experiment was self-paced.

Each participant received a starting capital of £3. To practice the trial structure, participants first went through 10 mandatory practice trials, after which their total was reset to £3. They then completed another 152 experimental trials.

Trial Generation To ensure that each participant had a comparable trajectory of stimuli, we fixed the order in which wins and losses occurred for each participant. That is, each participant went through the same string of outcomes, which was fixed beforehand. This information was, however, not available to the participants. Instead, they were told that the probability of winning and losing was equal on each trial. One main advantage of this design is that individual differences in affect dynamics can be more easily picked up. Given that all stimuli are the same, individual differences in affect dynamics cannot be attributed to a difference in the encountered stimuli, but are due to how participants affectively react to these stimuli. One disadvantage, however, relates to the motivation of the participants. Participants may (wrongly) assume that there is a pattern to be learned about the outcomes associated with different doors, and they may become frustrated over time when their internally generated hypotheses are proven wrong.

We generated the trials in several steps. First, wins and losses were semi-randomly drawn from the intervals [0.5,1.5] and [3,4], and the latter were multiplied by -1.[1] Then, these wins and losses were manipulated so that participants would build up a total gain of £5 during the course of

[1] The quantitative difference between both intervals was chosen to create a hypothesized qualitative difference between outcomes from these intervals.

the experiment, which together with the £3 starting capital makes a total reward of £8. Lastly, we randomized the order of all outcomes and saved this order to be used for each participant.

Computational Models

The main goal of our study was to evaluate whether nonlinearity is an important characteristic of affect dynamics, or whether it is an artifact caused by a lack of accounting for affective stimuli. For this, we assessed the performance of several computational models, each composed of two building blocks: a dynamical model and an input function. The former governs affective changes over time while the latter describes how the experimental stimuli are related to affect, thus linking the stimuli to the dynamical model. In what follows, we discuss both building blocks more thoroughly.

Dynamical Models The AIM and the BOU belong to the class of drift-diffusion models and consequently share several properties. First, the models assume that affective change happens continuously over time, which relates closely to how affective fluctuations are believed to be experienced (Boker & Nesselroade, 2002). Second, the models describe change in the affective system in terms of the current state, inherent stochastic fluctuations, and the presence of one (or more) *attractors* (Strogatz, 2018). An attractor can be interpreted as a baseline state to which the current affective state is regulated. In other words, every deviation from the attractor or baseline will be regulated back towards it. Due to the stochasticity of the models, variability around the attractor or baseline is to be expected, even in the absence of any immediate contextual influences.

To develop some intuition about how the dynamical models work, consider a marble that is released inside a bowl. As to our intuition, the marble tends to roll down towards the bottom of the bowl, which serves as the attractor of the system. However, stochastic influences will perturb the marble, so that it is flicked uphill again. Given these down- and uphill movements, the marble will stay in continuous motion, as is assumed to be the case with affect.

Modeling the dynamics of a process in terms of a baseline and variability around it (either in terms of regulation of perturbations, or in terms of oscillations) is common, both within and outside of psychology (see, e.g., Guastello et al., 2009; Boker & Laurenceau, 2006; Goldbeter, 2011), and both within a discrete- and continuous-time framework (e.g., VAR-models; Bringmann et al., 2018; Ariens et al., 2020). These models allow for capturing several affective phenomena, such as affective regulation, affective reactivity, and affective variability. Furthermore, these models allow for the identification of different regulatory patterns, each with their own defining trajectory (Strogatz, 2018; Vanhasbroeck et al., 2021).

In what follows, we zoom in on the AIM and the BOU separately. We provide an intuitive explanation of the models rather than a technical one. We refer the interested reader to the Supplementary Materials for a technical description of the models.

Affective Ising Model The AIM is a continuous-time, nonlinear drift-diffusion model that captures affective fluctuations through an affective surface. This surface is a person-specific hilly landscape through which an individual's affective state roams, in close resemblance to the marble metaphor above.

It is through this surface that the AIM displays its nonlinearities, of which we describe two in detail (visualized in Fig. 4). First, the AIM is able to create nonlinear "canyons" as seen from a bird-view perspective. These canyons may, for example, form an L-shape, having their minimum at neutral states (low PA-low NA) with protrusions towards positive (high PA-low NA) and negative states (low PA-high NA; see Fig. 4). Affect is expected to move within this canyon, thus mostly displaying neutral and occasionally positive or negative states (although mixed states are not excluded). Another nonlinear feature is what we previously referred to as multistability. This property refers to the presence of more than one (local) attractor or minimum in the affective surface. This means that while the affective state may initially be regulated towards one

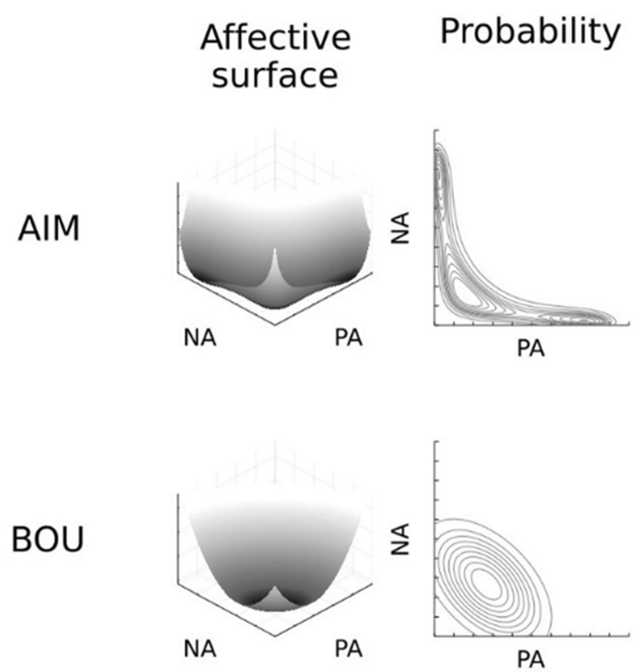

Fig. 4 Visualization of the affective surfaces and corresponding probability distributions of the AIM and the BOU. The affective surface of the AIM shows nonlinearities such as an L-shaped relationship between PA and NA, and multiple attractors, which can also be seen in the probability distribution. Such nonlinearities cannot be accommodated by the BOU, which assumes a linear relationship between PA and NA, and one single attractor

attractor, it may at some point jump towards another attractor, thus accounting for discrete jumps in affective states in a natural way (see again Fig. 1).

Bounded Ornstein-Uhlenbeck Model In contrast to the AIM, the BOU is a continuous-time, linear drift-diffusion model. The BOU's affective surface is much simpler than the AIM's, as it corresponds to a paraboloid bowl (see Fig. 4). This bowl may vary between being circular to elliptical, depending on the BOU's parameters (in particular those that capture the relationship between PA and NA). Unlike the AIM, the BOU cannot accommodate several nonlinear features of affect dynamics, such as the L-shaped relationships between PA and NA, and multistability in its affective surface.

As mentioned earlier, the BOU is a bounded version of the OU, in which reflecting boundaries are imposed on the model. Like the BOU, the OU is a linear drift-diffusion model that has received some attention in the literature. Unfortunately, the OU is an unbounded model, and it may thus produce nonsensical predictions that fall outside of the data range. As a consequence, the OU is naturally disadvantaged in comparison to the AIM, as the latter is naturally shielded against such nonsensical predictions. The reflecting boundaries within the BOU are imposed to alleviate this difficulty (see also Loossens et al., 2020).

Importantly, these boundaries make the BOU capable of accommodating skew in PA and NA, which stands in stark contrast to the strictly Gaussian patterns that the OU can create. By comparing the AIM with the BOU, we are therefore investigating whether the additional nonlinear features of the AIM (i.e., the L shape and/or multistability) provide a better means to describing affect dynamics than solely accounting for skew in the affective measurements, as done by the BOU. For simplicity's sake, we will refer to the BOU as a linear model throughout this text.

Input Functions On each trial of the experiment, participants encountered two experimental stimuli that are treated here as input to the dynamical models. These stimuli are the trial outcome (wins and losses) and the trial total (the updated total at a given trial) and their relationship to affect is captured by an input function. Such an input function links the input to the parameters of the dynamical models so that a change in the input produces a change in the affective surface and, ultimately, in the location (BOU) or strength of the attractors (AIM; see Fig. 5). For both dynamical models, the parameters β_1 and β_2 represent the input for PA and NA, respectively. Higher values of these parameters generally represent greater contextual change in the attractors, while a value of 0 represents the absence of any contextual influence. More specifically, this means that the zero vector $[0,0]^T$ represents the absence of any contextual influences, while the vectors $[\beta_1,0]^T$, $[0,\beta_2]^T$, and $[\beta_1,\beta_2]^T$ represent contextual influences on PA, NA, and PA and NA combined, respectively.

One crucial question, however, concerns the way in which the stimuli from the experiment determine the values of β_1 and β_2. To accommodate this issue, we constructed seven exploratory input functions, each differing in two aspects. First, input functions differed in which experimental stimuli they account for, namely only immediate trial outcomes, only the trial total (i.e., the total accumulated up to a specific trial), or both trial outcome and trial total. Second, the input functions differed in how the input plays a role, which was operationalized as a linear relationship between the parameters of the models and either the real values of the variables or dummy-coded versions thereof. Dummy-coding of the stimuli consisted of either differentiating between wins and losses (for trial outcomes) or between a positive and negative total (for trial total). Furthermore, we assumed that positive/negative events only changed PA/NA (relative to the neutral state). For completeness, we also considered the case in which the stimuli do not influence affect. The input functions that were used in our analyses are displayed in Table 1 with their respective names.

Combining these input functions with the AIM and the BOU, a total of 14 models were compared to each other. In the remainder of this paper, we will refer to these combinations with a label that specifies the dynamical model and the input function (using the labels from Table 1). For example, an AIM coupled with an input function that accounts for wins and losses in a dummy-like way will be referred to as AIM O_2. We note that the letter O stands for the inclusion of the *outcome* and the letter T for the inclusion of the *total*. The subscript *2* denotes a dummy-like effect of these variables and the subscript L denotes a linear effect. The single exception to these rules is the input function O_4, which includes a dummy-like effect for the outcomes drawn from each separate outcome generation interval (see the "Trial Generation" section).

Statistical Analyses

Using the data, models were tested both in a relative and an absolute sense. We describe both strategies in turn, but first detail how parameters were estimated as an intermediate step.

Parameter Estimation The parameters of the models were estimated for each individual separately through the GradientDiffusion package (Loossens et al., 2021). Estimates were obtained by minimizing the min-log-likelihood ℓ using the differential evolution heuristic (DE; Storn & Price, 1997; also see Supporting Information of Loossens et al., 2020). The general procedure went as follows. First, an individual's affective time series was divided into pairs of measurements, with each pair containing the affective states at trials t and $t-1$. Then, P parameter sets were randomly

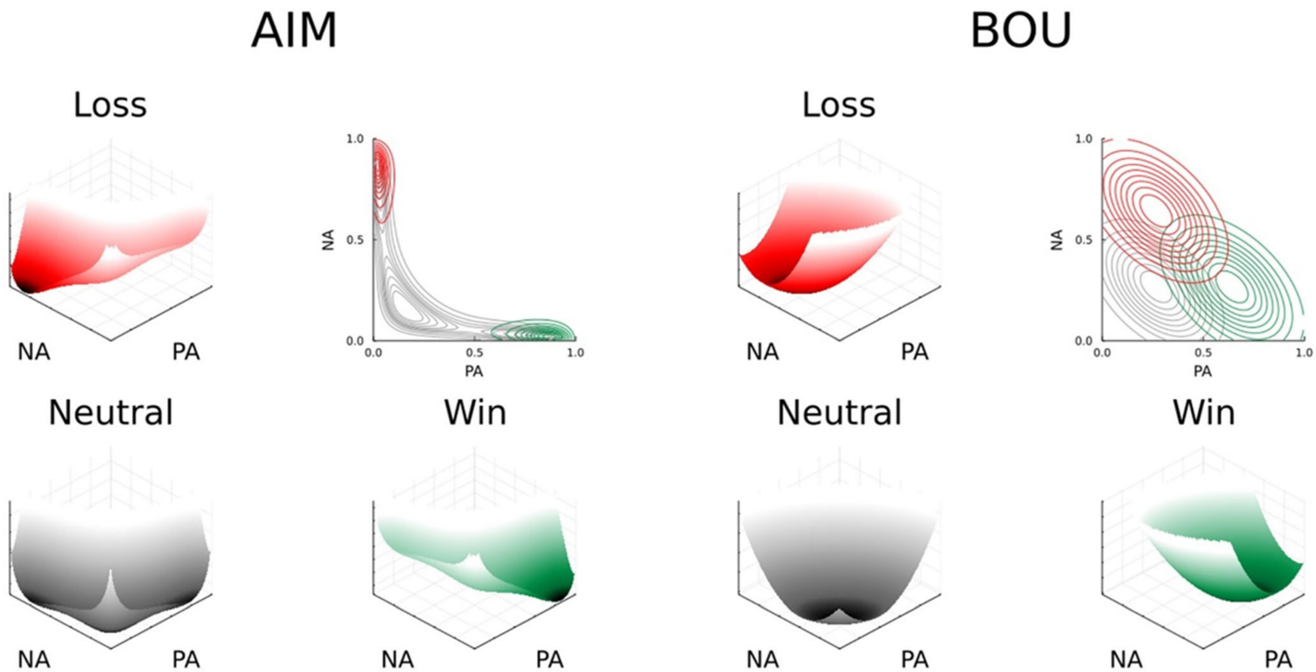

Fig. 5 Visualization of the context-related changes to the AIM (left) and BOU (right). For the AIM, the affective surface is tilted in a specific direction. For the BOU, this surface is not tilted, but the attractor changes location more directly. Both manipulations lead to an increased probability of experiencing certain affective states, which is visualized in the probability plots

generated and served as starting points from which optimization would occur. At each step of the DE, the probability distribution at time t conditional on the affective state at trial t–1 was computed for each parameter set. Importantly, these probability distributions were also conditional on the proposed parameters, the time that had passed since the previous affective measurement, and the contextual input at trial t. Next, the min-log-likelihood ℓ was computed for all measurement pairs and summed together, providing us with an overall measure of fit of the model to the data, using specific

Table 1 Values of the parameters β_1 and β_2, as determined by the input functions. Importantly, β_1 primarily changes the attractors of the models in the PA direction (i.e., higher or lower PA), while β_2 is primarily concerned with NA. The values of the ω's (for trial outcomes) and τ's (for trial total) are free parameters to be estimated for each individual within each dynamical model, and determined the strength with which the β's influenced the attractors

Input name	Code	When?	β_1	β_2
Empty		Always	0	0
Dummy outcome	O_2	Win	ω_1	0
		Loss	0	ω_2
Linear outcome	O_L	Win	$\omega_1 \times \|O_t\|$	0
		Loss	0	$\omega_2 \times \|O_t\|$
Dummy total	T_2	Positive total	τ_1	0
		Negative total	0	τ_2
Dummy low and high outcome	O_4	Low win	ω_{11}	0
		High win	ω_{12}	0
		Low loss	0	ω_{21}
		High loss	0	ω_{22}
Dummy outcome, dummy total	O_2T_2	Win, positive total	$\omega_1 + \tau_1$	0
		Win, negative total	ω_1	τ_2
		Loss, positive total	τ_1	ω_2
		Loss, negative total	0	$\omega_2 + \tau_2$
Linear outcome, dummy total	O_LT_2	Win, positive total	$\omega_1 \times \|O_t\| + \tau_1$	0
		Win, negative total	$\omega_1 \times \|O_t\|$	τ_2
		Loss, positive total	τ_1	$\omega_2 \times \|O_t\|$
		Loss, negative total	0	$\omega_2 \times \|O_t\| + \tau_2$

parameters. Finally, the parameter sets were combined to form P new parameter sets that were evaluated in the same way. The ℓ's of this "new" generation (children) were compared to the ℓ's of the "old" generation (parents). Whenever a child outperformed its parent, the child was selected for further use, and otherwise, the parent was used in the next part. The selected parameter sets were then used to create the next generation, and so the cycle continues. After N iterations of this combination and evaluation procedure, the parameter set with the lowest ℓ was selected to be the result of the estimation procedure.

As this procedure is a heuristic, it was not guaranteed that the global minimum would be found, possibly leading to suboptimal parameter estimates (Myung, 2003). The procedure was therefore run five times, each with 2,500 iterations and 100 starting values. To make sure that the parameters of the contextualized models converged, we ran the parameter estimation an additional five times with 5,000 iterations and 100 starting values when the ℓ of the models did not meet a basic sanity check. This sanity check consisted of comparing the ℓ for more complex models to the ℓ of a simpler, nested version thereof. When two models are nested, the ℓ of the simpler model should be greater than the ℓ of the more complex model, as a complex model will always fit better to the data than a more restricted model. Following this idea, we compared the ℓ of more complex models to the ℓ of simpler models, checking whether the former was indeed lower than the latter. For input functions with two parameters, we compared their ℓ to the ℓ of the *Empty* input function. For input functions with four parameters, we compared their ℓ to the O_2 input function, as the more complex input functions were all nested within this input function.

In the Supplementary Materials, we report on a recovery study in which we assessed whether the models' parameters could be adequately recovered with this procedure on a subset of all models under investigation. In short, we find evidence that this is indeed the case.

Relative Model Fit All models were compared to each other with regard to their relative fit to the data. This comparison may reveal whether the added complexity of the nonlinear AIM is necessary to fit the data. If this is not the case — that is, if the linear BOU outperforms the AIM — then we conclude that the observed nonlinearity can be attributed to the stimuli of the experiment. However, if the AIM were to outperform the BOU, then this nonlinearity is not completely dependent on the experimental stimuli, and it may thus possibly reside within the affective system itself. In essence, this analysis thus shows whether nonlinearity must be accounted for when fitting models to the data.

We compared the relative fit of the models in two ways, namely by using the AIC as a general measure of model fit and by comparing the predictive performance of the models using a cross-validation procedure.

General Fit As a measure of fit of a model to the data, we used Akaike's Information Criterion (AIC; Akaike, 1974). This measure balances the fit of a model to the data and the complexity of the model by penalizing fit according to the number of parameters of the model. As such, the AIC gives an indication of goodness-of-fit while compensating for overfit.

To ensure that we can adequately recover the relative performance of the models by comparing the AIC, we performed a distinguishability study using the AIC as a measure of fit (see Supplementary Materials). The results of this study suggest that we can adequately distinguish between very similar models. However, for practical reasons, we could only assess the distinguishability of a subset of all models and subsequently, results of the model comparison should be considered accordingly.

Predictive Performance Predictive performance was assessed using a blocked 15-fold cross-validation procedure (see Arlot & Celisse, 2010; Hastie et al., 2009). First, we transformed an individual's time series to contain the affective state at time t and time $t-1$ (see also *Parameter estimation*), ensuring relative independence between measurement pairs. Then, this transformed affective time series was repeatedly separated in two parts: a training sample and a test sample. The mapping of the data to samples was determined based on temporal proximity, so that data that lie closer together are kept together (see Fig. 6), ensuring relative independence between training and test samples (Roberts et al., 2017). The parameters of the models were estimated on the training sample and subsequently used to predict the data of the test sample. As we know the data of the test sample, we could assess how accurately the models predicted these data. We used the ℓ of the data given the conditional probability distributions as a measure of predictive accuracy.[2] This procedure was repeated until every data point served once in the test sample. The measure of predictive accuracy was averaged across iterations to obtain an overall assessment of predictive performance.

Absolute Model Fit While the relative fit based on AIC and predictive accuracy give a relative indication of how well a model fits the data, this same model may not necessarily be able to reproduce the characteristics of the data (Palminteri et al., 2017). This is important, as a computational model should be able to reproduce the data to answer the questions we posed above, namely the question of what the affective

[2] While not a conventional measure, we decided upon this alternative because the models were inherently stochastic in nature. This stochasticity limits the interpretation of deterministic point-predictions, which are typically used in other measures of predictive accuracy (e.g., the mean squared error).

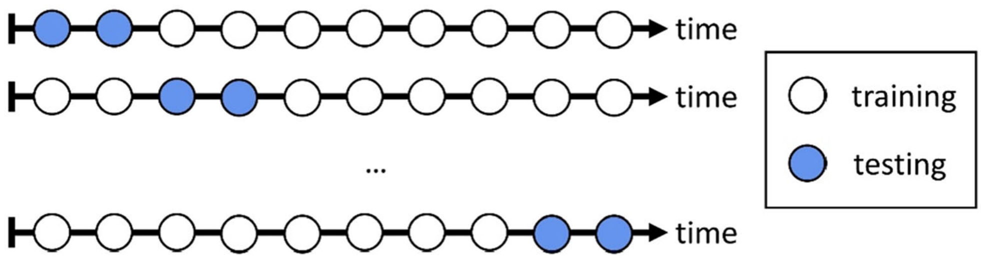

Fig. 6 Visualization of the selection of training and test data within the blocked *k*-fold cross-validation procedure

system looks like. Relative performance fails to say anything about whether a model is able to produce data that resemble the observed data, and thus of whether the model resembles the data-generating mechanism.

To assess the models' ability to reproduce the data characteristics, we used a parametric bootstrap procedure.

Procedure The procedure is visualized in Fig. 7. First, we fitted all models on each of the 178 participants' data. Next, we identified each participant's best-fitting AIM and BOU (i.e., the model with the lowest AIC). Based on the optimal parameter sets, we then simulated 1,000 new datasets per model per participant. We then computed a number of statistics of interest (described below) for both the observed data and these simulated data. Lastly, we compare the observed value of these statistics to the distribution of the simulated values. This comparison provides us with some indication of whether the models can adequately reproduce this observed value, and with it the characteristic of the data associated with the statistic.

Ideally, the models are able to perfectly reproduce the values of the statistics. However, this is unlikely to occur, urging us to account for the uncertainty associated to the estimates of the statistics. Therefore, we assessed the percentage of participants for whom the observed value of the statistic falls within the 95%CI of the simulated values. We will refer to this percentage as coverage.

We applied the parametric bootstrap on the best-fitting version of the AIM and BOU separately. With this procedure, we could assess the models' ability to reproduce the statistics of interest in both an absolute and a relative manner. As such, differences in the expected behavior between both models

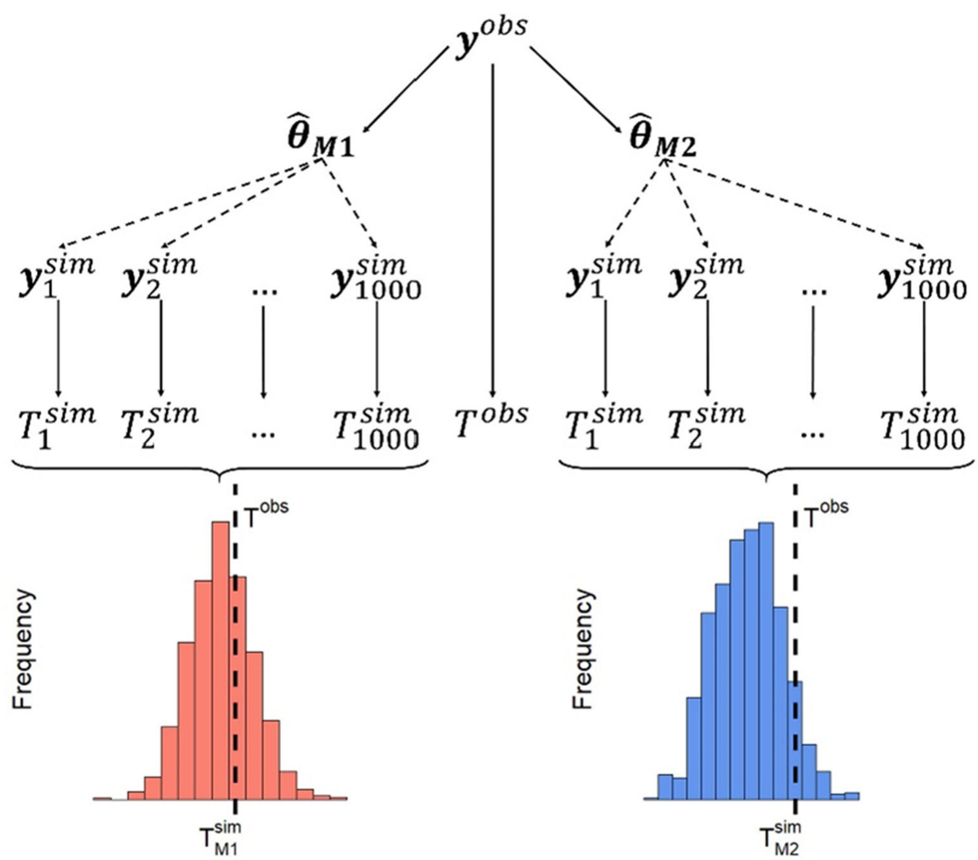

Fig. 7 Visualization of the parametric bootstrap procedure. The analysis starts from the observed data $\boldsymbol{y}^{obs}$, from which the parameters of model 1 and model 2 are estimated, resulting in $\hat{\boldsymbol{\theta}}_{M1}$ and $\hat{\boldsymbol{\theta}}_{M2}$, respectively. Next, N new datasets $\boldsymbol{y}_i^{sim}$ are simulated using the estimates retrieved through the previous step. Finally, a given statistic T is computed for both the observed data (T^{obs}) and the simulated data (${T_M}^{sim}{}_{1,i}$ and ${T_M}^{sim}{}_{2,i}$). This procedure leads to two distributions of T^{sim} to which you can compare the observed value of T (black line), as shown at the bottom of the plot. Comparing the observed value with the simulated values may give an idea of whether there is a mismatch between the models' dynamics and the observed affect dynamics

could be assessed, and divergence between expectation and observation could be identified for both models separately. This comparison also allowed for determining the cases in which the nonlinearity of the AIM paid off over and above the inclusion of an input function. In other words, this analysis could again indicate whether nonlinearity is due to the experimental stimuli or due to the way in which affect is structured. In the former case, both the AIM and the BOU should reproduce the statistics to a similar degree. In the latter case, however, the AIM should reproduce the statistics more closely than the BOU. To check whether there were differences in coverage between the two models, we used a nonparametric bootstrap procedure in which we constructed 10,000 new datasets of each 178 individuals from the original dataset. For each of these samples, we then assessed the coverage of the AIM, the coverage of the BOU, and the difference in coverage between the two.

Statistics For the choice of the statistics, we focused on several interrelated goals. First, we wanted to establish whether the computational models were able to reproduce some basic characteristics of the observed data. To answer this question, we used the following statistics:

(a) Autocorrelation of affect (r_{auto}). This statistic provides an indication of whether the models adequately account for the (linear) relation between two subsequent affective states. We calculated this measure for PA and NA separately;
(b) R^2 of a linear regression of the outcomes and total on affect. This statistic provides an indication of whether the models can adequately account for the stimuli of this study. Again, we calculated this measure separately for PA and NA;
(c) Correlation between PA and NA ($r_{\text{pa-na}}$). This statistic provides a linear measure of the association between PA and NA.

Additionally, we focused on the most important nonlinear characteristics of the data and evaluated the extent to which they could be reproduced by the models (see Loossens et al., 2020):

(a) Curvature of the relationship between PA and NA (κ). This statistic is a measure of the nonlinear association between PA and NA. More specifically, κ is aimed at capturing a (possible) L-shaped relationship between PA and NA, such as the one described in the introduction. For a formal description of this statistic, we refer the interested reader to Loossens et al. (2020);
(b) Bimodality coefficient of the data (*overall BC*). The bimodality coefficient is a measure of bimodality in the data, such that higher values of this statistic indicate the presence of more than one mode. In their study, Loossens et al. (2020) found that the BC was another statistic better captured by the AIM, probably due to its ability to incorporate multiple attractors in the affective surface. For a formal description of this statistic, we refer the interested reader to Loossens et al. (2020);
(c) Bimodality coefficient conditional on the affective state at trial $t-1$ and the stimuli at trial t (*Dynamical BC*). While the *overall BC* may identify general multistability in the data, multistability may also have its effect in a dynamical way. To illustrate, consider a participant who feels really good at a given trial $t-1$, displaying high PA at that trial. Now consider what might happen when a loss is received at trial t. The participant may suffer from this loss and feel less well, resulting in a lowered PA at trial t. However, the participant may also continue to feel good, despite the loss. Such an *affective lingering* is a dynamic kind of multistability and it implies that people may sometimes be inert in their affect dynamics, regardless of the environmental stimuli. To capture this, we define the *dynamical BC* to be the *BC* of a selection of affective data, where the selection is based on the extremity of the affective state at trial $t-1$ and the outcome at trial t. Extremity of the affective state is defined as the affective state being higher than the 75th quantile of all affective data. This measure is calculated for PA and NA, and for wins and losses on trial t separately.

Finally, we included some statistics that reflect some common assumptions of both dynamical models. These assumptions were tested with the following statistics:

(a) Autocorrelation of the residuals from an AR(1) model (r_ε). This statistics was used to address the Markovian assumption of both models, which holds that the current affective state only depends on its immediate predecessor.[3] In other words, the affective state at more distant lags (i.e., further in the past) are ignored. To test this assumption, we first filtered out the lag-1 temporal relations of PA and NA using an AR(1) model and then computed the autocorrelation of the filtered data. If the Markovian assumption holds, no temporal relations should be found in these filtered data, and the r_ε should be close to 0;

[3] While the Markov assumption is aimed at the lag of variables, it should be mentioned that the current affective state does not only depend on the previous affective state at lag 1, but also on the parameters of the model, the time that has passed since the previous measurement, and the stimuli of the experiment (depending on the input function). Furthermore, for our models, the Markov assumption applies to the conditional probability density function rather than an exact value of the current affective state (see the "Parameter Estimation" section).

(b) Variability of the value of a statistic estimated within a moving window (SD_T). This statistic was used to assess the stationarity assumption, which dictates that the parameters of the model should remain constant over time, or in other words that the model itself is time-independent. To test this assumption, we used a moving window of 100 data points to compute the value of a few statistics T repeatedly. Given all values of T, we then computed its standard deviation SD_T , revealing whether these statistics change in value over time. If the stationarity assumption holds, the SD_T should be close to 0. We computed two statistics this way: (a) the mean of PA and NA, as a general statistic that may change over time, and (b) the slope of the trial outcomes from a multiple regression on PA and NA, capturing the presence of temporal changes in the sensitivity to the stimuli of the experiment.

Results

Descriptive Results

Figure 8 displays the between-person averaged distribution of PA and NA scores, accompanied by an averaged time series (panel A). Red dots denote the affective state after a participant lost money, while green dots denote the affective state after a win. On average, the affective states of participants became more positive after a win (higher PA, lower NA) and vice-versa after a loss (lower PA, higher NA). Furthermore, the average result showed a largely linear relationship between PA and NA. However, this average does not imply that nonlinear features were absent from the data. To show this, we included PA-NA distributions of three participants (panel B). The first plot shows a linear pattern, resembling the average PA-NA patterns displayed in panel A. The other plots show nonlinear L-shaped patterns, with the right-most plot closely resembling the multistability pattern of Fig. 1. Interestingly, the latter pattern does not fully coincide with wins and losses, indicating that this participant may have displayed affective lingering in their time series.

The stimulus dependence and individual differences in affect were also confirmed by a number of descriptive results. Across participants, mean PA and NA scores were equal to 0.40 and 0.52 respectively, showing interindividual differences (SD_{PA} = 0.12; SD_{NA} = 0.15). Participants' affective states showed intra-individual variability (for PA: M_{SD} = 0.30; for NA: M_{SD} = 0.30), with some participants showing a larger variability than others. Some of this variability could be explained by the trial outcomes and trial total, although the contingency differed across individuals (for PA: M_{R2} = 0.57, SD_{R2} = 0.20; for NA: M_{R2} = 0.52, SD_{R2} = 0.22). Taken

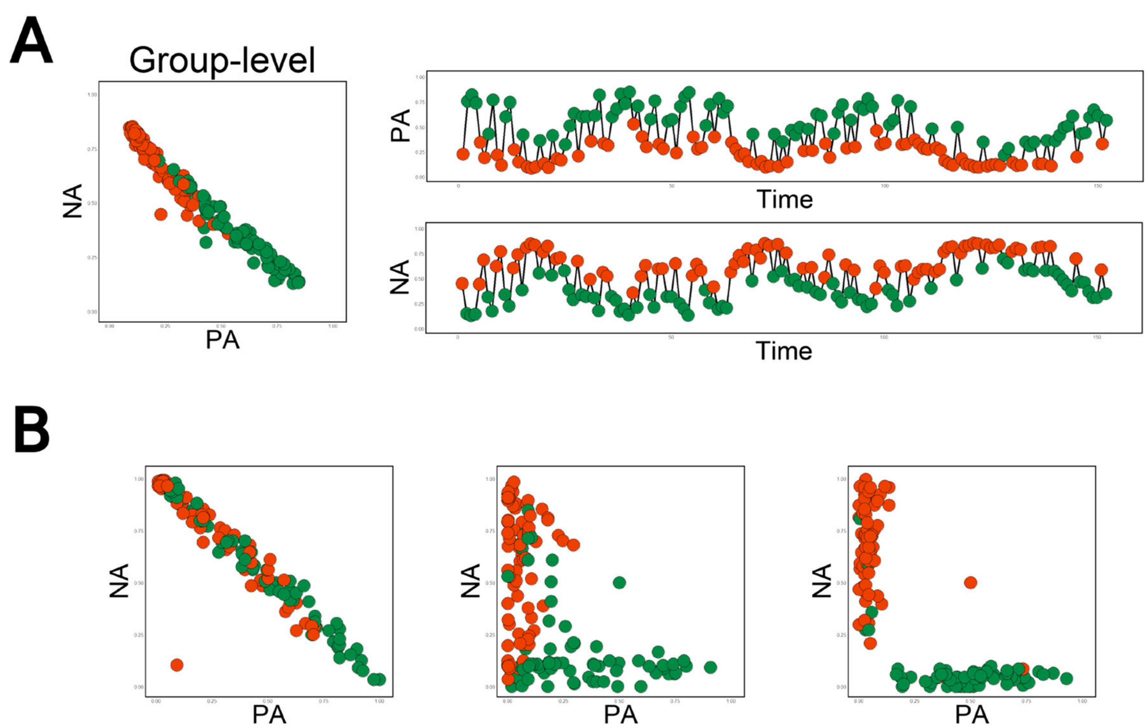

Fig. 8 Visualization of the data from our experiment. Panel A visualizes the distribution of PA and NA scores, and their time series averaged over all participants. In panel B, the distribution of PA and NA scores for several participants are displayed. These participants were selected based on their affective distributions to exemplify both the linear and nonlinear patterns that were observed in this study

together, these results indicate that affect fluctuated throughout the experiment and that these fluctuations were, at least partially, related to the experimental stimuli.

Relative Model Fit

General Fit Figure 9 summarizes the results of the model comparison based on the AIC of each model. In panel A, the frequency with which a model was selected as the within-person best (left) and worst (right) model is displayed. Overall, models that accounted for the stimuli of the experiment performed best (about 99% of the cases), indicating that the stimuli should be accounted for when modeling affect dynamics. This conclusion is further corroborated by the fact that models that did not account for these stimuli performed worst overall (about 84% of the cases). These results suggest that one should account for affective stimuli. However, models with input functions did not always perform better than models without input function (16% of the cases). Given that most of these cases were resolved by using another input function, this result shows that one should closely consider which stimuli to account for. Accounting for affective stimuli does not necessarily imply a better relative model fit compared to ignoring these stimuli.

The latter conclusion is further corroborated by a closer inspection of panel A, which reveals that different individuals responded differently to the stimuli. More concretely, there is not one input function (nor dynamical model) that performed best across all participants. Rather, participants showed differences in the kind of stimuli that they affectively respond to (only outcomes: 29%; only total: 2%; both outcome and total: 69%). Despite these individual differences, the majority of participants responded to both the trial outcomes and the trial total. Overall, participants were thus affected by all information shown to them.

In Fig. 9, panel B, the difference in AIC between the best-fitting version of the BOU and the AIM are displayed. As can be seen, the performance of the AIM and the BOU is similar, with the former accounting for 54% of the best-performing models. Nonlinearity is thus not needed to capture the affect dynamics of all participants, but at the same time, nonlinearity cannot be explained away by the experimental stimuli for all of them. One should thus consider accounting for nonlinearity when analyzing affective time series.

Predictive Performance Figure 10 summarizes the results of the model comparison based on the predictive performance of each model. Models that accounted for the stimuli of the experiment performed best for 99% of the participants, while

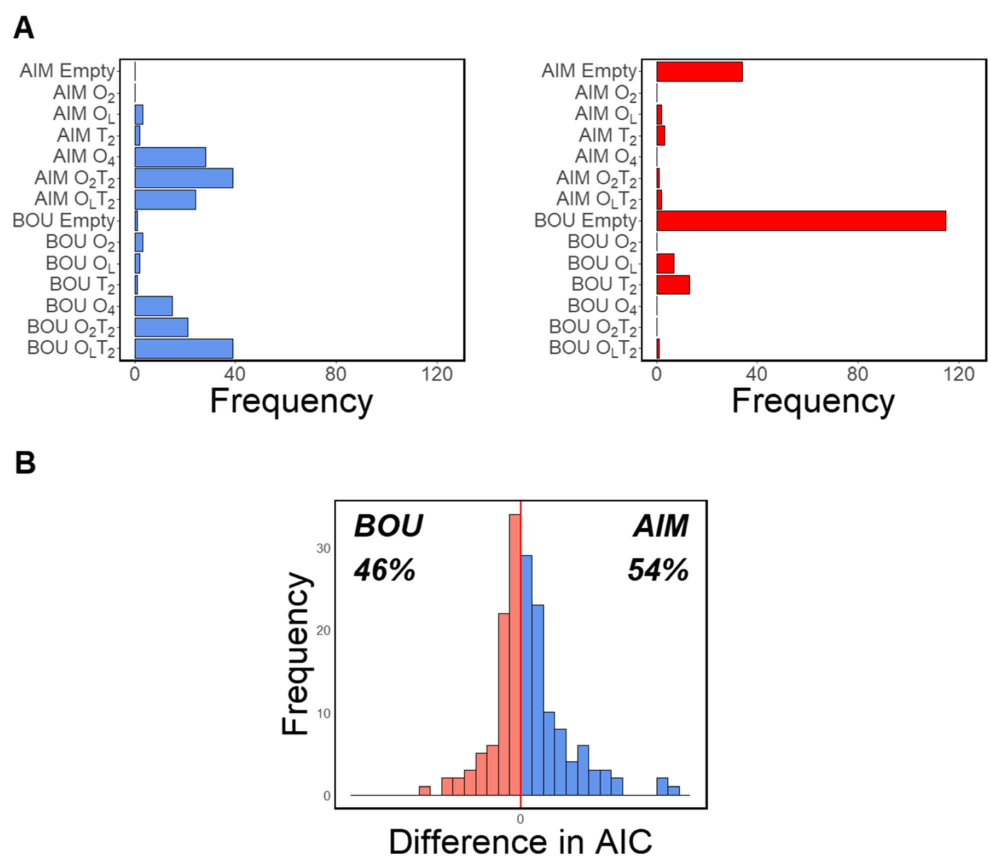

Fig. 9 Visualization of the AIC-related results. In panel A, the number of times a model was found to perform best (left; blue) or worst (right; red) is shown. In general, we see a preference for models that account for the experimental stimuli. Furthermore, an equal number of participants seem to prefer a BOU or an AIM. This result is confirmed in panel B, where the within-person differences in AIC of the best AIM versus the best BOU are visualized. Bars of the histogram that fall to the left of the red line show a preference of the BOU, and vice-versa for bars that fall to the right of this line

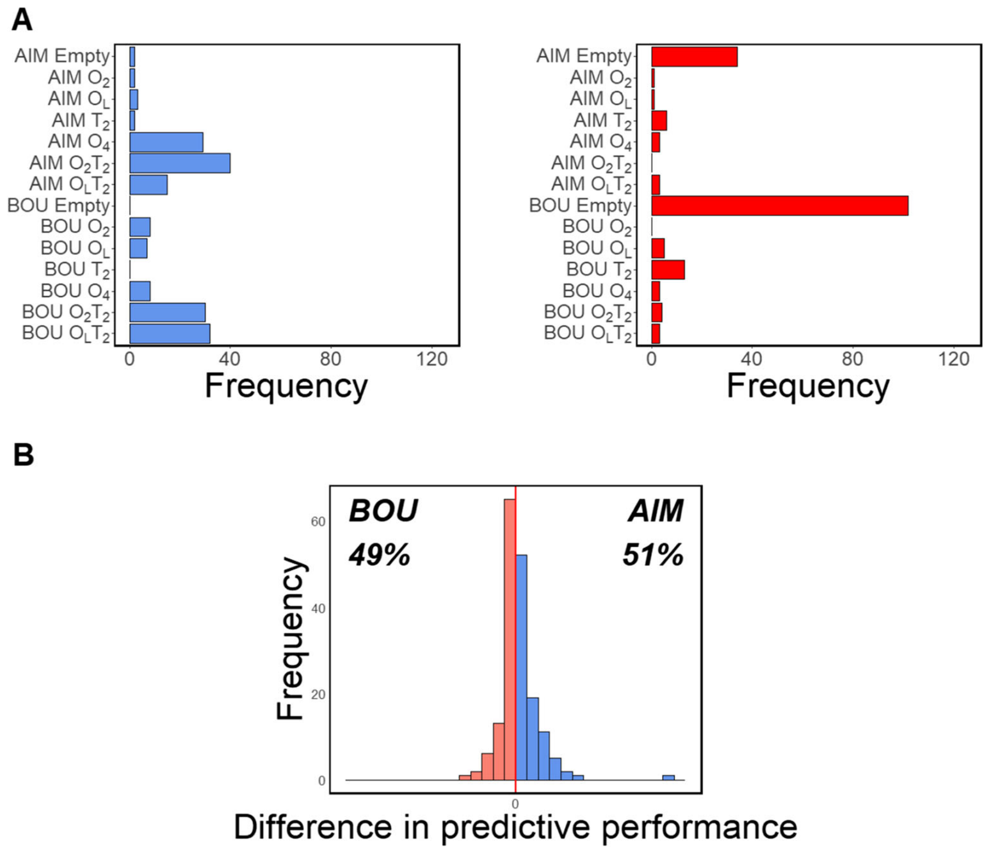

Fig. 10 Visualization of the results from the cross-validation procedure in a similar way to the AIC-based results. In general, we see a preference for models that account for the experimental stimuli. Furthermore, an equal number of participants seem to prefer a BOU or an AIM

models that did not account for these stimuli performed worst for 76% of the participants. Participants again showed differences in the kind of stimuli that they affectively responded to, with 32% of them responding only to outcomes, 1% only to the total, and 66% to both outcomes and total. Furthermore, the predictive performance of the best-fitting versions of the AIM and the BOU were again comparable, with the former accounting for 51% of the best-performing models.

Overall, the results from both measures of the relative model fit tell a similar story: (a) affective stimuli should be accounted for, (b) different individuals were sensitive to different kinds of stimuli, and (c) nonlinearity paid off for several, but not all, participants, even when accounting for these stimuli.

Absolute Model Fit

Basic Features Figure 11 displays the mean simulated values against the observed values of the statistics aimed at capturing some basic features of the data, using each participant's best-fitting AIM and BOU. In the top-left corner of each plot, the rank-order correlation (Spearman's ρ) is shown.

For the autocorrelation of affect (r_{auto}), we found that both models captured the observed value well, both in terms of sensitivity to the observed value and in terms of coverage, with little consistent difference in performance between the AIM and the BOU (PA: $\rho_{\text{AIM}} = 0.96$, $\rho_{\text{BOU}} = 0.92$, $C_{\text{AIM}} = 85$, $C_{\text{BOU}} = 83$, 95%CI $C_{\text{BOU–AIM}} = [-7, 4]$; NA: $\rho_{\text{AIM}} = 0.90$, $\rho_{\text{BOU}} = 0.91$, $C_{\text{AIM}} = 85$, $C_{\text{BOU}} = 76$, 95%CI $C_{\text{BOU–AIM}} = [-16, -3]$). The models were thus both well-suited to capture the temporal dependence of affect.

With respect to the R^2, both models seemed sensitive to its value, so that weaker or stronger associations between the stimuli and affect were reflected in the model-based simulated values of the R^2 (PA: $\rho_{\text{AIM}} = 0.85$, $\rho_{\text{BOU}} = 0.79$; NA: $\rho_{\text{AIM}} = 0.83$, $\rho_{\text{BOU}} = 0.75$). However, both models almost consistently underestimated its value, resulting in a low coverage rate (PA: $C_{\text{AIM}} = 34$, $C_{\text{BOU}} = 46$, 95%CI $C_{\text{BOU–AIM}} = [5, 20]$; NA: $C_{\text{AIM}} = 47$, $C_{\text{BOU}} = 48$, 95%CI $C_{\text{BOU–AIM}} = [-7, 8]$). The computational models thus produced data that show a weaker linear association between the stimuli and affect. This is surprising, as the best-performing models almost all account for these stimuli, and should thus be intuitively successful in reproducing this statistic. One reason for this misfit was related to the fact that the R^2 was computed for a linear regression of the stimuli on affect, without accounting for previous affective states. Indeed, coverage, but not sensitivity, rises when accounting for these affective states (PA: $\rho_{\text{AIM}} = 0.88$, $\rho_{\text{BOU}} = 0.78$, $C_{\text{AIM}} = 51$, $C_{\text{BOU}} = 54$, 95%CI $C_{\text{BOU–AIM}} = [-3, 11]$;

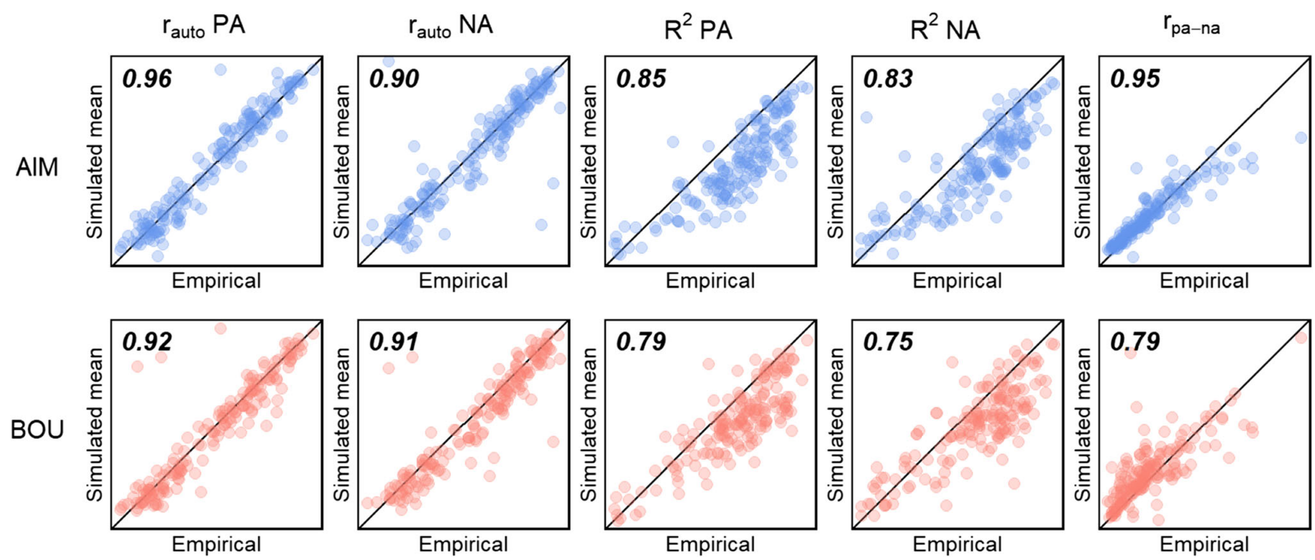

Fig. 11 Mean simulated values of the linear statistics for the within-person best BOU (bottom) and AIM (top) plotted against the observed value. In the upper left corner, the Spearman correlation (ρ) between observed and simulated values is indicated

NA: $\rho_{\mathrm{AIM}} = 0.82$, $\rho_{\mathrm{BOU}} = 0.75$, $C_{\mathrm{AIM}} = 51$, $C_{\mathrm{BOU}} = 47$, 95%CI $C_{\mathrm{BOU-AIM}}$ = [−11, 3]). However, misfit was still observed, which might indicate that some part of the input was not fully represented by the models. One source of such misfit was a restriction of the effect of the trial outcomes and trial total on affect. More specifically, these two variables could not interact with each other. However, such interaction was observed for a minority of the participants.

Finally, both models were sensitive to the value of the correlation between PA and NA ($r_{\text{pa-}na}$), although it was better reproduced by the nonlinear AIM (C_{AIM} = 76, C_{BOU} = 59, 95%CI $C_{\mathrm{BOU-AIM}}$ = [−27, −10]). It should be noted that the AIM can only produce negative correlations between PA and NA due to its architecture. In this sample, about 94% of the participants showed such a negative relationship, warranting this assumption for the vast majority of them.

Overall, it thus seems that both models capture these basic characteristics of the data — namely temporal dependence of affect, association between affect and stimuli, and the linear relationship between PA and NA —satisfactorily. The AIM outperformed the BOU with regard to some of the statistics, and otherwise performed similarly well.

Nonlinear Features Figure 12 displays results of the parametric bootstrap for the nonlinear features in a similar way as in Fig. 11.

For the curvature (κ), we again see that both models were able to deal with the interindividual variation in the data (ρ_{AIM}

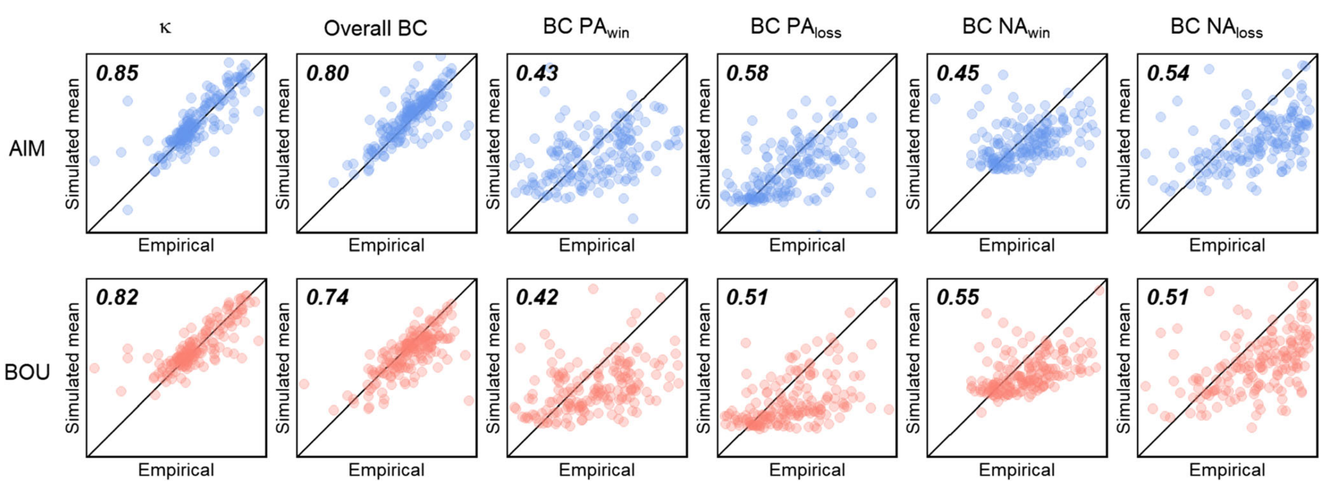

Fig. 12 Mean simulated values of the nonlinear statistics for the within-person best BOU (bottom) and AIM (top) plotted against the observed value. In the upper left corner, the Spearman correlation (ρ) between observed and simulated values is indicated

= 0.85, ρ_{BOU} = 0.82). This sensitivity was further corroborated by the coverage rates, as there was no difference between the two models in the ability to reproduce the value of κ (C_{AIM} = 75, C_{BOU} = 75, 95%CI $C_{BOU-AIM}$ = [−7, 6]). This is a remarkable result, as it was previously found that the BOU was unable to reproduce this statistic to the same extent as it does here (see Loossens et al., 2020). This result thus suggests that the BOU benefits from being coupled with an input function, so that it can reproduce this nonlinear data pattern closely.

For the bimodality coefficient (*overall BC*), we found that the AIM and the BOU were both sensitive to the observed values, although the AIM showed smaller deviations from the diagonal (ρ_{AIM} = 0.80, ρ_{BOU} = 0.74). Furthermore, we found that the AIM had a higher coverage rate than the BOU (C_{AIM} = 65, C_{BOU} = 53, 95%CI $C_{BOU-AIM}$ = [−21, −2]). Despite these coverage rates not being ideal yet, this result suggests that even after taking the stimuli into account, there was still some multistability left, which the AIM could then (partially) account for.

This observation is somewhat corroborated by the results for the dynamical bimodality coefficients (*dynamical BC*), which were also better reproduced by the AIM, regardless of the combination of affect and stimulus (PA_{win}: C_{AIM} = 72, C_{BOU} = 65, 95%CI $C_{BOU-AIM}$ = [−12,−1]; PA_{loss}: C_{AIM} = 80, C_{BOU} = 69, 95%CI $C_{BOU-AIM}$ = [−18,−5]; NA_{win}: C_{AIM} = 75, C_{BOU} = 63, 95%CI $C_{BOU-AIM}$ = [−17,−6]; NA_{loss}: C_{AIM} = 58, C_{BOU} = 52, 95%CI $C_{BOU-AIM}$ = [−12,−1]). However, the models were only slightly sensitive to the observed values of the statistic (range ρ_{AIM} = [0.48, 0.74]. range ρ_{BOU} = [0.46, 0.67]). These conflicting results can be explained by the relatively low number of left-over data points after conditioning on the data the way we did. This low sample size led to an increased uncertainty of the estimated value of the statistic, which in turn led to larger confidence intervals.

In sum, we found that the input functions made the BOU capable of partially reproducing the nonlinear statistics, suggesting that affective stimuli can create nonlinear features in affect dynamics. However, the AIM outperformed the BOU with regard to reproducing multistability within the data, suggesting that not all nonlinearity can be attributed to the experimental stimuli.

Model Assumptions Figure 13 displays the results of the parametric bootstrap for the statistics aimed at testing the model assumptions.

Starting with the Markov assumption, one can see that the models were not sensitive to the observed value of the autocorrelation of the residuals of an AR(1) (r_ε) at all, as evidenced by the horizontal cloud of simulated values (PA: ρ_{AIM} = 0.10, ρ_{BOU} = 0.19; NA: ρ_{AIM} = 0.12, ρ_{BOU} = 0.04). Nonetheless, coverage rates were acceptable (PA: C_{AIM} = 74, C_{BOU} = 78, 95%CI $C_{BOU-AIM}$ = [−2, 10]; NA: C_{AIM} = 67, C_{BOU} = 57, 95%CI $C_{BOU-AIM}$ = [−16, −4]), suggesting that the Markov assumption was a fair assumption to take for a majority of the participants. However, if the goal is to accurately capture the affective data of all participants, one should consider the use of multiple lags.

Moving to the stationarity assumption, we found that it is often warranted. First, the standard deviation of the temporal means (SD_M) were relatively well reproduced by the models (PA: C_{AIM} = 83, C_{BOU} = 78, 95%CI $C_{BOU-AIM}$ = [−10, 0]; NA: C_{AIM} = 78, C_{BOU} = 70, 95%CI $C_{BOU-AIM}$ = [−13, −2]). This was mostly due to the stimulus dependence of these temporal changes within the mean. While similar coverage rates were found for the standard deviation of the temporal effect of the trial outcomes (SD_β; PA: C_{AIM} = 78, C_{BOU} = 76, 95%CI $C_{BOU-AIM}$ = [−6, 2]; NA: C_{AIM} = 80, C_{BOU} = 75, 95%CI $C_{BOU-AIM}$ = [−9, 0]), the models were less sensitive

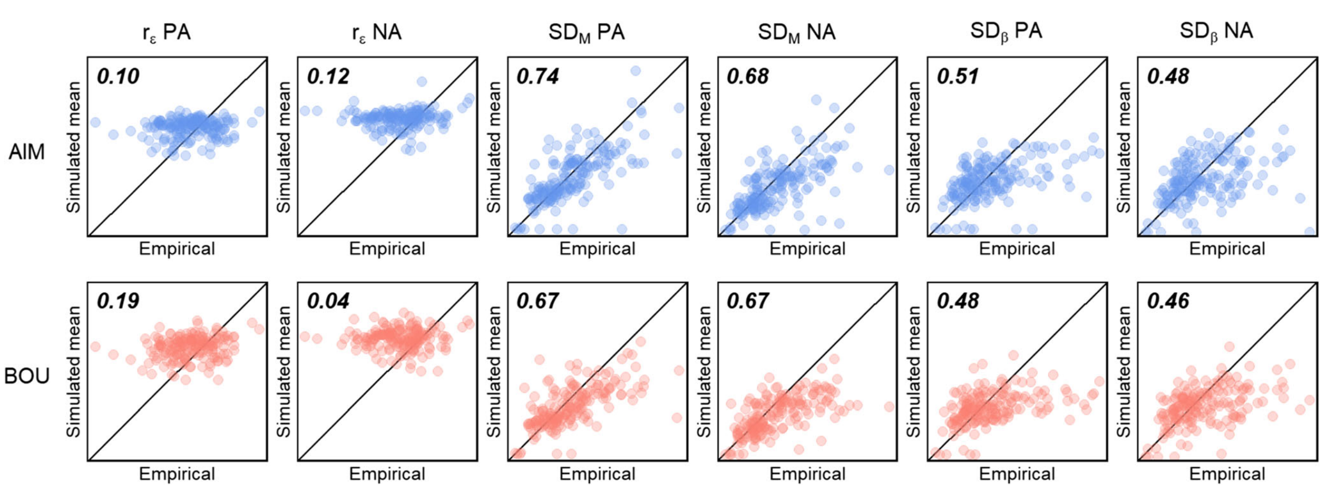

Fig. 13 Mean simulated values of the statistics aimed at testing the model assumptions plotted against the observed value. In the upper left corner, the Spearman correlation (ρ) between observed and simulated values is indicated

to the statistic's observed values (PA: ρ_{AIM} = 0.51, ρ_{BOU} = 0.48; NA: ρ_{AIM} = 0.48, ρ_{BOU} = 0.46). This slight deviation in results is probably due to the low values of this statistic for the data, indicating that an individual's affective reaction towards the trial outcomes did not really fluctuate throughout the experiment. However, when this affective reaction did change throughout the experiment, the models could not pick it up.

Discussion

In this study, we evaluated whether affect dynamics display signs of nonlinearity beyond those that are induced by affective stimuli. This question was answered by comparing the relative and absolute performance of two dynamical models — the linear BOU and the nonlinear AIM — coupled with several input functions. We found that the nonlinear AIM outperformed the linear BOU with regard to model fit and predictive performance for about half of the participants. This suggests that for these participants, nonlinearity was present in their affective time series, even after accounting for the influence of the experimental stimuli. We furthermore found that the nonlinear AIM could, in general, better reproduce the characteristics of the data. Together, the results suggest that nonlinearity may be an affective feature that cannot be fully attributed to affective stimuli, and should thus be taken into account when analyzing affective time series.

While nonlinearity was observed for one half of the sample, we found no evidence of affect-related nonlinearity for the other half. This means that some of the observed nonlinearity was indeed stimulus driven, which raises the question of just how much of this nonlinearity can be accounted for by affective stimuli. In other words, is it possible that all nonlinearity can be accounted for with a complex enough input function? We believe that there is indeed merit in investigating more complex input functions. The focus of these input functions should, however, go to identifying those characteristic features of the environment that fully or partially determine people's affective states. Inspiration for such features can be found in theoretical accounts of affect, such as appraisal theories (Scherer, 2009), goal-directed theories (Moors, 2017), and predictive processing theories (Van de Cruys, 2017; see also Moors et al., 2021). Based on these frameworks, future research may, for example, consider the expectedness of a stimulus as such a determinant, given its potential relevance in explaining affective states (Moors, 2017; Rutledge et al., 2014; Van de Cruys, 2017). Future research should thus attempt to create an input function that, while remaining sparse, accounts for most of an individual's affective environment. As a potential example of a such model, we refer the interested reader to Bennett et al. (2021).

The results of this study should be considered in light of its limitations. A first limitation involves the fact that we did not use the computational models to optimize the experiment with regard to answering our questions, that is with regard to detecting nonlinear patterns such as affective lingering. This may potentially show in two ways. First, the experiment may be limited in the extent to which temporal dependence of affect was measured and/or observed. More specifically, the constant perturbations of affect may have occluded some of the dynamical characteristics of the affective system, including temporal dependence. However, we found evidence that dynamical features of affect still played a role for the majority of participants (see Supplementary Materials), suggesting that the constant perturbation may not pose a major problem to the experiment. Second, every participant encountered the same string of outcomes. While this is in itself not necessarily problematic (see the "Trial Generation" section), using several such predetermined strings may have been more informative. Future research should thus consider the use of several strings of outcomes and compare results across these strings to adequately investigate the link between the sequence of outcomes and affect dynamics.

A second limitation involves the use of a single self-report measure of affect. While self-report measures have their strengths, they come with their own shortcomings, such as individual-specific interpretation of the scale and a risk for demand effects. In our study, individuals varied widely in their use of the ESG, with some reporting affect in an L shape, others in a diagonal line, and still others in fixed, dedicated regions of the grid. While these individual-specific patterns might reveal interesting features of one's affective system, they may also convey between-person noise. Unfortunately, it is impossible to detect such noise as the ESG is the sole measure used in this study. Future studies could benefit from the use of multiple measures of affect and from comparing results across them.

A final limitation resides in our input functions. In our analyses, we have accounted for the experimental stimuli in a highly exploratory way, as it was difficult to predict what affective stimuli any given individual would react to. Unfortunately, we were still unable to account for all possible stimulus-driven effects on affect. More specifically, we found that some of the misfit in the test for absolute model fit was attributable to the restriction that trial outcomes and trial total could not interact with each other, while such an interaction was observed for some participants. Given that this is only so for a minority of the participants, we contend that the results would not change when adjusting our input functions. This limitation does, however, underscore the importance of considering a wider array of input functions, as one may not know beforehand which affective stimuli an individual might react to.

In summary, we found evidence that some of the nonlinearity that has been observed in affective data may be due to unknown environmental influences. However, evidence also suggests that not all nonlinearity is attributable to such environmental effects and thus that nonlinearity may be an inherent feature of the affective system.

Acknowledgements We would like to thank Femke Drijkoningen and Tom Delvaux for their help with collecting the pilot data that were used to design this study. We would furthermore like to thank Stijn Verdonck for his useful suggestions regarding the analyses.

Additional Information

Funding This work was supported by the Research Fund of KU Leuven under Grant C14/19/054; and FWO under Grant G074219N. The funders had no role in study design, data collection, analysis, decision to publish or preparation of the manuscript.

Ethical Approval This study was approved by the local ethics committee (Social and Societal Ethics Committee at the KU Leuven; case number G-2020-2772-R2(MIN)). The study was performed in accordance to the ethical standards as laid out in the 1964 Declaration of Helsinki.

Conflict of Interest The authors declares no competing interests.

Informed Consent All participants had to complete an informed consent before their participation to the experiment.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s42761-022-00118-5.

Author Contributions Niels Vanhasbroeck, Sigert Ariens, and Nil Anarat created the experimental design, with valuable insights from Agnes Moors and Francis Tuerlinckx. Niels Vanhasbroeck, Tim Loossens, Nil Anarat, and Francis Tuerlinckx conceptualized the analyses, which were then carried out by Niels Vanhasbroeck. The article was written and reviewed by all authors.

Open Practice Statement This study was not preregistered. The data that were used in this article are available on OSF on the following link: https://osf.io/g6ske/.

References

Adolf, J. K., Voelkle, M. C., Brose, A., & Schmiedek, F. (2017). Capturing context-related change in emotional dynamics via fixed moderated time series analysis. *Multivariate Behavioral Research, 52*(4), 499–531. https://doi.org/10.1080/00273171.2017.1321978

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control, AC-19*(6), 716–723.

Albers, C. J., & Bringmann, L. F. (2020). Inspecting gradual and abrupt changes in emotion dynamics with the time-varying change point autoregressive model. *European Journal of Psychological Assessment, 36*(3), 492–499. https://doi.org/10.1027/1015-5759/a000589

Ariens, S., Cuelemans, E., & Adolf, J. K. (2020). Time series analysis of intensive longitudinal data in psychosomatic research: A methodological overview. *Journal of Psychosomatic Research, 137*, 110191. https://doi.org/10.1016/j.jpsychores.2020.110191

Arlot, S., & Celisse, A. (2010). A survey of cross-validation procedures for model selection. *Statistics Surveys, 4*, 40–79. https://doi.org/10.1214/09-SS054

Asutay, E., Genevsky, A., Feldman-Barrett, L., Hamilton, J. P., Slovic, P., & Västfjäll, D. (2021). Affective calculus: The construction of affect through information integration over time. *Emotion, 21*(1), 159–174. https://doi.org/10.1037/emo0000681

Asutay, E., Genevsky, A., Hamilton, P., & Västfjäll, D. (2020). Affective context and its uncertainty drive momentary affective experience. *Emotion.* https://doi.org/10.1037/emo0000912

Bennett, D., Davidson, G., & Niv, Y. (2021). A model of mood as integrated advantage. *Psychological Review.* https://doi.org/10.1037/rev0000294

Boker, S. M., & Laurenceau, J.-P. (2006). Dynamical systems modeling: An application to the regulation of intimacy and disclosure in marriage. In T. A. Walls & J. L. Schafer (Eds.), *Models for intensive longitudinal data*. Oxford University Press.

Boker, S. M., & Nesselroade, J. R. (2002). A method for modeling the intrinsic dynamics of intraindividual variability: Recovering the parameters of simulated oscillators in multi-wave panel data. *Multivariate Behavioral Research, 37*(1), 127–160. https://doi.org/10.1207/S15327906MBR3701_06

Bonsall, M. B., Wallace-Hadrill, S. M. A., Geddes, J. R., Goodwin, G. M., & Holmes, E. A. (2012). Nonlinear time-series approaches in characterizing mood stability and mood instability in bipolar disorder. *Proceedings of the Royal Society B, 279*, 916–924. https://doi.org/10.1098/rspb.2011.1246

Bringmann, L. F., Ferrer, E., Hamaker, E. L., Borsboom, D., & Tuerlinckx, F. (2018). Modeling nonstationary emotion dynamics in dyads using a time-varying vector autoregressive model. *Multivariate Behavioral Research, 53*(3), 293–314. https://doi.org/10.1080/00273171.2018.1439722

Congard, A., Dauvier, B., Antoine, P., & Gilles, P.-Y. (2011). Integrating personality, daily life events and emotion: Role of anxiety and positive affect in emotion regulation dynamics. *Journal of Research in Personality, 45*, 372–384. https://doi.org/10.1016/j.jrp.2011.04.004

Dejonckheere, E., Mestdagh, M., Houben, M., Rutten, I., Sels, L., Kuppens, P., & Tuerlinckx, F. (2019). Complex affect dynamics add limited information to the prediction of psychological well-being. *Nature Human Behaviour, 3*, 478–491. https://doi.org/10.1038/s41562-019-0555-0

Diener, E., & Iran-Nejad, A. (1986). The relationship in experience between various types of affect. *Journal of Personality and Social Psychology, 50*(5), 1031–1038.

Driver, C. C., & Voelkle, M. C. (2018). Understanding the time course of interventions with continuous-time dynamic models. In K. van Montfort, J. H. L. Oud, & M. C. Voelkle (Eds.), *Continuous Time Modeling in the Behavioral and Related Sciences*.

Ebner-Priemer, U. W., Houben, M., Santagelo, P., Kleindienst, N., Tuerlinckx, F., Oravecz, Z., et al. (2015). Unraveling affective dysregulation in borderline personality disorder: A theoretical model and observed evidence. *Journal of Abnormal Psychology, 124*(1), 186–198. https://doi.org/10.1037/abn0000021

Goldbeter, A. (2011). A model for the dynamics of bipolar disorders. *Progress in Biophysics and Molecular Biology, 105*, 119–127. https://doi.org/10.1016/j.pbiomolbio.2010.11.007

Guastello, S. J., Koopmans, M., & Pincus, D. (2009). *Chaos and complexity in psychology: The theory of nonlinear dynamical systems*. Cambridge University Press.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction*. Springer. https://doi.org/10.1007/978-0-387-84858-7

Hollenstein, T. (2015). This time, it's real: Affective flexibility, time scales, feedback loops, and the regulation of emotion. *Emotion Review, 7*(4), 308–315. https://doi.org/10.1177/1754073915590621

Holmes, E. A., Bonsall, M. B., Hales, S. A., Mitchell, H., Renner, F., Blackwell, S. E., et al. (2016). Applications of time-series analysis to mood fluctuations in bipolar disorder to promote treatment innovation: A case series. *Translational Psychiatry, 6*, e720. https://doi.org/10.1038/tp.2015.207

Houben, M., Mestdagh, M., Dejonckheere, E., Obbels, J., Sienaert, P., van Roy, J., & Kuppens, P. (2021). The statistical specificity of emotion dynamics in borderline personality disorder. *Journal of Personality Disorders, 35*(6), 819–840. https://doi.org/10.1521/pedi_2021_35_509

Kalokerinos, E. K., Murphy, S. C., Koval, P., Bailen, N. H., Crombez, G., Hollenstein, T., et al. (2020). Neuroticism may not reflect emotional variability. *Proceedings of the National Academy of Science, 117*(17), 9270–9276. https://doi.org/10.1073/pnas.1919934117

Kalokerinos, E. K., Erbas, Y., Ceulemans, E., & Kuppens, P. (2019). Differentiate to regulate: Low negative emotion differentiation is associated with ineffective use but not selection of emotion regulation strategies. *Psychological Science, 30*(6), 863–879. https://doi.org/10.1177/0956797619838763

Kuppens, P., Oravecz, Z., & Tuerlinckx, F. (2010). Feelings change: Accounting for individual differences in the temporal dynamics of affect. *Journal of Personality and Social Psychology, 99*(6), 1042–1060. https://doi.org/10.1037/a0020962

Kuppens, P., Sheeber, L. B., Yap, M. B. H., Whittle, S., Simmons, J. G., & Allen, N. B. (2012). Emotional inertia prospectively predicts the onset of depressive disorder in adolescence. *Emotion, 12*(2), 283–289. https://doi.org/10.1037/a0025046

Kuppens, P., Tuerlinckx, F., Russell, J. A., & Barrett, L. F. (2013). The relation between valence and arousal in subjective experience. *Psychological Bulletin, 139*(4), 917–940. https://doi.org/10.1037/a0030811

Kuppens, P., & Verduyn, P. (2017). Emotion dynamics. *Current Opinion in Psychology, 17*, 22–26. https://doi.org/10.1016/j.copsyc.2017.06.004

Lapate, R. C., & Heller, A. S. (2020). Context matters for affective chronometry. *Nature: Human Behaviour, 4*, 688–689. https://doi.org/10.1038/s41562-019-0555-0

Larsen, J. T., Norris, C. J., McGraw, A. P., Hawkley, L. C., & Cacioppo, J. T. (2008). The evaluative space grid: A single-item measure of positivity and negativity. *Cognition & Emotion, 23*(3), 453–480. https://doi.org/10.1080/02699930801994054

Loossens, T., Meers, K., Vanhasbroeck, N., Anarat, N., Verdonck, S., & Tuerlinckx, F. (2021). Efficient estimation of bounded gradient-drift diffusion models for affect on CPU and GPU. *Behavior Research Methods.* https://doi.org/10.3758/s13428-021-01674-7

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), e1007860. https://doi.org/10.1371/journal.pcbi.1007860

Moors, A. (2017). The integrated theory of emotional behavior follows a radically goal-directed approach. *Psychological Inquiry, 28*(1), 68–75. https://doi.org/10.1080/1047840X.2017.1275207

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. https://doi.org/10.1016/j.cobeha.2021.03.015

Myung, I. J. (2003). Tutorial on maximum likelihood estimation. *Journal of Mathematical Psychology, 47*, 90–100. https://doi.org/10.1016/S0022-2496(02)00028-7

Norris, C. J., Gollan, J., Berntson, G. G., & Cacioppo, J. T. (2010). The current status of research on the structure of evaluative space. *Biological Psychology, 84*(3), 422–436. https://doi.org/10.1016/j.biopsycho.2010.03.011

Oravecz, Z., Tuerlinckx, F., & Vandekerckhove, J. (2011). A hierarchical latent stochastic differential equation model for affective dynamics. *Psychological Methods, 16*(4), 468–490. https://doi.org/10.1037/a0024375

Palminteri, S., Wyart, V., & Koechlin, E. (2017). The importance of falsification in computational cognitive modeling. *Trends in Cognitive Sciences, 21*(6), 425–433. https://doi.org/10.1016/j.tics.2017.03.011

Roberts, D. R., Bahn, V., Ciuti, S., Boyce, M. S., Elith, J., Guillera-Arroita, G., Hauenstein, S., Lahoz-Monfort, J. J., Schröder, B., Thuiller, W., Warton, D. I., Wintle, B. A., Hartig, F., & Dormann, C. F. (2017). Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography, 40*, 913–929. https://doi.org/10.1111/ecog.02881

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. https://doi.org/10.1080/02699930902928969

Schimmack, U. (2001). Pleasure, displeasure, and mixed feelings: Are semantic opposites mutually exclusive? *Cognition and Emotion, 15*(1), 81–97. https://doi.org/10.1080/02699930126097

Sperry, S. H., Walsh, M. A., & Kwapil, T. R. (2020). Emotion dynamics concurrently and prospectively predict mood psychopathology. *Journal of Affective Disorders, 261*, 67–75. https://doi.org/10.1016/j.jad.2019.09.076

Steinacher, A., & Wright, K. A. (2013). Relating the bipolar spectrum to dysregulation of behavioural activation: A perspective from dynamical modelling. *PLoS ONE, 8*(5), e63345. https://doi.org/10.1371/journal.pone.0063345

Storn, R., & Price, K. (1997). Differential Evolution - A simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization, 11*(4), 341–359. https://doi.org/10.1023/A:1008202821328

Strogatz, S. (2018). *Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering*. CRC Press.

Trull, T. J., Lane, S. P., Koval, P., & Ebner-Priemer, U. W. (2015). Affective dynamics in psychopathology. *Emotion Review, 7*(4), 355–361. https://doi.org/10.1177/1754073915590617

Uhlenbeck, G. E., & Ornstein, L. S. (1930). On the theory of Brownian motion. *Physical Review, 46*, 823–841.

Van de Cruys, S. (2017). Affective value in the predictive mind. In Metzinger, T., & Wiese, W. (Eds.), *Philosophy and Predictive Processing.* Chapt. 24. Frankfurt am Main: MIND Group. https://doi.org/10.15502/9783958573253

Vanhasbroeck, N., Ariens, S., Tuerlinckx, F., & Loossens, T. (2021). Computational models for affect dynamics. In Waugh, C. H., & Kuppens, P. (Eds.), *Affect Dynamics.* Chapt. 10, (213-260). Cham: Springer. https://doi.org/10.1007/978-3-030-82965-0_10

Villano, W. J., Otto, A. R., Ezie, C. E. C., Gillis, R., & Heller, A. S. (2020). Temporal dynamics of real-world emotions are more strongly linked to prediction error than outcome. *Journal of Experimental Psychology: General, 149*(9), 1755–1766. https://doi.org/10.1037/xge0000740

Voelkle, M. C., & Oud, J. H. L. (2013). Continuous time modelling with individually varying time intervals for oscillating and non-oscillating processes. *British Journal of Mathematical and Statistical Psychology, 66*, 103–126. https://doi.org/10.1111/j.2044-8317.2012.02043.x

Wendt, L. P., Wright, A. G. C., Pilkonis, W. C., Denissen, J. J. A., Kühnel, A., & Zimmermann, J. (2020). Indicators of affect dynamics: Structure, reliability, and personality correlates. *European Journal of Personality, 34*, 1060–1072. https://doi.org/10.1002/per.2277