To cite this article:

Vanhasbroeck, N., Niemeijer, K., & Tuerlinckx, F. (in press). Nonlinearity in affect dynamics persists after accounting for the valence of daily-life events. *Emotion*.

Nonlinearity in affect dynamics persists after accounting for the valence of daily-life events

Niels Vanhasbroeck[1], Koen Niemeijer[1], and & Francis Tuerlinckx[1]

[1] Research Group of Quantitative Psychology and Individual Differences

KU Leuven

Leuven

Belgium

Author Note

All questions regarding this paper should be addressed to the first author of this article.

Correspondence concerning this article should be addressed to Niels Vanhasbroeck, Tiensestraat 102, 3000 Leuven. E-mail: niels.vanhasbroeck@kuleuven.be

Abstract

In recent years, increased attention has gone to studying nonlinear characteristics of affective time series. An example of such nonlinear features is multimodality – the presence of more than one mode in an affective time series –, which might mark the presence of discrete-like transitions between one and another affective state. In an attempt to capture these nonlinear features, Loossens et al. (2020) proposed the Affective Ising Model (AIM) as a model of affect dynamics. This model was validated on daily-life data, but these data did not contain any information on potential environmental factors that might have influenced a participant's affective state. Unfortunately, this omission may have lead to erronously concluding that nonlinearity is a defining characteristic of the affective system, even when it is solely driven by extrensic influences. To accommodate this limitation, we applied the AIM on daily-life data in which the valence of such external events was measured. Overall, we found that nonlinearity persisted after accounting for the valence of daily-life events, suggesting that nonlinearity is a defining characteristic of affect and should thus be accounted for. Interestingly, this effect was more pronounced for composite compared to single item measures of affect. While in line with previous research, these results should be replicated in a larger, more representative sample.

Keywords: affect dynamics, computational modeling, nonlinear, events, daily-life

Word count: article only: 7119, with appendix: 10,232

Nonlinearity in affect dynamics persists after accounting for the valence of daily-life events

Emotions are an integral part of the human experience, shaping how we perceive the world around us. In doing so, emotions do not remain stable, but change over time. These fluctuations in emotional experiences – and especially in its subjective feelings component *affect* – have been studied extensively in hopes of understanding individual differences in people's affective lives. True to this goal, *affect dynamics* have been studied in both healthy (Burns & Ma, 2015; Kalokerinos et al., 2020; Verduyn et al., 2009; Villano et al., 2020) and non-healthy populations (Bonsall et al., 2012; Hosenfeld et al., 2015; Myin-Germeys et al., 2001; Wright et al., 2015; see Trull et al., 2015 for a review) using both experimental (Asutay et al., 2020; Rutledge et al., 2014; Vanhasbroeck, Devos, et al., 2021) and daily-life methods (e.g., experience sampling method [ESM], Myin-Germeys & Kuppens, 2022).

One way to study affect dynamics is through the use of computational models (Vanhasbroeck, Ariens, et al., 2021; for examples see Bennett et al., 2022; Schweitzer & Garcia, 2010; Steinacher & Wright, 2013). In this approach, researchers construct formal statistical models that embody a theory of how the affective system works. By comparing these models with regard to their fit to the data, researchers can evaluate the viability of a set of theories in a formal way (Palminteri et al., 2017).

The extent to which a computational model performs well heavily depends on the model's ability to account for the observed data patterns. Unfortunately, most affect dynamical researchers make use of computational models that cannot account for what we will call *nonlinear* affective features. Consider, for example, the often encountered skewed distributions of affective measurements (see e.g., Heller et al., 2021; Jones et al., 2017; Larsen & Augustine, 2009) or the nonlinear, L-shaped relationship between affective

variables such as positive and negative affect (PA and NA resp., Diener & Iran-Nejad, 1986; Schimmack, 2001) or valence and arousal (Kuppens et al., 2013). Furthermore consider a nonlinear feature such as *multimodality*, which implies the existence of multiple modes within affective data (see e.g., Bonsall et al., 2012; Haslbeck et al., 2023; Hosenfeld et al., 2015). Multimodality is often associated with discrete-like jumps in one's affective state, which has been hypothesized to signal the development of psychopathology (see e.g., Borsboom, 2017; Bos et al., 2022; Schreuder et al., in press; van de Leemput et al., 2014). The prevalence of these nonlinear features in affective data calls for theoretical models that can adequately deal with such features.

As an answer to this call, Loossens et al. (2020) proposed the nonlinear Affective Ising Model (AIM) as a model for the dynamics in PA and NA. The AIM starts from a biologically inspired architecture of two pools of binary neurons, one for PA and one for NA. The activation of the neurons continually changes based on a random updating process, changing the activation level of each pool as a consequence. Despite this process being stochastic in nature, each update also accounts for several systematic forces that determine the probability with which a neuron's activity changes (Verdonck & Tuerlinckx, 2014). More specifically, active neurons within a pool will try to activate the other neurons within the same pool (excitation) and will try to decrease the neuronal activity within the other pool (inhibition). This excitation and inhibition implies that high levels of PA and NA are unlikely to co-occur, though this possibility is not excluded (Cacioppo & Berntson, 1994, 1999). Whether excitation is successful depends on the pool-specific threshold of the neurons, which determines each neuron's tendency to become active. The higher this threshold, the more difficult it will be to activate a single neuron, and by extension the pool as a whole. With this system in place, the AIM keeps track of the average activation in both pools, in effect modeling how PA and NA change over time.

Loossens et al. (2020) tested the viability of the AIM by comparing it to two related linear models, namely the Ornstein-Uhlenbeck model (OU, Driver & Voelkle, 2018; Kuppens et al., 2010) and a bounded version thereof, called the Bounded Ornstein-Uhlenbeck model (BOU). When applied to daily-life data, the AIM outperformed the OU and BOU with regard to relative fit, demonstrating the added value of accounting for nonlinear affective features in our computational models. Loossens et al. (2020) furthermore found that the AIM was better at mimicking observed linear and nonlinear data patterns, further validating the AIM as a model. Combined, these results corroborate the notion that affective time series display the aforementioned nonlinearities, and provides evidence for the nonlinear AIM as a model of affect dynamics.

While these results look promising, there was an important limitation to the study of Loossens et al. (2020): The data did not contain any information on external events that might have taken place between two consecutive measurement occasions. It is nevertheless important to recognize that such events may create nonlinear patterns within the data. To see this, we invite the reader to take a look at Figure 1 in which we illustrate three distinct multimodal affective time series. What is important for our argument is that the source of the multimodality is different within each time series.

Starting at the first time series, the person starts out with higher levels of the affective variable, but a little later this suddenly drops without any clear eliciting event. In this time series, the nonlinearity can therefore not be attributed to an external event, meaning we have to assume that this signifies a naturally occurring phase transition. Such a conclusion is not warranted for the time series in the middle, in which the transition from high to low affective states is completely dependent on an external event: The person feels low when the event happens, but immediately reverts to their original affective state once the event has passed. Given that the observed multimodality can be

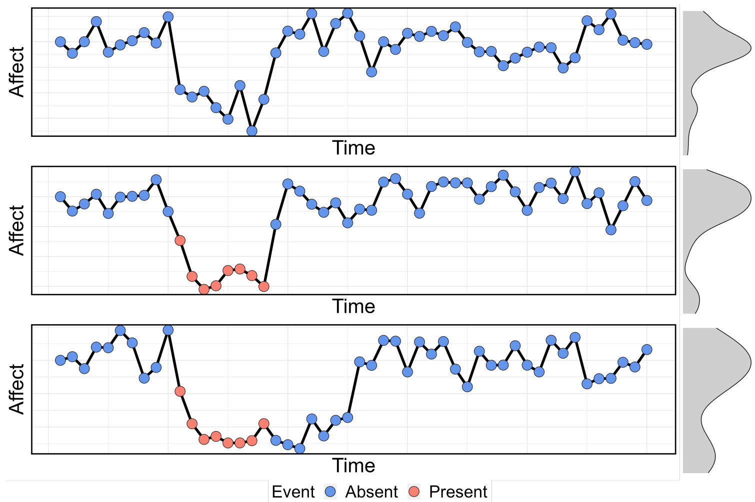

Figure 1

Visualization of multimodality in affective time series together with how the multimodality was installed. In the top panel, no known events occurred, requiring us to assume that multimodality came directly from the affective system. In the middle panel, the multimodality is completely dependent on an external event that took place. As a consequence of this dependence, the affective states recovered quickly after the event passed. The bottom panel visualizes a scenario in which the initial push towards lower values of affect was elicited by an external event, but affect remains low even when the external event has passed. While multimodality can be explained away by external events in the middle panel, this is not possible for the affective time series in the top and bottom panels.

explained away by an external event, this feature cannot be considered as evidence for the nonlinearity of the affective system. In the third time series, yet another pattern is shown. For this person, the sudden jump is elicited by an external event, but the affective state lingers on in a nonlinear way: A phenomenon that is sometimes referred to as hysteresis (Borsboom, 2017). Importantly, while the event initiated the lower affective state, its presence alone cannot fully capture this lowered affective state. Because multimodality cannot be completely explained away by the external events, we thus again

have a case of nonlinearity residing within the affective system.

Nonlinear features might thus arise both as a feature of affect or as a consequence of affective events. Unfortunately, the study of Loossens et al. (2020) did not allow us to distinguish between both kinds of nonlinearity. It is therefore unclear to which extent the nonlinear features that the AIM describes can be attributed to such external events.

In a first attempt to alleviate this difficulty, Vanhasbroeck et al. (2022) did an experimental study in which participants were repeatedly exposed to positive and negative decision outcomes. Using the affective time series that came from this procedure, the authors could compare stimulus-informed version of the BOU and the AIM in a similar way to Loossens et al. (2020). Results of this study were in favor of the AIM, corroborating the notion that affect evolves in a nonlinear fashion over time.

Given that the study of Vanhasbroeck et al. (2022) is a laboratory study, one might question its generalizability to daily life (see e.g., Koval et al., 2015). With this study, we attempt to answer this critique. In particular, we investigated the extent to which nonlinearity in daily-life affective fluctuations can be explained away by the affective events that happen to an individual. Like earlier, we will use a model-based approach, comparing informed versions of the BOU and AIM on both relative and absolute fit to the data.

Method

This study relied on secondary data from an ESM study conducted in Belgium between June 2021 and December 2021. In this study, Niemeijer et al. (2023) set out to evaluate the performance of the mobile sensing component of their ESM app (m-Path: https://m-path.io/). While they also gathered self-report data the authors did not examine these data yet. As the data were not primarily gathered for the purpose of this paper, the sample characteristics and known demographics are dependent on the choices of Niemeijer et al. (2023).

Participants

In their study, Niemeijer et al. (2023) recruited 462 participants from various Facebook groups and from an experiment recruitment ' system affiliated with the University of Leuven. The selection criteria excluded non-native Dutch speakers, those below the age of 18, and those with phones that did not meet the technical requirements of the ESM app that was used for data collection. In total, 104 participants were chosen based on their availability throughout the study period and their level of neuroticism as measured by the relevant Big Five Inventory-2 items (Denissen et al., 2008; Soto & John, 2017). The latter condition consisted of choosing individuals to maximize the variety of neuroticism scores in the sample and was imposed to ensure adequate diversity in emotional functioning.

Of the in total $21,851$ measurement occasions that were send out to the participants' phones – henceforth referred to as *beeps* –, $19,133$ were filled out. Of these, 189 values were found to be missing due to technical issues and 2 instances of duplicate responses were removed. We furthermore excluded 310 item responses with extremely short (less than 300 msec) and 67 beeps with extremely long response times (longer than 15 minutes) from the analysis (Delespaul, 1995).

For the current study, four participants were excluded from the dataset because they were recruited in a pilot phase in which data about external events was not yet collected. The final dataset thus consisted of 100 participants, 88 of whom identified as 'female', 11 as 'male', and 1 as 'other'. The average age of the participants was 21 years ($SD = 4$), with the youngest and oldest participant being 18 and 52 years old, respectively.

Materials

The ESM questionnaire used in the study of Niemeijer et al. (2023) included several questions aimed at assessing positive and negative emotions as well as contextual events.[1] Specifically, participants were prompted to rate their general level of PA and NA ("To what extent do you currently experience [positive/negative] feelings?") as well as specific emotions in the format "How [emotion] do you feel right now?", each question with a slider scale ranging from 0 ("not at all") to 100 ("very"). The assessed emotions were the current level of happiness, relaxation, anxiety, sadness, stress, loneliness, tiredness, and feelings of being down. For our analysis, we used the single items as well as a derived composite measure for PA and NA. This composite consisted of an average of the responses on happiness and relaxation for PA, and of the responses on anxiety, sadness, stress, loneliness, tiredness, and feeling down for NA. We furthermore transformed these composite measures to fall between 0 and 1, as the models assume this range for the affective data.

We included the composite measures of PA and NA for two reasons. First, we wanted to minimize the role of measurement error on our conclusions by aggregating over multiple items. Additionally, we wanted to use one of the standard approaches taken in the field to increase generalizability to other ESM studies (for similar approaches, see e.g. Dejonckheere et al., 2021; Sperry et al., 2020; Villano et al., 2020).

[1] In this study, mobile sensing data was also collected. However, we don't use these mobile sensing data in this analysis, and we will thus not describe them any further (see Niemeijer et al., 2023 for more information).

Table 1

Average within-person descriptive statistics for the variables used in this study. The descriptive variables range from 0 to 100.

Variable	*M*	*SD*	Min	Max	Skew	Kurtosis
Positive Affect (Single-item)	54.32	17.27	5.50	89.93	-0.70	1.50
Negative Affect (Single-item)	24.26	17.47	1.40	83.06	1.44	3.99
Positive Affect (Composite)	57.62	15.14	13.08	88.52	-0.57	0.40
Negative Affect (Composite)	25.10	11.01	5.32	62.72	0.97	2.02
Positive Context	25.46	21.93	0.77	87.62	1.24	2.97
Negative Context	12.84	15.58	0.27	82.69	2.70	9.79

The ESM questionnaire also included two context items that measure the valence of daily-life events. More specifically, we asked "To what extent has something [positive/negative] happened since the last beep?" to which participants had to respond by using another slider scale ranging from 0 ("nothing happened") to 100 ("something very [positive/negative] happened"). Like for the PA and NA composite scores, the responses on the context items were also transformed to fall within the $[0, 1]$ range. **Importantly, these context items measure the valence of daily-life events, but not the events that happened themselves. We will therefore refer to the variable measured by these context items as such in the remainder of this text.**

Tables 1 and 2 report the descriptive statistics and within-person correlations of the variables used in our reanalysis. Significance was corrected for multiple testing using Holm's method (Holm, 1979). One might notice a few things in these tables. First, there is an adequate amount of within-individual variation in all affect variables and the context

Table 2

Correlations between the variables used in this study.All correlations in this table were corrected for multiple testing.

Variable	2.	3.	4.	5.	6.
1. Positive Affect (Single item)	-.50*** [-.51,-.49]	.65*** [.64, .66]	-.54*** [-.55,-.53]	.39*** [.38, .41]	-.31*** [-.32,-.29]
2. Negative Affect (Single item)	-	-.56*** [-.57,-.55]	.68*** [.68, .69]	-.23*** [-.24,-.21]	.47*** [.46, .48]
3. Positive Affect (Composite)		-	-.61*** [-.62,-.60]	.35*** [.34, .37]	-.36*** [-.37,-.34]
4. Negative Affect (Composite)			-	-.27*** [-.29,-.26]	.41*** [.39, .42]
5. Positive Context				-	-.13*** [-.14,-.11]
6. Negative Context					-

Note. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

items, suggesting that both affect and the immediate environment of the participants changed over time. Additionally, and as expected, some of the variation in affect was related to external events, as the responses on the context items correlated significantly with both single item and composite PA and NA. Finally, there was considerable positive skew in the data, and more specifically in the NA scores and the context items. This result might be explained by a tendency of the responses for these variables to be low, close to the lower bound of 0.

Procedure

In the original study, participants were prompted to answer questions about their current affective state and the intensity of positive and negative events within a specified time window by means of push notifications from a mobile app. Participants received 10 notifications per day for 21 days throughout a 12-hour time window that they selected beforehand. Notifications were scheduled in 10 blocks of 72 minutes each and participants had 30 minutes to respond to the beep before it expired. If they did not respond within the first 5 minutes, they were sent a reminder notification. The interval between two consecutive measurements varied randomly, with a minimum of 15 minutes between consecutive notifications. There were, however, 7 exceptions to this rule, either because of a delayed or a duplicate transmission of beeps due to technical issues.

Participants received compensation for their time, either through university course credits or monetary compensation of up to 10 credits or €70, depending on their compliance rate. Participants received full remuneration if they completed at least 75% of the beeps. Once below this 75% threshold, each 10% loss in compliance led to a 1 credit or €10 reduction in their compensation.

After receiving information about the study protocol, remuneration, and obligations and benefits of participation, all participants granted informed consent. They were informed that their participation was entirely voluntary and that they could opt out at any moment. To protect privacy and confidentiality, the data were pseudonymized, and personal information was kept separate from study data. The study was approved by KU Leuven's Social and Societal Ethics Committee (G-2020-2200-R3[AMD]) and was conducted in accordance with ethical principles for human subject research.

Computational models

To answer our research question, we compared several computational models of affect dynamics. In this section, we will succinctly explain these models and refer the interested reader to the *Appendix* for a technical description of the models.

Each model consisted of two building blocks, namely a dynamical model and an input function. Of these, the former formalized how affect changes over time and the latter described how affect is related to the valence of external events. Both building blocks are now attended to more thoroughly.

Dynamical models

In this study, we compared two dynamical models, namely the linear BOU and the nonlinear AIM. Both are drift-diffusion models, that is, they are stochastic continuous-time models that describe affective fluctuations as they evolve continuously over time (Gardiner, 1985; Strogatz, 2018). To describe affect dynamics, these models thus not only take the previously observed affective state into account, but also the time that has passed since this observation. The models furthermore expect affect to evolve towards one or more baselines or *attractors*, which can be interpreted as the affective state that comes naturally to the person, or to which affect gets regulated after a perturbation. Both models are furthermore stochastic in nature, implying that variability around the attractor(s) is to be expected, even in the absence of known external influences.

To explain the fundamentals of both dynamical models, we use the following analogy. Imagine a skater on a ramp, and futher assume that at any time we are tracking the position of this skater at the ground level. Intuitively, one expects that the skater moves towards the lowest point of the ramp, where they come to a halt. Imagine however that some people around the ramp continually push the skater in random directions, so that the skater keeps moving and never quite reaches this minimum.

This analogy is meant to visualize how a typical drift-diffusion model works: In the context of affect dynamics, the skater's location on the ramp represents one's affective state, which is subject to two forces, namely (a) a regulatory pull of the affective state towards its baseline and (b) random perturbations to the affective state due to the inherent stochasticity of the system. Together, these influences keep affect evolving in a continuous manner over time. The skater analogy also helps in interpreting the two models of interest, as the primary difference between these models is the shape of their ramp (or surface, see Figure 2). Next, we will explain the peculiarities of both models and show the difference between both.

Bounded Ornstein-Uhlenbeck Model. The BOU is a linear model of affect dynamics that uses a paraboloid surface to capture affective fluctuations, the exact form of which can vary between being circular or elliptic (depending on the relationship between PA and NA). Consequently, this model has only one baseline, or only one minimum to which the skater can move. This is in very close correspondence to typically used

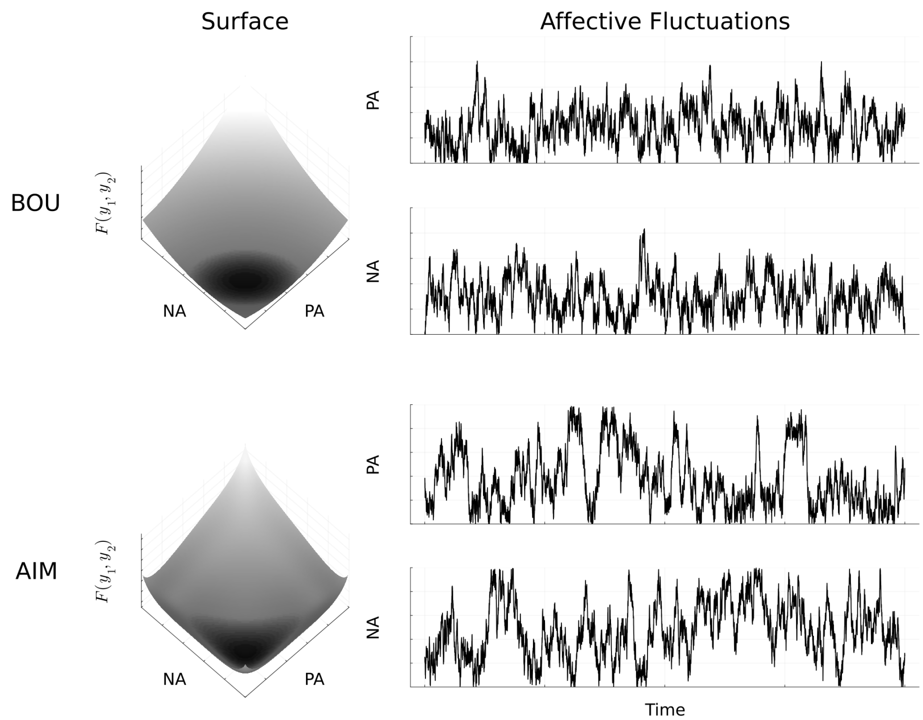

Figure 2

Visualization of the affective surfaces and continuous evolution of affect over time for the BOU (top) and the AIM (bottom). Note that the affective surface of the BOU has a paraboloid shape, so that affective fluctuations are always regulated back towards one baseline. This is not the case for the AIM, which has a nonlinearly shaped surface, leading to a nonlinear time series. In this figure, this nonlinearity leads to affect remaining stable in locations other than its primary baseline.

discrete-time linear models that describe affect dynamics in terms of a single baseline to which affect gets regulated (e.g., the vector autoregressive model [VAR], Hamilton, 1994).

As mentioned before, the BOU is a bounded version of the linear OU, so that PA and NA can only take on values that fall within 0 and 1. This restriction was imposed to make the comparison between the OU and the AIM more fair: As an unbounded model, the OU runs the risk of generating predictions that fall outside of the data range, which disadvantages it against a model that is shielded from making such nonsensical predictions.

To make sure that (non)linearity lies at the basis of the difference in performance between the models, we used the BOU rather than the more often used OU (e.g., Driver & Voelkle, 2018; Oravecz et al., 2011).

Importantly, these bounds make the BOU capable of accommodating skew in our measurements, a feat that the unrestricted OU cannot pull off. Given that the BOU cannot capture any other nonlinear features in the data and given that it still predicts linear drift of affect towards its single attractor, we will refer to the BOU as being a linear model.

Affective Ising Model. In contrast to the BOU, the AIM describes affective fluctuations using a more intricate shape than a paraboloid. This shape can accommodate many forms allowing it to capture the nonlinear phenomena of interest, namely the nonlinear relationship between PA and NA and multimodality in the PA-NA space (see again Figure 2). The L-shaped relationship between PA and NA is immediately introduced through the creation of nonlinear shapes in the surface, such that high PA states are unlikely to occur with high NA states. Multimodality is introduced through the creation of multiple local attractors to which affect can be regulated. This allows the model to capture discrete-like changes between different affective states, for example jumping from a generally neutral state (low PA and low NA) to generally negative states (low PA and high NA; see Figure 2).

With regard to the latter characteristic, it is important to note that the AIM expects these discrete jumps to be driven by randomness. More specifically, random perturbations may cause an individual to jump from one attractor to another. To be able to capture non-stochastic discrete changes, one is in need of the second building block: The input function.

Input functions

External events can influence a participant's affective system at any time. The extent to which an external event does so is determined by a mathematical function that

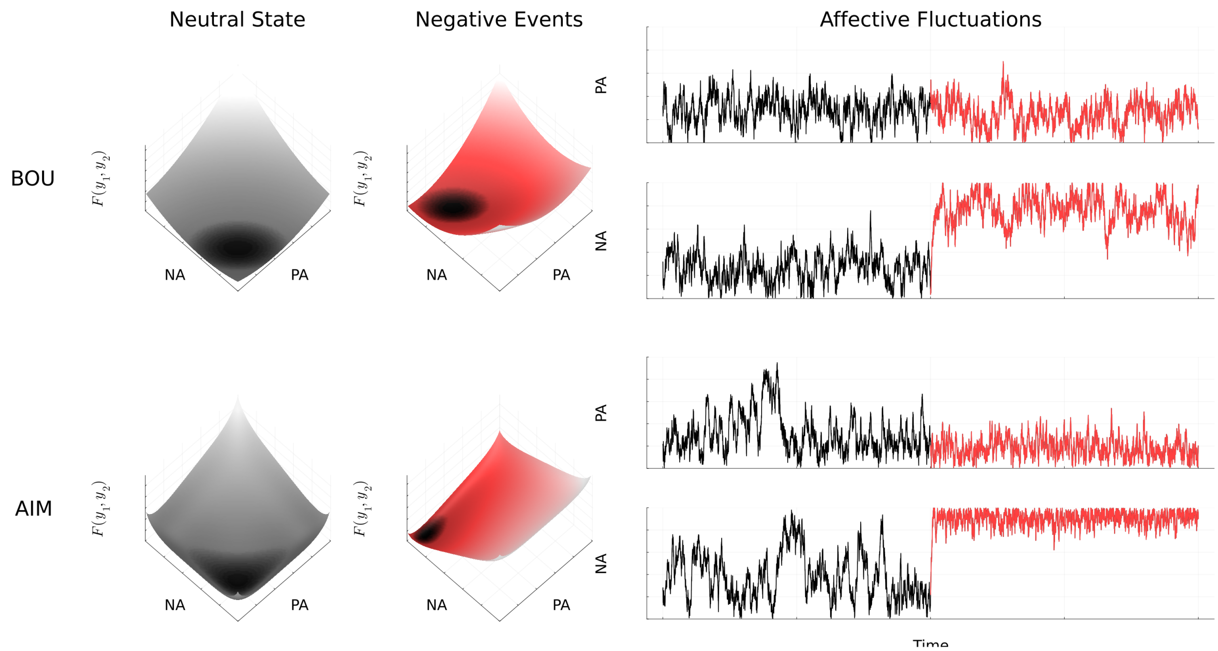

Figure 3

Visualization of input-related changes in the dynamical models in case of a negative event (red indicates the start of the event). In response to the event, both models shift their expected affective state towards higher values of NA. It needs to be noted that this change is not always as abrupt as in this figure: Changes can also be more gradual. The abrupt change here is thus illustrative

we denote as the input function (see also Vanhasbroeck et al., 2022). This input function takes in whatever happens to an individual and translates it to quantitative input to the models. This input then leads to a change in the surface of the dynamical models so that an affective reaction to the event is allowed (see Figure 3). The extent of the change in the dynamical models in response to an event is determined by the value of the input function and is estimated together with the other model parameters.

Exactly how the input changes the surfaces of the two models differs between the BOU and the AIM. For the BOU, an external event moves the surface to another location, leading to a direct relocation of the attractor or baseline. This is different from the AIM, for which an external event tilts the surface in a specific direction, leading to an increased or decreased elevation of the local minima in the surface. Importantly, this difference between the models makes the BOU more versatile in its accommodation of affective events. More concretely, the BOU is free to put its attractor anywhere it desires. The AIM does not enjoy such freedom and is therefore limited in exactly where it wants to have its attractors in response to the input (Loossens, 2021).

As summarized in Table 3, we defined three input functions based on the responses on the two context items described above. The first input function will be referred to as *Empty*, as it assumes that external events have no influence on one's affect. The second function is the *Main* input function, which assumes that positive events have an influence on PA and negative events have an influence on NA, but there are no cross-effects of the two kinds of events on the two kinds of affect. For the BOU, this restriction means that the attractor can only move in the PA- or the NA-direction in response to positive and negative events respectively. The restriction is similar for the AIM, but in this case the surface can only be tilted in one or the other direction. One might consider this assumption to be too restrictive, which is why we also included an *Interaction* input function in which cross-effects are allowed.

Combining the dynamical models with the input functions allows us to investigate both the influence of the valence of external events on and the dynamical structure of affect. Both parts are thus indispensable for answering our question. To ease comparison of the models, we will refer to the computational models with a label that combines the dynamical and input model. For example, pairing the AIM with the *Empty* input function will lead to the computational model *AIM Empty*.

Statistical analyses

All statistical analyses build upon the estimated parameters for the computational models, and we will therefore briefly describe the estimation procedure. Individual-specific

Table 3

Definition of the input functions for changes in the PA- and NA-direction. In this table, P_t refers to a participant's answer on the positive context item at time t and N_t refers to the response on the negative context item. In the Event column, one can see that the input functions only kick in the moment that the positive and/or negative context items P_t and N_t have a value greater than 0, indicating that an external factor has influenced the participant's affective state at that time t. The values of the parameters γ determine the extent to which the attractor is relocated (for BOU) or the extent to which the surface is tilted (for the AIM) and are estimated for each dynamical model within each individual.

Name	Event	Influence on PA	Influence on NA
Empty		0	0
Main	$P_t > 0$	$\gamma_1 P_t$	0
	$N_t > 0$	0	$\gamma_2 N_t$
Interaction	$P_t > 0$	$\gamma_1 P_t$	$\gamma_2 P_t$
	$N_t > 0$	$\gamma_3 N_t$	$\gamma_4 N_t$

parameters were estimated for each of the possible combinations between dynamical model and input function through the GradientDiffusion package in Julia (Bezanson et al., 2017; Loossens et al., 2021). This package uses a global minimizer to perform parameters estimation through the simultaneous minimization of the min-log-likelihood for several parameter sets (see *Supporting Information* to Loossens et al., 2020). Using this approach, we optimized 100 parameter sets in 5000 iterations for all models. To ensure convergence, we furthermore repeated this procedure 5 times and selected the best-fitting parameter set to be used in the following steps of our analyses.

Based on the estimated parameters, we conducted two kinds of analyses. First, we compared the models in terms of fit to the data, providing us with an indication of relative model performance. While useful, this comparison does not tell us how well the model fits to the data in an absolute way (Palminteri et al., 2017). Therefore, we also included a parametric bootstrap procedure to check how well the models could mimic the real data. Each of these analyses will now be handled in more detail.[2]

Relative model fit

As a measure of the relative fit of the models to the data, we used the Akaike's Information Criterion (AIC, Akaike, 1974) and the Bayesian Information Criterion (BIC, Schwarz, 1978). Both measures balance a model's fit and complexity by penalizing the fit according to the number of parameters in the model. As a result, we could compare model performance with a lower risk of overfitting the data.

Absolute model fit

One of the goals of computational modeling is to capture data patterns as closely as possible. In this study, a parametric bootstrap procedure was used to assess the models' ability to do so. As visualized in Figure 4, this procedure consisted of (a) identifying an individual's best-fitting BOU and AIM by using the AIC as criterion, (b) simulating 1000 datasets for each model, and (c) computing several statistics of interest on these simulated data. By computing the same statistics on the observed data, we were able to check the correspondence between the simulated and observed values of the statistics, giving an idea of whether the models were able to mimic the data.

Correspondence between simulated and observed values of the statistics was assessed in two ways. First, we used the Spearman correlation ρ between the average simulated and the observed values of the statistics to assess the sensitivity of the models to the data characteristics, with higher values of ρ indicating higher sensitivity. We preferred the Spearman ρ over the Pearson correlation coefficient for the formers ability to handle

[2] Code for all analyses, including the estimation, is available on Gitlab through the following link: https://gitlab.kuleuven.be/ppw-okpiv/researchers/u0123135/daily-events.

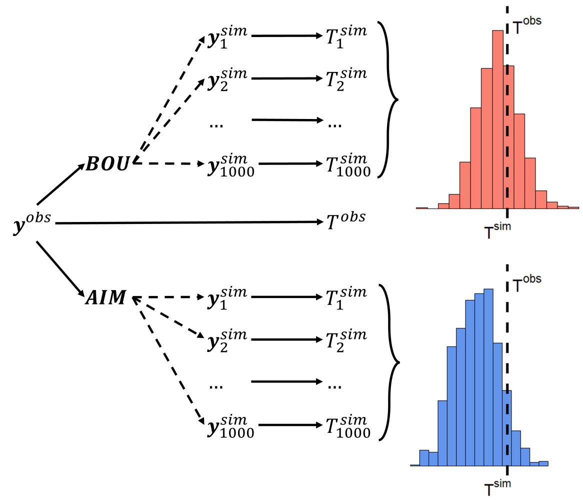

Figure 4

Visualization of the parametric bootstrap procedure for the data of a single participant. First, we estimated the parameters of the BOU and the AIM based on $\boldsymbol{y}^{obs}$*. Next, we simulated 1000 new datasets* $\boldsymbol{y}_i^{sim}$*, for each of which we computed the value of a given statistic* T *(*T_i^{sim}*). This procedure led to two distributions of* T^{sim} *to which we compared the observed value of* T *(*T^{obs}*; black vertical line).*

outliers in the data, allowing us to interpret its value even if the computed statistics deviate strongly for a single individual. While the Spearman ρ quantifies the rank-order association between simulated and observed statistical values, it cannot tell us how well the models picked up on the exact observed statistical values. To assess this, we computed the proportion of participants for whom the observed value of the statistic fell within the simulated 95% confidence interval (CI) of the same statistic. This will be referred to as the *coverage rate* in the remainder of the article. For well-fitting models, we expect the coverage rate to be around 95%.

All that's left for us now is to specify the statistics we used in the parametric bootstrap procedure.[3] First, we used the autocorrelation of PA and

[3] While we only report the results of three statistics in this text, we used several other statistics in the parametric bootstrap. The interested reader can find a complete treatment

NA (r_{auto}) as a measure of the temporal relationship between current and previous affective states. This statistic thus quantifies the dynamical component of affect, which is central to both the BOU and the AIM. We furthermore assessed the models' ability to capture two previously mentioned nonlinear features, them being multimodality and the L-shaped relationship between PA and NA. Of these, the latter was assessed by the curvature between the two affective variables ($\kappa_{\mathbf{PANA}}$), while the former was captured by the bimodality coefficient (BC; see the *Appendix* for a formal introduction to both statistics).

Especially the last two statistics are of interest, as these are directly relevant to our research question. To see this, consider the earlier finding that the BOU is incapable of capturing these nonlinear features (Loossens et al., 2020). Despite this inability, coupling the BOU with an input function could enable it to mimic such features (see again Figure 1). In such a case, comparing the BOU with the AIM could determine to which extent an observed nonlinear feature can be explained away by the valence of external events. More specifically, we conclude that a nonlinear feature is context-dependent if the BOU captures the corresponding statistic to the same extent as the AIM. If on the other hand the AIM still outperforms the BOU, this suggests that the observed nonlinearity cannot be entirely explained away by the immediate context.

Transparency and openness

This study was not preregistered. The data are only available on request, as they are still being used in another study. The code to analyze these data is available online through the following link:

---

of these in the *Appendix*.

https://gitlab.kuleuven.be/ppw-okpiv/researchers/u0123135/daily-events

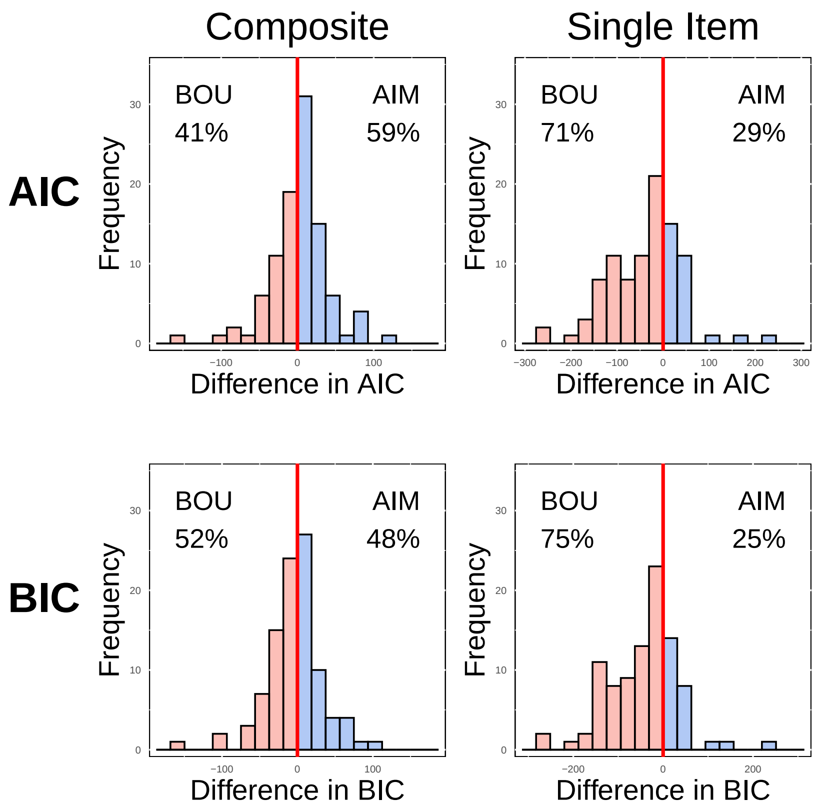

Figure 5

Visualization of the AIC- and BIC-related results (rows) for the composite and single item measures of PA and NA (columns). Each of the plots shows the within-person difference in fit of the best-fitting AIM versus the best-fitting BOU. Bars of the histogram that fall to the left of the vertical line show a preference for the BOU, and vice versa for bars that fall to the right of this line.

Results

Relative model fit

The results of the relative model comparison are displayed in Figure 5. **This plot shows the within-person differences in AIC and BIC between the best-performing BOU and AIM. We furthermore distinguish between the results for the composite and the single item measures of PA and NA.**

The results of the AIC and BIC conveyed a similar story and will

therefore be discussed together. For the composite measures we found that for 59% (AIC) and 48% (BIC) of the participants an AIM provided the best fit. Surprisingly, we found an opposite pattern for the single item measures, such that for 71% (AIC) and 75% (BIC) of the participants not the AIM, but the BOU provided the better fit to the data. **Together, these results suggest that nonlinear features in the data were still picked up by the models, even after accounting for the valence of external events. However, the extent to which this is the case depends on how affect is measured; Nonlinearity is more often picked up when affect is measured with multiple items as opposed to one single item. This observation can be explained in several ways, as we will show in the *Discussion*.**

Absolute model fit

In Figure 6, the observed values of the statistics of interest are plotted against the simulated values of the same statistics for the BOU and AIM separately. The results associated to each of these statistics will be described in detail below.

First, we found that both the BOU and the AIM were sensitive to the value of the autocorrelation to a similar extent, and that this sensitivity was lower for the single item measures of PA and NA (Composite: Spearman $\rho \in [0.83, 0.94]$; Single item: Spearman $\rho \in [0.57, 0.63]$). With regard to coverage rates, both models recovered the autocorrelation in composite PA scores outstandingly well (BOU: coverage $= 93\%$; AIM: coverage $= 89\%$) but consistently underestimated the value for the composite NA scores (BOU: coverage $= 29\%$; AIM: coverage $= 34\%$). It is not clear why this difference between PA and NA was observed, but it might be related to the low variation in the composite NA scores compared to the PA scores. For the single item measures, we found an adequate performance of both models, though less ideal

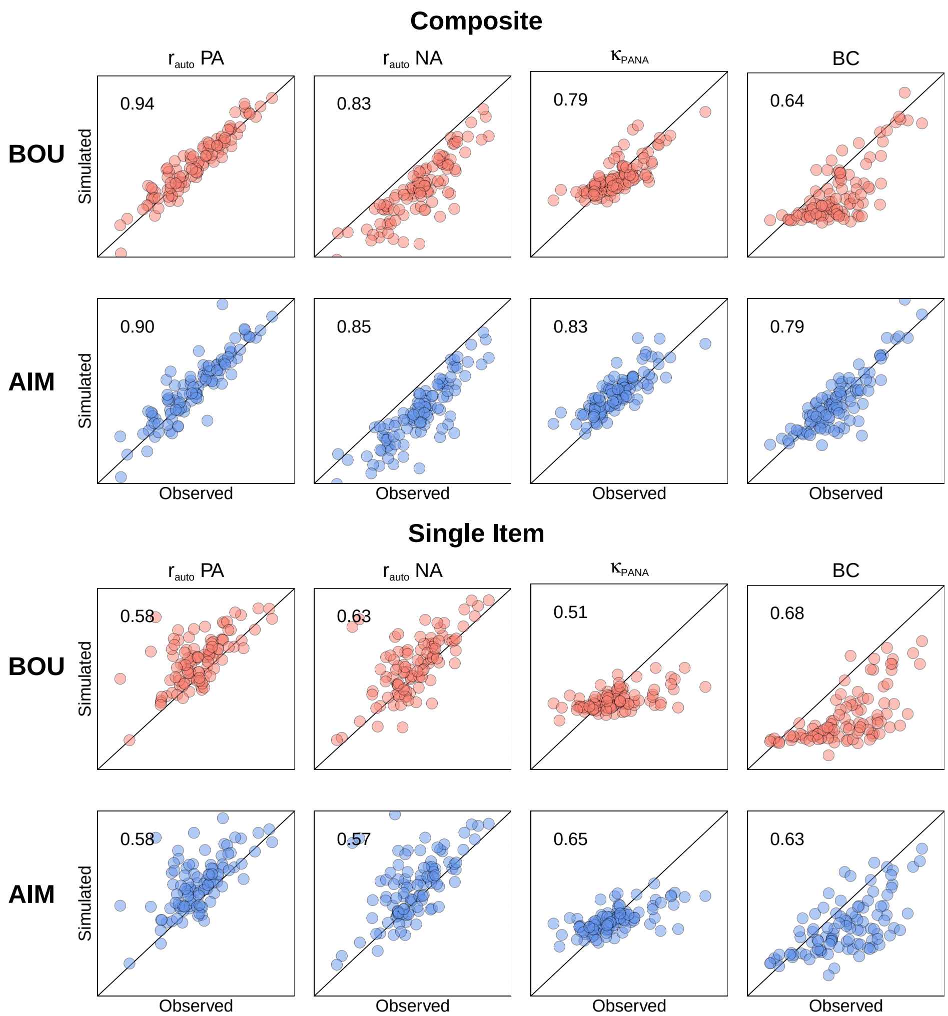

Figure 6

Visualization of the results of the parametric bootstrap. Observed values of the statistics (i.e., computed on the observed data for a participant) are plotted against simulated values of the same statistic (i.e., computed on simulated data for a participant). Data were simulated using the best-performing BOU and AIM (rows) for each participant's composite and single item PA and NA scores (top and bottom resp.). In the top-left corner of each plot, the Spearman correlation coefficient is given.

than for the composite scores (PA: coverage $\in$ [51%, 55%]; NA: coverage $\in$ [58%, 60%]). Contrary to the composite results, we did not find a difference between PA and NA for the single item measurements.

Moving on to the L-shaped relationship between PA and NA, we found that both the BOU and the AIM were adequately sensitive to its observed value (BOU: Spearman $\rho \in [0.51, 0.79]$; AIM: Spearman $\rho \in [0.65, 0.83]$). With regard to coverage, we did not find a difference in performance between the models, but rather found one between the composite measures (BOU: coverage $= 78\%$; AIM: coverage $= 72\%$) and the single items (BOU: coverage $= 38\%$; AIM: coverage $= 41\%$). Both models thus covered this statistic to a similar extent, which is in line with previous results (Vanhasbroeck et al., 2022). This suggests that the BOU is able to reproduce the nonlinear relationship between PA and NA to a similar extent as the AIM. In other words, whenever the AIM outperformed the BOU, it was not due to the AIMs ability to naturally account for the L-shaped relationship between PA and NA, but due to something else.

This "something else" might be bimodality, for which we used the BC as a statistic. Results indicate that the AIM (Composite: Spearman $\rho = 0.79$, coverage $= 79\%$; Single item: Spearman $\rho = 0.63$, coverage $= 26\%$) captured this statistic much better than the BOU (Composite: Spearman $\rho = 0.64$, coverage $= 45\%$; Single item: Spearman $\rho = 0.68$, coverage $= 15\%$). This suggests that the AIM could mimic multimodal data patterns better than its competitor, even if the latter is supplied with information on the valence of external events. In combination with the results of $\kappa_{\mathbf{PANA}}$, we can therefore hypothesize that when the AIM outperformed the BOU, this was most likely related to the AIM's ability and the BOU's inability to capture input-unrelated bimodalities in affect dynamics. Importantly, this conclusion is only warranted for the composite measures of affect, as both models performed equally well on the measurements from the single items.[4]

[4] This result also suggests that there is some multimodality left in the affective time series that is not attributable to the external events. This in turn might indicate indicate the presence of hysteresis in our

In sum, while the AIM and BOU seem to capture the dynamical features of and nonlinear relationship between PA and NA equally well, the AIM outperformed the BOU with regard to capturing multimodality in the data. Together, these results suggests that curvature, but not multimodality, could be mostly explained by the valence of external events. We furthermore found that the models did a better job of mimicking the composite data rather than the affective data from the single item measures, an observation to which we will return in the *Discussion*. Importantly, we add that there is still some ground to cover: Coverage rates are not what one would expect them to be, indicating that there is still some room for improvement.

data. Despite our interest in these kinds of patterns, we were not able to directly test the presence of hysteresis in an objective, statistical way.

Discussion

In this study, we followed up on our previous efforts and investigated the presence of nonlinearity in daily-life affect dynamics (Loossens et al., 2020; Vanhasbroeck et al., 2022). More specifically, we asked to which extent nonlinearity in affect dynamics persisted after accounting for the valence of external events. To investigate our question, we compared how well input-sensitive versions of the linear BOU and the nonlinear AIM fit data that was gathered in daily life. Our study's evidence for the persistence of nonlinearity is mixed. For composite measures of affect, we found that the AIM outperformed the BOU with regard to both relative and absolute fit for a slight majority of the participants. In contrast, the BOU outperformed the AIM with regard to relative fit and both models performed to a similar extent with regard to absolute fit for the single item measures of affect.

These results seem to be inconsistent and require an explanation. A first potential explanation might be that our two measures of PA and NA are not actually measuring the same construct.[5] While the single items explicitly and directly measure positivity and negativity, this is not the case for the composite scores, which consist of an average of several emotion items. Concretely, it is possible that an average of, say, *happiness* and *relaxation* does not measure PA to the same extent as the single item for positivity does (for a similar perspective, see Cloos et al., 2023). From this, one can then conclude that the extent to which nonlinearity is present in the data depends on the construct that you are measuring.

While we do not dismiss this possibility, we do note that it leaves some observations unexplained. For one, this explanation is difficult to reconcile with

[5] We thank our reviewers for bringing up this possibility.

our experimental results, which indicated the presence of nonlinear features in single item measurements of PA and NA (Vanhasbroeck et al., 2022). Additionally, both models experienced great difficulty in reproducing the characteristics of the single item data compared to the composite data. This observation was consistent across all statistics (cfr. *Appendix*), meaning that neither the AIM, nor the BOU captured the data adequately. So while it is possible that the composite and single item measures of PA and NA do not measure the same thing, it does not necessarily explain our current observations.

A more likely culprit for the different results for the composite and single item measures is measurement error. Single item measurements are more susceptible to noise, which in turn might occlude the characteristic features of affect dynamics, both linear and nonlinear. Composite measures, on the other hand, are less susceptible to such error and might therefore reveal such characteristics more successfully. Especially the results of the parametric bootstrap seem to corroborate this explanation, as neither of the two models captured the data well. Such a finding closely fits to what one would expect if measurement error prevails in one's measurements.

Unfortunately, however, we are not able to formally assess either of the explanations given, and are at this point left at speculating about the reason for our contradictory results. Nevertheless, we can cautiously conclude that for about 25% to 59% of the participants in our sample nonlinearity persisted even after accounting for the valence of external events. Consistent with our results of the parametric bootstrap, we furthermore conclude that primarily multimodality explained the difference in performance between the two models.

Implications

The results of this study are in line with our previous results, in which we found nonlinearity in the daily-life and experimental affect dynamics of several individuals (Loossens et al., 2020; Vanhasbroeck et al., 2022). Taken together, our results show that we should take nonlinearity of affect more seriously. Consider, for example, that the traditional models used in the field of affect dynamics are linear in nature. If we realize that computational models are a lens through which we look at and understand our data, then the linear model lens is not focused enough to detect subtle data patterns that are the result of nonlinear processes. These nonlinear features represent a very unique and specific signature of the underlying affective mechanism and are therefore worth investigating.

Think about multimodality, for example, which might signal the presence of hysteresis (see again Figure 1). Classically speaking, hysteresis occurs when a change in a given variable elicits a discrete-like jump towards another state, but reversing this change does not necessarily lead to a discrete jump back. This is a phenomenon that we are very familiar with. Think about the anger that you feel when a biker almost runs you over, which does not immediately wane when the biker disappears into the distance. Similarly but on a different time scale, once pushed in a depressive state, it is difficult to get out. While it is easy enough to come up with variables that might create hysteresis in affective states, it is more difficult to come up with ways to counteract its effects. This is an area that future research needs to put some more attention to, as it has implications for clinical interventions.

This illustration shows the importance of investigating nonlinear affective features more thoroughly. For this, we need nonlinear computational models

which, ideally, represent formal theories and which allow us to design thoroughly informed studies for answering our research questions (Palminteri et al., 2017). Only then might we better understand how such features come about.

Limitations and Constraints on Generality

The results of this study should be considered in light of its limitations. **First, this study is self-report heavy, in that participants were asked to rate both their affective states and the valence of the external events happened since the last beep. While we contend that self-report measures have several strengths, they also come with their own limitations and biases. For example, we have no information on the thought-process of participants behind their responses on the questionnaires. This limitation is of special relevance to the context items, as it is inclear what the measured events were and whether participants considered multiple events or just a single event when filling out the items. It might therefore be worthwhile to replicate these findings with not only self-report, but also other measures of affect and external events.**

A second limitation pertains to our explanations of the deviation in our results between the composite and single item measures of affect. These explanations highlight little knowledge about the validity and the reliability of our measures, both of which might explain the diverging results between the two kinds of measures. Future research could greatly benefit from investigating the psychometric properties of the measures employed here and in other daily-life studies.

A final limitation of this study pertains to the sample that was studied. More specifically, this study made use of a convenience sample, the participants mostly being affiliated to the University of Leuven. The sample furthermore consisted of mostly women, with few men and even fewer individuals who identified as neither. These limitations to the sample restrict the generalizability of our results to other, more diverse populations. These results should thus be replicated in a bigger, more representative sample.

Conclusion

To conclude, we found some evidence for the persistence of nonlinear features of affect dynamics after accounting for the valence of daily-life events. Given some of the limitations of this study, however, more research is necessary to give a more definite idea of the extent to which these nonlinear features are present in daily-life affective data and how they come about.

Additional Information

Author contribution. KN conceptualized the original study from which we used the data here, with a small contribution of NV. NV conceptualized and performed the analyses with valuable insight of FT. NV, KN, and FT all wrote and reviewed the article.

Conflict of interest. The authors declare no conflict of interest.

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

Ethical approval. As stated in the article, this study was approved by the local ethics committee at the Psychological department of the KU Leuven (the Social and Societal Ethics Committee) under case number G-2020-2200-R3[AMD]. The study was performed in accordance with the ethical standards as laid out in the 1964 Declaration of Helsinki.

Transparency and openness. This study was not preregistered. The data are only available on request, as they are still being used in another study. The code to analyze these data is available online through the following link: https://gitlab.kuleuven.be/ppw-okpiv/researchers/u0123135/daily-events

Acknowledgements. The analyses performed in this work were performed using resources and services of the VSC (Flemish Supercomputer Center), funded by the Research Foundation - Flanders (FWO) and the Flemish Government.

Appendix

Appendix A: Formal description of the models

Here, we discuss the models used in the article more technically. While we always discussed the BOU before the AIM in the article, we will switch the order here. This is because many of the inner workings of the AIM can be readily translated to the BOU, but the story becomes more difficult the other way around.

Affective Ising Model

In the article, we used the analogy of a skater to explain how the AIM works. Here, we will provide some more detail on the mathematical workings of the AIM and how it relates to this analogy.

First and foremost, the analogy mentioned a ramp which guided the skaters movements. This ramp is actually a potential function, referred to as the *free energy function* (Loossens et al., 2020), which is defined as follows:

$$\begin{aligned} F(y_1, y_2) = &\sum_{i=1}^{2} \Big(-\lambda_i y_i^2 + (\theta_i - f_i(x))\, y_i\Big) + \lambda_{12} y_1 y_2 + \\ &\sum_{i=1}^{2} \nu_i \left(y_i \ln(y_i) + (1 - y_i) \ln(1 - y_i)\right), \end{aligned} \tag{A1}$$

where $i \in \{1, 2\}$ determines the pool to which the parameters refer, so that y_1 and y_2 denote PA and NA respectively. The shape of the free energy function is determined by these parameters, which are all related to the forces described in the article (see Verdonck & Tuerlinckx, 2014 for a more elaborate explanation). More specifically, λ_i is the parameter related to the within-pool excitation of pool i, while λ_{12} relates to the between-pool inhibition. The thresholds of the neurons are also pool-specific and are captured by the parameter θ_i. Finally, the ν_i parameter scales the entropy of the system and is related to the size of the pools.

Now that we have defined the ramp, we should still describe the movement of the skater. This is achieved through the definition of the following two stochastic differential equations:

$$\begin{aligned} dy_1(t) &= -\delta \frac{\partial F(y_1(t), y_2(t))}{\partial y_1} dt + \sqrt{2\delta}\ dW_1(t) \\ dy_2(t) &= -\delta \frac{\partial F(y_1(t), y_2(t))}{\partial y_2} dt + \sqrt{2\delta}\ dW_2(t). \end{aligned} \tag{A2}$$

The first term on the right is called the *drift term* and guides the movement of the skater towards the bottom of the ramp. As one can see, this term relies on the free energy function. Indeed, it is the drift term that captures the regulatory movement of affect towards (one of) the baselines within the system. The second term is called the *diffusion term* and captures the perturbations to the skater on his way to the bottom of the ramp. This term describes the stochasticity of the affective state and keeps affect moving continuously over time. For the AIM, it consists of two independent scaled Wiener processes $dW_i(t)$.

Note that the parameter δ – referred to as the *diffusion constant* – controls the pace at which affect moves across the free energy function, thus governing the inertia of the system. Importantly, it scales both the drift and diffusion term, which is contrary to the BOU (and by extension the OU). In other words, it fixes the relative influence of the regulatory and stochastic processes that describe affect dynamics.

The solution of Equation A2 is a probability distribution for affect within the PA-NA space. This distribution is conditional on the previously observed affective state and the time that has passed since its measurement. However, after an adequately long time between observations, one might assume that both conditions (previous affect and time that has passed) do not relate to the observed affective state anymore. If, for example, one goes to sleep at night, the affective state at wake-up might not relate to the affective state of the previous night. In this case, we define the *equilibrium probability* *distribution* as the affective probability distribution that is not conditional on previous affect nor the time that has passed since this previous observation. It is defined as:

$$p(y_1, y_2) = \frac{e^{-F(y_1,y_2)}}{\mathcal{Z}},$$

where $\mathcal{Z}$ is a normalizing constant. Note that Equation A2 diffuses towards this equilibrium probability distribution, so that if time between observations Δt becomes very large, the conditional probability distribution approximates the equilibrium probability distribution. More formally:

$$lim_{\Delta t \to \infty}(p(y_1(\Delta t), y_2(\Delta t))) = p(y_1, y_2).$$

Bounded Ornstein-Uhlenbeck model

The BOU is a bounded version of the OU model, a linear continuous-time drift-diffusion model that has seen use in the literature (e.g., Driver & Voelkle, 2018; Kuppens et al., 2010). The OU is defined as:

$$d\boldsymbol{y}(t) = \Theta \left(\boldsymbol{\mu} - \boldsymbol{f}(\boldsymbol{x}(t)) - \boldsymbol{y}(t)\right) dt + \Gamma d\boldsymbol{W}(t), \tag{A3}$$

where $\boldsymbol{y}(t) = [y_1(t), y_2(t)]^T$ is a vector containing the affective state at time t. In this equation, drift is again contained in the first term on the right while diffusion is described by the second term. Within the OU, the drift towards the baseline $\boldsymbol{\mu}$ is defined by the drift matrix Θ and the diffusion consists of a Wiener process scales by the lower-triangular matrix Γ.

The BOU makes use of the same principles as the OU, defining drift and diffusion in the way described in Equation A3. However, the key difference between both models is their domain: While the OU is defined on the real domain, the BOU limits this domain to the AIMs ($y \in (0, 1)$). This difference in domain is installed to make comparison of the

AIMs dynamics to the dynamics of the OU fairer: The OU would almost definitely place probability mass in regions that fall outside of the PA-NA space, giving it a disadvantage in the model comparison.

The way in which the BOU limits its domain is by imposing reflecting bounds on the boundaries of the PA-NA space. This means that whenever probability falls outside of this space, it is reflected back within the grid, making the BOU capable of creating skewed distributions.

Input functions

With the formal description of both dynamical models out of the way, this only leaves the input functions to explain. We refer back to Equations A1 and A3, where one might see the function $\boldsymbol{f}$ that takes in argument $\boldsymbol{x}(t)$. This is the input function $\boldsymbol{f}$ that defines the effect of the input $\boldsymbol{x}(t)$ on affect $\boldsymbol{y}(t)$ at time t.

The output of this function – previously referred to as B (Loossens et al., 2020) and $\beta(t)$ (Supplementary Materials, Vanhasbroeck et al., 2022) – consists of two values that depend on both the input and the estimated parameters of $\boldsymbol{f}$. These values are subsequently used to alter some of the parameters in the models (see Equations A1 and A3). In the AIM, the input function changes the values of θ_1 and θ_2, which leads to the input-dependent tilting of the free energy function. For the BOU, the input function has a more direct effect, as it changes the values of the baseline $\boldsymbol{\mu}$, creating an input-dependent baseline to which affect will be regulated.

This difference in influence of the input function has two related consequences. First, the BOU is more flexible than the AIM in its accommodation of the input. This is so because the AIM has to first define a “neutral” free energy function which can then be tilted to one or another direction to accommodate input-related data. This means that the AIM can only indirectly influence the location of its baselines. The BOU has no such restriction, in that it can place its baseline anywhere within and outside of the PA-NA space.

Second, for the input-sensitive AIM the difference between θ_i and $f_i(x(t))$ is a positive number, or $\theta_i - f_i(x(t)) > 0$. This is a restriction inherent to the AIM (see Loossens et al., 2020) which implies a restriction on the values $f_i(x(t))$ can take. Such a restriction is, again, not present in the BOU, so that the BOU has another input-related advantage over the AIM.

Appendix B: Measuring curvature and bimodality

In the article, we use two statistics to measure the extent to which PA and NA are nonlinearly related to each other and the extent to which there is multimodality in our affective data. These statistics will be formally defined here.

Curvature

In a nutshell, we estimated the nonlinear relationship between PA and NA by first rotating the data by 45 degrees, then estimating a parabola on these rotated data, and finally computing the maximal curvature of this parabola as a measure of the curved relationship between the affective variables. This procedure is visualized in Figure A1.

Our procedure requires that we define the curvature of the parabola in a mathematical way. To do so, we first define the general formula for curvature of parametric curves as (see Stewart, 2008):

$$\kappa(x) = \frac{\left|\dfrac{d^2 f(x)}{dx}\right|}{\left(1 + \left(\dfrac{dy}{dx}\right)^2\right)^{\frac{3}{2}}}, \tag{A4}$$

where $|.|$ denotes the absolute value and $\frac{df(x)}{dx}$ and $\frac{d^2 f(x)}{dx}$ are the first and second derivative of $f(x)$ with regards to x. Defining the parabola as:

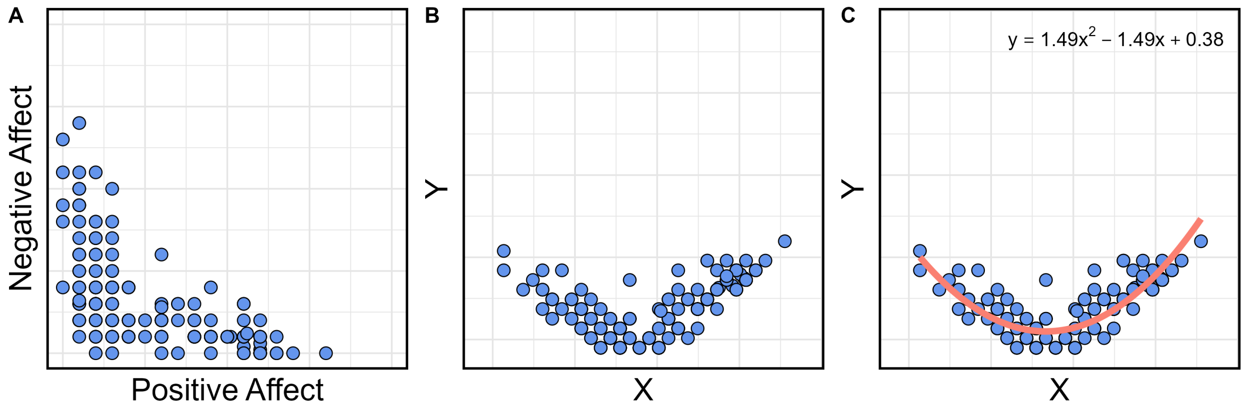

Figure A1

Visualization of the procedure for estimating the nonlinear relationship between PA and NA. We begin in Panel A, where the raw data of one of our participants is visualized. In Panel B, we rotate these data by 45 degrees and then fit a parabola on these rotated data in Panel C. The parameters of this parabola are then used to estimate the curvature.

$$f(x) = ax^2 + bx + c,$$

we can compute these two derivatives as:

$$\frac{df(x)}{dx} = 2ax + b$$
$$\frac{d^2f(x)}{dx} = 2a,$$

which we then fill out in Equation A4 to get:

$$\kappa(x) = \frac{|2a|}{(1 + (2ax + b)^2)^{\frac{3}{2}}}, \tag{A5}$$

which is the formal definition of the curvature of a parabola.

One may notice, however, that this definition still depends on the value of x. For our purposes, we are interested in that value of x for which $\kappa(x)$ is maximal, so that we measure the maximal amount of curvature in the relationship between PA and NA.

To find this value, we first have to find the local extrema of κ. These extrema can be found by finding the zero(s) of its derivative. The derivative of $\kappa(x)$ is:

$$\frac{d\kappa(x)}{dx} = -\frac{12a|a|(2ax+b)}{(1+(2ax+b)^2)^{\frac{5}{2}}}$$

Given that the denominator is always greater than 0, this function is continuous in x. We are then left with finding the zero(s) of the numerator, which are equal to:

$$\begin{aligned} 2ax + b &= 0 \\ 2ax &= -b \\ x &= -\frac{b}{2a}, \end{aligned}$$

which coincidentally is the top of the parabola and to which we will refer as T.

To find out whether this extremum is a minimum or a maximum, we compute the value for $\kappa(x)$ for T as well as T plus or minus an arbitrary small positive number ϵ. Filling out these values in Equation A5, we find:

$$\begin{aligned} x &= -\frac{b}{2a} &&\Rightarrow \quad \kappa(x) = |2a| \\ x &= -\frac{b}{2a} \pm \epsilon &&\Rightarrow \quad \kappa(x) = \frac{|2a|}{(1+\epsilon^2)^{\frac{3}{2}}}. \end{aligned}$$

Given that ϵ^2 is always a positive number, we know that $|2a| > \frac{|2a|}{(1+\epsilon^2)^{\frac{3}{2}}}$, which proves that for T $\kappa(x)$ reaches its maximal value. We therefore use T to compute κ in our analyses and consequently define κ as:

$$\kappa = |2a|$$

Using this definition, we are only able to measure the value of the curvature of the parabola, but are not able to measure its direction. More specifically, using the current definition of κ, we are not able to say whether the L-shape is in the expected direction – that is, so that high PA and high NA values are not likely to co-occur – or in the opposite direction – that is, so that such mixed feelings are very likely to occur. While the latter option is unlikely to be found, we did not want to restrict our analyses. We therefore used a version of κ in which the absolute value is ommitted:

$$\kappa = 2a$$

Importantly, using the top T as the value for x on which we compute κ is still warranted, as it will always yield the maximal possible value in the positive or negative direction.

Multimodality

In our article, we used the bimodality coefficient (BC) to measure multimodality in the data. This statistic is defined as (Pfister et al., 2013):

$$BC = \frac{m_3^2 + 1}{m_4 + 3\dfrac{(n-1)^2}{(n-2)(n-3)}}, \tag{A6}$$

where m_3 and m_4 are the skew and excess kurtosis of a single variable. Importantly, both moments should be corrected for sample bias. We define g and k as the sample skew and kurtosis:

$$g = \frac{\frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x})^3}{s^{\frac{3}{2}}}$$
$$k = \frac{\frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x})^4}{s^2},$$

where s is the sample corrected variance, defined as:

$$s = \frac{1}{n-1}\sum_{i=1}^{n}(x_i - \bar{x})^2.$$

Using these definitions, we can then define m_3 and m_4 as:

$$
\begin{aligned}
m_3 &= \frac{n^2}{(n-1)(n-2)}g \\
m_4 &= \frac{(n+1)n}{(n-1)(n-2)(n-3)}k - 3\frac{(n-1)^2}{(n-2)(n-3)}.
\end{aligned}
$$

Given that the BC is only defined for one dimension, we computed its value for all relevant dimensions within the PA-NA space. These dimensions corresponded to the axes and the diagonals, defined as PA, NA, PA + NA, and PA - NA respectively. We then selected the highest BC as a measure of multimodality in the data.

Appendix C: Additional results

Within the article we focused on the results that were directly applicable to our research question. Given that these are important with respect to validating the AIM as a model, we discuss them here.

Statistics

In this section, we will describe the statistics we used in our parametric bootstrap but that were not mentioned in the article. We will divide these statistics in two groups, based on the purpose for which we included the statistics.

The first group consists of statistics that assessed the models' ability to recover both linear and nonlinear indices that are frequently used in affective research. This group includes the following statistics:

(a) Correlation between PA and NA (r_{PANA}), which measures the linear relationship between PA and NA;

(b) R^2 of a linear regression of the events on affect, which measures the extent to which affective changes were related to these external events. This was again assessed for PA and NA separately;

(c) Skew, which was computed for PA and NA separately.

The second group of statistics were aimed at testing some of the implicit assumptions of both dynamical models, namely the Markovian and stationarity assumptions. In the context of our models, the Markovian assumption dictates that current affect should only depend on its previous state, and not on any lags prior to the previous one. More formally, it states that affect at time t only depends on the affective state at time $t-1$, but not on the affective state on time $t-2$, $t-3$, etc. In our data, this assumption holds if only the previous reported affective state carries information with respect to the current affective state. A violation of this assumption indicates misspecification of the model, which might come from a number of sources. For example, there might be some input variables that we are not accounting for, but that have an important influence on affect (e.g., cognitive processes like rumination).

With the stationarity assumption, one assumes that the affect dynamics of an individual do not change over (the studied) time. Within a model-based framework, this comes down to assuming that the parameters of the model should remain the same over time. Within this study, it is conceivable that participant's affective lives underwent minimal changes throughout the limited time that this study spans, thereby justifying the stationarity assumption. A violation of this assumption might indicate that the way in which affect evolves does not remain constant over time, so that models should allow parameters to change over time (e.g., Albers & Bringmann, 2020; Bringmann et al., 2018).

Given that an argument can be made for the validity of both assumptions, many researchers employ statistical models that imply these assumptions for the analysis of affective data (e.g., VAR, Hamilton, 1994). Regrettably, though, both assumptions are seldomly verified, thereby warranting an examination of their plausibility. Importantly, the degree to which these assumptions hold are also important for substantive research, as it may indicate the presence of internal or external processes that are as yet not accounted for within our models.

Both assumptions were tested with the following statistics:

(a) Residual autocorrelation from an AR(1) model ($r_{residual}$), which measures the left-over autocorrelation after accounting for both external events and lag-1 dynamics of affect. This statistic was aimed at testing the Markovian assumption, as the residuals of the specified AR(1) model should be independent over time. This statistic was assessed for PA and NA separately;

(b) Variability of a statistic T across a moving window (SD_T), which measures how much the value of a statistic changes over time. This statistic was aimed at testing the stationarity assumption, as variability across windows should be low. To estimate this statistic, we computed the values of T for several windows of 50 data points that moved over the complete data (each time moving by 1 datapoint). As T, we used the mean of PA and NA (separately) and the R^2 that was defined above.

Results

Statistical indices. In Figure A2, the observed values of the first group of statistics are plotted against the simulated values of the same statistics for the BOU and AIM separately.

Starting with the correlation between PA and NA, this statistic was captured excellently by both the BOU and the AIM, and for both the composite and single item measures of affect. This is evident in both the correlations (Composite: Spearman $\rho \in [0.93, 0.97]$; Single item: Spearman $\rho \in [0.71, 0.73]$) and the coverage rates (BOU: coverage = 96% for the composite measures, 40% for the single items; AIM: coverage = 83% for the composite measures, 35% for the single items). Note that the coverage rates were much lower for the single items than for the composite scores. This corroborates our in-text observation that both models had a harder time capturing the observed features for the single item measurements.

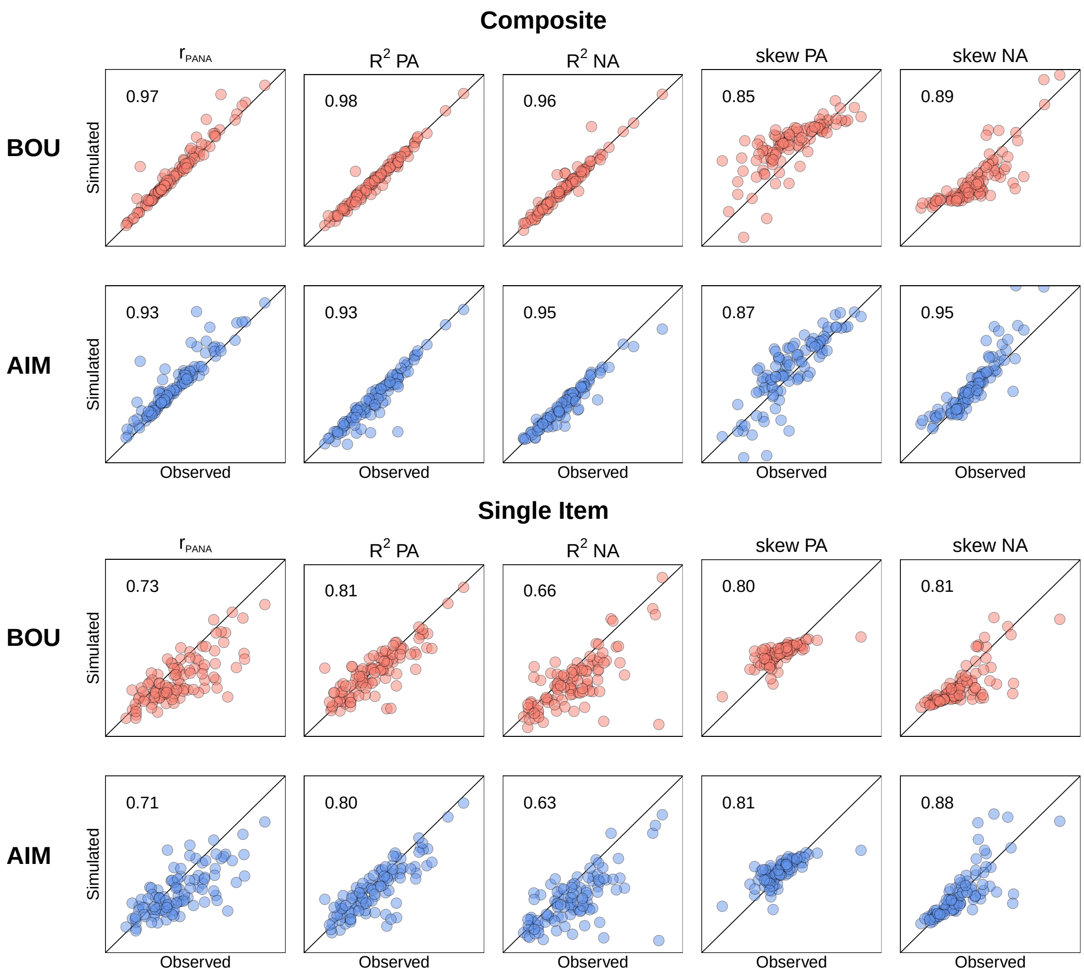

Figure A2

Visualization of the results of the parametric bootstrap for the first group of statistics. Like in the article, observed values of the statistics are plotted against simulated values of the same statistic. Data were simulated using the best-performing BOU and AIM (rows) for each participant's composite and single item PA and NA scores (top and bottom resp.). In the top-left corner of each plot, the Spearman correlation coefficient is given.

Moving on to the R^2, the AIM and the BOU again show excellent performance, both with regard to sensitivity (Composite: Spearman $\rho \in [0.93, 0.98]$; Single item: Spearman $\rho \in [0.63, 0.81]$) and with regard to coverage (BOU: coverage = [98%, 100%] for the composite measures, [58%, 72%] for the single items; AIM: coverage = [90%, 94%] for the composite measures, [52%, 68%] for the single items). We again found a difference in performance for the composite measures compared to the single item measures of affect, showing that the models were not able to capture the affective features well based on only single item measurements.

Finally turning our attention to the skew in PA and NA, we found that the AIM captured this feature of the data better than the BOU, especially with regard to coverage. More specificially, we found that the AIM showed a good sensitivity (Composite PA: Spearman $\rho = 0.87$; Single item PA: Spearman $\rho = 0.81$; Composite NA: Spearman $\rho = 0.95$; Single item NA: Spearman $\rho = 0.88$) and good coverage of the skew (Composite PA: coverage = 61%; Single item PA: coverage = 35%; Composite NA: coverage = 81%; Single item NA: coverage = 58%). This was especially the case for the composite measures and was, once again, attenuated for the single item measures. The BOU did not perform that well, showing a decreased success in covering this statistic (Composite PA: coverage = 53%; Single item PA: coverage = 30%; Composite NA: coverage = 46%; Single item NA: coverage = 44%) despite being sensitive to it (Composite PA: Spearman $\rho = 0.85$; Single item PA: Spearman $\rho = 0.80$; Composite NA: Spearman $\rho = 0.89$; Single item NA: Spearman $\rho = 0.81$). This result suggests that the AIMs natural way of incorporating skew is preferred over imposing bounds on an otherwise skew-less distribution (as is the case for the BOU).

Model assumptions. Figure A3 displays the observed values of the statistics aimed at testing the Markovian and stationarity assumptions, as explained earlier. Results are in line with our previously reported results (see Vanhasbroeck et al., 2022).

First, the models were insensitive to the residual autocorrelation (Spearman $\rho \in [-0.06, 0.10]$) but its values were extremely well recovered (coverage $\in [88\%, 96\%]$). In model-related terms, this indicates that the Markovian assumption seemed to be violated in our data, but that this violation did not pose too much trouble for our models due to a large uncertainty around the statistic. This violation might however signal that another unaccounted for process influenced affect.

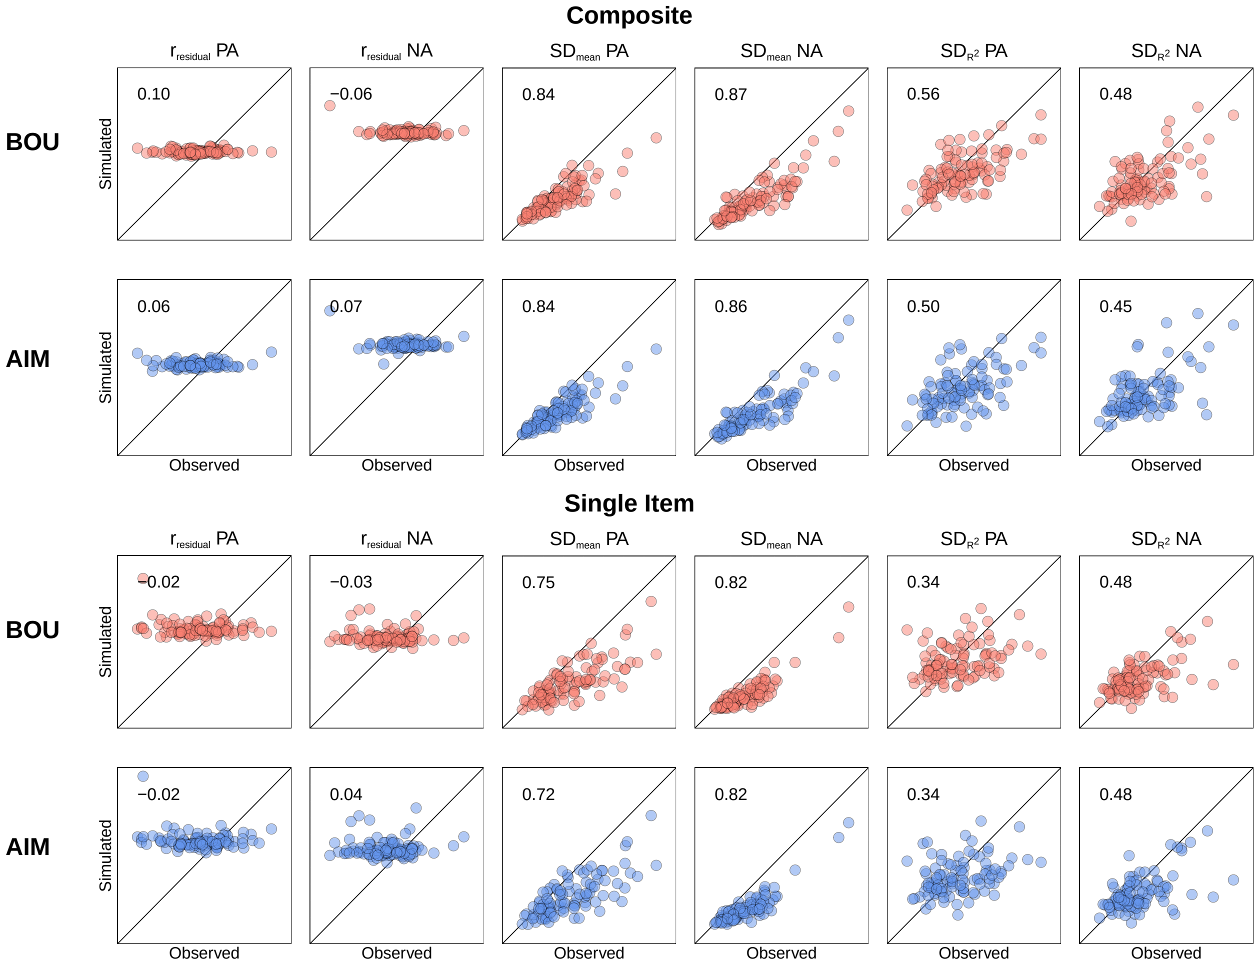

Figure A3

Visualization of the results of the parametric bootstrap for the statistics aimed at testing the model assumptions in the same way as for the linear statistics. In the top-left corner of each plot, the Spearman correlation coefficient is given.

Results were more mixed for the stationarity assumption: When examining variability in the mean, both models were sensitive to its observed value (Spearman $\rho \in [0.72, 0.87]$) but the coverage rates were not always as good (coverage $\in$ [18%, 63%]). This suggests that the models were partially able to account for changes in the mean over time, probably through accounting for the valence of external events. However, there were still some changes in these means that the models could not account for, suggesting that our models are missing another influence on affect.

Focusing now on variability in the extent to which events influenced affect, we find the opposite pattern. More specifically, both models were relatively insensitive to the observed value of the statistic (Spearman $\rho \in [0.34, 0.56]$) but covered its value quite well (coverage $\in [71\%, 92\%]$). This result suggests that while the models were unable to account for nonstationarity, this should not be taken as a definite sign of model misfit, given the large degree of uncertainty in the statistics. However, the source of nonstationarity should be better understood if one wishes to capture affect dynamics to the fullest.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *AC-19*, 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*, 492–499. https://doi.org/10.1027/1015-5759/a000589

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. (2022). A model of mood as integrated advantage. *Psychological Review*, *129*, 513–541. https://doi.org/10.1037/rev0000294

Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. (2017). Julia: A fresh approach to numerical computing. *SIAM Review*, *59*(1), 65–98. 10.1137/141000671

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

Borsboom, D. (2017). A network theory of mental disorders. *World Psychiatry*, *16*, 5–13. https://doi.org/10.1002/wps.20375

Bos, F. M., Schreuder, M. J., George, S. V., Doornbos, B., Bruggeman, R., van der Krieke, L., Haarman, B. C. M., Wichers, M., & Snippe, E. (2022). Anticipating manic and depressive transitions in patients with bipolar disorder using early warning signals. *International Journal of Bipolar Disorders*, *10*. https://doi.org/10.1186/s40345-022-00258-4

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*, 293–314. https://doi.org/10.1080/00273171.2018.1439722

Burns, R. A., & Ma, J. (2015). Examining the association between psychological wellbeing with daily and intra-individual variation in subjective wellbeing. *Personality and Individual Differences*, *82*, 34–39. https://doi.org/10.1016/j.paid.2015.02.023

Cacioppo, J. T., & Berntson, G. G. (1994). Relationship between attitudes and evaluative space: A critical review, with emphasis on the separability of positive and negative substrates. *Psychological Bulletin*, *115*, 401–423.

Cacioppo, J. T., & Berntson, G. G. (1999). The affect system: Architecture and operating characteristics. *Current Directions in Psychological Science*, *8*, 133–137.

Cloos, L., Ceulemans, E., & Kuppens, P. (2023). Development, validation, and comparison of self-report measures for positive and negative affect in intensive longitudinal research. *Psychological Assessment*, *35*, 189–204. https://doi.org/1037/pas0001200

Dejonckheere, E., Mestdagh, M., Verdonck, S., Lafit, G., Ceulemans, E., B., B., & Kalokerinos, E. K. (2021). The relation between positive and negative affect becomes more negative in response to personally relevant events. *Emotion*, *21*, 326–336. https://doi.org/10.1037/emo0000697

Delespaul, P. A. E. G. (1995). *Assessing schizophrenia in daily life : The experience sampling method* [{PhD} {Thesis}, Maastricht University]. https://doi.org/10.26481/dis.19950504pd

Denissen, J. J. A., Geenen, R., van Aken, M. A. G., Gosling, S. D., & Potter, J. (2008). Development and Validation of a Dutch Translation of the Big Five Inventory (BFI). *Journal of Personality Assessment*, *90*(2), 152–157. https://doi.org/10.1080/00223890701845229

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

Gardiner, C. (1985). *Handbook of stochastic methods, for physics, chemistry and the natural sciences* (2nd ed.). Springer-Verlag.

Hamilton, J. D. (1994). *Time series analysis.* Princeton University Press. https://doi.org/10.2307/j.ctv14jx6sm

Haslbeck, J., Ryan, O., & Dablander, F. (2023). Multimodality and skew in emotion time series. *Emotion.*

Heller, A. S., Stamatis, C. A., Puccetti, N. A., & Timpano, K. R. (2021). The distribution of daily affect distinguishes internalizing and externalizing spectra and subfactors. *Journal of Abnormal Psychology*, *130*, 319–332. https://doi.org/10.1037/abn0000670

Holm, S. (1979). A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics*, *6*, 65–70. http://www.jstor.org/stable/4615733

Hosenfeld, B., Bos, E. H., Wardenaar, K. J., Conradi, H. J., van der Maas, H. L. J., Visser, I., & de Jonge, P. (2015). Major depressive disorder as a nonlinear dynamic system: Bimodality in the frequency distribution of depressive disorder over time. *BMC Psychiatry*, *15.* https://doi.org/10.1186/s12888-015-0596-5

Jones, A. B., Brown, N. A., Serfass, D. G., & Sherman, R. A. (2017). Personality and density distributions of behavior, emotions, and situations. *Journal of Research in Personality*, *69*, 225–236. https://doi.org/10.1016/j.jrp.2016.10.006

Kalokerinos, E. K., Murphy, S. C., Koval, P., Bailen, N. H., G., C., Hollenstein, T., Gleeson, J., Thompson, R. J., Van Ryckeghem, D. M. L., Kuppens, P., & Bastian, B. (2020). Neuroticism may not reflect emotional variability. *Proceedings of the National Academy of Science*, *117*, 9270–9276. https://doi.org/10.1073/pnas.1919934117

Koval, P., Brose, A., Pe, M. L., Houben, M., Erbas, Y., Champagne, D., & Kuppens, P. (2015). Emotional inertia and external events: The roles of exposure, reactivity, and recovery. *Emotion*, *15*, 625–636. https://doi.org/10.1037/emo0000059

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*, 1042–1060. https://doi.org/10.1037/a0020962

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

Larsen, R. J., & Augustine, A. A. (2009). A process approach to emotion and personality: Using time as a facet of data. *Cognition and Emotion*, *23*, 1407–1426. https://doi.org/10.1080/02699930902851302

Loossens, T. (2021). *Toward parsimonious modeling of affect dynamics in daily life: An analysis of continuous time, non-linearities and sudden jumps.* Unpublished doctoral dissertation.

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

Myin-Germeys, I., & Kuppens, P. (2022). *The open handbook of experience sampling methodology: A step-by-step guide to designing, conducting, and analyzing ESM studies* (2nd ed.). Center for Research on Experience Sampling; Ambulatory Methods Leuven.

Myin-Germeys, I., van Os, J., Schwartz, J. E., Stone, A. A., & Delespaul, P. A. (2001). Emotional reactivity to daily life stress in psychosis. *Archives of General Psychiatry*, *58*, 1137–1144.

Niemeijer, K., Mestdagh, M., Verdonck, S., Meers, K., & Kuppens, P. (2023). Combining experience sampling and mobile sensing for digital phenotyping with m-Path Sense:

Performance study. *JMIR Formative Research*, *7*, e43296. https://doi.org/10.2196/43296

Oravecz, Z., Tuerlinckx, F., & Vandekerckhove, J. (2011). A hierarchical latent stochastic differential equation model for affective dynamics. *Psychological Methods*, *16*, 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*, 425–433. https://doi.org/10.1016/j.tics.2017.03.011

Pfister, R., Schwarz, K. A., Janczyk, M., Dale, R., & Freeman, J. B. (2013). Good things peak in pairs: A note on the bimodality coefficient. *Frontiers in Psychology*, *4*, 700. https://doi.org/10.3389/fpsyg.2013.00700

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

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

Schreuder, M. J., Hartman, C. A., Groen, R. N., Smit, A. C., Wichers, M., & Wigman, J. T. W. (in press). Anticipating transitions in mental health in at-risk youths: A 6-month daily diary study into early-warning signals. *Clinical Psychological Science*. https://doi.org/10.1177/21677026221103138

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, *6*, 461–464.

Schweitzer, F., & Garcia, D. (2010). An agent-based model of collective emotions in online communities. *The European Physical Journal B*, *77*, 533–545. https://doi.org/10.1140/epjb/e2010-00292-1

Soto, C. J., & John, O. P. (2017). The next Big Five Inventory (BFI-2): Developing and assessing a hierarchical model with 15 facets to enhance bandwidth, fidelity, and predictive power. *Journal of Personality and Social Psychology*, *113*(1), 117–143. https://doi.org/10.1037/pspp0000096

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*, e63345. https://doi.org/10.1371/journal.pone.0063345

Stewart, J. (2008). *Calculus early transcendentals* (6th ed.). Thomson Brooks/Cole.

Strogatz, S. (2018). *Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering* (2nd ed.). CRC Press.

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

van de Leemput, I. A., Wichers, M., Cramer, A. O. J., Borsboom, D., Tuerlinckx, F., Kuppens, P., van Nes, E. H., Viechtbauer, W., Giltay, E. J., Aggen, S. H., Derom, C., Jacobs, N., Kendler, K. S., van der Maas, H. L. J., Neale, M. C., Peeters, F., Thiery, E., Zachar, P., & Scheffer, M. (2014). Critical slowing down as early warning for the onset and termination of depression. *Proceedings of the National Academy of Sciences*, *111*, 87–92. https://doi.org/10.1073/pnas.1312114110

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

Vanhasbroeck, N., Devos, L., Pessers, S., Kuppens, P., Vanpaemel, W., Moors, A., & Tuerlinckx, F. (2021). Testing a computational model of subjective well-being: A preregistered replication of rutledge et al. (2014). *Cognition and Emotion*, *35*, 822–835.

https://doi.org/10.1080/02699931.2021.1891863

Vanhasbroeck, N., Loossens, T., Anarat, N., Ariens, S., Vanpaemel, W., Moors, A., & Tuerlinckx, F. (2022). Stimulus-driven affective change: Evaluating computational models of affect dynamics in conjunction with input. *Affective Science*, *3*, 559–576. https://doi.org/10.1007/s42761-022-00118-5

Verdonck, S., & Tuerlinckx, F. (2014). The Ising Decision Maker: A binary stochastic network for choice response time. *Psychological Review*, *121*, 422–462. https://doi.org/10.1037/a0037012

Verduyn, P., Van Mechelen, I., Tuerlinckx, F., Meers, K., & Van Coillie, H. (2009). Intensity profiles of emotional experience over time. *Cognition and Emotion*, *23*, 1427–1443. https://doi.org/10.1080/02699930902949031

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*, 1755–1766. https://doi.org/10.1037/xge0000740

Wright, A. G. C., Hopwood, C. J., & Simms, L. J. (2015). Daily interpersonal and affective dynamics in personality disorder. *Journal of Personality Disorders*, *29*, 503–525. https://doi.org/10.1521/pedi.2015.29.4.503