



Published in final edited form as:

Nat Hum Behav. 2023 April ; 7(4): 596–610. doi:10.1038/s41562-023-01519-7.

A Highly Replicable Decline in Mood During Rest and Simple Tasks

David C. Jangraw^{1,2,*}, Hanna Keren³, Haorui Sun², Rachel L. Bedder^{4,5}, Robb B. Rutledge^{4,5,6}, Francisco Pereira¹, Adam G. Thomas¹, Daniel S. Pine¹, Charles Zheng¹, Dylan M. Nielson^{1,+}, Argyris Stringaris^{7,8,+}

¹National Institute of Mental Health, Bethesda, MD, USA

²Department of Electrical and Biomedical Engineering, University of Vermont, Burlington, VT, USA

³Azrieli Faculty of Medicine, Bar-Ilan University, Safed, Israel

⁴Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London, London, UK

⁵Wellcome Centre for Human Neuroimaging, University College London, London, UK

⁶Departments of Psychology and Psychiatry, Yale University, New Haven, CT, USA

⁷Department of Psychiatry, National and Kapodistrian University of Athens, Greece

⁸Faculty of Brain Sciences, Division of Psychiatry, University College London, London, UK

Abstract

Does our mood change as time passes? This question is central to behavioural and affective science, yet it remains largely unexamined. To investigate, we intermixed subjective momentary mood ratings into repetitive psychology paradigms. We demonstrate that task and rest periods lowered participants' mood, an effect we call "Mood Drift Over Time." This finding was replicated in 19 cohorts totaling 28,482 adult and adolescent participants. The drift was relatively large (−13.8% after 7.3 minutes of rest, Cohen's $d = 0.574$) and was consistent across cohorts. Behaviour was also impacted: participants were less likely to gamble in a task that followed a rest period. Importantly, the drift slope was inversely related to reward sensitivity. We show that accounting for time using a linear term significantly improves the fit of a computational model of mood. Our work provides conceptual and methodological reasons for researchers to account for time's effects when studying mood and behaviour.

*Corresponding author: djangraw@uvm.edu.

+These authors contributed equally to this work

Author Contributions

D.C.J., H.K., D.M.N., and A.S. devised the task. D.C.J. wrote the online experiments. D.C.J. and H.S. collected the online data. R.L.B. and R.B.R. provided data and information from the mobile app experiments. C. Z. and F.P. devised the computational model. D.C.J., C.Z., and D.M.N. wrote analysis code. D.C.J. and D. M.N. ran the analyses. D.C.J., D.M.N., and A.S. wrote the manuscript. All authors provided revisions and finalized the text.

Competing Interests

The authors declare no competing interests.

Introduction

An important but implicit notion amongst behavioural and affective scientists is that each participant has a baseline mood or affective state that will remain constant during an experiment or only vary with emotionally salient events.¹ Mood is modelled as a discounted sum of rewards and punishments,^{2,3} but many models hold that the time scale over which these events unfold is irrelevant and the passage of time itself has no effect on mood.

This assumption of a constant affective background has profound methodological implications for psychological experiments. First, consider a “resting state” functional brain scan in which a participant is asked to stare at a fixation cross. Based on the constant affective background assumption, comparisons of resting-state neuroimaging data between (for example) depressed and non-depressed participants are thought to reveal differences in their task-general traits, rather than their response to experimentally imposed rest periods. Second, consider an event-related design, such as a gambling or face recognition task, during which participants experience stimuli (wins or losses) that elicit emotional reactions. When analysing these data, responses to task stimuli are thought to occur on top of (and are often contrasted to) the affective baseline, which is presumed to be time-invariant.

Whilst convenient, this assumption of a constant affective background contradicts evidence from multiple fields that time impacts mood and behaviour. Affective chronometry research has demonstrated that affect changes systematically with time after an affective stimulus,⁴⁻⁷ and that individuals vary in the rates at which positive or negative affect decays after an event.^{8,9} Such individual differences may be linked to mental health. For instance, psychopathologists theorise that anhedonia, a symptom of both depression and schizophrenia, arises from a failure to sustain reward responses for a normative period of time.¹⁰ And studies of ADHD suggest that hyperactivity’s impulsive behaviour results from delay aversion, the idea that a delay is itself unpleasant and impulsivity is simply a rational choice to avoid it.¹¹⁻¹³

Economists speak of the opportunity cost of time, suggesting that time spent performing one activity incurs the cost of other alternatives they might have chosen instead (such as paid work or leisure).¹⁴⁻¹⁶ This idea is fundamental to the explore/exploit question that has recently preoccupied neuroscientists.¹⁷⁻¹⁹ Affect is central to this question: it is currently thought that negative affective states (such as boredom) building over time provide the subjective motivation to switch to a different activity.^{20,21}

When participants are engaged in a psychological task or rest period, they are committed to exploiting that task environment and are unable to explore other activities. This sense of constraint, or reduced agency, is considered central to feelings of boredom and its associated negative affect.²² We might therefore conceive of a psychological task’s behavioural constraint as a sort of negative affective stimulus that could gradually draw mood downward.

If this is true and the constant affective background assumption is violated, this could be problematic given evidence that spontaneous affective changes vary systematically between the individuals and groups being compared in affective science. For example, spontaneous negative thoughts are known to occur and vary substantially between humans, as highlighted

by extensive work in mind-wandering.^{23–26} Similarly, it is well known from occupational psychology that periods of low or relatively constant stimulation (as occurs in rest or repetitive experimental tasks) can induce varying levels of boredom.^{27,28} These insights raise the possibility that mood states will follow a similar pattern of inter-individual variability, creating potential confounds for resting-state and event-related experiments. But the size, stability, and clinical correlates of this variability remain unexplored.

In order to answer these fundamental questions, we examine how the passage of time affects mood in a variety of experiments across studies, participants, and settings. We find that participants' mood worsened considerably during rest periods and simple tasks, an effect we call "Mood Drift Over Time" ("mood drift" for short). This downward mood drift was replicated in 19 large and varied cohorts, totaling 116 healthy and depressed adolescents recruited in person, 1,913 adults recruited online from across the United States, and 26,896 participants performing a gambling task in a mobile app. It was not observed when participants freely chose their own activities. We show that mood drift is related to, but not a trivial extension of, the existing constructs of boredom and thought content (including the task-unrelated thought often considered central to mind-wandering). We show that mood drift slopes are positively correlated with reward sensitivity and that this relationship is moderated by overall life happiness. These findings may have profound implications for experimental design and interpretation in affective science.

Results

Characterising the Effect

The results to follow characterise the average person's gradual decline in mood during rest and simple tasks, a phenomenon we call "Mood Drift Over Time" ("mood drift" for short). This effect was initially observed in a task where participants were periodically asked to rate their mood (Figure 1A). Between these mood ratings, the initial cohort was first asked to stare at a central fixation cross. They were told that the rest period would last up to 7 minutes and that they would be asked to rate their mood "every once in a while". The mood ratings observed during this rest period inspired a number of slightly modified tasks to better characterise the effect and eliminate methodological confounds. Each modification was presented to a new cohort of naive participants so that memory and expectations would not affect their mood ratings. Each cohort also played a gambling game at some point in the task, in which they chose between an uncertain gamble or a certain outcome. This task is a standard one commonly used to examine mood.^{3,29–31} It was included to observe the effects of rest on rational behaviour, to maintain links with previous studies of mood and reward,^{2,3,32} and to enable related analyses on a large cohort of participants ($n=26,896$) playing a similar game on their smartphones³³ (Figure 1B). A list of the cohorts we examined is in Extended Data Table 1).

To quantify time's effect on mood, we created a linear mixed effects (LME) model with terms for initial mood and mood slope (i.e., change in mood per unit time) as random effects that were fitted to each subject's data. The factors of interest described in the following sections were included in the model as fixed effects (see Methods). One factor of particular interest is a depression risk score for each participant, a continuous value defined as their

score on the Mood and Feelings Questionnaire (MFQ, for adolescents) or the Center for Epidemiologic Studies Depression Scale (CES-D, for adults) divided by a clinical cutoff, i.e., MFQ/12 or CES-D/16. The model was fitted to the cohort of all participants who experienced an opening period of rest, visuomotor task, or random gambling. The slope parameter learned for each participant was used to quantify that participant's mood drift. The distribution of slopes was assumed to be Gaussian,³⁴ but LME models are robust to violations of this assumption.³⁵ All statistical tests used were two-sided unless otherwise specified.

Because the smartphone game cohort was large enough to fit hyperparameters in a held-out set of participants, this cohort's mood ratings were also fitted to a computational model that estimates each participant's initial mood and their sensitivity to rewards, reward prediction, and time (See Methods). The model's time sensitivity parameter for each participant was used to quantify their mood drift.

Mood Drift Over Time Is Sizeable During Rest

Our first objective was to estimate the size of the effect. In our initial cohort (called 15sRestBetween in Extended Data Table 1) of 40 adults recruited on Amazon Mechanical Turk (MTurk), we asked whether mood would change consistently during a rest period that preceded a gambling game. We observed a gradual decline in mood over time (Figure 2A, blue line). After 9.7 minutes of rest, the change in mood was considerable (*Mean* \pm standard error (SE) = 22.4% \pm 4.15% of the mood scale). We replicated this in 5 other adult MTurk cohorts that received shorter opening rest periods (Figure 2A, other lines).

Mood Drift Over Time Is Robust to Methodological Choices

To examine possible methodological confounds, we created slightly modified versions of the task to see whether the observed decline in mood ratings might be due to the following:

1. The aversive nature of rating one's mood: we did not find evidence that more frequent ratings changed mood drift (inter-rating-interval \times time interaction = -0.0103 %mood, 95% CI = $(-0.0267, 0.0061)$, $t_{810} = -1.23$, $p = 0.219$, 2-sided, Extended Data Figure 1).
2. The method of rating mood and its susceptibility to fatigue: we did not find evidence that making every mood rating require an equally easy keypress changed mood drift (-2.22 vs. -2.45 %mood/min, 95% CI = $(-0.772, 1.23)$, $t_{70} = 0.427$, $p = 0.671$, 2-sided).
3. The expected duration of the rest period: groups expecting different rest durations did not have different mood drift (-1.47 vs. -1.53 %mood/min, 95% CI = $(-0.613, 0.743)$, $t_{104} = 0.185$, $p = 0.854$, 2-sided).
4. Multitasking or task switching: participants moved their mood rating slider on 97.7% of trials.

The results of these control analyses suggested that mood drift cannot be explained by these methodological factors (Supplementary Note C).

Mood Drift Over Time Occurs During Tasks

To see whether this decline was specific to rest or more generally linked to time on task, we administered two variants of the task. The first variant (cohort Visuomotor-Feedback, $n = 30$) was designed to mimic rest very closely while requiring the participant to respond regularly and giving feedback on their performance. Specifically, a fixation cross moved back and forth periodically across the screen, the participant was asked to press a button whenever it crossed the centerline, and each response would make the cross turn green if the response was accurate or red if it was too early or late (see Methods). In the second variant (cohort Daily-Random-01, $n = 66$), the subject played a random gambling game in which gambling outcomes and reward prediction errors (RPEs) were both random with mean zero. Both of these tasks produced similar mood timecourses, and we did not find evidence of a difference between the LME slope parameters of this group and those of the original cohort (-2.19 vs. -2.45 %mood/min, 95% CI = $(-0.876, 1.40)$, $t_{68} = 0.437$, $p = 0.663$ for visuomotor task, -1.91 vs. -2.45 , 95% CI = $(-0.453, 1.52)$, $t_{104} = 1.07$, $p = 0.287$ for random gambling, both 2-sided) (Figure 2B).

Mood Drift Over Time Is Generalizable

We next investigated the generalizability of this result across age groups and recruitment methods. To do this, we collected similar rest + gambling data via an online task from adolescent participants recruited in person at the National Institute of Mental Health in Bethesda, MD and asked to complete the task online via their home computers (see Methods). This group (Adolescent-01, $n=116$) showed a pattern of declining mood similar to that observed in the MTurk cohort (Figure 2C) (-1.69 vs. -1.93 %mood/min, 95% CI = $(-0.122, 0.599)$, $t_{884} = 1.09$, $p = 0.275$, 2-sided).

To more precisely characterise the effect, we fitted a large LME model to the complete cohort of online participants (both adults and adolescents) completing rest or simple tasks in the first block (Extended Data Table 2). The mood drift parameter (rate of mood decline with time) for these 886 participants was $Mean \pm SE = -1.89 \pm 0.185$ %mood/min, which was significantly less than 0 ($t_{864} = -10.3$, $p < 0.001$). After 7.3 minutes (the mean duration of the first block of trials), the mean decrease in mood estimated by this LME model was 13.8% of the mood scale. This corresponds to a Cohen's $d = 0.574$, with a 95% CI = $(0.464, 0.684)$.³⁶

Mood Drift Is Diminished in a Mobile App Gambling Game

We next tested whether mood drift could be observed in a large dataset ($n = 26, 896$) of mood ratings during a similar gambling task played on a mobile app. All analyses were applied to an exploratory cohort of 5,000 of these participants, then re-applied to the confirmatory cohort of all remaining participants after preregistration (<https://osf.io/paqf6>). We applied the LME modeling procedure to this confirmatory cohort and again found a slope parameter that was significantly below zero at the group level ($Mean \pm SE = -0.881 \pm 0.0613$ %mood/min, $t_{22804} = -14.4$, $p < 0.001$).

It is notable that even in this relatively engaging game (in which tens of thousands of participants completed the task despite not being paid for participating or penalised for failing to finish), mood tended to decrease with time spent on task.

We note, however, that mood drift was significantly smaller in this cohort (median= -0.752 , inter-quartile range (IQR)= 2.10 %mood/min) than in the combined cohort of online participants (median= -1.53 , IQR= 2.34 %mood/min, 2-sided Wilcoxon rank-sum test, $W_{21761} = -14.5$, $p < 0.001$, Extended Data Figure 2). 87.5% of online participants had negative slopes in the LME analysis, whereas only 70.2% of mobile app participants did. A histogram of the LME slope parameters for online and mobile app participants is plotted in Figure 3. This shows that, as one might expect, mood drift is sensitive to task context.

Next, to disentangle mood drift from the effects of reward and reward prediction error in this dataset, we fitted the computational model described in the Methods section to the mobile app data. Including the mood slope parameter in the model decreased the mean squared error on testing data (the last two mood ratings of the task) from 0.336% to 0.325% of the mood scale for the median subject across regularizations, a significant improvement (IQR= 0.00197% , 2-sided Wilcoxon signed-rank test, $W_{499} = 0$, $p < 0.001$). This suggests that time on task affected a participant's mood beyond the impacts of reward and expectation, and did so in a way that was stable within individuals because improved fits were observed in held-out data. Fits and parameter distributions can be seen in Extended Data Figures 3 and 4. The distribution of participants' time sensitivity parameters β_T (which can be interpreted as mood drift independent of reward effects) was centered significantly below zero ($Mean \pm SE = -0.128 \pm 0.00668$ %mood/min, 2-sided Wilcoxon signed-rank test $W_{21895} = 1.00 * 10^8$, $p < 0.001$).

Mood Drift Over Time Is Absent in Freely Chosen Activities

After the surprising finding that mood drift appeared during an engaging mobile app game, we wondered whether this phenomenon would be observed in daily life, outside the context of a psychological task. We therefore designed and preregistered (<https://osf.io/gt7a8>) a task in which the initial rest period was replaced with 7 minutes of free time, during which the participant could pursue activities of their choice. Participants completing this task (cohort Activities, $n=450$) were asked to rate their mood just before and just after the break period. They were then asked to report what they did. The most frequent activities reported were thinking, reading the news, and standing up (Supplementary Table 3).

This group was the first sample investigated in this study that did not exhibit mood drift. The mood ratings just after the free period were not statistically different from the mood ratings before the free period (66.6% vs. 65.7%, 95%CI = $(-2.15, 97)$, $t_{449} = -1.33$, *one-tailed* $P_{H0:decrease} = 0.0918$, $P_{H0:increase} = 0.908$). This change in mood was significantly greater than that of a cohort who received the standard rest period with interspersed mood ratings (cohort BoredomAfterOnly, $n=150$) (0.909% vs. -8.11% , 95%CI = $(5.95, 12.1)$, $t_{598} = 6.28$, $p < 0.001$, 2-sided). This shows that, perhaps unsurprisingly, mood drift is not universal to all activities. However, the nominal increase in mood during this period (0.130% mood/min) was much smaller than the decrease in mood observed during a typical rest period

(-1.89% mood/min). Each minute in which participants could choose their activity raised their collective mood less than 10% of the mood decline experienced during a minute of rest.

Inter-Individual Differences

Having characterised the effect at the group level, we next turned our attention to the individual. The motivation for this line of analysis is that if an individual's mood slope is different from that of others in a way that remains stable over days or weeks, it may be linked to traits of clinical and theoretical interest. While the group average mood drift is negative during rest and simple tasks, there is considerable variation across participants (2.5th - 97.5th percentile of subject-level mood drift for online participants: -7.23 - 1.79% mood/min) (Figure 3). Using an intraclass correlation coefficient (ICC) on cohorts that completed the task more than once, we found that these individual differences had moderate, statistically significant stability across blocks ($ICC(2,1) = 0.465, p < 0.001$), days ($ICC(2,1) = 0.343, p = 0.0031$), and weeks ($ICC(2,1) = 0.411, p < 0.001$, one-sided since ICC values are expected to be positive) (Extended Data Figure 5, Supplementary Note D). We therefore investigated the relationship between this variability and other traits of clinical and theoretical interest.

Mood Drift Is Associated with Sensitivity to Rewards—Mood is central to depression, which is thought to relate etiologically to reward responsiveness.^{37,38} The idea that mood drift might be related to this responsiveness prompted us to investigate the relationship between participants' mood drift, reward sensitivity, and life happiness in our computational model fits. The time sensitivity/mood drift parameter β_T was anticorrelated with the reward sensitivity parameter β_A ($r_s = -0.106, p < 0.001$, 2-sided) (Figure 4, left). This anticorrelation was weaker in participants with life happiness below the median (i.e., those at greater risk of depression) than it was in those at/above it ($r_s = -0.0513$ vs. -0.14 , $Z = 6.41, p < 0.001$, 2-sided) (Figure 4, right). This suggests that people more sensitive to the passage of time are also more sensitive to rewards, and that this relationship is less pronounced in those with greater depression risk.

The direct relationship between depression risk and mood drift was significant, but its effect on model fit was very small. In our online participant LME model, higher depression risk score was significantly associated with less negative mood drift (depression-risk * time interaction, $Mean \pm SE = 0.515 \pm 0.109\% \text{ mood/min}$, $t_{869} = 4.75, p < 0.001$, Extended Data Figure 6). Whilst the model fit improved, the within-individual variance explained by the addition of this interaction term was very small ($f^2 = 0.00289$).^{39,40} Nevertheless, the interaction term's significance was replicated in two more independent cohorts (including the mobile app cohort, where time sensitivity and life happiness were weakly anticorrelated, Extended Data Figure 7, bottom right) and was robust to methodological artefacts such as floor effects (Supplementary Notes E-G).

Taken together, these results demonstrate relationships between mood drift and other important individual differences: depression risk, life happiness, and reward sensitivity.

Impact on Behaviour

Participants Are Less Likely to Gamble After Rest Periods—To investigate whether mood drift's effects extend to behaviours beyond subjective mood reports, we examined the impact of rest and mood drift on behaviour in the gambling tasks. Past research has shown that a participant's choice between a certain outcome and a more exciting but uncertain gamble is affected by mood as induced by unexpected gifts,^{41,42} music,⁴³ and feedback.³¹ We asked whether mood drift would influence this behaviour in a similar way.

We observed that gambling (specifically positive closed-loop gambling, in which participants tended to receive positive RPEs) participants who had a preceding rest or visuomotor task block had significantly lower mood at gambling onset than those who did not (median 0.55 vs. 0.66, IQR 0.28 vs. 0.31, 2-sided Wilcoxon rank-sum test, $W_{722} = 2.08$, $p = 0.0377$) (Figure 5, top). This effect was no longer significant at the next mood rating, which took place around trial 4 of gambling. We therefore examined gambling behaviour in these first 4 trials. Those who had experienced either a short (350–450 s) or long (500–700 s) opening rest period were significantly less likely to gamble than those who had not (median=3, IQR=2 for both short- and long-rest, 2-sided Wilcoxon rank-sum test, no-rest vs. short-rest: $W_{469} = 4.85$, $p < 0.001$; no-rest vs long-rest: $W_{344} = 4.79$, $p < 0.001$; both $< 0.05/3$ controlling for multiple comparisons). (Figure 5, bottom). However, we did not find evidence of a difference between the long and short rest groups (2-sided $W_{629} = 0.52$, $p = 0.603$). Trial-wise gambling behaviour differences between rest and no-rest groups are most pronounced in the first four trials, much like the differences observed in mood (Figure 5, middle). However, no significant correlation was observed between an individual's mood drift parameter during the preceding rest block and the number of times they chose to gamble in the first 4 trials ($r_s = 0.0317$, $p = 0.427$, 2-sided).

Relationship to Boredom and Thought Content

We next examined whether the existing construct of boredom or mind-wandering (MW) could trivially explain mood drift. In a preregistered (<https://osf.io/gt7a8>) data collection and analysis, we examined the relationship between mood drift and these more established constructs at the state level, state change level, and trait level (Supplementary Notes L-M). Participants were randomised to a boredom, MW, or Activities cohort (described previously) at the time of participation.

Mood Drift Over Time is Weakly Related to State Boredom—We assessed whether mood drift could be explained by boredom. Participants completed a rest block with interspersed mood ratings, plus a state boredom questionnaire (the Multidimensional State Boredom Scale's short form, MSBS-SF)⁴⁴ afterwards (cohort BoredomAfterOnly, $n = 150$), or before and afterwards (cohort BoredomBeforeAndAfter, $n = 150$), and a trait-boredom questionnaire (the short boredom proneness scale, SBPS).⁴⁵

In our LME model of mood, we added a factor for final state boredom (i.e., at the end of the rest block). We then compared this baseline model to one that further added the interaction between final-boredom and time. The difference represents the ability of boredom to

account for mood drift. Whilst the model fit improved, the added within-individual variance explained by the addition of this new interaction term was very small ($f^2 = 0.00578$). The change in state boredom across the rest block produced similar results ($f^2 = 0.0111$).

Including time's interaction with trait boredom in the model did not explain significant additional variance in mood (Likelihood ratio test: $\chi^2(1, N = 16) = 0.0253, p = 0.874$).

Mood Drift Over Time is Weakly Related to Thought Content—We also assessed whether mood drift could be explained by the content of ongoing thought, including the task-unrelated thought, stimulus-independent thought, and spontaneity often considered in definitions of MW.⁴⁶ We note that such content-based definitions of MW are controversial and does not capture the dynamics-based definition espoused by some researchers.^{47,48} New participants completed a rest block with interspersed mood ratings, plus a Multidimensional Experience Sampling (MDES) questionnaire⁴⁹ afterwards (cohort MwAfterOnly, $n = 150$), or before and afterwards (cohort MwBeforeAndAfter, $n = 150$), and a trait-MW questionnaire (the mind-wandering questionnaire (MWQ)⁵⁰). MDES results produce 13 principal components that attempt to capture the content of ongoing thought. We investigated how well this complete collection of components explains within-individual mood variance.

In our LME model of mood, we added 13 factors for “final” MDES components (i.e., at the end of the rest block). We then compared this baseline model to one that further added the 13 interactions between these final-MDES components and time. The difference represents the ability of MDES components to account for mood drift. Whilst the model fit improved, the within-individual variance explained by the addition of these new interaction terms was small ($f^2 = 0.0227$). The change in MDES components across the rest block produced similar results ($f^2 = 0.0380$). Including time's interaction with trait MW in the model did not explain significant additional variance in mood ($\chi^2(1, N = 16) = 0.305, p = 0.581$).

Discussion

In this study, we describe the discovery of a highly replicable and relatively large effect which we call Mood Drift Over Time: the average participant's mood gradually declined with time as they completed simple tasks or rest periods. Mood's sensitivity to the passage of time is a long-intuited phenomenon that is widely acknowledged in literature^{51–53} and philosophy.^{54–56} Our results provide robust empirical evidence for this phenomenon and reveal its temporal structure, its variability across individuals, and its level of stability. These results call into question the long-held constant affective background assumption in behavioural and affective science.

The mechanism that enables mood to be sensitive to the passage of time is not yet known. One possibility is that humans store expectations about the rate of rewards and punishments in the environment and that prolonged periods of monotony violate such expectations. Such a view aligns with the recently articulated theoretical progress in integrating opportunity cost across time to guide behaviour.²¹ Lower mood could function as an estimate of that opportunity cost, making mood drift an adaptive signal that informs decisions to exploit (stay on task) or explore (switch task).²⁰

Supporting this reward/cost-based interpretation of our findings is our observation that depressed participants showed less negative mood drift. This would at first seem paradoxical since phenomena such as boredom have traditionally been linked to melancholia and depression (e.g., by Schopenhauer⁵⁷ and Kierkegaard⁵⁸). Yet it has been argued cogently⁵⁹ that such a view conflates negative affect as a trait (e.g., proneness to boredom) with negative affect as a state (a momentary experience). Since valuation of reward is thought to be reduced in depression,^{37,38} it is possible that misalignment with one's goals and violation of reward expectations—and resultant downward mood drift—will be less pronounced in depression. This interpretation is supported by our finding that mood drift is less pronounced in those with lower reward sensitivity, and that the relationship between reward sensitivity and mood drift was moderated by depression risk (Figure 4). It is tempting to speculate that reduced mood drift could contribute to reduced motivation for action or environmental change in those with depression.

We found that mood declined during rest and tasks (including a mobile app more engaging than most experiments) but not freely chosen activities. This suggests that researchers are subjecting their participants to an unnatural stressor in their experiments without accounting for it in their analyses or interpretations. Changes in mood on the scale of tens of minutes prevent these longer blocks of time from being truly interchangeable. This means that variations in experimental procedures that might seem inconsequential could still introduce confounds.

For example, let's consider a large collaborative study that is based on multisite imaging data collection, such as ENIGMA.⁶⁰ In this dataset, centres vary in the duration of the resting-state fMRI scan and whether it takes place at the start or end of the scan session.⁶¹ This could lead to high variability between sites simply because patients at sites with longer or later scans spent more of the scan in a bad mood. At best, the neural correlates of that decreased mood will be uncorrelated with the effect of interest, increasing noise and reducing statistical power. At worst, they could be mistaken for neural correlates of a certain genotype that is more common in the country where the longer scans took place. (We do not imply that mood drift lowers reliability in resting-state MRI;^{62–64} we simply point out its role as a potential confound when drawing inferences about mood and brain states during/after rest.)

In this paper, we introduce the new term Mood Drift Over Time for the following reasons. First, the phenomenon is highly replicable; second, it is of considerable effect size; third, it is relevant to both everyday situations and to scientific experiments; fourth, mood drift does not seem to be captured by existing terms such as boredom or mind-wandering. We employ the term mood drift in the spirit of describing a mental phenomenon,^{65–67} as a first step before explaining or categorising it. It is possible that mechanisms for mood drift are reward sensitivity and opportunity cost, yet the subjective experience and its influence on the outcome of experimental studies seem to require the separate term that we have introduced.

The distinction between mood drift and boredom requires special consideration due to their apparent similarities. State boredom assessed using the MSBS-SF⁴⁴ accounted for modest variance beyond other factors. Of course, the MSBS is only one (relatively well established)

way of measuring boredom; moreover, there is debate about the very conceptualisation of boredom and its heterogeneity.^{22,59,68} Therefore, we cannot conclude purely from these results that boredom is not driving mood drift. Future work might instead ask participants to directly report their boredom,⁶⁹ enabling more frequent assessment of boredom as an emotion.⁷⁰

Importantly, we show that accounting for time using a linear term significantly improves the fit of a computational model of mood. A linear term may be unrealistic as we expect that on a bounded mood scale, the effect will eventually saturate. However, we propose that until alternative models have been established, the linear term may be a good-enough way to account for the substantial effects of mood drift on the time scale of most experiments.

Our study has several strengths, including adherence to good data analysis practices such as preregistration and replication, the addition of a longitudinal design to test reliability, and the use of rigorous computational modeling (including train-test splits and regularisation). Our study demonstrated the effect in adolescents as well as adults and showed how the effect differs in people with varying reward sensitivity and depression risk. We used control experiments to eliminate potential confounds and test alternative explanations (Supplementary Notes C-G).

Yet our study should also be seen in light of some shortcomings.

First, this study uses self-reported momentary mood ratings as in previous studies with similar methodology.^{2,3} Such ratings can be criticised as being subjective and difficult to interpret. However, mood is a well-established construct of central importance to affective science. Its definition as a long-duration affective state that is not immediately responsive to stimuli^{71,72} makes it central to the study of mood disorders defined by long-term affect.⁷³ Mood is distinct from emotion, in part, by being less temporally responsive.^{74–76} Mood's links to long-term context makes it the more useful construct to describe gradual changes in affect.

Despite its subjectivity, self-report remains the gold standard for the measurement of mood and emotion.^{76–78} It is widely used in clinical,⁷⁹ epidemiological,⁸⁰ and psychological research (including ecological momentary assessment⁸¹). Other physiological “markers” of affect are typically benchmarked against these self-reports. And evidence suggests that these candidates lack the reliability of self-reports: different emotions cannot be distinguished by their autonomic nervous system signatures,⁸² facial expressions,^{83,84} or neural activity.⁸⁵ In our experiments, initial mood ratings showed strong association with trait mood ratings, underscoring their psychometric validity (Extended Data Figure 8).

Our study cannot conclusively determine mood drift's behavioural consequences. On average, rest induces downward mood drift (Figure 2) and decreases gambling behaviour (Figure 5). However, a significant correlation between and individual's mood drift and gambling behaviour was not observed. Our results are not able to discern whether the change in behaviour is directly linked to mood drift or to some other consequence of rest.

Our study's limited set of tasks, all of which induced mood drift, makes it difficult to discern the phenomenon's key contributing factors. We chose to focus on a category that is extremely common in neuroscience: long, neutral, low-stimulation tasks. Most researchers would see these qualities as unobjectionable or even desirable. We hope that the results of this study will lead researchers to reexamine this idea in their own research.

Methods

Participants

Online Adult Participants—Online adult participants were recruited using Amazon Mechanical Turk ([Amazon.com](https://www.amazon.com), Inc., Seattle, WA), a service that allows a person needing work done (a “requester”) to pay other people (“workers”) to do computerised tasks (“jobs”) from home.⁸⁶ Requesters can use “qualifications” to require certain demographic or performance criteria in their participants. We required that our participants be adults living in the United States, that they have completed over 5,000 jobs for other requesters, and that over 97% of their jobs have been satisfactory to the requester. We also required that participants had not performed any of our tasks (which were relatively similar to the ones in this study) before.

Every online participant received the same written instructions and provided informed consent on a web page where they were required to click “I Agree” to participate. Because we did not obtain information by direct intervention or interaction with the participants and did not obtain any personally identifiable private information, our MTurk studies were classified as not human subjects research and were determined to be exempt from IRB review by the NIH Office of Human Subjects Research Protections (OHSRP). The consent process and task/survey specifics were approved by the OHSRP. For data to be included in the final analyses, participants were required to complete both a task and a survey (described below). Participants submitted a 6-to-10-digit code revealed at the end of each one to prove that they had completed it. Both the task and survey had to be completed in a 90-minute period starting when they accepted the job on Amazon Mechanical Turk.

The consent form included a description of the tasks they were about to perform, but cohorts were blinded to the specific cohort to which they had been assigned. Most cohorts were collected in series, but some were randomised to a cohort at the time of participation (we have specified these in the Methods or Results). In the initial cohorts, no statistical methods were used to pre-determine sample sizes, but our cohort sample sizes are similar to those reported in,² and our combined cohorts are much larger.

914 participants completed the task online. Some data files did not save properly due to technical difficulties or the participant closing the task window before being asked to do so. 44 participants whose task or survey data did not save were excluded. Of the 870 remaining Mechanical Turk participants, 390 were female (44.8%). Participants had a mean age of 37.6 years (range: 19–74).

A subset of the online adult participants were invited to return the following day to repeat the same task and survey a second time. Of the 66 individuals who completed both the task

and the survey on the first day, 53 (80.3%) completed the task and survey on the second day. Gambling trials were randomised independently so that the subject was not seeing the exact same trials both times. Participants could complete the second task and survey any time in the following three days, but the task and survey had to be done together in the same 90-minute period.

Similarly, a different cohort was invited to return a week after their first run to repeat the same task and survey. These participants could complete the second task and survey any time in the following six days, but the task and survey had to be done together in the same 90-minute period. This cohort was then invited to complete the same task and survey a third time, two weeks after their first run. 196 individuals completed the task and survey the first week. 163 (83.2%) of these completed the task and survey the second week and 158 (80.6%) completed the task and survey the third week. 149 (76.0%) individuals completed the task and survey in all three weeks.

Online Adolescent Participants—Adolescent participants recruited in person at the National Institute of Mental Health were also invited to participate by completing a similar task on their computer at home. These participants completed a different set of questionnaires, developed for adolescents, about their mental health. Every participant received the same scripted instructions and provided informed consent to a protocol approved by the NIH Institutional Review Board.

There were 230 adolescents enrolled in the NIMH depression characterization study who were offered to complete tasks for this study. 129 agreed, a participation rate of 56.1%. 10 adolescents who had not completed all three questionnaires were excluded from the results, as were 3 participants who declined to allow their data to be shared openly. Of the remaining 116 adolescent participants, 77 were female (66.4%). They had a mean age of 16.3 years (range: 12 – 19). 56 participants (48.2%) had been diagnosed with major depressive disorder (MDD) by a clinician at the NIH, and 4 were determined to have sub-clinical MDD (3.4%). Participants had a mean depression score of MFQ = 6.5 (\pm 5.5 SD) and a mean anxiety score of SCARED = 2.2 (\pm 3.0 SD).

To assess the stability of findings in this population, the in-person adolescent participants were invited to return each week to complete the same task again, up to three times. 82 (70.6%) individuals completed the task a week later and 4 (3.4%) completed the task a third time the following week. The analyses presented in this paper use only the first run from this cohort.

Boredom, Mind-Wandering, and Activities Participants—In response to reviewer comments, a preregistered follow-up analysis included five new cohorts of MTurk participants who received similar tasks that also included mood ratings, rest periods, and the gambling game. This group was recruited to investigate the impacts of boredom and mind-wandering on mood changes, so they completed surveys about these traits in addition to the demographics, CES-D, and SHAPS questions. Participants were randomised to one of these 5 “follow-up cohorts,” summarised in Extended Data Table 1:

- BoredomBeforeAndAfter (n=150), who received a boredom state questionnaire both before and after a 7-minute rest period with 15 s of rest between mood ratings.
- BoredomAfterOnly (n=150), who received a boredom state questionnaire only after a 7-minute rest period with 15 s of rest between mood ratings.
- MwBeforeAndAfter (n=150), who received a multidimensional experience sampling (MDES) questionnaire both before and after a 7-minute rest period with 15 s of rest between mood ratings.
- MwAfterOnly (n=150), who received an MDES questionnaire only after a 7-minute rest period with 15 s of rest between mood ratings.
- Activities (n=450), who received instructions to leave the task for 7 minutes and perform activities of their choice, completing mood ratings just before and after this period.

After the rest periods described above, each group completed a block of negative closed-loop gambling trials and a block of positive closed-loop gambling trials (as described in the “Gambling Blocks” section). Details of the cohorts’ tasks are found in the following sections. A full description of the preregistered tasks and analyses can be found at <https://osf.io/gt7a8>, registered on November 18, 2021. 1143 participants completed these tasks online. 93 participants were excluded because their task or survey data was incomplete or did not save, because they completed the task more than once despite instructions to the contrary, or because they failed to answer one or more “catch” questions correctly on the survey. Of the 1050 remaining participants, 463 were female (44.1%). Participants had a mean age of 39.3 years (range: 20–80).

The above sample sizes were selected using power calculations described in detail in the preregistration. For the scale validation experiments, a sample size of 150 in each group with an alpha of 0.01 gives 99.02% power to detect a medium effect ($d = 0.5$) and 83.04% power to detect an intermediate effect ($d = 0.3$) assuming the effect truly is null at a population level. Power for linear multiple regression tests were calculated in G*Power.⁸⁷ In the boredom and MW cohorts, samples of 150 participants were selected to provide 80% power to detect a 7.99% increase in variance explained with the inclusion of a single parameter (alpha = 0.01, total predictors) and a 95% power to detect a 12.18% change in variance explained. In analyses using a pair of cohorts, 300 participants gives 80% power to detect a 3.93% increase in variance explained and a 95% power to detect a 6.01% increase in variance explained. An Activities cohort of 450 participants was chosen to provide 80% power to detect a difference between the Activities and MTurk cohorts of Cohen’s $d = 0.2$, and it also provides 80% power to detect a decrease in mood in the Activities cohort of Cohen’s $d = 0.15$.

Mobile App Participants—Gambling behaviour and mood rating data were collected from a mobile app called “The Great Brain Experiment”, described in.³ The Research Ethics Committee of University College London approved the study. When participants opened the app for the first time, they gave informed consent by reading a screen of information about

the research and clicking “I Agree.” They then rated their life satisfaction as an integer between 0 (not at all) and 10 (completely). Any time they used the app after this, participants could then choose between several games, including one called “What makes me happy?” that was used in this research. We used a subset of 26,896 people, primarily from the US and UK, in our analyses. The median life satisfaction of the included participants, which will be used as a proxy for depression risk in this cohort, was 7/10. Age for this cohort was provided in bands. These are the bands and number of individuals in each band in the subset of data used in our analysis: 18–24 (6,500), 25–29 (4,522), 30–39 (7,190), 40–49 (4,829), 50–59 (2,403), 60–69 (1,158), and 70+ (294). 13,168 were female (49.0%).

Mobile app participants were randomly split into an exploratory cohort of 5,000 participants and a confirmatory cohort of all remaining participants. All analyses and hyperparameters involving mobile app participants were optimised using only the exploratory cohort, then tested on the confirmatory cohort. These confirmatory analyses were preregistered on the Open science Framework (<https://osf.io/paqf6>, registered on January 29, 2021).

In the linear mixed effects model described below, we made an effort to exclude participants who were outliers in the time they took to complete the task. Such outliers would have a large effect on the LME model’s mood slope term, where non-zero slopes would lead to large errors in these outlier participants. Outlier completion times also suggest that the participant was not fully paying attention to the task, either by responding without thinking or leaving the app for an extended period. Mobile app participants with an average task completion time that was less than $Q1 - 1.5 * IQR$ or greater than $Q3 + 1.5 * IQR$ (where $Q1$ is the 25th percentile, $Q3$ is the 75th percentile, and $IQR = Q3 - Q1$) were excluded from this linear mixed effects analysis. 4.65% of participants were excluded based on these criteria, leaving $n = 20,877$ mobile app participants.

Task and Survey

The online tasks were created using PsychoPy3 (v2020.1.2) and were uploaded to the task hosting site Pavlovia for distribution to participants. Pavlovia used the javascript package PsychoJS to display tasks in the web browser. Each task used the latest version of Pavlovia and PsychoJS available at the time of data collection. A list of all cohorts collected can be seen in Extended Data Table 1.

Mood Ratings—The task given to online participants is outlined in Figure 1A. Periodically during all tasks, participants were asked to rate their mood. Participants first saw the question “How happy are you at the moment?” for 3 seconds. Then a slider appeared below the question, with a scale whose ends were labeled “unhappy” and “happy.” A red circle indicated the current slider position, and it started in the middle for each rating. Participants could press and hold the left and right arrow keys to move the slider, then spacebar to lock in their response. If the spacebar was not pressed in 4.5 seconds, the current slider position was used as their mood rating.

As part of the instructions at the start of each run, the participant was asked to rate their overall “life happiness” in a similar (but slightly slower) rating. In this case, participants first

saw the question “Taken all together, how happy are you with your life these days?” for 4 seconds. The slider then appeared, and the participant had 6.5 seconds to respond.

In one alternative version of the task, participants were asked to rate their mood with a single keypress instead of a slider. They could press a key 1–9 to indicate their current mood, where 1 indicated “very unhappy” and 9 indicated “very happy.” This alternative version was used to investigate the possibility that mood effects could be an artefact of the rating method, where participants’ ratings converged to the middle because this rating required the least effort.

Rest Blocks—In some blocks, participants were asked to simply rest in between mood ratings. These rest periods consisted of a central fixation cross presented on the screen. The duration of the rest period was 15 seconds for most versions of the experiment. For some versions, this duration was made longer or shorter to disentangle the impacts of rating frequency and elapsed time on mood, investigating the possibility that the mood ratings themselves were aversive.

Thought Probes and Activities Questions—Follow-up versions of the task included thought probes about state boredom or the emotional valence of ongoing thought (including mind-wandering). These groups received rest blocks as described above, but with additional questions just before and/or after it.

Two cohorts were collected to quantify the relationship between mood drift and boredom. Each received a rest period with mood ratings 20 seconds apart, followed by the Multidimensional State Boredom Scale’s short form (MSBS-SF), an 8-item scale of state boredom.⁴⁴ Participants rated statements like “I feel bored” on a 7-point Likert scale from 1 (“Strongly Disagree”) to 7 (“Strongly Agree”). Their level of boredom was quantified as the sum of their ratings on the 8 questions. The first (cohort BoredomBeforeAndAfter, $n = 150$) completed the MSBS-SF both before and after the rest period. The second (cohort BoredomAfterOnly, $n = 150$) completed the MSBS-SF only after the rest period.

Two other cohorts were collected to quantify the relationship between mood drift and the emotional valence of ongoing thought (including mind-wandering). Each participant in the two mind-wandering cohorts received a rest period with mood ratings 20 seconds apart, followed by a 13-item Multidimensional Experience Sampling (MDES) as described by Turnbull et al.⁴⁹ Participants were asked to respond to a set of questions by clicking on a continuous slider. Most questions, like “my thoughts were focused on the task I was performing”, were rated from “not at all” (scored as -0.5) to “completely” (scored as 0.5). The first (cohort MwBeforeAndAfter, $n = 150$) completed the MDES only after the rest period. The second (cohort MwAfterOnly, $n = 150$) completed the MDES only after the rest period.

As described by Ho et al.,⁸⁸ we used principal components analysis (PCA) to quantify the affective valence of thought at each administration of MDES. We first compiled the MDES responses of all participants in the MwAfterOnly group into a matrix with 13 (the number of items in each administration) columns and 450 (the number of administrations) rows.

We then used scikit-learn's PCA function to find 13 orthogonal dimensions explaining the MDES variance. The use of PCA orthogonalises the MDES responses, which is desirable for their use as explanatory variables in an LME.³⁵

For a preregistered analysis, we focused on the emotional content of ongoing thought (this approach was later abandoned in favour of examining the collective predictive power of all 13 MDES components, Supplementary Notes L-M). By examining the component matrix, we identified the component that loaded most strongly onto the "emotion" item of the MDES (in which they reported their thoughts as being negative or positive). The "emotion dimension" of each MDES (in both MW cohorts) was then quantified as the amplitude of this component, calculated by applying this prelearned PCA transformation to the data and extracting the corresponding column. The sign of PCA components is not meaningful, so we arbitrarily chose that increased emotion dimension would represent more negative thoughts.

Another follow-up task investigated the impact on mood of a break period where participants were released to do whatever they wanted. Just before this break period, an alarm sound was played on repeat, and participants were asked to increase the volume on their computer until they could hear the alarm clearly. Participants were informed that they would have 7 minutes to put the task aside and do something else but should be ready to come back when the alarm sounded at the end. After these instructions and before the break, they rated their mood. During the break, the task window displayed a message saying "this is the break. An alarm will sound when the break is over" After the alarm sounded and participants returned, they rated their mood again. They were then asked 27 questions about how much of the break they spent doing various activities. They were asked to rate each by clicking on a 5-point Likert scale with options labeled "not at all" (scored at 0%), "a little" (scored at 25%), "about half the time" (scored at 50%), "a lot" (scored at 75%), or "the whole time" (scored at 100%). These scores were used to roughly describe the most common activities performed by the participants during the break.

Participants were randomised to one of the follow-up cohorts described in this section at the time of participation.

Task Blocks—In some blocks, participants completed a simple visuomotor task. In this task, the fixation cross moved back and forth across the screen in a sine wave pattern (peak-peak amplitude: 1x screen height, period: 4 seconds). Participants were asked to press the spacebar at the exact moment when the cross was in the center of the screen (as denoted by a small dot). In some blocks, they received feedback on their performance: each time they responded, the white cross turned green for 400 ms if the spacebar was pressed within the middle 40% of the sine wave's position amplitude (i.e., less than 0.262 seconds before or after the actual center crossing).

Gambling Blocks—In each trial of the gambling task, participants saw a central fixation cross for 2 seconds. Three boxes with numbers in them then appeared. Two boxes on the right side of the screen indicated the possible point values they could receive if they chose to gamble (the "win" and "loss" values). On the left side, a single number indicated the points they would receive if they chose not to gamble (the "certain" value). Participants had

3 seconds to press the right or left arrow key to indicate whether they wanted to gamble or not. If no choice was made, gambling was chosen by default. After making their choice, the option(s) not chosen would disappear. If they chose to gamble, both possible gambling outcomes appeared for 4 seconds, then the actual outcome appeared for 1 second. If they chose not to gamble, the certain outcome appeared for 5 seconds. The locations (top/bottom) of the higher and lower gambling options were randomised.

The gambling outcome values were calculated according to several rules depending on the version of the experiment. In each version, the “base” value was a random value between -4 and 4 points. The other value was this base value plus a positive or negative reward prediction error (RPE). If they chose to gamble, participants would always receive the base value + RPE option. To encourage gambling, the “certain” value was set to $(win + 2 * loss)/3$, or $1/3$ of the way from the loss value to the win value. (Note that this rule was the same for every subject and was therefore unlikely to drive individual differences in gambling behaviour.)

In the “random” version, the RPE was a random value with uniform distribution between -5.0 and 5.0 . RPE magnitudes of less than 0.03 were increased to 0.03 . If 3 trials in a row happened to have the same outcome (win or loss), the next trial was forced to have the other outcome.

In the “closed-loop” version, RPEs were calculated based on the difference between a participant’s mood and a “target mood” of 0 or 1 . Some blocks of trials were “positive” blocks in which the participant had a 70% chance of winning on each trial (“positive congruent trials”) and a 30% chance of losing (“positive incongruent trials”). Other blocks were “negative” blocks in which the participant had a 70% chance of losing on each trial (“negative congruent trials”) and a 30% chance of winning (“negative incongruent trials”). If there had been 3 incongruent trials in a row, the next trial was forced to be congruent. The RPE was calculated as in a Proportional-Integral (PI) controller: a weighted sum of the current difference and the integral across all such differences reported so far in the block. The weightings were different for congruent and incongruent trials. Specifically, the RPE was set to:

$$RPE(t) = \begin{cases} 14 * (M(t-1) - M_{target}) + \sum_{j=1}^{t-1} (M(j) - M_{target}) & \text{congruent trial} \\ -3.5 * (M(t-1) - M_{target}) + \sum_{j=1}^{t-1} (M(j) - M_{target})/12 & \text{incongruent trial} \end{cases}$$

Where $t = 1, 2, \dots, n$ is the trial index relative to the start of the block, $M(t)$ is the mood reported after trial t , and M_{target} is the target mood for the current block. RPEs with a magnitude of less than 0.03 were assigned a magnitude of 0.03 .

During gambling blocks, mood ratings occurred after every 2 or 3 trials (on average, 1 rating every 2.4 trials). Every subject received mood ratings after the same set of trials.

At the end of the task, participants were presented with their overall point total. These point totals were translated into a cash bonus of \$1–6 depending on their performance. Bonus cutoffs were determined based on simulations such that any value 1–6 were possible to

achieve, but a typical subject gambling at every opportunity could be expected to receive approximately \$3. Upon payment, participants received \$8 for their participation (this was later increased to \$10) plus this bonus.

Survey—After performing the task, online adult participants were asked to complete a series of questionnaires. In the demographics portion, they were asked for their age, gender and location (city and state). They were also asked to indicate their overall status using the MacArthur Scale of Subjective Social Status.⁸⁹ Shown a ten-rung ladder, participants clicked on the rung that represented their overall status relative to others in the United States. This scale is a widely used indicator of subjective social status, and in certain cases, it has been shown to indicate health status better than objective measures of socioeconomic status.⁹⁰

After the demographics portion, online adult participants completed questionnaires including the Center for Epidemiologic Studies Depression Scale (CES-D), a 20-item scale of depressive symptoms.⁹¹ They also completed the Snaith-Hamilton Pleasure Scale (SHAPS), a 14-item scale of hedonic capacity.⁹²

In-person adolescent participants completed a different set of questionnaires, selected to be age-appropriate and maintain consistency with other ongoing research projects. These questionnaires included the Short Child Self-Report Mood and Feelings Questionnaire (MFQ), a 13-item scale of how the participant has been feeling and acting recently.^{79,93} They also included the Screen for Child Anxiety Related Emotional Disorders (SCARED), a 41-item scale of childhood anxiety.⁹⁴ These questionnaires were completed before the subject began completing the online tasks described above.

Participants recruited for follow-up investigations of boredom, mind-wandering, and free time activities also completed the short boredom proneness scale (SBPS), an 8-item scale of an individual's proneness to boredom in everyday life.⁴⁵ They also completed the 5-item mind-wandering questionnaire (MWQ), which quantifies a person's proneness to mind-wandering in everyday life.⁵⁰ The SBPS and MWQ were used to quantify trait-level boredom and mind-wandering, respectively.

Mobile App—The task given to mobile app participants is outlined in Figure 1B. Mobile app participants completed 30 trials of a gambling game. In each trial, participants chose between a certain option and a gamble, represented as a spinner in a circle with two possible outcomes. If the participant chose to gamble, the spinner rotated for approximately 5 seconds before coming to rest on one of the two outcomes. Participants were equally likely to win or lose if they chose to gamble. The points were added to or subtracted from the participant's total during an approximately 2-second inter-trial interval before the game advanced to the next trial. After every 2–3 trials (12 times per play), the participant rated their mood. They were presented with the question, “How happy are you right now?”. A slider was presented with a range from “very unhappy” to “very happy.” The participant could select a value by moving their finger on the slider and tapping “Continue”. No limit was placed on their reaction times.

Each participant received 11 gain trials (with gambles between one positive outcome and one zero), 11 loss trials (one negative outcome and one zero), and 8 mixed trials (one positive and one negative outcome). The possible gambling outcomes were randomly drawn from a list of 60 gain trials, 60 loss trials, and 30 mixed trials. Participants played one of two versions of the app, between which the only difference was the precise win, loss, and certain amounts in these lists. The amounts in the first version are described in detail in the supplementary material of.³ In the second version, gain trials had 3 certain amounts (35, 45, 55) and 15 gamble amounts (59, 66, 72, 79, 85, 92, 98, 105, 111, 118, 124, 131, 137, 144, 150). As in the first version, the set of loss trials was identical to the gain trials except that the values were negative. Mixed trials has 3 prospective gains (40, 44, 75) and 10 prospective losses (-10, -19, -28, -37, -46, -54, -63, -72, -81, -90). Both versions are described further in.³³ The median participant played the game for approximately 5 minutes.

After playing the game, participants saw their score plotted against those of other players, and they were told if their score was a “new record” for them. They could then choose to play again and try to improve their score. We reasoned that introducing the notion of a “new record” would significantly change participants’ motivations and behaviour on subsequent runs, and we therefore limited our analysis to the first run from each participant.

Linear Mixed Effects Model

Analyses and statistics were performed using custom scripts written in Python 3. Participants’ momentary subjective mood ratings were fitted with a linear mixed effects (LME) model with rating time as a covariate using the Pymer4 software package (<http://eshinjolly.com/pymer4/>).⁹⁵ Rating times were converted to minutes to satisfy the algorithm’s convergence criteria while maintaining interpretability. This method resulted in each participant’s data being modelled by a slope and intercept parameter such that:

$$M(t) = M_0 + \beta_T * T(t) \quad (1)$$

where M_0 is the estimated mood at block onset (intercept), β_T is the estimated change in mood per minute (slope), and $T(t)$ is the time in minutes from the start of the block. The LME modeling algorithm also produced a group-level slope and intercept term as well as confidence intervals and statistics testing against the null hypothesis that the true slope or intercept was zero.

The first block of the first run for all online adult and in-person adolescent cohorts experiencing rest or random gambling first were fitted together in a single model, with factors:

$$Mood \sim 1 + Time * (isMale + meanIRIOver20 + totalWinnings + meanRPE + fracRiskScore + isAge0to16 + isAge16to18 + isAge40to100) + (Time | Subject) \quad (2)$$

isMale is 1 if the participant reported their gender as “male,” 0 otherwise. meanIRIOver20 is the mean inter-rating interval across the block(s) of interest (in seconds) minus 20 (a round number near the mean). totalWinnings is the total points won by the participant in the block(s). meanRPE is the mean reward prediction error across the block(s). totalWinnings

and meanRPE will be zero for participants who were experiencing rest instead of gambling. fracRiskScore is the participant's clinical depression risk score divided by a clinical cutoff: i.e., their MFQ score divided by 12 or their CES-D score divided by 16.

While the bounded mood scale prevents the error term of our mood models from being truly Gaussian, LMEs are typically robust to such non-Gaussian distributions.³⁵

For reliability analyses, the first block of each run was modelled separately for each cohort/run with the same model shown above. An intraclass correlation coefficient quantifying absolute agreement (ICC(2,1)) between the runs of each cohort, was calculated using R's "psych" package, accessed through the python wrapper package rpy2.

To measure the psychometric validity of the subjective momentary mood ratings, we correlated the initial mood (or "Intercept") parameter of this model with the life happiness ratings. The correlation was highly significant ($r_s = 0.548$, $p < 0.001$, 2-sided, Extended Data Figure 8, left).

For comparisons with the online data, the same model was also employed in the initial analysis of the mobile app data.

LME Model Comparisons

To compare the ability of additional terms like depression risk and state boredom to explain variance in our model of mood, we employed an ANOVA that compared two models: a reduced model with the factor but without its interaction with time, and an expanded model with both the factor and its interaction with time. All factors in Equation 2 were included in both models (except in the case of depression risk, where the reduced model contained fracRiskScore but not its interaction with Time). We then used R's ANOVA function to compare the expanded and reduced model. The degrees of freedom were quantified as the difference in the number of parameters in the two models.

To examine the impact of including a factor(s) on mood variance explained, we used the within-individual and between-individual variance explained (R_1^2 and R_2^2) as defined in.^{96,97} This calculation required a null model including only an intercept and random effects, which we defined as:

$$\text{Mood} \sim 1 + (1 + \text{Time} \mid \text{Subject}) \quad (3)$$

The within-individual variance R_1^2 of each model was defined as:

$$R_1^2 = 1 - \frac{\sigma_\epsilon^2 + \sigma_\alpha^2}{\sigma_{\epsilon_0}^2 + \sigma_{\alpha_0}^2} \quad (4)$$

where σ_ϵ^2 is the variance of the residuals of the model, σ_α^2 is the variance of the random effects, $\sigma_{\epsilon_0}^2$ is the variance of the residuals of the null model, and $\sigma_{\alpha_0}^2$ is the variance of the random effects in the null model. The variance of the random effects in a model was

calculated using R's MuMIn library,⁹⁸ taking into account the correlation between model factors.

The between-individual variance R_2^2 of each model was defined as:

$$R_2^2 = 1 - \frac{\sigma_\epsilon^2 + \sigma_a^2/k}{\sigma_{\epsilon 0}^2 + \sigma_{a0}^2/k} \quad (5)$$

where k was defined as the harmonic mean of the number of mood ratings being modelled for each participant.

Because the depression risk, boredom, and mind-wandering factors were constant for each subject, we focus primarily on the between-individual variance explained R_2^2 .

To compare the variance explained by the expanded and reduced models as a measure of effect size, we used Cohen's f^2 statistic,^{39, 40} defined as:

$$f^2 = \frac{R_{AB}^2 - R_A^2}{1 - R_{AB}^2} \quad (6)$$

Where R_{AB}^2 is the variance explained by the expanded model and R_A^2 is the variance explained by the reduced model. Separate f^2 values can be calculated using the within-individual or between-individual variances. Using Cohen's guidelines,³⁹ $f^2 = 0.02$ is considered a small effect, $f^2 = 0.15$ is considered a medium effect, and $f^2 = 0.35$ is considered a large effect.

Computational Model

When examining the effect of time on mood during random gambling in the mobile app data, we next attempted to disentangle time's effects from those of reward and expectation using a computational model. The model is based on one described in detail by² that has been validated on behavioural data from a similar gambling task. The authors found that changes in momentary subjective mood were predicted accurately by a weighted combination of current and past rewards and RPEs in the task. Quantifying RPEs relies on subjective expectations that are formulated according to a "primacy model," in which expected reward is more heavily influenced by early rewards than it is by recent ones.

The model described in² was modified to include a coefficient β_T that linearly relates time and mood. Our modified model is defined as follows:

$$\widehat{M}(t) = M_0 + \beta_A \sum_{u=1}^t \lambda^{t-u} A(u) + \beta_E \sum_{u=1}^t \lambda^{t-u} E(u) + \beta_T T(t) \quad (7)$$

In the above equation, $t = 1, 2, \dots, n$ is the trial index, and $\widehat{M}(t)$ is the estimated mood rating from trial t . M_0 (the estimated mood at time 0), λ (an exponential discounting factor), and the β s are learned parameters of the model. $A(t)$ is the actual outcome (in hundreds of

points) of trial t , $T(t)$ is the time of trial t in minutes, and $E(t)$ is the primacy model of the subject's reward expectation in trial t , defined as:

$$E(t) = \frac{1}{t-1} \sum_{u=1}^{t-1} A(u) \quad (8)$$

If we remove the influence of time (i.e., set our $\beta_T = 0$), the full mood model in² is equivalent to this one as long as its reward prediction error coefficient is less than its expectation coefficient (i.e., $\beta_R^{Keren} < \beta_E^{Keren}$) and $\beta_E^{Keren} > 0$, where β_R^{Keren} and β_E^{Keren} denote the values β_R and β_E defined in²). The values in our model can be derived from the values in theirs by setting $\beta_A = \beta_R^{Keren}$ and $\beta_E = \beta_E^{Keren} - \beta_R^{Keren}$.

We used the PyTorch package⁹⁹ on a GPU to fit 500 models simultaneously for each subject. β_T was initialised to random values with distribution $\mathcal{N}(0, 1)$. β_E and β_A were initialised to random values with distribution *Lognormal*(0, 1) and capped to the interval [0,10] on every iteration. M_0 and λ were initialised to random values with normal distributions $\mathcal{N}(0, 1)$, then sigmoid-transformed (to facilitate optimization and conform to the interval [0, 1]) using the standard logistic function:

$$y = \frac{1}{1 + e^{-x}} \quad (9)$$

At the end of 100,000 iterations, the model with the lowest sum of squared errors (SSE) (i.e., $\sum_{t=1}^N (\widehat{M}(t) - M(t))^2$) was selected. The time coefficient β_T learned by the model could then be used as a measure of the influence of time on that participant's mood, disentangled from the effects of rewards and RPEs.

End-to-end optimization was carried out using ADAM¹⁰⁰ with a learning rate of $\alpha = 0.005$. L2 penalty terms were placed on the β terms and added to the sum of squared errors. This meant that the objective function being minimised was:

$$L = \sum_{t=1}^n (\widehat{M}(t) - M(t))^2 + \lambda_{EA} * (\beta_A^2 + \beta_E^2) + \lambda_T * \beta_T^2 \quad (10)$$

The regularization hyperparameters λ_{EA} and λ_T were determined from a tuning step, in which the model was trained on the first 10 mood ratings and tested on the last two in each of 5,000 exploratory participants. One model was trained with each combination of λ_{EA} and λ_T ranging from 10^{-4} to 10^3 in 20 steps (evenly spaced on a log scale). The testing loss (median across participants) across penalty terms was fitted to a third degree polynomial using Skikit-Learn's kernel ridge regression with regularization strength $\alpha = 10.0$. The best fitting regularization hyperparameters were defined as those that minimised this smoothed testing loss.

As in the LME, the bounded mood scale prevents the error term of our mood models from being truly Gaussian. Our computational model attempted to mitigate the effect of non-Gaussianity by capping mood predictions to the allowable range, initialising parameters to non-normal distributions, and restricting parameters to feasible ranges on every iteration.

As in the online cohort's LME model, the initial mood parameter M_0 showed psychometric validity. It was significantly correlated with life happiness ($r_s = 0.362$, $p < 0.001$, Extended Data Figure 8, right).

Control Model—To quantify the effect of including the time-related term, we fitted a control model without β_T . This control model is defined as follows:

$$\widehat{M}(t) = M_0 + \beta_A \sum_{u=1}^t \lambda^{t-u} A(u) + \beta_E \sum_{u=1}^t \lambda^{t-u} E(u) \quad (11)$$


As in the primary model, the regularization hyperparameter λ_{EA} in this control model was tuned using the method described above.

Data Availability

All data used in the manuscript have been made publicly available. Online Participants' data can be found on the Open Science Framework at <https://osf.io/km69z/>. Mobile App Participants' data can be found on Dryad at <https://doi.org/10.5061/dryad.prr4xgxkk>.¹⁰¹


Code Availability

The code for the task and survey is available on GitLab at <https://gitlab.pavlovia.org/mooddrift>. Our data analysis software, as well as the means to create a Python environment that automatically installs it on a user's machine, has been made available online at <https://github.com/djangraw/MoodDrift>.

Extended Data**Extended Data Fig. 1.**

Mood rating frequency does not affect mood drift slope

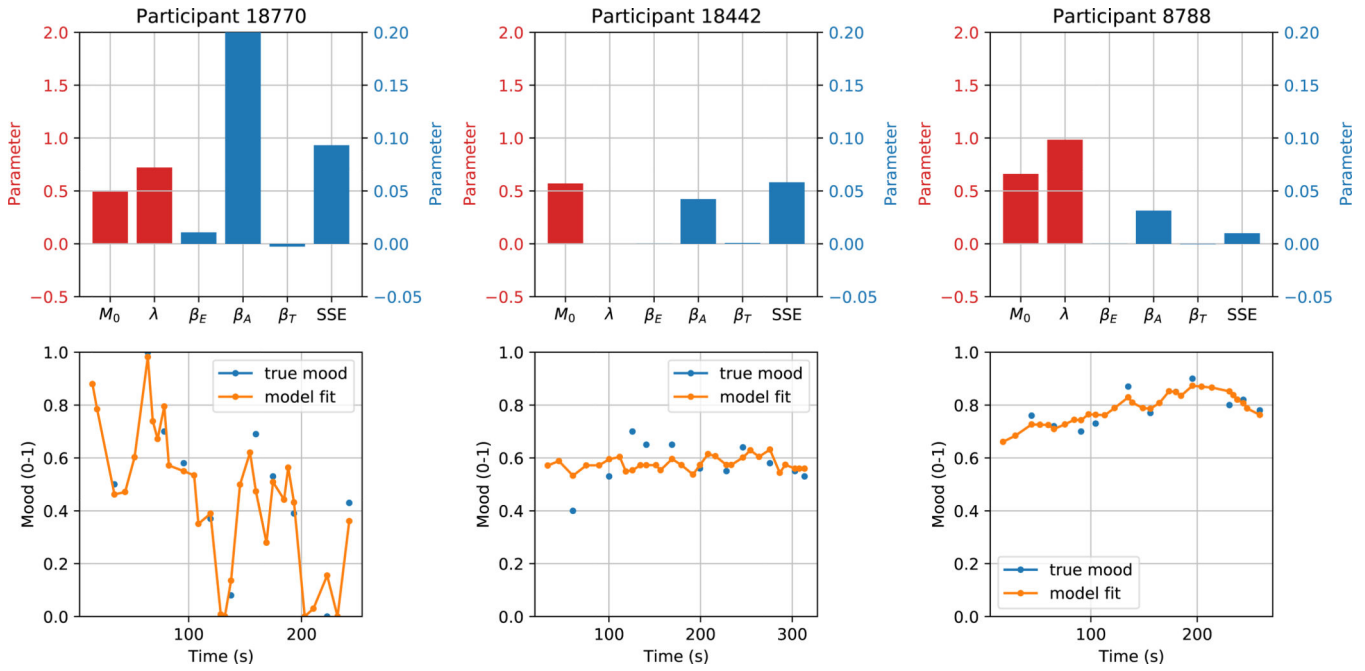
Mean \pm STE mood rating at each time in the 4 cohorts with 60 s, 30 s, 15 s, and 7.5 s of rest between mood ratings (cohorts 60sRestBetween, 30sRestBetween, 15sRestBetween, and 7.5sRestBetween, respectively). The magnitude of mood drift did not vary with the frequency of mood ratings.



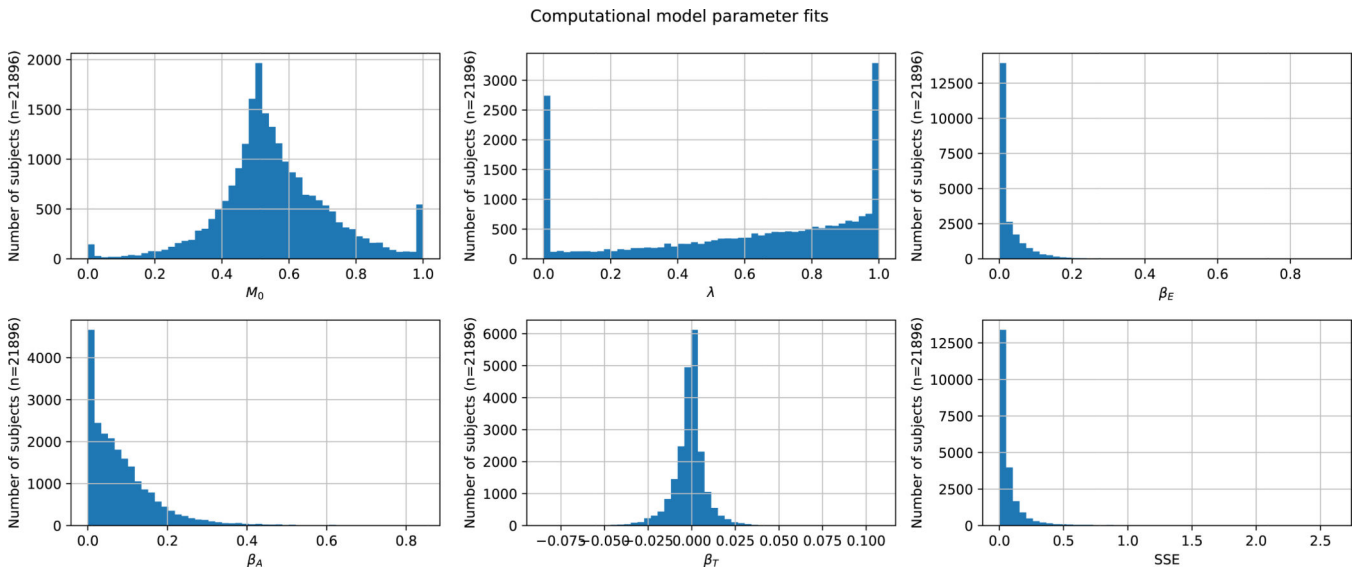
Extended Data Fig. 2.

Mood slope parameter distributions vary with analysis choice

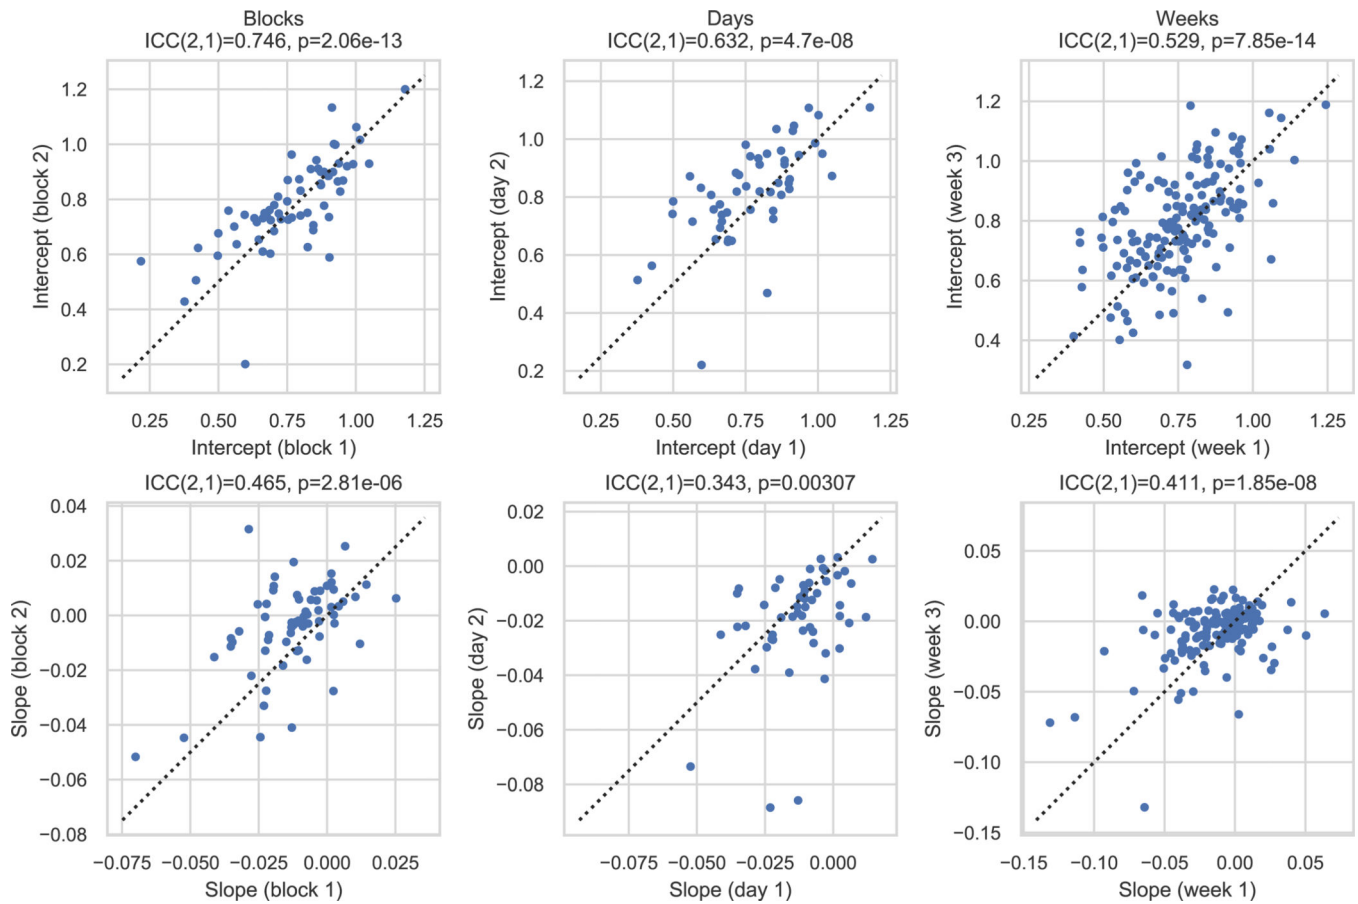
Histogram of the LME mood slope parameters for the online cohort (blue) and the confirmatory mobile app cohort (orange), along with the computational model time sensitivity parameter for the confirmatory mobile app cohort (green). Mobile app participants with outlier task completion times were excluded from the LME analysis (see Methods). Note that the use of LME modeling to analyze the mobile app data significantly lowered the distribution of slopes compared to when the computational model was used (median = -0.752 vs. -0.0408 , IQR = 2.10 vs. 0.764 %mood/min, 2-sided Wilcoxon rank-sum test, $W_{42771} = -54.2$, $p < 0.001$), but the LME slopes from the mobile app were still significantly greater than those of the online cohort (median = -1.53 vs. 0.752 , IQR = 2.34 vs. 2.1 %mood/min, 2-sided Wilcoxon rank-sum test, $W_{21761} = 14.5$, $p < 0.001$). Vertical lines represent group medians. Stars indicate $p < 0.05$. P values were not corrected for multiple comparisons.



Extended Data Fig. 3.
 Sample fits of the computational model
 Sample fits of the computational model for three random subjects in the confirmatory mobile app cohort. SSE = sum squared error, a measure of goodness of fit to the training data. In the top plots, the red bars are in units of the left-hand y axis, and the blue bars are in units of the right-hand y axis.



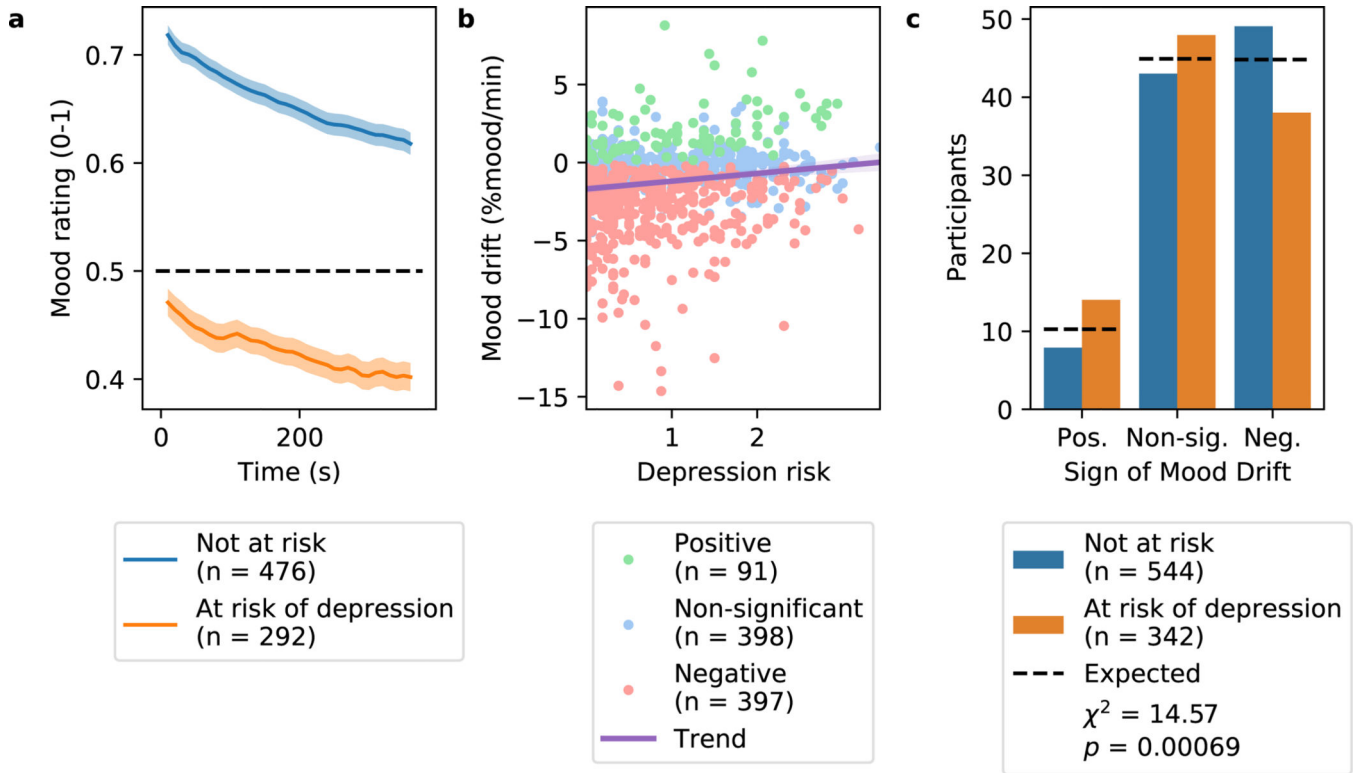
Extended Data Fig. 4.
 Histogram of computational model parameters
 Histogram of computational model parameters across the 21,896 confirmatory mobile app subjects.



Extended Data Fig. 5.

Mood drift stability over blocks, days, and weeks

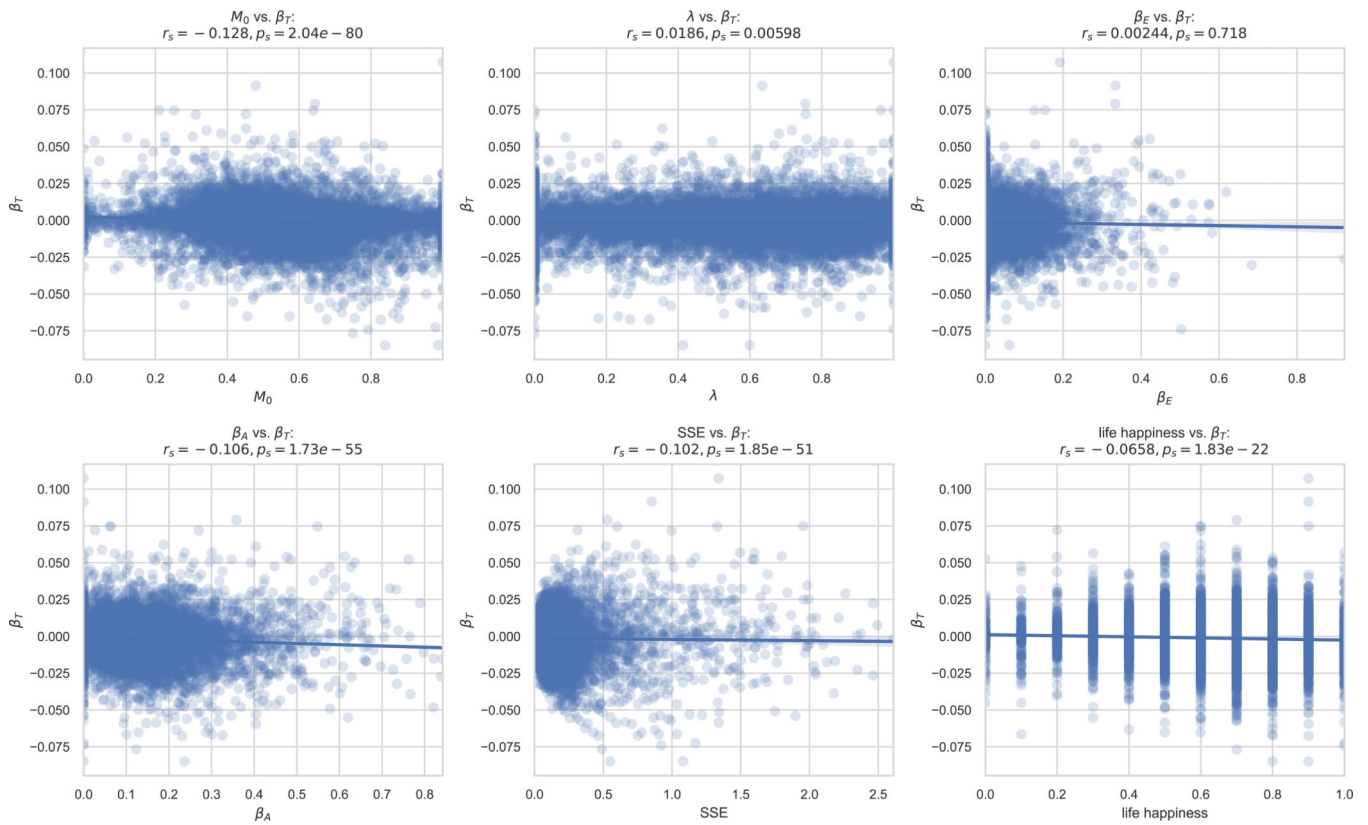
Stability of LME coefficients estimating the initial mood (top) and slope of mood over time (bottom) for each participant across rest periods one block apart (left), 1 day apart (middle), and 2 weeks apart (right). ICC denotes the intra-class correlation coefficient for each comparison. P values shown are one-sided (since ICC values are expected to be positive) with no correction for multiple comparisons.



Extended Data Fig. 6.

Relationship between mood drift and depression risk

Relationship between mood drift and depression risk. (a) Mood ratings over time of online participants at risk of depression (defined as MFQ>12 or CES-D>16) vs. those not at risk for the 768 participants with at least 6 minutes of resting mood data (error bars are SEM). The dotted line represents the mean initial rating (mean of cohort means). (b) We fitted simple regressions of time versus mood within each individual and determined significance of the time term with BenjaminiHochberg false-discovery rate correction (2-sided $\alpha = 0.5$, $p < 0.05$) to better understand the relationship between depression risk and the change in mood over time. Depression risk is operationalised as score on the CES-D or MFQ divided by the threshold for depression risk on each measure (16 and 12 respectively). The line is a linear best fit, and the patch shows the 95% confidence interval of this fit. (c) Proportion of individuals with or without risk of depression (i.e., depression risk >1 or <1) with positive (significantly greater than zero), non-significant (no evidence of a significant difference from zero), and negative (significantly less than 0) slopes of mood over time. 13 more individuals at risk of depression have a positive slope than the 35 expected based on the rates in individuals not at risk of depression, $\chi^2(1, N=886)=14.57$, $p < 0.001$ (2-sided Pearson's chi-squared statistic with no correction for multiple comparisons).

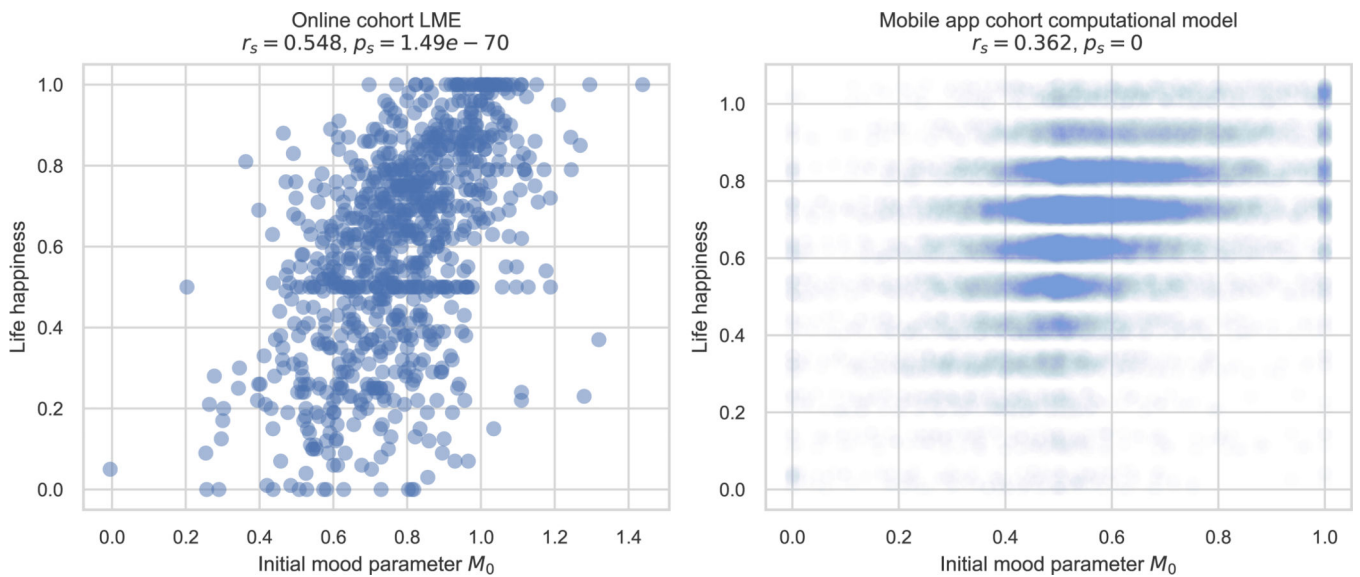


Extended Data Fig. 7.

Mood drift's relation to other computational model parameters

Time sensitivity parameter β_T vs. other parameters in the confirmatory mobile app cohort.

Each dot is a participant ($n=21,896$). Each line is a linear best fit, and patches show the 95% confidence interval of this fit. r_s denotes Spearman correlation coefficient. P values shown are 2-sided with no correction for multiple comparisons.



Extended Data Fig. 8.

Initial mood parameter’s relation to life happiness

Initial mood parameter vs. life happiness rating in the online cohort (left) and the confirmatory mobile app cohort (right). Life happiness ratings were always multiples of 0.1; small positive random values were added during plotting to reduce overlap between data points. Each dot is a participant (left: n=886, right: n=21,896). r_s denotes Spearman correlation coefficient. P values shown are 2-sided with no correction for multiple comparisons.

Extended Data Table 1

List and description of cohorts collected

A list and description of cohorts collected. nParticipants contains the number of participants who completed both the task and survey in this cohort. The columns beginning with “Block” denote the type, parameter, and number of cycles used in that block. “Rest” denotes looking at a fixation cross, and “task” denotes a simple visuomotor task in which a cross moves predictably across the screen and the subject is asked to press a button when it crosses the center line. The number that follows these labels is the time in seconds between mood ratings. “Break” denotes a free period where participants could leave to do anything they chose. “Closed” and “random” denote the closed-loop and random gambling task conditions described in the Methods section. (“open” denotes open-loop gambling not described in this paper; these blocks were not used in analyses). The + or – after the “closed” label indicates whether mood was being manipulated upwards (+) or downwards (–). The number after the * indicates how many cycles of this type were included in the block. Certain cohort names also contain information. The AlternateRating cohort rated their mood with a single button press rather than moving a slider. The Expectation cohorts received opening instructions stating that the upcoming rest period would be up to 7 minutes or 12 minutes. Groups beginning with “Daily” or “Weekly” returned 1 day or 1 week apart to complete a similar task again (e.g., the Daily-Rest-02 cohort is the same participants as Daily-Rest-01, returning to complete the same task one day later). The Adolescent-01 cohort is a group of adolescents recruited in person rather than on Amazon Mechanical Turk.

Opening Rest Cohort	nParticipants	Block 0	Block 1	Block 2	Block 3
15sRestBetween	40	rest15 * 30	closed+ * 54		
30sRestBetween	37	rest30 * 18	closed+ * 54		
7.5sRestBetween	38	rest7.5 * 45	closed+ * 54		
60sRestBetween	39	rest60 * 10	closed+ * 54		
AlternateRating	32	rest15 * 30	closed+ * 54		
Expectation-7mRest	64	rest15 * 18	random * 22	closed– * 22	closed+ * 22
Expectation-12mRest	67	rest15 * 18	random * 22	closed– * 22	closed+ * 22
RestDownUp	58	rest15 * 18	closed– * 33	closed+ * 33	
Daily-Rest-01	66	rest15 * 18	closed+ * 18	rest15 * 18	closed+ * 18
Daily-Rest-02	53	rest15 * 18	closed+ * 18	rest15 * 18	closed+ * 18
Weekly-Rest-01	196	rest15 * 18	closed+ * 22	closed– * 22	closed+ * 22
Weekly-Rest-02	164	rest15 * 18	open+ * 22	open– * 22	open+ * 22

Opening Rest Cohort	nParticipants	Block 0	Block 1	Block 2	Block 3
Weekly-Rest-03	160	rest15 * 18	open+ * 22	open- * 22	open+ * 22
Adolescent-01	116	rest15 * 18	closed+ * 22	closed- * 22	closed+ * 22
Opening Task Cohort					
Visuomotor	37	task15 * 30	closed+ * 54		
Visuomotor-Feedback	30	task15 * 30	closed+ * 54		
Opening Gambling Cohort					
RestAfterWins	25	closed+ * 54	rest15 * 30		
Daily-Closed-01	68	closed+ * 32	closed- * 32	closed+ * 32	
Daily-Random-01	66	random * 32	random * 32	random * 32	
App-Exploratory	5000	random * 30			
App-Confirmatory	21896	random * 30			
Follow-Up Cohorts					
BoredomBeforeAndAfter	150	rest15 * 18	closed- * 33	closed+ * 33	
BoredomAfterOnly	150	rest15 * 18	closed- * 33	closed+ * 33	
MwBeforeAndAfter	150	rest15 * 18	closed- * 33	closed+ * 33	
MwAfterOnly	150	rest15 * 18	closed- * 33	closed+ * 33	
Activities	450	break420 * 1	closed- * 33	closed+ * 33	

Extended Data Table 2

Linear mixed effects (LME) model results

Results of the LME model trained on all naïve online adult and adolescent participants who received opening rest, visuomotor task, or random gambling periods; as produced by the pyme software package. The first column lists each factor in the model as described in the Methods section. Factors beginning with “is” are binary (0 or 1). “Time” is the mood slope parameter we use to quantify mood drift. Mood ratings ranged from 0–1, and time was in minutes. totalWinnings and meanRPE were in points, whose monetary value is unknown to naïve subjects. fracRiskScore was the score on a clinical depression questionnaire divided by a clinical cutoff. Age was in years. Factors preceded by “Time:” indicate the interaction of that parameter and the elapsed time. The next four columns describe the effect size: “Estimate” is the estimated coefficient of each factor in the model, 2.5 and 97.5 ci are the 95 percent confidence interval of the estimate, and SE is its standard error. DF is the degrees of freedom, T-stat is the t statistic, and P-val is the 2-sided p value. All values are rounded to 3 decimal places. The Sig (significance) column contains.

Factor	Estimate	2.5_ci	97.5_ci	SE	DF	T-stat	P-val	Sig
(Intercept)	0.784	0.756	0.812	0.0141	875	55.6	< 10 ⁻⁶	*
Time	-0.0189	-0.0226	-0.0153	0.00185	864	-10.3	< 10 ⁻⁶	*
isMale	-0.0144	-0.0395	0.0107	0.0128	877	-1.12	0.262	

Factor	Estimate	2.5_ci	97.5_ci	SE	DF	T-stat	P-val	Sig
meanIRIOver20	0.000698	-0.000585	0.00198	0.000655	901	1.07	0.287	
totalWinnings	-0.000332	-0.00435	0.00369	0.00205	898	-0.162	0.872	
meanRPE	0.158	-0.0104	0.326	0.0859	898	1.84	0.0662	
fracRiskScore	-0.186	-0.202	-0.169	0.00828	877	-22.4	< 10 ⁻⁶	*
isAge0to16	-0.0456	-0.108	0.0168	0.0318	879	-1.43	0.152	
isAge16to18	-0.0883	-0.144	-0.0325	0.0285	879	-3.1	0.002	*
isAge40to100	-0.00712	-0.0351	0.0208	0.0143	877	-0.5	0.617	
Time:isMale	0.00159	-0.00171	0.00488	0.00168	869	0.944	0.345	
Time:meanIRIOver20	-0.000103	-0.000267	6.1 * 10 ⁻⁵	8.4 * 10 ⁻⁵	810	-1.23	0.219	
Time:totalWinnings	-1.9 * 10 ⁻⁵	-0.000566	0.000529	0.00028	1.04 * 10 ³	-0.0664	0.947	
Time:meanRPE	-0.00743	-0.0304	0.0155	0.0117	1.05 * 10 ³	-0.634	0.526	
Time:fracRiskScore	0.00515	0.00303	0.00728	0.00109	869	4.75	2 * 10 ⁻⁶	*
Time:isAge0to16	-0.00144	-0.00967	0.00678	0.0042	895	-0.344	0.731	
Time:isAge16to18	0.00869	0.00131	0.0161	0.00376	898	2.31	0.0212	*
Time:isAge40to100	0.00302	-0.000638	0.00668	0.00187	865	1.62	0.106	

* if p<0.05

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

This research was supported in part by the Intramural Research Program of the National Institute of Mental Health, part of the National Institutes of Health (NIH) (Grant Nos. ZIAMH002957 [to AS], ZICMH002968 [to FP], ZIAMH002871 [to DSP], ZIAMH002872 [to DSP], and ZICMH002960 [to AGT]). This work used the computational resources of the NIH high-performance computing (HPC) Biowulf cluster (<http://hpc.nih.gov>). Data collection for the mobile app dataset was supported by the Wellcome Trust (Grant No. 101252/Z/13/Z). The online adolescent sample in this study was collected under NIH IRB protocol number 18-M-0037, registered on clinicaltrials.gov as NCT03388606. The online adult sample was collected under NIH Office of Human Subjects Research Protection protocol P194594. The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript. The views expressed in this article do not necessarily represent the views of the National Institutes of Health, the Department of Health and Human Services, or the United States Government.

References

1. Penny WD, Friston KJ, Ashburner JT, Kiebel SJ & Nichols TE Statistical Parametric Mapping: The Analysis of Functional Brain Images (Elsevier Science, 2011).
2. Keren H et al. The temporal representation of experience in subjective mood. *eLife* 10, 1–24 (2021). 10.7554/elife.62051.
3. Rutledge RB, Skandali N, Dayan P & Dolan RJ A computational and neural model of momentary subjective well-being. *Proceedings of the National Academy of Sciences of the United States of America* 111 (33), 12252–12257 (2014). 10.1073/pnas.1407535111. [PubMed: 25092308]
4. Frijda N, Mesquita B, Sonnemans J & Goozen S The duration of affective phenomena or emotions, sentiments and passions, Vol. 1 187–225 (1991).

5. Scherer KR & Wallbott HG Evidence for universality and cultural variation of differential emotion response patterning. *Journal of Personality and Social Psychology* 66 (2), 310–328 (1994). 10.1037//0022-3514.66.2.310. [PubMed: 8195988]
6. Davidson RJ Affective Style and Affective Disorders: Perspectives from Affective Neuroscience. *Cognition and Emotion* 12 (3), 307–330 (1998). 10.1080/026999398379628.
7. Davidson RJ Comment: Affective Chronometry Has Come of Age. *Emotion Review* 7 (4), 368–370 (2015). 10.1177/1754073915590844.
8. Gilboa E & Revelle W *Personality and the Structure of Affective Responses* (Psychology Press, 1994).
9. Hemenover SH Individual differences in rate of affect change: studies in affective chronometry. *Journal of personality and social psychology* 85 (1), 121 (2003). [PubMed: 12872888]
10. Kring AM & Barch DM The motivation and pleasure dimension of negative symptoms: Neural substrates and behavioral outputs. *European Neuropsychopharmacology* 24 (5), 725–736 (2014). 10.1016/J.EURONEURO.2013.06.007. [PubMed: 24461724]
11. Sonuga-Barke EJS, Taylor E, Sembi S & Smith J Hyperactivity and delay aversion—I. The effect of delay on choice. *Journal of Child Psychology and Psychiatry* 33 (2), 387–398 (1992). [PubMed: 1564081]
12. Solanto MV et al. The ecological validity of delay aversion and response inhibition as measures of impulsivity in AD/HD: A supplement to the NIMH multimodal treatment study of AD/HD. *Journal of Abnormal Child Psychology* 29 (3), 215–228 (2001). 10.1023/A:1010329714819. [PubMed: 11411784]
13. Sonuga-Barke EJS, Cortese S, Fairchild G & Stringaris A Annual Research Review: Transdiagnostic neuroscience of child and adolescent mental disorders—differentiating decision making in attention-deficit/hyperactivity disorder, conduct disorder, depression, and anxiety. *Journal of Child Psychology and Psychiatry* 57 (3), 321–349 (2016). [PubMed: 26705858]
14. McRae TW Opportunity and Incremental Cost: An Attempt to Define in Systems Terms. *The Accounting Review* 45 (2), 315–321 (1970). <https://www.jstor.org/stable/244383>.
15. Hoskin RE Opportunity Cost and Behavior. *Journal of Accounting Research* 21 (1), 78–95 (1983).
16. Palmer S & Raftery J Opportunity cost. *BMJ* 318 (7197), 1551–1552 (1999). [PubMed: 10356019]
17. Cohen JD, McClure SM & Yu AJ Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences* 362 (1481), 933–942 (2007).
18. Constantino SM & Daw ND Learning the opportunity cost of time in a patch-foraging task. *Cognitive, Affective, & Behavioral Neuroscience* 15 (4), 837–853 (2015).
19. Addicott MA, Pearson JM, Sweitzer MM, Barack DL & Platt ML A primer on foraging and the explore/exploit trade-off for psychiatry research. *Neuropsychopharmacology* 42 (10), 1931–1939 (2017) [PubMed: 28553839]
20. Geana A, Wilson R, Daw ND & Cohen JD Boredom, Information-Seeking and Exploration. *Proc. 38th Annual Conference of the Cognitive Science Society* (2016).
21. Agrawal M, Mattar MG, Cohen JD & Daw ND The temporal dynamics of opportunity costs: A normative account of cognitive fatigue and boredom. *Psychological review* 129 (3), 564–585 (2022). 10.1037/rev0000309. [PubMed: 34383523]
22. Eastwood JD, Frischen A, Fenske MJ & Smilek D The Unengaged Mind: Defining Boredom in Terms of Attention. *Perspectives on Psychological Science* 7 (5), 482–495 (2012). 10.1177/1745691612456044. [PubMed: 26168505]
23. Robison MK, Miller AL & Unsworth N A multi-faceted approach to understanding individual differences in mind-wandering. *Cognition* 198 (September 2019), 104078 (2020). 10.1016/j.cognition.2019.104078. [PubMed: 32062086]
24. Killingsworth MA & Gilbert DT A wandering mind is an unhappy mind. *Science* 330 (6006), 932 (2010). 10.1126/science.1192439. [PubMed: 21071660]
25. Fox KC, Thompson E, Andrews-Hanna JR & Christoff K Is thinking really aversive? A commentary on Wilson et al.'s “Just think: The challenges of the disengaged mind”. *Frontiers in Psychology* 5 (DEC), 10–13 (2014). 10.3389/fpsyg.2014.01427. [PubMed: 24478742]

26. Fox KC et al. Affective neuroscience of self-generated thought. *Annals of the New York Academy of Sciences* 1426 (1), 25–51 (2018). 10.1111/NYAS.13740.
27. van Hooff ML & van Hooft EA Boredom at work: Proximal and distal consequences of affective work-related boredom. *Journal of Occupational Health Psychology* 19 (3), 348–359 (2014). 10.1037/a0036821. [PubMed: 24885686]
28. Miner AG & Glomb TM State mood, task performance, and behavior at work: A within-persons approach. *Organizational Behavior and Human Decision Processes* 112 (1), 43–57 (2010). 10.1016/j.obhdp.2009.11.009.
29. Camille N et al. The involvement of the orbitofrontal cortex in the experience of regret. *Science* 304 (5674), 1167–1170 (2004). 10.1126/SCIENCE.1094550/SUPPL_FILE/CAMILLE_SOM.PDF. [PubMed: 15155951]
30. Eldar E, Rutledge RB, Dolan RJ & Niv Y Mood as representation of momentum. *Trends in cognitive sciences* 20 (1), 15–24 (2016). [PubMed: 26545853]
31. Vinckier F, Rigoux L, Oudiette D & Pessiglione M Neuro-computational account of how mood fluctuations arise and affect decision making. *Nature Communications* 9 (1708) (2018). 10.1038/s41467-018-03774-z.
32. Liuzzi L et al. Magnetoencephalographic correlates of mood and reward dynamics in human adolescents. *Cerebral Cortex* 32 (15), 3318–3330 (2022). 10.1093/cercor/bhab417. [PubMed: 34921602]
33. Bedder RL, Vaghi MM, Dolan RJ & Rutledge RB Risk taking for potential losses but not gains increases with time of day. *psyarxiv* (2020). 10.31234/osf.io/3qdnxarXiv:psyarxiv.com/3qdnx/.
34. Grilli L & Rampichini C Specification of random effects in multilevel models: a review. *Quality & Quantity* 49 (3), 967–976 (2015).
35. Schielzeth H et al. Robustness of linear mixed-effects models to violations of distributional assumptions. *Methods in Ecology and Evolution* 11 (9), 1141–1152 (2020).
36. Feingold A Confidence interval estimation for standardized effect sizes in multilevel and latent growth modeling. *Journal of consulting and clinical psychology* 83 (1), 157 (2015). [PubMed: 25181028]
37. Pizzagalli DA, Iosifescu D, Hallett LA, Ratner KG & Fava M Reduced hedonic capacity in major depressive disorder: Evidence from a probabilistic reward task. *Journal of Psychiatric Research* 43 (1), 76–87 (2008). 10.1016/J.JPSYCHIRES.2008.03.001. [PubMed: 18433774]
38. Halahakoon DC et al. . Reward-processing behavior in depressed participants relative to healthy volunteers: a systematic review and meta-analysis. *JAMA Psychiatry* 77, 1286–1295 (2020). [PubMed: 32725180]
39. Cohen J *Statistical Power Analysis for the Behavioral Sciences* (Routledge, 2013). <https://www.taylorfrancis.com/books/9781134742707>.
40. Selya AS, Rose JS, Dierker LC, Hedeker D & Mermelstein RJ A Practical Guide to Calculating Cohen's f^2 , a Measure of Local Effect Size, from PROC MIXED. *Frontiers in Psychology* 3, 111 (2012). 10.3389/fpsyg.2012.00111. [PubMed: 22529829]
41. Isen AM & Patrick R The effect of positive feelings on risk taking: When the chips are down. *Organizational behavior and human performance* 31 (2), 194–202 (1983).
42. Arkes HR, Herren LT & Isen AM The role of potential loss in the influence of affect on risk-taking behavior. *Organizational Behavior and Human Decision Processes* 42 (2), 181–193 (1988). 10.1016/0749-5978(88)90011-8.
43. Schulreich S et al. Music-evoked incidental happiness modulates probability weighting during risky lottery choices. *Frontiers in psychology* 4, 981 (2014). [PubMed: 24432007]
44. Hunter JA, Dyer KJ, Cribbie RA & Eastwood JD Exploring the utility of the Multidimensional State Boredom Scale. *European Journal of Psychological Assessment* 32 (3), 241–250 (2016). 10.1027/1015-5759/a000251.
45. Struk AA, Carriere JSA, Cheyne JA & Danckert J A short boredom proneness scale: Development and psychometric properties. *Assessment* 24 (3), 346–359 (2017). [PubMed: 26467085]
46. Seli P et al. Mind-Wandering as a Natural Kind: A Family-Resemblances View. *Trends in Cognitive Sciences* 22 (6), 479–490 (2018). 10.1016/J.TICS.2018.03.010. [PubMed: 29776466]

47. Christoff K et al. Mind-Wandering as a Scientific Concept: Cutting through the Definitional Haze. *Trends in Cognitive Sciences* 22 (11), 957–959 (2018). 10.1016/j.tics.2018.07.004. [PubMed: 30220476]
48. Seli P et al. The Family-Resemblances Framework for Mind-Wandering Remains Well Clad. *Trends in Cognitive Sciences* 22 (11), 959–961 (2018). 10.1016/j.tics.2018.07.007. [PubMed: 30220475]
49. Turnbull A et al. The ebb and flow of attention: Between-subject variation in intrinsic connectivity and cognition associated with the dynamics of ongoing experience. *Neuroimage* 185 (September 2018), 286–299 (2019). 10.1016/j.neuroimage.2018.09.069. [PubMed: 30266263]
50. Mrazek MD, Phillips DT, Franklin MS, Broadway JM & Schooler JW Young and restless: validation of the Mind-Wandering Questionnaire (MWQ) reveals disruptive impact of mind-wandering for youth. *Frontiers in psychology* 4, 560 (2013). [PubMed: 23986739]
51. Nunokawa J The Importance of Being Bored: The Dividends of Ennui in” The Picture of Dorian Gray”. *Studies in the Novel* 28 (3), 357–371 (1996).
52. Shattuck R Proust’s way: A field guide to in Search of Lost Time (WW Norton & Company, 2001).
53. Proust M Swann’s Way: in Search of Lost Time Vol. 1 (Yale University Press, 2013).
54. Ciocan C Heidegger and the Problem of Boredom. *Journal of the British Society for Phenomenology* 41 (1), 64–77 (2010).
55. Ratcliffe M. in *The Cambridge Companion to Heidegger’s Being and Time* (ed. Wrathall MA) 157–176 (Cambridge University Press, 2013).
56. Heidegger M *The fundamental concepts of metaphysics: World, finitude, solitude* (Indiana University Press, 1995).
57. Schopenhauer A in *Parerga und Paralipomena*, Vol. 1 217 (1851). https://cedires.com/wp-content/uploads/2019/11/Schopenhauer_Arthur_Parerga-und-Paralipomena_full-text-but-modern-typeface.pdf.
58. Kierkegaard S *Either/Or: A Fragment of Life* (Penguin Classics, 1992).
59. Elpidorou A The bright side of boredom. *Frontiers in psychology* 5, 1245 (2014). [PubMed: 25404925]
60. Thompson PM et al. The ENIGMA Consortium: large-scale collaborative analyses of neuroimaging and genetic data. *Brain Imaging and Behavior* 8 (2), 153–182 (2014). 10.1007/s11682-013-9269-5. [PubMed: 24399358]
61. Adhikari BM et al. A resting state fMRI analysis pipeline for pooling inference across diverse cohorts: an ENIGMA rs-fMRI protocol. *Brain Imaging and Behavior* 13 (5), 1453–1467 (2019). 10.1007/s11682-018-9941-x. [PubMed: 30191514]
62. Birn RM et al. The effect of scan length on the reliability of resting-state fMRI connectivity estimates. *Neuroimage* 83, 550–558 (2013). [PubMed: 23747458]
63. Noble S et al. Influences on the Test-Retest Reliability of Functional Connectivity MRI and its Relationship with Behavioral Utility. *Cerebral Cortex* 27 (11), 5415–5429 (2017). 10.1093/CERCOR/BHX230. [PubMed: 28968754]
64. Noble S, Scheinost D & Constable RT A decade of test-retest reliability of functional connectivity: A systematic review and meta-analysis. *Neuroimage* 203, 116157 (2019). 10.1016/J.NEUROIMAGE.2019.116157. [PubMed: 31494250]
65. Jaspers K in *Die abnorme Seele in Gesellschaft und Geschichte (Soziologie und Historie der Psychosen und Psychopathien)* 594–623 (Springer, 1973).
66. Schneider K *Klinische Psychopathologie* 14 edn (Georg Thieme Verlag, Stuttgart, 1992).
67. Berrios GE Phenomenology, psychopathology and Jaspers: a conceptual history. *History of psychiatry* 3 (11), 303–327 (1992). [PubMed: 11612879]
68. Westgate EC & Wilson TD Boring thoughts and bored minds: The MAC model of boredom and cognitive engagement. *Psychological Review* 125 (5), 689 (2018). [PubMed: 29963873]
69. Barrett LF Feelings or words? Understanding the content in self-report ratings of experienced emotion. *Journal of Personality and Social Psychology* 87 (2), 266–281 (2004). 10.1037/0022-3514.87.2.266. [PubMed: 15301632]

70. Westgate EC & Steidle B Lost by definition: Why boredom matters for psychology and society. *Social and Personality Psychology Compass* 14 (11), e12562 (2020).
71. Frijda NH in *Mood* (eds Sander D & Scherer KR) *The Oxford Companion to Emotion and the Affective Sciences* 258–259 (Oxford University Press, New York, 2009).
72. Ekkekakis P. *The Measurement of Affect, Mood, and Emotion: A Guide for Health-Behavioral Research* (Cambridge University Press, 2013).
73. Rottenberg J. Mood and emotion in major depression. *Current Directions in Psychological Science* 14, 167–170 (2005).
74. Nowlis V & Nowlis HH *The Description and Analysis of Mood*. *Annals of the New York Academy of Sciences* 65 (4), 345–355 (1956). 10.1111/J.1749-6632.1956.TB49644.X. [PubMed: 13373215]
75. Ekman P An argument for basic emotions. *Cognition & emotion* 6 (3–4), 169–200 (1992).
76. Watson D *Mood and temperament* (Guilford Press, 2000).
77. Diener E Subjective well-being: The science of happiness and a proposal for a national index. *American psychologist* 55 (1), 34 (2000). [PubMed: 11392863]
78. Robinson MD & Clore GL Belief and feeling: Evidence for an accessibility model of emotional self-report. *Psychological Bulletin* 128 (6), 934–960 (2002). 10.1037/0033-2909.128.6.934. [PubMed: 12405138]
79. Costello EJ & Angold A Scales to Assess Child and Adolescent Depression: Checklists, Screens, and Nets. *Journal of the American Academy of Child & Adolescent Psychiatry* 27 (6), 726–737 (1988). 10.1097/00004583-198811000-00011. [PubMed: 3058677]
80. Pavot W & Diener E The affective and cognitive context of self-reported measures of subjective well-being. *Social Indicators Research* 28 (1), 1–20 (1993). 10.1007/BF01086714.
81. Ebner-Priemer UW & Trull TJ. Ecological momentary assessment of mood disorders and mood dysregulation. *Psychological Assessment* 21, 463 (2009). [PubMed: 19947781]
82. Siegel EH et al. Emotion fingerprints or emotion populations? A meta-analytic investigation of autonomic features of emotion categories. *Psychological bulletin* 144 (4), 343 (2018). [PubMed: 29389177]
83. Gendron M, Roberson D & Barrett LF Cultural Variation in Emotion Perception Is Real: A Response to Sauter, Eisner, Ekman, and Scott (2015). *Psychological Science* 26 (3), 357–359 (2015). 10.1177/0956797614566659. [PubMed: 25608863]
84. Barrett LF, Adolphs R, Marsella S, Martinez AM & Pollak SD Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological science in the public interest* 20 (1), 1–68 (2019). [PubMed: 31313636]
85. Lindquist KA, Wager TD, Kober H, Bliss-Moreau E & Barrett LF The brain basis of emotion: a meta-analytic review. *The Behavioral and brain sciences* 35 (3), 121–143 (2012). 10.1017/S0140525X11000446. [PubMed: 22617651]
86. Paolacci G, Chandler J & Ipeirotis PG Running experiments on Amazon mechanical turk. *Judgment and Decision Making* 5 (5), 411–419 (2010).
87. Faul F, Erdfelder E, Lang A-G & Buchner A. G*Power 3: a flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavioral Research Methods* 39, 175–191 (2007).
88. Ho NSP et al. Facing up to why the wandering mind: Patterns of off-task laboratory thought are associated with stronger neural recruitment of right fusiform cortex while processing facial stimuli. *Neuroimage* 214 (March), 116765 (2020). 10.1016/j.neuroimage.2020.116765. [PubMed: 32213314]
89. Adler NE, Epel ES, Castellazzo G & Ickovics JR Relationship of subjective and objective social status with psychological and physiological functioning: Preliminary data in healthy white women. *Health Psychology* 19 (6), 586–592 (2000). 10.1037/0278-6133.19.6.586. [PubMed: 11129362]
90. Singh-Manoux A, Marmot MG & Adler NE Does subjective social status predict health and change in health status better than objective status? *Psychosomatic Medicine* 67 (6), 855–861 (2005). 10.1097/01.psy.0000188434.52941.a0. [PubMed: 16314589]
91. Radloff LS *The CES-D Scale: A Self-Report Depression Scale for Research in the General Population*. *Applied Psychological Measurement* 1 (3), 385–401 (1977). 10.1177/014662167700100306.

92. Snaith RP et al. A scale for the assessment of hedonic tone. The Snaith-Hamilton Pleasure Scale. *British Journal of Psychiatry* 167 (JULY), 99–103 (1995). 10.1192/bjp.167.L99.
93. Angold A, Costello EJ, Messer SC & Pickles A Development of a short questionnaire for use in epidemiological studies of depression in children and adolescents. *International Journal of Methods in Psychiatric Research* 5, 237–249 (1995).
94. Birmaher B et al. Psychometric properties of the screen for child anxiety related emotional disorders (SCARED): A replication study. *Journal of the American Academy of Child and Adolescent Psychiatry* 38 (10), 1230–1236 (1999). 10.1097/00004583-199910000-00011. [PubMed: 10517055]
95. Jolly E Pymer4: Connecting R and Python for linear mixed modeling. *Journal of Open Source Software* 3 (31), 862 (2018).
96. Snijders TAB & Bosker RJ Modeled variance in two-level models. *Sociological methods & research* 22 (3), 342–363 (1994).
97. Nakagawa S & Schielzeth H A general and simple method for obtaining R² from generalized linear mixed-effects models. *Methods in ecology and evolution* 4 (2), 133–142 (2013).
98. Barton K MuMIn: multi-model inference. <http://r-forge.r-project.org/projects/mumin/> (2009).
99. Paszke A. et al. . Pytorch: an imperative style, high-performance deep learning library. *Advanced Neural Information Processing Systems* 32, 8026–8037 (2019).
100. Kingma DP & Ba J Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014).
101. Rutledge RB Risky decision and happiness task: The Great Brain Experiment smartphone app (2021). 10.5061/dryad.prr4xgxxk.

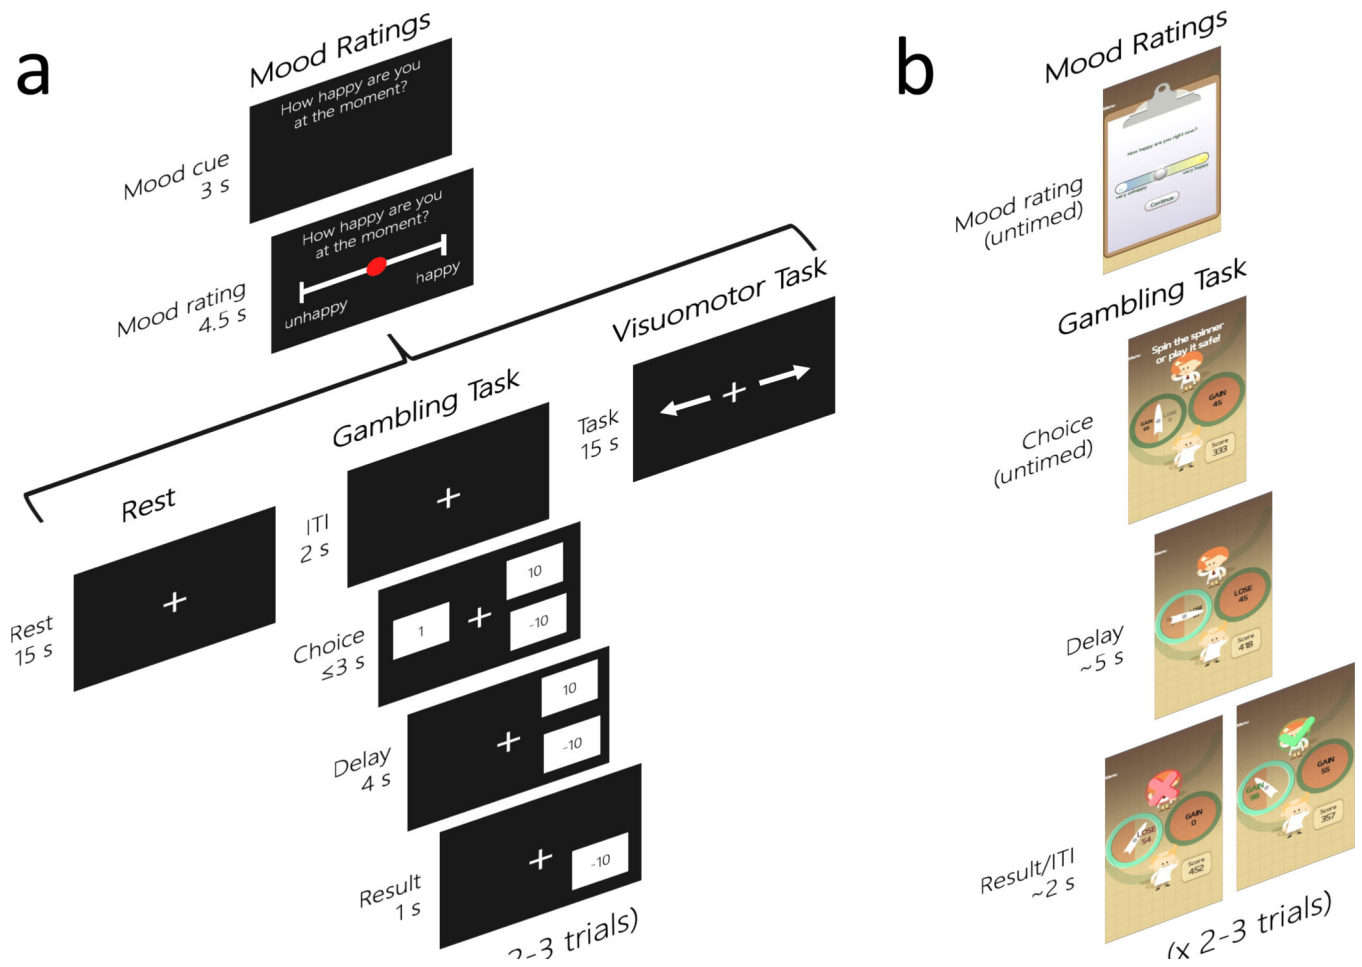


Figure 1: One cycle (mood rating + task) of the administered to (A) online participants and (B) mobile app participants. After completing their first mood rating, participants completed one cycle of the rest, gambling, or visuomotor task, then completed another mood rating, and so on. In the case of the rest and visuomotor tasks, the cycle duration was determined by time. In the case of the gambling task, it was determined by the time taken to complete 2 or 3 (randomised) trials of the gambling task.

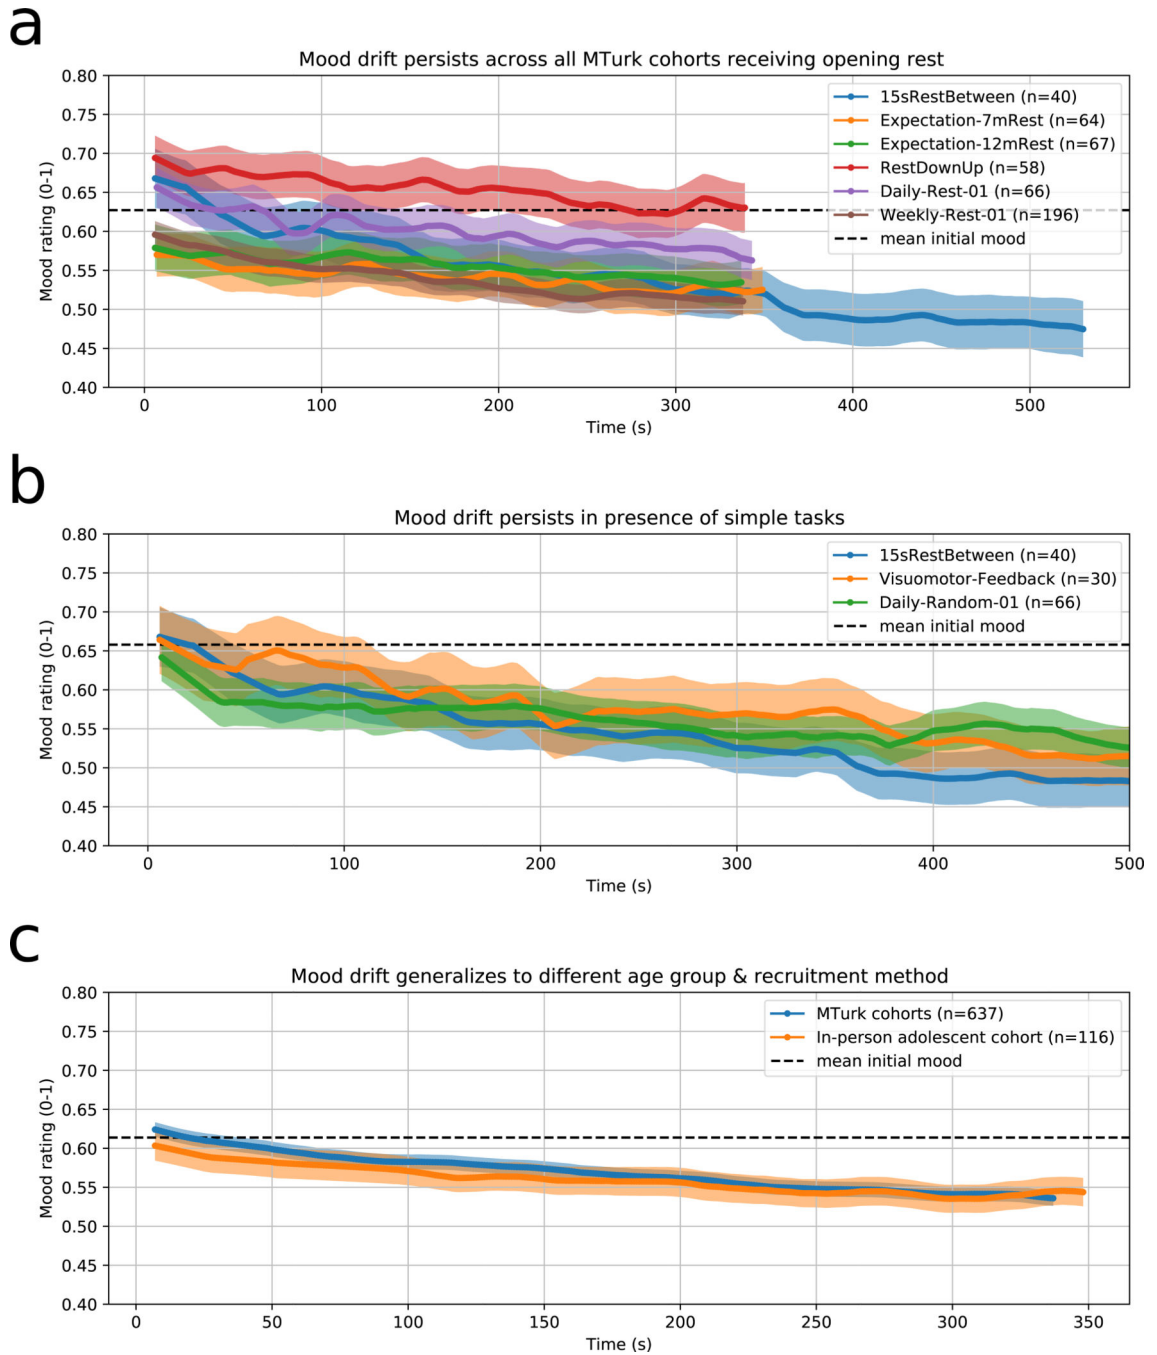


Figure 2: The timecourse of mood drift is consistently present across many cohorts and task modulations. These plots each show the mean timecourse of mood across participants in various online cohorts for the first block of the task. Each participant’s mood between ratings was linearly interpolated before averaging across participants. The shading around each line represents the standard error of the mean. Each name in the legend corresponds to a cohort completing a slightly different task (Extended Data Table 1). Mean initial mood refers to the mean of cohort means, not the mean of subject means. (a) Mean timecourse

of mood ratings during an opening rest period in all Amazon Mechanical Turk (MTurk) cohorts that received it. Mood drift was discovered in one cohort (blue line) and replicated in five independent naive cohorts. (b) Mood drift was observed not only in rest periods (blue), but also in a simple task requiring action and giving feedback (orange), and in a random gambling task with 0-mean reward prediction errors and winnings (green). (c) Mood drift was observed both in adults recruited on MTurk (combining across all MTurk participants that received opening rest or visuomotor task periods) (blue) and in adolescents recruited in person (orange).

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

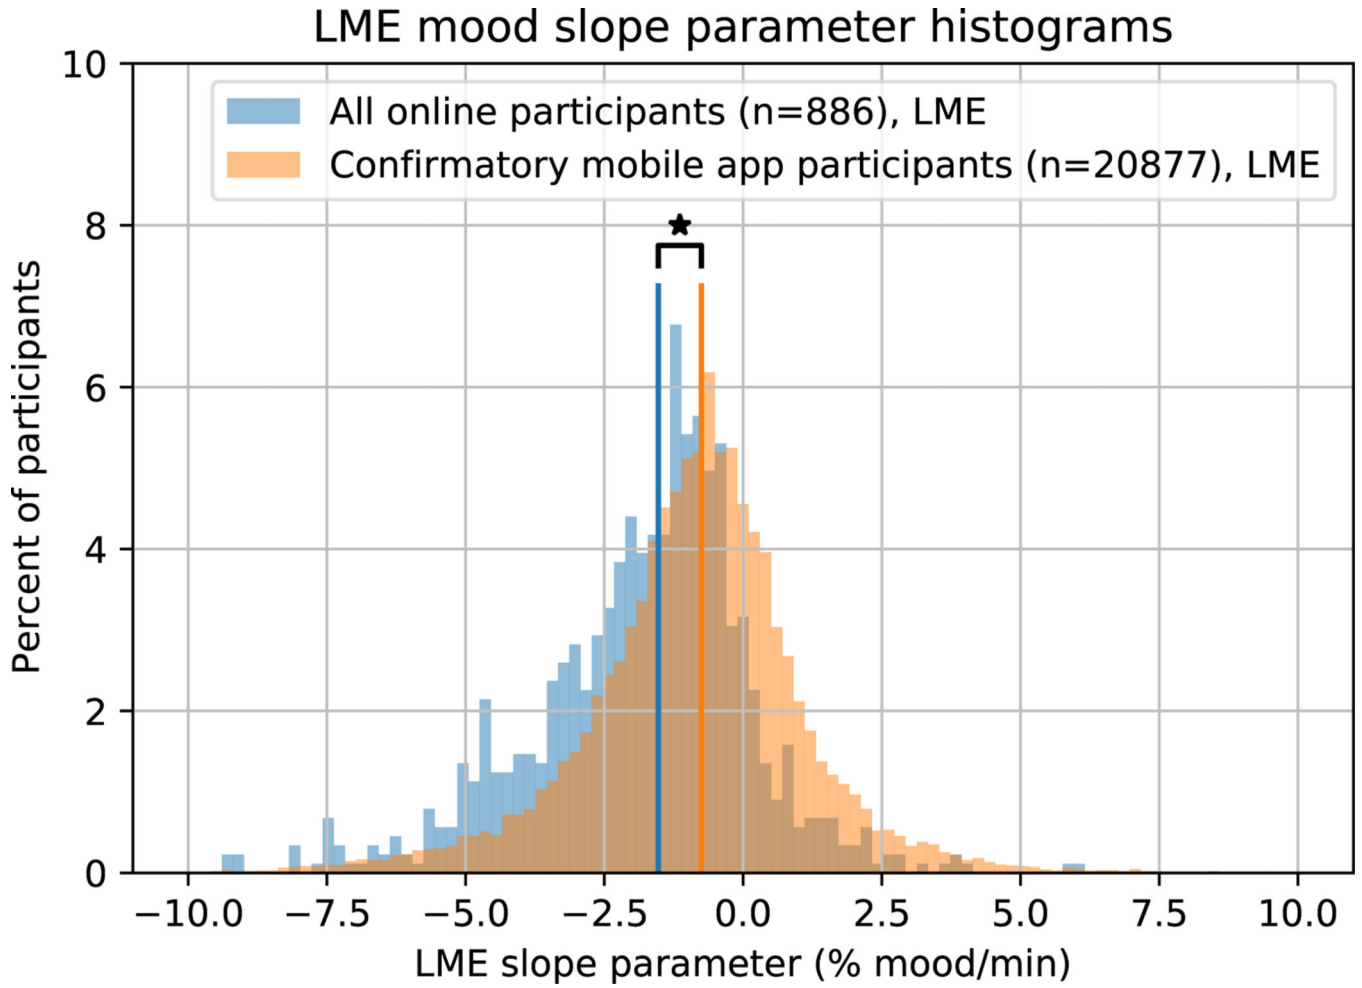


Figure 3: Individual subject LME slope parameters for online participants (blue) and mobile app participants (orange). The online participants had slopes below zero on average ($Mean \pm SE = -1.89 \pm 0.185 \% mood/min, t_{864} = -10.3, p < 0.001$), as did the mobile app participants ($Mean \pm SE = -0.881 \pm 0.0613 \% mood/min, t_{22804} = -14.4, p < 0.001$). Mood drift was significantly less negative in the mobile app participants (median=-0.752, IQR=2.10 % mood/min) than in the online participants (median=-1.53, IQR=2.34 % mood/min, 2-sided Wilcoxon rank-sum test, $W_{21761} = -14.5, p < 0.001$). Vertical lines represent group medians. Stars indicate $p < 0.05$.]

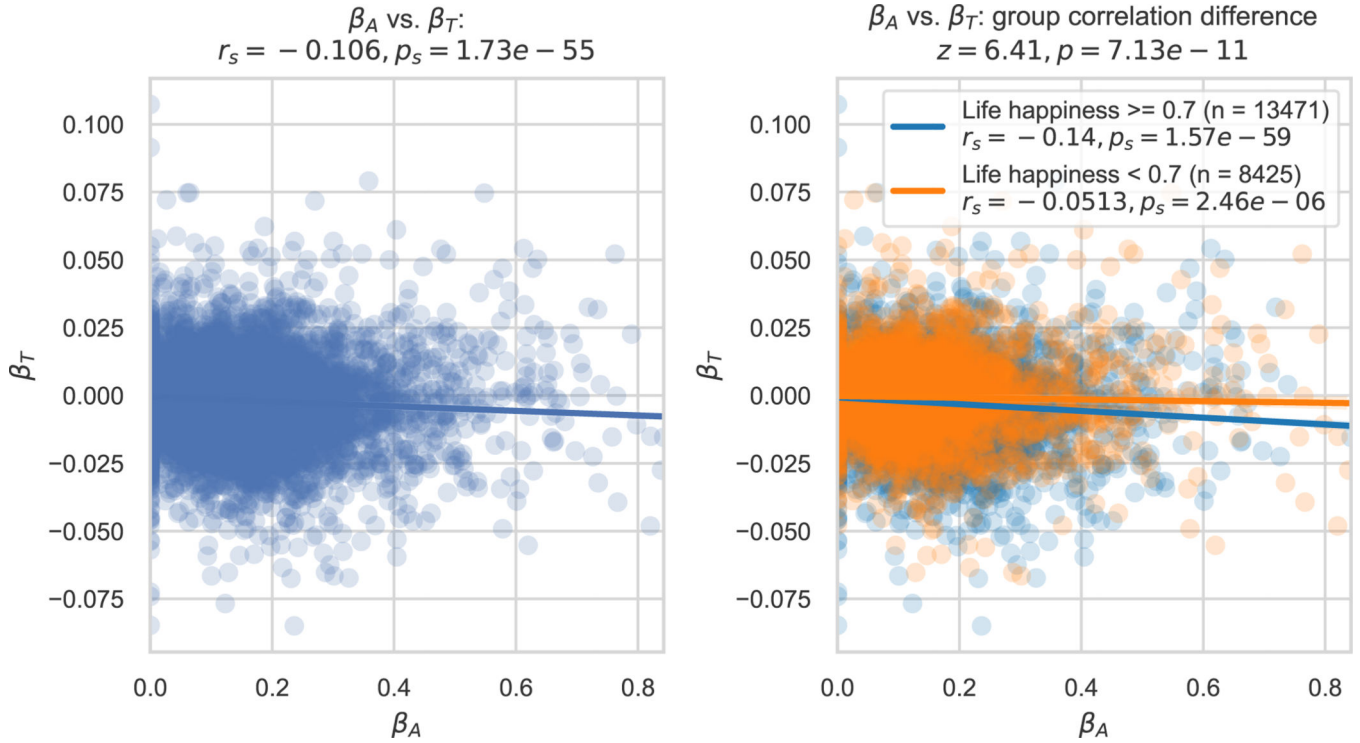


Figure 4: Individual differences in sensitivity to the passage of time relate to other individual differences in the mobile app cohort. The computational model’s time sensitivity parameter β_T for each participant in the mobile app cohort is plotted against that participant’s reward sensitivity parameter β_A (left). r_s and p_s denote Spearman correlation coefficient and corresponding P value. When grouped by life happiness, participants with happiness at or above the median had a stronger $\beta_T - \beta_A$ anticorrelation than participants with happiness below the median (right). Each dot is a participant ($n=21,896$). Each line is a linear best fit, and patches show the 95% confidence interval of this fit. The group difference in Spearman correlations was statistically confirmed using a z statistic. P values shown are 2-sided with no correction for multiple comparisons.

Opening rest period is associated with reduced gambling choices

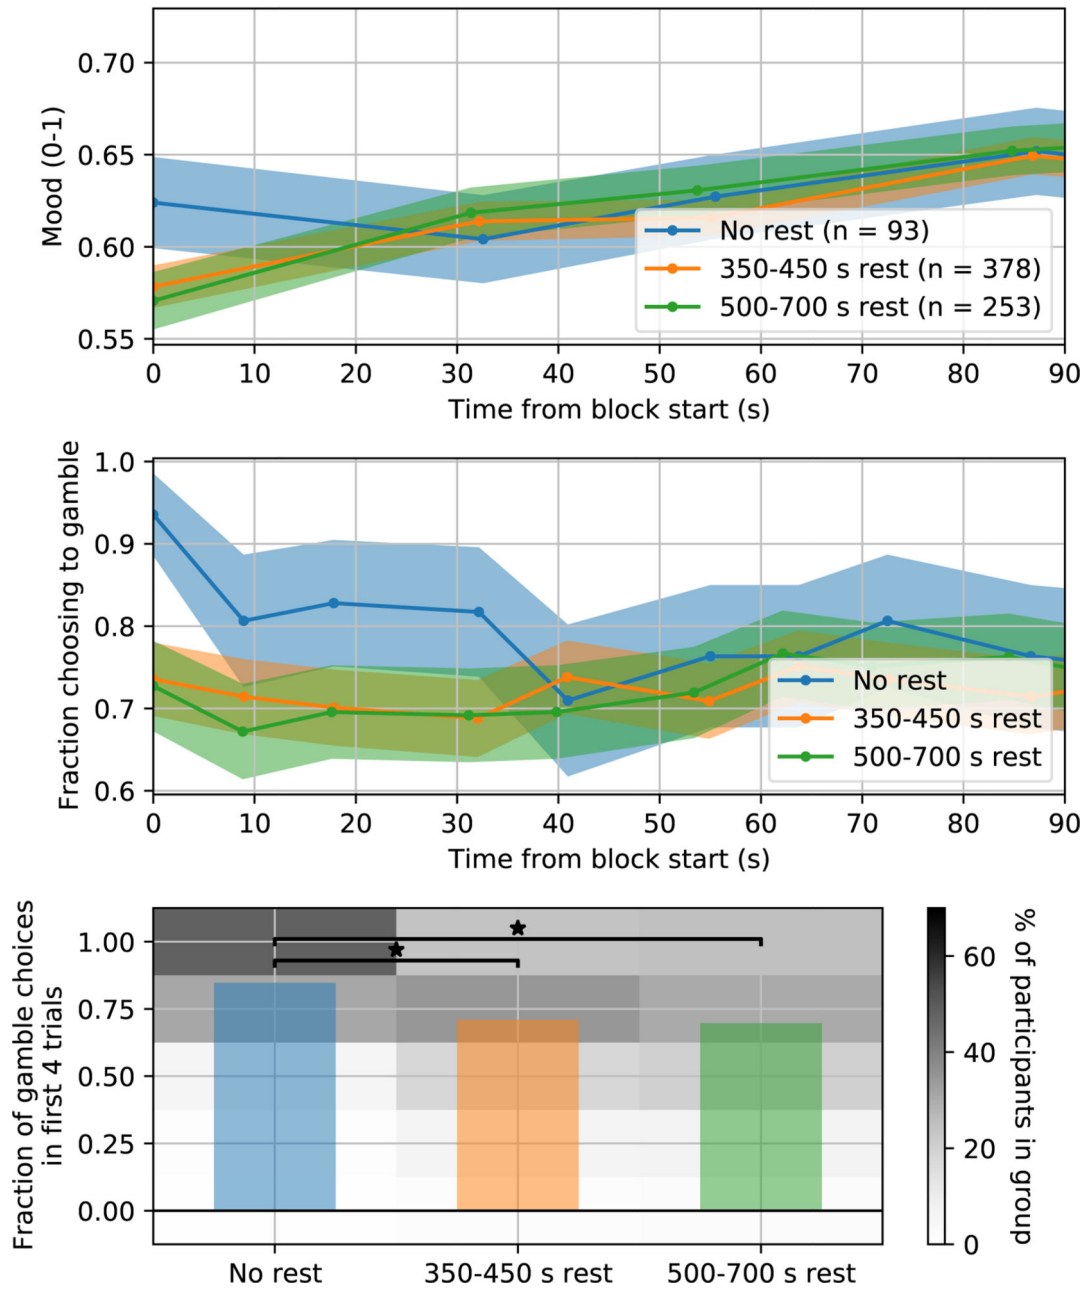


Figure 5: Rest periods decreased the likelihood of choosing to gamble in the first 4 trials after rest ended. Top: mean \pm standard error mood ratings across participants in their first block of (positive closed-loop) gambling preceded by different rest period durations. Middle: fraction of participants in each group that chose to gamble on each trial of this first gambling block (mean \pm 95 percent confidence intervals derived from a binomial distribution). Bottom: bars show mean across participants of the fraction of the first 4 trials of this first gambling block that participants chose to gamble. Histogram shows the distribution of choices (i.e., to

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

gamble on 0, 1, 2, 3, or 4 trials) within each group. Stars indicate that a pair of groups was significantly different (2-sided Wilcoxon rank-sum test, no-rest vs. short-rest: $W_{469} = 4.85$, $p < 0.001$; no-rest vs long-rest: $W_{344} = 4.79$, $p < 0.001$; both $< 0.05/3$ controlling for multiple comparisons). Sample sizes are: No rest group: $n = 93$ participants, 350–450 s rest group: $n = 378$ participants, 500–700 s rest group: $n = 253$ participants).