eLife

Reviewed Preprint
v1 • January 13, 2026
Not revised

For correspondence:
darro015@umn.edu

* These authors contributed equally to this work

† Co-senior authors

Competing interests: No competing interests declared

Funding Statement: None

Reviewing editor: Roshan Cools, Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, Netherlands

Impaired Adaptive Learning in Chronic Pain Contributes to Apathy

Xinyuan Yan[1,2,*], Crina M Peterson[1,*], Lisa M Schmidt[3], Seth Koenig[2], Donald R Nixdorf[4], Alexander Herman[2,†], David P Darrow[1,2,†]

[1]Department of Neurosurgery, University of Minnesota, Minneapolis, United States • [2]Department of Psychiatry, University of Minnesota, Minneapolis, United States • [3]The TMJ Association, Milwaukee, United States • [4]School of Dentistry, University of Minnesota, Minneapolis, United States

eLife Assessment

This study provides a **useful** application of computational modelling to examine how people with chronic pain learn under uncertainty, contributing to efforts to link pain with motivational processes. However, the evidence supporting the main claims is **incomplete**, as the modelling differences are not reflected in observable behaviour or pain measures, and the interpretation extends beyond what the data can substantiate. The conclusions would benefit from a clearer explanation of the behavioural differences that underlie the computational findings.

https://doi.org/10.7554/eLife.109185.1.sa3

Abstract

The cognitive mechanisms linking chronic pain to motivational symptoms remain poorly understood. We demonstrate that individuals with chronic temporomandibular disorder (TMD), a common cause of chronic pain, exhibit a specific deficit in adaptive learning in uncertain environments, characterized by failure to reduce uncertainty over time and maintain efficient learning rates. Using a probabilistic reward task, we pioneered the application of a novel volatile Kalman filter to model behavior in 26 TMD participants and 39 matched controls, uniquely tracking trial-wise updates in uncertainty, volatility, and learning rate. Although surface-level performance did not differ across groups, model-based analysis revealed that those with TMD failed to reduce uncertainty and adapt their learning over time. TMD participants also reported significantly greater apathy, depression, and pain catastrophizing, as well as lower health-related quality of life. Mediation analysis confirmed that impaired uncertainty adaptation partially mediated the relationship specifically between TMD and apathy. These findings identify a computational signature of disrupted uncertainty adaptation in people with TMD and provide evidence for a mechanistic link between chronic pain and motivational dysfunction. This work lays a foundation for future studies examining how belief-updating deficits contribute to broader affective and cognitive symptoms in chronic pain.

Introduction

Chronic pain affects 21% of the population and causes substantial morbidity and disability. Although often treated as peripheral in origin, a growing body of evidence suggests that chronic pain involves central mechanisms linked to emotion, cognition, and motivation (Apkarian et al., 2009; Baliki and Apkarian, 2015). In particular, perceptions of uncertainty may play an important role in the burden experienced by people with chronic pain (Brown et al., 2008; Wiech and Tracey, 2009; Yoshida et al., 2013). Uncertainty about diagnosis, prognosis, and treatment efficacy can amplify pain perception, increase anxiety, and lead to maladaptive coping mechanisms. For many patients, not knowing what causes their symptoms or when pain might intensify creates a state of hypervigilance that can exacerbate muscle tension and pain. This

uncertainty often leads to catastrophizing thoughts, where patients anticipate the worst possible outcomes, creating a self-perpetuating cycle of stress and pain that can leave patients feeling helpless and apathetic (Quartana et al., 2009; Vlaeyen and Linton, 2000). Clinicians working with people who have chronic pain have long known that addressing altered perceptions of these uncertainties may help reduce anxiety and improve outcomes.

One particularly disabling form of chronic pain, temporomandibular disorders (TMD), affects 5% to 10% of the population and is the second most common cause of chronic musculoskeletal pain (National Academies of Sciences, Engineering, and Medicine et al., 2020). Anxiety and depression, conditions associated with low motivation, are found in 20-60% of patients with TMD (De La Torre Canales et al., 2018), similar to the prevalence rates seen in chronic pain more generally (Aaron et al., 2025). Patients with TMD exhibit executive function deficits related to the functional disconnection of the cognitive, motivational, and affective areas of the brain (Weissman-Fogel et al., 2011). Whether these deficits reflect generalized impairment or disruption in specific cognitive processes that represent potential therapeutic targets remains unknown despite the use of cognitive behavioral therapy as a standard treatment (Shivakumar et al., 2025). Previously, we found that perceptions of uncertainty influence motivation levels in a large normative sample (Yan et al., 2025). Here we provide behavioral and computational evidence that individuals with TMD exhibit a core deficit in adaptive learning in response to uncertainty, which contributes to motivational dysfunction, particularly apathy.

Results

We recruited 26 individuals with chronic TMD (age±SD=33.69±14.45) and 40 healthy controls (age±SD = 32.60±13.68). One control participant reported unexpectedly high pain (VAS=68/100) despite no diagnosed pain conditions and was excluded as an outlier, yielding a final control sample of 39 participants (age±SD=32.75±13.81). After exclusion, control participants showed minimal pain levels (VAS=6.44±9.76) compared to TMD participants (VAS=36.58±20.64, $t(63)=7.91$, $p<0.001$). All TMD participants had a muscle pain diagnosis (9 local myalgia, 17 myofascial pain with referral), and 96% had a joint pain diagnosis (7 unilateral, 18 bilateral arthralgia). The distribution of Graded Chronic Pain Scale (GCPS) was I=11 (42%), II= (27%), III=6 (23%), and IV=2 (8%), which is similar to the distribution in U.S. dental practice settings(Velly et al., 2022) and is composed of primarily mild to moderate severity. Participants completed a 300-trial three-armed bandit task (Figure 1, Method) designed to probe learning under uncertainty. Participants with TMD reported higher pain at baseline ($t(63)=7.91$, $p<0.001$) (Figure 1D) and completed self-report assessments of apathy, mood, and pain catastrophizing (details see Method). Task behavior was analyzed using a volatile Kalman filter model (Figure 1C, Method), which estimates latent learning parameters on a trial-by-trial basis.

TMD participants failed to adapt to uncertainty

Although participants with TMD performed similarly (all $p > 0.400$, see Method) to healthy controls in model-free measures (e.g., accuracy, performance, win-stay/lose-shift), model-based analyses revealed persistent deficits in learning (Figure 2). Compared to controls (HC), participants with TMD failed to reduce their uncertainty estimates and adapt their learning rates over time (Figure 2A–C). Regression analysis showed that controls decreased their learning rates over time, with significant negative slopes in uncertainty and volatility estimates (all $p<0.001$), while TMD participants exhibited flat or non-significant trajectories (Figure 2D–F). This pattern suggests TMD participants maintain a chronically heightened sensitivity to environmental changes. First, they fail to decrease their estimate of environmental volatility despite substantial experience. Second, they maintain high uncertainty about reward expectations throughout the experiment. Third, their persistently elevated learning rates indicate that they continually prioritize recent feedback over prior beliefs. This computational signature resembles a trauma-like state where the individual anticipates ongoing volatility and remains hypervigilant to environmental changes.

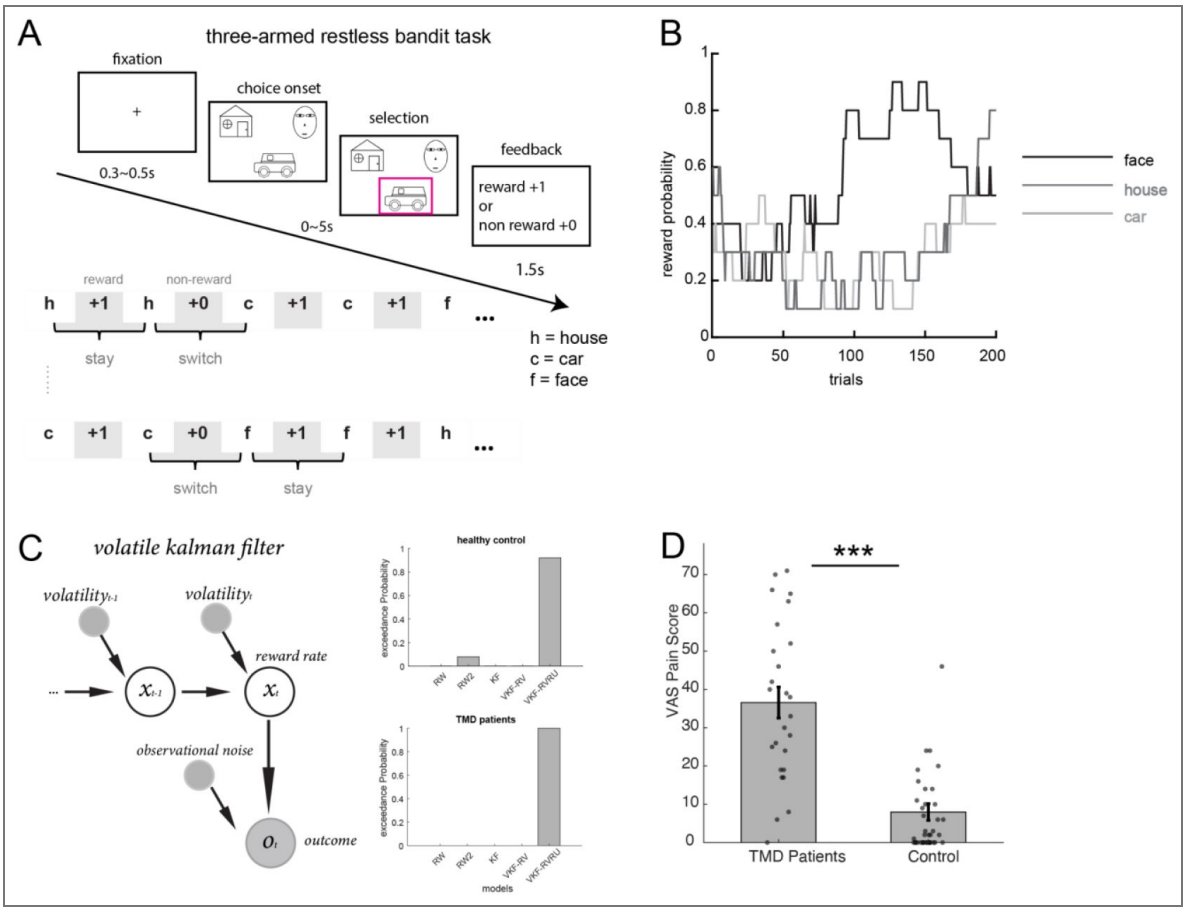

Figure 1. Task, computational models, and pain ratings.

(A) Three-armed restless bandit task (Method). Patients chose one of three options (face, car, house), followed by reward or no-reward feedback. (B) Example reward probability trajectories across trials. (C) Volatile Kalman filter model implementation (left panel) and Bayesian model comparison results (right two panels). The model computes trial-by-trial value and uncertainty estimates for each bandit. The volatile Kalman filter with relative value and relative uncertainty incorporated into the softmax function outperformed other models. Model abbreviations: RW = Rescorla-Wagner; Dual RW = Rescorla-Wagner with positive and negative learning rates; KF = Kalman filter; VKF-RV+RU = volatile Kalman filter with relative value and uncertainty in softmax function; VKF-RV = volatile Kalman filter with relative value in softmax function (model validation analyses see Method). Relative value (RV) and relative uncertainty (RU) were calculated as the ratio of the chosen option's value/uncertainty to the sum across all options (see Method). (D) TMD patients show significantly higher pain ratings on the visual analog scale (VAS) than controls. ***p<0.001.

Such a persistent state of uncertainty about the environment may undermine the ability to form stable expectations, contributing to learned helplessness and apathy observed in chronic pain conditions (Samwel et al., 2006).

To ensure the robustness of these findings, we also tested exponential decay models (see Methods & Table S2 for details). This analysis revealed that TMD participants not only failed to show expected exponential decay but often exhibited paradoxical patterns. For volatility, 82% of controls showed adaptive exponential decay versus only 42% of TMD participants (Fisher's exact p = 0.001), with 39% of TMD participants showing exponential growth. The decay rate parameter for uncertainty differed significantly between groups (HC: b = -0.00094 vs TMD: b = +0.00012, two-sample t test, p = 0.002), with TMD's positive value indicating increasing uncertainty over time. For learning rate, 69% of controls versus 35% of TMD participants showed adaptive decay (Fisher's exact p = 0.01). These converging analyses, whether using linear slopes or exponential parameters, demonstrate that TMD participants exhibit qualitatively aberrant learning trajectories incompatible with adaptive uncertainty processing.

Together, this computational profile, where uncertainty remained unchanged rather than resolved with experience, aligns with predictive-processing (Bayesian) accounts of pain that emphasize aberrant precision (uncertainty) weighting and misestimation of environmental volatility, and with evidence that people learn the statistics of pain in Bayesian-like ways (Onysk et al., 2024). It is also consistent with patients' lived reports that chronic pain is increasingly unpredictable, which fosters ongoing uncertainty in day-to-day functioning (Little et al., 2025).

Uncertainty learning dynamics predict psychological outcomes in TMD participants

To assess the clinical relevance of these findings, we examined the psychiatric symptom burden across groups. TMD participants showed significantly higher apathy (Apathy Motivation Index, $t(63)=-3.63$, $p<0.001$) (Ang et al., 2017), depression (PHQ-9, $t(63)=-6,63$, $p<0.001$), pain catastrophizing scores ($t(63)=-4.15$, $p<0.001$), and lower quality of life (EQ-5D; $t(63)=-3.70$, $p<0.001$, Figure 3A–D). Mediation analysis demonstrated that reduced uncertainty adaptation significantly mediated the relationship between TMD status and apathy indirect effect (ab), effect = 0.175, 95%CI = [0.006, 0.484], 5000 sample bootstrap) (Figure 3E), suggesting that disrupted uncertainty learning may underlie motivational symptoms in chronic TMD pain. We did not find any other significant mediation results for PHQ-9 (95%CI = [-0.153, 0.283]), Pain Catastrophizing score (95%CI = [-0.123, 0.395]), or EQ-5D (95%CI = [-0.003, 0.447]).

The selective mediation of apathy, but not depression or pain catastrophizing, by uncertainty adaptation deficits reveals a specific computational pathway linking chronic pain to motivational dysfunction. This persistent uncertainty state mirrors the clinical phenomenology of chronic pain, where patients cannot predict when symptoms will worsen, which may lead to learned helplessness and withdrawal.

Discussion

Our findings identify a computational mechanism that links chronic pain to apathy: impaired adaptation of uncertainty over time. Healthy controls reduced volatility and uncertainty with experience, showing robust exponential decay, whereas participants with temporomandibular disorder (TMD) often failed to reduce uncertainty, and nearly 40% showed paradoxical increases in volatility or uncertainty. This qualitative breakdown suggests that the pain state disrupts belief updating itself, not merely the pace of learning. Such results align with predictive-processing accounts of pain, which propose that maladaptive precision-weighting of sensory and contextual signals sustains hypervigilance and unpredictability(Apkarian et al., 2009; Wiech and Tracey, 2009). Clinically, patients often describe their pain as arising "out of nowhere," and our findings provide a computational basis for this lived experience.

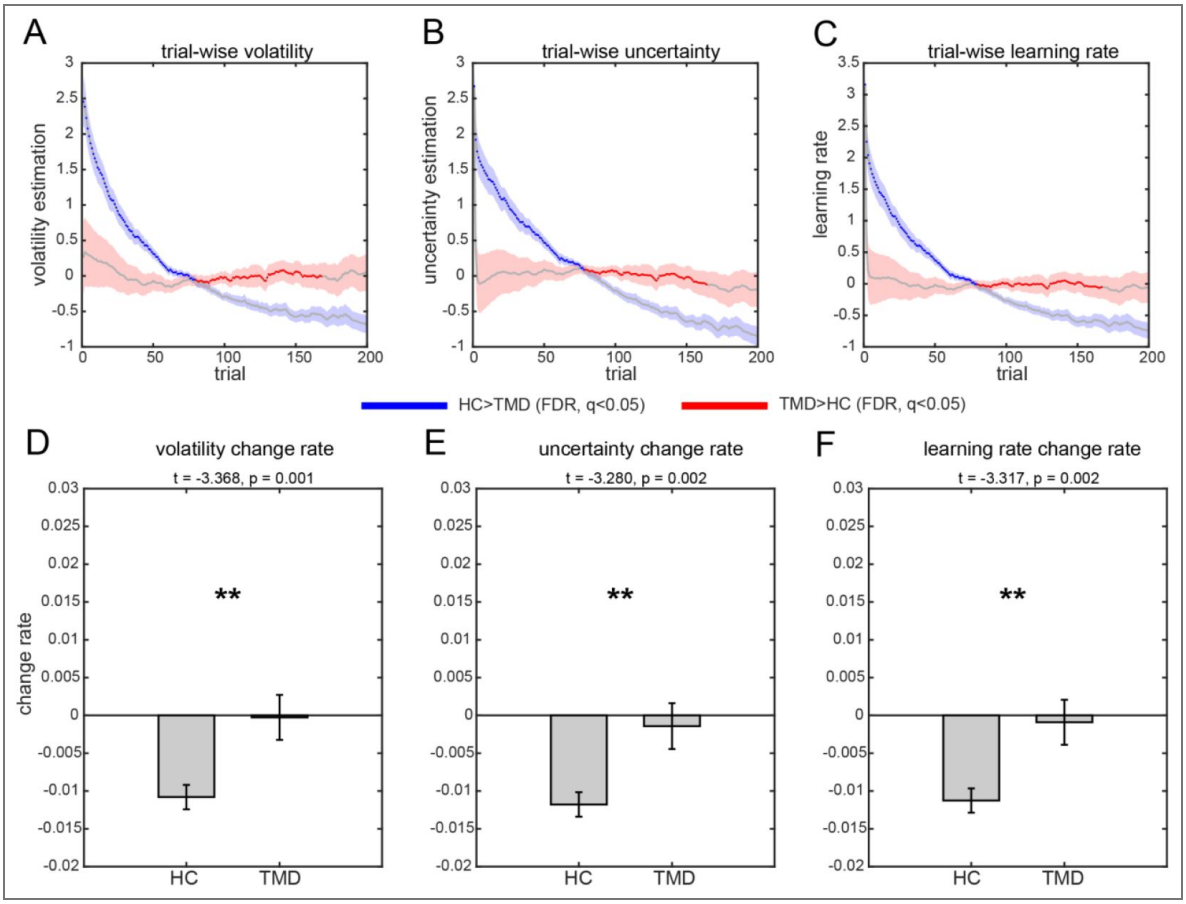

Figure 2. Temporal dynamics of model parameters.

(A–C) Trial-by-trial estimates of volatility, uncertainty, and learning rate in TMD and control groups. (D–F) Linear regression slopes for each parameter over time. Controls showed adaptive decreases in parameter values (all p<0.001); TMD participants did not (all p>0.05). Controls also showed significantly decreasing trends than TMD patients as well (all p < 0.01). **p<0.01, ***p<0.001. n.s, non-significant.

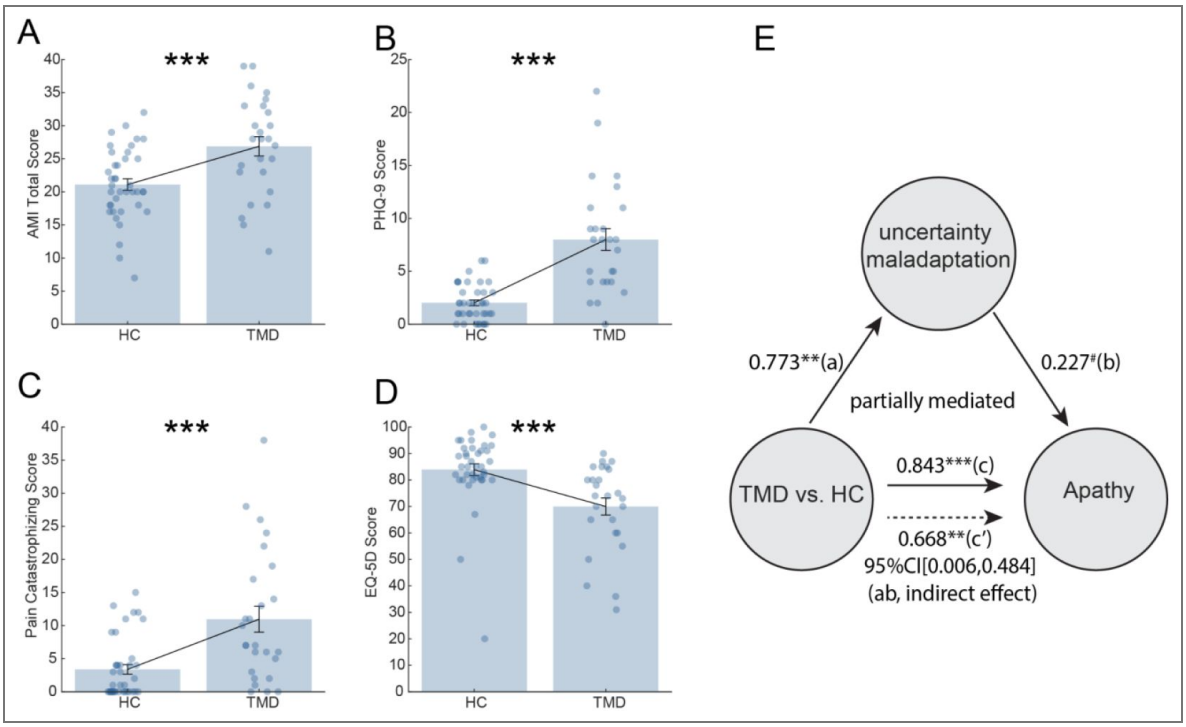

Figure 3. Psychiatric symptom burden and mediation analysis.

(A–D) Group differences in apathy, depression, pain catastrophizing, and EQ-5D. (E) The mediation model shows that impaired uncertainty adaptation mediates the relationship between TMD status and apathy (the indirect effect was significant. p<0.06 (marginally significant), *p < 0.05, **p<0.01, ***p < 0.001.

Mediation analysis further underscored the specificity of this mechanism. Impaired uncertainty adaptation predicted apathy, but not depression, catastrophizing, or quality of life. Apathy is frequently conflated with depression, yet evidence points to distinct neural substrates(Ang et al., 2017; Husain and Roiser, 2018). Our results extend this dissociation into chronic pain, highlighting apathy as the behavioral outcome most directly tied to uncertainty misestimation and suggesting it may serve as a mechanistic bridge between disrupted learning and functional disability.

At the same time, surface performance was preserved. TMD participants tracked rewards above chance, exploited optimal options, and showed normal win–stay/lose–shift behavior. Their impairment emerged only at the computational level, where latent learning parameters failed to adapt. This dissociation explains how patients can appear cognitively intact while struggling with motivation. Together, these findings reframe motivational symptoms in chronic pain as arising from domain-general failures of belief updating rather than nociception or mood symptoms alone. They also suggest a testable therapeutic direction: interventions that recalibrate uncertainty, through cognitive-behavioral strategies or neuromodulation of prefrontal–striatal circuits, may help restore adaptive learning and motivation in patients with chronic pain.

Materials and Methods

Three-armed restless bandit task

In this three-armed restless bandit task (Ebitz et al., 2018; Yan et al., 2025), participants completed 200 trials where they freely selected between three targets on each trial to earn potential rewards of 1 point. Each target is associated with a hidden reward probability that randomly and independently changes throughout the task. We seeded each participant's reward probability randomly to prevent biases due to particular kinds of environments. Specifically, on each correct trial, there was a 10% chance that the reward probability for each target would either increase or decrease by 1.0, with these probabilities bounded between 0.1 and 0.9. Due to the variable and independent nature of the rewards, participants could only estimate the probabilities by actively sampling from the targets and accumulating their reward experiences over time. Trials were excluded if participants failed to respond within 5 seconds. The task was programmed in MATLAB using Psychtoolbox 3 (*http://psychtoolbox.org/*). Stimuli were presented on an HP laptop positioned approximately 70 cm from the participant's face and centered in their field of view while they sat upright in their hospital bed.

Self-reported measures

Participants completed a comprehensive battery of validated self-report measures. The Graded Chronic Pain Scale (GCPS) (Korff et al., 1992) assessed pain intensity and pain-related disability. The Patient Health Questionnaire-9 (PHQ-9) (Kroenke et al., 2001) was administered to evaluate the presence and severity of depressive symptoms. The EuroQol Five Dimension (EQ-5D)(Brooks et al., 2020) questionnaire assessed health-related quality of life, with a higher score indicating better health status. Motivational deficits were evaluated using the Apathy Motivation Index (AMI) (Ang et al., 2017).

Model-free analyses

We adopted some widely used model-free measures, including win-stay and lose-switch (den Ouden et al., 2013), as the direct measurement for this learning task.

Win-stay. Win-stay is the percentage of times that the choice in trial t-1 was repeated on trial t following a reward.

Lose-switch. In contrast, the lose-switch equals the percentage of trials in which the choice was shifted or changed when the outcome of trial t-1 was non-reward.

Initial analysis of model-free behavioral metrics, including stay rate, switch rate, win-stay/lose-switch behaviors, overall performance (corrected for chance level), and response times, revealed no significant differences between TMD patients and controls (all $p > 0.400$; see Figure S1). Given the non-stationary nature of our restless bandit task, we assessed learning through maintenance of above-chance performance and adaptive strategy (see below) use rather than traditional improvement curves, as reward probabilities changed continuously throughout the experiment.

To comprehensively characterize task behavior, we computed seven model-free metrics examining reward sensitivity, choice consistency, and performance optimization (see below).

We first examined choice autocorrelation to assess systematic decision patterns independent of reward feedback. The lag-1 autocorrelation coefficient revealed positive values for both groups (HC: 0.270 ± 0.051; TMD: 0.169 ± 0.064; $t(63) = 1.24$, $p = 0.221$), calculated as $r = \Sigma[(choice_t - mean)\times(choice_{t+1} - mean)] / \Sigma(choice - mean)^2$. These positive autocorrelations indicate that both groups exhibited systematic choice patterns rather than random responding, with TMD participants showing a non-significant trend toward lower perseveration that did not reach the statistical threshold.

To assess reward-driven learning, we evaluated how well participants tracked and exploited the objectively best option. Both groups performed well above the chance level of 33.3%, with TMD participants (52.1 ± 3.4%) tracking the best option as effectively as controls (50.4 ± 1.9%; $t(63) = -0.49$, $p = 0.624$). Furthermore, we quantified reward-seeking behavior by correlating each participant's choice probability with each option's average reward rate. Both groups showed strong positive correlations (HC: 0.522 ± 0.101; TMD: 0.643 ± 0.089; $t(63) = -0.84$, $p = 0.405$), with TMD participants displaying numerically higher reward sensitivity, directly contradicting any hypothesis of disengagement or reward insensitivity.

To assess learning success, we calculated individualized chance levels for each participant based on the mean reward probabilities of their specific bandit instantiation. Binomial tests revealed that 66.7% of controls and 61.5% of TMD participants performed above their personalized baseline, with group-level analyses confirming both groups performed significantly above chance (both $p < 0.0001$). This comparable achievement rate demonstrates that TMD participants successfully learned the task structure and exploited reward opportunities as effectively as controls.

Across all behavioral metrics spanning reward sensitivity, choice consistency, and performance optimization, TMD participants showed normal task behavior with no significant group differences (all $p > 0.14$). This behavioral equivalence, including reward tracking, systematic choice patterns, and above-chance performance, makes our computational finding particularly specific and interpretable. Despite demonstrating intact trial-by-trial learning and reward processing at the behavioral level, TMD participants uniquely fail to reduce their uncertainty estimates over time, revealing a selective deficit in uncertainty resolution rather than a general learning impairment. This dissociation between preserved behavioral performance and impaired uncertainty updating highlights a subtle but critical difference in how TMD patients process and integrate information under uncertainty.

Temporal dynamics of model-free behavioral metrics

To examine learning dynamics across the task, we divided each participant's 200 trials into consecutive 10-trial bins (20 bins total). Within each bin, we calculated win-stay rate (proportion of rewarded trials followed by staying with the same choice) and lose-shift rate (proportion of non-rewarded trials followed by switching to a different choice). These metrics were averaged across participants within each group to generate temporal trajectories. Group differences at each time point were assessed using independent samples t-tests. To examine learning phases, we aggregated data into early (trials 1-50), middle (trials 51-100), and late (trials 151-200) learning periods.

Analysis of behavioral dynamics across 10-trial bins revealed remarkably stable performance in both groups throughout the task (Figure S2). Win-stay rates remained consistently high across all learning phases (early: Controls = 64.0 ± 6.1%, TMD = 67.1 ± 4.2%; late: Controls = 72.6 ± 1.3%, TMD = 71.5 ± 3.8%), with no significant group differences at any time point (0/20 bins, all $p > 0.05$). Similarly, lose-shift rates showed no group differences across bins (0/20 bins, all $p > 0.05$), though both groups showed increasing lose-shift rates from early (Controls = 67.8 ± 7.6%, TMD = 71.7 ± 7.0%) to middle learning (Controls = 75.7 ± 5.2%, TMD = 65.3 ± 3.9%).

Computational modeling

Of note, our approach here is essentially the same as that taken by Piray and Daw ((Chakroun et al., 2020; Piray and Daw, 2020). Here, we briefly described the model details as follows.

Kalman filter

The Kalman filter (KF) model has been widely applied in psychology and neuroscience to study various aspects of learning and decision-making (Cheng et al., 2022; Dayan et al., 2000).

In the Kalman filter model for a multi-armed bandit task, *process noise*, and *observation noise* refer to two distinct sources of uncertainty that affect the learning and decision-making process. Process noise represents the uncertainty in the evolution of the hidden state (reward mean) over time. It accounts for how the true state evolves from one point in time to the next. In mathematical terms, process noise is part of the state transition equation in the Kalman Filter:

$$x_t = x_{t-1} + e_t$$

x_t is the state at time t

e_t is the process noise t, which is assumed to be drawn from a normal distribution with zeros mean and **process noise variance** v. Where the $e_t \sim N(0, v)$.

The process noise captures the idea that the reward means for each arm can change from one trial to the next, even in the absence of any observations. A higher process noise variance v indicates a more volatile environment, where the reward means are expected to change more rapidly.

In contrast, **observation noise** represents the uncertainty in the observed rewards, given the current hidden state (reward mean). Which is assumed to be Gaussian with zero mean and a fixed variance σ^2.

The observation noise captures the idea that the observed rewards are noisy and can deviate from the true reward mean due to random fluctuations or measurement errors.

A higher measurement noise variance indicates a more stochastic environment, where the observed rewards are less reliable and informative about the underlying reward means.

The Kalman Filter operates optimally when the statistical properties of the process noise and the measurement noise are accurately known.

When observation noise variance (σ^2) is high relative to the process noise variance (v), the Kalman gain will be small, and the model will rely more on its prior beliefs and less on noisy observations. Conversely, when the observation noise variance (v), is high relative to the process noise variance (σ^2), the Kalman gain will be large, and the model will update its beliefs more strongly based on the observed rewards.

Extended Kalman filter for the three-armed bandit task

The Kalman filter model can be extended to capture the effects of both volatility and stochasticity in a multi-armed bandit task (Chakroun et al., 2020; Piray and Daw, 2020). In the current study, process noise variance (v) and observation noise variance (σ^2) represent volatility and stochasticity, respectively.

A traditional assumption of the Kalman filter is that the process noise variance, v, as well as the observation noise variance, $\sigma^2$are constant.

Reward means update:

$$m_t = m_{t-1} + k_t(O_t - m_{t-1})$$

Where m_t is the estimated mean or value of the chosen arm at time t and O_t is the observed reward at time t.

The mean update is driven by the prediction error, which is the difference between the observed reward and the previous estimate.

Kalman gain is defined as:

$$k_t = (w_{t-1} + v_{t-1})/(w_{t-1} + v_{t-1} + \sigma^2)$$

Here, k_t represents the Kalman gain or learning rate, which adjusts the weight given to new information based on the relative uncertainty of the prior estimate (w_{t-1}) and the total noise (v + σ^2). When the stochasticity (σ^2) is high relative to the volatility (v), the Kalman gain (learning rate) will be small, and the model will rely more on its prior beliefs and less on the observations. Conversely, when the volatility (v), is high relative to the stochasticity (σ^2), the Kalman gain (learning rate) will be large, and the model will update its beliefs more strongly based on the observed rewards.

Variance update equation:

$$w_t = (1 - k_t)(w_{t-1} + v)$$

This equation updates the posterior variance (w_t), which represents the estimate's uncertainty after observing O_t.

Volatile kalman filter for three-armed bandit task

The key difference between a standard Kalman filter and a volatile Kalman filter (VKF) is the variance of the process noise, a stochastic variable that changes with time. In other words, the VKF introduces parameters to handle the volatility in the process noise. Specifically, it allows the process noise variance v to vary with the observed prediction errors, reflecting changes in environmental volatility.

Kalman gain:

$$k_t = (w_t + v_{t-1})/(w_{t-1} + v_{t-1} + \sigma^2)$$

where W is a noise parameter specific to binary observations. σ^2 capture the observation noise. Update for the reward means:

$$m_t = m_{t-1} + k_t(O_t - m_{t-1})$$

Update for posterior variance w_t:

$$w_t = (1 - k_t)(w_{t-1} + v_{t-1})$$
$$w_{t-1,t} = (1 - k_t)w_{t-1}$$

Update for volatility:

$$v_t = v_{t-1} + \lambda((m_t - m_{t-1})^2 + w_{t-1} + w_t - 2w_{t-1,t} - v_{t-1})$$

The model has three free parameters:

Where the λ represents the volatility learning rate, which is constrained to the unit range [0,1], determines how quickly volatility estimates update

Soft-max choice-probability function for (volatile) kalman filter models

Decisions were modeled using a soft-max choice-probability function in which the probability of selecting a particular bandit i depends on its utility.

In the first soft-max function, we only included decision weights β_V for expected-value. The probability P of selecting bandit i in trial t was modeled as: $P_i(t) = \frac{e^{\beta_V * V_i(t)}}{\sum_j e^{\beta_V * V_j(t)}}$

In the second response model, we only included decision weights β_U for outcome uncertainty to estimate the probability of choosing bandit *i*. So the probability P of selecting bandit *i* in trial *t* was modeled as:

$$P_i(t) = \frac{e^{\beta_U * U_i(t)}}{\sum_j e^{\beta_U * U_j(t)}}$$

We included the third response model's decision weights for both expected value and outcome-uncertainty. So the probability P of selecting bandit *i* in trial *t* was modeled as:

$$P_i(t) = \frac{e^{\beta_V * V_i(t) + \beta_U * U_i(t)}}{\sum_j e^{\beta_V * V_j(t) + \beta_U * U_j(t)}}$$

The definition for relative value and relative uncertainty

For each trial t, we computed the relative value (RV) and relative uncertainty (RU) of the chosen option compared to the unchosen options. Relative value was defined as the ratio of the chosen option's estimated mean reward probability to the sum of all options' value:

$$RV(t) = \frac{m_{t,chosen}}{\sum m_{t,all\ options}}$$

where $m_{t,chosen}$ is the estimated mean of the chosen option and $\sum m_{t,all\ options}$ represents the sum of estimated means across all options.

Similarly, relative uncertainty was calculated as the ratio of the chosen option's posterior variance to the sum of all posterior variances:

$$RU(t) = \frac{w_{t,chosen}}{\sum w_{t,all\ options}}$$

where $w_{t,chosen}$ is the posterior variance of the chosen option and $\sum w_{t,all\ options}$ is the sum of posterior variances across all options.

Alternative models: Rescorla-Wagner models

We also fitted the data to the classical Rescorla-Wagner model. Successful adaptation in a dynamic situation requires the appropriate feedback-based learning process where individuals integrate the feedback (reward or non-reward) into the stimulus-outcome association (Forstmann et al., 2016). The basic reinforcement learning model, the Rescorla-Wagner model can address this process well. So the first model (RW1) was defined as:

$$v_t = v_{t-1} + a \times (R_{t-1} - v_{t-1})$$

where v_t is the value of the option on trial *t*.

a represents the general learning rate from feedback.

To verify whether participants employed distinct or shared computational responses to positive and negative feedback, we built another model with two learning rates, one for positive feedback and the other for negative feedback (den Ouden et al., 2013). This model (RW2) can be defined as:

$$v_t = v_{t-1} + \alpha^{pos} \times (R_{t-1} - v_{t-1}), positive\ feedback$$
$$v_t = v_{t-1} + \alpha^{neg} \times (R_{t-1} - v_{t-1}), negative\ feedback$$

Where v_t is the value of the option on trial *t*. α^{pos} and α^{neg} represents the learning rates from positive and negative feedback, respectively.

For these two models, $R_{t-1} \in \{0,1\}$ represent the feedback received in response to participants' choice on trial *t*-1. And $R_{t-1} - v_{t-1}$ represents prediction error in trial *t*-1.

Soft-max choice-probability function for Rescorla-Wagner models

We used a softmax choice function to map the value into choice. The softmax function for these four models can be defined as:

$$P_i(t) = \frac{e^{\beta_V * V_i(t)}}{\sum_j e^{\beta_V * V_j(t)}}$$

Where the β_V represents the inverse temperature with choice value.

Model fitting and comparison

Hierarchical Bayesian inference (HBI) is a powerful method for model fitting and comparison in group studies (Piray et al., 2019). Unlike traditional approaches such as maximum likelihood estimation (MLE) or maximum a posteriori (MAP) estimation, which fit models to each subject independently, HBI simultaneously fits models to all subjects while constraining individual fits based on group-level statistics (i.e., empirical priors). This approach yields more robust and reliable parameter estimates, particularly when individual subject data is noisy or limited. In our study, we employed HBI to fit models to choice data. The method quantifies group-level mean parameters and their corresponding hierarchical errors. To ensure that parameter estimates remain within appropriate bounds during the fitting process, we used the sigmoid function to transform parameters bounded in the unit range or with an upper bound and the exponential function to transform parameters bounded to positive values. The initial parameters of all models were obtained using a MAP procedure, with the initial prior mean and variance for all parameters set to 0 and 6.25, respectively, based on previous research (Piray and Daw, 2020). This initial variance allows parameters to vary widely without substantial influence from the prior.

We used Bayesian model selection (26) for model comparison, specifically employing the exceedance probability (XP) to select the winning model. The XP quantifies the probability that a given model is more frequent in the population than all other models under consideration while accounting for the possibility that the observed differences in model evidence may be due to chance (see Table S1).

Model validation

To validate our model fitting procedure, we conducted a comprehensive parameter recovery analysis that strictly follows the standard method described in the original methodology paper of VKF (Piray and Daw, 2020) (page 21/26). We generated synthetic data from 30 artificial subjects using the binary VKF combined with a softmax choice model. We employed the same Hierarchical Bayesian Inference (HBI) approach used in our empirical data analysis for parameter estimation. We repeated this procedure 20 times with different random seeds to ensure robust estimation and minimize the effects of random variation.

The recovery analysis revealed strong correlations between the true and recovered parameters, with median correlation coefficients across the 20 simulations: r_λ = 0.852, r_{v0} = 0.473, r_σ^2 = 0.919, $r_{\beta1}$ = 0.918, $r_{\beta2}$ = 0.782 (Figure S3).

Control analyses for initial volatility parameter(v_o)

While our main analyses treated the initial volatility parameter (v_o) as a free parameter, we observed a relatively lower recovery rate compared to the other parameters. Here we demonstrate that this lower recovery rate does not compromise the model's reliability or our main conclusions. We validate our parameter by systematically analyzing parameter recovery across different fixed values (=[1, 3, 5, 7, 9]). Using the same set of simulated data, we found that the recovery rates for the key parameters (λ, σ^2, β1, β2) remained stable regardless of the fixed value. Specifically, the correlations between true and recovered parameters maintained consistent levels (λ= 0.852, σ^2= 0.949, β1= 0.937, β2= 0.831) across all values, suggesting that it does not substantially interact with the recovery of other parameters.

Exponential model analysis

In addition to linear regression analysis, we fitted exponential models to characterize the temporal dynamics of learning parameters.

$$f(x) = ae^{b*x}$$

Where:

f(x) = parameter value (volatility, uncertainty, or learning rate) at trial x; a = initial amplitude coefficient (parameter value at trial 0); e = Euler's number (≈ 2.71828); b = growth/decay rate parameter (b < 0: exponential decay (adaptive learning); b = 0: no change (flat trajectory); b > 0: exponential growth (maladaptive trajectory))

x = trial number (1 to 200)

The time constant τ = -1/b (for b < 0) represents the number of trials required for the parameter to decay to 1/e (≈37%) of its initial value. Exponential models were fitted using MATLAB's Curve Fitting Toolbox (fit function with "exp1" model type).

See Table S2 for results.

Control analyses

No significant correlations were detected between GCPS severity scores and uncertainty adaptation indices in TMD patients (p=0.765). Further correlation analyses revealed no significant associations between GCPS and several psychological and pain-related measures, including Apathy-Motivation Index (AMI) scores(p=0.552), Patient Health Questionnaire-9 (PHQ-9) depression severity(p=0.161), Pain Catastrophizing Scale total scores(p=0.385), and EuroQol-5D (EQ-5D) health utility indices(p=0.480). These findings suggest that, within this TMD patient sample, pain severity and disability, as measured by GCPS, operate independently from uncertainty processing mechanisms and related psychological factors. Correlation analyses were performed using Pearson's correlation coefficient (Figure S4).

Data availability

Anonymized behavioral, model-derived data, and analysis code are available upon request.

Acknowledgements

This work was supported by NIMH under award #R21MH127607, NIDA under award #K23DA050909, and the University of Minnesota's MnDRIVE (Minnesota's Discovery, Research and Innovation Economy) initiative. We thank the participants, staff, and providers at the University of Minnesota TMD, Orofacial Pain and Dental Sleep Medicine Clinic for their contributions.

Additional files

Supplementary Information

Additional information

Author ORCID iDs

Xinyuan Yan: https://orcid.org/0000-0003-3373-330X
Alexander Herman: https://orcid.org/0000-0001-6229-433X
David P Darrow: https://orcid.org/0000-0001-9335-0584

References

Aaron RV, Ravyts SG, Carnahan ND, Bhattiprolu K, Harte N, McCaulley CC, Vitalicia L, Rogers AB, Wegener ST, Dudeney J. (2025) Prevalence of depression and anxiety among adults with chronic pain: A systematic review and meta-analysis: A systematic review and meta-analysis. *JAMA Netw Open* **8**:e250268

Ang Y-S, Lockwood P, Apps MAJ, Muhammed K, Husain M. (2017) Distinct subtypes of apathy revealed by the apathy Motivation Index. *PLoS One* **12**:e0169938

Apkarian AV, Baliki MN, Geha PY (2009) Towards a theory of chronic pain. *Prog Neurobiol* **87**:81-97

Baliki MN, Apkarian AV (2015) Nociception, pain, negative moods, and behavior selection. *Neuron* **87**:474-491

Brooks R, Boye KS, Slaap B. (2020) EQ-5D: a plea for accurate nomenclature. *J Patient Rep Outcomes* **4**:52

Brown CA, Seymour B, Boyle Y, El-Deredy W, Jones AKP (2008) Modulation of pain ratings by expectation and uncertainty: Behavioral characteristics and anticipatory neural correlates. *Pain* **135**:240-250

Chakroun K, Mathar D, Wiehler A, Ganzer F, Peters J. (2020) Dopaminergic modulation of the exploration/exploitation trade-off in human decision-making. *eLife* **9** https://doi.org/10.7554/eLife.51260

Cheng S, Zhou Y, Zhang W, Wu D, Yang C, Li B, Wang W. (2022) Uncertainty-aware and multigranularity consistent constrained model for semi-supervised hashing. *IEEE Trans Circuits Syst Video Technol* **32**:6914-6926

Dayan P, Kakade S, Montague PR (2000) Learning and selective attention. Nat. *Neurosci* **3 Suppl**:1218-1223

De La Torre Canales G, Câmara-Souza MB, Muñoz Lora VRM, Guarda-Nardini L, Conti PCR, Rodrigues Garcia RM, Del Bel Cury AA, Manfredini D. (2018) Prevalence of psychosocial impairment in temporomandibular disorder patients: A systematic review. *J Oral Rehabil* **45**:881-889

den Ouden HEM, Daw ND, Fernandez G, Elshout JA, Rijpkema M, Hoogman M, Franke B, Cools R. (2013) Dissociable effects of dopamine and serotonin on reversal learning. *Neuron* **80**:1572

Ebitz RB, Albarran E, Moore T. (2018) Exploration disrupts choicepredictive signals and alters dynamics in prefrontal cortex. *Neuron* **97**:450-461

Forstmann BU, Ratcliff R, Wagenmakers E-J. (2016) Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions. *Annu Rev Psychol* **67**:641-666

Husain M, Roiser JP (2018) Neuroscience of apathy and anhedonia: a transdiagnostic approach. *Nat Rev Neurosci* **19**:470-484

Korff V, Ormel M, Dworkin KF (1992) Grading the severity of chronic pain. *Pain* **50**:133-149

Kroenke K, Spitzer RL, Williams JB (2001) The PHQ-9: validity of a brief depression severity measure. *J Gen Intern Med* **16**:606-613

Little C, Naim N, Grogan T, Yu S, Chia PA (2025) Proof of concept study of the posterior quadratus lumborum block for laparoscopic myomectomy: A randomized controlled trial. *PLoS One* **20**:e0321890

National Academies of Sciences, Engineering, and Medicine, Health and Medicine Division, Board on Health Care Services, Board on Health Sciences Policy, Committee on Temporomandibular Disorders (TMDs): From Research Discoveries to Clinical Treatment (2020) *Temporomandibular disorders: Priorities for research and care* Washington D.C.: DC: National Academies Press.

Onysk J, Gregory N, Whitefield M, Jain M, Turner G, Seymour B, Mancini F. (2024) Statistical learning shapes pain perception and prediction independently of external cues. *eLife* **12** https://doi.org/10.7554/elife.90634.3

Piray P, Daw ND (2020) A simple model for learning in volatile environments. *PLoS Comput Biol* **16**:e1007963

Piray P, Dezfouli A, Heskes T, Frank MJ, Daw ND (2019) Hierarchical Bayesian inference for concurrent model fitting and comparison for group studies. *PLoS Comput Biol* **15**:e1007043

Quartana PJ, Campbell CM, Edwards RR (2009) Pain catastrophizing: a critical review. *Expert Rev Neurother* **9**:745-758

Samwel HJA, Evers AWM, Crul BJP, Kraaimaat FW (2006) The role of helplessness, fear of pain, and passive pain-coping in chronic pain patients. *Clin J Pain* **22**:245-251

Shivakumar S, Abdul NS, Jyoti B, Kalburgi V, Cicciù M, Minervini G. (2025) Comparative evaluation of cognitive behavioural therapy versus standard treatment in temporomandibular disorders: A systematic review. *J Oral Rehabil* **52**:521-530

Velly AM, Anderson GC, Look JO, Riley JL, Rindal DB, Johnson K, Wang Q, Fricton J, Huff K, Ohrbach R, *et al.* (2022) Management of painful temporomandibular disorders: Methods and overview of The National Dental Practice-Based Research Network prospective cohort study. *J Am Dent Assoc* **153**:144-157

Vlaeyen JWS, Linton SJ (2000) Fear-avoidance and its consequences in chronic musculoskeletal pain: a state of the art. *Pain* **85**:317-332

Weissman-Fogel I, Moayedi M, Tenenbaum HC, Goldberg MB, Freeman BV, Davis KD (2011) Abnormal cortical activity in patients with temporomandibular disorder evoked by cognitive and emotional tasks. *Pain* **152**:384-396

Wiech K, Tracey I. (2009) The influence of negative emotions on pain: behavioral effects and neural mechanisms. *Neuroimage* **47**:987-994

Yan X, Ebitz RB, Grissom N, Darrow DP, Herman AB (2025) Distinct computational mechanisms of uncertainty processing explain opposing exploratory behaviors in anxiety and apathy. *Biol Psychiatry Cogn Neurosci Neuroimaging* https://doi.org/10.1016/j.bpsc.2025.01.005

Yoshida W, Seymour B, Koltzenburg M, Dolan RJ (2013) Uncertainty increases pain: Evidence for a novel mechanism of pain modulation involving the periaqueductal gray. *J Neurosci* **33**:5638-5646

Peer reviews

Reviewer #1 (Public review):

Summary:

This study investigates how individuals with chronic temporomandibular disorder (TMD) learn from uncertain rewards, using a probabilistic three-armed bandit task and computational modelling. The authors aim to identify whether people living with chronic pain show altered learning under uncertainty and how such differences might relate to psychological symptoms.

Strengths:

The work addresses an important question about how chronic pain may influence cognition and motivation. The task design is appropriate for probing adaptive learning, and the modelling approach is novel. The findings of altered uncertainty updating in the TMD group are interesting.

Weaknesses:

Several aspects of the paper limit the strength of the conclusions. The group differences appear only in model-derived parameters, with no corresponding behavioural differences in task performance. Model parameters do not correlate with pain severity, making the proposed mechanistic link between pain and learning speculative. Some of the interpretations extend beyond what the data can directly support.

https://doi.org/10.7554/eLife.109185.1.sa2

Reviewer #2 (Public review):

Summary:

In this paper, the authors report on a case-control study in which participants with chronic pain (TMD) were compared to controls on performance of a three-option learning task. The authors find no difference in task behavior, but fit a model to this behavior and suggest that differences in the model-derived metrics (specifically, change in learning rate/estimated volatility/model estimated uncertainty) reveal a relevant between-group effect. They report a mediation effect suggesting that group differences on self-report apathy may be partially mediated by this uncertainty adaptation result.

Strengths:

The role of sensitivity to uncertainty in pathological states is an interesting question and is the focus of a reasonable amount of research at present. This paper provides a useful assessment of these processes in people with chronic pain.

Weaknesses:

(1) The interpretation of the model in the absence of any apparent behavioral effect is not convincing. The model is quite complex with a number of free parameters (what these parameters are is not well explained in the methods, although they seem to be presented in the supplement). These parameters are fitted to participant choice behavior - that is, they explain some sort of group difference in this choice behavior. The authors haven't been able to demonstrate what this difference is. The graphs of learning rate per group (Figure 2) suggest that the control group has a higher initial learning rate and a lower later learning rate. If this were actually the case, you would expect to see it reflected in the choice data (the control group should show higher lose-shift behavior earlier on, with this then declining over time, and the TMD group should show no change). This behavior is not apparent. The absence of a clear effect on behavior suggests that the model results are more likely to be spurious.

(2) As far as I could see, the actual parameters of the model are not reported. The results (Figure 2) illustrate the trial-level model estimated uncertainty/learning rate, etc, but these differ because the fitted model parameters differ. The graphs look like there are substantial differences in v0 (which was not well recovered), but presumably lambda, at least, also differs. The mean(SD) group values for these parameters should be reported, as should the correlations between them (it looks very much like they will be correlated).

(3) The task used seems ill-suited to measuring the reported process. The authors report the performance of a restless bandit task and find an effect on uncertainty adaptation. The task does not manipulate uncertainty (there are no periods of high/low uncertainty) and so the only adaptation that occurs in the task is the change from what appears to be the participants' prior beliefs about uncertainty (which appear to be very different between groups - i.e. the lines in Figure 2a,b,c are very different at trial 0). If the authors are interested in measuring adaptation to uncertainty, it would clearly be more useful to present participants with periods of higher or lower uncertainty.

(4) The main factor driving the better fit of the authors' preferred model over listed alternatives seems to be the inclusion of an additive uncertainty term in the softmax-this differentiates the chosen model from the other two Kalman filter-based models that perform less well. But a similar term is not included in the RW models-given the uncertainty of a binary outcome can be estimated as p(1-p), and the RW models are estimating p, this would seem relatively straightforward to do. It would be useful to know if the factor that actually drives better model fit is indeed in the decision stage (rather than the learning stage).

https://doi.org/10.7554/eLife.109185.1.sa1

Reviewer #3 (Public review):

This paper applies a computational model to behavior in a probabilistic operant reward learning task (a 3-armed bandit) to uncover differences between individuals with temporomandibular disorder (TMD) compared with healthy controls. Integrating computational principles and models into pain research is an important direction, and the findings here suggest that TMD is associated with subtle changes in how uncertainty is represented over time as individuals learn to make choices that maximize reward. There are a number of strengths, including the comparison of a volatile Kalman filter (vKF) model to some standard base models (Rescorla Wagner with 1 or 2 learning rates) and parameter recovery analyses suggesting that the combination of task and vKF model may be able to capture some properties of learning and decision-making under uncertainty that may be altered in those suffering from chronic pain-related conditions.

I've focused my comments in four areas: (1) Questions about the patient population, (2) Questions about what the findings here mean in terms of underlying cognitive/motivational processes, (3) Questions about the broader implications for understanding individuals with TMD and other chronic pain-related disorders, and (4) Technical questions about the models and results.

(1) Patient population

This is a computational modelling study, so it is light on characterization of the population, but the patient characteristics could matter. The paper suggests they were hospitalized, but this is not a condition that requires hospitalization per se. It would be helpful to connect and compare the patient characteristics with large-scale studies of TMD, such as the OPPERA study led by Maixner, Fillingim, and Slade.

(2) What cognitive/motivational processes are altered in TMD

The study finds a pattern of alterations in TMD patients that seems clear in Figure 2. Healthy controls (HC) start the task with high estimates of volatility, uncertainty, and learning rate, which drop over the course of the task session. This is consistent with a learner that is initially uncertain about the structure of the environment (i.e., which options are rewarded and how the contingencies change over time) but learns that there is a fixed or slowly changing mean and stationary variance. The TMD patients start off with much lower volatility, uncertainty, and learning rate - which are actually all near 0 - and they remain stable over the course of learning. This is consistent with a learner who believes they know the structure of the environment and ignores new information.

What is surprising is that this pattern of changes over time was found in spite of null group differences in a number of aspects of performance: (1) stay rate, (2) switch rate, (3) win-stay/lose-switch behaviors, (4) overall performance (corrected for chance level), (5) response times, (6) autocorrelation, (7) correlations between participants' choice probability and each option's average reward rate, (7) choice consistency (though how operationalized is not described?), (8) win-stay-lose-shift patterns over time. I'm curious about how the patterns in Figure 2 would emerge if standard aspects of performance are essentially similar across groups (though the study cannot provide evidence in favor of the null). It will be important to replicate these patterns in larger, independent samples with preregistered analyses.

The authors believe that this pattern of findings reveals that TMD patients "maintain a chronically heightened sensitivity to environmental changes" and relate the findings to predictive processing, a hallmark of which (in its simplest form) is precision-weighted updating of priors. They also state that the findings are not related to reduced overall attentiveness or failure to understand the task, but describe them as deficits or impairments in calibrating uncertainty.

The pattern of differences could, in fact, result from differences in prior beliefs, conceptualization of the task, or learning. Unpacking these will be important steps for future work, along with direct measures of priors, cognitive processes during learning, and precision-weighted updating.

(3) Implications for understanding chronic pain

If the findings and conclusions of the paper are correct, individuals with TMD and perhaps other pain-related disorders may have fundamental alterations in the ways in which they make decisions about even simple monetary rewards. The broader questions for the field concern (1) how generalizable such alterations are across tasks, (2) how generalizable they are across patient groups and, conversely, how specific they are to TMD or chronic pain, (3) whether they are the result of neurological dysfunction, as opposed to (e.g.) adaptive strategies or assumptions about the environment/task structure.

It will be important to understand which features of patients' and/or controls' cognition are driving the changes. For example, could the performance differences observed here be attributable to a reduced or altered understanding of the task instructions, more uncertainty about the rules of the game, different assumptions about environments (i.e., that they are more volatile/uncertain or less so), or reduced attention or interest in optimizing performance? Are the controls OVERconfident in their understanding of the environment?

This set of questions will not be easy to answer and will be the work of many groups for many years to come. It is a judgment call how far any one paper must go to address them, but my view is that it is a collaborative effort. Start with a finding, replicate it across labs, take the replicable phenomena and work to unpack the underlying questions. The field must determine whether it is this particular task with this model that produces case-control differences (and why), or whether the findings generalize broadly. Would we see the same findings for monetary losses, sounds, and social rewards? Tasks with painful stimuli instead of rewards?

Another set of questions concerns the space of computational models tested, and whether their parameters are identifiable. An alteration in estimated volatility or learning rate, for example, can come from multiple sources. In one model, it might appear as a learning rate change and in another as a confirmation bias. It would be interesting in this regard to compare the "mechanisms" (parameters) of other models used in pain neuroscience, e.g., models by Seymour, Mancini, Jepma, Petzschner, Smith, Chen, and others (just to name a few).

One immediate next step here could be to formally compare the performance of both patients and controls to normatively optimal models of performance (e.g., Bayes optimal models under different assumptions). This could also help us understand whether the differences in patients reflect deficits and what further experiments we would need to pin that down.

In addition, the volatility parameter in the computational model correlated with apathy. This is interesting. Is there a way to distinguish apathy as a particular clinical characteristic and feature of TMD from apathy in the sense of general disinterest in optimal performance that may characterize many groups?

If we know this, what actionable steps does it lead us to take? Could we take steps to reduce apathy and thus help TMD patients better calibrate to environmental uncertainty in their lives? Or take steps to recalibrate uncertainty (i.e., increase uncertainty adaptation), with benefits on apathy? A hallmark of a finding that the field can build off of is the questions it raises.

(4) Technical questions about the models and results

Clarification of some technical points would help interpret the paper and findings further:

(a) Was the reward probability truly random? Was the random walk different for each person, or constrained?

(b) When were self-report measures administered, and how?

(c) Pain assessments: What types of pain? Was a body map assessed? Widespreadness? Pain at the time of the test, or pain in general?

(d) Parameter recovery: As you point out, r = 0.47 seems very low for recovery of the true quantity, but this depends on noise levels and on how the parameter space is sampled. Is this noise-free recovery, and is it robust to noise? Are the examples of true parameters drawn from the space of participants, or do they otherwise systematically sample the space of true parameters?

(e) What are the covariances across parameter estimates and resultant confusability of parameter estimates (e.g., confusion matrix)?

(f) It would be helpful to have a direct statistical comparison of controls and TMD on model parameter estimates.

(g) Null statistical findings on differences in correlations should not be interpreted as a lack of a true effect. Bayes Factors could help, but an analysis of them will show that hundreds of people are needed before it is possible to say there are no differences with reasonable certainty. Some journals enforce rules around the kinds of language used to describe null statistical findings, and I think it would be helpful to adopt them more broadly.

(h) What is normatively optimal in this task? Are TMD patients less so, or not? The paper states "aberrant precision (uncertainty) weighting and misestimation of environmental volatility". But: are they misestimates?

(i) It's not clear how well the choice of prior variance for all parameters (6.25) is informed by previous research, as sensible values may be task- and context-dependent. Are the main findings robust to how priors are specified in the HBI model?

https://doi.org/10.7554/eLife.109185.1.sa0