

Computational Growth Modeling: Assessing Longitudinal Change with Joint Models

Jeremy Haynes

jmhaynes@georgiasouthern.edu

Georgia Southern University

Nathaniel Haines

Bayesian Beginnings

Holly Sullivan-Toole

University of Minnesota

Thomas Olino

Temple University

Method Article

Keywords: joint modeling, computational modeling, longitudinal growth modeling, hierarchical Bayesian analysis, reinforcement learning models, Iowa Gambling Task

Posted Date: October 21st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11269872/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Computational Growth Modeling: Assessing Longitudinal Change with Joint Models

Jeremy M. Haynes^a, Nathaniel Haines^b, Holly Sullivan-Toole^c, & Thomas M. Olino^d

^aDepartment of Psychology, Georgia Southern University, Statesboro, GA 30458 USA, ORCID: 0000-0002-0811-3849

^bBayesian Beginnings, Columbus, OH 43201 USA, ORCID: 0000-0001-6654-9010

^cDepartment of Pediatrics, Masonic Institute for the Developing Brain, University of Minnesota, Minneapolis, MN 55414, USA, ORCID: 0000-0001-9617-9276

^dDepartment of Psychology & Neuroscience, Philadelphia, PA 19122 USA, ORCID: 0000-0001-5139-8571

Author Note

Correspondence regarding this manuscript should be addressed to Jeremy Haynes, Department of Psychology, Georgia Southern University 2670 Southern Dr, Statesboro, GA 30458 (jmhaynes@georgiasouthern.edu) or Thomas Olino (thomas.olino@temple.edu).

Some of the data presented in this manuscript have been presented elsewhere (Haynes et al., 2025, 2026). Details of these data are described in the Method.

Statements and Declarations

Funding

This work was supported by National Institute of Mental Health grants (R01 MH107495 & R21 MH130792) awarded to TMO. The authors have no known competing interests to declare.

Data & Code Availability Statement

Code used in the following sections are available on the Open Science Framework (OSF) at https://osf.io/xn6g8/overview?view_only=3fab1d0fad404f258bb860dc94b2ecce. Data from the *Real Data Application* available upon request.

Ethics Approval

All procedures were approved by the Temple University institutional review board (IRB Protocol #23174 originally approved on 08/31/2015). All participants provided informed consent.

Abstract

Computational models involve fitting a behavioral model to data from a behavioral task to estimate parameters characterizing task performance. Frequently, computational models are used to clarify associations between psychological constructs (e.g., reward learning) and mental illness (e.g., depression). A common method for assessing the etiology of mental illness involves longitudinal research designs that evaluate how behavioral processes change over time. In this paper, we illustrate how hierarchical computational modeling can be extended to longitudinal data via joint modeling. We first provide a brief overview of common approaches to longitudinal modeling and then describe how to integrate computational and longitudinal analyses using joint modeling; an approach we refer to as *computational growth modeling*. Next, using simulation, we show how computational growth modeling improves upon more traditional approaches for analyzing longitudinal change in parameters from a simple reinforcement learning model. In particular, the computational growth modeling approach estimates person-level longitudinal growth parameters better than a more traditional, two-stage approach whereby computational parameters and growth trajectories are modeled separately. After this initial proof-of-concept simulation, we show how computational growth modeling can be applied to real longitudinal data collected from five waves of adults tested on the play-or-pass Iowa Gambling Task.

Ultimately, our goal is to illustrate how to use the computational growth modeling approach, which shows promise for more precisely modeling changes in behavioral processes across time, and has particular relevance to understanding how the development of these behavioral processes may play an etiological role in psychopathology.

Keywords: joint modeling, computational modeling, longitudinal growth modeling, hierarchical Bayesian analysis, reinforcement learning models, Iowa Gambling Task

Computational Growth Modeling: Assessing Longitudinal Change with Joint Models

Studies assessing behavior across multiple occasions, either across time or across tasks, provide a means of characterizing a wide repertoire of behavior compared to studies that employ single assessments of behavior. In the past 15 years, the importance of employing multi-assessment approaches has become increasingly recognized following the introduction of the National Institutes of Mental Health's Research Domain Criteria (RDoC) which aim to identify patterns of dysfunction in mental illness based on individual differences in both biological and behavioral processes (Insel, 2009). As research designs become more complex using multi-assessment approaches, more sophisticated strategies to characterize patterns of behavior on these assessments have been developed. Here, we discuss an implementation of joint modeling for assessing longitudinal change in parameters from behavioral models that characterize complex patterns of behavior (e.g., learning rates from a reinforcement learning model). Joint modeling is a form of hierarchical modeling that integrates behavioral and statistical models and provides a means of more adequately accounting for uncertainty in how parameters are estimated at both the person and group levels. We begin with a brief introduction to standard practices in longitudinal modeling. Next, we introduce how to integrate a behavioral model with a growth model using a joint modeling approach to assess change across time in computationally derived parameters.

Longitudinal Modeling

Longitudinal research takes many forms, with some focusing on test-retest reliability¹, and others focusing on change across wider spans of development. In its simplest form, test-retest reliability studies measure behavior across two timepoints to examine mean-level

¹ Here, we include studies measuring subjects on two occasions within longitudinal research; however, typically, longitudinal research involves measuring subjects on three or more occasions (Ployhart & MacKenzie, 2014).

differences and rank-order consistency of behavioral measures from session 1 to session 2 (e.g., Buelow & Barnhart, 2018; Schmitz et al., 2020). Studies employing more than two assessments frequently focus on assessing developmental change, course of treatment, or intervention outcomes that unfold over time. These latter studies are critical for identifying how behavioral processes play a role in socially relevant individual differences (e.g., substance abuse; Cicchetti & Toth, 2009).

With summary measures of behavior (e.g., session-wide choice proportions), standard general linear models are frequently employed to assess change over time. General linear models include the simple paired-samples *t*-test and repeated measures analysis of variance (RM-ANOVA) which estimate mean-level differences across sessions (Oehlert, 2010). More complex modeling strategies, such as hierarchical linear models (HLMs; also known as multilevel models) and latent growth curve models (LGCs), provide additional flexibility for considering underlying trajectories of change that further explain the mean-level changes in outcomes via the addition of covariates (Bryk & Raudenbush, 1992). Although HLMs and LGCs differ in terms of the organization of data structures, estimation of these models is identical when the same constraints are applied across models (Curran, 2003). For these models, change across time is characterized by estimates of a group-level intercept (i.e., starting point) and slope (i.e., rate of change), with random effects capturing individual differences in intercepts and slopes. Hierarchical linear and latent growth curve models are indispensable approaches for analyzing longitudinal data because of their ability to account for the non-independence of data collected from repeated measures and the flexibility in characterizing change across time at both the group and person levels simultaneously.

Although HLMs and LGCMs are valuable tools for assessing change in molar measures of behavior (i.e., those aggregated at the session level), standard implementations of these approaches (e.g., lmer in R; Bates et al., 2015) do not allow us to model more complex patterns of behavior within and across sessions simultaneously. As an example, consider a study that assesses performance on the Iowa Gambling Task (IGT) across multiple sessions. The IGT is a commonly used behavioral task in which participants choose between decks of cards that vary in their short- and long-term gains and losses. Aggregate measures of behavior at the session level (e.g., proportion of choices for the advantageous decks) can be computed to characterize how well people identify which decks are advantageous and disadvantageous. As part of a longitudinal study, we can assess changes in these session-level behavioral measures with one or more of the approaches described above (e.g., HLM). However, more recent approaches to analyzing IGT performance rely on computational models that characterize patterns of behavior on the task via parameters measuring trial-level learning, typically in the form of a reinforcement learning (RL) model. To employ such a model, the researcher is now faced with fitting two models: the computational model to characterize within-task behavior (e.g., with an RL model) and the statistical model to characterize across-task behavior (e.g., with an HLM to assess longitudinal change).

A common method of working with two models to characterize within- and across-task behavior is the two-stage approach. In Stage 1, the computational model (e.g., RL model) is fit to the task data for each session individually. In Stage 2, the statistical model (e.g., HLM) is fit to the computational parameters extracted from Stage 1. A disadvantage of this approach is that, in the second stage, person-level parameters are implicitly treated as true scores (i.e., perfectly measured) when they are actually contaminated with measurement error (Haines et al., 2023,

2025). That is, the uncertainty in estimating the person-level parameters from the within-task data does not propagate to the statistical model which can bias inferences drawn from the statistical model. This is not a new issue — many researchers have noted the shortcomings of two-stage approaches including reduced power, poor reliability, convergence failures, among others (Haines et al., 2023; Kaplan et al., 2021; Young, 2018). Although much can still be learned using two-stage approaches (e.g., developmental changes in delay discounting & reward-related processing; Klein et al., 2022; Sullivan-Toole et al., 2026), we can improve our modeling of longitudinal data from behavioral tasks by further integrating longitudinal growth models and computational models via joint modeling.

Joint Modeling

Broadly speaking, joint modeling is a form of hierarchical modeling that integrates a behavioral model and a statistical model to estimate parameters from each model simultaneously. By estimating parameters simultaneously, information about a parameter value from the behavioral model informs parameter estimation for the statistical model and vice versa. This approach has been successfully used to improve test-retest reliability (Brown et al., 2020; Haines et al., 2025; Matzke et al., 2017) as well as for assessing how patterns of behavior across different behavioral tasks are related (Kvam et al., 2021; Haynes et al., 2026). In general, a joint model involves specifying a person-level model to characterize performance on a behavioral task and a group-level model to characterize how performance varies across assessments:

$$\begin{aligned} Y_{i,c,j} | C_{c,j}, \theta_{i,j} &\sim \text{person-level model} \\ \theta_{i,j} | T_j, \Psi &\sim \text{group-level model} \end{aligned} \tag{Equation 1}$$

For the person-level model, the observed data ($Y_{i,c,j}$) are assumed to vary in response to task stimuli ($C_{c,j}$) according to a parameter ($\theta_{i,j}$) that summarizes how an individual responds to the

task stimuli depending on the structure of the behavioral model. For the group-level model, person-level parameters are assumed to vary across time (T_j) according to group-level parameters (Ψ) that summarize how the computational parameters change across time.

Equation 1 is a very general description of joint modeling such that any person or group-level model may be specified based on a researcher's assumptions about task behavior. For example, in the applications that follow, we use reinforcement learning models as the person-level models, but the specific behavioral model used will depend on the behavioral task (e.g., the hyperbolic discounting model for intertemporal choice task data, Mazur, 1987). Likewise, the group-level model can take many forms. Prior joint modeling work has frequently specified the group-level model as a multivariate normal distribution or a linear model across two assessments (e.g., Haynes et al., 2024; Kvam et al., 2021; Sullivan-Toole et al., 2022). This application of joint modeling has been shown to improve rank-order consistency of reinforcement learning parameters — e.g., reward learning rates from $r = .39$ to $r = .73$ (Sullivan-Toole et al., 2022) — as well as for identifying how patterns of behavior are related across tasks that presumably measure similar behavioral processes (e.g., sensitivity to delayed consequences on an intertemporal choice task & the IGT; Haynes et al., 2026). Here, we extend this work by integrating a longitudinal growth model with a computational model to assess longitudinal change in person-level (i.e., computationally derived) parameters, an approach we refer to as computational growth modeling. To do this, we first illustrate the mathematical foundation of this integration. Next, we show proof-of-concept of the computational growth modeling approach using simulation with a simple model and an example with a more complex model applied to real data.

Computational Growth Modeling

To integrate computational and growth models, we specify a joint model that will allow us to simultaneously estimate computational parameters from a behavioral model (to capture within-task performance) with growth parameters from a longitudinal model (to capture change across time in the computational parameters). The approach is similar to that used by Casabianca et al. (2017) for analyzing longitudinal change in rater scores using Item Response Theory; thus, we use similar notation for consistency. In its simplest form, without covariates, the group-level model is identical to a standard unconditional longitudinal growth model:

$$\theta_{i,j} = \beta_{0i} + \beta_{1i}T_{i,j} + E_{i,j} \quad \text{Equation 2}$$

where $\theta_{i,j}$ is a computationally derived parameter for person i on occasion j , β_{0i} is the person-level intercept, β_{1i} is the person-level slope, $T_{i,j}$ represents time, and $E_{i,j}$ is a residual. As is common practice for growth models, we assume normally distributed residuals with a mean of 0 and standard deviation of σ_1 (Bryk & Raudenbush, 1992). We specify person-level growth parameters (i.e., β_{0i} & β_{1i}) as drawn from a multivariate normal distribution:

$$\begin{pmatrix} \beta_{0i} \\ \beta_{1i} \end{pmatrix} \sim \mathcal{N} \begin{pmatrix} \gamma_0 & \sigma_{20}^2 & \sigma_{01} \\ \gamma_1 & \sigma_{01} & \sigma_{21}^2 \end{pmatrix} \quad \text{Equation 3}$$

$$\sigma_{01} = \rho\sigma_{20}\sigma_{21}$$

where γ_0 and γ_1 represent the group-level intercept and slope, respectively; σ_{20}^2 and σ_{21}^2 represent intercept and slope variances, respectively; and ρ represents the correlation between intercepts and slopes.

Equations 2 and 3 can be further extended to include covariates that moderate intercepts and slopes, as in standard conditional growth models. For this, we use an approach similar to that of Vandekerckhove and Tuerlinckx (2007) for developing joint models to assess between-

condition differences for diffusion models. Using matrix notation, we can express the group-level model equations as the following:

$$\begin{aligned}
 \boldsymbol{\theta} &= \boldsymbol{\beta}_0 + \boldsymbol{\beta}_1 \mathbf{T} + \mathbf{E}_1 & \text{Equation 4} \\
 \boldsymbol{\beta}_0 &= \mathbf{X}_0 \boldsymbol{\gamma}_0 + \mathbf{E}_{20} \\
 \boldsymbol{\beta}_1 &= \mathbf{X}_1 \boldsymbol{\gamma}_1 + \mathbf{E}_{21} \\
 \mathbf{E}_1 &\sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_1) \\
 \begin{pmatrix} \mathbf{E}_{20} \\ \mathbf{E}_{21} \end{pmatrix} &\sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_2) \\
 \boldsymbol{\Sigma}_1 &= \sigma_1^2 \mathbf{I} \\
 \boldsymbol{\Sigma}_2 &= \begin{pmatrix} \sigma_{20}^2 & \mathbf{0} \\ \mathbf{0} & \sigma_{21}^2 \end{pmatrix} \mathbf{R} \begin{pmatrix} \sigma_{20}^2 & \mathbf{0} \\ \mathbf{0} & \sigma_{21}^2 \end{pmatrix} \\
 \mathbf{R} &= \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}
 \end{aligned}$$

where $\boldsymbol{\theta}$ is an $N \times T$ matrix of computational parameters, $\mathbf{X}_0$ and $\mathbf{X}_1$ are $N \times M_0$ and $N \times M_1$ design matrices with covariates of interest (N indicating number of subjects & M indicating number of covariates), $\boldsymbol{\gamma}_0$ and $\boldsymbol{\gamma}_1$ are M_0 - and M_1 -length coefficient vectors capturing covariate effects on the intercept and slope growth components, respectively, and $\mathbf{R}$ is a correlation matrix capturing the correlation between intercept and slope parameters across participants per the level 2 residual variation (i.e. $\mathbf{E}_{20}$ & $\mathbf{E}_{21}$ which are M_0 - & M_1 -length vectors, respectively). $\mathbf{E}_1$ is then a residual matrix of dimension $N \times T$, which represents level 1 residual variation of the person-session-level computational parameters from the person-level growth models implied by $\boldsymbol{\beta}_0$ and $\boldsymbol{\beta}_1$ —we assume these residuals are normal and independently distributed across people and sessions.

The structure of the design matrices, $\mathbf{X}_0$ and $\mathbf{X}_1$, in Equation 4 allow us to examine longitudinal growth in multiple ways. In the simplest case, we could specify both matrices with a

single column vector of 1s, in which case $\boldsymbol{\gamma}_0$ and $\boldsymbol{\gamma}_1$ become the group-level intercepts and slopes across all participants, and Equation 4 reduces to Equations 2 and 3. Such a model is appropriate in cases when we have no hypothesized predictors to explain variation in growth parameters across people. Conversely, if we expect variation in growth across people as a function of individual differences (e.g., history of mental illness), we can specify $\mathbf{X}_0$ and $\mathbf{X}_1$ with a column vector of 1s plus a column vector for the covariate(s) (e.g., a 1 indicator for if a person has a history of depression, otherwise 0). Then, $\gamma_{0,1}$ and $\gamma_{1,1}$ become intercepts for the level 2 growth model terms (in this case, representing the control group), and $\gamma_{0,2}$ and $\gamma_{1,2}$ become slopes for the level 2 growth model terms (representing the difference between the control group & depression group). It is worth emphasizing here that there are two different sets of intercepts and slopes in the model—the level 2 intercepts and slopes (i.e., $\boldsymbol{\gamma}_0$ & $\boldsymbol{\gamma}_1$) that capture group-level effects (a.k.a. “fixed effects”) given covariates of interest, and the level 1 intercepts and slopes (i.e., $\boldsymbol{\beta}_0$ & $\boldsymbol{\beta}_1$) that capture each person’s unique trajectory of computational model parameters across sessions.

Group- and person-level growth parameters from Equation 4, or Equations 2-3 when no covariates are present, are estimated simultaneously with computational parameters in a joint model, expressed as the following:

$$\begin{aligned} \mathbf{Y}|\mathbf{C}, \boldsymbol{\theta} &\sim \text{computational model} \\ \boldsymbol{\theta}|\mathbf{X}_0, \mathbf{X}_1, \mathbf{T}, \boldsymbol{\Psi} &\sim \text{growth model} \end{aligned} \tag{Equation 5}$$

where computational model parameters ($\boldsymbol{\theta}$) characterize task performance ($\mathbf{Y}$) based on the task-related variables ($\mathbf{C}$) and the growth model parameters ($\boldsymbol{\Psi}$) capture changes in the computational parameters across time ($\mathbf{T}$) and as a function of covariates (embedded within $\mathbf{X}_0$ & $\mathbf{X}_1$), if

present. Note, Equation 5 is simply a recapitulation of Equation 1, but specifically with longitudinal growth as the group-level model.

Now that the framework for the joint model is specified, we can now illustrate how such a model could be implemented. We do this first with simulation to show how the computational growth modeling approach improves upon a two-stage approach for assessing change across time on a simple decision-making task. Next, we show how the computational growth modeling approach could be implemented for a more complex model used to characterize performance on the play-or-pass version of the Iowa Gambling Task using real longitudinal data. Both implementations involve the use of reinforcement learning algorithms to characterize trial-level behavior. Our goal here is not to make a case for reinforcement learning models; rather, our goal is to illustrate how one could apply computational growth modeling to characterize performance on a behavioral task assessed longitudinally with the hopes that researchers can generalize the approach to their specific research context.

Simulation-Based Application of Computational Growth Modeling

In this first application, we simulated data from a simple two-armed bandit task according to a single-parameter reinforcement learning model. Reinforcement learning (RL) models are a frequently-used class of computational models for analyzing trial-level data from decision-making tasks (Doll et al., 2012). With reinforcement learning models, we estimate parameters that characterize how the value of choosing an option changes across sequential trials based on gains and losses from prior choices according to a Rescorla-Wagner learning rule (Rescorla & Wagner, 1972)²:

$$v(t + 1) = v(t) + A \cdot (C(t) - v(t)) \quad \text{Equation 6}$$

² This is a simplified version of the Rescorla-Wagner learning rule.

where $v(t)$ is the value of choosing the option on trial t , A is a learning rate, and $C(t)$ is the obtained outcome on trial t . In words, Equation 6 states that after trial t , the value of choosing the option is updated by the weighted difference between the obtained outcome on trial t and the value of that outcome on trial t . The difference between the obtained and current (i.e., expected) value of the outcome is typically referred to as the prediction error and the weighting parameter A (i.e., learning rate) characterizes the degree to which the prediction error changes the value of the option. Estimating learning rates typically involves fitting a form of Equation 6 to trial-level data using a Bernoulli distribution:

$$Y(t) \sim \text{Bernoulli}\left(\frac{1}{1 + \exp(-v(t))}\right) \quad \text{Equation 7}$$

where $Y(t)$ is the observed choice (chosen: $Y(t) = 1$; not chosen: $Y(t) = 0$) and $v(t)$ is the value of the option, updated from trial to trial according to Equation 6. Equation 7 may be thought of as a logistic regression model in which choices are modeled as a function of the sequentially updated value of the option.

For the simulation, we use a similar framework to illustrate the construction of a computational growth model (CGM). To do this, we simulated 100 datasets across nine conditions (thus, 900 datasets total). For each dataset, we generated data from a 100-trial two-armed bandit task for 50 “participants”, assessed across three sessions (thus, 300 trials worth of data per participant). Each trial consisted of a choice between two options resulting in either reward (+1) or no reward (0), with a 10% and 90% probability of reward for options 1 and 2, respectively. Conditions differed according to a parametric manipulation of three group-level slope effects (γ_1), characterizing overall change across time in learning rates, crossed with three intercept-slope correlations (ρ), characterizing the correlation between person-level intercepts (i.e., starting learning rates) and slopes (i.e., rate of change in learning rates; see Table 1).

Parameter values were chosen to produce observable changes in behavior across time as well as to explore performance of the models across various growth trajectories (i.e., group-level slope effects) and individual difference profiles (i.e., intercept-slope correlations).

We used the group-level parameters to simulate all subsequent parameters which were then used to generate trial-level data. Specifically, for each condition, person-level intercepts and slopes (i.e., growth parameters) were randomly drawn from a multivariate normal distribution based on the group-level parameters. The person-level growth parameters were used to create session-specific learning rates for each participant, and these learning rates were subsequently used to generate trial-level choice data. Next, we fit models to the trial-level choice data using a two-stage approach and the CGM approach. For both approaches, we assessed recovery of RL parameters (i.e., learning rates) and growth parameters (i.e., intercepts & slopes).

Table 1
Conditions for Simulations

Condition						
Slope	Correlation	γ_0	γ_1	σ_{20}^2	σ_{21}^2	ρ
-Moderate	-Moderate	-1.64	-.5	1.0	1.0	-.3
-Moderate	None	-1.64	-.5	1.0	1.0	0
-Moderate	+Moderate	-1.64	-.5	1.0	1.0	+.3
None	-Moderate	-1.64	0	1.0	1.0	-.3
None	None	-1.64	0	1.0	1.0	0
None	+Moderate	-1.64	0	1.0	1.0	+.3
+Moderate	-Moderate	-1.64	+.5	1.0	1.0	-.3
+Moderate	None	-1.64	+.5	1.0	1.0	0
+Moderate	+Moderate	-1.64	+.5	1.0	1.0	+.3

Note. True parameter values used in simulations.

Simulation Method

As described above, we simulated learning rates for the two-armed bandit task according to a longitudinal growth model (cf. Equation 2):

$$\gamma_{i,j} = \beta_{0i} + \beta_{1i}T_{i,j} + E_{i,j} \quad \text{Equation 8}$$

where A_{ij} is an unbounded learning rate for person i on session j , β_{0i} is the person-level intercept, β_{1i} is the person-level slope, T_{ij} is the session number ranging from 1 to 3, and a residual given by $E_{ij} \sim \mathcal{N}(0, 1)$. Learning rates (A_{ij}) are bounded between 0 and 1; thus, we transformed the unbounded learning rates using the cumulative distribution function for the standard normal distribution such that $A_{ij} = \Phi(A_{ij})$. Group-level parameters for each condition are displayed in Table 1. Person-level growth parameters (i.e., β_{0i} & β_{1i}) were simulated from these group-level parameters by drawing from a multivariate normal distribution (Equation 3). Note that, as shown in Table 1, the variances were all 1.0; thus, the covariance between intercepts and slopes (σ_{01}) also represents the correlation (ρ) between intercepts and slopes.

After simulating learning rates (A_{ij}), we generated trial-level choice data according to a Bernoulli distribution:

$$Y(t)_{i,j} \sim \text{Bernoulli} \left(\frac{1}{1 + \exp(-(v_2(t)_{i,j} - v_1(t)_{i,j}))} \right) \quad \text{Equation 9}$$

where $Y(t)_{i,j}$ is the choice for option 1 ($Y = 0$) or option 2 ($Y = 1$) on trial t during the j th session for participant i , and $v_1(t)_{i,j}$ and $v_2(t)_{i,j}$ are the expected values associated with choosing option 1 or 2, updated from trial to trial according to a Rescorla-Wagner learning rule (cf. Equation 6):

$$v_k(t + 1)_{i,j} = A_{i,j} \cdot (C(t)_{i,j} - v_k(t)_{i,j}) \quad \text{Equation 10}$$

where $k = 1$ for option 1 and $k = 2$ for option 2, $A_{i,j}$ is the learning rate, and $C(t)_{i,j} = +1$ for reward and 0 for no reward.

Model Implementation

For each dataset, we estimated learning rates and growth parameters using a two-stage approach and the computational growth modeling (CGM) approach. Both approaches involved hierarchical Bayesian estimation implemented in Stan (v. 2.32.7; Stan Development Team,

2025b), a probabilistic programming language which estimates parameters using Hamiltonian Monte Carlo, a variant of Markov chain Monte Carlo (MCMC) to sample from high-dimensional probabilistic models. We used R (v. 4.4.3; R Core Team, 2025) to interface with Stan via the RStan package (v. 2.32.7; Stan Development Team, 2025a). For both approaches, we sampled each model using 4 chains, with each chain consisting of 500 warmup samples and 500 target samples (thus, 1000 iterations/chain), with weakly informative priors (Table S1)³.

For the two-stage approach, we first estimated learning rates (i.e., A_{ij}) for each session individually using models based on Equations 9 and 10 (first stage). We calculated posterior means of the unbounded learning rates (i.e., $\hat{A}_{ij}$) for each participant at each session to be used as point estimates of the person-level learning rates. Unbounded learning rates were used to ensure that the estimand was the same across the two-stage and CGM approaches. We used the person-level point estimates of $\hat{A}_{ij}$ to estimate growth parameters from Equations 3 and 8 with brms (Bürkner, 2017), an R package for fitting multilevel models with Stan (second stage). Note that because growth parameters were estimated from point estimates of the unbounded learning rates, uncertainty in those learning rates does not carry over into estimating the growth parameters. In contrast, for the CGM approach, we estimated learning rates with Equations 9 and 10 for all sessions simultaneously with growth model parameters from Equations 3 and 8 in a joint model (represented by Equation 5), thus allowing uncertainty in estimating the learning rates to inform estimation of the growth parameters, and vice versa.

Data Analysis

To compare the two-stage and CGM approaches, we assessed parameter recovery for each approach. First, we assessed the discrepancy between the true value of a parameter and the

³ We used the same priors for the two-stage and CGM approach.

estimated value of the parameter. To do this, we computed posterior means of the group- and person-level parameters. For the group-level parameters, we plot the distributions of these estimated parameters for each condition across both approaches and then visually assess the degree to which these distributions centered around the true parameter values and the direction of bias in the estimated parameters. To summarize the magnitude of these discrepancies, we computed mean absolute deviations by calculating the difference between the true and estimated parameter, taking the absolute value of that difference, and then averaging those absolute differences for each parameter across both approaches. For person-level parameters, comparisons of true and estimated parameters consisted of calculating the correlation between true and estimated person-level parameters for each simulation (i.e., parameter recovery; Wilson & Collins, 2019), resulting in 100 correlations per condition for each person-level parameter. We summarize these correlations by plotting the distributions of correlations for each condition and by calculating the average correlation for each condition across the two-stage and CGM approaches. In addition, we also calculated mean absolute deviations between the true and estimated person-level parameters using the same approach described for the group-level parameters.

Next, we calculated the degree to which each approach captured the true parameter values within the respective posterior distributions. For this, we calculated 95% coverage probabilities for each parameter in each condition. Assuming a well-fitting model, the 95% highest density intervals (HDIs; cf. confidence intervals) for each parameter should capture the true parameter value approximately 95% of the time in the long run. We calculated 95% coverage probabilities as the percentage of simulations in which the 95% HDI of the posterior distributions encompassed the true parameter value. To summarize these probabilities, we plot

the coverage probabilities for each parameter and each condition across both approaches and we computed mean absolute deviations by calculating the difference between the obtained coverage probability (i.e., % of simulated HDIs that contained the true parameter) and the target 95% probability, taking the absolute value of this difference, and then averaging the absolute deviations across conditions for each parameter.

The Online Resource consists of additional analyses that expand upon the above. Specifically, we show percentile plots which are an extension of assessing coverage probabilities that further characterize bias (e.g., overdispersion) in the posterior distributions (cf. simulation-based calibration; Talts et al., 2018). In addition, we show how long estimation took to complete for each approach as well as R-hat values to assess chain convergence across approaches (Gelman & Rubin, 1992).

Results

Group Level Parameters

Overall, the computational growth modeling (CGM) approach estimated the group-level growth parameters better than the two-stage approach. To illustrate, Figure 1 depicts true and estimated group-level parameters across the two-stage and CGM approaches. First, the two-stage approach overestimated the group-level intercept (γ_0) in conditions predicting decreasing learning rates across time ($\gamma_1 = -0.30$) and underestimated the group-level intercept in conditions predicting increasing learning rates across time ($\gamma_1 = +0.30$). The two-stage approach also tended to overestimate the group-level slope coefficient (γ_1) in conditions predicting decreasing learning rates ($\gamma_1 = -0.30$) but otherwise estimated these parameters similar to the CGM approach. Both approaches showed bias in estimating the correlation between intercepts and slopes (ρ) but typically in opposite directions. Specifically, the two-stage approach systematically

underestimated the intercept-slope correlations whereas the CGM approach frequently overestimated these correlations in conditions with negative ($\rho = -.30$) or no correlation ($\rho = 0$) between intercepts and slopes. In addition, the two-stage approach systematically underestimated the standard deviations of the intercept and slope coefficients (σ_{20} & σ_{21}) as well as the residual standard deviation (σ_1), whereas the CGM approach generally estimated these parameters well. Table 2 summarizes these findings with the mean absolute deviations (MAD) between the true and estimated growth parameters for both approaches. With the exception of the residual standard deviation (σ_1), the CGM approach had either similar or lower MADs than the two-stage approach.

Table 2
Mean Absolute Deviations for Group-Level Parameters

		True vs. Estimate		Coverage	
		Two-Stage	CGM	Two-Stage	CGM
Intercept:	γ_0	0.33 (0.24)	0.19 (0.14)	28.33	1.33
Slope:	γ_1	0.17 (0.13)	0.17 (0.13)	27.89	2.00
Standard Deviation of Intercepts:	σ_{20}	0.48 (0.20)	0.25 (0.17)	45.11	2.56
Standard Deviation of Slopes:	σ_{21}	0.60 (0.14)	0.25 (0.18)	93.78	2.56
Intercept-Slope Correlation:	ρ	0.37 (0.26)	0.22 (0.16)	6.67	2.89
Residual:	σ_1	0.14 (0.08)	0.19 (0.15)	35.56	2.78

Note. Mean absolute deviations (MADs; ± 1 SD) between true and estimated parameters and between the 95% target probability and the obtained coverage probability for the group-level parameters.

Figure 1
Distributions of Group-Level Growth Parameters




Fig. 1 Distributions of posterior means for group-level growth model parameters. Rows correspond to different simulation conditions and columns correspond to different group-level parameters. Vertical red lines indicate true values and datapoints represent the average of the posterior means across simulations.

Figure 2 shows coverage probabilities for the group-level parameters across both conditions. In general, the CGM approach was closer to the 95% target coverage probability than the two-stage approach and showed a similar pattern as in Figure 1. The two-stage approach

showed lower coverage for the group-level intercept (γ_0) and slope (γ_1) coefficients in the conditions with learning rates that either decreased ($\gamma_1 = -0.30$) or increased ($\gamma_1 = +0.30$) across time. In general, both approaches showed adequate coverage for the intercept-slope correlations (ρ), although the two-stage approach had systematically lower coverage than the target 95%. In addition, the two-stage approach had lower coverage for the standard deviations of the intercepts (σ_{20}) and particularly poor coverage for the standard deviations of the slopes (σ_{21}) across all conditions. Finally, the two-stage approach showed lower coverage for the residual standard deviations (σ_1). These findings are summarized in Table 2 which displays mean absolute deviations (MADs) between observed coverage probabilities and the target 95% probability for each parameter. Consistent with Figure 2, MADs for the two-stage approach were systematically higher (i.e., worse) than for the CGM approach across group-level parameters.

Figure 2
95% Coverage for Group-Level Growth Parameters




Fig. 2 95% coverage probabilities for group-level growth parameters from each condition across the two-stage and CGM approaches. Columns correspond to different levels of the group-level slope coefficient (γ_1). Rows correspond to different levels of the correlation between intercepts and slopes (ρ). The target 95% probability is indicated by the dashed red line.

Person Level Parameters

The computational growth modeling (CGM) approach generally estimated the person-level computational and growth parameters more accurately than the two-stage approach. Figure 3 shows distributions of the correlations between true and estimated person-level parameters across both approaches. The two-stage approach showed poor recovery of the person-level intercepts (β_0 ; all $r_s \leq .51$) whereas the CGM approach recovered these parameters well (all $r_s \geq .59$). Although the two-stage approach recovered the person-level slopes (β_1) worse than the CGM approach, correlations between true and estimated person-level slopes were all strong (all $r_s \geq .59$). Likewise, although the CGM approach consistently recovered the person-level learning rates (A_s) better than the two-stage approach, correlations between true and estimated parameters across both approaches were all strong (all $r_s > .73$). Table 3 shows MADs between true and estimated person level parameters and reflects a similar pattern of findings as shown in Figure 3 (i.e., less accuracy for β_0 using two-stage approach).

Table 3
Mean Absolute Deviations for Group-Level Parameters

		True vs. Estimate		Coverage	
		Two-Stage	CGM	Two-Stage	CGM
Intercept:	β_0	1.67 (0.87)	0.19 (0.14)	69.12	3.06
Slope:	β_1	0.74 (0.58)	0.17 (0.13)	48.10	2.15
Unbounded Learning Rate:	$'A$	0.84 (0.87)	0.25 (0.17)	7.13	0.19
Learning Rate:	A	0.12 (0.16)	0.19 (0.15)	2.04	0.27

Note. Mean absolute deviations (MADs; $\pm 1SD$) between true and estimated parameters and between the 95% target probability and the obtained coverage probability for the person-level parameters.


Figure 3*Correlation Between True & Estimated Person-Level Parameters*

Fig. 3 Distributions of the correlation between true person-level parameters and point estimates (posterior means) of the person-level parameters (columns) for each condition (rows). Vertical lines represent mean correlations across simulations, also represented with text in upper-left corner of each panel. The y-axis has been log10 transformed to better illustrate differences between estimation approaches, with 1.0 added to the frequencies to avoid a log-transformation on 0.

Finally, Figure 4 shows coverage probabilities for the person-level parameters. Here, the differences between the two-stage and CGM approaches are more readily apparent. The two-

stage approach systematically showed lower coverage than the target 95% for the person-level growth parameters (β_0 & β_1) relative to the CGM approach, suggesting that the two-stage approach frequently fails to capture the true starting points and rates-of-change in learning rates at the person level. For the session-specific learning rates (A_s), coverage probabilities across both approaches were close to the target 95%. Table 3 supports this conclusion, showing that MADs for the coverage probabilities were higher for the person-level growth parameters from the two-stage approach than the CGM approach. Similarly, MADs for the coverage probabilities were higher for the learning rates from the two-stage approach than for the CGM approach.

Figure 4
95% Coverage for Person-Level Parameters




Fig. 4 95% coverage for person-level parameters from each condition across the two-stage and CGM approaches. Rows correspond to different levels of the correlation between intercepts and slopes (ρ). Columns correspond to different levels of the slope effect (γ_1). The target 95% probability is indicated by the dashed red line.

Discussion

Our results illustrate that, generally, parameter recovery was better using the computational growth modeling (CGM) approach compared to the two-stage approach. Comparisons between true and estimated parameters show that the CGM approach outperformed the two-stage approach in estimating the growth parameters at both the group and person levels (Figures 1 & 3; Table 2). Likewise, posterior distribution coverage for group- and person-level growth parameters was closer to the target 95% for the CGM approach than for the two-stage approach (Figures 2 & 4; Table 3). Although a 95% target for coverage is arbitrary, we used this value as it relates to the convention of using 95% confidence intervals with frequentist statistics. Other intervals would be expected to show similar results (cf. percentile plots in the Online Resource).

For the learning rates, the CGM approach tended to outperform the two-stage approach; however, this difference in performance was smaller in comparison with the growth parameters. For example, although the CGM approach consistently showed stronger correlations between true and estimated parameters, both approaches showed very strong correlations between true and estimated learning rates ($r > .7$; see Figure 2). Similarly, coverage for the learning rates was comparable between both approaches (see Figure 4). Thus, the CGM and two-stage approaches work well in estimating parameters at individual timepoints; however, the CGM approach is particularly advantageous for estimating change in those parameters across time.

Real Data Application of Computational Growth Modeling

Here, we extend a model designed for the play-or-pass Iowa Gambling Task to a computational growth modeling (CGM) framework. The play-or-pass IGT is a decision-making task in which participants are presented with decks of cards across a series of trials (usually 120

trials) and they can choose to either “play” or “pass” on the presented deck (Bechara et al., 1994; Cauffman et al., 2010; IGT; Peters & Slovic, 2000). When participants play, they receive either a monetary gain, a monetary loss, or nothing. When participants pass, they receive nothing and move onto the next trial. There are four decks presented one at a time: two decks are advantageous, meaning that playing on them consistently results in net gains, and the other two decks are disadvantageous, meaning that playing on them consistently results in net losses (see Online Resource for further details of the task). The IGT is used to measure how participants learn which options (i.e., decks) are advantageous or disadvantageous in the long-run, and reinforcement learning models are commonly used models for characterizing patterns of behavior on the task.

We used the play-or-pass Outcome Representation Learning model (PP-ORL) to characterize performance on the play-or-pass IGT (Haynes et al., 2024). The PP-ORL consists of four parameters: a reward learning rate (A_{rew}), describing updating in response to better than expected positive outcomes (i.e., gains); a punishment learning rate (A_{pun}), describing updating in response to worse than expected negative outcomes (i.e., losses); win/loss frequency sensitivity (βf), describing bias towards decks resulting in more frequent gains or away from decks resulting in more frequent losses, regardless of magnitude; and response bias (βb), describing a general tendency towards playing (when positive) or passing (when negative), regardless of the outcomes associated with the decks. The four parameters inform the value of playing on each deck across trials within the play-or-pass IGT (cf. Equations 6 & 7; for details on the trial-by-trial algorithm, see Haynes et al., 2024).

To model longitudinal change in the PP-ORL parameters, we used the computational growth model scheme from Equation 5 to estimate a longitudinal growth model for each parameter of the PP-ORL:

$$\mathbf{Y}|\mathbf{C}, \mathbf{A}_{rew}, \mathbf{A}_{pun}, \boldsymbol{\beta f}, \boldsymbol{\beta b} \sim \text{PPORL model} \quad \text{Equation 11}$$

$$\mathbf{A}_{rew}|\mathbf{X}_0, \mathbf{X}_1, \mathbf{T}, \boldsymbol{\Psi}_{\mathbf{A}_{rew}} \sim \text{growth model on reward learning rate}$$

$$\mathbf{A}_{pun}|\mathbf{X}_0, \mathbf{X}_1, \mathbf{T}, \boldsymbol{\Psi}_{\mathbf{A}_{pun}} \sim \text{growth model on punishment learning rate}$$

$$\boldsymbol{\beta f}|\mathbf{X}_0, \mathbf{X}_1, \mathbf{T}, \boldsymbol{\Psi}_{\boldsymbol{\beta f}} \sim \text{growth model on win/loss frequency sensitivity}$$

$$\boldsymbol{\beta b}|\mathbf{X}_0, \mathbf{X}_1, \mathbf{T}, \boldsymbol{\Psi}_{\boldsymbol{\beta b}} \sim \text{growth model on go bias}$$

$$\boldsymbol{\Psi} = \gamma_0, \gamma_1, \sigma_1^2, \sigma_{20}^2, \sigma_{21}^2, \rho$$

where the ORL parameters are defined by the covariate ($\mathbf{X}_0$ & $\mathbf{X}_1$) and time ($\mathbf{T}$) predictors, and the level 1 and level 2 residual variation and growth parameters ($\boldsymbol{\Psi}$) from Equation 4 for each ORL parameter, with $\mathbf{A}_{rew}$ and $\mathbf{A}_{pun}$ estimated on the unbounded scale ($A = \Phi(\cdot A)$) as in the simulation. After building the CGM version of the PP-ORL model, we fit the model to a dataset consisting of five waves of play-or-pass IGT data from adults collected across 3 years as part of a larger longitudinal study (see Online Resource for details of the participants).

Covariate Specification

For illustration, we fit the CGM variant of the PP-ORL with two different covariate specifications. For the first model, we did not include any covariates by specifying the design matrices as follows:

$$\mathbf{X}_0 = \mathbf{X}_1 = \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} \quad \text{Equation 12}$$

Equation 12 assumes all person-level growth parameters follow the same group-level distribution as in a standard unconditional growth model. This simple model is useful to characterize growth

trajectories for all PP-ORL parameters across the full sample, which can serve as a baseline to reference in comparison to a more complex model with covariates of interest. We fit the second model assuming growth trajectories varied across people according to history of diagnosed anxiety, depression, or substance use disorder and according to sex:

$$\mathbf{X}_0 = \mathbf{X}_1 = \begin{bmatrix} 1 & ANX_1 & DEP_1 & SUD_1 & SEX_1 \\ 1 & ANX_2 & DEP_2 & SUD_2 & SEX_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & ANX_i & DEP_i & SUD_i & SEX_i \end{bmatrix} \quad \text{Equation 13}$$

where all diagnostic covariates were indicator variables with 0 indicating a negative history and 1 indicating a positive history of the given diagnosis. For the sex covariate, we coded 0 to indicate female and 1 to indicate male participants. Unlike the unconditional growth model, this *conditional* growth model allowed us to examine associations between covariates and person-level growth trajectories.

Data Analysis

As with the simulation, we fit the two specifications of the PP-ORL using hierarchical Bayesian estimation implemented in Stan (v. 2.32.7; Stan Development Team, 2023b), interfaced with R (v. 4.3.3; R Core Team, 2022) via RStan (v. 2.32.7; Stan Development Team, 2023a). We sampled the models with 1000 warmup iterations and 2000 target iterations across 8 chains (thus, 24000 iterations/chain). For both the unconditional (i.e., baseline) and conditional (i.e., covariate) models, we present both the level 2 growth model coefficients (i.e. γ_0 & γ_1) in addition to the implied level 2 growth trajectories (a.k.a. the group-level growth model predictions). We report posterior means $\pm 95\%$ highest density intervals (HDIs), unless otherwise stated.

Results

Figure 5 shows trial-level choice data (datapoints) for each deck on the play-or-pass Iowa Gambling Task (IGT) across the five testing sessions. As is typical, participants start the task by playing on each deck about equally often and then transition to playing more frequently on decks C and D, the advantageous decks, and less frequently on decks A and B, the disadvantageous decks (Case & Olino, 2020; Haynes et al., 2024; Peters & Slovic, 2000). Figure 5 also shows model predictions from the unconditional growth version of the ORL model. In general, the model captured broad trends in the trial-level data, although misfit is apparent. For example, the model overestimated how frequently participants chose deck B and underestimated how frequently participants chose deck C, particularly in the latter sessions. This may be due to the decreasing number of participants across sessions, making estimation more difficult in the latter sessions. Despite this, the model characterizes the consistent trend in which participants play more frequently on the advantageous decks relative to the disadvantageous decks and that choice of the disadvantageous decks systematically decreases across trials within the task after experiencing the losses on those decks.

Figure 5
Observed & Model-Predicted Play Proportions




Fig. 5 Observed (datapoints) and model-predicted (datapaths) play proportions from the play-or-pass Iowa Gambling Task (IGT). Model-predicted values represent posterior predictions, averaged across participants and MCMC sample, derived from the unconditional model. Data have been smoothed over 5-trial bins to better illustrate trends.

Figure 6 shows credible group level trends from the unconditional model for all PP-ORL parameters across sessions. Overall, both learning rates (A_{rew} & A_{pun}) decay across sessions, whereas both win/loss frequency sensitivity (βf) and response bias (βb) increase across sessions. The decrease in the reward learning rates across sessions, $\gamma_1 = -0.16 [-0.20, -0.11]$, captures participants shifting away from playing on the disadvantageous decks (A & B) which consist of the largest trial-level wins. The increase the win/loss frequency sensitivity, $\gamma_1 = 1.04 [0.72, 1.35]$, captures participants shifting away from deck A as this deck is associated with the most frequent losses (15 losses) compared to the other decks (3-6 losses). We also observed increases in response bias across sessions, $\gamma_1 = 0.09 [0.03, 0.15]$, which would suggest that participants played more frequently regardless of the deck across sessions; however, this is not readily apparent in Figure 5. We have previously speculated that the response bias parameter characterizes exploratory behavior at the beginning of the task which would allow participants to identify which decks are better and worse and behave accordingly (Haynes et al., 2025), but this has yet to be empirically tested. Finally, we observed decreases in punishment learning rates across sessions, $\gamma_1 = -0.06 [-0.10, -0.02]$. This finding was surprising given that elevated punishment learning rates typically capture avoidance of deck B's large trial-level losses (i.e., a \$1150 loss); thus, we would expect to see punishment learning rates increase over time if participants were avoiding that deck as they appear to in Figure 5.

Figure 6

Level 2 Growth Trajectories from the Unconditional Model




Fig. 6 Posterior predictive growth trajectories for the level 2 component of the unconditional model. Black datapaths show posterior predictive means with errorbands representing highest density intervals (HDIs). We generated these posterior predictions by simulating forward Equation 11 after fitting the unconditional model, including the level 1 error (E_1) and then averaging over the simulated draws (see Online Resource for details).

Figures 7 and 8 show the posterior means of the group-level growth parameters ($\pm 95\%$ HDIs) and the trajectories of the ORL parameters across time from the conditional model. Focusing first on the intercept term, male participants show both a credibly higher baseline punishment learning rate compared to females. Additionally, people with a lifetime history of depression tend to show a higher baseline win/loss frequency sensitivity compared to those without such a history. For the slope term, there was a trend in which a lifetime history of

depression was associated with decreasing punishment learning rates across sessions. This steeper reduction is readily apparent in Figure 8 where those with depression show similar punishment learning rates at session 1, and by session 5, show the lowest learning rates across the different groups. Thus, the conditional growth model shows meaningful group-level differences in IGT performance between male and female participants and between those with and without depression.

Figure 7

Level 2 Growth Model Coefficients from the Conditional Model




Fig. 7 Posterior distributions for the level 2 growth model terms of the conditional model (a.k.a. fixed effects). For learning rates, the parameter values are still on the untransformed scale (i.e. before the Φ transform). Thick and thin black intervals indicate the 80% and 95% HDIs, respectively. Black datapoints indicate posterior means.

Figure 8
Level 2 Growth Trajectories from the Conditional Model




Fig. 8 Posterior predictive growth trajectories for the level 2 component of the conditional model. Reference group refers to participants without a history of anxiety, depression, or substance abuse disorder. We only show the 50% HDI for interpretability purposes, as using wider intervals makes it difficult to visualize individual trajectories and associated uncertainty. We generated these posterior predictions by simulating forward as in Figure 6.

Discussion

Here, we showed that the computational growth modeling (CGM) approach can be applied to real-world data involving three or more timepoints. Such models can reveal important patterns of decision-making across time. For example, results from the unconditional growth model showed that, among adults completing the play-or-pass Iowa Gambling Task (IGT), reward and punishment learning rates decreased and win/loss frequency sensitivity and response

bias increased across repeated assessments. Our previous work assessing IGT performance across two timepoints suggest that these patterns of change should be associated with better performance, such that decreases in reward learning rate, coupled with increases in bias, are associated with better performance on the play-or-pass IGT from session 1 to session 2 (Haynes et al., 2024). In the present study, participants' net "earnings" increase across sessions supporting this finding; however, the decrease in punishment learning rates and increase in win/loss frequency sensitivity were surprising as this was not detected in our prior study with two timepoints. This could indicate that changes in punishment learning rates and win/loss frequency sensitivity may be associated with longer-term changes in decision-making on the task that can only be captured by more than two timepoints worth of data.

Importantly, however, the CGM approach allows for seamless integration of covariates within the model, permitting testing whether trajectories are conditional on socially relevant individual differences (i.e., history of psychopathology). The conditional growth version of the PP-ORL model showed that depression was associated with higher baseline win/loss frequency sensitivity and sharper decreases in punishment learning rates across time (conditional on a history of anxiety & SUD diagnosis, & sex). This could indicate that individuals with depression show sharper decreases in punishment learning beyond what we would expect from typical age-related decreases (Lerner et al., 2018).

We also found that male participants showed elevated punishment learning rates compared to female participants (conditional on a history of anxiety, depression, & SUD diagnosis). This finding replicates our previous results with this sample (see Haynes et al., 2025). The reasons for sex-differences on the IGT are not fully understood; however, previous reviews suggest that women/female participants may be more sensitive to losses that occur more

frequently whereas men/male participants may be more sensitive to losses of larger magnitude (Grissom & Reyes, 2019; van den Bos et al., 2013) which could be related to the reduced punishment learning rates in female participants that we observed.

General Discussion

Assessing longitudinal change of many behavioral processes frequently relies on growth models, under an assortment of labels. Typically, these models are fit to summary measures of behavior (e.g., choice proportions), but for many behavioral tasks, there are behavioral models that can characterize more complex patterns of behavior on those tasks via computational parameters (e.g., learning rates, win-frequency sensitivities). Assessing changes in these parameters across time, however, requires integration of computational models with growth models. Here, we illustrate how to integrate these models using a joint modeling framework that allows for simultaneous estimation of parameters at the person level, for the computational model, and the group level, for the growth model — what we call computational growth modeling (CGM). We demonstrated how computational growth modeling could be applied in two contexts: first, with a simulation to compare the CGM approach with the more traditional two-stage approach and second, with real data from a longitudinal study on reward and punishment learning in adults. We discuss the results and broader implications of these applications below.

First, the simulation-based results showed that although the two-stage and CGM approaches estimated computational parameters (i.e., learning rates) well cross-sectionally (i.e., within each session individually), the CGM estimated the person-level changes across time in those computational parameters better. The advantage of the CGM approach, over the two-stage approach, was a result of the differences in the availability of information across levels of

analysis (i.e., the computational model & the growth model). For the CGM, learning rates and growth parameters were estimated simultaneously, allowing uncertainty in the learning rates to inform estimation of the growth parameters (Brown et al., 2020; Haines et al., 2025; Matzke et al., 2017). This resulted in better performance of the CGM approach in recovering the growth parameters. In contrast, the two-stage approach relied on estimating growth parameters from point estimates (i.e., posterior means) of the learning rates, thus ignoring the uncertainty of those learning rates. By ignoring this uncertainty, the two-stage approach resulted in posterior distributions of growth parameters that systematically failed to capture the true parameter values (cf. poor coverage of the growth parameters in Figures 2 & 4). This is a direct result of the two-stage approach estimating the growth parameters with too much certainty (i.e., without measurement error). Combined, these findings suggest that CGM may be a more robust approach to assess longitudinal change in computational parameters.

Second, the real data application illustrates how to implement the CGM approach to assess longitudinal changes in reward and punishment learning. By extending the play-or-pass Outcome Representation Learning (PP-ORL) model (Haines et al., 2018; Haynes et al., 2024) as a CGM, we found that both reward and punishment learning rates decreased, and win frequency sensitivity and response bias increased across five repeated administrations of the play-or-pass IGT over the course of three years. It is beyond the scope of this paper to speculate to what degree changes in these parameters reflect changes in cognitive processing (Lerner et al., 2018) or practice effects; however, the CGM provides a framework to explore both possibilities in future work using the PP-ORL or with extensions of other computational models (e.g., Ahn et al., 2008; Busemeyer & Stout, 2002).

In addition, the flexibility of the CGM-implemented PP-ORL allowed us to characterize how change across time in PP-ORL parameters differed as a function of multiple covariates. To do this, we used a conditional growth form of the CGM and included lifetime history of anxiety, depression, and substance use disorder, as well as sex as covariates in the model. We chose these covariates as they are the same as those used in our previous study examining baseline performance on the IGT using the PP-ORL model (Haynes et al., 2025). The conditional ORL model revealed that ORL parameters meaningfully vary across diagnoses and sex and suggest that change-over-time among the diagnoses may not be equal (cf. depression showing decreases in punishment learning rates over time). In general, these findings are consistent with our own findings examining IGT performance at the first timepoint in these data (Haynes et al., 2025)⁴, particularly the observed sex-differences which aligns with prior literature in this area (Grissom & Reyes, 2019; van den Bos et al., 2013; Zanini et al., 2024). The goal of fitting the conditional growth model was to illustrate the flexibility with which we could extend the CGM approach to various research questions. The specific findings presented here likely require replication in separate samples, particularly as data from the latter timepoints were collected during the Covid 19 pandemic and could have contributed to greater variability in IGT performance and thus less certainty in estimating some parameters. Importantly, the real-world example illustrates the utility of the CGM approach with a relatively simple set of variables.

Although the computational growth modeling approach shows promise in modeling data across time, a number of questions remain. First, joint modeling is a form of hierarchical

⁴ Our previous study found elevated punishment learning rates in anxiety at the first timepoint, and results from the CGM are consistent with that finding (i.e., higher A_{pun} intercept in anxiety). However, uncertainty in the effect of anxiety on punishment learning was greater using the CGM approach — the 95% HDIs for the A_{pun} anxiety intercept overlap with 0. This could indicate that the effect of anxiety on punishment learning rates is less reliable than we initially found (Haynes et al., 2025).

modeling that relies on shrinkage to pool person-level parameters toward the group-level mean (Haines et al., 2023, 2025; Matzke et al., 2017). This provides a number of benefits (e.g., improved reliability, increased power; Brown et al., 2020; Haynes et al., 2026); however, this comes at the cost of compressing between-subject variability. Specifically, by hierarchically estimating person-level parameters, those parameters are assumed to be more extreme representations of participants' "typical" behavior and thus biased by the individual's data alone. Shrinkage corrects for this bias by pooling person-level parameters closer to the group-level mean, thereby compressing between-subject variability. For behavior analysts, whose main focus is on understanding the behavior of individual organisms (Sidman, 1960), approaches like hierarchical modeling may appear contrary to the field's values. Studies illustrating that hierarchically-estimated person-level parameters outperform their non-hierarchically-estimated counterparts in predicting out-of-sample performance at the *individual* level could shed light on how this approach may be advantageous for behavior analysts. Second, from a practical perspective, joint modeling is a tedious and time-consuming approach to analyzing data. Specifically, computational modeling, and as such computational growth modeling, is bespoke by nature and thus typically requires researchers to program their own models depending on the behavioral task and the model one chooses to characterize behavior on a task. We fit our joint models using Stan, which relies on Bayesian estimation procedures. There are a number of texts available for illustrating how to build and implement models programmed in Stan (e.g., McElreath, 2018; Sorensen et al., 2016); however, other estimation approaches (e.g., maximum likelihood) may be used for joint modeling (for applications, see Ballard & McClure, 2019; Brown et al., 2020).

Further extensions of the CGM approach could be used to answer questions across many different types of data. For example, we only included time-invariant binary variables (e.g., history of psychopathology) as covariates within the model, but continuous (e.g., dimensional assessments of psychopathology) and/or time-varying covariates (e.g., current vs. remitted status) could be included in the model, providing further nuance in understanding state influences on computational parameters. Moreover, given enough power, the model could be used to examine interactions between variables. Similarly, time could be used flexibly in the model to examine curvilinear or non-linear changes in growth based on theories of change (Casey et al., 2011; Singer & Willett, 2003). Furthermore, although our lab has primarily focused on joint modeling of human subjects data, the approach could be equally applied to animal data. This may be particularly advantageous for the animal researcher as studies with animals typically measure behavior repeatedly across multiple sessions. This would naturally lend itself to computational growth modeling, particularly if the goal is to characterize how performance on a task changes over time based on parameters derived from a theoretically-informed behavioral model. Thus, a number of applications are possible with the CGM approach that would allow researchers to characterize behavior across time.

In sum, we present a framework to understand the longitudinal trajectories of computationally derived parameters using joint modeling. We illustrated the framework as applied to two reinforcement learning models: a simple 1-parameter model fit to simulated data and a more complex 4-parameter model fit to real data. Our hope is that this approach, along with further extensions, can advance our modeling, and understanding, of how behavioral processes change over time.

References

- Ahn, W.-Y., Busemeyer, J. R., Wagenmakers, E.-J., & Stout, J. C. (2008). Comparison of decision learning models using the generalization criterion method. *Cognitive Science*, 32(8), 1376–1402. <https://doi.org/10.1080/03640210802352992>
- Ballard, I. C., & McClure, S. M. (2019). Joint modeling of reaction times and choice improves parameter identifiability in reinforcement learning models. *Journal of Neuroscience Methods*, 317, 37–44.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, 50(1), 7–15. [10.1016/0010-0277\(94\)90018-3](https://doi.org/10.1016/0010-0277(94)90018-3)
- Brown, V. M., Chen, J., Gillan, C. M., & Price, R. B. (2020). Improving the Reliability of Computational Analyses: Model-Based Planning and Its Relationship with Compulsivity. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 5(6), 601–609. <https://doi.org/10.1016/j.bpsc.2019.12.019>
- Bryk, A. S., & Raudenbush, S. W. (1992). *Hierarchical linear models: Applications and data analysis methods*. Sage Publications Newbury Park, CA.
- Buelow, M. T., & Barnhart, W. R. (2018). Test–Retest Reliability of Common Behavioral Decision Making Tasks. *Archives of Clinical Neuropsychology*, 33(1), 125–129. <https://doi.org/10.1093/arclin/acx038>

- Bürkner, P. C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80, 1–28. <https://doi.org/10.18637/jss.v080.i01>
- Busemeyer, J. R., & Stout, J. C. (2002). A contribution of cognitive decision models to clinical assessment: Decomposing performance on the Bechara gambling task. *Psychological Assessment*, 14(3), 253. <https://doi.org/10.1037/1040-3590.14.3.253>
- Casabianca, J. M., Junker, B. W., Nieto, R., & Bond, M. A. (2017). A hierarchical rater model for longitudinal data. *Multivariate Behavioral Research*, 52(5), 576–592. <https://doi.org/10.1080/00273171.2017.1342202>
- Case, J. A., & Olino, T. M. (2020). Approach and avoidance patterns in reward learning across domains: An initial examination of the Social Iowa Gambling Task. *Behaviour Research and Therapy*, 125, 103547. <https://doi.org/10.1016/j.brat.2019.103547>
- Casey, B. J., Jones, R. M., & Somerville, L. H. (2011). Braking and accelerating of the adolescent brain. *Journal of Research on Adolescence*, 21(1), 21–33. <https://doi.org/10.1111/j.1532-7795.2010.00712.x>
- Cauffman, E., Shulman, E. P., Steinberg, L., Claus, E., Banich, M. T., Graham, S., & Woolard, J. (2010). Age differences in affective decision making as indexed by performance on the Iowa Gambling Task. *Developmental Psychology*, 46(1), 193.
- Cicchetti, D., & Toth, S. L. (2009). The past achievements and future promises of developmental psychopathology: The coming of age of a discipline. *Journal of Child Psychology and Psychiatry*, 50(1-2), 16–25. <https://doi.org/10.1111/j.1469-7610.2008.01979.x>
- Curran, P. J. (2003). Have multilevel models been structural equation models all along? *Multivariate Behavioral Research*, 38(4), 529–569. https://doi.org/10.1207/s15327906mbr3804_5

- Doll, B. B., Simon, D. A., & Daw, N. D. (2012). The ubiquity of model-based reinforcement learning. *Current Opinion in Neurobiology*, 22(6), 1075–1081.
<https://doi.org/10.1016/j.conb.2012.08.003>
- Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4), 457–472. <https://doi.org/10.1214/ss/1177011136>
- Grissom, N. M., & Reyes, T. M. (2019). Let's call the whole thing off: Evaluating gender and sex differences in executive function. *Neuropsychopharmacology*, 44(1), 86–96.
<https://doi.org/10.1038/s41386-018-0179-5>
- Haines, N., Kvam, P. D., Irving, L., Smith, C. T., Beauchaine, T. P., Pitt, M. A., Ahn, W.-Y., & Turner, B. M. (2025). A tutorial on using generative models to advance psychological science: Lessons from the reliability paradox. *Psychological Methods*.
<https://doi.org/10.1037/met0000674>
- Haines, N., Sullivan-Toole, H., & Olino, T. (2023). From classical methods to generative models: Tackling the unreliability of neuroscientific measures in mental health research. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 8(8), 822–831.
<https://doi.org/10.1016/j.bpsc.2023.01.001>
- Haines, N., Vassileva, J., & Ahn, W.-Y. (2018). The outcome-representation learning model: A novel reinforcement learning model of the Iowa Gambling Task. *Cognitive Science*, 42(8), 2534–2561. <https://doi.org/10.1111/cogs.12688>
- Haynes, J. M., Haines, N., Sullivan-Toole, H., & Olino, T. M. (2024). Test-retest reliability of the play-or-pass version of the Iowa Gambling Task. *Cognitive, Affective, & Behavioral Neuroscience*, 24(4), 740–754. <https://doi.org/10.3758/s13415-024-01197-6>

- Haynes, J. M., Kvam, P., & Olino, T. (2026). Associations Between Steep Delay Discounting and Punishment Learning. *Computational Psychiatry*, 10(1), 159–183. <https://doi.org/10.5334/cpsy.172>
- Haynes, J. M., Sullivan-Toole, H., Haines, N., & Olino, T. (2025). Reward and Punishment Learning Among People with a Lifetime History of Anxiety, Depression, and Substance Use Disorder. *Cognitive, Affective, & Behavioral Neuroscience*. <https://doi.org/10.3758/s13415-025-01331-y>
- Insel, T. R. (2009). Translating scientific opportunity into public health impact: A strategic plan for research on mental illness. *Archives of General Psychiatry*, 66(2), 128–133. <https://doi.org/10.1001/archgenpsychiatry.2008.540>
- Kaplan, B. A., Franck, C. T., McKee, K., Gilroy, S. P., & Koffarnus, M. N. (2021). Applying mixed-effects modeling to behavioral economic demand: An introduction. *Perspectives on Behavior Science*, 44(2), 333–358. <https://doi.org/10.1007/s40614-021-00299-7>
- Klein, S. D., Collins, P. F., & Luciana, M. (2022). Developmental trajectories of delay discounting from childhood to young adulthood: Longitudinal associations and test-retest reliability. *Cognitive Psychology*, 139, 101518. <https://doi.org/10.1016/j.cogpsych.2022.101518>
- Kvam, P. D., Romeu, R. J., Turner, B. M., Vassileva, J., & Busemeyer, J. R. (2021). Testing the factor structure underlying behavior using joint cognitive models: Impulsivity in delay discounting and Cambridge gambling tasks. *Psychological Methods*, 26(1), 18. <https://doi.org/10.1037/met0000264>
- Lerner, I., Sojitra, R., & Gluck, M. (2018). How age affects reinforcement learning. *Aging (Albany NY)*, 10(12), 3630–3631. <https://doi.org/10.18632/aging.101649>

- Matzke, D., Ly, A., Selker, R., Weeda, W. D., Scheibehenne, B., Lee, M. D., & Wagenmakers, E.-J. (2017). Bayesian inference for correlations in the presence of measurement error and estimation uncertainty. *Collabra: Psychology*, 3(1), 25.
<https://doi.org/10.1525/collabra.78>
- Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In *Quantitative analyses of behavior: Vol. 5. The effect of delay and of intervening events on reinforcement value* (pp. 55–73). Erlbaum.
- McElreath, R. (2018). *Statistical rethinking: A Bayesian course with examples in R and Stan*. Chapman and Hall/CRC.
- Oehlert, G. W. (2010). *A First Course in Design and Analysis of Experiments*. University of Minnesota Press. <https://doi.org/10.4324/9780203885659>
- Peters, E., & Slovic, P. (2000). The springs of action: Affective and analytical information processing in choice. *Personality and Social Psychology Bulletin*, 26(12), 1465–1475.
<https://doi.org/10.1177/01461672002612002>
- Ployhart, R. E., & MacKenzie, W. I. (2014). Two waves of measurement do not a longitudinal study make. In *More statistical and methodological myths and urban legends* (pp. 85–99). Routledge.
- R Core Team. (2025). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org>
- Rescorla, R., & Wagner, A. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical Conditioning II: Current Research and Theory: Vol. 2*.

Schmitz, F., Kunina-Habenicht, O., Hildebrandt, A., Oberauer, K., & Wilhelm, O. (2020).

Psychometrics of the Iowa and Berlin Gambling Tasks: Unresolved Issues With

Reliability and Validity for Risk Taking. *Assessment*, 27(2), 232–245.

<https://doi.org/10.1177/1073191117750470>

Sidman, M. (1960). *Tactics of scientific research: Evaluating experimental data in psychology*.

Basic Books.

Singer, J. D., & Willett, J. B. (2003). *Applied longitudinal data analysis: Modeling change and event occurrence*. Oxford university press.

Sorensen, T., Hohenstein, S., & Vasishth, S. (2016). Bayesian linear mixed models using Stan: A

tutorial for psychologists, linguists, and cognitive scientists. *The Quantitative Methods*

for Psychology, 12(3), 175–200. <https://doi.org/10.20982/tqmp.12.3.p175>

Stan Development Team (2025a). RStan: The R interface to Stan. *R Package Version 2.26.22*.

<http://mcstan.org>

Stan Development Team (2025b). Stan modeling language users guide and reference manual.

Technical Report Version 2.32. <http://mcstan.org>

Sullivan-Toole, H., Haines, N., Dale, K., & Olino, T. M. (2022). Enhancing the psychometric

properties of the iowa gambling task using full generative modeling. *Computational*

Psychiatry, 6(1), 189–212. <https://doi.org/10.5334%2Fcpsy.89>

Sullivan-Toole, H., Haynes, J. M., Schmidt, H., Larsen, B., Haines, N., & Olino, T. M. (2026).

Reward-specific learning parameters change across normative adolescent development

and are blunted in youth with high risk for depression. *Journal of Child Psychology and*

Psychiatry, 67(6), 963–975. <https://doi.org/10.1111/jcpp.70086>

- Talts, S., Betancourt, M., Simpson, D., Vehtari, A., & Gelman, A. (2018). Validating Bayesian inference algorithms with simulation-based calibration. *arXiv Preprint arXiv:1804.06788*.
- van den Bos, R., Homberg, J., & de Visser, L. (2013). A critical review of sex differences in decision-making tasks: Focus on the Iowa Gambling Task. *Behavioural Brain Research*, 238, 95–108. <https://doi.org/10.1016/j.bbr.2012.10.002>
- Vandekerckhove, J., & Tuerlinckx, F. (2007). Fitting the Ratcliff diffusion model to experimental data. *Psychonomic Bulletin & Review*, 14(6), 1011–1026. <https://doi.org/10.3758/BF03193087>
- Wilson, R. C., & Collins, A. G. (2019). Ten simple rules for the computational modeling of behavioral data. *Elife*, 8, e49547. <https://doi.org/10.7554/eLife.49547>
- Young, M. E. (2018). Discounting: A practical guide to multilevel analysis of choice data. *Journal of the Experimental Analysis of Behavior*, 109(2), 293–312. <https://doi.org/10.1002/jeab.316>
- Zanini, L., Picano, C., & Spitoni, G. F. (2024). The Iowa Gambling Task: Men and women perform differently. A meta-analysis. *Neuropsychology Review*, 1–21. <https://doi.org/10.1007/s11065-024-09637-3>

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SupplementComputationalgrowthmodeling.docx](#)