Online Supplement to
Learning to be Proficient? A Structural Model of User Dynamic Engagement in eHealth Behavioral Interventions

Tongxin Zhou
W. P. Carey School of Business, Arizona State University, tongxin.zhou@asu.edu

Yingfei Wang
Michael G. Foster School of Business, University of Washington, yingfei@uw.edu

Lu (Lucy) Yan
Kelley School of Business, Indiana University, yanlucy@iu.edu

Yong Tan
Michael G. Foster School of Business, University of Washington, ytan@uw.edu

A1. Reduced-Form Analysis

We conduct a reduced-form analysis to test whether users' previous weight-loss experience has any impact on their participation decisions. If there is such an impact, it would indicate that individuals can potentially adjust their behaviors based on their past weight-loss outcomes, providing supporting evidence for learning behaviors. Specifically, we consider the following regression:

$$Y_{it} = \eta_1 \cdot \text{intercept} + \eta_2 \cdot \text{gender}_i + \eta_3 \cdot \text{age}_i + \eta_4 \cdot \text{initial\_weight}_i + \eta_5 \cdot \text{tenure}_{it} + \eta_6 \cdot \text{weight\_change}_{it}. \tag{A1}$$

The dependent variable represents an individual's participation decision at a specific week. It is measured in two ways: (1) whether a user opts for any interventions (a binary measure) and (2) the number of interventions chosen (a discrete count). The first measure reflects the individual's propensity to participate, while the second captures the extent of their participation.

The independent variable $\text{weight\_change}_{it}$ represents the user's weight-loss experiences prior to the given time. We measure it by whether the user has experienced weight loss or weight maintenance in the most recent period. A learning effect is implied if this variable positively influences the dependent variable, as it could suggest that users who have experienced weight loss or weight maintenance are more likely to form positive perceptions of the interventions and participate more actively. We also account for the influence of personal characteristics such as gender, age, initial weight, and duration of participation.

The estimation results are presented in Table A1. Model 1 presents the results for the logit model in which the dependent variable is binary, and Model 2 shows the results for the Poisson regression model in which the dependent variable is the number of interventions chosen. The results from both models suggest a positive relationship between $\text{weight\_change}_{it}$ and users' participation decisions. These findings suggest that users are adapting their engagement based on the outcomes they experience, which indicates a feedback loop where outcomes inform future actions. Such a pattern provides initial evidence for learning, where individuals modify their behavior in response to the consequences of their previous actions. We also examined the variety of interventions chosen by individuals over time, under the hypothesis that if users are learning, they will reduce the diversity of their selections. Our results support this conjecture. Specifically, we calculated the average variance in individuals' intervention choices, represented as one-hot encoded vectors, across different time points. The analysis revealed a clear trend of decreasing choice variability, indicating that as users progress in their understanding of intervention effectiveness, they are likely to focus on a narrower set of interventions. This reduction in variety provides further evidence of learning behaviors, as users refine their decision-making and increasingly select interventions they perceive as more effective.

Table A1 Regression of Participation Decision

	Model 1 (Logit) Binary		Model 2 (Poisson) Counts	
gender	−0.0055	(0.0380)	0.0013	(0.0114)
age	−0.0215	(0.0166)	−0.0210***	(0.0051)
initial_weight	−0.0683*	(0.0386)	−0.0276**	(0.0115)
tenure	−0.0579*	(0.0342)	−0.0166	(0.0104)
weight_change	0.0452***	(0.0159)	0.0155***	(0.0048)
intercept	1.5635***	(0.1166)	0.8549***	(0.0346)

Note: Standard errors in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Number of observations: 15,735.

A2. Annotation of Intervention Categories

We classify interventions based on their titles and descriptions. Following the framework in Section 2.1.2, we determine an intervention to be *BehaviorSpecific* if it targets a concrete, repeatable health behavior (e.g., dietary change, exercise). We determine an intervention to be *OutcomeOriented* if it emphasizes achieving a measurable weight-loss outcome. An intervention is classified as *SelfRegulation* if it promotes reflective, self-monitoring, or motivational practices. Two researchers independently coded the interventions. Coding disagreements were discussed and resolved by consensus. The annotated set comprised 102 *BehaviorSpecific* interventions, 37 *OutcomeOriented* interventions, and 26 *SelfRegulation* interventions. Table A2 presents several illustrative examples of the annotated interventions.

Table A2 Annotated Intervention Categories

Intervention Title	Description (Excerpt)	Primary Target Identified	Meta-Level Category
4.5 kg for the Season	4.5 kgs down by the first day of fall, September 22. No pressure, no rules, no judgement; just healthy choices!	Outcome goal	Outcome-oriented
5 kg Challenge	Let's see who can get to the 5kg loss in 4 weeks.	Outcome goal	Outcome-oriented
2 miles a day on treadmill	Walk 2 miles on treadmill every day, that's it nothing else.	Physical activity	Behavior-specific
1500-Calorie Daily Challenge	Track food intake and stay under 1,500 kcal/day.	Calorie intake	Behavior-specific
Low Carb - High (Natural) Fat LifeStyle!	Cut sugar/carbs! Eat natural fats, protein, and fiber.	Nutrition	Behavior-specific
The First Step - Starting Out	Be more mindful of what you put into your body. What do you eat and drink? Try to make a conscious effort to be more active.	Self-monitoring & reflection	Self-regulation-focused
The Power of One	Decrease your caloric intake. Not by limiting your food choices, but by limiting your serving to one.	Calorie intake	Behavior-specific
30 Day Buns, Guns & Abs Exercise Challenge	30 Day BUNS GUNS ABS Exercise Challenge. Do a set number of pushups, squats, and leg raises daily for 30 days.	Physical activity	Behavior-specific
Mindful Eating Reflection	Daily journaling about eating habits.	Self-monitoring & reflection	Self-regulation-focused

A3. Additional Evidence Supporting the Belief-Updating Framework

To further evaluate the empirical plausibility of local belief updating, we conduct additional analyses of intervention participation behavior. Notably, the purpose of this analysis is not to prove that users never engage in forward-looking behavior, but rather to examine whether longer-horizon planning variables provide substantial incremental explanatory power beyond short-horizon local updating variables in the focal context. The dependent variable equals one if user i participates in an intervention category c in week t, and zero otherwise.

Specifically, we estimated three nested logit specifications at the user-week-category level:

- **Model 1 (Short-Horizon Model)** includes only short-horizon local updating variables, including recent participation behavior, recent realized outcomes, and intervention popularity. This specification captures the idea that users rely primarily on recent experiences and current feedback signals when making participation decisions.

- **Model 2 (Long-Horizon Augmented Model)** additionally incorporates variables associated with forward-looking behavior, including cumulative weight-loss progress, accumulated intervention exposure, and user tenure. If users engage in strong forward-looking optimization, these longer-horizon variables should substantially improve explanatory power beyond short-term feedback signals.

- **Model 3 (Interaction Model)** further introduces interaction terms between recent outcome feedback and long-horizon variables. The rationale is that strongly forward-looking users should react less strongly to temporary short-term feedback fluctuations because such fluctuations are interpreted within a broader long-term optimization framework. Accordingly, users with greater cumulative progress or longer tenure should exhibit weaker responses to recent feedback shocks if forward-looking behavior is central.

Table A3 summarizes the results of the three models.

The results suggest that adding longer-horizon planning proxies yields only negligible improvements in explanatory power. Relative to Model 1, the pseudo-R^2 increases only from 0.1655 to 0.1656 after introducing long-horizon variables, while the interaction specification contributes no additional improvement. Although Model 2 generates a marginal improvement in log-likelihood, the magnitude of this improvement is very small relative to the additional model complexity,

Table A3 Additional Analysis Supporting the Belief-Updating Framework

Parameter	(1) Short-Horizon Model	(2) Long-Horizon Augmented Model	(3) Interaction Model
Recent weight loss	0.0198*	0.0220*	0.0203
	(0.0120)	(0.0134)	(0.0183)
Lagged participation	0.5914***	0.5844***	0.5842***
	(0.0236)	(0.0240)	(0.0240)
Intervention popularity	0.5748***	0.5733***	0.5733***
	(0.0095)	(0.0097)	(0.0097)
Cumulative weight loss	—	-0.0067	-0.0065
		(0.0157)	(0.0157)
Accumulated intervention exposure	—	0.0199	0.0199
		(0.0181)	(0.0182)
		(0.0293)	(0.0293)
High tenure	—	0.0630	0.0632
		(0.0594)	(0.0594)
Recent weight loss × cumulative weight loss	—	—	0.0097
			(0.0087)
Recent weight loss × high tenure	—	—	0.0038
			(0.0246)
Category fixed effects	Included	Included	Included
Week fixed effects	Included	Included	Included
User-level controls	Included	Included	Included
Log-likelihood	-26514.61	-26511.91	-26511.14
Pseudo-R^2	0.1655	0.1656	0.1656
AIC	53075.22	53077.81	53080.29
BIC	53276.76	53314.39	53334.39

Notes: Standard errors clustered at the user level are reported in parentheses. *, **, and *** indicate significance at the 0.1, 0.05, and 0.01 levels, respectively. The dependent variable equals one if the user participates in an intervention category during a given week. Accumulated intervention exposure measures the cumulative number of prior participation experiences within a category before the current decision period and serves as a proxy for longer-horizon investment, familiarity, and habit formation associated with that intervention category. All specifications include category fixed effects, week fixed effects, and baseline user-level controls such as gender, age, initial weight.

with AIC and BIC worsening. We also find that the interaction terms associated with forward-looking moderation are statistically insignificant. This suggests that users with greater cumulative progress or greater tenure do not respond differently to recent outcomes. These findings suggest that observed week-to-week participation behavior in this highly uncertain healthcare environment is more consistent with locally adaptive experiential learning under uncertainty than with fully forward-looking intertemporal optimization.

A4. Descriptive Intervention Features

We measure the length of each intervention by the total number of words contained in its title and description. As such, the length variable gauges the amount of information that users need to process when they participate in an intervention. We calibrate the sentiment polarity score for each intervention through VADER (Hutto and Gilbert 2014), which is a popular technique for sentiment analysis in social media. The calculated sentiment scores are values between 0 and 1. There are three dimensions: positive, negative, and neutral, and the scores for these dimensions sum to one. In our model, we include the positive and negative sentiment, as neutral would be derived from these two scores. In Table A4, we show two examples of the interventions that are detected as positive and negative. The first intervention contains more positive words (i.e., words associated with positive meanings, e.g., "start," "big," "success"), which explains why the positive sentiment score is higher. In comparison, the second intervention uses more negative words (i.e., words associated with negative meanings, e.g., "difficult," "problems," "not") and therefore has a higher negative sentiment score.

Table A4 Examples of Intervention Sentiment

Intervention Description	Positive	Negative
"We all probably want to lose more than just 5 lbs., but you got to start somewhere! Small steps will lead to big success."	0.223	0.057
"Life is difficult, always problems for different things, but eat just for eat, it's not the solution."	0.117	0.248

We further measure the intensity level of interventions for each intervention type. Consistent with (Zhou et al. 2023), the degree of intensity can be determined by the number of constraints imposed on calorie intake or expenditure. For example, if the intervention is about improving diet, we evaluate the extent to which individuals need to change their recipes or dietary habits. Minor diet-structure changes typically require simple adjustments, such as increasing vegetable intake. Moderate diet structure changes involve altering a main meal, such as having a salad for dinner. Intense diet structure changes require individuals to completely modify their food choices and recipes, such as consuming only high-protein, high-fiber foods. Similarly, the intensity for an exercise-related intervention is determined by required efforts in calorie consumption. According to metabolic equivalent of task (MET), walking and strolling are considered low-intensity, while bicycling, yoga, and brisk walking (over 3 mph) are moderate-intensity activities, and jogging, running, and jumping rope are high-intensity activities. For outcome-oriented interventions, we define the intensity based on the average weekly weight loss rate. According to CDC, a weight loss of 1-2 lbs. per week is achievable by burning 500-1000 calories more than consumed each day. Interventions with weight-loss targets within this range are considered moderate, while targets below 1 lb. per week are low-intensity and targets exceeding 2 lbs. per week are high-intensity. For self-regulation-focused interventions, intensity level is determined based on the type of behavior support provided. Those that aim to motivate users through mental therapies are considered low intensity, while interventions that require users to take certain actions such as meditation and/or regular weigh-in are classified as moderate intensity. Finally, interventions that require significant effort such as journaling of detailed movements and calorie intake are classified as high intensity.

A5. Estimation Procedure and Identification

Table A5 provides a summary of the parameter definitions. The parameter vector consists of user-specific intervention-effectiveness parameters, user-specific prior beliefs, user-specific category fixed effects, and a set of common parameters. Let

$$\Theta = \left\{ \{\boldsymbol{\beta}_i, \boldsymbol{\theta}_{0i}, \boldsymbol{\zeta}_i\}_{i=1}^N, \mathbf{B}, \boldsymbol{\Sigma}_\beta, \boldsymbol{\mu}_{\theta_0}, \boldsymbol{\Sigma}_\mu, \boldsymbol{\mu}_\zeta, \boldsymbol{\Sigma}_\zeta, \boldsymbol{\phi} \right\},$$

where $\boldsymbol{\beta}_i$ denotes user i's true intervention effectiveness across categories, $\boldsymbol{\theta}_{0i}$ denotes user i's initial beliefs, $\boldsymbol{\zeta}_i$ denotes user-specific category fixed effects in the mixed logit utility, and $\boldsymbol{\phi}$ collects the common utility and variance parameters. In the current specification, the common parameter vector is

$$\boldsymbol{\phi} = \left(\log \sigma_{y1}^2, \ldots, \log \sigma_{yC}^2, \log \alpha, \boldsymbol{\gamma}^\top, \boldsymbol{\lambda}^\top, \rho, \log \sigma_w^2\right)^\top,$$

where $C = 3$ which is the total number of intervention categories in our context.

The individual-level likelihood combines two components: the weight-outcome likelihood and the intervention-choice likelihood. For individual i, the complete-data likelihood conditional on parameters Θ is:

$$L_i(D_i \mid \Theta) = \prod_{t=1}^{T} \underbrace{\Pr(\Delta w_{it} \mid \boldsymbol{\beta}_i, \sigma_w^2, \mathbf{d}_{i,t-1})}_{\text{weight-outcome component}} \cdot \underbrace{\Pr(\mathbf{d}_{it} \mid \tilde{\boldsymbol{\theta}}_{it}, \boldsymbol{\zeta}_i, \boldsymbol{\phi})}_{\text{choice component}}. \tag{1}$$

Weight-outcome component. Given the weight equation $\Delta w_{it} = \boldsymbol{\beta}_i^\top \mathbf{x}_{s_{i,t-1}} + e_{it}$ with $e_{it} \sim \mathcal{N}(0, \sigma_w^2)$, the contribution from observed weight changes is:

$$\Pr(\Delta w_{it} \mid \boldsymbol{\beta}_i, \sigma_w^2, \mathbf{d}_{i,t-1}) = \frac{1}{\sqrt{2\pi\,\sigma_w^2}} \exp\left(-\frac{\left(\Delta w_{it} - \boldsymbol{\beta}_i^\top \mathbf{x}_{s_{i,t-1}}\right)^2}{2\,\sigma_w^2}\right). \tag{2}$$

Choice component. Under the independent multiple-discrete logit specification, the probability of observing the vector of participation decisions $\mathbf{d}_{it} = (d_{i1t}, \ldots, d_{iKt})^\top$ is:

$$\Pr(\mathbf{d}_{it} \mid \tilde{\boldsymbol{\theta}}_{it},\, \boldsymbol{\zeta}_i,\, \boldsymbol{\phi}) = \prod_{k=1}^{K} p_{ikt}^{d_{ikt}}\,(1 - p_{ikt})^{1-d_{ikt}}, \tag{3}$$

where

$$p_{ikt} = \frac{1}{1 + \exp(-U_{ikt})} \tag{4}$$

is the choice probability for intervention k, with utility U_{ikt} as defined in Equation (1) in the main text. Note that because the belief vector $\tilde{\boldsymbol{\theta}}_{it}$ entering U_{ikt} is itself a function of the learning parameters $(\boldsymbol{\theta}_{0i},\, \boldsymbol{\Sigma}_{\theta_0},\, \{\sigma_{yc}^2\}_{c=1}^{C})$ and the history of signals received through period $t-1$, the dependence on $\boldsymbol{\theta}_{0i}$, $\boldsymbol{\Sigma}_{\theta_0}$, and $\{\sigma_{yc}^2\}_{c=1}^{C})$ is implicit through the recursive Bayesian updating equations specified in Section 4.1.2.

Integration over unobserved heterogeneity. The beliefs $\tilde{\boldsymbol{\theta}}_{it}$ evolve stochastically over time, as they depend on the sequence of noisy signals $y_{ict} \sim \mathcal{N}(\beta_{ic},\, \sigma_{yc}^2)$. Because the signal realizations are unobserved to the researcher, the choice component requires integrating over the distribution of signal paths. In practice, this integral is approximated via Monte Carlo simulation. For each individual i, we draw S independent signal sequences, propagate beliefs forward through the Bayesian updating equations starting from $\boldsymbol{\theta}_{0i}$, compute the resulting choice probabilities along each simulated belief path, and average over simulation draws:

$$\Pr(\mathbf{d}_{it} \mid \boldsymbol{\theta}_{0i},\, \boldsymbol{\zeta}_i,\, \boldsymbol{\phi}) \approx \frac{1}{S} \sum_{s=1}^{S} \prod_{k=1}^{K} p_{ikt}^{(s)\,d_{ikt}} \left(1 - p_{ikt}^{(s)}\right)^{1-d_{ikt}}, \tag{5}$$

where $p_{ikt}^{(s)}$ denotes the choice probability computed under the s-th simulated belief path.

Full-sample likelihood. The likelihood for the entire sample is:

$$L(D \mid \Theta) = \prod_{i=1}^{N} L_i(D_i \mid \Theta). \tag{6}$$

In the MCMC estimation, we work with the log-likelihood $\ell(D \mid \Theta) = \sum_{i=1}^{N} \log L_i(D_i \mid \Theta)$, where

$$\begin{aligned}\log L_i = &\sum_{t=1}^{T} \left[-\frac{1}{2} \log(2\pi\, \sigma_w^2) - \frac{\left(\Delta w_{it} - \boldsymbol{\beta}_i^\top \mathbf{x}_{s_{i,t-1}}\right)^2}{2\, \sigma_w^2} \right] \\ &+ \sum_{t=1}^{T} \log\left(\frac{1}{S} \sum_{s=1}^{S} \prod_{k=1}^{K} p_{ikt}^{(s)\, d_{ikt}} \left(1 - p_{ikt}^{(s)}\right)^{1-d_{ikt}} \right). \end{aligned} \tag{7}$$

The first sum is the weight-outcome log-likelihood, computed in closed form. The second sum is the simulated choice log-likelihood. In our implementation, we use $S = 100$ simulation draws per individual per MCMC iteration.

A5.1. Estimation Steps

Let $\mathbf{B} = (\boldsymbol{b}_0, \boldsymbol{b}_1, \boldsymbol{b}_2, \boldsymbol{b}_3)^T$ denote the coefficient matrix for observed characteristics, which has dimension 4×3. Let $\mathbf{X}$ be the $n \times 4$ design matrix containing the intercept and user characteristics. Let $\mathbf{BI}$ denote the random-coefficient matrix, with each row being $\boldsymbol{\beta}_i$. $\mathbf{BI}$ has dimension $n \times 3$ and it is centered around $\mathbf{BI}_0 = \mathbf{XB}$. We iterate the following steps until convergence. We use superscript r to denote the parameter values at iteration round r.

Step 1: To generate conditional posterior mean $\mathbf{BI}_0^r$, we first generate $\mathbf{B}^r$, as $\mathbf{BI}_0^r = \mathbf{XB}^r$.

Assume conjugate hyper priors:

$$\begin{gathered} p(\mathbf{B}, \boldsymbol{\Sigma}_\beta) = p(\mathbf{B} \mid \boldsymbol{\Sigma}_\beta) \cdot p(\boldsymbol{\Sigma}_\beta), \\ \boldsymbol{\Sigma}_\beta \sim \mathcal{IW}(d_0, \mathbf{V}_0), \\ \mathrm{vec}(\mathbf{B}) \mid \boldsymbol{\Sigma}_\beta \sim \mathcal{N}\left(\mathrm{vec}(\mathbf{B}_0),\, \boldsymbol{\Sigma}_\beta \otimes \mathbf{A}_0^{-1}\right), \end{gathered} \tag{A1}$$

where $\mathcal{IW}$ represents the inverse-Wishart distribution with degree of freedom d_0 and scale matrix $\mathbf{V}_0$, vec($\mathbf{B}$) denotes vectorized $\mathbf{B}$, and $\otimes$ denotes the Kronecker product. Since we have no prior knowledge on $\mathbf{B}$ and $\boldsymbol{\Sigma}_\beta$, we assume diffuse hyper priors to let the estimation be driven by the data. Specifically, $d_0 = 8$, $\mathbf{V}_0 = 8 \cdot \mathbf{I}_{npar}$, $\mathrm{vec}(\mathbf{B}_0) = \mathbf{0}_{nspar \cdot npar}$, and $\mathbf{A}_0 = 0.1 \cdot \mathbf{I}_{nspar}$, where $\mathbf{I}$ denotes

Table A5 Summary of Parameter Notations

Notation	Type	Description
$\boldsymbol{\theta}_{0i}$	Learning-related parameter	Users' initial beliefs on intervention effectiveness; the jth dimension (θ_{ij0}) denotes users' prior beliefs for intervention category j.
$\boldsymbol{\Sigma}_{\theta_0}$	Learning-related parameter	The variance-covariance matrix of users' prior beliefs. The diagonal is normalized to one for identification.
σ^2_{yc}	Learning-related parameter	The variance of signals generated from participating in intervention category c, capturing the noise level in the delivered signal.
$\boldsymbol{\beta}_i$	Signal-generating parameter	Individual-specific true intervention effectiveness; the c-th dimension (β_{ic}) denotes the true mean effectiveness of intervention category c.
σ^2_w	Signal-generating parameter	The variance in the weight-loss distribution.
α	Utility parameter	Individuals' preferences for perceived intervention effectiveness.
$\boldsymbol{\gamma}$	Utility parameter	Individuals' preferences for intervention category-level popularity and attainability.
$\boldsymbol{\lambda}$	Utility parameter	Individuals' preferences for descriptive intervention features, such as intervention length and sentiment polarity.
ρ	Utility parameter	Individuals' preferences for continuous participation.
$\boldsymbol{\zeta}_i$	Utility parameter	User-specific random effects capturing baseline tastes for each category.
$\boldsymbol{\mu}_\zeta$	Hyper parameter	Mean of user-specific random effects.
$\boldsymbol{\Sigma}_\zeta$	Hyper parameter	Variance-covariance matrix of user-specific random effects, capturing correlations between individuals' intervention category preferences.
$\boldsymbol{b}_0$	Hyper parameter	Baseline effect on true mean intervention effectiveness (constant term); the jth dimension (b_{0j}) corresponds to type-j interventions.
$\boldsymbol{b}_1$	Hyper parameter	Effect of gender on true mean intervention effectiveness; the jth dimension (b_{1j}) corresponds to type-j interventions.
$\boldsymbol{b}_2$	Hyper parameter	Effect of age on true mean intervention effectiveness; the jth dimension (b_{2j}) corresponds to type-j interventions.
$\boldsymbol{b}_3$	Hyper parameter	Effect of initial weight on true mean intervention effectiveness; the jth dimension (b_{3j}) corresponds to type-j interventions.
$\boldsymbol{\Sigma}_\beta$	Hyper parameter	Variance-covariance matrix of intervention effectiveness β_i.
$\boldsymbol{\mu}_{\theta_0}$	Hyper parameter	Mean of individuals' log-transformed initial beliefs.
$\boldsymbol{\Sigma}_\mu$	Hyper parameter	Variance-covariance of individuals' log-transformed initial beliefs.

an identity matrix. We use *nspar* to denote the number of parameters controlled for shared characteristics (including the constant term). In our setting, $nspar = 4$. We use *npar* to denote the number of parameters in $\boldsymbol{\beta}_i$; in our setting, $npar = 3$.

The conditional posterior $\mathbf{B}^r$ and $\boldsymbol{\Sigma}_\beta^r$ can be generated as follows:

$$\begin{aligned}\boldsymbol{\Sigma}_\beta^r \mid \mathbf{BI}^{*r-1}, \mathbf{X}^{*r-1} &\sim \mathcal{IW}(d_0 + m, \mathbf{V}_0 + \boldsymbol{\Gamma}),\\ \text{vec}(\mathbf{B}^r) \mid \boldsymbol{\Sigma}_\beta^r, \mathbf{BI}^{*r-1}, \mathbf{X}^{*r-1} &\sim \mathcal{N}\Big(\text{vec}(\tilde{\mathbf{B}}^r),\, \boldsymbol{\Sigma}_\beta^r \otimes \big((\mathbf{X}^{*r-1})^T \mathbf{X}^{*r-1} + \mathbf{A}_0\big)^{-1}\Big),\end{aligned} \tag{A2}$$

where

$$\begin{aligned}\hat{\mathbf{B}} &= \big((\mathbf{X}^{*r-1})^T \mathbf{X}^{*r-1}\big)^{-1} (\mathbf{X}^{*r-1})^T \mathbf{BI}^{*r-1},\\ \tilde{\mathbf{B}}^r &= \big((\mathbf{X}^{*r-1})^T \mathbf{X}^{*r-1} + \mathbf{A}_0\big)^{-1} \Big((\mathbf{X}^{*r-1})^T \mathbf{X}^{*r-1} \hat{\mathbf{B}} + \mathbf{A}_0 \mathbf{B}_0\Big),\\ \boldsymbol{\Gamma} &= (\mathbf{BI}^{*r-1} - \mathbf{X}^{*r-1}\tilde{\mathbf{B}}^r)^T (\mathbf{BI}^{*r-1} - \mathbf{X}^{*r-1}\tilde{\mathbf{B}}^r) + (\tilde{\mathbf{B}}^r - \mathbf{B}_0)^T \mathbf{A}_0 (\tilde{\mathbf{B}}^r - \mathbf{B}_0)\end{aligned}$$

Here, $\mathbf{BI}^{*r-1}$ contains the accepted $\boldsymbol{\beta}_i$ in round $r-1$, m is the number of $\boldsymbol{\beta}_i$ that were accepted, and $\mathbf{X}^{*r-1}$ is the design matrix for the accepted instances.

Step 2: Generate $\boldsymbol{\beta}_i^r$ from the posterior distribution:

$$\boldsymbol{\beta}_i \propto |\boldsymbol{\Sigma}_\beta|^{-\frac{1}{2}} \exp\bigg(-\frac{1}{2}(\boldsymbol{\beta}_i - \mathbf{BI}_{0i})^T \boldsymbol{\Sigma}_\beta^{-1} (\boldsymbol{\beta}_i - \mathbf{BI}_{0i})\bigg)\, L_i\big(D_i \mid \boldsymbol{\beta}_i, \sigma_w^2, \boldsymbol{\theta}_{0i}, \boldsymbol{\sigma}_y^2, \alpha, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \rho, \boldsymbol{\zeta}_i\big)\,. \tag{A3}$$

where $\mathbf{BI}_{0i}$ is the ith row of $\mathbf{BI}_0$. Since the posterior distribution does not have a closed form, we use the Metropolis-Hastings method to make the draws. The proposal distribution is a normal distribution with mean $\boldsymbol{\beta}_i^{r-1}$ (i.e., $\boldsymbol{\beta}_i$ in the last iteration) and variance $\mathbf{v}$ (e.g., $\mathbf{v} = 0.2 \cdot \mathbf{I}$). The acceptance rate is:

$$\text{acc} = \frac{\exp\big(-\frac{1}{2}(\boldsymbol{\beta}_i^* - \mathbf{BI}_{0i}^r)^T \boldsymbol{\Sigma}_\beta^{-1} (\boldsymbol{\beta}_i^* - \mathbf{BI}_{0i}^r)\big)\, L_i\big(D_i \mid \boldsymbol{\beta}_i^*, \sigma_w^2, \boldsymbol{\theta}_{0i}, \boldsymbol{\sigma}_y^2, \alpha, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \rho, \boldsymbol{\zeta}_i\big)}{\exp\big(-\frac{1}{2}(\boldsymbol{\beta}_i^{r-1} - \mathbf{BI}_{0i}^r)^T \boldsymbol{\Sigma}_\beta^{-1} (\boldsymbol{\beta}_i^{r-1} - \mathbf{BI}_{0i}^r)\big)\, L_i(D_i \mid \boldsymbol{\beta}_i^{r-1}, \sigma_w^2, \boldsymbol{\theta}_{0i}, \boldsymbol{\sigma}_y^2, \alpha, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \rho, \boldsymbol{\zeta}_i)}, \tag{A4}$$

where $\boldsymbol{\beta}_i^*$ is the parameter drawn from the proposal distribution. If accepted, $\boldsymbol{\beta}_i^r = \boldsymbol{\beta}_i^*$; otherwise, $\boldsymbol{\beta}_i^r = \boldsymbol{\beta}_i^{r-1}$.

Step 3: Generate the hierarchical parameters for initial beliefs, $\boldsymbol{\mu}^r_{\theta_0}$ and $\boldsymbol{\Sigma}^r_{\mu}$. Assume

$$\boldsymbol{\Sigma}_\mu \sim \mathcal{IW}(d_{0\mu}, \mathbf{D}_\mu), \qquad \boldsymbol{\mu}_{\theta_0} \sim \mathcal{N}(\boldsymbol{m}_\mu, \mathbf{S}_\mu), \tag{A5}$$

where $\boldsymbol{m}_\mu$ and $\mathbf{S}_\mu$ denote the prior mean and covariance matrix for $\boldsymbol{\mu}_{\theta_0}$, respectively. Let $\boldsymbol{O}^{*r-1}$ denote the matrix containing the accepted $\boldsymbol{\theta}_{0i}$ in round $r-1$, and let m be the number of accepted $\boldsymbol{\theta}_{0i}$. Then the covariance matrix is drawn as

$$\boldsymbol{\Sigma}^r_\mu \mid \boldsymbol{O}^{*r-1}, \boldsymbol{\mu}^{r-1}_{\theta_0} \sim \mathcal{IW}(d_{0\mu} + m, \boldsymbol{\Gamma}_\mu), \tag{A6}$$

where $\boldsymbol{\Gamma}_\mu = \mathbf{D}_\mu + (\boldsymbol{O}^{*r-1} - \mathbf{1}_m(\boldsymbol{\mu}^{r-1}_{\theta_0})^T)^T(\boldsymbol{O}^{*r-1} - \mathbf{1}_m(\boldsymbol{\mu}^{r-1}_{\theta_0})^T)$.

Conditional on $\boldsymbol{\Sigma}^r_\mu$, the posterior distribution of $\boldsymbol{\mu}^r_{\theta_0}$ is

$$\boldsymbol{\mu}^r_{\theta_0} \mid \boldsymbol{\Sigma}^r_\mu, \boldsymbol{O}^{*r-1} \sim \mathcal{N}(\tilde{\boldsymbol{\mu}}^r_{\theta_0}, \mathbf{V}^r_\mu), \tag{A7}$$

where $\mathbf{V}^r_\mu = \left(m(\boldsymbol{\Sigma}^r_\mu)^{-1} + \mathbf{S}^{-1}_\mu\right)^{-1}$, and $\tilde{\boldsymbol{\mu}}^r_{\theta_0} = \mathbf{V}^r_\mu \left((\boldsymbol{\Sigma}^r_\mu)^{-1} \sum_{i=1}^m \boldsymbol{\theta}^{*r-1}_{0i} + \mathbf{S}^{-1}_\mu \boldsymbol{m}_\mu\right)$.

Step 4: Generate $\boldsymbol{\theta}^r_{0i}$ from the posterior distribution:

$$\boldsymbol{\theta}_{0i} \propto |\boldsymbol{\Sigma}_\mu|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(\boldsymbol{\theta}_{0i} - \boldsymbol{\mu}_{\theta_0})^T \boldsymbol{\Sigma}^{-1}_\mu (\boldsymbol{\theta}_{0i} - \boldsymbol{\mu}_{\theta_0})\right) L_i\left(D_i \mid \boldsymbol{\beta}_i, \sigma^2_w, \boldsymbol{\theta}_{0i}, \boldsymbol{\sigma}^2_y, \alpha, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \rho, \boldsymbol{\zeta}_i\right). \tag{A8}$$

Since the posterior distribution does not have a closed form, we use the Metropolis-Hastings method to make the draws. The proposal distribution is constructed in log space. Specifically, we propose $\log \boldsymbol{\theta}^*_{0i} \sim \mathcal{N}\left(\log \boldsymbol{\theta}^{r-1}_{0i},\, v_\theta \mathbf{I}\right)$, and then transform back as $\boldsymbol{\theta}^*_{0i} = \exp(\log \boldsymbol{\theta}^*_{0i})$. This guarantees positivity of $\boldsymbol{\theta}_{0i}$, and the acceptance ratio includes the corresponding Jacobian correction:

$$acc = \frac{\exp\left(-\frac{1}{2}(\boldsymbol{\theta}^*_{0i} - \boldsymbol{\mu}^r_{\theta_0})^T (\boldsymbol{\Sigma}^r_\mu)^{-1} (\boldsymbol{\theta}^*_{0i} - \boldsymbol{\mu}^r_{\theta_0})\right)\ L_i\left(D_i \mid \boldsymbol{\beta}^r_i, \sigma^{2,r}_w, \boldsymbol{\theta}^*_{0i}, \boldsymbol{\sigma}^{2,r}_y, \alpha^r, \boldsymbol{\gamma}^r, \boldsymbol{\lambda}^r, \rho^r, \boldsymbol{\zeta}^r_i\right)\ \prod_{j=1}^C \theta^*_{0ij}}{\exp\left(-\frac{1}{2}(\boldsymbol{\theta}^{r-1}_{0i} - \boldsymbol{\mu}^r_{\theta_0})^T (\boldsymbol{\Sigma}^r_\mu)^{-1} (\boldsymbol{\theta}^{r-1}_{0i} - \boldsymbol{\mu}^r_{\theta_0})\right)\ L_i\left(D_i \mid \boldsymbol{\beta}^r_i, \sigma^{2,r}_w, \boldsymbol{\theta}^{r-1}_{0i}, \boldsymbol{\sigma}^{2,r}_y, \alpha^r, \boldsymbol{\gamma}^r, \boldsymbol{\lambda}^r, \rho^r, \boldsymbol{\zeta}^r_i\right)\ \prod_{j=1}^C \theta^{r-1}_{0ij}}. \tag{A9}$$

If accepted, $\boldsymbol{\theta}^r_{0i} = \boldsymbol{\theta}^*_{0i}$; otherwise, $\boldsymbol{\theta}^r_{0i} = \boldsymbol{\theta}^{r-1}_{0i}$

Step 5: Update the individual-specific category fixed effects and their hyperparameters. The individual-level parameters $\boldsymbol{\zeta}^r_i$ are generated from the posterior distribution:

$$\boldsymbol{\zeta}_i \propto |\boldsymbol{\Sigma}_\zeta|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(\boldsymbol{\zeta}_i - \boldsymbol{\mu}_\zeta)^T \boldsymbol{\Sigma}^{-1}_\zeta (\boldsymbol{\zeta}_i - \boldsymbol{\mu}_\zeta)\right) L_i\left(D_i \mid \boldsymbol{\beta}_i, \sigma^2_w, \boldsymbol{\theta}_{0i}, \boldsymbol{\sigma}^2_y, \alpha, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \rho, \boldsymbol{\zeta}_i\right). \tag{A10}$$

Similarly, we use a Metropolis-Hastings method with a multivariate normal proposal centered at $\boldsymbol{\zeta}_i^{r-1}$.

After updating $\boldsymbol{\zeta}_i^r$, the hyperparameters are updated using the accepted draws. Let $\mathcal{I}^*$ denote the set of accepted individuals in the current iteration, with cardinality m. The mean is updated as $\boldsymbol{\mu}_\zeta^r = \frac{1}{m}\sum_{i\in\mathcal{I}^*}\boldsymbol{\zeta}_i^r$, and the covariance matrix is updated as $\boldsymbol{\Sigma}_\zeta^r = \frac{1}{m}\sum_{i\in\mathcal{I}^*}(\boldsymbol{\zeta}_i^r - \boldsymbol{\mu}_\zeta^r)(\boldsymbol{\zeta}_i^r - \boldsymbol{\mu}_\zeta^r)^T$.

Step 6: In this step, we generate the parameters that are common across individuals $\boldsymbol{\phi}$. We assume a diffuse prior for $\boldsymbol{\phi}$: $\boldsymbol{\phi} \sim \mathcal{N}(\boldsymbol{\phi}_0, \boldsymbol{\Sigma}_0), \boldsymbol{\phi}_0 = \mathbf{0}, \boldsymbol{\Sigma}_0 = 10 \cdot \mathbf{I}$.

The posterior distribution can be expressed as:

$$\boldsymbol{\phi} \propto |\boldsymbol{\Sigma}_0|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(\boldsymbol{\phi} - \boldsymbol{\phi}_0)^T \boldsymbol{\Sigma}_0^{-1}(\boldsymbol{\phi} - \boldsymbol{\phi}_0)\right) \cdot L(D \mid \boldsymbol{\beta}_i, \boldsymbol{\theta}_{0i}, \boldsymbol{\phi}, \boldsymbol{\zeta}_i)\,. \tag{A11}$$

Using a Metropolis-Hastings method, the proposed parameters are accepted with the rate:

$$\mathrm{acc} = \frac{\exp\left(-\frac{1}{2}(\boldsymbol{\phi}^* - \boldsymbol{\phi}_0)^T \boldsymbol{\Sigma}_0^{-1}(\boldsymbol{\phi}^* - \boldsymbol{\phi}_0)\right) \cdot L(D \mid \boldsymbol{\beta}_i, \boldsymbol{\theta}_{0i}, \boldsymbol{\phi}^*, \boldsymbol{\zeta}_i)}{\exp\left(-\frac{1}{2}(\boldsymbol{\phi}^{r-1} - \boldsymbol{\phi}_0)^T \boldsymbol{\Sigma}_0^{-1}(\boldsymbol{\phi}^{r-1} - \boldsymbol{\phi}_0)\right) \cdot L\left(D \mid \boldsymbol{\beta}_i, \boldsymbol{\theta}_{0i}, \boldsymbol{\phi}^{r-1}, \boldsymbol{\zeta}_i\right)}, \tag{A12}$$

where $\boldsymbol{\phi}^*$ is the proposed parameter vector centered around the parameter values in the last round. If accepted, $\boldsymbol{\phi}^r = \boldsymbol{\phi}^*$; otherwise, $\boldsymbol{\phi}^r = \boldsymbol{\phi}^{r-1}$.

We follow the literature to carry out a two-step estimation (Ho et al. 2017). In the first step, we execute the MCMC for 40000 rounds, with the first half discarded as the burn-in samples. In the second step, we use the sample means derived from the first step as the initial values and run a separate MCMC sampler for 20000 rounds. We use 20 as the thinning window to generate the posterior distribution. In Figure A1, we show the likelihood value by each round. As can be seen, the likelihood become steady quickly. To quantitatively test the MCMC convergence, we perform a Gelman-Rubin diagnostic (Gelman and Rubin 1992) by running another parallel sample chain with different seeds and initial values. Our results show that the potential scale reduction factors are bounded by range 0.99-1.2, which suggests that the MCMC has converged to the targeted distribution.

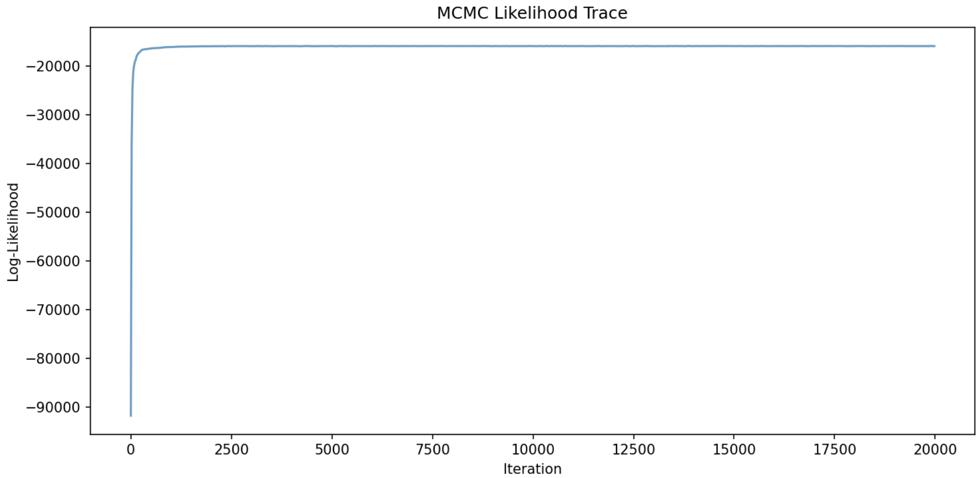

Figure A1 **Likelihood Trajectory.**

A5.2. Identification

Identification comes from three sources. First, variation in intervention choices across available alternatives and across time identifies the utility parameters $(\alpha, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \rho)$ as well as the user-specific category fixed effects. Second, the joint observation of weight changes and selected intervention categories identifies the true intervention-effectiveness parameters $\boldsymbol{\beta}_i$, because the weight equation provides a separate outcome-generating process that links realized weight change to category-specific intervention exposure. Third, learning-related parameters are identified from the divergence between initial participation patterns and subsequent choices after users accumulate experience. In particular, the model distinguishes prior beliefs from true effectiveness because posterior beliefs enter the utility function, while true effectiveness governs the signal-generation and weight-outcome processes. As shown in the Bayesian learning framework, only the relative scale of the prior variance and the signal variance is identified under Gaussian Bayesian updating; therefore, we normalize the prior variance to one and interpret σ^2_{yc} as the relative noisiness of the feedback signal for category c compared to the prior variances.

A6. Additional Estimation Results

A6.1. Estimated Intervention Effectiveness

In Table A6, we report the parameter estimates for the true mean effectiveness for each intervention category (i.e., b_0, b_1, b_2, b_3). The first column reports the constant (intercept) effects, while columns 2-4 shows the effects of individuals' characteristics (gender, age, initial weight) on the true mean intervention effectiveness. Our results show that interventions that focus on self-regulation have higher effectiveness at the baseline level. In addition, outcome-oriented interventions tend to be more effective for individuals who have a higher initial weight. The variance of intervention effectiveness is reported in Table A7.

Table A6 Estimated Intervention Effectiveness

	Constant (b_0)	**Gender (b_1)**	**Age (b_2)**	**Initial Weight (b_3)**
Type 1 (*BehaviorSpecific*)	−0.0116	0.0728	0.0030	−0.0375
Type 2 (*OutcomeOriented*)	0.0539	−0.0409	−0.0177	0.0802*
Type 3 (*SelfRegulation*)	0.1450**	−0.0980	−0.0220	−0.0334

Note: * 90%, ** 95%, *** 99% credible intervals exclude zero.

Table A7 Estimated Variance-Covariance Matrix Σ_β

Var-Cov	β_{i1}	β_{i2}	β_{i3}
β_{i1}	0.4504***	−0.0657	−0.0712
β_{i2}	−0.0657	0.7773***	0.0124
β_{i3}	−0.0712	0.0124	0.8148***

Figure A2 visualizes the overall individual heterogeneity for each intervention type. The figure shows that users may experience varied outcomes from intervention participation, with effects spanning from negative to positive. This variability aligns with the expectation that not all interventions suit all users. For instance, interventions can cause psychological stress for certain individuals, potentially triggering disordered eating or even resulting in weight gain (Jung et al. 2023). These results further underscore the importance of individuals' self-observation and self-learning, given that health management is a highly personal experience. This necessitates the development of useful strategies to help individuals evaluate and comprehend the potential benefits and drawbacks associated with intervention participation. In this paper, we identify possible obstacles and propose several actionable denoising strategies to improve user engagement.

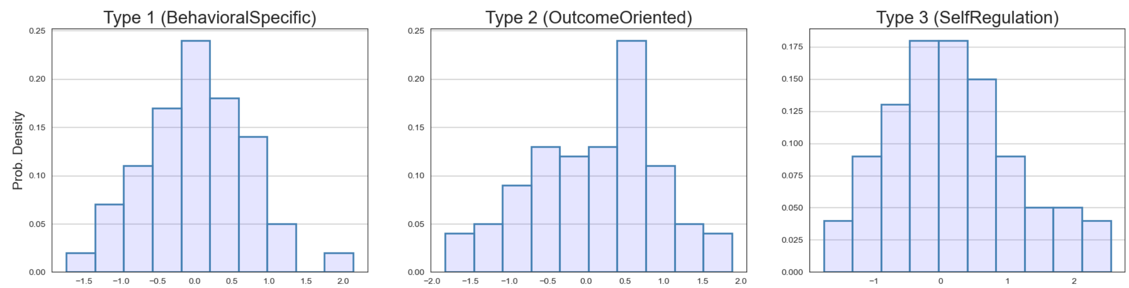

Figure A2 Distribution of True Intervention Effectiveness in the User Population.

A6.2. Parameter Estimates in the Utility Function

Table A8 shows that, in general, users prefer interventions that are popular within the community ($\gamma_1 > 0$). While users prefer interventions with longer instructions ($\lambda_1 > 0$), excessively lengthy instructions could decrease their likelihood of selection ($\lambda_2 < 0$). The results further suggest that users are attracted to interventions expressed with positive sentiment ($\lambda_3 > 0$) and are repelled by negative sentiment ($\lambda_4 < 0$). Higher-intensity interventions are less appealing ($\lambda_5 < 0$). Lastly, users who have participated in an intervention in the previous period are more likely to continue with it ($\rho > 0$), indicating a propensity to sustain with ongoing interventions.

Table A8 Parameter Estimates in the Utility Function

Parameter	Description	Posterior Mean
α	Coefficient of expected weight loss	0.0257***
γ_1	Coefficient of intervention-type popularity	0.3176***
γ_2	Coefficient of intervention-type attainability	−0.0253
λ_1	Coefficient of intervention description length	0.1442*
λ_2	Coefficient of squared intervention description length	−0.1658*
λ_3	Coefficient of positive sentiment	0.0388*
λ_4	Coefficient of negative sentiment	−0.1058***
λ_5	Coefficient of intervention intensity	−0.1161***
ρ	Coefficient of continued participation	7.7057***

Note: * 90%, ** 95%, *** 99% credible intervals exclude zero.

A6.3. Robustness Checks

We conduct two sets of robustness checks to assess the sensitivity of our results. The first set examines alternative model specifications, including removing the attainability variable and estimating a two-stage decision model. These tests address concerns that attainability may introduce double counting with learned effectiveness and that the choice structure may affect the results. The second set of robustness checks focuses on the normalization of prior variance in the Bayesian learning process. Since prior variance and signal variance are not separately identified, we explore alternative values of prior variance and assess whether the main findings are sensitive to this normalization.

A6.3.1. Alternative Specifications We conduct two specification-based robustness checks to address potential concerns regarding model structure and variable construction. First, we re-estimate the model by removing the attainability variable from the utility function. This test is motivated by the concern that attainability may overlap with the learned effectiveness component, potentially introducing a form of double counting. By excluding this variable, we assess whether our estimates of learning-related parameters are sensitive to this specification.

Second, we estimate a two-stage decision model to account for the possibility that users make hierarchical choices. In the first stage, users select intervention categories using a multivariate probit model, allowing for correlated preferences across categories. In the second stage, conditional on the selected category, users choose specific interventions within that category using a binary logit model. This structure captures an alternative decision process in which users first determine the category of intervention and then select among alternatives within that category.

Table A9 reports the estimated learning-related parameters across these alternative specifications. The results show that the relative magnitudes of signal variances remain stable across specifications, with self-regulation interventions exhibiting the highest level of noise, followed by outcome-oriented and behavior-specific interventions. This consistency suggests that our main conclusions are not driven by the inclusion of the attainability variable or by the assumed decision structure.

A6.3.2. Sensitivity to Prior Variance Normalization Another concern is that individuals may have different confidence in intervention effectiveness, reflected as different prior variances of their beliefs. However, in the Bayesian learning framework, prior variance is not separately identified from signal variance, as learning dynamics depend on their relative precision. Therefore, our main

Table A9 Robustness Checks on Choice Model Specifications

	Main Model	No Attainability	Two-Stage
Prior Mean			
Type 1 (BehaviorSpecific)	0.1924***	0.1805***	0.1987***
Type 2 (OutcomeOriented)	0.1855***	0.1775***	0.1669***
Type 3 (SelfRegulation)	0.1158***	0.1218***	0.1046***
Signal Variance			
Type 1 (BehaviorSpecific)	0.1991***	0.1867***	0.4065***
Type 2 (OutcomeOriented)	0.5185***	0.9474***	1.1519***
Type 3 (SelfRegulation)	2.8091***	1.7732***	2.0740***

Note: * 90%, ** 95%, *** 99% credible intervals exclude zero.

analysis follows the existing literature by normalizing the prior variance to one for identification (Ho et al. 2017, Zhang et al. 2019). This normalization implies that differences in belief updating are captured through variation in signal noise.

To assess how sensitive our results are to this identification strategy, we conduct a series of robustness checks. First, we re-estimate the model under alternative fixed values of prior variance (0.5 and 2), representing lower and higher levels of prior uncertainty. Second, we allow prior variance to vary across individuals using tenure as a proxy for experience, assigning lower prior variance (0.5) to long-tenure users (i.e., users whose tenure is longer than the median) and higher prior variance (2) to short-tenure users (i.e., users whose tenure is shorter than the median). The results (Table A10) show that while the magnitudes of the estimated signal variances adjust with changes in prior variance, the relative ranking across intervention categories remains stable. In particular, self-regulation interventions consistently exhibit the highest level of noise, followed by outcome-oriented and behavior-specific interventions. These findings suggest that our main conclusions are robust to alternative specifications of prior uncertainty and are not driven by the normalization assumption.

A6.3.3. Alternative Intervention Categorization Our main specification groups interventions into three categories: behavior-specific, outcome-oriented, and self-regulation-focused interventions. This categorization is based on the conceptual distinctions derived from social cognitive theory, and is intended to capture differences in learning conditions, such as feedback clarity, consistency, and interpretability, which are central to our theoretical framework.

Table A10 Sensitivity to Prior Variance Normalization

	Main Model	PV = 0.5	PV = 2	Tenure-based Proxy
Prior Mean				
Type 1 (BehaviorSpecific)	0.1924***	0.2057***	0.1797***	0.1851***
Type 2 (OutcomeOriented)	0.1855***	0.1879***	0.1929***	0.1714***
Type 3 (SelfRegulation)	0.1158***	0.1357***	0.1054***	0.1275***
Signal Variance				
Type 1 (BehaviorSpecific)	0.1991***	0.4020***	0.2901***	0.1111***
Type 2 (OutcomeOriented)	0.5185***	0.9622***	0.5053***	0.4034***
Type 3 (SelfRegulation)	2.8091***	1.5420***	1.1483***	1.4429***

Note: * 90%, ** 95%, *** 99% credible intervals exclude zero.

To examine whether our findings depend on this level of aggregation, we re-estimate the model using a more fin-grained five-category specification as a robustness check. Specifically, the behavior-specific category contains three distinct behavioral targets: (1) Nutrition improvement – improving the quality and composition of diet; (2) Caloric restriction – reducing overall caloric intake through portion/frequency control; (3) Exercise – increasing energy expenditure via structured movement. It might be meaningful to retain this subdivision, as these behaviors differ in mechanisms of action, engagement patterns, and health impacts. The purpose of this specification is not to redefine the main conceptual structure of the paper, but to verify that the main conclusions are robust to a more disaggregated classification of interventions.

The results are reported in Table A11. We find that the qualitative patterns remain consistent with the main specification. Outcome-oriented and self-regulation interventions continue to show substantially higher signal variance, suggesting noisier learning conditions and greater difficulty in belief updating.

A7. Additional Counterfactual Analysis Results

In this counterfactual analysis, we examine individuals' engagement contingent on their initial beliefs as well as persistence levels. Users' persistence level can be reflected in their preferences of continued participation (parameter ρ), as a stronger inclination toward ongoing engagement suggests a greater degree of persistence. We add random noises to the parameter ρ to simulate a range of users with different persistence levels. We simulate and plot user engagement over 30 periods to see how their behaviors change across time.

Table A11 Estimates of Learning-Related Parameters

Parameter	Posterior Mean	
Prior Mean		
Category 1 (*NutritionImprovement*)	0.0595	***
Category 2 (*CalorieRestriction*)	0.1891	***
Category 3 (*Exercise*)	0.6637	***
Category 4 (*OutcomeOriented*)	0.7841	***
Category 5 (*SelfRegulation*)	0.6286	***
Signal Variance		
Category 1 (*NutritionImprovement*)	0.7747	***
Category 2 (*CalorieRestriction*)	1.3736	***
Category 3 (*Exercise*)	0.7020	***
Category 4 (*OutcomeOriented*)	2.9894	***
Category 5 (*SelfRegulation*)	6.5706	***

Notes: * 90%, ** 95%, *** 99% credible interval does not contain 0.

Our results, as depicted in Figure A3, indicate that users with a high level of persistence are able to sustain a high degree of engagement, whereas those with low persistence may experience a decline in engagement over time. The decrease rate is greater when users are optimistic. In addition, low-persistence users display more fluctuation patterns that may suggest a greater vulnerability to negative signals. When exposed to such signals, these users may discontinue their participation in the current intervention and switch to alternative types of intervention. This phenomenon could potentially account for the wave-like patterns observed in their engagement curves. Our findings indicate that noise reduction is the most effective for users with low levels of persistence, as these users are more susceptible to being negatively impacted by uncertain feedback. By reducing the likelihood of biased signals, denoising can help to mitigate the risk of rapid disengagement among these users.

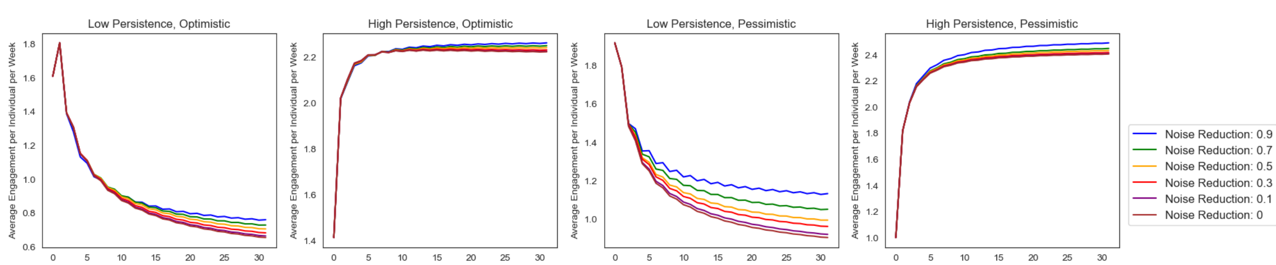

Figure A3 User Engagement by Initial Belief and Persistence Levels.

A8. Denoising Strategies to Improve Individuals' Learning Efficiency

In this section, we propose several denoising strategies that healthcare platforms can utilize in practice to improve individuals' learning performance. We consider two learning schemes that seek to reduce signal noises from different sources. In the first learning scheme, we consider the scenario in which individuals are able to aggregate signals from multiple engagement periods. This learning scheme may help reduce the noises resulting from individuals' short-term health fluctuations. In the second learning scheme, we examine how the accessibility of similar peers' engagement experiences can impact learning efficiency. Learning from peers could help users reduce the noises brought by their own behavior courses, as it enables them to combine personal experience with the experiences of others to derive a more comprehensive understanding. In the following, we provide details of these learning strategies.

A8.1. Learning from Multiple Periods

Suppose that online healthcare interventions facilitate individuals' learning of healthcare signals from multiple periods. For example, interventions can exhibit individuals' weight-loss progress not only for the previous period, but also for the past n $(n > 1)$ periods, empowering them to integrate more signals from their engagement experiences. Specifically, the Bayes updating mechanism can be expressed as: $\tilde{\theta}_{ict} \sim \mathcal{N}(\theta_{ict}, \sigma^2_{\theta ict})$, where $\theta_{ict} = \left(\frac{\theta_{ic,t-1}}{\sigma^2_{\theta ic,t-1}} + \frac{m_{ict} \cdot \bar{y}_{ict}}{\sigma^2_{yc}}\right)\left(\frac{1}{\sigma^2_{\theta ic,t-1}} + \frac{m_{ict}}{\sigma^2_{yc}}\right)^{-1}$, and $\sigma^2_{\theta ict} = \left(\frac{1}{\sigma^2_{\theta ic,t-1}} + \frac{m_{ict}}{\sigma^2_{yc}}\right)^{-1}$. Here, $\bar{y}_{ict}$ denotes the average signals that user i receives for category c during the past n $(n = 2, 3, 4, \ldots)$ weeks prior to time t, and m_{ict} is the number of signals that user i receives. Under this new learning scheme, we simulate users' learning trajectories and calculate the average learning rate. Our results, as shown in Table A12, indicate that the more engagement periods that individuals use to update their posterior beliefs, the higher the learning rate. In Figure A4, we plot the simulated average user engagement when different numbers of periods are used. The results demonstrate that improvement in learning performance further translates into an increase in user engagement. These results support the efficacy of our proposed scheme. In chronic condition management, individuals may experience frequent health fluctuations, such as repeated weight loss and regain. Therefore, it can be critical for individuals to aggregate multiple periods of engagement experiences to obtain a more reliable observation.

Table A12 Users' Average Learning Rate: Learning from Multiple Periods

	$n=2$	$n=3$	$n=4$	$n=5$
Category 1 (*BehaviorSpecific*)	85.52	85.97	86.15	86.24
Category 2 (*OutcomeOriented*)	44.05	46.10	47.13	47.71
Category 3 (*SelfRegulation*)	50.02	56.66	60.53	63.04

Note: Values of learning rates are percentages.

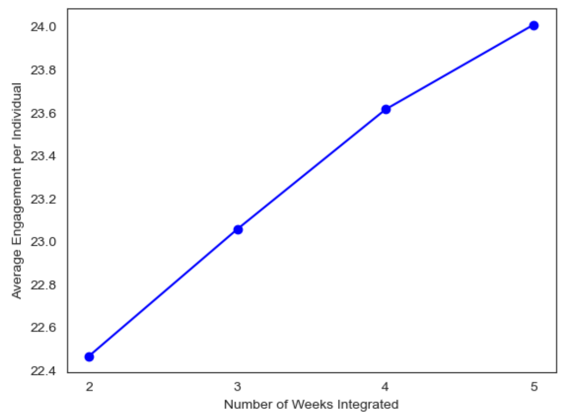

Figure A4 Average User Engagement: Learning from Multiple Periods.

A8.2. Learning from Similar Peers

During our data collection, the focal platform did not provide specific functions for users to identify similar peers and learn from their experiences. In this section, we consider the scenario in which users can observe not only their own engagement experiences, but also those of their peers. Suppose each user can observe other (anonymous) users' profiles (e.g., gender, age, initial weight, etc.) as well as their past intervention participation experiences (e.g., how an intervention worked on their health). Based on this additional information, each user can identify peers with similar characteristics (e.g., same gender, similar age and body weight) and incorporate their engagement experiences to update their posterior beliefs.

Denote the set of similar users of user i as S_i. Individuals' learning scheme is expressed as: $\tilde{\theta}_{ict} \sim \mathcal{N}(\theta_{ict}, \sigma^2_{\theta ict})$, where $\theta_{ict} = \left(\frac{\theta_{ic,t-1}}{\sigma^2_{\theta ic,t-1}} + \frac{n_{ict} \cdot \bar{y}_{ict}}{\sigma^2_{yc}}\right)\left(\frac{1}{\sigma^2_{\theta ic,t-1}} + \frac{n_{ict}}{\sigma^2_{yc}}\right)^{-1}$, and $\sigma^2_{\theta ict} = \left(\frac{1}{\sigma^2_{\theta ic,t-1}} + \frac{n_{ict}}{\sigma^2_{yc}}\right)^{-1}$. Here, $\bar{y}_{ict}$ denotes the average of the signals that user i and his/her similar peers received for category-c interventions; that is, $\bar{y}_{ict} = n_{ict}^{-1}\left(l_{ict} \cdot y_{ict} + \sum_{u \in S_i} l_{uct} \cdot y_{uct}\right)$, in which l_{ict} represents the number of signals that user i has received for type-c interventions in week t, and $n_{ict} = l_{ict} +$

$\sum_{u \in S_i} l_{uct}$ is the total number of signals. In determining the set of similar users S_i, we calculate the cosine similarity between two users' observed demographic features and include the users whose cosine similarity surpasses a certain threshold z. The smaller the value of z is, the more users will be included in the similar user set. Therefore, we can vary the value of z to control the degree of learning among peers.

We simulate individuals' learning trajectories based on the learning scheme described above. In Table A13, we report our results of users' average learning rates under different threshold levels. We find that when individuals are allowed to learn from similar others, their average learning rate increases dramatically. This is because individuals can now gather experiences even if they do not participate in an intervention themselves. However, it is not necessarily the case that the more peers that individuals learn from (i.e., a lower threshold), the significantly better their learning performance becomes. This phenomenon is potentially due to the tradeoff between the number of the experience signals and the accuracy or "quality" of the signals. When the similarity threshold takes a lower value, on the one hand, individuals incorporate more users as the "similar peers" and use more observations to update their posterior beliefs; on the other hand, individuals are referencing more users who are not that similar to themselves, so using their engagement experiences could be misleading. These findings suggest that although learning from peers is an overall effective strategy to improve users' learning performance, individuals still need to consider who they learn from (i.e., high-similarity users vs. low-similarity users) in order to achieve the best improvement. Figure A5 illustrates the simulated average user engagement under different similarity thresholds.

Table A13 Users' Average Learning Rate: Learning from Similar Peers

	$z = 0.6$	$z = 0.7$	$z = 0.8$	$z = 0.9$	$z = 0.95$
Avg No. of Similar Users	19.92	16.36	11.98	7.30	4.08
Average Learning Rate					
Category 1 (*BehaviorSpecific*)	93.36	93.31	93.15	93.01	92.78
Category 2 (*OutcomeOriented*)	92.93	88.95	88.37	82.31	75.44
Category 3 (*SelfRegulation*)	67.22	70.79	69.93	61.69	52.26

Note: Values of learning rates are percentages.

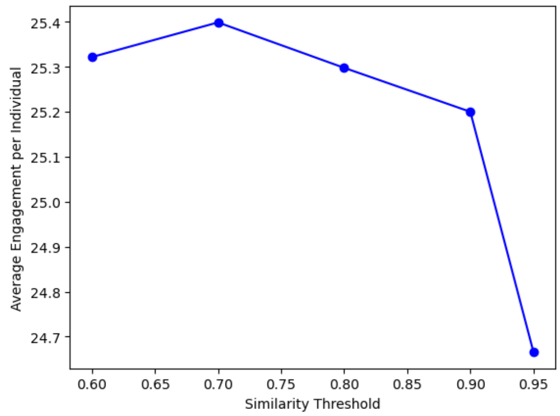

Figure A5 Average User Engagement: Learning from Similar Peers.

References

Gelman A, Rubin DB (1992) Inference from iterative simulation using multiple sequences. *Statistical science* 7(4):457–472.

Ho YC, Wu J, Tan Y (2017) Disconfirmation effect on online rating behavior: A structural model. *Information Systems Research* 28(3):626–642.

Hutto C, Gilbert E (2014) Vader: A parsimonious rule-based model for sentiment analysis of social media text. *Proceedings of the international AAAI conference on web and social media*, volume 8, 216–225.

Jung FU, Riedel-Heller SG, Luck-Sikorski C (2023) The relationship between weight history and psychological health—differences related to gender and weight loss patterns. *PLoS One* 18(2):e0281776.

Zhang S, Singh PV, Ghose A (2019) A structural analysis of the role of superstars in crowdsourcing contests. *Information Systems Research* 30(1):15–33.

Zhou T, Wang Y, Yan L, Tan Y (2023) Spoiled for choice? personalized recommendation for healthcare decisions: A multiarmed bandit approach. *Information Systems Research* 34(4):1493–1512.