

© Copyright 2021

Tongxin Zhou

Digital Healthcare: Individual Engagement and IT-Enabled Design

Tongxin Zhou

A dissertation

submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

University of Washington

2021

Reading Committee:

Yong Tan, Chair

Ming Fan

Yingfei Wang

Program Authorized to Offer Degree:

Business Administration

University of Washington

Abstract

Digital Healthcare: Individual Engagement and IT-Enabled Design

Tongxin Zhou

Chair of the Supervisory Committee:
Michael G. Foster Endowed Professor Yong Tan
Information Systems and Operations Management

In this dissertation, I study the transformational effect of information technology on healthcare. I take two main perspectives to investigate the topic: individual engagement and IT-enabled design. Along the first perspective, I study individuals' engagement process in online healthcare communities. Individuals' engagement in online healthcare communities provides opportunities for healthcare providers to deliver services in a timely and cost-effective manner. In an aim to better understand and facilitate individuals' online engagement, I apply econometric structural approaches to disentangle individuals' online behavior dynamics in my first two essays. Specifically, the first essay studies individuals' participation states that drive their self-regulatory behaviors online, and the second essay investigates individuals' learning behaviors during their health-management process. Along the second perspective, I follow the design-science paradigm

to study personalized healthcare recommendation systems in my third essay. To address several key challenges in healthcare recommendations, I propose a bandit-based recommendation framework that is enhanced by deep-learning feature engineering and diversity promotion constraint. I evaluate the proposed framework through extensive simulation experiments, and the results justify the superiority of the model as compared to a wide range of state-of-the-art recommendation systems.

TABLE OF CONTENTS

List of Figures	vi
List of Tables	vii
Chapter 1. Introduction	1
Chapter 2. Literature Review	5
2.1 Personal Health Management	5
2.2 Online Healthcare Communities.....	7
2.3 Healthcare Recommendations	8
Chapter 3. Multi-Dimensional Participation in Online Weight Management	12
3.1 Theoretical Framework.....	18
3.2 Models.....	20
3.2.1 The Framework of Continuous-Time Markovian Process.....	22
3.2.2 MCTHMM Framework	24
3.2.3 Model Specification.....	26
3.2.4 Model Identification and Estimation	28
3.3 Context and Data.....	30
3.3.1 Focal Online Weight-Loss Community	30
3.3.2 Data	31
3.4 Empirical Results and Findings	35
3.4.1 Model Selection and Model Estimates	36
3.4.2 Dynamics Associated with Self-Regulatory Dimensions	38

3.4.3	Dynamics of Efficacy State	40
3.4.4	Robustness Checks.....	43
3.5	Additional Analysis	44
3.5.1	Perspective of Information Synergy	44
3.5.2	Model Performance Comparisons.....	45
3.5.3	Individuals' Regular Engagement.....	48
3.6	Discussion and Conclusion	49
Chapter 4. A Structural Model of User Dynamic Engagement in Online Healthcare Interventions		
		53
4.1	Context.....	55
4.1.1	Research Context and Data.....	55
4.1.2	Model-Free Evidence.....	57
4.2	Model	58
4.2.1	Challenge Selection	58
4.2.2	Learning	60
4.2.3	Signal Generating Process	61
4.2.4	Likelihood	62
4.2.5	Estimation	63
4.2.6	Identification	64
4.3	Results.....	65
4.3.1	Main Estimation Results	65
4.3.2	Learning Curve	69
4.4	Conclusion and Discussion	71

Chapter 5. Personalized Recommendation for Healthcare Decisions: A Multi-Armed Bandit

Approach.....	74
5.1 Design Theories	79
5.2 Model Development.....	80
5.2.1 Optimizing Intervention with Multi-Armed Bandit.....	82
5.2.2 Diversity Constraint	83
5.2.3 Deep-Learning-Based Feature Engineering.....	85
5.3 Evaluation Design.....	88
5.3.1 Focal Evaluation Content.....	88
5.3.2 Data	90
5.3.3 Operationalizing Model Components: Embeddings and Diversity Constraint.....	93
5.4 Performance Analysis and Result	94
5.4.1 Experiment 1: t-SNE Visualization of Context Space	94
5.4.2 Experiment 2: Ablation Analysis of Model Components.....	98
5.4.3 Experiment 3: Comparison with Benchmark Recommendation Systems	102
5.4.4 Experiment 4: On Dynamic User Preferences	105
5.4.5 Experiment 5: Diversity Analysis	107
5.4.6 Experiment 6: User Improvement.....	108
5.4.7 Experiment 7: Using Our Recommendation Framework to Improve Weight Loss	109
5.5 Conclusion and Discussion	111
Chapter 6. Concluding Remarks	115
Bibliography	117

Appendix A.....	130
A.1 Continuous or Discrete	130
A.2 User Descriptive Statistics	131
A.3 Estimation Procedure	132
A.3.1 Latent Class Model	132
A.3.2 Discount Factor in Exponential Learning	134
A.3.3 Selecting the Number of Supports for the Random-Effect Variables.....	134
A.4 Robustness Checks.....	135
A.4.1 Alternative State-Dependent Outcome Distribution	135
A.4.2 Alternative Variable Measures.....	136
A.4.3 Sampling Seeds	138
A.5 Single-Dimensional Models.....	139
A.6 Prediction Procedure	141
Appendix B	143
Appendix C	146
C.1 Deep Learning Embeddings	146
C.1.1 Challenge Embedding Model	146
C.1.2 User Embedding Model	147
C.2 Screenshot of Weight-loss Challenges on the Focal Website	152
C.3 Challenge Meta-Data Annotation	154
C.4 Data Summary Statistics	157
C.5 Top-K Precision, Recall, and F-Measure	157

C.6 Evaluation Approach.....	158
C.6.1 Doubly-Robust Estimation.....	158
C.6.2 Offline Precision	161
C.6.3 Simulation	162
C.7 Benchmark Models	162
C.7.1 Baseline MABs	162
C.7.2 Benchmark State-of-the-Art Recommendation Systems	163
C.8 Supplementary Results of Experiment 4.....	168

LIST OF FIGURES

Figure 3.1 Online Weight Management: A Multi-Dimensional Dynamic Process	20
Figure 3.2 A CTHMM Stochastic Process	25
Figure 3.3 Hidden State Trajectories	49
Figure 4.1 Individuals' Participation in Online Weight-Loss Challenges	56
Figure 4.2 Distribution of True Challenge Effectiveness in the User Population	68
Figure 4.3 Average Learning Curve	70
Figure 5.1 An MAB-Driven Framework for Healthcare Recommendation	81
Figure 5.2 Illustration of the Base Architecture.....	87
Figure 5.3 Challenge Embedding t-SNE Visualization	96
Figure 5.4 User Embedding t-SNE Visualization.....	97
Figure 5.5 Learning Curve of Our Model and Baseline MABs.....	102
Figure 5.6 Performance Variation for Dynamic Users	107
Figure 5.7 Diversity Distribution.....	108
Figure 5.8 User Improvement (Using PMF as Benchmark)	109

LIST OF TABLES

Table 3.1 Variable Description and Summary Statistics	35
Table 3.2 Selection on the Number of Hidden States	36
Table 3.3 Estimates for the Three-State MCTHMM	37
Table 3.4 Model Comparison (1): Model Fit.....	46
Table 3.5 Model Comparison (2): Prediction Accuracy	47
Table 4.1 Summary Statistics of Data.....	57
Table 4.2 Regression of Participation Decision.....	58
Table 4.3 Coding of Challenge Types	60
Table 4.4 Summary of Notations	64
Table 4.5 Estimated Learning Parameters	66
Table 4.6 Estimated Factorized Challenge Effectiveness.....	67
Table 4.7 Estimated Variance-Covariance Matrix Σ_{β}	68
Table 4.8 Parameter Estimates in Utility Function	69
Table 5.1 Example of Healthcare Interventions	82
Table 5.2 Challenge Meta Feature Annotation	91
Table 5.3 Comparison of Baseline MABs	101
Table 5.4 Summary of Comparison Benchmark Recommendation Systems	103
Table 5.5 Comparison with State-of-the-Art Recommendation Systems.....	104
Table 5.6 Results for Dynamic Test Users	106
Table 5.7 Weight Loss as Recommendation Goal.....	111

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to my advisor, Professor Yong Tan. Prof. Tan has been a great mentor and role model for me. I am grateful for his continuous support and guidance through my Ph.D. study. His wisdom, humility, and enthusiasm for research are admirable and have greatly influenced me. His caring for students has been a great consolation, especially in the most difficult times during the COVID-19 pandemic. I feel fortunate to have Prof. Tan as my academic advisor.

My thanks also go to the other committee members, Professor Ming Fan, Professor Yingfei Wang, Professor Shi Chen, and Professor Quan Wen. Their insightful comments and constructive suggestions have helped me to improve this dissertation.

I would like to also express my appreciation to Shawna M. Reimers, Beau Kirkeby, Nuzulita Budhiari, Jaime Banaag, Jessica Aceves, and other staff at the Foster School of Business for their generous support during my study. My gratitude extends to all fellow doctoral students who have shared this journey and provide me with long-standing help.

Last but not the least, I sincerely thank my parents, Sen Zhou and Xiuling Wang, for their unconditional love and support. They are the permanent source of inspiration and courage for me. I also thank my boyfriend, Teng Zhang, who have been by my side for the past five years, giving me help and love. Without them I wouldn't achieve this far.

Chapter 1. INTRODUCTION

Personal health management and health-oriented care are becoming increasingly important themes in current healthcare systems. Different from disease-oriented care, which relies largely on healthcare professionals' expertise, personal health management focus on every day's behavior changes and relies on individuals' effort. Research have shown that individuals' engagement in health management can be crucial for wellness promotions and chronic disease management. It is argued that active participation in health management enables patients to make better healthcare decisions and potentially lead to better health outcomes (Borders et al. 2005; Hibbard et al. 2007). However, health management can be challenging for most individuals, and the reasons are multi-folds. First, health management is inherently difficult, as individuals need to consistently invest efforts in a long term to regulate their behaviors. Second, individuals may lack awareness or understanding of the association between behavior patterns and health outcomes (Wortley 2015).

To better support personal health management, many online healthcare communities (OHCs) are designed to provide individuals with health-related information and health-management tools. Researchers have argued that OHCs play an important role in personal health promotion and chronic disease management (Hamine et al. 2015). The 24/7 accessibility of online healthcare communities allows individuals to acquire information in a timely, cost-effective way (Yan and Tan 2014). In addition, OHCs bring together people with similar health interest or concerns and facilitate exchange of knowledge and support among community peers, which help individuals to stay engaged and motivated in the health-management process (van der Eijk et al. 2013).

While OHCs have attracted a large body of research, several challenges remain in understanding the functionalities of OHCs and facilitating users' continuous engagement. In this

dissertation, I take two different angles to answer the following questions progressively: 1) How should we understand users' dynamic participation patterns in OHCs? 2) How should platforms utilize knowledge on users' participation pattern to deliver more accurate service to individuals? The first question stands from individual perspective to rationalize their decision-making process, whereas the second question stands from service providers' perspective to come up with actionable, prescriptive support for individuals. Understanding these questions can be of important implications. First, it helps service providers to understand users' choice-making patterns so that they can improve platform design to better sustain users. Second, it helps policymakers to stand from individuals' perspective to study their objectives, preferences, and values, which supports the launch of patient-centered medicine (Sacristán 2013).

I use three separated but connected chapters to answer the above questions. In Chapter 3, I seek to understand individuals' self-regulation process in an online weight-loss community. Weight management has become a popular theme in many online healthcare communities, due to the widespread of obesity and overweight in developed countries. Online weight-loss communities provide various weight-management tools to help individuals to manage their body weight along with their behavioral routines. To comprehensively characterize individuals' dynamic engagement in different health aspects, I develop a multi-dimensional continuous-time hidden Markov model to structurally rationalize users' engagement process. Through empirical estimation, I find that individuals' may shift their engagement focus along time; particularly, through participating in social activities, individuals' focuses on behavioral routine can be strengthened whereas their focus on weight can be lessened. These findings manifest the effectiveness of online weight-loss communities, as they facilitate more self-management of behavior course, which has traditionally

taken the form of tedious diaries and were often neglected by individuals (Johnson and Wardle 2011; Saperstein et al. 2007).

In Chapter 4, I study individuals' dynamic engagement in online healthcare interventions. Online healthcare interventions are behavioral treatment programs or plans that help individuals to conduct lifestyle regulation. Examples include weight-loss challenges that users can find in an online weight-loss platform, which help users to restrict calorie intake and/or expenditure, or nutrition campaigns that instruct users to improve their dietary behaviors. I seek to understand how users dynamically make choices of intervention to participate in OHCs. I use a Bayesian learning framework to characterize users' learning behaviors in the participation process, which captures individuals' posterior beliefs on intervention effectiveness as a function of their prior beliefs, intervention selection histories, and observed health conditions. The analysis results provide evidence for individuals' learning capability and limitation and reveal users' preference structures on intervention features.

In Chapter 5, I propose a personalized recommendation framework that helps OHCs to provide tailored intervention suggestions to users. This research is motivated by our observation that OHCs typically provide a large variety of intervention options to users, which may hinder users from making effective health-management decisions as they may face choice-overload issues. The proposed recommendation framework adopts an online learning scheme to account for the fact that users' health management is dynamically evolving and their healthcare needs can be constantly changing. To better adapt the recommendation model to the healthcare context, I synthesize two innovative design components based on prominent health theories. The first component is a deep-learning-based feature engineering procedure, which is designed to learn crucial recommendation contexts in regard to users' sequential health histories, health-

management experiences, preferences, and intrinsic attributes of healthcare interventions. The second component is a diversity constraint, which structurally diversifies recommendations in different dimensions to provide well-rounded support for individuals' health management. Our evaluation results on massive simulation experiments demonstrate that each of the design components is effective and that our recommendation design outperforms a wide range of state-of-the-art recommendation systems, including sequential-based deep recommenders, context-aware recommendation systems, and content-based filtering algorithms.

Chapter 2. LITERATURE REVIEW

This dissertation is mainly concerned with three streams of literature: personal health management, online healthcare communities, and recommendation systems. In the following, I introduce the related theories and studies in each stream.

2.1 PERSONAL HEALTH MANAGEMENT

Individuals' health behaviors play a critical role in wellness promotion and disease prevention. According to W.H.O. (2016), improvement in lifestyles can help individuals to prevent or mitigate multiple chronic conditions, such as obesity, cardiovascular diseases, and type-2 diabetes. Unlike disease-oriented care performed by professionals, personal health management is primarily relied on individuals' efforts in everyday life. Typically, individuals need to continuously engage in behavior regulation in order to keep fit. Prior studies suggest that individuals' health behaviors need to be understood with respect to specific health and social context. It is argued that individuals may frequently adapt their health behaviors based on their personal health condition and health-management experiences, such as treatment compliance, self-monitoring, and healthcare-knowledge seeking (Johnson et al. 2002). In addition, the social environment may dynamically transform individuals' health behaviors by affecting mental well-being (King et al. 2006; Yan and Tan 2014). For instance, the exchange of emotional or informational support among peers may encourage optimism and self-esteem of individuals (DiMatteo 2004), which can help them to better comply with a treatment plan and make a behavior change (Johnson and Wardle 2011; Krukowski et al. 2008; Wang et al. 2012).

In the social-science literature, psychosocial theories have extended our understanding of how cognitive and social factors contribute to personal health, among which the social cognitive theory

(Bandura 1991; Bandura 1998) has long been recognized as a classic paradigm for understanding individuals' self-regulation behaviors, such as health management. The social cognitive theory proposes a personal-influence-based self-regulation model, in which individuals exert control over their motivation and behaviors to achieve better health outcomes. The theory suggests that individuals' self-regulation contains multi-dimensional effort. First, individuals need to set proper health goals to motivate themselves toward a desirable health outcome. Second, individuals need to operationalize their goals into actual behavioral aspects so that they can gain behavior-management skills and strategies to tackle challenges and fulfill expectations effectively.

In addition, the social cognitive theory suggests that individuals' self-regulation is dynamically affected by individuals' cognitive status. Specifically, the theory proposes a causal structure in which individuals' self-efficacy determines how efforts are made to execute the self-regulation loop. Under this framework, self-efficacy refers to a belief or perception of one's capability to exert control over the environment and his or her behavior. Individuals' strong self-efficacy means that they have a strong belief in their ability to make a behavior change; thus, they tend to be more confident and motivated to observe and manage their conditions and behaviors and to react more positively when facing a negative situation (Baum et al. 2012). Individuals' self-efficacy can be altered, in turn, by their behavior patterns. For example, people who observe their adherence to health behaviors generally have a high confidence level in following subsequent medical plans (i.e., high self-efficacy) (Gherman et al. 2011). In contrast, people who attend merely to health-management outcomes may get discouraged by temporary failures, which could damage their sense of efficacy (Bandura 1998).

Our work extends this stream of literature by investigating the role of technology in the process of health management. It has become clear that individuals' health management nowadays,

as assisted by modern technology, is more often affected by social environments. Social-media-based online healthcare communities, for instance, assist individuals in frequently exchanging health-management-related knowledge and experience with similar others. In addition, due to the idiosyncrasy of the Internet, healthcare interventions can be provided in a more flexible way through online channels to better satisfy users' healthcare needs. We characterize individuals' use of OHCs to manage their personal health and, thus, provide insights into the integration of healthcare and technology.

2.2 ONLINE HEALTHCARE COMMUNITIES

With the advent of the Internet and social media, online healthcare communities (OHCs) have become popular tools for individuals to manage personal health. OHCs are online platforms that gather users with similar health-management interests or concerns. To better support individuals' health management, many OHCs provide users with various self-regulation tools, such as self-monitoring functions and/or online behavior-intervention programs. It is widely acknowledged that online healthcare communities (OHCs) play an important role in chronic condition management. Compared to offline interventions that usually generate huge costs due to geographical and temporal restrictions, OHCs provide individuals with convenient access to healthcare resources (van der Eijk et al. 2013). In addition, OHCs provide users with social functions, such as discussion forums and social-networking tools, which help individuals to exchange health-management knowledge and experience with similar others.

In the literature, there is a large body of research on online healthcare platforms (also known as Health 2.0) and user participation in personal health management. It is found that online healthcare platforms can help empower individuals and encourage more health activities (Holman and Lorig 2000; van der Eijk et al. 2013). Researchers argued that online healthcare platforms

provide a channel for individuals to communicate with each other about how to discern and manage specific health problems so that they can become familiar with a health-management routine (Yan and Tan 2014). Social influence theories also suggest a potential for OHCs to affect individuals' health behaviors. Researchers found that social support can improve individuals' mental well-being (Clark 2005; Umberson and Karas Montez 2010) by encouraging optimism and self-esteem (DiMatteo 2004) and, thus, can potentially help individuals to comply with treatment and achieve better health outcomes (Johnson and Wardle 2011; Krukowski et al. 2008; Wang et al. 2012). On the other hand, there is evidence that suggests a negative effect of social support, which likely occurs when individuals experience stress or low self-esteem (Buunk and Gibbons 2007). It is found that when individuals receive more support than what they can offer to others, they tend to develop feelings such as weakness, incompetence, and indebtedness (Greenberg 1980; Newsom 1999), which further damages their self-efficacy (Yan 2018).

Our study builds upon these theories to capture the effect of online social environment on individuals' engagement in OHCs. We apply structural models to investigate individuals' different engagement mechanisms in OHCs, including multi-faceted engagement process, cognitive state transitions, and learning behaviors. As such, we contribute to the Health 2.0 literature by providing insights into user dynamic engagement patterns and platform sustainability.

2.3 HEALTHCARE RECOMMENDATIONS

As my third essay concerns personalized healthcare recommendations, in this section, I review existing literature on recommendation systems and discuss potential challenges in healthcare recommendations.

Recommendation systems are intelligence-based decision-making algorithms that can help users to filter information or product choices based on their own preferences or interests, especially

when there is information or product overload (Konstan and Riedl 2012; Vozalis and Margaritis 2003). During the last few decades, recommendation systems have garnered considerable attention from both academia and industry for their capability in delivering personalized services and generating benefits for service providers and customers (Isinkaye et al. 2015; Pathak et al. 2010; Pu et al. 2011). In the literature, a large body of research has focused on batch-learning-based recommendation systems, such as collaborative filtering, content-based filtering, and hybrid models (Zhang et al. 2019). These recommendation systems generally adopt a “first learn, then earn” recommendation scheme. That is, they first learn users’ preference patterns based on a series of historical data, and then they fully exploit the learned knowledge to make future recommendations. For example, collaborative filtering makes recommendations based on similarities in users’ item-selection histories (Adomavicius and Tuzhilin 2005; Sedhain et al. 2014), and content-based filtering leverages the content attributes of users’ previously selected items (Bieliková et al. 2012; Pon et al. 2007). Previous studies have proposed a variety of techniques to learn users’ preference patterns from historical data, such as context-aware recommendation systems (CARS) that model contextual dependency of users’ behaviors, and deep-learning-based techniques to learn latent user representations.

The “first learn, then earn” scheme, however, is based on the assumption that users’ preferences have a static pattern that can be well represented by the historical data (Adomavicius and Tuzhilin 2005; Sahoo et al. 2012). When users’ preferences are constantly changing, such recommendation methods may become less effective in adapting to individuals’ behavior dynamics, as it is likely that individuals’ preference patterns will not be fully captured by the data. In addition, prior studies have generally shown that the batch-learning-based models tend to over-specialize recommendations in the long run (Yu et al. 2009), as they tend to focus on well-known

items that already have accumulated adequate historical information, whereas the items with limited historical data will be overlooked (Fleder and Hosanagar 2009; Pariser 2011).

In addition to the batch-learning algorithms, researchers have proposed an online learning scheme, multi-armed bandit (MAB), which can be used to offer recommendations on an adaptive basis (Auer et al. 2002; Gittins 1979; Lai and Robbins 1985). Specifically, an MAB models a sequential decision making problem in which the underlying reward distribution for each action is unknown and data can be obtained in a sequential order to update posterior of the reward distribution. An MAB solves this problem by adaptively balancing between obtaining immediate rewards through exploiting the current best option and exploring potential better alternatives (i.e., the “exploitation-versus-exploration” tradeoff) (Kim and Lim 2015; Li et al. 2010; Tang et al. 2014). By doing so, an MAB aims to maximize the cumulative reward in the long run while improving knowledge of the environment; that is, *earning while learning* (Misra et al. 2019).

It is worth noting that the online learning scheme is in stark contrast with batch-learning algorithms. The former acts actively to collect data to learn the environment, with a forward-looking goal to maximize long-term rewards. In other words, online learning may deviate from the current “best” knowledge from time to time in exchange of a potential better learning performance and higher rewards collected in the future. In contrast, batch-learning algorithms fully exploit the current knowledge without exploring potential better opportunities that are not shown in the historical data. As such, they tend to interact with the environment in a passive and myopic manner and may not learn effectively when the environment contains many uncertainties that cannot be represented by current data. Researchers showed that MABs are suitable for tackling decision making problems in noisy and changing environment (Speekenbrink and Konstantinidis 2015). For example, Misra et al. (2019) applied MAB to a pricing problem in which the volume of demand

is uncertain. Schwartz et al. (2017) used MAB to improve advertising design when online advertisers cannot identify targeted users.

In the healthcare recommendation context, individuals' healthcare preferences can be constantly changing, as they need to adapt their health-management behaviors according to their evolving health conditions, health histories, and engagement experiences. The varying health preference patterns create the first challenge in healthcare recommendations, as the behavior variations may not be fully captured and observed by online service providers. Another challenge of healthcare recommendation is that individuals' healthcare needs can be diverse. Based on the social cognitive theory, individuals need to jointly manage their health outcomes and their behavior processes and thus form multiple health-management focuses. An effective healthcare recommendation system should provide well-rounded suggestions to users on each health-management aspect. We propose that MABs can serve to address these challenges in healthcare recommendations, as they can help service providers to effectively explore users' preference variations while improving users' online engagement during the process. In addition, through exploration, an MAB increases choice stochasticity and thus can better promote recommendation diversity (Qin et al. 2014) to satisfy users' diverse healthcare needs. In this dissertation, I propose a bandit-based recommendation system that is customized based on the healthcare context.

Chapter 3. MULTI-DIMENSIONAL PARTICIPATION IN ONLINE WEIGHT MANAGEMENT

Obesity has become one of the most prominent health problems in the world. According to the World Health Organization (2018), in the US, nearly 40% of adults aged 18 years and above were overweight in 2016, and 13% were obese. Due to its progressive nature, obesity is recognized as chronic conditions by the American Medical Association (American Medical Association 2013). To manage weight, individuals need to constantly self-regulate their daily dieting behaviors and exercise habits by paying close attention to their ongoing lifestyles (Franz 2001; Hall and Kahan 2018). Traditionally, such self-attention to lifestyles is gained through writing paper diaries about daily weight-loss routines, which is a tedious and distractive procedure that individuals usually find hard to engage in. With the advent of Internet and social media, individuals nowadays are able to self-monitor and self-manage their personal health in a more convenient way. One particular type of user-centric, self-help online healthcare platforms, online weight-loss communities (OWC), for example, allow individuals to digitally self-track their weight-loss history through integrated online databases. These communities bring together people who aim to lose or maintain weight and facilitate timely information sharing and support exchange among community members (Harvey-Berino et al. 2002; Harvey-Berino et al. 2004; Perri et al. 1993; Tate et al. 2003). It is found that OWCs can help individuals to stay motivated for daily weight control (Hwang et al. 2010) and increase adherence behaviors regarding diet and exercise activities (Arora et al. 2012; McKay et al. 2002).

There are various weight-management tools provided by OWCs to help individuals to manage multiple health aspects easily and concurrently. Examples include self-weighing tools for individuals to self-monitor weight condition and e-journals for individuals to self-manage weight-

management behavioral routines in regard to diet and/or physical activities. According to the social cognitive theory, both health conditions and health-management behaviors are important dimensions for self-regulation. Whereas the former provides individuals with goal-directed motivation, the latter enables individuals to self-analyze the correlation between a behavior and the corresponding outcome (Bandura 1991; Krukowski et al. 2008; Thomas et al. 2017). Prior research, however, has separately evaluated individuals' self-regulatory behaviors in OWCs from a partial dimension. For instance, some studies have focused on individual's self-tracking of weight-loss conditions (Thomas et al. 2017) while some studies have examined individual's engagement in improving a behavioral routine (Johnson and Wardle 2011). To date, how individuals jointly manage different self-regulatory dimensions in OWCs remains far from obvious.

In this study, we are motivated to investigate individuals' engagement in different self-regulatory dimensions in an OWC. Understanding this research problem is of great significance. First, as each self-regulatory dimension is an inseparable component for weight management, by studying this research problem we are able to characterize individuals' behavior patterns in OWCs in a more comprehensive manner. Second, an OWC provides a social environment in which individuals can frequently interact with each other to exchange resources such as information or support. Social influence theories suggest that peer interactions may affect individuals' self-regulatory dimensions differently. On the one hand, individuals may acquire additional weight-management skills through remote communications with community members. On the other hand, individuals may update their perceptions about their body weight conditions by comparing personal weight-loss experiences with those of others. Studying individuals' engagement in

different self-regulatory dimensions can, thus, help us to fully reveal the operation mechanisms of OWCs and the effectiveness of social media on individuals' weight management.

It is worth noting that there are several challenges in the examination of individuals' self-regulatory dimensions. First, according to the social cognitive theory (Bandura 1991; Bandura 1998), individuals' self-regulatory behaviors are dynamically adjusted through a cognitive process. As a result, how individuals perceive their weight-loss status in terms of weight-management performances and/or capabilities affect the amount of efforts being produced. Although individual's cognitive perception is unobserved, it is important to control such hidden condition when analyzing individuals' engagement patterns in each self-regulatory dimension. Second, individuals' multiple self-regulatory dimensions need to be jointly calibrated to capture the potential correlations between any two dimensions. To address these challenges, we propose a structural model, a multi-dimensional continuous-time hidden Markov model (MCTHMM). Our framework is based on the classic approach, hidden Markov model (HMM), which is capable to control for the unobserved stochastic process that is relevant to the generation of observed data. We take advantage of such a modeling structure to account for individuals' cognitive perceptions about their weight-loss status, which is unobservable to researchers and yet is suggested in theory to be an important determinant of individuals' self-regulatory behaviors. Different from a traditional HMM, our MCTHMM uses an outcome vector to describe individuals' weight-management behaviors, so that different self-regulatory dimensions can be measured at the same time. In addition, our model evaluates individuals' weight management using a continuous time frame, so that the timing process of engagement can be captured and individuals' behavior patterns can be analyzed in a more accurate way.

Our analysis uses data collected from one of the leading non-commercial OWCs in the US. The focal community enables users to record their weight changes and to keep digital journals to self-track their weight-management routines such as dietary behaviors, participation in physical activities, and/or general thoughts related to weight-loss. Users are also provided with social-interaction tools such as forum post and social-tie functions. We use the number of self-weighing activities to measure individuals' engagement in managing weight conditions and the number of journals to calibrate individuals' engagement in managing weight-loss behavioral routines. Through model estimation, we find that individuals tend to increase their journal-recording behaviors while decreasing self-weighing behaviors after they have participated in social activities in the OWC. Given that individuals usually have limited energy to invest for weight management, this result suggests that individuals may shift their focus from weight-loss outcomes (i.e., changes in weight) to weight-loss behavioral routines through social interactions.

This result also indicates that an incomplete characterization of self-regulatory dimensions would lead to biased understanding about individuals' online engagement. For example, if we only account for individuals' self-weighing behaviors while neglecting journal-recording behaviors, we would likely underestimate individuals' engagement in OWCs. We further show through model comparisons that investigating multiple self-regulatory dimensions jointly can actually help decision makers to better describe users' participation patterns in an OWC and to make more accurate predictions about users' future online weight-management behaviors.

Our analysis results further derive several insights regarding individuals' cognitive process for weight management. We show that social interactions among community members can play a significant role in shaping individuals' perceptions about their weight-loss status. In particular, we investigate two types of posting behaviors, i.e., initiating a post thread and replying to others'

posts. We find that the former can be more effective than the latter in helping users to boost their self-believes and to stay confident and motivated. As individuals can freely lead a discussion topic by initiating a post thread whereas by replying to a post individuals follow the discussion topics led by others, this finding indicates the necessity for individuals to proactively seek for social support to match their needs. We also find that individuals' participation frequency in OWCs can affect how quickly individuals change their self-believes. Finally, we investigate individuals' self-believes corresponding to two distinct self-regulation modes. In the first mode individuals regularly conduct self-weighing and journal-recording, whereas individuals only self-weigh regularly in the second mode. We find that individuals in the first self-regulation mode generally have a stronger self-believe compared to those in the second.

Our study contributes to the extant literature in the following ways. First, we contribute to the stream of literature on online healthcare communities. Although there has been a great body of research discussing users' engagement in online healthcare communities, to the best of our knowledge, this study is among the first to identify and theorize individuals' different self-regulatory dimensions in an OWC. Our empirical results demonstrate that individuals may dynamically adjust their focus from one self-regulatory dimension to another in the process of online participation. This finding provides evidence for individuals' joint management of multiple self-regulatory dimensions and how OWCs serve to address individuals' different needs for weight management. Second, we contribute to the social influence literature. Different from prior studies that have focused on the discrepancy between proactive and passive behaviors of individuals (Ballantine and Stephenson 2011; Crant 2000), we distinguish and analyze two proactive posting behaviors (i.e., initiating and replying a post). Our results show that although both posting behaviors are proactive, their effects can still be different, potentially due to the fact that

individuals play different roles in matching social support. As such, we demonstrate that the effect of social support can depend on individuals' actual engagement strategies, and, thus, enrich the mechanisms on social support.

Third, we contribute to the social cognitive literature by incorporating and integrating various behavior aspects in depicting individuals' health management. The social cognitive theory has been frequently used to describe individuals' self-regulatory behaviors such as health management. Little empirical research, however, has examined the mechanisms of self-regulation in healthcare contexts. In this study, we operationalize the social-cognitive construct in an online weight-management setting, and we extend the classic personal-influence-based self-regulation framework by infusing social influence to capture the impact of social media on individuals' weight-management behaviors. In this process, we also make methodological contribution by inventing a new structural model, i.e., MCTHMM. This model synthesizes innovative model structures into a traditional HMM, including the multi-dimensional outcome vector and the continuous time frame. These model structures can be used to improve characterization of individuals' health-management behaviors. Particularly, in managing a chronic condition, individuals usually need to simultaneously manage multiple self-regulatory dimensions, such as health outcomes, daily health behaviors, and/or medication. These dimensions can be jointly captured by our model.

From a practical perspective, our study offers significant implications as well. First, our study provides evidence for the effectiveness of OWCs. We find that the social-interaction tools provided by OWCs can help individuals to increase engagement in managing weight-loss behavioral routines, which is traditionally in the form of tedious diaries and often neglected by individuals (Johnson and Wardle 2011; Saperstein et al. 2007). Second, our study provides an

effective approach in analyzing individuals' online weight-management behaviors, which can be utilized to improve platform design and to better engage users in OWCs. In particular, as our study identifies different effects of online weight-loss tools across individuals depending on their participation styles, platforms can use these results as a guide for providing personalized recommendations to individuals to improve their weight-loss performances. For example, platforms can send notifications to inactive users, encouraging them to participate in more social activities, such as initiating discussion topics in a community forum. Moreover, our study shows that individuals' regular engagement in both self-weighing and journal recording can be desirable; thus, OWCs can launch new platform design to facilitate individuals' daily self-monitoring behaviors.

3.1 THEORETICAL FRAMEWORK

Weight management requires individuals to constantly self-regulate their behaviors such as diet and physical exercise to establish a healthy lifestyle. To establish a theoretical foundation for our analysis on individuals' online weight-management behaviors, we explore theories related to mechanisms on self-regulation. In the literature, the social cognitive theory has been widely used to characterize individuals' self-regulatory behaviors. This is because the social cognitive theory highlights personal influence in achieving a behavior change, such as how individuals cognize their goals and performances through the process of self-regulation. Thus, it provides the theoretical foundation for our study. We further extend and enhance our theoretical framework by introducing social influence to infuse the social impact into the traditional personal-influence-based self-regulation model.

In particular, we characterize individuals' online weight-management behaviors as a multifaceted dynamic process. According to the social cognitive theory on self-regulation, a

promotion for health behavior requires individuals to continuously self-observe and self-analyze their health conditions along with their behavior processes. Likewise, in weight management, individuals need to pay attention to both their weight conditions and their weight-management routines such as dietary behaviors and physical activities. OWCs provide various tools to help individuals manage these aspects, including self-weighing function, calorie counter, and/or digital weight-loss diaries. Therefore, we identify the multifaceted nature of individuals' online weight-management behaviors, and we theorize these aspects as different self-regulatory dimensions in our conceptual framework.

Core to the social cognitive theory, individuals' self-regulatory behaviors are inherently driven by their self-efficacy. That is, the stronger self-efficacy individuals perceive, the more confident and motivated individuals are to make a behavior change, and the more efforts tend to be produced for self-regulation (Baum et al. 2012). In our weight-management context, we consider that the mechanism of self-efficacy can operate on each of the self-regulatory dimensions, as these dimensions are all critical components of self-regulation for weight-loss, although they reflect individuals' efforts for making a behavior change from different perspectives.

Moreover, individuals' self-efficacy can be dynamically changed throughout their participation process in OWCs. This is theorized by the social influence theories that individuals can update their self-believes about their performances and capabilities through exchanging experiences and support with similar others (Buunk and Gibbons 2007; Gallant 2003; Yan 2018). Individuals' self-monitoring activities in OWCs can also help themselves to develop a better understanding about their weight-loss progress and thereby shift their self-efficacy (Bandura 1998; Gherman et al. 2011). The changes in self-efficacy will be further reflected in individuals' subsequent online weight-management behaviors.

Self-efficacy determines the baseline of individuals' self-regulatory behaviors. Conditional on self-efficacy, there may be other factors that affect individuals' self-regulatory dimensions differently. For example, individual-level characteristics may contribute to personal variations in weight-management patterns. The recent online activities may also exert a short-term impact on individuals' self-regulatory dimensions. According to social influence theories, social support from others can help individuals to comply with a treatment plan. Therefore, individuals may invest more efforts in managing their behavioral routines after participating in social activities in OWCs. At the same time, individuals may reduce their attention to weight-loss outcome, as the amount of effort that individuals can produce is usually limited. We include a set of control variables to capture these effects. Figure 3.1 is the graphic representation of our proposed conceptual framework.




Figure 3.1 Online Weight Management: A Multi-Dimensional Dynamic Process

3.2 MODELS

To model individuals' dynamic participations in OWCs, we need to consider at least two challenges that are imposed by the process of self-regulation as theorized in Section 3.1. First, as suggested by the social cognitive theory, self-efficacy is an important determinant of individuals' weight-management behaviors. It is unobservable, yet it is crucially associated with individuals' participation patterns in OWCs and, thus, needs to be controlled for. Second, individuals can

simultaneously manage multiple self-regulatory dimensions online, such as weight and weight-management behavioral routines. These dimensions need to be jointly calibrated by the model so that the potential correlations/dynamics between any two dimensions can be captured.

To address these challenges, we establish an MCTHMM. An HMM is able to capture the underlying stochastic process that determines the observed outcome sequence (Rabiner 1989). In the literature, HMMs have been extensively applied to pattern recognition problems, such as speech recognition and bioinformatics (Krogh et al. 1994; Schuller et al. 2003), to describe the evolution of observed events that depend on internal factors, i.e., hidden states, that are not directly observable. In this study, we extend an HMM to characterize individuals' dynamic self-regulation process in OWCs, which are inherently driven by their underlying cognitive status, self-efficacy. Our proposed MCTHMM augments a conventional HMM in two ways. First, as individuals' self-regulation in OWCs usually contains multiple dimensions and each dimension is fundamentally driven by the inherent self-efficacy, we allow the outcome variable to be a multi-dimensional vector so that individuals' engagement in different self-regulatory dimensions can be jointly captured. Second, we enhance a conventional HMM by allowing the stochastic process to follow a continuous-time evolution scheme.

The continuous-time HMM framework can improve the characterization of users' weight-management patterns from multiple perspectives. First, from a conceptual perspective, as individuals' self-management is naturally a continuous process, a continuous-time HMM can track the timing process of individuals' behaviors and, thus, depict individuals' behavior patterns in a more accurate way. Second, from a practical perspective, a continuous-time HMM can address the potential zero-inflation issue raised by using the conventional HMM in the estimation procedure. Specifically, in the conventional HMM framework, the underlying stochastic process is specified

as a discrete-time Markovian process. This modeling framework assumes that the emission of outcomes occurs at every discrete time unit; if there is no outcome observed, it is assumed that the outcome value is zero. In the procedure of model estimation, treating a null value as zero would cause zero inflation in the outcome distribution, especially when a small time unit is adopted, leading to potential bias in estimation results. A continuous-time HMM is able to address this issue by tracking the time between two outcome emissions. We provide a more detailed discussion on this in Appendix A.1. In the following, we first provide a brief discussion of the baseline methodology framework, continuous-time Markovian process, which helps to explain the modeling structure of an MCTHMM and its inference method. We then elaborate our multi-dimensional CTHMM (MCTHMM) approach and its inference method.

3.2.1 *The Framework of Continuous-Time Markovian Process*

The classic discrete-time Markovian process produces a sequence of random variables $\{X_n : n \in N\}$ at discretized time unit $n \in N$ from a countable state space $S = \{1, 2, \dots, K\}$. A continuous-time Markovian process, in contrast, considers a stochastic process that can spend a continuous amount of time at any state. That is, a continuous-time Markovian process, $\{X(t) : t \geq 0\}$, allows random variable X to be produced from S at any moment of time $t \in [0, +\infty)$. We use waiting time Δt to refer to the time spans between two adjacent state transitions.

The evolution of a continuous-time Markovian process obeys the Markov property; that is,

$$P\{X(t+s) = j \mid X(t) = i, X(u), u < t\} = P\{X(t+s) = j \mid X(t) = i\} = p_{ij}(s), \quad (3.1)$$

where $p_{ij}(s)$ represents the probability of transition from state i to state j within time interval s ($s \geq 0$) (assume the stochastic process to be homogeneous). This property suggests that the transition to the future state is dependent on the current state only (Bass 2011; Durrett 1999; Karlin

2014). Given that the current state is i at time t , the probability of transition to state j within an infinitesimal time interval h can be expressed as follows:

$$P\{X(t+h)=j \mid X(t)=i\} = \delta_{ij} + q_{ij}h + o(h), \quad h \rightarrow 0, \quad (3.2)$$

where δ_{ij} equals to 1 if $i=j$ and 0 otherwise. q_{ij} can be seen as a measure for the speed of transition (or, transition rate) from state i to state j .

Corresponding to the transition matrix that specifies a discrete-time Markov chain, the key quantities that determine a continuous-time Markov chain are the transition rates $q_{ij} (q_{ij} \geq 0, i \neq j)$.

Define transition rate matrix Q as follows:

$$Q = \begin{bmatrix} -q_0 & q_{01} & q_{02} & \cdots \\ q_{10} & -q_1 & q_{12} & \cdots \\ q_{20} & q_{21} & -q_2 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}, \quad (3.3)$$

where $q_i = \sum_{j \neq i} q_{ij} = -q_{ii}$. q_i is the rate of transition to all other states, that is,

$$P\{\text{a change of state in } (t, t+h)\} = q_i h + o(h) = 1 - p_{ii}(h). \quad (3.4)$$

Each continuous-time Markovian process contains an embedded discrete-time Markov chain. It is the sequence of state that $\{X(t)\}$ arrives in at each transition. Formally, let τ_n denote the time at which the n^{th} state transition occurs, the embedded discrete-time Markov chain can be defined as $\{X(\tau_n+)\} \triangleq \{X_n\}$. It can be shown that the transition probabilities of the imbedded discrete-time Markov chain can be calculated based on matrix Q in the following way,

$$P\{X(t)=j \mid X(s)=i, s < t\} = p_{ij} = \frac{q_{ij}}{q_i}, \quad \forall j \neq i. \quad (3.5)$$

In addition, it is shown that the waiting time at state i is exponentially distributed with rate q_i^{-1} .


The relationship established above helps us to identify a continuous-time Markovian process from its embedded discrete-time Markov chain. Specifically, a continuous-time Markovian process can be fully determined by an initial state distribution and a transition rate matrix Q , or, equivalently, an initial state distribution, a waiting time distribution, and a probability transition matrix P of the imbedded discrete-time Markov chain.

3.2.2 *MCTHMM Framework*

HMMs are generalized Markovian processes in which the stochastic states are not directly observable. The identification of HMMs relies on another sequence of observable outcomes that are produced by the hidden states (Rabiner 1989). Thus, an HMM contains two stochastic processes, one is the hidden Markovian process, and the other is the outcome generation process. Different from prior studies that assume the hidden Markovian process to be discrete time, in this study, we allow the underlying Markovian process to follow a continuous time frame. In other words, we assume that a sequence of states $\{X(t): t \geq 0\}$ is generated continuously from a finite space $S = \{1, 2, 3, \dots, K\}$. This state sequence is unobservable; however, it produces a sequence of observable outcomes $\{Y_{X(t)}\}$ based on the hidden states.

To identify a CTHMM, given the data about the outcomes $\{Y_{X(t)}\}$, the key challenge lies in the fact that we do not observe when hidden state transitions occur. Prior HMM studies have proposed two solutions for overcoming this challenge. The first solution is to assume that an outcome is produced if and only if a state transition occurs (Montgomery et al. 2004). This assumption, however, may not be feasible in practice as outcomes will also likely be produced even if individuals remain in the same hidden state. The second solution assumes that outcomes are produced at every discrete time unit (e.g., day, hour) within waiting time, so that the number

of outcome observations during waiting time is dependent on the length of waiting time (Yu 2010). However, it is possible that outcomes may not be produced at every time unit that predetermined by researchers. In our analysis, we generalize these assumptions on the outcome generation process by allowing an outcome to be produced at any time point during waiting time. As a demonstration, let Figure 3.2 denotes a simple observed outcome process from time t_1 to t_7 . While there is an observed outcome Y_{s_1} at t_1 and an observed outcome Y_{s_2} at t_2 , there is a transition of hidden state. In contrast, there is no hidden state change from t_2 to t_3 , and the state change happens at t_4 . Yet, at t_3 , there is an observed outcome Y_{s_2} that when an individual's hidden state remains unchanged. We propose an identification and estimation strategy for this framework, and present the details in Section 3.2.4.



Note: Circles represent hidden states, and squares denote the observed outcomes

Figure 3.2 A CTHMM Stochastic Process

In our research context, the outcome sequence $\{Y_{x(t)}\}$ corresponds to individuals' online weight-management behaviors. As individuals' self-regulation for weight-loss usually contains multiple dimensions, we allow the outcome variable to be a multi-dimensional vector. As such, we further extend a CTHMM to an MCTHMM. Based on the social cognitive theory, we identify individuals' self-efficacy for online weight management as the hidden state $\{X(t)\}$ that affects individuals' online weight-management behaviors. It is defined as individuals' perceptions about

their weight-management performances and/or capabilities by using an OWC, which will be referred to as *efficacy state* hereafter.

3.2.3 Model Specification

Similar to a continuous-time Markovian process, an MCTHMM can be fully characterized by an initial hidden state distribution ω , a state-transition matrix of the imbedded discrete-time Markovian process P , a waiting time distribution f , and a state-dependent outcome distribution Π . For the estimation purpose, we specify the major model components for an MCTHMM as follows.

The first major model component is state-dependent outcome. As noted, the state-dependent outcome in this study is a multi-dimensional vector that calibrates individuals' online weight-management behaviors. Formally, let $Y_{it} = (Y_{1it}, Y_{2it}, \dots, Y_{nit})$ be the multi-dimensional outcome vector for individual i generated at time t . The j -th element, Y_{jit} , represents the number of activities produced in self-regulatory dimension j . For instance, in the dimension of self-regulation for weight-loss outcome, we can observe the number of self-weighing activities conducted by individuals and use it to calibrate individuals' efforts invested in that dimension. We model each dimension as a Poisson distribution, that is,

$$\Pi(Y_{jit} | s_{it} = s) = e^{-\lambda_{jsit}} \cdot \lambda_{jsit}^{Y_{jit}} / Y_{jit}!, \quad j = 1, 2, \dots, n, \quad (3.6)$$

where λ_{jsit} can be further expressed as $\exp(\beta_{js} W_{it} + \varepsilon_{ji})$, in which W_{it} denotes a set of variables that can affect individual i 's outcome at time t depending on the efficacy state s , such as individuals' personal characteristics and the short-term impacts of recent online activities. β_{js} is the vector of coefficients, which is state-dependent. We use ε_{ji} to denote individual-level random effect in dimension j , which accounts for unobserved individual heterogeneity. Moreover, by modeling

the joint distribution of random coefficients, we are able to capture the potential correlations among different self-regulatory dimensions raised by individual characteristics.

The second major model component is waiting time. Waiting time refers to the time until the hidden state transition into a different value. That is, during waiting time, individuals' efficacy state remains unchanged. We use the following exponential distribution to model the waiting time distribution,

$$f(\Delta t) = \theta_{sit} \cdot e^{-\theta_{sit} \cdot \Delta t}. \quad (3.7)$$

Parameter θ_{sit} measures the rate for individual i to transit to a different state at time t , given that she is in state s . The average length of individual i 's waiting time in state s at time t is θ_{sit}^{-1} . Considering that transition rate can be potentially affected by other factors, we further model it by $\theta_{sit} = e^{\alpha_s \cdot D_{it}}$, where D_{it} captures a set of panel variables that may affect individual i 's transition rate at time t , and α_s is a vector of state-dependent coefficients.

The third major model component, the state-transition matrix for the embedded Markov chain, contains the probabilities with which individuals' efficacy state changes. In a continuous-time framework, the transition to the same state (i.e., no transition) is already captured by waiting time; thus, the diagonal of the state-transition matrix is left undefined (or, defined as zero). To model the transition probabilities among different efficacy states, we adopt an ordered logit model. Specifically, we assume that there are n efficacy states, with 1 denoting the lowest self-efficacy level and n denoting the highest self-efficacy level. An ordered logit model allows us to model the transition between two non-adjacent efficacy levels. Let $p(s, s')$ denote the probability of transition from state s to state s' ($s' \neq s$), and

$$p_{it}(s, s') = \begin{cases} \frac{e^{\mu_{s \rightarrow s'+1} - \rho_s R_{it}}}{1 + e^{\mu_{s \rightarrow s'+1} - \rho_s R_{it}}} - \frac{e^{\mu_{s \rightarrow s'} - \rho_s R_{it}}}{1 + e^{\mu_{s \rightarrow s'} - \rho_s R_{it}}}, & s' \neq s-1, s \\ \frac{e^{\mu_{s \rightarrow s'+2} - \rho_s R_{it}}}{1 + e^{\mu_{s \rightarrow s'+2} - \rho_s R_{it}}} - \frac{e^{\mu_{s \rightarrow s'} - \rho_s R_{it}}}{1 + e^{\mu_{s \rightarrow s'} - \rho_s R_{it}}}, & s' = s-1 \end{cases}, \quad (3.8)$$

where $\mu_{s \rightarrow s'}$ defines the threshold for transition from state s to state s' . Due to the order of the hidden states, the thresholds should follow an order of $-\infty \triangleq \mu_{s \rightarrow 1} < \mu_{s \rightarrow 2} < \dots < \mu_{s \rightarrow n} < \mu_{s \rightarrow n+1} \triangleq +\infty$. R_{it} contains the covariates that may affect transitions in individuals' efficacy state, such as those to measure social influence. ρ_s is a vector of coefficients to be estimated. If the propensity value $\rho_s R_{it}$ exceeds or falls below a threshold needed for a transition to a higher or lower state, individuals' efficacy state will transition accordingly; otherwise, individuals will remain in the same state (Kim and Krishnan 2019). Obviously, $0 \leq p_{it}(s, s') \leq 1$, $\forall s, s'$. For each state s , $\sum_{s' \neq s} p_{it}(s, s') = 1$. That is, given that a transition to a different state happens, the probability of transitioning to all other states sums to one.

3.2.4 Model Identification and Estimation

Estimation for an MCTHMM can be achieved through likelihood maximization. The likelihood function, L , should consist the probability density for the initial hidden state, state transition, waiting time, and state-dependent outcome. The biggest difference between estimating an MCTHMM and estimating a standard discrete-time HMM lies in the inference of waiting time distribution. In an MCTHMM, waiting time is unobservable. We need to develop an inference method that enables us to identify waiting time from individuals' observed behavior sequence. Suppose that individuals' behavior outcomes arrive at $T_1, T_2, T_3, \dots$. A direct inference of waiting time using the time interval between two adjacent behavior outcomes, i.e., $\Delta T = T_k - T_{k-1}$, can be biased, as we assume that an outcome may be produced during a waiting time. In other words, a

state transition may or may not occur at an outcome observation, and, thus, the actual waiting time length may be longer than ΔT . To revamp this bias, we adjust the likelihood in the following way.

Specifically, given that a behavior outcome is observed, if a state transition occurs, the waiting time is then the time interval between two adjacent behavior outcomes (ΔT), and we calculate the likelihood using the PDF of the waiting time distribution; if the state remains unchanged, meaning that the waiting time surpasses ΔT , we calculate the complementary CDF of the waiting time distribution. Particularly, the likelihood for individual i conditional on the random effect can be expressed as follows:

$$L_i(Y_i | \varepsilon_i) = \sum_{x(T_1)=1}^n \sum_{x(T_2)=1}^n \cdots \sum_{x(T_{K_i})=1}^n \omega(x(T_1)) \prod_{k=2}^{K_i} p_{iT_k}(x(T_{k-1}), x(T_k)) Z(Y_{iT_k} | x(T_k)) f(T_k - T_{k-1}), \quad (3.9)$$

where T_k denotes the time point at which the k -th outcome is observed; $x(T_k)$ is the corresponding hidden state at time T_k ; K_i represents the number of outcomes observed for individual i , and it can be different across individuals. We define $p_{it}(s, s) \triangleq \theta_{sit}^{-1}$, so that we can automatically adjust for the potential inference bias as mentioned above. Function $Z(\cdot)$ is a product of probability density for each outcome dimension, i.e., $Z(Y_{iT_k} | x(T_k)) = \prod_{j=1}^n \Pi(Y_{jiT_k} | x(T_k))$.

With the conditional likelihood defined as above, the unconditional likelihood for individual i can be calculated by integrating over the random effect variable ε_i ; that is,

$$L_i(Y_i) = \int_{\varepsilon_i} L_i(Y_i | \varepsilon_i) dG(\varepsilon_i), \quad (3.10)$$

where $G(\varepsilon)$ is the joint distribution of the random coefficient and can be estimated non-parametrically. Finally, the overall likelihood can be calculated as $L(Y) = \prod_i L_i(Y_i)$.

3.3 CONTEXT AND DATA

3.3.1 *Focal Online Weight-Loss Community*

In this study, we investigate a leading non-commercial OWC in the United States. The focal OWC is intended for people who have overweight or obesity issues and want to lose weight or maintain a healthy weight. Users can have access to this community by visiting its website on computers. To join the community, an individual needs to register and create a personal profile, in which weight or weight-related information such as calorie intake/expenditure can be recorded. By continuously updating the personal profile, individuals can have a better knowledge about their weight-management performance and analyze the desirable weight-loss behaviors. In general, there are two self-monitoring tools to help users self-observe and self-analyze their weight-management status: weigh-ins and journals. Weigh-ins allow individuals to record their weights with time stamps. Compared with weigh-ins, journals in the focal OWC provide more flexibility in the content that individuals report. Individuals can record detailed information about their weight-management progress such as diet, physical exercise, thoughts, and plans. Thus, writing journals can help individuals to get familiar with a weight-management routine.

In addition, the focal platform provides a community forum and social networking tools for individuals to interact with similar others. In the forum, individuals can discuss weight-loss-related knowledge or experience through publishing posts. Individuals can either initiate a new post thread, leading the ongoing topics that they are particularly interested in, or reply to others' existing posts, providing personal thoughts or support. Different from forum posts, the social networking tools, such as friends and supporters, enable individuals to exchange social support through social ties. The establishment of friendship is based on mutual consent. Users can interact with their friends by discussing weight-loss experience and/or weight-loss-related knowledge in a private

chat room. Supporters, in contrast, only involve actions from one side. Gaining supporters can emotionally motivate individuals; however, this emotional support is passively received, as individuals cannot determine how many supporters they gain or who becomes a supporter.

Moreover, individuals can participate in weight-loss challenges on the platform, which are structured programs that help individuals to focus on specific weight-loss goals in a short time period. For instance, individuals can join challenge with goals on dietary habits (e.g., eat at least one salad daily), participation in physical activities (e.g., run or jog 30 minutes every day), or monitor weight changes (e.g., losing 20 lbs. in 2 months). Individuals can obtain challenge information from the website, such as challenge target, instructions, starting date, and duration. Usually a challenge lasts for a couple of weeks. Any individual can choose to participate in a challenge before its starting date. By participating in weight-loss challenges, individuals aim to break down their long-term goals into smaller pieces and to stay motivated in the process of weight management.

3.3.2 *Data*

From the focal platform, we collected data about individuals' online activities on a daily basis, including the number of weigh-ins and journal-recording behaviors that individuals documented, the number of challenges that individuals participated in, the number of posts that individuals published in online forums, and the number of friends and supporters that individuals gained. We also collected data about individuals' gender and membership duration. User descriptive statistics are provided in Appendix A.2. Following the convention of previous behavioral weight-management studies (Baker and Kirschenbaum 1993; Boutelle et al. 1999; Helsel et al. 2007; Tate et al. 2001), which typically investigate individuals' behavior in a period of time ranging from 2 weeks to 6 months, we collected 4 months of data from January 2014 to April 2014.

Users' engagement in self-regulation on the focal platform is mainly reflected in two activities, weigh-ins and journals. To operationalize individuals' self-regulatory dimensions, we construct a two-dimensional outcome variable that consists of each users' daily number of weigh-ins and daily number of journals. Specifically, we use the daily number of weigh-ins to measure individuals' engagement in self-regulating their weight-loss outcome; this is because the weigh-in function can facilitate users' self-attention to their weight variations. We use the daily number of journals to measure users' efforts toward self-regulating their behavioral routines, as the users use the journal function to reflect on and record their behavior details in a weight-management process. Together, these two dimensions provide a summary of individuals' efforts in regard to conducting self-regulation in a comprehensive manner.

Apart from the state-dependent outcome, three sets of covariates are constructed. The first set of covariates are those that affect state transitions, i.e., R_{it} in equation (3.8). Prior studies suggest that individuals' self-efficacy can be dynamically shaped by their personal observations of their past self-regulation patterns (Bandura 1998; Gherman et al. 2011) and social influences (Leahey et al. 2011; Miller and DiMatteo 2013). We measure the effect of self-observation through individuals' self-monitoring activities on the platform, i.e., weigh-in and journal recording, which reflect individuals' attention paid to their previous weight-management processes. To capture the social influences brought by the OWC, we construct variables to inclusively measure individuals' engagement in different types of social activities on the platform, including posts that are initiated by users and posts that are replied to by users, friends, and supporters. The first two covariates measure the amount of social support that is exchanged through forum posts. In particular, the number of posts that are initiated by users is able to capture individuals' dominant role in leading post topics, whereas the number of posts that are replied to by users is a proxy for individuals'

following behavior, as the ongoing post topics are led by others. Distinguishing between these two posting behaviors allows us to identify the mechanisms of different support-exchange strategies. The last two covariates measure the amount of social support that is exchanged through social ties. Because a friendship is established by mutual consent, we use it to calibrate support availability on the platform (Shumaker and Brownell 1984; Yan 2018). In contrast, supporters are one-directional relationships that do not require mutual consent; thus, we use them to capture the amount of social support that is passively received. Finally, according to Bandura (1991), individuals' self-efficacy is continuously shaped with temporal proximity; that is, events that happen closer in time tend to have a stronger impact whereas distal events have a less salient effect. Thus, we calibrate the above activities by their cumulative number to capture the enduring impact on self-efficacy and we apply exponential smoothing to account for the time-discounted effects. Specifically, for the activity that occurs Δt time unit ago, we discount its effect by $e^{-r \cdot \Delta t}$, where r is the discounting factor, which needs to be optimized.

The second set of covariates corresponds to W_{it} in equation (3.6), which are the control variables that cast immediate effects on state-dependent outcomes. As noted, given an efficacy state, individuals' engagement in self-regulation can be affected or shocked by other factors, such as personal characteristics and direct impacts of recent online activities. To capture these immediate effects, we first include two covariates to capture individuals' personal characteristics: gender and membership duration. These two variables are generally shown to be influential on individuals' participation patterns in online healthcare communities (Gooden and Winefield 2007; Yan and Tan 2014; Yan and Tan 2017). In addition, five variables are constructed to capture the short-term shocks of individuals' recent activities: the number of posts initiated, number of posts replied to, number of new friends, number of new supporters since the last outcome observation,

and the number of weight-loss challenges that a user has engaged in during the most recent week. The first four covariates are able to capture the amount of social support that has been exchanged recently, while the fifth covariate measures individuals' short-term goal-setting behaviors. Again, we apply exponential smoothing on these variables to adjust for temporal proximity.

Finally, the third set of covariates contains the effects on state-dependent waiting time, that is, D_{it} in equation (3.7). We consider that individuals' participation frequency in OWCs might be an influential factor for the average length of waiting time. A direct intuition is that individuals who frequently participate in online activities are more likely to be affected by the online environment and, consequently, may shift to other states more quickly. To control the effect of participation frequency on waiting time at time t , we calculate individuals' average participation frequency until time t for both social activities and self-monitoring activities. In Table 3.1, we provide variable descriptions and their statistics.

Table 3.1 Variable Description and Summary Statistics

Variable	Description	Mean	Std. Dev.	Min	Max
Variables that Affect State-dependent Outcome					
Gender	Female = 1, Male = 0.	0.52	0.50	0	1
Duration	The membership duration since registration (in days)	342.90	161.18	22	678
Challenge	The no. of challenges undertaken by individuals during the most recent week	1.93	1.67	0	13
Post_init_rec	The no. of posts initiated by individuals since last outcome occurrence	0.14	0.39	0	4
Post_re_rec	The no. of posts replied to by individuals since last outcome occurrence	0.15	0.40	0	4
Friends_rec	The no. of newly-made friends since last outcome occurrence	0.55	0.89	0	16
Supporters_rec	The no. of new supporters since last outcome occurrence	0.58	0.98	0	13
Variables that Affect State Transitions					
Post_init_cum	The cumulative no. of posts initiated by individuals	2.20	2.06	0	18
Post_re_cum	The cumulative no. of posts replied to by individuals	2.25	2.13	0	22
Friends_cum	The cumulative no. of friends	8.12	6.57	0	60
Supporters_cum	The cumulative no. of supporters	8.94	7.38	0	66
Weigh-in_cum	The cumulative no. of weigh-ins	8.63	7.12	0	55
Journal_cum	The cumulative no. of journals	8.94	6.76	0	66
Variables that Affect State-dependent Waiting Time					
Freq_monitor	Individuals' average engagement frequency for self-monitoring activities till time t	0.36	0.17	0.07	3
Freq_social	Individuals' average engagement frequency for social activities till time t	0.38	0.15	0	3

3.4 EMPIRICAL RESULTS AND FINDINGS

As noted, we obtain the model estimates by maximizing the likelihood function specified in Section 3.2.4. We randomly sampled 300 individuals from the dataset to conduct our analysis for the purpose of improving computation efficiency. Robustness checks are performed on the sample construction. The initial hidden state distribution can be estimated by a latent class model. Before estimating the full model, we conduct a one-dimensional search to determine the optimal time discount factor. This method is based on evaluating the likelihood function at different points within a certain interval, which can provide us with a decent estimate while lowering the computational burden. The optimal discounting factor that we obtain is 0.2. We report these details

in Appendix A.3. In the following, we first report the model selection results and the model estimates. We then discuss our main findings derived from the MCTHMM framework, including dynamics associated with individuals' self-regulatory dimensions and dynamics associated with efficacy-state transitions. Finally, we conduct several robustness checks on our main findings.

3.4.1 *Model Selection and Model Estimates*

Because we cannot observe individuals' efficacy states directly, we run a set of models with different numbers of efficacy states to determine the optimal state number based on certain model selection criteria. In particular, we consider the Akaike information criterion (AIC) and the Bayesian information criterion (BIC). Both AIC and BIC are commonly used model selection criteria that account for model fit as well as the number of parameters to prevent overfitting. According to the results presented in Table 3.2, the optimal number of hidden states is three. The three states correspond to low (L), medium (M), and high (H) efficacy levels for online weight management.

Table 3.2 Selection on the Number of Hidden States

Number of States	Log-likelihood	No. of Parameters	AIC	BIC
$n = 1$	-15654.95	16	31341.90	31474.14
$n = 2$	-15335.42	38	30746.84	31060.90
$n = 3$	-15020.44	78	30196.88	30841.54
$n = 4$	-16011.21	108	32238.42	33131.03

Table 3.3 presents the estimation results for a three-state model. As can be seen, the estimates for state-dependent outcomes are presented in two separate columns in each state, with each column displaying estimates for one particular self-regulatory dimension. To prevent the identification problem of label switching (Jasra et al. 2005; Netzer et al. 2008; Yan 2018), we added constraints on the constant terms for both dimensions of the outcome distribution; that is,

$\beta_{0jL} < \beta_{0jM} < \beta_{0jH}$, $j=1,2$. This constraint assumes that, without other influences, a higher self-efficacy level tends to produce more weigh-ins or journal-recordings.

Table 3.3 Estimates for the Three-State MCTHMM

Parameters	State 1 (L)		State 2 (M)		State 3 (H)	
Variables Affecting State-Dependent Outcome						
	Weigh-in	Journal	Weigh-in	Journal	Weigh-in	Journal
Gender	-2.1210 (2.7722)	3.2190 * (1.9033)	1.6992 (1.5257)	-3.5121 * (1.9039)	-0.5090 (1.5408)	-0.3075 (1.2146)
Duration	-1.5102 (2.3955)	3.2176 * (1.8472)	3.0201 (2.1658)	3.1969 * (1.8709)	-1.1386 (1.3950)	2.2579 * (1.3138)
Challenge	-4.1674 (4.1245)	4.2957 * (2.3484)	4.8599 ** (2.4057)	-4.2863 ** (2.1840)	0.5497 (1.8363)	0.9084 (1.4472)
Post_init_rec	1.7233 (2.7545)	4.7871 ** (2.3758)	-4.5581 * (2.7504)	3.6018 * (2.0076)	-0.8377 (1.8262)	0.7198 (1.1278)
Post_re_rec	0.4551 (2.4133)	-1.6322 (1.7259)	-6.3258 * (3.8369)	6.3360 * (3.3461)	-3.9083 * (2.2285)	2.6684 * (1.3998)
Friend_rec	-1.4224 (1.3654)	-0.1189 (1.7623)	-0.9738 (1.2801)	1.5382 (1.3459)	1.0664 (1.5708)	-0.8014 (1.5009)
Supporter_rec	-7.9549 *** (3.0332)	12.2663 ** (6.0148)	-23.0426 * (12.2373)	25.5467 ** (13.0224)	-16.8050 * (9.5764)	18.0086 * (9.7945)
Constant	-0.6558 *** (0.1736)	-0.3773 (1.0062)	-0.6022 (5.2540)	-0.3206 (4.0141)	-0.2027 (5.0607)	-0.3117 (2.2054)
Variables Affecting State Transition						
Post_init_cum	10.6556 ** (4.3429)		9.5925 ** (4.5876)		-0.7617 (1.2313)	
Post_re_cum	-10.0809 ** (4.6225)		9.2963 * (5.0236)		-0.7433 (1.1199)	
Friend_cum	-8.5900 (6.0917)		-16.3472 *** (5.0632)		4.2385 (3.4471)	
Supporter_cum	12.8867 *** (4.3154)		-4.2921 (3.9516)		-6.3624 *** (2.1134)	
Weighin_cum	5.5579 ** (2.4648)		3.2244 (2.9086)		-0.6762 (1.0832)	
Journal_cum	4.4346 ** (2.1617)		3.9706 ** (2.0149)		-0.2730 (1.4310)	
Thresholds						
μ_1 (L->H)	18.6753 ** (8.1912)					
μ_3 (M->H)			16.2998 * (8.4318)			
μ_3 (H->M)					-1.3338 (1.3603)	
Variables Affecting Waiting Time						
Freq_social	6.9825 (6.3119)		2.5203 (7.5597)		-3.5819 * (2.1064)	
Freq_monitor	18.8466 * (11.4834)		26.2941 * (13.6510)		28.7292 (22.1341)	
Constant	-2.1326 (1.7229)		-2.0826 * (1.2152)		-2.0156 ** (0.8649)	
Individual-level Unobserved Heterogeneity						
Normalizing constants: $C_{\varepsilon_{1i}} = 0.0017$, $C_{\varepsilon_{2i}} = 0.0103$						
	$\varepsilon_{2i}^1 = 0$		$\varepsilon_{2i}^2 = 0.5854$		$\varepsilon_{2i}^3 = 1$	
$\varepsilon_{1i}^1 = 0$	0.0178		0.7094		0.1962	
$\varepsilon_{1i}^2 = 1$	0.0610		0.0122		0.0034	

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Note: Standard errors are reported in parentheses. We took logarithms on all the variables to address the over-dispersion issue.

The estimates of thresholds set the baseline value for a state transition to occur. As the state-transition probabilities in an MCTHMM account only for the scenario that individuals' efficacy state changes to a different value, we set only one threshold for each state in a three-state model. For example, in state L, the one threshold corresponds to the baseline value that an individual needs to achieve to arrive at state H; if the threshold is not met at the end of the waiting time, the individual's efficacy will shift to state M.

In the last two rows of Table 3.3, we report the estimated joint distribution of the random-effect variables. The number of supports for each random-effect variable is selected by model fit. We use superscript to label the supports and subscript to denote the distribution dimension. For example, ε_{1i}^2 denotes the second support for the random-effect variable controlled in the first outcome dimension. The parameter estimates suggest that the marginal distribution of each dimension is mutually dependent, indicating that our model captures the correlation between the two self-regulatory dimensions.

3.4.2 *Dynamics Associated with Self-Regulatory Dimensions*

The parameter estimates for state-dependent outcomes in Table 3.3 illustrate the dynamics associated with individuals' different self-regulatory dimensions. One of the key findings that we derive is that individuals who have recently participated in social activities on the platform tend to increase their journal-recording behaviors while decreasing self-weighing behaviors, regardless of what efficacy state they have. For instance, our results show that the number of posts that are initiated by users has a significant negative effect on the daily number of weigh-ins ($-4.5581, p < 0.1$) but a significant positive effect on the daily number of journals ($3.6018, p < 0.1$) for users who are in an M efficacy state; for users in an L efficacy state, the effect on the daily number of journals remains significantly positive ($4.7871, p < 0.05$) but the effect on the daily number of

weigh-ins becomes insignificant; for users in an H state, both effects are insignificant. Similar patterns (i.e., negative effect on weigh-ins and positive effect on journals) persist for other social activity-variables, such as the number of posts that are replied to and the number of supporters.

Because individuals typically expend limited effort toward managing their weight-loss progress, these results indicate that individuals may shift their focus from weight-loss outcomes to weight-management behavioral routines through participating in social activities in an OWC. As such, our study provides evidence that the dynamics of individuals' self-regulation can be continuously shaped by the online social environment. The results are consistent with our observations that users typically discuss their weight-loss experience and share their weight-management skills with each other through social interactions. According to the social comparison theories (Buunk and Gibbons 2007; Festinger 1954; Suls 2000), this process can help users to establish a reference for their own weight-management performance, with which users can reason through their own behavior patterns. For instance, through communicating with similar others, individuals can gain knowledge about others' weight-loss experience, on which they are likely to reflect in regard to their own behavioral routines.

Our results also show that gender, membership duration, and challenge participation can generate direct impacts on users' engagement in self-weighing and journal recording and that the effects may vary, depending on the state. Specifically, female users who are in an L state tend to put more effort into journal recording than do males in the same state (3.2190, $p < 0.1$). In contrast, male users tend to write more weight-loss journal entries than do females when the users' efficacy state is M (-3.5121, $p < 0.1$). In addition, users tend to increase their journal-recording behaviors as their membership duration grows, regardless of which state they are in (3.2176, $p < 0.1$ for L; 3.1969, $p < 0.1$ for M; 2.2579, $p < 0.1$ for H). Finally, challenge participation tends to increase the

number of weigh-ins (4.8599, $p < 0.05$) while decreasing the number of journal-recording behaviors (-4.2863, $p < 0.05$) when individuals are in an M state; for individuals in an L efficacy state, challenge participation has a positive effect on journaling behaviors (4.2957, $p < 0.1$). These results further demonstrate unique dynamics for each self-regulatory dimension.

Collectively, the above findings illustrate how self-weighing and journal recording jointly describe individuals' self-regulation dynamics. Failing to fully capture these behavioral dimensions would likely result in a biased understanding of users' online engagement. For example, if we consider only individuals' self-weighing behaviors but neglect journaling behaviors, it is likely that we would notice a decrease in users' engagement in self-monitoring after users have participated in social activities on the platform. In Section 3.5, we present the results of several additional analyses to further illustrate the importance of the inclusion of multiple self-regulatory dimensions.

3.4.3 *Dynamics of Efficacy State*

In this section, we introduce our estimation results on the efficacy state dynamics. These results correspond to the parameter estimates in the waiting time distribution and the state-transition matrix in our proposed model. In particular, the parameter estimates for waiting time describe the influences on the average time length that an individual stays in an efficacy state. We find that individuals' average participation frequency for social activities has a negative effect on the transition rate when individuals are in an H efficacy state (-3.5819, $p < 0.1$). That is, when users are at a high self-efficacy level, the more frequently they participate in social activities, the longer they tend to stay in the H state. Individuals' average participation frequency for self-monitoring activities, in contrast, tends to positively affect the transition rate when individuals' efficacy state is L or M (18.8466, $p < 0.1$ for L; 26.2941, $p < 0.1$ for M), meaning that the more frequently users

engage in self-monitoring activities, the more quickly they will transition to other states. It is worth noting that a quicker transition, or a shorter waiting time, is always desirable for an L-state user; however, it may not always be preferable for an M-state user, as that user may likely transition to a lower state.

The parameter estimates for state transitions describe the effects on individuals' transition to a different state. In general, we find that all types of social interactions on the platform can generate an impact on individuals' self-efficacy, although the direction/magnitude of the impact may vary, depending on the efficacy state. Specifically, we find that the cumulative number of posts that are initiated by individuals has a positive effect on state transition when individuals' efficacy state is L or M (10.6556, $p < 0.05$ for L; 9.5925, $p < 0.05$ for M). In contrast, the cumulative number of posts that are replied to can negatively affect state transition when individuals are in an L state (-10.0809, $p < 0.05$) while positively affecting state transition when individuals are in an M state (9.2963, $p < 0.1$). These differences in results can be potentially explained by the support exchange procedure of initiating posts and replying to posts. When individuals reply to others' posts, they react to others' opinions and follow others' leads in ongoing topics, whereas when individuals initiate a post thread, they take a dominant role in deciding the ongoing discussion topics so that they can match their interests or needs more effectively. Thus, these results indicate that individuals' actual support-exchange strategies can influence the effectiveness of social support and that it is important for individuals to match social support with their needs, especially for users who are in an L efficacy state. Users who are in an L efficacy state do not have much confidence in themselves; thus, they usually need strong support from the social environment to help them manage difficulties. When they cannot effectively obtain the support they need, they may develop more negative feelings about their performance and further reduce their self-efficacy.

We also find that the cumulative number of friends can negatively affect state transition when users are in an M state ($-16.3472, p < 0.01$) and that it has an insignificant effect when individuals' efficacy state is L or H. These results suggest that when individuals have a medium level of self-efficacy, the more friends they make, the more likely that they transition to a lower efficacy state. The negative effect may be due to the peer pressure embedded in friendships (Umberson and Karas Montez 2010). That is, after observing the weight-loss effort made by their friends, individuals may lower their perception about their capability. Our results suggest that users are subjected to such peer pressure when their self-efficacy level is moderate, possibly because users in this state have more confidence than the low-efficacy users and are thus more likely to compare their capability with others, increasing their tendency to be affected by peers; at the same time, they do not have as strong a belief about themselves to keep motivated as the high-efficacy users do. These results are consistent with the social comparison theories (Buunk and Gibbons 2007).

In addition, we find that the cumulative number of supporters can positively affect state transition when users are in an L state ($12.8867, p < 0.01$) and negatively affect state transition when users are in an H state ($-6.3624, p < 0.01$), whereas the effect of supporter number is insignificant for users in an M state. Supporters represent the amount of emotional support that users passively receive. Our results suggest that such emotional support can significantly encourage users when their self-efficacy level is low, as it can help the users to build self-esteem and improve mental wellbeing; however, it can have a negative effect for users at a high self-efficacy level, as those users do not need much emotional support from others to stay confident, and the excessive support they gain may potentially lead to feelings of indebtedness and/or incompetence, harming their self-efficacy (Greenberg 1980; Newsom 1999). Finally, our results suggest that the cumulative number of self-monitoring activities, such as weigh-ins and journal-

recordings, can have a positive effect on state transition when individuals' efficacy state is L or M (weigh-in: 5.5579, $p < 0.05$ for L; journal: 4.4346, $p < 0.05$ for L and 3.9706, $p < 0.05$ for M). These results indicate that individuals can improve their confidence by continuously observing their engagement in self-management when they do not have strong self-efficacy. These findings are in line with the literature on health management (Bandura 1998; Gallant 2003; Yan 2018).

3.4.4 Robustness Checks

We make several efforts to ensure the robustness of our results. First, we examine an alternative distribution for the state-dependent outcome. Instead of the Poisson distribution that we deploy in the main model, we consider a zero-inflated Poisson to examine whether there is an additional zero-generating process that needs to be captured. Specifically, we add two inflation parameters (i.e., ω_j , $j = 1, 2$) to our model:

$$\Pi(Y_{jit} | s_{it} = s) = \omega_j I_{[Y_{jit}=0]} + (1 - \omega_j) \cdot e^{-\lambda_{jsit}} \cdot \lambda_{jsit}^{Y_{jit}} / Y_{jit}!, \quad (3.11)$$

where $I_{[Y_{jit}=0]}$ is an indicator function that equals 1 when $Y_{jit} = 0$. Our estimates for the “inflated probabilities,” that is, ω_j ($j = 1, 2$), are not statistically different from 0, indicating that there is no zero inflation in our data; thus, a Poisson distribution is sufficient to capture the data variation pattern.

We then examine whether different variable measures can affect our results. In particular, we investigate measures for individuals' participation frequency. In the main model, we calculate individuals' average participation frequency until time t for both social activities and self-monitoring activities to capture the effect of participation intensity on the average length of waiting time. As a robustness check, we explore a different measure: individuals' average participation frequency during the most recent week, considering that the time for individuals to stay in a state

may be affected only by their recent engagement intensity. We replace *Freq_social* and *Freq_monitor* by the new measures and re-estimate the model. Our estimation results remain statistically consistent.

Finally, as noted, we randomly sampled 300 individuals from the full dataset to conduct our analysis. To examine whether our findings are robust to the sample selection, we resampled data multiple times by using different sampling seeds to reconduct our analysis. Our results remain largely the same, indicating that the sampled data are representative for the full dataset. The full estimation results for the robustness checks are provided in Appendix A.4.

3.5 ADDITIONAL ANALYSIS

So far, we have discussed the main findings derived from the model estimation. In this section, we perform several additional analyses to further reveal the importance of joint inclusion of multiple self-regulatory dimensions from the perspective of information synergy, model performance, and individuals' regular engagement. Implications associated with these analyses are discussed.

3.5.1 *Perspective of Information Synergy*

As noted, the MCTHMM framework we propose leverages multiple self-regulatory dimensions in calibrating individuals' online weight-management behaviors. In this subsection, we investigate the scenario in which individuals' self-regulatory dimensions are partially calibrated. Specifically, we estimate two single-dimensional CTHMMs, using the number of weigh-ins and the number of journal-recording behaviors as the state-dependent outcome, respectively. The covariates are kept the same as those in our main model. The model estimation results are reported in Appendix A.5.

We find that the results produced by the single-dimensional models are not entirely the same as those of our two-dimensional model, which is not surprising, as the single-dimensional models

utilize only a partial set of information to calibrate individuals' online weight-management behaviors in the estimation procedure. In addition, we notice that the estimation results produced by the two single-dimensional models can be different from each other as well, even though both weigh-ins and journal-recording behaviors are valid proxies for individuals' engagement in self-regulation. For instance, the effect of initiating posts on state transition is suggested to be positive in state H in the first single-dimensional model but negative in the other, whereas the corresponding result is insignificant in our two-dimensional model. These inconsistencies provide evidence that, although both weigh-ins and journal-recording behaviors reflect individuals' efforts to conduct self-regulation, they provide heterogeneous measurements. In practice, researchers and practitioners may rely on their own domain knowledge to choose measurements for individuals' health-management behaviors. It is likely that researchers and practitioners partially evaluate individuals' health-management behaviors and derive completely different implications. Thus, these results imply that it is important to combine the information revealed from different self-regulatory dimensions to improve the characterization of individuals' online behavior patterns.

3.5.2 *Model Performance Comparisons*

In this subsection, we present the results of several model comparisons to further illustrate the effectiveness of our proposed MCTHMM framework to describe and predict individuals' online weight-management behaviors. The alternative models we consider are the single-dimensional CTHMMs and a two-dimensional discrete-time HMM (DTHMM) in which the time frame is aggregated to a week. Again, we keep all of the covariates the same so that we are able to compare how different model structures contribute to our understanding of users' participation patterns. By comparing our model with the single-dimensional models, we are able to examine the effectiveness of combining multiple self-regulatory dimensions. The comparison between our model and the

two-dimensional DTHMM can provide us with the value of the timing process of individuals' online engagement.

We first compare the model fit, which is an important indicator of a model's efficiency in describing data patterns. The comparison results are reported in Table 3.4. As seen in the table, our two-dimensional CTHMM has the best model fit (it has the highest log-likelihood and the lowest AIC/BIC). These results suggest that both the inclusion of multiple self-regulatory dimensions and the timing process can help us to improve the description of individuals' online behavior patterns.

Table 3.4 Model Comparison (1): Model Fit

Model Fit	One-Dimensional Models		Two-Dimensional Models	
	Weigh-in	Journal	DT	CT
Log-likelihood	-15380.35	-16506.07	-16509.04	-15020.44
AIC	30868.70	33120.14	33162.08	30196.88
BIC	31279.55	33535.54	33718.51	30841.54

Then, we investigate the model performance from a prediction perspective. Knowing and predicting users' behaviors can be of significance in practical scenarios, as platform designers and policymakers can make prompt adjustments to their decision making based on the forecasts. The prediction is conducted on a test dataset, which contains 100 individuals who are randomly selected from the remainder of the dataset that is not used for estimation. We predict individuals' self-regulatory dimensions on a weekly basis, as typical weight-loss therapy includes weekly consultations (Ballantine and Stephenson 2011; Moisio and Beruchashvili 2009). We use binary-valued outcome variables to calculate the prediction accuracy in percentage. A detailed description of the prediction procedure can be found in Appendix A.6.

The prediction results are summarized in Table 3.5. For the two-dimensional models, we calculate the prediction accuracy not only for each dimension but also for the sum of both

dimensions. For the single-dimensional models, we can predict for only one self-regulatory dimension at a time; however, we can combine the prediction results from each model to calculate the prediction accuracy for the sum of both dimensions. The results show that our two-dimensional CTHMM outperforms all of the alternative models in predicting users' overall engagement in self-regulation on the platform, reaching a prediction accuracy of 87.31%, and that it predicts more accurately than the two-dimensional DTHMM for each self-regulatory dimension. Compared to the single-dimensional models, our two-dimensional CTHMM has a higher prediction accuracy for journal-recording behaviors and a slightly lower prediction accuracy for weigh-ins. These results, again, demonstrate the importance of including the timing process and multiple self-regulatory dimensions.

Table 3.5 Model Comparison (2): Prediction Accuracy

Prediction Accuracy	One-Dimensional Models		Two-Dimensional Models	
	Weigh-in	Journal	DT	CT
Dimension 1 (Weigh-in)	86.83%	NA	79.92%	86.67%
Dimension 2 (Journal)	NA	81.83%	76.16%	83.84%
Sum of Both Dimensions	84.44%		80.76%	87.31%

All of the above comparison results provide supporting evidence for the superiority of the modeling approach we proposed in this study. As explained in Shmueli and Koppius (2011), model fitness and prediction power not only embody practicality but also may provide support for the underlying theoretical framework. Thus, our approach is not only considered to make a methodological contribution by extending the traditional HMM to account for multi-dimensional outcomes and continuous-time processes, but also promotes the establishment of new theory. So far, Section 3.5.1 and 3.5.2 illustrate, from researchers' or healthcare practitioners' perspective, that the inclusion of multiple self-regulatory dimensions is essential for describing individuals' online weight-management behaviors. In Section 3.5.3, we further investigate from individuals'

perspective whether joint management of multiple self-regulatory dimensions in OWCs is indeed associated with better health-management performance.

3.5.3 *Individuals' Regular Engagement*

Evidence in the literature suggests that regular engagement in health management is generally associated with better health outcomes. For instance, Butryn et al. (2007) found that consistent self-weighing can help individuals maintain their weight by allowing them to make timely adjustments to weight-management behavior. In practice, self-weighing can be easy, as individuals need only to step on a scale and record numbers. Compared to self-weighing, keeping a weight-loss journal requires more effort, as individuals need to record in detail their weight-loss routines. Considering that not only researchers but also individuals themselves may ignore certain self-regulatory dimensions, we are particularly interested in investigating two synthetic scenarios: (1) when individuals are able to regularly weigh in and make a journal record on a daily basis and (2) when individuals are able to regularly weigh in but do not have time to regularly make a journal record. We aim to explore whether regular engagement in both self-regulatory dimensions is linked to higher efficacy levels.

Corresponding to the synthetic scenarios, we simulate two self-regulation modes. In the first mode, individuals weigh in and make a journal record every day, whereas, in the second mode, individuals only weigh in each day but make a journal record occasionally. Using the model estimates that we obtain, we can then simulate individuals' hidden state trajectories based on the filtering approach proposed by Hamilton (1989). This approach relies on the Bayes rule to calculate the posterior probabilities of being classified into one state. In Figure 3.3, we plot the average hidden state trajectories across individuals recovered for each self-regulation mode. We also plot the average hidden state trajectory recovered from the real data. It is clearly shown that,

on average, individuals who weigh in and make a journal record each day tend to have a higher state trajectory compared to those who only weigh in each day. This finding indicates that individuals who regularly engage in both self-regulatory dimensions generally have a higher self-efficacy level. In addition, the results show that the average hidden state trajectory inferred from the real data is much lower than those recovered from the synthetic data. This result further confirms the importance of regular engagement in both self-regulatory dimensions, as individuals in our dataset on average do not weigh in or write journals as regularly as on a daily basis.




Figure 3.3 Hidden State Trajectories

3.6 DISCUSSION AND CONCLUSION

As obesity has become a major health concern globally, weight management has been gathering increasing attention from both academia and healthcare practitioners. OWCs are believed to play a critical role in assisting individuals' weight management. Individuals can easily find multiple weight-management tools in OWCs, such as digital weight recorders, weight-loss journals, and calorie counters, which help them to engage in different self-regulatory dimensions in regard to weight-loss outcome and weight-management behavioral routines. Prior research on self-regulation has suggested that individuals' self-management of both health outcomes (e.g., body

weight) and health behaviors (e.g., weight-management behavioral routines such as diet and physical activity) are important (Bandura 1991; Bandura 1998). In particular, through observing and managing their weight, individuals can acquire directional feedback on their weight-loss performance, whereas through managing behavioral routines, individuals can better understand and conceptualize the correlation between their behaviors and the corresponding weight-loss outcomes (Bandura 1991; Krukowski et al. 2008; Thomas et al. 2017). However, prior studies have largely overlooked individuals' engagement in multiple self-regulatory dimensions in OWCs (Johnson and Wardle 2011; Thomas et al. 2017). In this study, we aim to investigate the dynamics associated with individuals' different self-regulatory dimensions in OWCs as a means to develop a comprehensive understanding of how the operating mechanisms of OWCs affect individuals' self-regulation.

We propose a novel structural model, an MCTHMM, for our research purpose. This model synthesizes innovative model structures into a traditional HMM to evaluate individuals' multi-dimensional self-regulatory dimensions in OWCs under a continuous-time framework. Through model estimation, we derive evidence for individuals' dynamic engagement patterns across different self-regulatory dimensions. In particular, we find that individuals may shift their participation focus from weight-loss outcomes to weight-management behavioral routines after they have participated in social activities on the platform. We also derive interesting insights regarding individuals' self-efficacy, which is a core cognitive driver for engagement in self-regulation. For instance, we investigate the effects of individuals' social activities on their efficacy states and find that different posting strategies and social ties may affect transitions in individuals' efficacy state in distinct ways. These findings are robust to several variations, including model specification, variable measure, and sampling procedure.

Our analysis results amplify the importance of accounting for individuals' multiple self-regulatory dimensions in OWCs. As individuals may dynamically shift their self-regulation focus throughout the process of online participation, overlooking certain self-regulatory dimensions would lead to biased estimation for individuals' online engagement, which may further underrate the effectiveness of OWCs. These implications can be applied to other health-management contexts, especially individuals' chronic disease management. Due to its complexity, a chronic condition usually requires individuals to jointly manage multiple health aspects, such as related health outcomes, daily health behaviors, and/or medication. An online healthcare community can easily facilitate individuals' management of these aspects by providing relevant tools. Thus, individuals' engagement in online healthcare communities may also exhibit similar dynamics across different self-regulatory dimensions.

Our study can be extended in several ways. First, our results illustrate how individuals engage in self-regulatory dimensions in an online setting. It is unclear whether individuals would exhibit similar engagement patterns in offline settings. Future research can explore this direction to reveal more insights regarding individuals' offline self-regulatory dimensions and online-offline weight-management dynamics if offline data can be acquired. Second, our study analyzes individuals' behavior patterns in weight management rather than the associated impact on weight outcomes. As prior studies have generally shown that individuals' engagement in self-regulation is a key factor for weight maintenance (Burke et al. 2011; Butryn et al. 2007; Hall and Kahan 2018), it is worth exploring whether (and how) individuals' engagement patterns in different self-regulatory dimensions could affect their actual weight-loss outcomes. Finally, the data we use were collected from a website-based OWC. Future research may consider an examination of how the utilization

of interactive technologies, such as mobile devices, further helps people to engage in weight management.

Chapter 4. A STRUCTURAL MODEL OF USER DYNAMIC ENGAGEMENT IN ONLINE HEALTHCARE INTERVENTIONS

Online healthcare platforms provide behavior intervention programs to help users make lifestyle changes and promote personal health. Typical behavior intervention programs offer structured behavior-treatment plans and/or instructions to guide users' behavior change. The online behavior intervention programs engage users through breaking long-term health-management goal into small pieces that can be accomplished every day. Prior studies have found that behavior intervention programs with specified short-term goals can help users to stay engaged and stick to daily routine (Norcross and Vangarelli 1988). In addition, the outlined behavior-change instructions help users to reflect what behaviors are beneficial to their health so that they can better conceptualize the association between lifestyles and health outcomes.

While there is a large body of research on online healthcare platforms and Health 2.0 (Holman and Lorig 2000; van der Eijk et al. 2013), less investigation is on individuals' dynamic engagement in the provided behavior intervention programs. In practice, users can sequentially choose intervention programs to participate in based on their past engagement experience. This opens an opportunity for individuals to learn about the effectiveness of the intervention programs and dynamically adjust their participation decisions. For example, after observing the health variations, users can learn about how different behavior interventions work on themselves so that they can adjust their participation decision based on their perception. To date, however, such learning behavior of users have not been examined. In this thesis, I am thus motivated to fill the research gap by characterizing individuals' learning process in participating in online behavior intervention programs.

We examine individuals' learning process using a 4-month panel data collected from a leading online weight-loss platform in the United States. The focal platform provides weight-loss challenges to help users regulate their lifestyles, such as diet-related challenges that help users modify their diet habits, and exercise-related challenges to encourage users be more physically active. Each challenges provide detailed instructions on how to achieve a particular weight-loss goal. We build a structural model to characterize individuals' learning behaviors on the platform, in which individuals update their perceptions of challenge effectiveness based on the weight changes they observe. Our model estimation results provide evidence for individuals' learning capabilities. We find that although individuals' initial priors can be very different from the uncovered true challenge effectiveness, after several periods' engagement, individuals are able to adjust their perceptions so that their posteriors are close to the true challenge effectiveness. In addition, we find that the efficiency of learning is largely dependent on the noise level contained in the delivered signal. In personal health management settings, individuals usually use their health variations (e.g., weight changes) as a direct feedback of intervention effectiveness. When the health signals contain much noises, individuals may not be able to learn the effectiveness of interventions well. We further conduct simulations to visualize individuals' average learning curve to demonstrate the overall learning efficiency on the platform.

Our study makes several contributions. First, our study is among the first to examine individuals' learning behavior in personal health management. We develop a health-signal-based hierarchical learning model to rationalize individuals' dynamic engagement patterns in online behavior intervention programs. Second, our results demonstrate the limits of individuals' learning capability in the healthcare context. We find that individuals' learning efficiency is restricted by the noise level of the delivered health signals. In many settings, such as managing a chronic

condition, the delivered health signals can contain much variance (e.g., body weight fluctuations), due to the complex nature of the disease. Thus, our study connects to clinical behavioral research and illustrates the importance of patient support. Third, our study adds to the emerging literature on patient empowerment and patient-centered medication. Patient-centered care is the core of modern healthcare systems and it emphasizes patient empowerment in the care delivery process. Patient empowerment means more participation, more activeness, and more controls over the decisions made from patients. Our study examines individuals’ dynamic action course stimulated by their past engagement experience. As such, it demonstrates patient-oriented motions and provides implications for patient empowerment.

4.1 CONTEXT

In this section, we discuss the research context and present descriptive statistics for the data. We also show a model-free reduced-form analysis to demonstrate the potential learning behaviors of individuals on online healthcare platforms.

4.1.1 *Research Context and Data*

In this study, we investigate an online weight-loss platform that provides weight-management interventions to users. The intervention programs are called weight-loss challenges, which are aimed to help users set specific weight-loss goals and provide instructions. Typical weight-loss challenges require users to modify their dietary behaviors (e.g., “avoid processed food”), increase participation in physical activities (e.g., “walk 10k steps every day”), or to lose certain amount of weight (e.g., “lose 5 pounds in two weeks”). Users can choose to participate in challenges by signing up and following the instructions provided. Users can find challenge information in their portal, such as challenge title, description, and duration. We were able to track users’ participation

record from their personal profile on a weekly basis, that is, in which week that a user participated in what challenge(s). Most challenges require users to report their body weight along the participation process. This reporting schema can help users track their weight variations with time so that they can develop a better understanding of their weight-loss progress through challenge engagement. In this study, we argue that users can gradually learn about the effectiveness of different challenge types and dynamically adjust their participation decisions. In Figure 4.1, we provide a diagram to show how users can retrieve challenge information and participate in challenges on the platform.




Figure 4.1 Individuals' Participation in Online Weight-Loss Challenges

We collected data on users' challenge participation and weigh-in records on a weekly basis. The dataset contains 1049 users' participation record from January 1st 2014 to April 30th 2014. During our data collection window, there were 165 challenges posted on the platform. For each challenge, we collected the challenge title information and the challenge descriptions/instructions. We also collected information on users' gender, age, tenure of participation, initial weight, and weekly weigh-in records. We provide the summary statistics of the data in Table 4.1.

Table 4.1 Summary Statistics of Data

Variable	Min	Max	Mean	Std. Dev.	Count
<i>Challenge information</i>					
The number of diet-related challenges (e.g., intervention on nutrition, calorie intake)					65
The number of exercise-related challenges (e.g., intervention on physical activities)					37
The number of weight-change-related challenges (e.g., intervention on body weight changes)					37
The number of other challenges (e.g., meditation, self-monitoring, etc.)					26
Total number of challenges					165
The number of challenges chosen in each week	0	5	1.54	1.20	
<i>User Information</i>					
Gender (Female = 1)	0	1	0.52	0.50	
Age	19	57	38.07	11.26	
Tenure of participation on the platform (in days)	21	672	341.61	161.23	
Initial weight when registering as a platform member	127	298	210.62	49.10	
Users' weigh-in records	108	320	210.74	49.73	
Total number of users					1049

4.1.2 Model-Free Evidence

Before we perform structural analysis on users' learning behavior, we use the following regression model to test whether users' past weight-loss information has any impact on their participation decisions:

$$Y_{it} \sim \text{Intercept} + \eta_1 \cdot \text{gender}_i + \eta_2 \cdot \text{age}_{it} + \eta_3 \cdot \text{initial_weight}_i + \eta_4 \cdot \text{tenure}_{it} + \eta_5 \cdot \text{ex_weightloss}_{it} \quad (4.1)$$

In equation (4.1), we control users' gender, age, initial weight when first registered on the platform, tenure of participation, and the number of experienced weight-loss periods. The last variable, *ex_weightloss*, is our focal of interests. It measures the number of periods that a user has experienced weight loss and, thus, captures the impact of users' past weight-loss information. The dependent variable is measured by two methods: whether a user choose any challenge to

participate (binary) and the number of challenges chosen (discrete). If there is any learning effect, we should expect that users' past weight-loss experience ($ex_weightloss$) has a significant effect on the dependent variable Y_{it} .

Table 4.2 reports the estimated regression coefficients. The results suggest a positive and significant effect of users' past weight-loss experience on both their participation propensity (model 1) and the number of challenges they choose (model 2). These results provide model-free evidence for the existence of individuals' learning behaviors.

Table 4.2 Regression of Participation Decision

Variable	Model 1 dependent variable: binary regression link: logit		Model 2 dependent variable: discrete regression link: Poisson	
	Coefficient Estimate	Std. Error	Coefficient Estimate	Std. Error
gender	-0.0120	0.0367	-0.0010	0.0112
age	-0.0187	0.0165	-0.0188***	0.0047
initial_weight	-0.0712*	0.0371	-0.0295***	0.0115
tenure	-0.0168	0.0113	-0.0047	0.0032
ex_weightloss	0.0269**	0.0118	0.0228***	0.0035
Intercept	1.4365***	0.1125	0.7731***	0.0337

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Total number of observations: 16784.

4.2 MODEL

In this section, we establish a structural model to rationalize users' learning behaviors in online healthcare platforms. Our model adopts a Bayesian learning framework to describe how users' health signal helps them to modify their perceptions of behavior interventions during the participation process. In the following, we first explain users' utility-based decision making, and then we describe the dynamic learning mechanism.

4.2.1 Challenge Selection

Users select weight-loss challenges based on the utility they will gain. As the goal for challenge engagement is to lose weight and improve wellness, the utility is largely dependent on the

effectiveness of challenges, that is, how well a challenge works on individual users. We define the utility function for each individual as follows:

$$U_{ijt} = \alpha \cdot EW_{ijt} + \lambda^T controls + \varepsilon_{it} . \quad (4.2)$$

The first term, EW_{ijt} , denotes individual i 's expected weight loss for participating in challenge j . This term reflects users' perceived effectiveness of challenges. As individuals may gradually form belief about challenge effectiveness, we allow EW_{ijt} to vary with time. The second term incorporates a set of control variables. We use this term to capture users' preferences over the challenge characteristics that are not directly related to weight-loss, such as the length of challenge instructions and the sentiment polarity of challenge statements. These characteristics are shown to affect users' preferences on text information (Almars et al. 2019; Das and Chen 2007) and thus influence users' participation behaviors. ε_{it} is a random shock that is private to users. Users' expected weight loss are contingent on specific challenge features. For instance, depending on their past weight-loss experience, users may form expectations for different types of challenges, such as diet-related challenges, exercise-related challenges, and etc. Therefore, we further express EW_{ijt} with the following equation:

$$EW_{ijt} = \tilde{\theta}_t^T x_j , \quad (4.3)$$

where x_j is a one-hot encoded vector denoting challenge types; $\tilde{\theta}_t$ represents user i 's perceived effectiveness of the challenge types at time t . In the focal context, the challenges can be roughly divided into 5 types. We provide the coding of the challenges types in Table 4.3.

Table 4.3 Coding of Challenge Types

Coding	Challenge Type	Description
Type 1	DietStructureChange	Challenges that require users to change their diet structure, e.g., eat high-protein, low-carbon foods.
Type 2	CalorieRestriction	Challenges that require users to restrict their calorie intake, e.g., eat 1200-1500 calories per meal, restrict serving to half, etc.
Type 3	Exercise	Challenges that require users to do physical exercise, e.g., jog 30 minutes per day
Type 4	WeightLoss	Challenges that require users to lose certain pounds within a period, e.g., lose 5 lbs. in two weeks
Type 5	BehavioralOther	Challenges that require users to make other behavioral changes, e.g., to meditate, reflect, establish motivation, etc.

As such, equation (4.3) allows us to model the formation of effectiveness perceptions for each challenge type, which is important in practice, as different behavior treatment tend to have discrepant effects on individuals.

Combining equation (4.2) and (4.3), the in-period utility can be formulated as follows:

$$U_{ijt} = \alpha \cdot \tilde{\theta}_{it}^T x_j + \lambda^T controls + \varepsilon_{it}. \quad (4.4)$$

For estimation simplicity, we assume the random shock ε_{it} to follow *i.i.d.* Gumbel distribution.

Thus, the probability for user i to choose challenge j at week t is

$$P_{ijt} = \frac{e^{E[U_{ijt}]}}{1 + \sum_j e^{E[U_{ijt}]}}, \quad (4.5)$$

where $E[U_{ijt}] = \alpha \cdot E[\tilde{\theta}_{it}]^T x_j + \lambda^T controls$.

4.2.2 Learning

Through continuous engagement, users gain experience on how each challenge type works on their health. We model users' dynamic update of their perceptions of challenge effectiveness using a Bayesian learning framework. Formally, suppose that each individual holds an initial belief about challenge effectiveness at the beginning, which follows a normal distribution: $\tilde{\theta}_0 \sim N(\theta_0, \Sigma_{\theta_0})$.

Each period, by participating in weight-loss challenges and observing their weight loss, users can

learn about the effectiveness of challenge types and update their beliefs. Specifically, we write users' experienced effectiveness of challenge types as follows:

$$y_i^k \sim N(\beta_i^k, \sigma_{yk}^2), \quad (4.6)$$

where β_i^k is the true effectiveness of type- k challenges for individual i , and σ_{yk}^2 captures the noise level in the delivered signal. The true effectiveness of challenge types are individual-specific, as healthcare is highly personal and the same challenge type may have different effect across individuals. Let $\Lambda = \text{diag}(\sigma_{y1}^2, \sigma_{y2}^2, \dots, \sigma_{yK}^2)$. Based on the Bayes rule, the updating mechanism is written as follows:

$$\tilde{\theta}_{it}^k \sim N(\theta_{it}^k, \sigma_{\theta_{it}^k}^2), \quad k=1, \dots, K, \quad (4.7)$$

$$\theta_{it}^k = \left(\frac{\theta_{i,t-1}^k}{\sigma_{\theta_{i,t-1}^k}^2} + \frac{y_i^k \delta_k}{\sigma_{yk}^2} \right) \left(\frac{1}{\sigma_{\theta_{i,t-1}^k}^2} + \frac{\delta_k}{\sigma_{yk}^2} \right)^{-1}, \quad (4.8)$$

$$\sigma_{\theta_{it}^k}^2 = \left(\frac{1}{\sigma_{\theta_{i,t-1}^k}^2} + \frac{\delta_k}{\sigma_{yk}^2} \right)^{-1}, \quad (4.9)$$

where θ_{it}^k denotes the k -th element of θ_{it} , δ_k is an indicator function which equals one when the chosen challenge belongs to type k . When a user has chosen a challenge of type k , he/she experienced y_i^k and use it to update his/her belief about the effectiveness of challenge type k ; the user's belief of other unchosen challenge types will not be updated.

4.2.3 Signal Generating Process

As challenge effectiveness is reflected in users' weight variations, we model the signal generating process as follows:

$$\Delta w_{it} = \beta_i^T x_{j^*(t-1)} + e_{it}, \quad (4.10)$$

where Δw_{it} denotes user i 's weight changes in period t , i.e., $\Delta w_{it} = w_{i,t-1} - w_{it}$;

$\beta_i = (\beta_i^1, \beta_i^2, \dots, \beta_i^K)$; $e_{it} \stackrel{i.i.d.}{\sim} N(0, \sigma_w^2)$ captures random shock. $j^*(t)$ denotes the indices of the challenge(s) chosen at period t , and $x_{j^*(t)}$ denotes the types of the chosen challenge(s). When users selected more than one challenges, $x_{j^*(t)} = \sum_{j \in j^*(t)} x_j$. We assume that individual-specific coefficient β_i follows the following distribution:

$$\beta_i \sim MVN(\beta_{0i}, \Sigma_\beta), \quad (4.11)$$

where $\beta_{0i} = b_0 + b_1 \cdot \text{gender}_i + b_2 \cdot \text{age}_i + b_3 \cdot \text{initial_weight}_i$, MVN denotes multivariate normal distribution. The formulation of β_{0i} as a linear function of users' gender, age, and initial weight enables us to capture the effects of shared individual characteristics on challenge effectiveness. For example, female users who are in certain weight range may find a particular type of challenges generally more effective than others. The variance-covariance matrix, Σ_β , captures the level of unobserved individual heterogeneity after controlling gender, age, and initial weight. In addition, it can model the correlations between any two dimensions of β_i and thus models the synergy among different challenge types.

4.2.4 Likelihood

The likelihood function is composed of two parts: the likelihood for observing individuals' in-period weight loss, and the likelihood for observing individuals' challenge-selection behaviors. Mathematically, the individual-level likelihood is as follows:

$$L_i(D_i | \Theta) = \prod_t \Pr(\Delta w_{it} | \beta_i, \sigma_w^2, d_{it}) \Pr(d_{it} | \theta_{it}), \quad (4.12)$$

where d_{it} denotes the choice of individual i at period t ; D_i denotes the observed data for individual i , i.e., $D_i = \{\Delta w_{it}, d_{it}\}_{t=1}^T$. The likelihood for the entire observed data D is

$$L(D | \Theta) = \prod_i L_i(D_i | \Theta).$$

4.2.5 Estimation

The parameters to be estimated in our model include

$$\Theta = \{\{b_0, b_1, b_2, b_3\}, \Sigma_\beta, \{\beta_i\}_{i=1, \dots, N}, \sigma_w^2, \theta_0, \{\sigma_{yk}^2\}_{k=1, \dots, K}, \alpha, \lambda\}.$$

Among these parameters,

$\{b_0, b_1, b_2, b_3, \Sigma_\beta\}$ are the hyper-parameters that generate β_i ; $\{\{\beta_i\}_{i=1, \dots, N}, \sigma_w^2\}$ are the parameters in the weight-loss distribution; $\{\theta_0, \{\sigma_{yk}^2\}_{k=1, \dots, K}\}$ are learning-related parameters; $\{\alpha, \lambda\}$ are the parameters in the utility function. For all the variance parameters, we estimate its logarithm transformation, since they only take positive values. We provide a summary of all the notations in Table 4.4.

As our model contains both individual-specific parameters and parameters that are common across individuals, we adopt a hierarchical Bayesian approach to estimate the model. Specifically, we use a Markov chain Monte Carlo (MCMC) method to iteratively make draws from the posterior distribution to form a series of samples that can be used to approximate the unknown parameter distributions. The estimation procedure is executed 40000 rounds, with the first half discarded as the burn-in samples. We discuss the technical details in Appendix B.

Table 4.4 Summary of Notations

Notation	Description
θ_0	Individuals' initial belief on challenge effectiveness; the k th dimension denotes individuals' prior beliefs for challenge type k .
Σ_{θ_0}	The variance-covariance matrix of individuals' prior beliefs. We normalize it to an identity matrix for identification purpose.
σ_{yk}^2	The variance of signals generated in participating in challenge type k , which measures the noise level in the delivered signal.
β_i	The individual-specific true challenge effectiveness; the k th dimension denotes the average true effectiveness of challenge type k .
α	Individuals' preference on perceived challenge effectiveness
λ	Individuals' preference on challenge statement features, such as challenge statement length and sentiment polarity.
σ_w^2	The variance in the signal generating process
b_0	The hyper parameter that captures the baseline effect on the true mean challenge effectiveness (the constant term)
b_1	The hyper parameter that captures the gender effect on the true mean challenge effectiveness
b_2	The hyper parameter that captures the age effect on the true mean challenge effectiveness
b_3	The hyper parameter that captures the effect of initial weight on the true mean challenge effectiveness
Σ_β	The variance-covariance of challenge effectiveness β_i

4.2.6 Identification

We briefly discuss the identification of the model parameters here. All the model parameters can be identified based on the data variation patterns associated with them. In our model, users' choice-making patterns uncovers how they prefer one item than the others. These patterns help us to identify the parameters in users' utility function. We set the utility of the outside option, i.e., choosing no challenges, to zero for identification purpose. The difference between users' prior choice-making patterns and their later-stage choice-making patterns allows us to infer the learning effect. That is, users' behavior changes are inherently driven by changes in their perception of

challenges. We normalize the variance-covariance matrix of the initial belief distribution to an identity matrix I in order to identify the noise level in the signals. As such, the estimated signal noise would represent the relative magnitude of noise level compared to the prior variance. The true effectiveness of challenges is identified through the signaling process.

4.3 RESULTS

4.3.1 *Main Estimation Results*

In Table 4.5, we report the estimated learning parameters, including individuals' initial prior on challenge types and signal noises. The results suggest that individuals' prior beliefs are different across challenge types. Specifically, on average, individuals tend to hold a negative prior on the first four types of challenges while holding a positive prior on the fifth challenge type. These results indicate that users in general believe that making other behavioral changes, such as self-monitoring and meditation, are more effective than diet- and exercise-related challenges at the beginning. These prior beliefs, however, can be changed during individuals' engagement process, as individuals continuously observe their health signals and learn about the true effectiveness of each challenge type. We report the estimated log noise levels (i.e., the variance) of the delivered signals for each challenge type. These parameter estimates reflect the heterogeneity in the signaling processes when individuals participate in different challenge types. We find that signal noises for certain challenge types can be large; for example, when individuals participate in Type-1 challenges, the experienced noise level in the delivered signal is $\exp(5.9490)$. A large signal variance may work against individuals' learning efficiency in the process of health management. In Section 4.3.2, we develop a more granular discussion on individuals' learning speed affected by signal variances.

Table 4.5 Estimated Learning Parameters

Parameter	Description	Estimate	95% Credible Interval
Initial priors of challenge effectiveness (θ_0)			
θ_0^1	Prior on Type-1 challenge	-2.4554	(-2.6600, -2.0424)
θ_0^2	Prior on Type-2 challenge	-8.8807	(-13.2187, -5.1733)
θ_0^3	Prior on Type-3 challenge	-16.6489	(-19.7755, -12.0719)
θ_0^4	Prior on Type-4 challenge	-19.2831	(-22.6199, -15.9726)
θ_0^5	Prior on Type-5 challenge	6.9256	(5.1202, 7.4220)
Variance of signals (σ_y^2)			
$\log(\sigma_{y1}^2)$	Log signal variance of Type-1 challenge	5.9490	(3.2259, 9.3833)
$\log(\sigma_{y2}^2)$	Log signal variance of Type-2 challenge	3.8467	(1.9342, 5.5773)
$\log(\sigma_{y3}^2)$	Log signal variance of Type-3 challenge	-5.4978	(-10.1173, -1.7055)
$\log(\sigma_{y4}^2)$	Log signal variance of Type-4 challenge	-3.4483	(-8.6227, 1.8348)
$\log(\sigma_{y5}^2)$	Log signal variance of Type-5 challenge	4.7922	(2.2647, 6.8215)

We report the estimated true mean effectiveness for each challenge type in Table 4.6. As modeled in equation (4.11), the true mean effectiveness is factorized by individuals' shared characteristics, such as gender, age, and initial weight. Each column in Table 4.6 contains the factorized effect in one particular dimension. Specifically, the first column b_0 contains the constant terms, which capture the baseline of effectiveness for each challenge type. The last three columns, b_1 , b_2 and b_3 , capture the effects of gender, age, and initial weight on average challenge effectiveness, respectively. The results suggest that depending on individuals' demographics, challenge effectiveness vary across different groups of people. For example, although gender do not generate significant difference for challenge type 1, 2, 3 and 5, challenge type 4 (WeightLoss) tend to be more effective for female users than male users ($b_1^4=6.0216$). In addition, it is found that type-5 challenges (BehavioralOther) tend to benefit people more when they have a larger age ($b_2^5=6.9754$), which is reasonable, as older people can have a difficult time in conducting intensive diet or exercise challenges due to their body condition. Type-1 challenges (DietStructureChange) are found to be more effective for people who are with a lower weight ($b_3^1=-3.1490$). This can be

explained by the fact that people who have a larger body weight usually need stricter calorie intake and expenditure methods to help them lose weight rather than just modifying their diet structures. These results provide insights into personalization of weight-loss programs.

Table 4.6 Estimated Factorized Challenge Effectiveness

Estimates (95% CI)	b_0 (constant)	b_1 (gender)	b_2 (age)	b_3 (initial weight)
Type 1	-1.0731 (-1.5063, -0.4171)	3.4153 (-3.1636, 7.7568)	5.3995 (-5.6475, 12.6935)	-3.1490 (-6.7730, -0.7577)
Type 2	3.0297 (1.9213, 3.7590)	-1.5737 (-3.8064, 1.8059)	3.6812 (-2.1084, 12.4501)	-2.4654 (-4.2260, 0.1979)
Type 3	10.4674 (8.3063, 13.7440)	-8.5546 (-15.1867, 1.4916)	-0.9766 (-3.7348, 3.1996)	-2.2412 (-7.3166, 5.4492)
Type 4	-2.3550 (-6.2491, 3.5447)	6.0216 (2.5217, 11.3207)	-0.8101 (-9.3759, 12.1632)	-3.5175 (-9.5667, 5.6449)
Type 5	2.4328 (-0.3458, 6.6402)	2.6030 (-0.0039, 4.3475)	6.9754 (6.6174, 7.2131)	-0.6636 (-1.9023, 1.2117)

In addition to the shared characteristics, individuals' unobserved heterogeneity may further create variations in challenge effectiveness at the personal level. We capture such unobserved heterogeneity by the variance-covariance matrix Σ_β and report the estimation results in Table 4.7. The results suggest that after controlling gender, age, and initial weight, there is still considerable amount of heterogeneity across individuals. These results indicate that the effectiveness of behavior treatments can be highly personal, which is in line with prior studies (Chawla and Davis 2013; Hopp et al. 2018). Thus, it is important for individuals to self-monitor and analyze their health-management process to derive a better understanding of what behavior treatments are favorable for themselves. In Figure 4.2, we plot the distribution of the true challenge effectiveness for all users. The plots demonstrate different degrees of variations across challenge types. In general, the effectiveness of type-1 challenges vary the most across individuals, followed by type-5 and type-3. Note that these variations are determined jointly by users' shared characteristics and unobserved individual-level heterogeneity.

Table 4.7 Estimated Variance-Covariance Matrix Σ_β

Var-Cov	β_i^1	β_i^2	β_i^3	β_i^4	β_i^5
β_i^1	0.1065	-0.0176	0.0535	-0.0004	-0.0396
β_i^2	-0.0176	0.0905	-0.0361	-0.0165	0.0265
β_i^3	0.0535	-0.0361	0.1888	0.0937	-0.0958
β_i^4	-0.0004	-0.0165	0.0937	0.1287	-0.0228
β_i^5	-0.0396	0.0265	-0.0958	-0.0228	0.1911




Figure 4.2 Distribution of True Challenge Effectiveness in the User Population

We report the estimated utility parameters in Table 4.8. We find that individuals' preference on expected weight loss is positive and significant. This finding is consistent with our assumption that users' challenge selection decision is driven by their underlying healthcare needs. Apart from expected weight loss, individuals may form preferences over explicit challenge features, such as the length of challenge description and the sentiment polarity. Specifically, we find that the effect of challenge description length is insignificant but the effect of squared description length is negative and significant. These results indicate that users do not prefer tedious intervention descriptions and the effect of intervention description length is nonlinear. In addition, we find that

the effects of positive sentiment and negative sentiment in challenge descriptions are positive and significant. These results suggest that users tend to choose challenges that have a strong sentiment polarity, which is reasonable, as those challenges typically attract much user attention. Our findings on user preference is in line with prior behavior-intervention and goal-setting theories (Doran 1981; Locke and Latham 1990; Ogbeiwi 2018), which suggest that the formulation of goal statement can affect how individuals perceive the goal and thus influence their behaviors.

Table 4.8 Parameter Estimates in Utility Function

Parameter	Description	Estimate	95% Credible Interval
α	Preference of expected weight loss	0.0497	(0.0400, 0.0624)
λ_1	Preference of challenge statement length	-0.6299	(-1.5391, 0.3790)
λ_2	Preference of squared challenge statement length	-0.8118	(-2.2930, -0.2506)
λ_3	Preference of positive sentiment	0.5133	(0.2621, 0.7593)
λ_4	Preference of negative sentiment	0.6167	(0.4534, 0.7995)

4.3.2 Learning Curve

The difference between individuals' initial prior beliefs and the estimated true challenge effectiveness indicate that individuals generally do not have adequate knowledge about healthcare behavioral interventions, and this gap demonstrates a potential for learning through continuous engagement in the health-management process. In the learning process, signal variances can determine how well individuals learn from their weight-loss records when they participate in each challenge type. To illustrate how signal variances affect individuals' learning, in Figure 4.3, we visualize individuals' average learning curve. In each plot, the y-axis represents individuals' average posterior beliefs of challenge effectiveness, and the x-axis represents time (in weeks). The recovered learning curves show that users can gradually shift their posterior beliefs towards the true effectiveness value. However, because the signal variances are different across challenge types, the speed of learning varies. For challenge types that have a small signal variance (e.g., type

3 and 4), the learning curve converges to the estimated true effectiveness value quickly; this means that users can effectively learn the true challenge effectiveness through participation. However, for challenge types that have a larger signal variance (e.g., type 1, 2 and 5), the learning appears to be much slower, as the average posterior deviates slightly from the initial prior. Together, these results indicate that despite the potential for learning, individuals may not learn intervention effectiveness well due to potential high noise level in the delivered signal.




Figure 4.3 Average Learning Curve

In reality, the noise of signal can come from multiple sources, such as individuals' behavior process and the unpredictability of health conditions. First, individuals' behavior execution is affected by their perception of the action course. For instance, two individuals may execute the same challenge differently, due to the diversity in their motivation levels, knowledge or understanding of the challenge instructions, habits, etc. Such personal variations in the behavior process create the first layer of noise in the delivered signal. Second, the complex nature of health introduces another layer of noise in the health-management process. Individuals' health condition

is jointly determined by many hidden health factors, including past disease histories, hormone level, metabolism rate, etc. As a result, health outcome can be unpredictable even if the same amount of effort is made. These noises further increase the difficulty of learning in the health-management process.

4.4 CONCLUSION AND DISCUSSION

Online behavior intervention programs are becoming increasingly popular in personal health management. Different from traditional caring processes, personal health management is centered on individuals' efforts in making a lifestyle change. Online behavior interventions provide detailed behavior-change plans and instructions to help users better engage in lifestyle regulation. Individuals' participation in online behavior intervention programs opens an opportunity for them to learn the association between behaviors and health outcomes and the effectiveness of interventions. In this study, we develop a structural model to characterize how individuals dynamically participate in online behavior intervention programs through learning.

The model we develop is a hierarchical learning model, which is able to capture the dynamic changes in individuals' posterior beliefs about behavior interventions. This modeling structure enable us to uncover individuals' sequential learning of how behavior interventions affect their health. In addition, our model incorporates both shared user characteristics and unobserved individual-level heterogeneity in the signaling process, which is able to well-capture personalized effect of health interventions.

We base our analysis on a leading online weight-loss platform in the US. The platform provides weight-loss challenges to help users intervene their lifestyles. Our investigation on the focal platform derives the following key findings. First, our estimation results show that individuals can form different prior beliefs on different intervention types, such as diet structure

change, calorie intake restriction, and exercise. Second, we find significant individual heterogeneity in intervention effectiveness, manifesting the importance of personalized healthcare. Third, we show that individuals can gradually learn the true intervention effectiveness and shift their beliefs during their participation; however, the learning speed is affected by the noise level in the delivered health signal. Fourth, in addition to expected intervention effectiveness, we find that individuals can form preferences on explicit intervention features, such as the length of intervention description and sentiment polarity. These findings provides important insights into user engagement and health promotion.

Our study contributes to the extant literature in the following ways. First, our study demonstrates that healthcare is not necessarily a credence goods (De Jaegher and Jegers 2001). Although prior studies have argued that the quality of medical care, such as treatment delivered by healthcare professionals, is usually difficult to infer, our study expands the existing literature by showing that the effectiveness of behavior interventions can be learned effectively, as behavior interventions emphasize individuals' efforts and typically involve continuous engagement. Second, our study adds to clinical health literature (Fulford 2011; Good et al. 2014) by showing individual differences in intervention effectiveness. Our results provide evidence of how individual heterogeneity and common demographics affect intervention effectiveness and, thus, support literature on healthcare personalization (Chawla and Davis 2013; Hopp et al. 2018). Third, our study connects to literature on online healthcare communities and Health 2.0. Over the recent decades, online healthcare has attracted great attention from not only academia but also industry. Researchers and practitioners are interested in understanding how users engage in online healthcare communities and how healthcare services and information can be better delivered through online channels. Our study propose a new perspective to investigate individuals'

engagement behaviors on online healthcare platforms; we show that individuals can dynamically form their perceptions on behavior interventions and adjust their participation strategies on the platform through learning. These findings provide insights into user engagement and empowerment.

Our study also provides practical implications. First, the learning process illustrated by our analysis indicates that user engagement is a dynamic course. As most users may not have adequate knowledge or information about the platform and its functionalities in the initial phase, it is important that online healthcare platforms provide decision-making support to help users engage and learn. Second, our findings reveal individuals' preference structures on healthcare intervention statements. Online healthcare platforms can use our findings to improve their behavior intervention design. For instance, as individuals are shown to dislike lengthy intervention descriptions, platforms can make their intervention design more concise; at the same time, platforms may increase the emotional intensity of the intervention descriptions to attract users' attention. Third, our analysis shows that individuals' learning efficiency is largely affected by noises in the delivered health signals. Thus, online healthcare platforms can develop tools to help users "de-noise" in the learning process. For instance, platforms can develop AI-based analysis tools to help users better understand the causes of their health variations or to provide more accurate learning signals (e.g., health status analysis report, etc.). Lastly, platforms may consider providing personalized intervention recommendations based on individuals' characteristics to help users achieve better health outcomes.

Chapter 5. PERSONALIZED RECOMMENDATION FOR HEALTHCARE DECISIONS: A MULTI-ARMED BANDIT APPROACH

Internet technologies enable information to be generated and disseminated at almost no cost, which accelerates the growth of information in online environments. Social media platforms, for example, allow users to share abundant content, including blogs, music, videos, and other formats, which individuals can freely choose to consume. Although various information options increase individuals' choice opportunities, having too many choices can be overwhelming and sometimes even confusing. Often, individuals experience difficulties in spotting content that is truly relevant to themselves or in which they are indeed interested (Konstan and Riedl 2012; Ricci et al. 2015). Such a choice overload issue can lessen users' experience and create barriers to individuals' engagement in online platforms.

Online healthcare communities (OHCs), which are social-media-based platforms that gather users with similar health-management interests, are no exception. Due to their easy access, OHCs are increasingly being used by individuals to learn about their illness, become familiar with treatment routines, and connect with others in similar circumstances. Typical OHCs provide users with various healthcare interventions to promote healthy behavior and improve adherence. Examples include behavioral treatment programs or plans that help individuals to establish healthy habits in regard to diet and physical exercise. During the recent COVID-19 pandemic, for example, individuals often engage in online work-out activities to relieve stress and stay healthy (Pew Research Center 2020). Individuals can freely choose interventions in which to participate in an online environment. When faced with too many choices, however, individuals may find it difficult to decide which option to take, as they may not know what would work or even what to expect,

especially when they are not healthcare professionals and do not have adequate experience in evaluating each choice. As a result, they may fall into analysis paralysis (Oulasvirta et al. 2009) and fail to engage in any health-management activities. This may harm their self-intervention adherence and outcome (Nutting et al. 2011; Snyderman and Dinan 2010).

The choice overload issue significantly affects one's participation experience or outcome in OHCs, leading to the pressing demand for services that can better fit individuals' healthcare needs. Therefore, in this study, we aim to follow the design-science paradigm to develop a personalized healthcare recommendation system as a means to support individuals' engagement in health management. Recommendation systems are intelligence-based algorithms that can help users to filter information and discover alternatives that they might not have found otherwise (Konstan and Riedl 2012; Vozalis and Margaritis 2003). Existing recommendation systems deploy various approaches to learn users' preferences from user-behavior data, such as collaborative filtering, content-based filtering, and hybrid models (Zhang et al. 2019). Research has shown that recommendation systems can effectively improve business performance and customer experience (Konstan and Riedl 2012; Pu et al. 2011) in ecommerce settings.

Despite their extensive use in ecommerce settings, whether and how recommendation systems can be integrated with online healthcare platforms has received little attention and remains largely underexplored. There are several unique patterns associated with users' health behaviors that create challenges in healthcare recommendations. First, previous health studies suggest that individuals' health behaviors are frequently affected by their evolving health status and health-management experiences (Johnson et al. 2002; King et al. 2006; Yan and Tan 2014). Thus, individuals' healthcare needs can exhibit strong temporal dynamics. Second, individuals' health management usually contains multi-dimensional effort, as promoting health requires individuals

to make a series of changes in all aspects of motivation and lifestyle. For instance, in weight management, individuals need to jointly monitor and manage different behavioral aspects, such as dietary behaviors and participation in physical activities. These patterns indicate that individuals' healthcare needs can be diverse, as individuals may need support for each type of health-management activity.

Given these unique patterns of individuals' health behaviors, conventional recommendation systems that are proven effective in ecommerce settings may not be effective in the healthcare context. This is because these algorithms generally exploit historical data to learn users' preferences. As individuals' health behaviors are continually changing, their health-behavior variations may not be fully captured by the historical data, especially when individuals' health-behavior data remain limited.¹ Thus, a mere exploitation of historical data may not be sufficient in healthcare recommendations. In addition, when individuals do not have well-established preferences about healthcare interventions, they may dynamically form their preferences based on the recommended items. The conventional recommendation systems do not take into account such interactions between users and recommendations and, thus, may not be effective in improving long-term recommendation performance (Liu et al. 2019). Finally, conventional recommendation systems are generally shown to over-specialize recommendations (Fleder and Hosanagar 2009; Pariser 2011). Thus, they may not well support users' diverse healthcare needs.

These research gaps motivate us to propose a novel recommendation design that utilizes a multi-armed bandit (MAB) as the main building block. An MAB is an online-learning framework in statistics and machine learning for solving decision-making problems in noisy or changing

¹ Different from ecommerce recommendation problems, data collection in healthcare recommendations can be a costly procedure, as healthcare is a long-term service, and individuals typically do not generate much health-behavior data when consuming healthcare goods.

environments (Auer et al. 2002; Chapelle and Li 2011). Specifically, when decision-makers (e.g., service providers) do not know the outcome of an action (e.g., recommendation), an MAB can help them to sequentially select choice alternatives while actively gathering information on each alternative's expected payoff (Zeng et al. 2016). In this process, an MAB strikes a balance between *exploiting* the learned knowledge to gain immediate rewards (reusing a highly rewarding alternative from the past) and *exploring* potential better alternatives (trying new or less-used alternatives to gather more information), which is known as the “exploitation-versus-exploration” tradeoff. By doing so, an MAB aims to maximize the cumulative reward during the entire decision-making period. In the healthcare recommendation context, service providers tend to have little knowledge about users' healthcare preferences, as individuals may constantly change their health behaviors and healthcare needs. Thus, the MAB framework can be used in such a setting to efficiently guide the learning of users' changing healthcare needs. In addition, through the exploration process, an MAB framework can promote the discovery of users' diverse healthcare needs, which may not be revealed by their historical behavior data.

To better adapt an MAB to the healthcare-recommendation context, we follow prominent health-behavior theories to further extend and enhance a standard MAB by synthesizing two model components, deep-learning-based feature engineering and diversity constraint. First, we design and implement two deep-learning models to extract user embeddings and item embeddings, which enables us to capture information that is critical to a healthcare decision-making context, such as users' health histories and health-behavior sequences (Johnson et al. 2002; King et al. 2006; Yan and Tan 2014) and intrinsic attributes of healthcare interventions. Taken together, the constructed user embeddings and item embeddings help to improve the personalization and contextualization of healthcare recommendations. The second model component is incorporated based on social

cognitive theory (SCT) (Bandura 1991; Bandura 1998). SCT proposes a classic paradigm for understanding individuals’ personal-influenced-based health-management behaviors. Based on SCT, we theorize the major dimensions of health management, and we use a diversity constraint to ensure that recommendations are structurally diversified along each of the health-management dimensions, so that individuals are provided with well-rounded support. To this end, we propose a Thompson sampling (TS)-based algorithm to solve this constrained recommendation task.

Our proposed recommendation framework is evaluated through a series of experiments, using data collected from a leading non-commercial online weight-loss platform in the United States. The focal platform provides weight-loss challenges to users, which are structured behavioral treatment programs to help users to manage short-term weight-loss goals, such as changing a dietary behavior, increasing physical exercise, and reducing weight in certain periods. We apply our recommendation framework to this weight-management setting to help users to find the most relevant challenges as a means to improve their engagement in weight-management activities. Our evaluation results suggest that each of our proposed model components is effective and that our recommendation framework significantly outperforms a wide range of benchmark models, including UCB, ϵ -greedy, and state-of-the-art conventional recommendation systems, such as context-aware collaborative filtering (CACF), probabilistic matrix factorization (PMF), and content-based filtering (CB). In addition, we demonstrate that our recommendation framework can more effectively learn the dynamics and the diversity distribution in users’ challenge choices. From users’ perspectives, we find that our recommendation design can serve to benefit a larger user population on the platform. Finally, we take a further step to evaluate our recommendation performance with respect to users’ weight-loss outcomes. The evaluation results suggest that our

proposed recommendation design can help users to achieve the highest average weight-loss rate compared to the benchmark models.

Our study makes several key contributions to the literature and practice. First, one major contribution of our study is the proposed healthcare recommendation framework, which demonstrates that prescriptive analytics can be integrated via a design-science artifact (Abbasi et al. 2016; Chen et al. 2012) to provide decision-making support for individuals' health management. The novel aspects of our recommendation framework include (1) a deep-learning-based feature engineering procedure, (2) a domain-knowledge-driven diversity constraint, and (3) a customized online-learning scheme. To the best of our knowledge, our study is among the first to combine an MAB with deep context representations and to introduce recommendation constraints for diversity promotion. Second, from a practical perspective, our recommendation framework can be applied to address real-world challenges in healthcare recommendations. Online healthcare platforms can adopt our recommendation design to improve users' health-management experience on the platform. Finally, the design of our recommendation framework can be further generalized to settings beyond healthcare. The online-learning scheme of an MAB enables decision-makers to adaptively adjust their strategies to minimize opportunity cost, and the deep-learning-based feature engineering procedure can help decision-makers to better understand the context-dependency of their decision results.

5.1 DESIGN THEORIES

We develop our recommendation algorithm based on prominent health theories. Prior health studies suggest that individuals' health behaviors are dynamically affected by a series of contexts, including their evolving health status, health-management experiences, and social context (Johnson et al. 2002; King et al. 2006; Yan and Tan 2014). These studies indicate that individuals

can form dynamic preference on healthcare interventions. In addition, according to the social cognitive theory, individuals' health management usually contain multiple dimensions, such as outcome-oriented dimension(s) and behavior-oriented dimension(s). The former influences individuals' motivation for health behaviors, whereas the latter affects the course of behavior execution. Corresponding to these dimensions, individuals may need different types of support to guide them through the self-regulation process. This indicates that individuals' healthcare needs can be diverse. Therefore, an effective health recommendation system should account for both the dynamics and the diversity in individuals' healthcare preferences.

5.2 MODEL DEVELOPMENT

In this study, we propose a deep-learning and diversity-enhanced MAB (DLDE-MAB) framework for offering healthcare recommendations. Specifically, we adopt an MAB as the main building block of our framework, as it can effectively explore variations in users' healthcare preferences and promote recommendation diversity at the same time. To better adapt MAB to the healthcare recommendation setting, we further enhance our framework by synthesizing two model components: deep-learning-based feature engineering and a diversity constraint. The deep-learning-based feature engineering procedure is designed to learn crucial recommendation contexts in regard to users' sequential health histories, health-management experiences, preferences, and intrinsic attributes of healthcare interventions. As individuals' health management is an evolving process, the sequential information embedded in individuals' health histories and health-behavior paths can play an essential role in shaping individuals' health behaviors. Deep-learning models can effectively capture patterns from dynamic temporal sequences and extract complex synergies between different features, thereby enabling the enhanced representation of variations in individuals' health behaviors. In addition, we theorize

individuals' major health-management dimensions based on the social cognitive theory and incorporate a diversity constraint to help us structurally diversify recommendations along different dimensions so that individuals are provided with well-rounded support.

In Figure 5.1, we provide a graphical illustration of our recommendation design. Each construct of the recommendation framework is intended to enhance healthcare recommendation performance. In a recommendation cycle, we first use deep-learning models to construct representations for users and items, i.e., the user embeddings and item embeddings. We then use the constructed embeddings to capture the contextual features of the recommendation environment, which enables us to learn the context-dependency of the recommendation results and generalize users' feedback. The MAB algorithm, shown on the right side of Figure 5.1, adaptively learns users' preferences by balancing the exploitation-versus-exploration tradeoff. The diversity constraint seeks to diversify the recommendations along the theorized health-management dimensions. We elaborate each of these constructs in the remainder of this section.




Figure 5.1 An MAB-Driven Framework for Healthcare Recommendation

5.2.1 *Optimizing Intervention with Multi-Armed Bandit*

OHCs provide healthcare interventions to encourage and instruct individuals' health behaviors. In Table 5.1, we provide several examples of typical healthcare interventions provided in OHCs. We consider the setting in which an online healthcare platform provides intervention suggestions to users on a regular (e.g., weekly) basis. It is unlikely that every individual user will prefer the same interventions, and there is interpersonal heterogeneity in terms of which interventions to adopt and to what extent. The goal of the platform is to adaptively suggest K interventions with the highest chance of improving individuals' engagement in their healthcare management.

Table 5.1 Example of Healthcare Interventions

1.	30 minutes jogging every day
2.	No food after 7 p.m.
3.	Avoid chips, soda and alcohol for 2 months

Formally, let T be the number of recommendation periods and I be the number of users on the platform. Suppose that each time the platform provides each user with K alternatives from a full available item set C_t . Let S_{it} denote the subset of items provided to user i at period t , $S_{it} \subseteq C_t$. C_t may vary with time. At the end of each period, the platform receives users' feedback on the recommendations, that is, whether they have adopted or engaged in the recommended interventions. Let $r_t(i, k)$ denote user i 's feedback on item k . Users' feedback serves as a reward for the platform's recommendation decisions; the platform may use the information of users' feedback to update its knowledge about users' preferences and adjust its subsequent recommendations.

We formulate the above recommendation problem as a contextual MAB. In a contextual algorithm, the decision of choice is leveraged upon a set of contextual features of the environment, such as the attributes of choice alternatives and user characteristics, so that the algorithm can

exploit the similarity between choice alternatives and deliver online personalized recommendations (Zeng et al. 2016). Contextual MABs learn to map the contexts into appropriate actions (Greenewald et al. 2017). Thus, they are able to personalize recommendations based on specific situations. In addition, as the pool of users and healthcare interventions may likely undergo frequent changes, it is desirable to learn a feature-based model that can generalize users' behavior histories to the user-item pairs that have never or rarely occurred in the past. To this end, in the healthcare recommendation context, we consider two sets of contextual information: individuals' health-management contexts $\mathbf{x}_{it}$ and attributes of healthcare interventions $\mathbf{z}_k$. We assume that users' feedback, i.e., the intervention engagement decision $r_t(i, k)$, is stochastically generated by an underlying probability that depends on the contexts. We model this probability as a logistic function:

$$E[r_t(i, k) | \mathbf{v}_{itk}] = (1 + \exp(-\boldsymbol{\theta}_*^T \mathbf{v}_{itk}))^{-1}, \quad (5.1)$$

where $\mathbf{v}_{itk} = \{\mathbf{x}_{it}, \mathbf{z}_k\}$, and $\boldsymbol{\theta}_*$ denotes the underlying coefficient vector, which can be learned adaptively in the recommendation process. The objective for the online healthcare platform is to maximize the expected cumulative user engagement during the entire course of recommendation, i.e.,

$$\max_{S_{it}} \sum_{t \in [T]} \sum_{i \in [I]} \sum_{k \in S_{it}} E[r_t(i, k) | \mathbf{v}_{itk}]. \quad (5.2)$$

5.2.2 Diversity Constraint

Modern recommendation systems should be well-diversified, motivated by the principle that recommending redundant items leads to diminishing returns on utility. In the context of healthcare recommendations, the major health-management dimensions that we identify based on SCT include the outcome-oriented dimension(s) and behavior-oriented dimension(s). Whereas the

former helps individuals to gain outcome-driven motivation, the latter enables them to acquire health-management skills and strategies. Thus, to provide individuals with well-rounded support, recommendations need to cover each of the health-management dimensions.

Although an MAB framework is able to promote recommendation diversity through the exploration process, we further incorporate a diversity-constraint MAB to ensure that the exploration is conducted in guided directions and that the recommendations are structurally diversified along each of the health-management dimensions. Formally, our diversity constraint can be expressed as

$$D = \{|S \cap \dim_{outcome1}| \geq 1, |S \cap \dim_{outcome2}| \geq 1, \dots, |S \cap \dim_{behavior1}| \geq 1, |S \cap \dim_{behavior2}| \geq 1, \dots\}, \quad (5.3)$$

where S denotes the recommendation set, $\dim_{outcome}$ denotes the outcome-oriented dimension(s), and $\dim_{behavior}$ denotes the behavior-oriented dimension(s). We subject the optimization problem in (5.2) to the diversity constraint D to ensure that the recommendation set S_{it} contains suggestions for each health-management dimension.

To solve this constrained recommendation task, we propose an algorithm that is adapted from Thompson Sampling (TS). TS is a machine-learning algorithm that addresses the exploitation-versus-exploration tradeoff presented in a bandit problem. TS is best understood in a Bayesian setting in which it computes the posterior distribution of the unknown parameters θ in the likelihood function, given the realized stochastic feedback. The rationale of TS is to encourage exploration through probability matching. That is, in each round, a TS algorithm randomly draws alternatives according to its probability of being optimal. Research has shown that TS generally has better empirical performance than do alternative bandit algorithms, such as UCB and ϵ -greedy (Chapelle and Li 2011). Our algorithm extends an ordinary TS by integrating a constrained

optimization problem to solve for the optimal recommendation decisions subject to the diversity constraint. We present the details of our algorithm below.

A TS Algorithm with Diversity Constraint

Input: prior mean m_j and prior variance ν_j for each parameter $\theta_j, j=1,2,\dots,d$.

For $t=1,2,\dots,T$ **do**

For $i=1,2,\dots,I$ **do**

Step 1 (Random Draw):

 Draw $\hat{\theta}_j$ from $N(m_j, \nu_j)$. $\hat{\boldsymbol{\theta}} = (\hat{\theta}_1, \hat{\theta}_2, \dots, \hat{\theta}_d)$.

For arm $k \in C_t$ **do**

 Compute expected reward: $\hat{r}_t(i, k) = (1 + \exp(-\mathbf{v}_{itk}^T \hat{\boldsymbol{\theta}}))^{-1}$.

End for

Step 2 (Optimization):

 Solve the following optimization problem:

$$\begin{aligned}
 & \max_{q_k, k \in C_t} \sum_{k=1}^K \hat{r}_t(i, k) q_k \\
 \text{s.t.} \quad & \sum_{k \in \text{dim}_{\text{outcome}}} q_k \geq 1, \quad \forall i, t \\
 & \sum_{k \in \text{dim}_{\text{behavior}}} q_k \geq 1, \quad \forall i, t \quad \text{Diversity constraints} \\
 & \sum_{k \in C_t} q_k \leq K \quad \text{Recommendation size constraint} \\
 & q_k = 0 \text{ or } 1, \quad \forall k \quad \text{Binary decision}
 \end{aligned}$$

 Offer item k to individual i if $q_k^* = 1$. $S_{it} = \{\forall k, q_k^* = 1\}$.

End for

Observe a new batch of data $(\mathbf{v}_{itk}, r_t(i, k))$, $i \in [I], k \in S_{it}$

Update the posterior mean by:

$$\mathbf{m} = \arg \min_{\mathbf{0}} \frac{1}{2} \sum_{j=1}^d \nu_j^{-1} (\theta_j - m_j)^2 + \sum_{i=1}^I \sum_{k \in S_{it}} \log(1 + \exp(-r_t(i, k) \boldsymbol{\theta}^T \mathbf{v}_{itk})).$$

Update the posterior variance by:

$$\begin{aligned}
 \nu_j^{-1} &= \nu_j^{-1} + \sum_{i=1}^I \sum_{k \in S_{it}} \nu_{jikt}^2 p_{ikt} (1 - p_{ikt}), \\
 \text{where } p_{ikt} &= (1 + \exp(-\mathbf{m}^T \mathbf{v}_{itk}))^{-1}, \nu_{jikt} \text{ is the } j\text{-th element of } \mathbf{v}_{itk}.
 \end{aligned}$$

End for

5.2.3 Deep-Learning-Based Feature Engineering

To improve the characterization of individuals' health-management contexts and enhance recommendation personalization, we design a deep-learning model to construct user embeddings. Specifically, our user-embedding model leverages information on users' attribute features (e.g., gender, age, etc.), health-status trajectories, and health-management behavioral sequences. For

each user, the attribute variables usually remain unchanged over time, whereas health status and behavioral sequences will vary with time. Hence, the user embeddings depend on both user i and time t to reflect the evolving dynamics. To properly guide the learning on these aspects, we propose a novel wide-and-deep neural network. The wide-and-deep structure was originally proposed for user response modeling in mobile apps. It combines two branches of user features (i.e., a “wide” branch and a “deep” branch) to facilitate user representation learning (Cheng et al. 2016). In this study, we design a “wide” branch to process users’ attribute features to take into account that certain intervention suggestions can be more actionable for specific users given their personal attributes, and we apply a fully-connected structure to account for possible interactions among the attribute features.

We then use a “deep” branch to learn sequence features, such as users’ health-status trajectories, historical healthcare-intervention adoptions, and other health-management experiences in regard to self-monitoring activities and social behaviors. The sequence of historical intervention adoptions is included to capture the dynamics in users’ preferences. Together with the health-status trajectories, it captures users’ evolving health histories and the corresponding changing healthcare preference. In addition, based on prior health theories mentioned in Section 2, health-management experiences, such as social supports and self-monitoring activities, may also affect individuals’ health behaviors and thus influence their preferences for healthcare interventions. Therefore, we further include related behavior paths to capture the effect of health-management experiences on users’ intervention-adoption behaviors. We use Long Short-Term Memory (LSTM) with a self-attention mechanism to capture the dynamically changed patterns in these features and their correlation with adopted healthcare interventions. To address the fact that different sequence features have different dimension scales (e.g., number of social activities vs.

87

In addition, we enhance the learning process of the deep branch by an auxiliary loss function with the healthcare outcome as the goal. We use the auxiliary loss function to incorporate the effects of healthcare outcomes on individuals' health behaviors and individuals' preferences on healthcare interventions. This is unique to the healthcare recommendation context, in which individuals' health behaviors are fundamentally driven by the goal of optimizing health outcomes. Meanwhile, the auxiliary loss branch can also guide the neural network to properly extract signals from the sequence features, as, otherwise, the gradient flow will not be balanced and the shallow structure will dominate the gradient flow. An illustration of the proposed deep learning architecture for user embedding construction is provided in Figure 5.2.




Figure 5.2 Illustration of the Base Architecture

Finally, as healthcare interventions provided online are usually in the form of textual information, we build another deep-learning model to construct item embeddings to extract the inherent semantic features embedded in intervention information. For example, a typical online healthcare intervention may contain a short title that highlights the main features of the intervention and a description or instruction that describes the detailed execution procedure. To capture the semantics embedded in this information, we propose a hybrid model, in which we apply LSTM to learn the semantics of the intervention descriptions, and we use average token-level embedding to extract signals from the intervention titles. The outputs are the meta attributes of the intervention (e.g., duration, category, and/or intensity of the intervention). We further fine-tune the token-level embedding in the procedure of representation learning to ensure low information loss.

In sum, the user embeddings and item embeddings help us to learn key contextual information for healthcare recommendations concerning users’ health behavior contexts and item attributes, which are then used as the input of our bandit recommendation model. Due to the space limit, more details on our user-embedding and item-embedding models are provided in Appendix C.1.

5.3 EVALUATION DESIGN

5.3.1 *Focal Evaluation Content*

To evaluate the performance of our recommendation framework, we collected data from a leading non-commercial online weight-loss platform in the United States. Weight management has garnered considerable attention from healthcare communities due to the increasing prevalence of overweight and obesity in the global populations. Research indicates that overweight and obesity are associated with multiple health risks, such as type 2 diabetes, cardiovascular diseases, and cancer (WHO 2018). Recently, the CDC (2020) announced that obesity is a risk factor for severe COVID-19 complications. Although there are genetic and metabolic influences on body weight,

overweight and obesity are likely to occur when individuals develop unhealthy behaviors, such as poor dietary habits and a lack of physical activities, which impede individuals from balancing their calorie intake and expenditure (Franz 2001; Hall and Kahan 2018).

To help users to establish a healthy living style, the focal platform provides weight-loss challenges, which are behavioral treatment programs that help users to focus on a specific weight-loss goal in a short time period. Examples of weight-loss challenges include diet-oriented challenges, such as “cut off processed carbs and include 100g of mixed veg in every meal,” and activity-oriented challenges, such as “30 minutes of jogging every day.” The diet-oriented and activity-oriented challenges provide behavioral guidance for individuals’ weight-management routine. Users can also find weight-loss-oriented challenges that help them to set goals directly for weight changes, such as losing a certain amount of weight during specific periods. Participation in weight-loss-oriented challenges can help users to establish outcome-driven motivation. Each weight-loss challenge is defined by a short title and description that contains information on the challenge goal, duration, and instructions. Users can choose to join any challenge as long as its starting date has not passed. In Appendix C.2, we provide a screenshot of the challenge webpage to show how users can retrieve challenge information from the online platform.

The focal platform does not incorporate any recommendation system to facilitate users’ challenge selection. During our investigation period, there were more than 100 challenges provided to users. Users may likely find it difficult to select challenges to join, as users may lack the ability to discern, from various choice alternatives, what challenges are suitable for their weight management. In addition, a significant search cost is expended, as users need to spend time reading the challenge descriptions and/or instructions before deciding which challenge to join. These problems can potentially be solved by providing personalized challenge recommendations to

support healthy behaviors. Our investigation on weight-loss challenge recommendations can help to improve the match between individuals and weight-loss challenges and, thus, improve individuals' weight-management performance. From the platform's perspective, the recommendations can enhance users' participation experience and, thus, contribute to user maintenance and platform sustainability.

5.3.2 *Data*

We collected three datasets to support our investigation of recommendation performance. The first dataset contains descriptive information for each challenge provided on the platform. During our data collection window,² there were 165 challenges provided on the platform in total. For each challenge, we collected the title, description, and duration. To capture challenge characteristics, we apply SMART metric to extract a series of meta attributes from the text data. The SMART metric (i.e., specific, measurable, attainable, relevant, and time-bound) has been widely used as a gold standard in areas such as education and healthcare for assessing the quality of goals (Doran 1981; Ogbeiwi 2018). This metric can help individuals to clearly identify the direction for logical action planning and implementation (Ogbeiwi 2017; Ogbeiwi 2018). Thus, SMART-based challenge features are critical in the current weight management context, in which users aim to manage short-term weight-loss goals.

Based on the SMART metric, we calibrate the following challenge meta attributes: whether the challenge is specifically defined (specificity), whether the challenge goal is measurable (measurability), the intensity level of the challenge (attainability), whether the challenge is related to diet or physical activity (relevancy), and the time span of the challenge (duration). In addition

² Our data collection period was January 1, 2014 to April 30, 2014.

to SMART-based attributes, we consider two other features that may affect individuals' engagement in challenge participation: *Motivational* and *Self-Monitoring*. *Motivational* characterizes challenges from the perspective of goal statement, which has been suggested to be important in helping individuals to build up inner motivation (Locke and Latham 1990). *Self-monitoring* is an important step in goal fulfillment, as it helps individuals to process their performance toward goal achievement (Bandura 1991). We use *Self-Monitoring* to indicate whether a challenge encourages individuals to regularly monitor and report their weight-loss progress. Together, these attributes help us to quantify the similarities among challenges. In Table 5.2, we provide a summary of the annotated challenge meta attributes, which will be used for learning challenge-embedding representation and downstream recommendation task. The detailed annotation procedure is provided in Appendix C.3.

Table 5.2 Challenge Meta Feature Annotation

Features	Definition
<i>Specific</i>	Whether a challenge is specifically defined (0 or 1)
<i>Measurable</i>	Whether a challenge goal is measurable (0 or 1)
<i>Diet</i>	Whether a challenge is related to dietary behaviors (0 or 1)
<i>Intensity_Diet</i>	Intensity level for a diet-oriented challenge (L, M, H)
<i>Activity</i>	Whether a challenge is related to physical activities (0 or 1)
<i>Intensity_Activity</i>	Intensity level for an activity-oriented challenge (L, M, H)
<i>Weight-Loss</i>	Whether a challenge contains a goal for weight changes (0 or 1)
<i>Intensity_Weight_Loss</i>	Intensity level for a weight-loss-oriented challenge (L, M, H)
<i>Motivational</i>	Whether a challenge contains motivational words/sentences (0 or 1)
<i>Self-Monitoring</i>	Whether a challenge requires individuals to regularly monitor and report their weight-loss progress, e.g., body weight, daily diet, running mileage (0 or 1)
<i>Duration</i>	Time span (in weeks) of a challenge

The second dataset is users' challenge-selection histories; that is, we recorded for each user the challenge(s) that he or she selected per week. This dataset enables us to learn users' preferences for weight-loss challenges. We find that users on the focal platform chose two challenges on average per week. When users chose any challenge, they chose multiple challenges about 70% of

the time. As noted, there are three major challenge types on the platform: weight-loss oriented, diet oriented, and exercise oriented. We find that users tend to choose different types of challenges whenever they choose multiple challenges. Specifically, users choose more than two challenge types 92% of the time when they choose multiple challenges; they choose all three types of challenges about 51% of the time. These results indicate the existence of diversity in users' preferences for weight-loss challenges. In addition, we find that users' selection of challenge types drifts over time. That is, users may have preferred certain challenge types at the beginning of a time period and gradually shift to other challenge types as time goes by. These findings provide evidence for the dynamics of users' preferences, which may be due to users' transitions to different weight-loss statuses, in which they need different types of support. It is also likely that users gradually establish their personal tastes in regard to weight-loss challenges during the process of challenge participation. These findings thus provide support for our recommendation design.

The third dataset contains auxiliary information for each user, such as gender, age, membership duration, initial weight when first joining the platform, online weigh-in activities, the number of friends, and the posts published in the community forum. We use this information to capture users' heterogeneous weight-management contexts. In particular, gender and age are two factors that directly affect individuals' weight status. Membership duration measures users' overall weight-management experiences on the platform. Initial weight and weekly weigh-in records help us measure individuals' weight-loss status and health histories. The number of friends and forum posts provide proxies for the amount of social support available to individuals (Shumaker and Brownell 1984; Yan 2018); thus, we use them to capture the social contexts of users' weight management. We provide a summary of key data statistics in Appendix C.4.

5.3.3 Operationalizing Model Components: Embeddings and Diversity Constraint

As noted in Section 5.2.3, our user-embedding model adopts a wide-and-deep network structure, in which we use the wide branch to capture users' attribute features and the deep branch to capture the sequence features. In our evaluation context, users' attribute features include gender, age, initial weight, and membership duration. The sequence features include three parts. The first part captures users' health status, that is, their historical weight variations. The second part is the sequence of historical challenges chosen by individuals, which we use to account for users' personal tastes. The third part concerns users' other behavioral sequences, such as their past social activities (e.g., establish friendships with other users and publish forum posts) and self-monitoring activities (e.g., weigh-in). The auxiliary loss head is designed to measure the weight loss in the next time period, where we choose a combined loss of MSE for absolute value prediction and cross-entropy loss for weight-loss sign prediction.

For our challenge-embedding model, we use challenge name and description as the inputs, and the annotated challenge meta attributes based on the SMART metric are the outputs. The annotated challenge meta attributes help us to depict the key characteristics of a weight-loss-related goal; thus, they provide a good standard for calibrating challenge similarity in our focal context.

In operationalizing our diversity constraint, we identify weight loss as the outcome-oriented dimension, as it is the health goal that individuals aim to achieve in our focal context. We identify diet and physical exercise as two behavior-oriented dimensions, as they are two essential behavioral-regulation aspects in weight management. With respect to these dimensions, we ensure that our recommendations cover all three challenge types, i.e., weight-loss oriented, diet oriented, and exercise oriented. Therefore, our diversity constraint is specified as

$$D = \left\{ \left| S \cap \dim_{\text{weightloss}} \right| \geq 1, \left| S \cap \dim_{\text{diet}} \right| \geq 1, \left| S \cap \dim_{\text{exercise}} \right| \geq 1 \right\}, \quad \text{where } S \text{ represents the}$$

recommended challenge set, $dim_{weightloss}$ represents the weight-loss-oriented dimension, dim_{diet} denotes the diet-oriented dimension, and $dim_{exercise}$ denotes the exercise-oriented dimension.

5.4 PERFORMANCE ANALYSIS AND RESULT

We apply our recommendation framework to the weight-management context described in Section 5.3, with the aim of promoting users’ engagement in weight-loss challenges on the platform. In particular, we implement the algorithm introduced in Section 5.2.2 to offer top-K challenges to users on a weekly basis.³ The time span of recommendation is 16 weeks (i.e., the same as our data collection window). To demonstrate the effectiveness of our recommendation design, we follow the design-science paradigm to rigorously evaluate our recommendation framework through a series of experiments. We first examine the effectiveness of our deep-learning embeddings in capturing user characteristics and challenge attributes. We then apply different evaluation approaches to test each of our model components as well as to compare our model against state-of-the-art recommendation systems.

5.4.1 *Experiment 1: t-SNE Visualization of Context Space*

To show how the construction of deep-learning models improves feature engineering, we use t-SNE to visualize the embeddings in a two-dimensional space (Maaten and Hinton 2008). t-SNE is a nonlinear dimensionality reduction technique that is well suited for deconstructing high-dimensional data. It can project high-dimensional vectors into lower dimensions without changing the data structure so that it helps us to understand data patterns in a more intuitive way. In a t-SNE plot, two points that are close to each other indicate that the corresponding embedding

³ In our main analysis, we provide the results for $K = 10$. In Appendix C.5, we provide the top-K ($K = 5, 10, 15$) precision, recall, and F-measure.

representation vectors are similar. We provide the visualization results for challenge embeddings and user embeddings in Figures 5.3 and 5.4, respectively.

Specifically, we present the t-SNE plot for our constructed challenge embeddings in Figure 5.3(1). In Figures 5.3(2) and 5.3(3), we provide the t-SNE plots for two state-of-the-art word2vec deep-learning models, BERT and FastText. We use these two models as a benchmark for evaluating the performance of our proposed challenge-embedding model. In the plot, we use different colors to indicate each challenge type, such as weight-loss oriented, diet oriented, and exercise oriented. We use the color degree to denote the intensity level of a challenge: a deeper color indicates a challenge of higher intensity. For example, diet-oriented challenges are denoted by orange points, and among the points, there are three color degrees: light orange, dark orange, and orange-red, representing low-, medium-, and high-intensity levels, respectively. We find that, in Figure 3(1), challenges that belong to the same type are tightly clustered. In addition, within each challenge-type cluster, challenges of the same intensity level tend to be close to each other. These patterns indicate that our challenge embedding model can well capture intrinsic challenge attributes, especially the ones that are key to goal-setting theory (Doran 1981). In comparison, the benchmark models cannot clearly distinguish these challenge patterns.




Figure 5.3 Challenge Embedding t-SNE Visualization

The t-SNE plot for our constructed user embeddings is presented in Figure 5.4(1). As our deep-learning model incorporates users' time-varying sequence features, such as their weight curves and activity paths, each user may generate different embeddings at different times, which are denoted by different points. To show how our model captures user patterns, we use different colors to distinguish users' gender and age groups.⁴ Specifically, the pink points in the plot denote female users, and the blue points denote male users; the deeper the color, the larger the age group to which users belong. In addition, we use different marker shapes to denote users' in-period weight-loss status, i.e., weight loss, weight maintenance, weight gain, and unknown status. Our visualization results show that users tend to gather into sequence clusters if they have similar features noted above. Specifically, each sequence cluster typically comprises the same gender. Users of a larger age group are typically clustered at the two sides of the plot (deeper colors), whereas users of a smaller age group generally locate in the middle (lighter colors). Finally, users whose weight remains unchanged tend to show up together, and users who have weight changes

⁴ Age groups are encoded as follows: Group 0: ≤ 29 , Group 1: 30–39, Group 2: 40–49, Group 3: 50–59, Group 4: 60–69, Group 5: > 69 .

tend to gather elsewhere. These results indicate that our model can well capture user patterns in terms of gender, age, and weight-loss status. Note that these features are most directly related to our weight-loss context; in particular, gender and age are two demographics that directly affect users' body weight, and weight-loss status reflects users' in-period weight variations.

We compare our proposed user-embedding model with two benchmark models, `collab_learner` and `tabular`. The results are presented in Figures 5.4(2) and 5.4(3), respectively. The `collab_learner` model and the `tabular` model are two encapsulated python learners provided in the `Fastai` library. In particular, the `collab_learner` model learns user representations from the historical challenge-selection data; however, it is not able to incorporate users' personal characteristics. The `tabular` model is a deep-learning model that learns user embeddings based on users' tabular attributes, such as gender and age, but does not extract signals from sequential data, such as users' health histories and behavior paths. As shown, these two models do not perform as well as our models, as there is no explicit user pattern displayed.




Figure 5.4 User Embedding t-SNE Visualization

Finally, the sequence shape of the user clusters produced by our model in Figure 5.4(1) motivates us to further investigate granular individual-level patterns, as the points in a sequence are likely generated by the same or similar users. We randomly sampled several individual users and plotted their corresponding embeddings in Figure 5.4(4). We find that the embeddings of the same user locate close to each other and tend to be concatenated into a trajectory and that the embeddings of different users are located relatively far apart. These results show that our proposed user-embedding model can effectively capture the sequential changes in users’ behaviors. Together, Figures 5.3 and 5.4 provide evidence for the effectiveness of our deep-learning models.

5.4.2 *Experiment 2: Ablation Analysis of Model Components*

We conduct an ablation analysis to compare our model with a series of baseline MABs. This allows us to show that each of our model components (i.e., the deep-learning-based feature engineering procedure and the diversity constraint) is effective in helping users to find the relevant challenges. The baseline MABs are the counterpart MAB models that partially incorporate or do not incorporate the proposed model components. Specifically, the baseline MABs that we investigate include the MAB without user embeddings, the MAB without challenge embeddings, the MAB without either embeddings, the MAB without diversity constraint, and the MAB without either embeddings or constraint. When user embeddings are not incorporated, we use users’ attribute features (e.g., gender, age, etc.) to account for recommendation personalization. When challenge embeddings are not incorporated, we use the annotated challenge features to capture the inherent attributes of challenges, namely, the variables listed in Table 5.2.

Evaluating an explore/exploit policy is difficult because we typically do not know the reward of an action that was not chosen. Possible solutions include doubly-robust estimation (Dudík et al. 2011), offline precision evaluated by preference set (Qin et al. 2014), and simulation. The first evaluation approach, doubly-robust estimation, is an offline data evaluation approach that utilizes

pre-collected historical data to evaluate policy performance. The historical data are assumed to contain three sets of information: action, context, and reward. As the data are pre-collected, we are able to observe only the rewards for the chosen action in the data. To adjust for the potential bias caused by the data collection process, the doubly-robust estimator combines two policy evaluation methods, direct simulation (DS) and inverse propensity score (IPS). Formally, let G denote the offline dataset, which contains action a , context $\mathbf{v}$, and reward r . In our context, action refers to the platform's provision of a challenge in the data, context $\mathbf{v}$ includes individuals' weight-management context $\mathbf{x}_{it}$ and the challenge features, and r is users' feedback, that is, whether a user selects a challenge. Let S_{it} denote the set of challenges recommended to user i at week t , and $|S_{it}| = K$. Our doubly-robust estimator can be expressed as follows:

$$\bar{R} = \frac{1}{|G|} \sum_{\{i,t,a\} \in G} \left\{ \frac{1}{K} \sum_{k \in S_{it}} \left[\hat{\phi}_{\mathbf{v}_{ik}} + \frac{(r_{a,\mathbf{x}_{it}} - \hat{\phi}_{\mathbf{v}_{ia}})I_{[k=a]}}{\hat{p}(a|\mathbf{x}_{it})} \right] \right\}, \quad (5.4)$$

where $\hat{\phi}$ is a reward simulator, and $\hat{p}$ is the propensity of challenge provision in the data. The rationale of this method is that, when data are not available, the method uses a pre-trained reward predictor to simulate the reward; otherwise, it applies a correction to the reward predictor using the actual data.

The second evaluation method measures recommendation precision. Following previous studies (Qin et al. 2014; Qin and Zhu 2013), we construct each user's preference set as the set of challenges selected by the user in the data. The recommendation precision is thus the overlap ratio between the recommendation set and the preference set. Finally, in light of previous theoretical MAB studies (Hertz et al. 2018; Sani et al. 2012), we examine our model performance through a simulated environment. Specifically, we construct a logistic predictor for users' binary challenge-selection decisions, that is, $r_t(i,k) = (1 + \exp(-\mathbf{v}_{itk}^T \boldsymbol{\zeta}))^{-1}$, where $\mathbf{v}$ is a concatenation of user

embeddings and challenge embeddings. The weight vector ζ could have been chosen arbitrarily, but it was in fact a perturbed version of the weight vector trained on a randomly constructed training set (Nguyen et al. 2017), and the performance evaluation is conducted on a test set. This simulator is omniscient, in the sense of full knowledge of users’ preferences and the actual amount of reward accrued by recommendations. We provide the details of these evaluation approaches in Appendix C.6.

The evaluation results are provided in Table 5.3. Each value in the table represents users’ average selection rate during the entire recommendation course. A superscripted asterisk denotes that a benchmark model performs significantly worse than the proposed model. Our results show that the MABs that include only one of the components have an inferior performance (i.e., lower average challenge-selection rate) than our proposed model. Specifically, the MAB without diversity constraint is suggested to be significantly worse by the doubly-robust estimation and simulation method. The MAB without user embeddings or challenge embeddings (or both) is shown to have a worse performance by all three evaluation methods. Finally, we find that the MAB without either embeddings or diversity constraint performs worse than the MABs that partially incorporate the model components. These results indicate that each of our proposed model components is effective. As compared to users’ attribute features, user embeddings can better capture the sequential information embedded in users’ health histories and behavior paths and, thus, are more effective. Challenge embeddings are more effective than the annotated challenge features, as they are able to extract semantic information from the textual challenge descriptions. In addition, as most annotated challenge features are categorical and one-hot encoded, they may not provide much information for the learning process. In comparison, challenge embeddings can better calibrate the similarity among challenges and make the learning more effective.

Table 5.3 Comparison of Baseline MABs

Models	Doubly-Robust Estimation	Offline Precision	Omniscient Simulator
proposed MAB	0.5115	0.6054	0.4467
<i>Baseline MABs</i>			
No user embeddings	0.4983 **	0.5933 *	0.4275 **
No challenge embeddings	0.4722 ***	0.5625 ***	0.4215 ***
No embeddings	0.4685 *	0.5519 ***	0.4206 ***
No diversity constraints	0.5028 ***	0.5974	0.4394 *
No embeddings, no constraints	0.4583 ***	0.5462 ***	0.4172 **

Note: Asterisk in superscript denotes that a benchmark model performs significantly worse than the proposed model. Significance levels are: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

In Figure 5.5, we plot the recommendation performance across time to show the learning curve of each model. The x -axis denotes recommendation rounds, and the y -axis denotes the average challenge selection rate up to round t . It is shown that our model achieves the highest learning rate across all periods, regardless of the evaluation approach taken. That is, our model can boost the average challenge-selection rate faster than the benchmark models can. For example, when evaluated by simulation, our model increases the average challenge selection rate from approximately 0.39 to approximately 0.45 after the 16-week recommendation phase, which is a 15% increase. This is followed by the MAB with no diversity constraint (~13%), the MAB with no user embeddings (~11%), the MAB with no challenge embeddings (~10%), the MAB without either embedding (~8%), and the MAB without either diversity constraint or embeddings (~8%). These results further highlight the effectiveness of our deep-learning-based feature engineering and diversity constraint in the learning procedure.




Figure 5.5 Learning Curve of Our Model and Baseline MABs

5.4.3 Experiment 3: Comparison with Benchmark Recommendation Systems

We compare our proposed recommendation framework against a wide range of benchmark models. In particular, for benchmark bandit models, we consider UCB and ϵ -greedy, which are two classic online-learning methods to solve the “exploitation-versus-exploration” tradeoff. For batch-learning-based models, we consider a variety of collaborative filtering methods, such as context-aware recommenders and matrix-factorization-based models. We also consider content-based filtering and hybrid filtering. Content-based filtering is able to offer recommendations based on item features. Hybrid filtering further combines content-based filtering with collaborative filtering to incorporate information embedded in users’ challenge selection histories. As the batch-learning-based models make recommendations by exploiting users’ item-selection histories, we use the first four weeks of data to train the algorithms. Finally, we consider pure exploitation and pure exploration, which are two recommendation schemes that do not seek a balance between exploitation and exploration. We summarize our benchmark models in Table 5.4. The implementation details of these benchmark models are provided in Appendix C.7. Similarly, we

calibrate recommendation performance by users' average challenge-selection rate and evaluate it through the aforementioned three evaluation approaches, i.e., doubly-robust estimation, offline precision, and simulation.

Table 5.4 Summary of Comparison Benchmark Recommendation Systems

Benchmark Models	Recommendation Algorithm Description
UCB	A bandit algorithm that incorporates the opportunity for exploration by choosing the arm with the highest upper bound of the reward confidence interval;
ε -greedy	A bandit algorithm that chooses the arm with the seemingly highest average reward with probability $1 - \varepsilon$ and explores a random arm with probability ε ;
CACF	Context-Aware Collaborative Filtering, which incorporates the contexts of users' item selections as weights into a normal collaborative filtering procedure (Chen 2005);
SCF	Social Collaborative Filtering, which formulates a neighborhood-based method for cold-start collaborative filtering in a generalized matrix algebra framework (Sedhain et al. 2014);
PMF	Probabilistic Matrix Factorization, a model-based collaborative filtering approach that uses matrix factorization under a probabilistic framework to estimate user-item interactions (Mnih and Salakhutdinov 2008);
CAMF	Context-Aware Matrix Factorization, which is an extension of the classic matrix factorization approach for incorporating contextual information (Baltrunas et al. 2011);
CB	Content-based filtering, an approach to offer recommendations based on content similarities of items (Bieliková et al. 2012; Pon et al. 2007);
hybrid_pure	A hybrid model that combines pure collaborative filtering with CB using mixed hybridization (Burke 2002);
hybrid_cacf	A hybrid model that combines CACF and CB using mixed hybridization;
pure exploitation	A model that selects the best option given current knowledge;
pure exploration	A model that fully randomizes recommendations.

Our evaluation results are provided in Table 5.5. The results show that our proposed model has the best recommendation performance under all evaluation measures. It achieves an average challenge selection rate of 51.15% when evaluated by doubly-robust estimation; 60.54%, by offline precision; and 44.67%, by simulation. UCB and ε -greedy are shown to have comparable performance (~42% for doubly-robust estimation, ~52% for offline precision, and ~39% for simulation), but both perform significantly worse than our model. Batch-learning-based models generally do not perform well (mainly below 40%). The performance inferiority may be due to users' dynamic preferences for weight-loss challenges. In other words, the batch-learning-based models assume that users' preferences can be well represented by their past behavior patterns

(Adomavicius and Tuzhilin 2005; Sahoo et al. 2012); thus, these models can be biased when users' preferences contain dynamic variations that are not fully captured by the historical behavior data.

Table 5.5 Comparison with State-of-the-Art Recommendation Systems

Models	Doubly-Robust Estimation	Offline Precision	Omniscient Simulator
proposed MAB	0.5115	0.6054	0.4467
<i>Baseline State-of-the-Art Bandit Models</i>			
UCB	0.4199 ***	0.5266 ***	0.3981 ***
ε -greedy	0.4193 ***	0.5242 ***	0.3848 ***
<i>Baseline State-of-the-Art Batch-Learning Models</i>			
CACF	0.4294 ***	0.3149 ***	0.4325 **
SCF	0.3696 ***	0.3970 ***	0.4005 ***
PMF	0.3618 ***	0.2124 ***	0.3943 ***
CAMF	0.3100 ***	0.2653 ***	0.2824 ***
CB	0.4015 ***	0.2660 ***	0.4330 **
hybrid_pure	0.2951 ***	0.2492 ***	0.3046 ***
hybrid_cacf	0.4109 ***	0.2897 ***	0.4338 **
<i>Pure Exploitation and Pure Exploration</i>			
pure exploitation	0.4785 **	0.5760 ***	0.4188 ***
pure exploration	0.4617 ***	0.5690 ***	0.4006 ***

Note: Asterisk in superscript denotes that a benchmark model performs significantly worse than the proposed model. Significance levels are: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Finally, our results show that pure exploitation and pure exploration achieve worse recommendation performance as compared to our model. This performance gap indicates the importance and necessity of balancing the exploitation-versus-exploration tradeoff. A pure exploitation method may be stuck in a worse local optimum. Pure exploration, in contrast, over-explores users' preferences; it fully randomizes the recommendations without utilizing or learning from information embedded in users' past behaviors. Note that pure exploitation and pure exploration can often be seen in the design of A/B testing. Specifically, in an A/B test, experimenters first spend a short time period for pure exploration, whereby they randomly assign users to different groups to examine the performance of policy variants. They then engage in a long period of pure exploitation, assigning all of the users to the group that achieves the best performance. In practice, the pure exploratory phase can be expensive or even infeasible to

implement. For example, in a health-management context, it is usually infeasible to arbitrarily assign individuals to a treatment plan. Instead of two distinct periods of pure exploration and pure exploitation, a bandit-driven design adaptively combines exploration and exploitation. Thus, it can reduce the opportunity cost incurred in the exploratory phase and help service providers to achieve better performance.

5.4.4 *Experiment 4: On Dynamic User Preferences*

In a typical healthcare context, individuals' behavior patterns likely change over time (Johnson et al. 2002; King et al. 2006). In this experiment, we investigate the recommendation results for the users whose choices tend to vary considerably as a means to examine whether our recommendation framework can well capture the dynamics in users' preferences. From the test-user set, we select the 30 users whose challenge choices have the largest embedding variance.⁵ We implement our recommendation algorithm for the selected users and compare the recommendation performance with that of using the full test-user set. Table 5.6 presents the results for the new test-user set. Our model is shown to outperform all of the benchmark models by all evaluation measurements. In particular, our model achieves an average challenge selection rate of 53.44% when evaluated by doubly-robust estimation; 60.90%, by offline precision; and 47.56%, by simulation.

⁵ As challenge embeddings are vectors, we define the variance of embeddings as the minimum value of the variances of the elements.

Table 5.6 Results for Dynamic Test Users

Method	Doubly-Robust Estimation	Offline Precision	Omniscient Simulator
proposed MAB	0.5344	0.6090	0.4756
<i>Benchmark Models</i>			
UCB	0.4287 ***	0.5266 ***	0.3981 ***
epsilon greedy	0.4415 ***	0.5242 ***	0.3848 ***
CACF	0.4183 ***	0.2278 ***	0.4378 ***
SCF	0.3524 ***	0.3964 ***	0.3783 ***
PMF	0.3494 ***	0.1975 ***	0.3894 ***
CAMF	0.3061 ***	0.2567 ***	0.2800 ***
CB	0.3694 ***	0.2317 ***	0.4350 ***
hybrid_pure	0.2772 ***	0.2208 ***	0.3005 ***
hybrid_cacf	0.3942 ***	0.2250 ***	0.4308 ***

Note: Asterisk in superscript denotes that a benchmark model performs significantly worse than the proposed model. Significance levels are: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

To better show the performance variations, we plot the differences between the new recommendation results and the original results on the full test set in Figure 5.6. Here, we present the performance changes measured by doubly-robust estimation. The performance changes measured by offline precision and simulation are similar and are provided in Appendix C.8. We find that bandit-driven models, such as our proposed model, UCB, and ϵ -greedy, have different degrees of performance increase. In contrast, batch-learning-based models generally have a performance decrease. These results suggest that the advantage of the online-learning scheme is further strengthened when evaluated on the dynamic users whose preferences tend to vary frequently. The performance decrease of batch-learning-based models indicates that the “first learn, then earn” recommendation scheme may perform even worse when users’ preferences exhibit strong dynamics. This may be because the dynamic patterns in users’ preferences cannot be fully captured by the historical data. Bandit-driven models, in contrast, are able to actively collect users’ feedback on recommendations and, thus, can capture changes in users’ preferences more promptly.




Figure 5.6 Performance Variation for Dynamic Users

5.4.5 Experiment 5: Diversity Analysis

We evaluate the recommendation diversity across challenge types to examine whether our model can effectively learn users' diverse challenge preferences in the data. Specifically, for our proposed model and each of the benchmark models, we calculate the recommendation frequency for each challenge type to construct a diversity distribution. We then compare the recommendation frequencies with users' challenge selection frequencies in the data, which reveal users' true preferences of challenge types. We use the Jensen-Shannon divergence (JSD) to measure the similarity between two diversity distributions (Endres and Schindelin 2003; Fuglede and Topsoe 2004). A small JSD value indicates high similarity between two diversity distributions.

In Figure 5.7, we visualize the diversity distribution for each recommendation model and present the corresponding JSD values. The first bar in the figure represents the diversity distribution in users' challenge-selection histories observed in the data. The second bar represents the diversity distribution in the recommendations provided by our proposed model. As can be seen, the first two bars are very similar to each other, indicating that the recommendations produced by our model are diversified in a way that is similar to users' actual challenge-selection histories. The

recommendation diversity distributions produced by the benchmark models generally have a larger difference from the observed challenge-selection data. For example, the bar that corresponds to CACF is quite different from the first bar. These observations are confirmed by our JSD results, that our proposed model has the smallest JSD value (i.e., 0.03) across all of the benchmark models. Combined with the results in the earlier-discussed experiments, these findings provide evidence that our proposed recommendation framework can well support users' diverse healthcare preferences and that the diversity constraint can further guide the recommendation system to explore along each of the weight-management dimensions and, thus, improves learning efficiency.




Figure 5.7 Diversity Distribution

5.4.6 Experiment 6: User Improvement

In this experiment, we aim to examine whether our recommendation framework can benefit more users. We define user improvement as the percentage of users who receive more preferred items from a focal recommendation algorithm than from a baseline algorithm. We use probabilistic matrix factorization (PMF) as our baseline algorithm. PMF models hidden user representations based on the users' challenge-selection histories. It does not, however, incorporate contextual

information of users' selection behaviors, and it is batch-learning-based. Thus, by comparing with PMF, we are able to assess the value of the recommendation context along with the online-learning scheme. We present our results for user improvement in Figure 5.8. It is shown that our proposed recommendation approach has the highest user improvement rate ($\sim 76.19\%$), suggesting that approximately 76% of users receive more preferred challenges from our recommendation framework than from PMF. Comparatively, the other recommendation approaches have a lower user improvement rate (all below 70%). These results indicate that our recommendation framework can serve to improve a larger user population on the platform.




Figure 5.8 User Improvement (Using PMF as Benchmark)

5.4.7 Experiment 7: Using Our Recommendation Framework to Improve Weight Loss

The former experiments evaluate the effectiveness of our recommendation framework in improving users' challenge-selection rates. In the healthcare context, it is also important that service providers take further steps to evaluate corresponding health-related outcomes. Although necessary, increasing users' engagement in interventions may not be enough, as it does not directly guarantee an improvement in health. Therefore, in this experiment, we further examine our recommendation performance in improving users' weight-loss outcomes. In particular, we explore

a different recommendation target: users’ in-period weight-loss rate.⁶ That is, we use users’ in-period weight-loss rate as feedback to guide the learning of our recommendation framework. Different from prior experiments, generating weight-loss feedback requires two steps of simulation. First, we use the logistic predictor described in Section 5.4.2 to simulate whether users will choose a particular item. Second, we simulate users’ in-period weight-loss status (e.g., weight gain or non-gain) based on their choices. We use users’ weigh-in data to train a logistic predictor for this simulation, which takes user embeddings and the average challenge embeddings of users’ choice as the input variables and users’ weight-loss status as the prediction target.

As shown in Table 5.7, our proposed recommendation framework has the best performance under the new recommendation target, achieving an average in-period weight-loss rate of 75.64%. In contrast, the benchmark models achieve a significantly worse average in-period weight-loss rate. These results indicate that our recommendation design is effective not only in helping users to find preferable challenges to engage in but also in further helping them to achieve better weight-loss performance.

⁶ In this study, we define in-period weight-loss rate as the percentage of periods in which users’ weight has decreased or remained unchanged.

Table 5.7 Weight Loss as Recommendation Goal

Method	Average In-Period Weight-loss Rate
proposed MAB	0.7564
<i>State-of-the-Art Bandit Models</i>	
UCB	0.7039 ***
epsilon greedy	0.6739 ***
<i>State-of-the-Art Batch-Learning Models</i>	
CACF	0.6318 ***
SCF	0.6294 ***
PMF	0.6459 ***
CAMF	0.6165 ***
CB	0.6472 ***
hybrid_pure	0.6049 ***
hybrid_cacf	0.6507 ***
<i>Pure exploitation and Pure Exploration</i>	
pure exploitation	0.7035 ***
pure exploration	0.4006 ***

Note: Asterisk in superscript denotes that a benchmark model performs significantly worse than the proposed model. Significance levels are: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

5.5 CONCLUSION AND DISCUSSION

In this study, we take a design-science perspective to develop a novel recommendation framework for providing personalized healthcare recommendations to users on online healthcare platforms. The design of our recommendation framework is motivated by several unique patterns in individuals' health behaviors. First, due to the evolving process of health management, users' health behaviors may continuously change with time. Second, users may need multiplex healthcare information to manage different health aspects. These characteristics indicate that users' healthcare preferences can be dynamic and diverse. To this end, we propose a deep-learning and diversity-enhanced MAB framework. MAB is able to adaptively learn users' changing behavior patterns while promoting diversity along the exploration process. To better adapt an MAB to the healthcare recommendation context, we further synthesize two model components into our framework based on prominent health-behavior theories. The first component is a deep-learning-based feature

construction procedure, aimed at capturing important healthcare recommendation contexts, such as healthcare intervention's inherent attributes and individuals' evolving health condition and health-behavior paths. The second component is a diversity constraint, which we use to ensure that recommendations are provided in each of the major health-management dimensions, so that individuals can receive well-rounded support for their health management. We conduct a series of experiments to test each of our model components as well as to compare our model against state-of-the-art recommendation systems, using data collected from a representative online weight-loss platform. The results of the experiments provide strong evidence for the effectiveness of our proposed recommendation framework.

Our study contributes to the emerging literature on the application of business intelligence (Abbasi et al. 2016; Chen et al. 2012). We demonstrate that prescriptive analytics can be integrated with IT artifacts to generate applicable insights. In particular, the innovative healthcare recommendation framework that we developed provides an important contribution to the literature on recommendation systems and online healthcare systems. To the best of our knowledge, we are among the first to combine MAB models with deep-learning-based embeddings to improve the characterization of recommendation contexts. In addition, the inclusion of diversity constraints demonstrates a way of promoting recommendation diversity according to pre-designed dimensions. This innovation can be of significance in professional industries, in which domain expertise needs to be incorporated to guide recommendation diversification.

From a practical perspective, our recommendation framework can be used to address real-world challenges in healthcare recommendation problems. The effectiveness of our framework, as demonstrated by our results, implies great potential for using our recommendation design to provide users with tailored engagement suggestions. Although our recommendation design is

proposed for assisting individuals' engagement in online health management, the framework can be extended to broader problem settings. For example, the combination of an MAB and deep-learning-based feature engineering can be used to solve other healthcare problems, such as drug discovery, disease diagnosis, clinical trials, and therapy development. Decision making in such healthcare problems usually involves complex contextual knowledge (e.g., drug structures, patient health histories, symptom development paths), and decision-makers usually do not have full knowledge of the environment (e.g., whether a drug is effective). The online-learning framework of an MAB can help decision-makers to better cope with uncertainties in the healthcare environment, and deep-learning models can be combined to improve the characterization of decision-making contexts. In addition, we demonstrate a way to formulate recommendation constraints, which can be used to incorporate domain expertise to guide the recommendation procedure.

Our recommendation framework can also be extended to fields beyond healthcare. Real-world decision-making problems, such as financial investment, product pricing, and marketing, usually contain different levels of uncertainty. The uncertainty may be because decision-makers do not gather enough data to guide their decision making, or the decision-making environment is frequently changing (e.g., market instability, technology change, policy environment fluctuation). Thus, it is of practical importance to develop an adaptive decision-making framework that can respond well to the uncertainty in the environment. Decision-makers may consider combining an MAB with deep-learning embeddings to learn the context-dependency of their decision results while adaptively adjusting their strategies to minimize opportunity cost. In addition, in real-world recommendation problems, it is usually desirable to recommend diversified content to maximize the coverage of the information that users find interesting to improve their engagement experience.

Our formulation of the diversity constraint can be used to strengthen recommendation diversification along with theory-guided dimensions in multiple application areas.

Chapter 6. CONCLUDING REMARKS

The recent years have witnessed the transition of healthcare from one that is disease-oriented to one that is health-oriented. Compared to disease-oriented care, which is often reactive and emphasizes healthcare professionals' expertise in treatment procedure, health-oriented care highlights individuals' efforts in disease prevention, diagnosis, and wellness promotion (Chawla and Davis 2013; Stange 2009). Patients who used to rely entirely on healthcare professionals now become active participants in their care. Individuals' increasing use of information systems has accelerated such transition process. According to Fox and Duggan (2013), approximately 35% of U.S. adults have utilized online resources to discern a medical condition. The usage of online healthcare services continues to increase during the COVID-19 pandemic period, in which people use online portals to access remote medical consultations and personal health-management support.

Online platforms enable convenient access of healthcare information and services and they serve as effective channels for individuals to store and manage personal health data. As a result, the integration of information technology and healthcare provides potentials not only for individuals to improve their health-management efficiency but also for service providers to improve care delivery. This dissertation seeks to understand the role of information technology on individuals' health behaviors and explore associated opportunities and challenges. The investigation is conducted on three different aspects. In the first essay, I study the underlying engagement state of users to disentangle individuals' dynamic participation patterns in online healthcare platforms. In the second essay, I study users' learning behaviors in participating in online behavior intervention programs as a means to understand individuals' perception of online intervention features. In the third essay, I develop a personalized recommendation system to help

users select online interventions, aiming to provide prescriptive decision-making support for users' health management. Together, these studies provide implications for users' online health-behavior patterns, online healthcare platform design, and user sustainability.

BIBLIOGRAPHY

Abbasi A, Sarker S, and Chiang R H. (2016) Big Data Research in Information Systems: Toward an Inclusive Research Agenda. *Journal of the Association for Information Systems* 17(2): 3.

Adomavicius G, and Tuzhilin A. (2005) Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. *IEEE Transactions on Knowledge and Data Engineering* 17(6): 734-749.

Almars A, Li X, and Zhao X. (2019) Modelling User Attitudes Using Hierarchical Sentiment-Topic Model. *Data & Knowledge Engineering* 119: 139-149.

American Medical Association. (2013) Recognition of Obesity as a Disease, Accessed Dec 06, 2018. <https://www.npr.org/documents/2013/jun/ama-resolution-obesity.pdf>.

Arora S, Peters A L, Agy C, and Menchine M. (2012) A Mobile Health Intervention for Inner City Patients with Poorly Controlled Diabetes: Proof-of-Concept of the Text-Med Program. *Diabetes technology & therapeutics* 14(6): 492-496.

Auer P, Cesa-Bianchi N, and Fischer P. (2002) Finite-Time Analysis of the Multiarmed Bandit Problem. *Machine Learning* 47(2-3): 235-256.

Baker R C, and Kirschenbaum D S. (1993) Self-Monitoring May Be Necessary for Successful Weight Control. *Behavior Therapy* 24(3): 377-394.

Ballantine P W, and Stephenson R J. (2011) Help Me, I'm Fat! Social Support in Online Weight Loss Networks. *Journal of Consumer Behaviour* 10(6): 332-337.

Baltrunas L, Ludwig B, and Ricci F. (2011) Matrix Factorization Techniques for Context Aware Recommendation. *Proceedings of the Fifth ACM Conference on Recommender Systems*: 301-304.

- Bandura A. (1991) Social Cognitive Theory of Self-Regulation. *Organizational Behavior and Human Decision Processes* 50(2): 248-287.
- Bandura A. (1998) Health Promotion from the Perspective of Social Cognitive Theory. *Psychology and Health* 13(4): 623-649.
- Bass R F (2011). *Stochastic Processes* (Cambridge University Press).
- Baum A, Revenson T A, and Singer J (2012) Handbook of Health Psychology. (Psychology Press).
- Bieliková M, Kompan M, and Zeleník D. (2012) Effective Hierarchical Vector-Based News Representation for Personalized Recommendation. *Computer Science and Information Systems* 9(1): 303-322.
- Borders T F, Xu K T, Heavner J, and Kruse G. (2005) Patient Involvement in Medical Decision-Making and Pain among Elders: Physician or Patient-Driven? *BMC health services research* 5(1): 4.
- Boutelle K N, Kirschenbaum D S, Baker R C, and Mitchell M E. (1999) How Can Obese Weight Controllers Minimize Weight Gain During the High Risk Holiday Season? By Self-Monitoring Very Consistently. *Health Psychology* 18(4): 364.
- Burke L E, Wang J, and Sevick M A. (2011) Self-Monitoring in Weight Loss: A Systematic Review of the Literature. *Journal of the American Dietetic Association* 111(1): 92-102.
- Burke R. (2002) Hybrid Recommender Systems: Survey and Experiments. *User Modeling and User-Adapted Interaction* 12(4): 331-370.
- Butryn M L, Phelan S, Hill J O, and Wing R R. (2007) Consistent Self-Monitoring of Weight: A Key Component of Successful Weight Loss Maintenance. *Obesity* 15(12): 3091-3096.
- Buunk A P, and Gibbons F X. (2007) Social Comparison: The End of a Theory and the Emergence of a Field. *Organizational Behavior and Human Decision Processes* 102(1): 3-21.

CDC. (2020) People with Certain Medical Conditions, Accessed Aug 10, 2020. <https://www.cdc.gov/coronavirus/2019-ncov/need-extra-precautions/people-with-medical-conditions.html#obesity>.

Chapelle O, and Li L (2011) An Empirical Evaluation of Thompson Sampling. *Advances in neural information processing systems*, 2249-2257.

Chawla N V, and Davis D A. (2013) Bringing Big Data to Personalized Healthcare: A Patient-Centered Framework. *Journal of general internal medicine* 28(3): 660-665.

Chen A (2005) Context-Aware Collaborative Filtering System: Predicting the User's Preference in the Ubiquitous Computing Environment. *International Symposium on Location-and Context-Awareness* (Springer), 244-253.

Chen H, Chiang R H, and Storey V C. (2012) Business Intelligence and Analytics: From Big Data to Big Impact. *MIS Quarterly*: 1165-1188.

Cheng H-T, Koc L, Harmsen J, Shaked T, Chandra T, Aradhye H, Anderson G, Corrado G, Chai W, and Ispir M. (2016) Wide & Deep Learning for Recommender Systems. *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*: 7-10.

Clark C M. (2005) Relations between Social Support and Physical Health. *Rochester Institute of Technology*.

Crant J M. (2000) Proactive Behavior in Organizations. *Journal of Management* 26(3): 435-462.

Das S R, and Chen M Y. (2007) Yahoo! For Amazon: Sentiment Extraction from Small Talk on the Web. *Management science* 53(9): 1375-1388.

De Jaegher K, and Jegers M. (2001) The Physician–Patient Relationship as a Game of Strategic Information Transmission. *Health Economics* 10(7): 651-668.

DiMatteo M R. (2004) Social Support and Patient Adherence to Medical Treatment: A Meta-Analysis. *Health Psychology* 23(2): 207.

- Doran G T. (1981) There's a Smart Way to Write Management's Goals and Objectives. *Management Review* 70(11): 35-36.
- Dudík M, Langford J, and Li L. (2011) Doubly Robust Policy Evaluation and Learning. *arXiv preprint arXiv:1103.4601*.
- Durrett R (1999). *Essentials of Stochastic Processes* (Springer).
- Endres D M, and Schindelin J E. (2003) A New Metric for Probability Distributions. *IEEE Transactions on Information Theory* 49(7): 1858-1860.
- Festinger L. (1954) A Theory of Social Comparison Processes. *Human Relations* 7(2): 117-140.
- Fleder D, and Hosanagar K. (2009) Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity. *Management Science* 55(5): 697-712.
- Fox S, and Duggan M (2013) *Tracking for Health* (Pew Research Center's Internet & American Life Project,
- Franz M J. (2001) The Answer to Weight Loss Is Easy—Doing It Is Hard! *Clinical Diabetes* 19(3): 105-109.
- Fuglede B, and Topsøe F (2004) Jensen-Shannon Divergence and Hilbert Space Embedding. *Proceedings of International Symposium on Information Theory*. (IEEE), 31.
- Fulford K. (2011) Bringing Together Values-Based and Evidence-Based Medicine: Uk Department of Health Initiatives in the 'Personalization' of Care. *Journal of Evaluation in Clinical Practice* 17(2): 341-343.
- Gallant M P. (2003) The Influence of Social Support on Chronic Illness Self-Management: A Review and Directions for Research. *Health Education & Behavior* 30(2): 170-195.
- Gherman A, Schnur J, Montgomery G, Sassu R, Veresiu I, and David D. (2011) How Are Adherent People More Likely to Think? A Meta-Analysis of Health Beliefs and Diabetes Self-Care. *The Diabetes Educator* 37(3): 392-408.

Gittins J C. (1979) Bandit Processes and Dynamic Allocation Indices. *Journal of the Royal Statistical Society: Series B (Methodological)* 41(2): 148-164.

Good B M, Ainscough B J, McMichael J F, Su A I, and Griffith O L. (2014) Organizing Knowledge to Enable Personalization of Medicine in Cancer. *Genome biology* 15(8): 1-9.

Gooden R J, and Winefield H R. (2007) Breast and Prostate Cancer Online Discussion Boards: A Thematic Analysis of Gender Differences and Similarities. *Journal of Health Psychology* 12(1): 103-114.

Greenberg M S (1980) A Theory of Indebtedness. *Social Exchange* (Springer),3-26.

Greenewald K, Tewari A, Murphy S, and Klasnja P. (2017) Action Centered Contextual Bandits. *Advances in Neural Information Processing Systems*: 5977-5985.

Hall K D, and Kahan S. (2018) Maintenance of Lost Weight and Long-Term Management of Obesity. *Medical Clinics of North America* 102(1): 183-197.

Hamilton J D. (1989) A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica: Journal of the Econometric Society*: 357-384.

Hamine S, Gerth-Guyette E, Faulx D, Green B B, and Ginsburg A S. (2015) Impact of Mhealth Chronic Disease Management on Treatment Adherence and Patient Outcomes: A Systematic Review. *Journal of Medical Internet Research* 17(2).

Harvey-Berino J, Pintauro S, Buzzell P, DiGiulio M, Gold B C, Moldovan C, and Ramirez E. (2002) Does Using the Internet Facilitate the Maintenance of Weight Loss? *International Journal of Obesity* 26(9): 1254.

Harvey-Berino J, Pintauro S, Buzzell P, and Gold E C. (2004) Effect of Internet Support on the Long-Term Maintenance of Weight Loss. *Obesity Research* 12(2): 320-329.

Helsel D L, Jakicic J M, and Otto A D. (2007) Comparison of Techniques for Self-Monitoring Eating and Exercise Behaviors on Weight Loss in a Correspondence-Based Intervention. *Journal of the American Dietetic Association* 107(10): 1807-1810.

Hertz U, Bahrami B, and Keramati M. (2018) Stochastic Satisficing Account of Confidence in Uncertain Value-Based Decisions. *PLOS One* 13(4): e0195399.

Hibbard J H, Mahoney E R, Stock R, and Tusler M. (2007) Do Increases in Patient Activation Result in Improved Self-Management Behaviors? *Health services research* 42(4): 1443-1463.

Holman H, and Lorig K. (2000) Patients as Partners in Managing Chronic Disease: Partnership Is a Prerequisite for Effective and Efficient Health Care. *British Medical Journal* 320(7234): 526.

Hopp W J, Li J, and Wang G. (2018) Big Data and the Precision Medicine Revolution. *Production and Operations Management* 27(9): 1647-1664.

Hwang K O, Ottenbacher A J, Green A P, Cannon-Diehl M R, Richardson O, Bernstam E V, and Thomas E J. (2010) Social Support in an Internet Weight Loss Community. *International Journal of Medical Informatics* 79(1): 5-13.

Isinkaye F, Folajimi Y, and Ojokoh B. (2015) Recommendation Systems: Principles, Methods and Evaluation. *Egyptian Informatics Journal* 16(3): 261-273.

Jasra A, Holmes C C, and Stephens D A. (2005) Markov Chain Monte Carlo Methods and the Label Switching Problem in Bayesian Mixture Modeling. *Statistical Science*: 50-67.

Johnson F, and Wardle J. (2011) The Association between Weight Loss and Engagement with a Web-Based Food and Exercise Diary in a Commercial Weight Loss Programme: A Retrospective Analysis. *International Journal of Behavioral Nutrition and Physical Activity* 8(1): 83.

Johnson P E, Veazie P J, Kochevar L, O'connor P J, Potthoff S J, Verma D, and Dutta P. (2002) Understanding Variation in Chronic Disease Outcomes. *Health Care Management Science* 5(3): 175-189.

Karlin S (2014). *A First Course in Stochastic Processes* (Academic Press).

Kim M J, and Lim A E. (2015) Robust Multiarmed Bandit Problems. *Management Science* 62(1): 264-285.

Kim Y, and Krishnan R. (2019) The Dynamics of Online Consumers' Response to Price Promotion. *Information Systems Research*.

King G, Willoughby C, Specht J A, and Brown E. (2006) Social Support Processes and the Adaptation of Individuals with Chronic Disabilities. *Qualitative Health Research* 16(7): 902-925.

Konstan J A, and Riedl J. (2012) Recommender Systems: From Algorithms to User Experience. *User Modeling and User-Adapted Interaction* 22(1-2): 101-123.

Krogh A, Brown M, Mian I S, Sjolander K, and Haussler D. (1994) Hidden Markov Models in Computational Biology. Applications to Protein Modeling. *Journal of molecular biology* 235(5): 1501-1531.

Krukowski R A, Harvey-Berino J, Ashikaga T, Thomas C S, and Micco N. (2008) Internet-Based Weight Control: The Relationship between Web Features and Weight Loss. *Telemedicine and e-Health* 14(8): 775-782.

Lai T L, and Robbins H. (1985) Asymptotically Efficient Adaptive Allocation Rules. *Advances in applied mathematics* 6(1): 4-22.

Leahey T M, LaRose J G, Fava J L, and Wing R R. (2011) Social Influences Are Associated with Bmi and Weight Loss Intentions in Young Adults. *Obesity* 19(6): 1157-1162.

Li L, Chu W, Langford J, and Schapire R E (2010) A Contextual-Bandit Approach to Personalized News Article Recommendation. *Proceedings of the 19th international conference on World wide web* (ACM), 661-670.

Liu Y, Zhang Y, Wu Q, Miao C, Cui L, Zhao B, Zhao Y, and Guan L. (2019) Diversity-Promoting Deep Reinforcement Learning for Interactive Recommendation. *arXiv preprint arXiv:1903.07826*.

Locke E A, and Latham G P (1990) *A Theory of Goal Setting & Task Performance* (Prentice-Hall, Inc.),

Maaten L v d, and Hinton G. (2008) Visualizing Data Using T-Sne. *Journal of Machine Learning Research* 9(Nov): 2579-2605.

McKay H G, Glasgow R E, Feil E G, Boles S M, and Barrera Jr M. (2002) Internet-Based Diabetes Self-Management and Support: Initial Outcomes from the Diabetes Network Project. *Rehabilitation Psychology* 47(1): 31.

Miller T A, and DiMatteo M R. (2013) Importance of Family/Social Support and Impact on Adherence to Diabetic Therapy. *Diabetes, Metabolic Syndrome and Obesity: Targets and Therapy* 6: 421.

Misra K, Schwartz E M, and Abernethy J. (2019) Dynamic Online Pricing with Incomplete Information Using Multiarmed Bandit Experiments. *Marketing Science*.

Mnih A, and Salakhutdinov R R. (2008) Probabilistic Matrix Factorization. *Advances in Neural Information Processing Systems*: 1257-1264.

Moisio R, and Beruchashvili M. (2009) Questing for Well-Being at Weight Watchers: The Role of the Spiritual-Therapeutic Model in a Support Group. *Journal of Consumer Research* 36(5): 857-875.

Montgomery A L, Li S, Srinivasan K, and Liechty J C. (2004) Modeling Online Browsing and Path Analysis Using Clickstream Data. *Marketing Science* 23(4): 579-595.

Netzer O, Lattin J M, and Srinivasan V. (2008) A Hidden Markov Model of Customer Relationship Dynamics. *Marketing Science* 27(2): 185-204.

Newsom J T. (1999) Another Side to Caregiving: Negative Reactions to Being Helped. *Current Directions in Psychological Science* 8(6): 183-187.

Nguyen K, Daumé III H, and Boyd-Graber J. (2017) Reinforcement Learning for Bandit Neural Machine Translation with Simulated Human Feedback. *arXiv preprint arXiv:1707.07402*.

Norcross J C, and Vangarelli D J. (1988) The Resolution Solution: Longitudinal Examination of New Year's Change Attempts. *Journal of Substance Abuse* 1(2): 127-134.

Nutting P A, Crabtree B F, Miller W L, Stange K C, Stewart E, and Jaén C. (2011) Transforming Physician Practices to Patient-Centered Medical Homes: Lessons from the National Demonstration Project. *Health Affairs* 30(3): 439-445.

Ogbeiwi O. (2017) Why Written Objectives Need to Be Really Smart. *British Journal of Healthcare Management* 23(7): 324-336.

Ogbeiwi O. (2018) General Concepts of Goals and Goal-Setting in Healthcare: A Narrative Review. *Journal of Management & Organization*: 1-18.

Oulasvirta A, Hukkinen J P, and Schwartz B. (2009) When More Is Less: The Paradox of Choice in Search Engine Use. *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval*: 516-523.

Pariser E (2011) *The Filter Bubble: How the New Personalized Web Is Changing What We Read and How We Think* (Penguin, UK)

Pathak B, Garfinkel R, Gopal R D, Venkatesan R, and Yin F. (2010) Empirical Analysis of the Impact of Recommender Systems on Sales. *Journal of Management Information Systems* 27(2): 159-188.

Perri M G, Sears S F, and Clark J E. (1993) Strategies for Improving Maintenance of Weight Loss: Toward a Continuous Care Model of Obesity Management. *Diabetes Care* 16(1): 200-209.

Pew Research Center. (2020) From Virtual Parties to Ordering Food, How Americans Are Using the Internet During Covid-19, Accessed August 31, 2020. <https://www.pewresearch.org/fact-tank/2020/04/30/from-virtual-parties-to-ordering-food-how-americans-are-using-the-internet-during-covid-19/>.

Pon R K, Cardenas A F, Buttlar D, and Critchlow T (2007) Tracking Multiple Topics for Finding Interesting Articles. *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 560-569.

Pu P, Chen L, and Hu R (2011) A User-Centric Evaluation Framework for Recommender Systems. *Proceedings of the Fifth ACM Conference on Recommender Systems (ACM)*, 157-164.

Qin L, Chen S, and Zhu X (2014) Contextual Combinatorial Bandit and Its Application on Diversified Online Recommendation. *Proceedings of the 2014 SIAM International Conference on Data Mining* (SIAM), 461-469.

Qin L, and Zhu X (2013) Promoting Diversity in Recommendation by Entropy Regularizer. *Twenty-Third International Joint Conference on Artificial Intelligence*.

Rabiner L R. (1989) A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. *Proceedings of the IEEE* 77(2): 257-286.

Ricci F, Rokach L, and Shapira B (2015) Recommender Systems: Introduction and Challenges. *Recommender Systems Handbook* (Springer), 1-34.

Sacristán J A. (2013) Patient-Centered Medicine and Patient-Oriented Research: Improving Health Outcomes for Individual Patients. *BMC medical informatics and decision making* 13(1): 1-8.

Sahoo N, Singh P V, and Mukhopadhyay T. (2012) A Hidden Markov Model for Collaborative Filtering. *MIS Quarterly*: 1329-1356.

Sani A, Lazaric A, and Munos R. (2012) Risk-Aversion in Multi-Armed Bandits. *Advances in Neural Information Processing Systems*: 3275-3283.

Saperstein S, Atkinson N, and Gold R. (2007) The Impact of Internet Use for Weight Loss. *Obesity reviews* 8(5): 459-465.

Schuller B, Rigoll G, and Lang M (2003) Hidden Markov Model-Based Speech Emotion Recognition. *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP'03)*. (IEEE), II-1.

Schwartz E M, Bradlow E T, and Fader P S. (2017) Customer Acquisition Via Display Advertising Using Multi-Armed Bandit Experiments. *Marketing Science* 36(4): 500-522.

Sedhain S, Sanner S, Braziunas D, Xie L, and Christensen J. (2014) Social Collaborative Filtering for Cold-Start Recommendations. *Proceedings of the 8th ACM Conference on Recommender Systems*: 345-348.

Shmueli G, and Koppius O R. (2011) Predictive Analytics in Information Systems Research. *MIS Quarterly*: 553-572.

Shumaker S A, and Brownell A. (1984) Toward a Theory of Social Support: Closing Conceptual Gaps. *Journal of Social Issues* 40(4): 11-36.

Snyderman R, and Dinan M A. (2010) Improving Health by Taking It Personally. *Jama* 303(4): 363-364.

Speekenbrink M, and Konstantinidis E. (2015) Uncertainty and Exploration in a Restless Bandit Problem. *Topics in cognitive science* 7(2): 351-367.

Stange K C. (2009) The Generalist Approach. *The Annals of Family Medicine* 7(3): 198-203.

Suls J (2000) Opinion Comparison. *Handbook of Social Comparison* (Springer),105-122.

Tang L, Jiang Y, Li L, and Li T (2014) Ensemble Contextual Bandits for Personalized Recommendation. *Proceedings of the 8th ACM Conference on Recommender Systems* (ACM), 73-80.

Tate D F, Jackvony E H, and Wing R R. (2003) Effects of Internet Behavioral Counseling on Weight Loss in Adults at Risk for Type 2 Diabetes: A Randomized Trial. *Jama* 289(14): 1833-1836.

Tate D F, Wing R R, and Winett R A. (2001) Using Internet Technology to Deliver a Behavioral Weight Loss Program. *Jama* 285(9): 1172-1177.

Thomas J, Raynor H, Bond D, Luke A, Cardoso C, Wojtanowski A, Vander Veur S, Tate D, Wing R, and Foster G. (2017) Weight Loss and Frequency of Body-Weight Self-Monitoring in an Online Commercial Weight Management Program with and without a Cellular-Connected ‘Smart’ Scale: A Randomized Pilot Study. *Obesity Science & Practice* 3(4): 365-372.

Umberson D, and Karas Montez J. (2010) Social Relationships and Health: A Flashpoint for Health Policy. *Journal of Health and Social Behavior* 51(1): S54-S66.

van der Eijk M, Faber M J, Aarts J W, Kremer J A, Munneke M, and Bloem B R. (2013) Using Online Health Communities to Deliver Patient-Centered Care to People with Chronic Conditions. *Journal of Medical Internet Research* 15(6).

Vozalis E, and Margaritis K G. (2003) Analysis of Recommender Systems Algorithms. *The 6th Hellenic European Conference on Computer Mathematics & Its Applications*: 732-745.

W.H.O. (2016) Noncommunicable Diseases, Accessed Dec 06, 2018. http://www.who.int/topics/noncommunicable_diseases/en/.

Wang Y-C, Kraut R, and Levine J M (2012) To Stay or Leave?: The Relationship of Emotional and Informational Support to Commitment in Online Health Support Groups. *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work* (ACM), 833-842.

WHO. (2018) Obesity and Overweight, Accessed May 01, 2019. <https://www.who.int/news-room/fact-sheets/detail/obesity-and-overweight>.

World Health Organization. (2018) Obesity and Overweight, Accessed May 01, 2019. <https://www.who.int/news-room/fact-sheets/detail/obesity-and-overweight>.

Wortley D (2015) Gamification and Lifestyle Technologies for Personal Health Management. *European conference on games based learning* (Academic Conferences International Limited), 762.

Yan L. (2018) Good Intentions, Bad Outcomes: The Effects of Mismatches between Social Support and Health Outcomes in an Online Weight Loss Community. *Production and Operations Management* 27(1): 9-27.

Yan L, and Tan Y. (2014) Feeling Blue? Go Online: An Empirical Study of Social Support among Patients. *Information Systems Research* 25(4): 690-709.

Yan L, and Tan Y. (2017) The Consensus Effect in Online Health-Care Communities. *Journal of Management Information Systems* 34(1): 11-39.

Yu S-Z. (2010) Hidden Semi-Markov Models. *Artificial Intelligence* 174(2): 215-243.

Zeng C, Wang Q, Mokhtari S, and Li T. (2016) Online Context-Aware Recommendation with Time Varying Multi-Armed Bandit. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*: 2025-2034.

Zhang S, Yao L, Sun A, and Tay Y. (2019) Deep Learning Based Recommender System: A Survey and New Perspectives. *ACM Computing Surveys (CSUR)* 52(1): 1-38.

APPENDIX A

A.1 CONTINUOUS OR DISCRETE

As noted in the manuscript, our data is obtained on a daily basis. A natural question to ask is how much a continuous-time HMM (CTHMM) differs from a discrete-time HMM (DTHMM) with a daily time frame. This concern is valid, as theoretically and mathematically, a CTHMM can be seen as a DTHMM when the time frame goes to infinitesimal. However, when we add in panel covariates to control the dynamics of waiting time, state transition, and the state-dependent outcome, a CTHMM will no longer be equivalent to a DTHMM with a small time unit.

In addition, there are several reasons for choosing a CTHMM over a DTHMM from the practical perspective. First, estimating a DTHMM with a small time unit can create computational issues. In reality, many individuals may not participate in online communities on a frequent basis. For instance, a community member may stop usage for a couple of months before restarting online participation. If we develop a DTHMM with a daily time frame, we will need to compute the state-transition probabilities on a daily basis, which will generate tremendous computation. Different from the DTHMM approach, a CTHMM captures individuals' inactive period by the waiting time distribution. This unique modeling structure can greatly reduce computation burden since we need to compute the state transitions only when online activities occur.

Second, a DTHMM may generate “sparsity issue” in the state-dependent outcome distribution when the time unit is small. As individuals may not participate in online communities frequently, it is very likely that researchers will observe a great amount of zeros in the outcome sequences if a small time unit (e.g., day) is adopted. Due to the zero inflation, the outcome sequence data can hardly be fitted by any distribution, as the probability for observing zero-valued outcomes is almost

1. In contrast, a CTHMM is able to capture the timing process of individuals' online participation and, thus, can better characterize individuals' behavior patterns.

A.2 USER DESCRIPTIVE STATISTICS

We provide distributions about users' gender, age, and membership duration in Figures A.1-A.3. As can be seen in Figure A.1, the gender ratio of the users is approximately 1:1. Figure A.2 shows that most of the users are middle-aged adults, with an average age of around 40. These observations are consistent with the offline overweight population in the United States, as suggested by the Centers for Disease Control and Prevention (2020). Figure A.3 shows that the average membership duration of the users is about 400 days, meaning that a typical user in our data stays in the community for approximately a year.




Figure A.1 Gender Distribution




Figure A.2 Age Distribution




Figure A.3 Membership Duration Distribution

A.3 ESTIMATION PROCEDURE

A.3.1 Latent Class Model

As noted in the manuscript, we use a latent class model (LCM) to determine the initial distribution for hidden states. A LCM describes individuals' behaviors by observed attributes and latent heterogeneity. Different from hidden Markov models (HMMs), LCMs are static. That is, state transitions are not captured by LCMs. Instead, LCMs use individual-specific characteristics to model the probabilities of class assignment. In the LCM that we build, we use individuals' average

participation frequency for self-monitoring activities and social activities to capture individual-level heterogeneity that affects state/class assignment. The variables that affect state-dependent outcomes are kept as the same as the ones in our main model.

Formally, the probability of individual i being assigned to class s is:

$$p_{is} = e^{c_s \cdot freq_i} / \left(1 + \sum_{j=1}^2 e^{c_j \cdot freq_i} \right), \quad s=1,2, \quad (\text{A.1})$$

where $freq_i$ represents a vector that contains individuals' average participation frequencies and a constant term, i.e., $freq_i = (freq_monitor_i, freq_social_i, 1)^T$. When $s=3$, $p_{is} = 1 / (1 + \sum_{j=1}^2 e^{c_j \cdot freq_i})$.

Given that individual i belongs to class s , the probability of observing behavior outcome $Y_{it} = (Y_{1it}, Y_{2it}, \dots, Y_{nit})$ is:

$$\Pi(Y_{jit} | s_{it} = s) = e^{-\lambda_{jsit}} \cdot \lambda_{jsit}^{Y_{jit}} / Y_{jit}!, \quad j=1,2,\dots,n, \quad (\text{A.2})$$

where $\lambda_{jsit} = e^{\beta_j W_{it} + \varepsilon_{ji}}$. W_{it} is a set of variables that affect individual's outcome at time t ; β_j is the coefficient vector; s represents state indices. ε_{ji} denotes individual-level random effect in dimension j .

The likelihood for patient i can be computed as $L_i = \sum_s p_{is} (\prod_t Z_{it|s})$, where $Z_{it|s} = \prod_{j=1}^n \Pi(Y_{jit} | s_{it} = s)$. The overall log-likelihood is $\ln L = \sum_i \ln L_i$. Through likelihood maximization, we are able to obtain the model estimates and compute the probability distribution of class assignment for each individual. The initial hidden state distribution can be calculated as the average class-assignment probabilities in the sample. The probabilities we obtain are (0.9413, 0.0266, 0.0321) for the three-state model.

A.3.2 Discount Factor in Exponential Learning

Exponential smoothing is a method of discounting the effects of prior events when the time is continuous. For example, for activity that occurs Δt time unit ago, we discount it by $e^{-r \cdot \Delta t}$, where r is the discounting factor to be optimized. In this study, we deploy a one-dimensional search to determine the optimal discount factor. Specifically, we evaluate the likelihood function at different points ranging from 0 to 1. The log-likelihood values at each point are reported in Table A.1. As can be seen, when $r = 0.2$, the log-likelihood reaches the highest value. Thus, we determine the optimal discount factor to be 0.2.

Table A.1 Selection on Discount Factor

r	Log-likelihood
0.0	-15020.44
0.1	-15018.84
0.2	-15016.93
0.3	-15023.33
0.4	-15026.08
0.5	-15030.71
0.6	-15035.67
0.7	-15041.16
0.8	-15037.22
0.9	-15038.26
1.0	-15040.43

A.3.3 Selecting the Number of Supports for the Random-Effect Variables

A non-parametric method is deployed to estimate the joint distribution of the random effects. That is, for each random-effect variable, we estimate a fixed number of supports and the corresponding probabilities. To decide the optimal number of supports, we run a set of models with different combinations of support numbers and select the one that generates the highest likelihood value. The selection results are provided in Table A.2.

Table A.2 Selection on the Number of Supports

Support Numbers	Log-likelihood
(2, 2)	-15029.14
(2, 3)	-15022.51
(3, 2)	-15025.14
(3, 3)	-15024.77

A.4 ROBUSTNESS CHECKS

The robustness checks on our results are performed from three perspectives. First, we examine a potential alternative distribution for the state-dependent outcome. Second, we investigate alternative measures for the frequency variables. Third, we examine whether our findings are robust to sample selection.

A.4.1 Alternative State-Dependent Outcome Distribution

In the main text, we model the state-dependent outcome by a Poisson distribution. As a robustness check, we consider a zero-inflated Poisson to see if there is zero inflation in our data. The zero-inflated Poisson distribution assumes that there are two zero-generating schemes: the first one is governed by a binary distribution that generates structural zeros, and the second one is governed by a Poisson process that generates counts that may contain zeros. We model the two zero-generating processes as follows:

$$\Pi(Y_{jit} | s_{it} = s) = \omega_j I_{[Y_{jit}=0]} + (1 - \omega_j) \cdot e^{-\lambda_{jsit}} \cdot \lambda_{jsit}^{Y_{jit}} / Y_{jit}!, \quad j=1,2,\dots,n, \quad (\text{A.3})$$

where $I_{[Y_{jit}=0]}$ is an indicator function that equals one when $Y_{jit}=0$, and ω_j is a “inflation” parameter that controls excess zeros generated from a binary distribution. Our estimates for ω_j ($j=1,2$) are very close to zero, indicating that there is no zero inflation in our data and that a Poisson distribution is thus sufficient to capture the data variation patterns. The full estimation

results are presented in Table A.3. The results show that our main findings stay robust after this change in model specification.

Table A.3 Robustness Check 1: Zero Inflated Poisson

Parameters	State 1 (L)		State 2 (M)		State 3 (H)	
Variables Affecting State-Dependent Outcome						
	Weigh-in	Journal	Weigh-in	Journal	Weigh-in	Journal
Gender	0.3161 (4.0874)	1.2438 (3.5608)	0.5270 (1.3517)	-2.4414 * (1.3891)	-2.5769 (2.8832)	1.9165 (4.1030)
Duration	0.8260 (3.6547)	2.2015 * (1.3312)	3.9098 (4.1002)	-3.9043 (3.5247)	-4.0749 (6.6210)	2.3982 (5.5142)
Challenge	-4.1320 *** (1.9484)	2.7123 (1.8325)	5.3054 (7.9680)	-3.8486 (5.2020)	-3.0640 (3.8977)	2.7080 * (1.5234)
Post_init_rec	-0.1952 (2.0987)	5.3521 (4.0907)	-5.3021 (4.7762)	3.9533 (3.0903)	1.0681 (3.6360)	-1.5049 (5.5852)
Post_re_rec	-2.1122 (3.7306)	0.5246 (3.9073)	-9.7581 *** (1.3750)	9.6536 *** (1.1703)	-0.3245 (2.6570)	1.3053 (3.2983)
Friend_rec	-2.1516 (2.4134)	0.8988 (3.3708)	2.1557 (5.1939)	-0.5349 (2.7649)	6.4320 (12.5698)	-5.0136 (10.4883)
Supporter_rec	-8.7545 *** (2.0107)	15.4147 * (8.6133)	-30.2211 * (16.8257)	34.8592 * (20.6219)	-19.2785 ** (8.9173)	18.2246 *** (1.3366)
Constant	-0.6551 (0.1839)	-0.4358 (1.1474)	-0.6182 (1.6342)	-0.2895 (2.1873)	-0.2483 (1.4281)	-0.2474 (11.2469)
Variables Affecting State Transition						
Post_init_cum	8.7147 * (4.8812)		9.4034 *** (2.2333)		-0.7832 (1.7566)	
Post_re_cum	-10.9859 *** (1.4127)		9.3715 *** (1.6232)		0.0837 (2.3397)	
Friend_cum	-13.0692 ** (5.8256)		-16.3720 *** (5.4588)		-6.5529 ** (2.7163)	
Supporter_cum	13.1741 (16.7846)		-4.1294 *** (1.1327)		-5.6416 * (3.3078)	
Weighin_cum	7.7380 * (4.3538)		3.9171 * (2.0481)		-2.3664 (8.6156)	
Journal_cum	4.1624 ** (2.0810)		3.6798 *** (1.1662)		3.0388 (4.6341)	
Thresholds						
mu1 (1->3)	17.2952 *** (6.0856)					
mu2 (2->3)			16.2631 *** (3.1508)			
mu3 (3->2)					0.4387 (6.3206)	
Variables Affecting Waiting Time						
Freq_social	3.6476 (5.1152)		-1.0178 (6.4788)		0.8984 (10.3984)	
Freq_monitor	19.2960 *** (3.6837)		24.9976 *** (1.9530)		28.4453 *** (3.9045)	
Constant	-2.0301 *** (0.7471)		-1.8845 ** (0.8959)		-2.1644 ** (1.0034)	

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Note: Standard errors are reported in the parentheses. We took logarithms on all the variables to address the over-dispersion issue.

A.4.2 Alternative Variable Measures

We further investigate whether different measures of variables can affect our findings. In particular, we investigate the frequency variables affecting waiting time lengths. As noted in the

main text, we measure the frequency variables by individuals' average participation frequency until time t . As a robustness check, we explore a different measure: individuals' average participation frequency during the most recent week, as it is possible that the time individuals stay in a state may just be affected by their recent engagement intensity. We replace *Freq_social* and *Freq_monitor* by the new measures and re-estimate the model. The estimation results are displayed in Table A.4. Our main findings, again, remain consistent.

Table A.4 Robustness Check 2: Different Variable Measures

Parameters	State 1 (L)		State 2 (M)		State 3 (H)	
Variables Affecting State-Dependent Outcome						
	Weigh-in	Journal	Weigh-in	Journal	Weigh-in	Journal
Gender	-5.2439 *** (1.7067)	4.0583 * (2.4161)	-1.0710 (3.4213)	-0.1207 (1.7373)	3.5767 *** (1.2680)	-3.5574 (2.5710)
Duration	-3.2981 *** (1.2416)	3.8628 (2.3728)	0.4697 (1.3906)	-0.4511 (1.1606)	3.1834 * (1.8336)	-3.1166 (2.5996)
Challenge	-7.0622 * (3.9995)	4.9570 (3.7910)	2.4077 (3.3782)	-1.3579 (1.4799)	1.6333 (5.0024)	-0.7433 (2.1383)
Post_init_rec	-3.0570 *** (1.1448)	-5.2256 (3.4384)	-2.1327 (1.4750)	0.6377 (1.2215)	-4.5882 ** (2.1592)	3.6167 (2.8112)
Post_re_rec	-2.1224 * (1.0969)	-2.0995 (1.5241)	-3.2502 * (1.9128)	3.4074 * (1.8694)	-7.5653 * (4.4157)	6.8168 (5.1447)
Friend_rec	-0.4232 (1.1466)	0.8526 (1.3171)	2.1574 (4.7489)	-2.5589 (4.5148)	-0.6562 (3.2761)	-0.0279 (1.9801)
Supporter_rec	-5.5452 *** (1.1007)	10.1088 *** (2.9513)	-21.4340 * (11.9304)	22.5399 * (12.4774)	-22.5798 * (12.2125)	23.8650 (15.8978)
Constant	-0.6695 *** (0.1167)	-0.3849 (0.3391)	-0.6694 (5.7849)	-0.3849 (5.2734)	-0.6694 (8.3547)	-0.3848 (7.5396)
Variables Affecting State Transition						
Post_init_cum	8.8205 *** (1.3900)		11.1596 *** (1.3054)		0.7145 (2.7949)	
Post_re_cum	-8.6176 *** (1.6199)		11.1220 *** (1.2152)		0.6454 (3.0843)	
Friend_cum	-14.8127 *** (1.8481)		-11.0073 *** (1.0042)		7.3763 *** (1.5113)	
Supporter_cum	-3.0438 (2.9947)		-15.5487 *** (4.4936)		-8.8852 *** (2.1790)	
Weighin_cum	4.9477 *** (1.2641)		3.0878 (2.1133)		0.9633 (2.3036)	
Journal_cum	3.2244 ** (1.2892)		6.8657 (4.6489)		0.3737 (2.4315)	
Thresholds						
mu1 (1->3)	15.2427 *** (2.5276)					
mu2 (2->3)			20.1942 *** (1.3501)			
mu3 (3->2)					1.1664 (4.2818)	
Variables Affecting Waiting Time						
Freq_social	6.1110 (6.1433)		0.8988 (10.3590)		1.7319 (5.9740)	
Freq_monitor	18.4171 ** (8.6364)		25.7160 ** (10.7387)		28.3734 *** (7.9934)	
Constant	-2.1116 *** (0.5450)		-2.0114 ** (0.8685)		-2.1554 *** (0.6855)	

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Note: Standard errors are reported in the parentheses. We took logarithms on all the variables to address the over-dispersion issue.

A.4.3 Sampling Seeds

Finally, to show that the random sampling procedure does not affect our results, we change our sampling seed and re-estimate our model as a robustness check. We report the estimates in Table

A.5. Our findings remain consistent, indicating that the sampled data do not bias our results.

Table A.5 Robustness Check 3: Different Sampling Seed

Parameters	State 1 (L)		State 2 (M)		State 3 (H)	
Variables Affecting State-Dependent Outcome						
	Weigh-in	Journal	Weigh-in	Journal	Weigh-in	Journal
Gender	3.8110 ***	-3.9947 **	5.0712 ***	-5.7351 **	-11.9773 ***	9.4621 ***
	(1.3348)	(1.7433)	(1.8488)	(2.2609)	(3.6216)	(3.4648)
Duration	5.1837 ***	-3.6761 **	5.8639 ***	-5.0943 **	-11.1274 ***	8.8560 ***
	(1.7535)	(1.6741)	(2.0930)	(2.0468)	(3.3651)	(3.3520)
Challenge	4.2054 ***	-2.7385	6.8138 ***	-4.6413 **	-12.9504 ***	9.2539 ***
	(1.3386)	(1.6675)	(1.9714)	(2.0873)	(3.9051)	(3.3802)
Post_init_rec	-6.6062 ***	3.7287 **	-7.1964 ***	5.1425 **	-11.6342 ***	9.3289 ***
	(2.0410)	(1.7160)	(2.3945)	(2.0533)	(3.5225)	(3.5019)
Post_re_rec	-8.2510 ***	6.6403 ***	-8.3019 ***	7.5659 ***	-10.5029 ***	7.4513 **
	(2.4166)	(2.3692)	(2.6391)	(2.6487)	(3.1527)	(3.0130)
Friend_rec	-3.5278 **	2.9327 *	-5.2080 ***	4.1875 **	8.1568	-3.2642
	(1.4205)	(1.5799)	(1.9340)	(1.8233)	(5.8797)	(1.9959)
Supporter_rec	-16.1660 ***	18.5099 ***	-13.6861 ***	16.2842 ***	-2.8997 **	5.3830 ***
	(3.8898)	(4.9793)	(3.3099)	(4.2934)	(1.3085)	(1.7252)
Constant	-6.7410 ***	-4.2599 ***	-6.7387 ***	-4.2391 ***	-6.7380 ***	-4.1759 *
	(1.2398)	(1.5961)	(1.0891)	(1.2715)	(1.3556)	(2.2927)
Variables Affecting State Transition						
Post_init_cum	-0.3138	(2.3742)	2.7962 **	(1.2927)	1.9385	(1.2549)
Post_re_cum	-2.1568	(2.6355)	2.0792 *	(1.2606)	1.7968	(1.2105)
Friend_cum	-0.2577	(0.5552)	-21.8327 **	(8.8554)	6.1023 ***	(2.1112)
Supporter_cum	1.0046	(0.8354)	35.5452 ***	(9.0504)	-1.8776	(2.1509)
Weighin_cum	-2.9523	(4.5575)	2.6208 **	(1.3120)	1.3523	(1.2162)
Journal_cum	11.0034 ***	(4.1618)	0.6186	(1.1008)	2.5201 *	(1.3693)
Thresholds						
mu1 (1->3)	3.0963	(3.5495)				
mu2 (2->3)			3.8692 **	(1.6518)		
mu3 (3->2)					3.0766 *	(1.5989)
Variables Affecting Waiting Time						
Freq_social	0.9503	(5.2238)	0.6331	(4.4324)	-3.2780 **	(1.6271)
Freq_monitor	26.3410 ***	(9.2655)	26.7601 ***	(6.6321)	18.1896 ***	(3.5153)
Constant	-2.0225 ***	(0.6265)	-2.0956 **	(1.0239)	-1.9362 ***	(0.5620)

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Note: Standard errors are reported in the parentheses. We took logarithms on all the variables to address the over-dispersion issue.

A.5 SINGLE-DIMENSIONAL MODELS

In this section, we report the model estimates for the single-dimensional CTHMMs. The first model uses the number of weigh-ins as the state-dependent outcome, whereas the second model uses the number of journals as the state-dependent outcome. We present the estimation results in Table A.6 and Table A.7, respectively. Table A.6 shows that individuals who have participated in posting activities tend to decrease their weigh-ins when they are at a low efficacy level. Table A.7 shows that individuals who have established social ties may tend to increase their journal-recording behaviors. These results are in line with our main findings. However, in terms of state-transition estimates, we find that the results for single-dimensional models can be different from our main results. For instance, the results in Table A.6 suggest that the effect of supporters is negative at L state, whereas in our main analysis the corresponding effect is positive; the results in Table A.7 show that the effect of replying to posts is positive at L state, whereas in our main analysis the corresponding effect is negative. We deem these inconsistencies reasonable, as the single-dimensional models leverage only a partial set of information to calibrate individuals' online weight-management behaviors.

In addition, we find that the estimation results in Table A.6 and Table A.7 can be different from each other as well, even though both weigh-ins and journals are valid measure for individuals' self-regulatory behaviors. For example, Table A.6 shows that initiating posts tends to positively affect individuals' efficacy state when at state H (6.0216, $p < 0.1$), whereas Table A.7 shows that the corresponding effect is negative (-1.2232, $p < 0.1$); Table A6 shows that weight-in can negatively affect efficacy state when at state H (-8.4579, $p < 0.01$), whereas the corresponding effect is suggested to be positive in Table A.7 (10.4300, $p < 0.1$). These differences indicate that weigh-ins and journals provide heterogeneous measurements for individuals' self-regulatory

behaviors, and that partially relying on one measurement may result in deriving completely different implications.

Table A.6 Single Dimensional CTHMM: Weigh-In

Parameters	State 1 (L)		State 2 (M)		State 3 (H)	
Variables Affecting State-Dependent Outcome (Weigh-in)						
Gender	-2.5003	(2.9894)	-1.5827	(1.7362)	1.2316	(1.2666)
Duration	2.7127 *	(1.3972)	0.3398	(1.5787)	1.7509	(1.7269)
Challenge	0.7350	(4.0839)	-0.2073	(1.3027)	0.3032	(1.1877)
Post_init_rec	-2.8995 **	(1.4374)	-0.5452	(1.5130)	-1.6667	(1.4898)
Post_re_rec	-2.5329 *	(1.5344)	-1.4736	(1.6189)	-1.9037	(1.6454)
Friend_rec	-1.3645	(1.4635)	-0.0380	(1.5612)	0.2119	(1.4448)
Supporter_rec	1.9980	(3.5998)	-0.5443	(1.2825)	-0.8420	(1.2339)
Constant	10.3965 **	(4.1783)	10.9446 ***	(3.6558)	11.0751 **	(4.4867)
Variables Affecting State Transition						
Post_init_cum	-0.4695	(1.0019)	-0.2500	(1.4728)	6.0216 *	(3.3351)
Post_re_cum	-0.3062	(1.0004)	-0.9648	(1.6004)	3.2795	(5.4134)
Friend_cum	-2.0943 **	(1.0609)	-11.0112	(9.2389)	1.1167	(2.2707)
Supporter_cum	-2.8911 ***	(1.0647)	-13.9992	(12.6385)	2.2601	(2.5511)
Weighin_cum	0.9871	(1.0121)	-8.9064 **	(4.0370)	-8.4579 ***	(3.1510)
Journal_cum	0.8796	(1.4910)	-1.9981 **	(1.0098)	-14.3489	(9.6558)
Thresholds						
mu1 (1->3)	-9.1373 ***	(1.7613)				
mu2 (2->3)			-4.8824	(10.5370)		
mu3 (3->2)					-0.1904	(3.3662)
Variables Affecting Waiting Time						
Freq_social	1.5657	(6.3211)	-0.6156	(1.6949)	-12.5037 ***	(2.5832)
Freq_monitor	3.9356 **	(1.7903)	0.8932	(1.4484)	4.8659 ***	(1.1021)
Constant	-3.6203 ***	(0.7324)	-1.4462 **	(0.7151)	-3.1158 ***	(0.6223)

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Note: Standard errors are reported in parentheses. We took logarithms on all the variables to address the over-dispersion issue.

Table A.7 Single Dimensional CTHMM: Journal

Parameters	State 1 (L)		State 2 (M)		State 3 (H)	
Variables Affecting State-Dependent Outcome (Journal)						
Gender	0.4061	(1.2846)	0.5039	(1.2217)	-0.6061	(1.6673)
Duration	2.6375 *	(1.5201)	0.1977	(1.0612)	2.5963 *	(1.5343)
Challenge	0.3331	(1.3085)	-0.8273	(1.4405)	0.4254	(1.1462)
Post_init_rec	-2.5937 **	(1.7315)	-0.4672	(1.1082)	-3.1514	(2.3574)
Post_re_rec	-0.4721	(1.4610)	1.2031	(1.7158)	-0.9435	(1.0722)
Friend_rec	1.7365	(1.2372)	1.7279 *	(1.0190)	0.6066	(1.3233)
Supporter_rec	4.6753 **	(2.3100)	3.1594	(2.4242)	3.4594 **	(1.5221)
Constant	6.0073 **	(2.5779)	6.0074	(7.8804)	6.0075	(4.8472)
Variables Affecting State Transition						
Post_init_cum	5.5520 *	(3.0494)	-0.5528	(1.1438)	-1.2232 *	(2.2130)
Post_re_cum	3.0367 *	(1.6079)	-1.0489	(1.0523)	-5.6103	(6.4440)
Friend_cum	0.2909	(3.6125)	-8.9824 *	(5.3708)	0.0698	(1.6198)
Supporter_cum	0.6586	(2.0224)	4.4819 *	(2.5257)	-0.1778	(1.4798)
Weighin_cum	-19.2451	(21.0761)	1.3649	(1.1109)	10.4300 *	(6.2658)
Journal_cum	-37.5502 **	(18.0771)	-0.0555	(1.1685)	24.6909	(16.5454)
Thresholds						
mu1 (1->3)	-0.6792	(3.4534)				
mu2 (2->3)			-0.7582	(1.1186)		
mu3 (3->2)					-1.5913	(1.8082)
Variables Affecting Waiting Time						
Freq_social	4.7753 ***	(1.5474)	9.7695	(7.9903)	8.6408 ***	(2.5134)
Freq_monitor	2.1036 ***	(0.8398)	2.4705 ***	(0.9567)	2.2841	(2.0601)
Constant	-2.6084 ***	(0.4314)	-3.2633 ***	(0.9044)	-2.8883 ***	(0.8544)

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Note: Standard errors are reported in parentheses. We took logarithms on all the variables to address the over-dispersion issue.

A.6 PREDICTION PROCEDURE

In our main text, we predict individuals' self-monitoring behaviors using different types of HMMs as a means to illustrate the effectiveness of our MCTHMM framework. The prediction procedure can be summarized as follows. First, we simulate individuals' initial efficacy state by using the initial hidden state distribution estimated by the LCM. We then simulate the state-dependent outcomes. For the two-dimensional models, we simulate not only the number of weigh-ins, but also the number of journals. For the one-dimensional models, we simulate only one type of the

self-monitoring activities. Further, based on the previous observations about individuals' online activities, we can compute the posterior hidden state distribution by applying Bayes rule (Hamilton 1989). Specifically, when calculating the posterior distribution at time t , we assume that only the activities before time t can be observed. Based on these observations, we can compute the posterior likelihood and update the hidden state distributions according to the Bayes rule. The entire process is repeated until the observation window ends.

Because a typical weight-loss therapy includes weekly consultations (Ballantine and Stephenson 2011; Moisio and Beruchashvili 2009), we consider individuals' self-monitoring behaviors on a weekly basis. Particularly, for the continuous-time models, we aggregate individuals' predicted and observed self-monitoring behaviors into weeks and use the binary-valued outcome to compute the prediction accuracy in percentage. For the discrete-time model, which naturally produces an outcome on a weekly basis, we directly compare the predicted results with the observed data.

References

- Ballantine P W, and Stephenson R J. (2011) Help Me, I'm Fat! Social Support in Online Weight Loss Networks. *Journal of Consumer Behaviour* 10(6): 332-337.
- Centers for Disease Control and Prevention. (2020) Prevalence of Obesity and Severe Obesity among Adults: United States, 2017–2018, Accessed March 19, 2020. <https://www.cdc.gov/nchs/products/databriefs/db360.htm>.
- Hamilton J D. (1989) A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica: Journal of the Econometric Society*: 357-384.
- Moisio R, and Beruchashvili M. (2009) Questing for Well-Being at Weight Watchers: The Role of the Spiritual-Therapeutic Model in a Support Group. *Journal of Consumer Research* 36(5): 857-875.

APPENDIX B

As noted in the main text, we use a MCMC approach to estimate the parameters. The estimation procedure is as follows. We first set initial values for the model parameters: $b_0, b_1, b_2, b_3, \Sigma_\beta, \beta_i, \theta_0, \sigma_w^2, \sigma_{y_k}^2, \alpha, \lambda$. Let $B = (b_0, b_1, b_2, b_3)^T$, which has dimension 4×5 . Let X be a $n \times 4$ design matrix. Each row of X corresponds to one user's record, including the intercept, gender, age, and initial weight. Let BI denote the random-coefficient matrix, with each row being β_i . BI has dimension $n \times 5$. BI is centered around $B_0 = XB$. We iterate the following steps until convergence. Let r denote the parameter values at round r .

Step 1: Generate conditional posterior for B^r and Σ_β^r .

Assume hyper priors for B : $B \sim \mathcal{MN}_{4,5}(Mu_0, Sig_0, \Sigma_\beta)$, where $\mathcal{MN}$ denotes matrix normal distribution. Since we have no prior knowledge on B , we assume diffuse hyper priors on B : $Mu_0 = J_{5 \times 4}$, where J is an all-ones matrix with dimension 5×4 ; $Sig_0 = 10 * I_{4 \times 4}$, where I denotes an identity matrix. The conditional posterior for B^r can be generated as follows:

$$B^r \sim \mathcal{MN}(Mu_1, Sig_1, \Sigma_\beta), \quad (B.1)$$

$$Sig_1 = ((X^{*r-1})^T X^{*r-1} + Sig_0^{-1})^{-1}, \quad (B.2)$$

$$Mu_1 = Sig_1(Sig_0^{-1}Mu_0 + (X^{*r-1})^T BI^{*r-1}), \quad (B.3)$$

where BI^{*r-1} contains the accepted β_i in round $r-1$, and X^{*r-1} is the design matrix for the accepted instances.

Assume hyper priors $\Sigma_\beta \sim \mathcal{IW}(d_0, \Delta_0)$, where $\mathcal{IW}$ represents the inverse-Wishart distribution with degree of freedom d_0 and scale matrix Δ_0 . To let the estimation results driven

by data, we set diffuse hyper priors: $d_0 = 10$, $\Delta_0 = I_{npar}$, where I denote the identity matrix, and $npar$ is the number of parameters in β_i . In our setting, $npar = 5$. The conditional variance Σ_β^r can be generated as follows:

$$\Sigma_\beta^r \sim IW(df, \Gamma), \quad (B.4)$$

$$\Gamma = ((BI^{*r-1} - B_0^{*r-1})^T (BI^{*r-1} - B_0^{*r-1}) + \Delta_0)^{-1}, \quad (B.5)$$

$$df = m + d_0. \quad (B.6)$$

BI^{*r-1} and B_0^{*r-1} contain the accepted entries in round $r-1$, m is the number of samples that are accepted.

Step 2: Generate β_i^r

β_i^r can be generated from the following posterior distribution:

$$\beta_i \propto |\Sigma_\beta|^{-\frac{1}{2}} \exp(-\frac{1}{2}(\beta_i - \beta_{0i})^T \Sigma_\beta^{-1} (\beta_i - \beta_{0i})) L_i(D_i | \beta_i, \sigma_w^2, \theta_0, \sigma_{yk}^2, \alpha, \lambda), \quad (B.7)$$

where β_{0i} is the i -th row of B_0 . Since the posterior distribution does not have a closed form, we use the Metropolis-Hasting method to make the draws. The proposal distribution is a normal distribution with mean β_i^{r-1} (i.e., β_i in the last iteration) and variance ν (e.g., $\nu = 0.2 * I$). The acceptance rate is calculated as

$$acc = \frac{\exp(-\frac{1}{2}(\beta_i^* - \beta_{0i})^T \Sigma_\beta^{-1} (\beta_i^* - \beta_{0i})) L_i(D_i | \beta_i^*, \sigma_w^2, \theta_0, \sigma_{yk}^2, \alpha, \lambda)}{\exp(-\frac{1}{2}(\beta_i^{r-1} - \beta_{0i})^T \Sigma_\beta^{-1} (\beta_i^{r-1} - \beta_{0i})) L_i(D_i | \beta_i^{r-1}, \sigma_w^2, \theta_0, \sigma_{yk}^2, \alpha, \lambda)}, \quad (B.8)$$

where β_i^* is the proposed parameter drawn from the proposal distribution. If accept, $\beta_i^r = \beta_i^*$; else, $\beta_i^r = \beta_i^{r-1}$.

Step 3: Generate $\varphi = \{\log(\sigma_w^2), \log(\sigma_{yk}^2), \theta_0, \alpha, \lambda\}$ from the posterior distribution:

$$\varphi \propto |\Sigma_0|^{-\frac{1}{2}} \exp(-\frac{1}{2}(\varphi - \varphi_0)^T \Sigma_0^{-1}(\varphi - \varphi_0))L(D | \beta_i, \varphi). \quad (\text{B.9})$$

Again, we assume a diffuse prior for φ , that is, we set $\varphi_0 = (0, 0, \dots, 0)^T$ and $\Sigma_0 = 10 * I$, so that the parameter drawn from the posterior distribution are mostly driven by data patterns and less affected by the prior knowledge. Similarly, we use a Metropolis-Hasting method to draw φ .

Proposed parameters are accepted with the rate:

$$acc = \frac{\exp(-\frac{1}{2}(\varphi^* - \varphi_0)^T \Sigma_0^{-1}(\varphi^* - \varphi_0))L(D | \beta_i, \varphi^*)}{\exp(-\frac{1}{2}(\varphi^{r-1} - \varphi_0)^T \Sigma_0^{-1}(\varphi^{r-1} - \varphi_0))L(D | \beta_i, \varphi^{r-1})}. \quad (\text{B.10})$$

If accept, $\varphi^r = \varphi^*$; else, $\varphi^r = \varphi^{r-1}$.

APPENDIX C

C.1 DEEP LEARNING EMBEDDINGS

C.1.1 Challenge Embedding Model

In this section, we formally demonstrate how challenge embeddings are constructed. Challenge embeddings will serve as a building block for the later user embedding model and the downstream challenge recommendation tasks. The goal for our challenge embedding model is to build a vector for each challenge such that similar challenges will have close embedding vectors. As long as the embedding learned from our model is able to properly distinguish challenges based on their relative similarity and intrinsic attributes, the user representation learning model will properly extract relevant signals to serve the challenge recommendation task.

There are 165 challenges, each with a challenge name, description, and the annotated meta attributes. The challenge name and description contain semantic information about the challenges, while the meta attributes summarize high-level latent properties of the challenges. To build challenge embeddings, we design a hybrid neural network, which takes challenge names and descriptions as inputs and challenge meta attributes as the output. As challenge name $n(i)$ is a short string of several words, we use average token-level embedding to extract signals. For challenge description $d(i)$, which contains many more words, we apply vanilla LSTM to capture the semantics and extract signals. Most of the challenge meta attributes are Boolean, such as *Measurable*, *Specific*, *Diet*, *Exercise*, etc., except *Duration*. We consider challenge duration as an integer variable and other attributes as categorical. We use $attr_{challenge}(i, j)$ to denote the j -th attribute of challenge i .

Formally, we first apply pre-trained token-level embedding to the tokenized name and description of a challenge to obtain $v_n(i) = (v_n(i,1), v_n(i,2), \dots, v_n(i, k_i))$ and $v_d(i) = (v_d(i,1), v_d(i,2), \dots, v_d(i, l_i))$, namely, the token-level embedding sequences for the name and description of challenge i . Then the flattened name vector is calculated as

$$v_n(i) = \frac{1}{k_i} \sum_{m=1}^{k_i} v_n(i, m), \quad (C.1)$$

and the flattened description is calculated as the vanilla LSTM processed token-sequence

$$v_d(i) = LSTM(v_d(i,1), v_d(i,2), \dots, v_d(i, l_i)). \quad (C.2)$$

We concatenate the flattened name vector and description vectors, followed by three fully connected layers. For the pre-trained token-level embedding, we adopt the FastText model, which is able to leverage sub-token information for token embedding. We further fine-tune the token-level embeddings during the challenge representation learning to ensure low information loss. The dimension of LSTM cell latent vector is set to 100, which equals the FastText dimension; hence, the contributions of the challenge name and description are balanced. The dimensions of the fully connected layers are 256, 64, and 16, respectively. The loss at the end is cross-entropy for the categorical meta attributes and mean-square-error for challenge duration; that is,

$$OBJ_{challenge} = \frac{1}{165} (MSE(attr_{challenge}(i, 0), F_{challenge}(v_n(0), v_d(0))) + \sum_{j \neq 0} CrossEntropy(attr_{challenge}(i, j), F_{challenge}(v_n(j), v_d(j))))), \quad (C.3)$$

where we use $attr_{challenge}(i, 0)$ to denote challenge duration. We use $c(i)$ to denote the learned embedding for challenge i , which is a 16-dimensional vector.

C.1.2 User Embedding Model

To better learn user representations and facilitate recommendation personalization, we focus on both users' attribute features (shallow features) and their sequence features regarding health

histories and behavior paths (deep features). Whereas users' attribute features provide a basic characterization of personal heterogeneity, the sequence features provide a deeper contextual understanding of users' health behaviors (Johnson et al. 2002). As such, these two types of features describe user interests from different aspects. For each user, the attribute variables (e.g., gender, age, etc.) usually remain unchanged over time, whereas health status and behavioral sequences will vary with time. Therefore, the user embeddings, which incorporate information of both user attributes and behavioral sequences, will depend on time. This is aligned with our objective of offering dynamic healthcare-intervention suggestions based on users' changing health-management contexts.

We use $a_1, a_2, \dots, a_n$ to denote the attribute features and we convert them into one-hot encoding if the attribute is categorical. The converted attribute features are $\tilde{a}_1, \tilde{a}_2, \dots, \tilde{a}_n$. We concatenate them to get a single vector $\tilde{a}$ as the input for the wide branch. Here $n = 4$, and the corresponding four shallow features are users' gender, age, initial weight, and membership duration. We apply a standard 3-layer fully connected neural network to handle shallow features, with details as follows:

$$\xi_1 = \text{Relu}(W_1^a \tilde{a} + b_1^a), \quad (\text{C.4})$$

$$\xi_2 = \text{Relu}(W_2^a \xi_1 + b_2^a), \quad (\text{C.5})$$

$$\xi_3 = \text{Relu}(W_3^a \xi_2 + b_3^a). \quad (\text{C.6})$$

After the fully connected layer, we obtain a vector ξ_3 , which is the endpoint of the shallow features.

We use $s_1^t, s_2^t, \dots, s_m^t$, to denote the deep features, where each of them is a sequence of numerical values or texts with time coordinate t . Here $m = 5$, and the corresponding features are the sequences of users' historical weight, the historical challenges chosen by users, the historical number of friends, the historical posts that users published in the community forum, and the

historical number of weigh-ins. We use BERT embedding for each text feature (i.e., users' posts) and one-hot encoding for other features, and the raw deep input features are converted to $\tilde{s}_1^t, \tilde{s}_2^t, \dots, \tilde{s}_m^t$. Note that the BERT encoding is set un-trainable since a lower-level understanding of specific meanings of each text, e.g., posts, is not the focus of this model. Instead, we focus on extracting the patterns from the sequences.

One challenge to handle these five sequences is that their feature dimensions and the richness of the signals vary significantly. It is problematic to apply five separate LSTM modules to handle each of the sequences individually. The coupling between these sequences at the same time is crucial for the understanding of the deep features. Therefore, we propose using a single but enhanced LSTM module, called eLSTM, to simultaneously process the five input sequences. Formally, we apply a transformation to each of the encoding sequences at the beginning:

$$r_1^t = \text{Relu}(W_1^S \tilde{S}_1^t + b_1^S), r_2^t = \text{Relu}(W_2^S \tilde{S}_2^t + b_2^S), \dots, r_m^t = \text{Relu}(W_m^S \tilde{S}_m^t + b_m^S). \quad (\text{C.7})$$

This transformation brings different features with different scales in different spaces together, followed by a concatenation operation

$$\tilde{r}^t = \text{Concat}(r_1^t, r_2^t, \dots, r_m^t), \quad (\text{C.8})$$

and then we apply another nonlinear transformation to further merge the information for different sequences to get a unified input sequence of encodings as:

$$x^t = \text{Relu}(\tilde{W}^S \tilde{r}^t + \tilde{b}^S). \quad (\text{C.9})$$

We apply a LSTM module to this unified input sequence via

$$f_t = \sigma(W_f \cdot [h_{t-1}, x^t] + b_f), \quad (\text{C.10})$$

$$i_t = \sigma(W_i \cdot [h_{t-1}, x^t] + b_i), \quad (\text{C.11})$$

$$\tilde{C}_t = \tanh(W_c \cdot [h_{t-1}, x^t] + b_c), \quad (\text{C.12})$$

$$C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t, \quad (C.13)$$

$$o_t = \sigma(W_o \cdot [h_{t-1}, x^t] + b_o), \quad (C.14)$$

$$h_t = o_t \cdot \tanh(C_t), \quad (C.15)$$

and eventually obtain a sequence of latent vectors $[h_1, h_2, \dots, h_t]$. Note that the length of the sequence t is also a variable because we build user representations given a history chunk with any size. It is difficult to directly aggregate the latent vector sequence, and naively using the last cell's output could potentially lose partial information of the earlier steps. To ensure that the neural network automatically detects the importance of each hidden vector and learns in favor of more important ones, we use a weighted average where the weight is the attention weight learned from a fully connected neural network:

$$u_t = \tanh(W^{att} \cdot h_t + b^{att}), \quad (C.16)$$

$$\chi_t = \exp(u_t \cdot v) / \sum_t \exp(u_t \cdot v), \quad (C.17)$$

$$\eta_1 = \sum_t \chi_t h_t. \quad (C.18)$$

Concatenating the endpoint of the shallow branch ξ_3 and the endpoint of the deep branch η_1 , we eventually obtain a latent layer combining the information from both the wide features and the deep features:

$$\theta_1 = \text{Concat}(\xi_3, \eta_1). \quad (C.19)$$

Now we start to construct the loss branches. The main branch combining signals from both the shallow features and the deep features is built for a challenge embedding prediction:

$$\theta_2 = \text{Relu}(W_1^{user} \theta_1 + b_1^{user}), \quad (C.20)$$

$$\theta_3 = W_2^{user} \theta_2 + b_2^{user}, \quad (C.21)$$

where θ_3 is the final prediction. By omitting its dependency on the neural network parameters, θ_3 becomes a function of user u and time t : $\theta_3 = \theta_3(u, t)$.

For each user u and time (i.e., week) t , if there exists a challenge selection, we get the average challenge embedding $c(u, t)$, and we calculate the cosine distance between the user embedding $f(u, t)$ and $c(u, t)$. We sum up the distance over all users and time to construct our objective loss function. Note that a user may not select any challenge during some time t . We do not include such cases in our objective function, as they do not provide signals of user preference. Formally, we define our objective function as

$$OBJ_{user} = -\frac{1}{|\Omega|} \sum_{(u,t) \in \Omega} \frac{f(u,t) \cdot c(u,t)}{\|f(u,t)\|_2 \|c(u,t)\|_2}, \quad (C.22)$$

where $\Omega = \{(u, t) \mid \text{there exists a challenge selection from user } u \text{ at time } t\}$, $|\Omega|$ denotes the cardinality of set Ω , and $\|\cdot\|_2$ is the l_2 norm. Note that there is a uniform denominator before the sum, regardless of how many challenges a user has selected during the time window $[1, T]$. Therefore, a more active user will contribute more terms in the objective function. In this way, the objective function will favor active users for more guidance during user representation learning.

For the auxiliary-loss branch, we apply fully connected operators to the endpoint of the deep features. In this way, the learning of the eLSTM module does not interfere with the wide features; hence, the gradients are well propagated to better learn the weights:

$$\eta_2 = \text{Relu}(W_1^{user} \theta_1 + b_1^{user}), \quad (C.23)$$

$$\eta_3 = \text{Relu}(W_2^{health} \eta_2 + b_2^{health}). \quad (C.24)$$

The auxiliary prediction η_3 is a prediction of the weight loss $\alpha(u, t)$, and we choose a combined loss of MSE for absolute value prediction and cross-entropy loss for the weight-loss sign prediction. Finally, the full loss function for user representation learning is written as

$$L = -\frac{1}{|\Omega|} \sum_{(u,t) \in \Omega} \frac{f(u,t) \cdot c(u,t)}{\|f(u,t)\|_2 \|c(u,t)\|_2} + \frac{1}{|\Omega|} \sum_{(u,t) \in \Omega} |\eta_3 - \alpha(u,t)|^2 + \text{CrossEntropy}(\text{sign}(\eta_3), \text{sign}(\alpha(u,t))) .$$

(C.25)

C.2 SCREENSHOT OF WEIGHT-LOSS CHALLENGES ON THE FOCAL WEBSITE

Figure C.1 is a screenshot of the challenge page on the focal weight-loss platform. As can be seen, six challenges are included in this screenshot; the first three had not been started at the time we took this screenshot, and the last three were already in progress and could not be joined. For each of the challenges, users can find the challenge title at the top, followed by a short description of the challenge guidelines and the duration. Users can consider this information when they make choices about the challenges in which to participate. To join a challenge, users need to click on the challenge title and enter the inner homepage of the challenge, where a “join challenge” button can be found. We provide Figure C.2, in which we show the internal homepage of the first challenge (i.e., “Sassy Summer Slimdown”) as an example.

[Home](#) > [Challenges](#)


Challenges

Find a Challenge that's right for you and start achieving your goals today!



Sassy Summer Slimdown

Whether you have a lot or just a little to lose, join us in getting in shape for the next six weeks.

36

starts in: **5 days**
duration: 6 weeks
day 1: Monday 27 May



Daily Accountability

NEW Start Weekly! Develop discipline leading to success in any plan! Do only the minis you want! Make friends and have fun!

14

starts in: **5 days**
duration: 7 days
day 1: Monday 27 May



Feeling good by LOOSING

Get ready for a summer holiday by losing a kilo a week.

2

starts in: **a day**
duration: 4 weeks
day 1: Thursday 23 May



36 hour fast

this challenge is about fasting for 36 hours, eat normally for 8 hours go back to 36 hours fast.

6

status: **In Progress**
currently: Week 1 of 4
end date: Monday 17 Jun



Daily Accountability

NEW Start Weekly! Develop discipline leading to success in any plan! Do only the minis you want! Make friends and have fun!

57

status: **In Progress**
currently: Day 2 of 7
end date: Monday 27 May



Swimsuit Slimdown PART 2 (SS#2)

Starting just b4 Memorial Day & going through July 4th - Get your Firecracker Body ready to make a splash at the pool!

79

status: **In Progress**
currently: Week 1 of 7
end date: Monday 08 Jul

Figure C.1 A Screenshot of Challenge Homepage




Figure C.2 A Screen Shot of the Challenge Inner Homepage

C.3 CHALLENGE META-DATA ANNOTATION

As noted, we use the SMART metric to guide our annotation of the challenge description data. SMART stands for five important elements for describing a challenge goal: specific, measurable, attainable, relevant, and time. Based on the SMART metric, we evaluate the following features for each challenge: whether the challenge is specifically defined (specificity), whether the challenge goal is measurable (measurability), the intensity level of the challenge (attainability), whether the challenge is related to diet or physical activity (relevancy), the time span of the challenge (duration). To illustrate how we annotate these challenge features, we provide several challenge examples in Table C.1.

We determine a challenge to be specific if it has a clearly defined goal (i.e., what to do), and we determine a challenge to be measurable if it can be easily measured with certain numeric criteria

(i.e., how much to do). For example, In Table C.1, the challenge “30 minutes on a treadmill for 6 weeks” is measurable but not specific, because individuals only know how much time to spend for exercising, but they do not know what exercise to do (walking or running). The challenge “Eat salad for dinner for 2 months” is specific but not measurable, because it clearly specifies the recipe but does not provide information about how much individuals should eat. The challenge “Jog 5 miles a week” is both specific and measurable, as it is clear in the activity form and the way to measure completion. Similarly, the challenge “Lose 10 lbs. in 4 weeks” is also both specific and measurable, as individuals can clearly identify the weight-loss goal and measure challenge completion.

Table C.1 SMART Challenge Features

Challenge Description	Diet	Activity	Weight-Loss	Specific	Measurable	Intensity	Duration
30 minutes on a treadmill for 6 weeks	0	1	0	0	1	moderate	6
Eat salad for dinner for 2 months	1	0	0	1	0	moderate	8
Jog 5 miles a week	0	1	0	1	1	high	1
Lose 10 lbs. in 4 weeks	0	0	1	1	1	high	4

As shown in Table C.1, a challenge can be diet-oriented (challenge 2), physical-activity-oriented (challenge 1, 3), and weight-loss-oriented (challenge 4). For each challenge type, we develop rules to determine their respective intensity levels. The general idea behind these rules is that the more restrictions a challenge places on individuals’ calorie intake or expenditure, the higher the intensity level it has. In particular, for a diet-oriented challenge, we determine its intensity level to be low if it requires individuals to change a minor eating or drinking habit, such as drinking more water or avoiding soda. We determine the intensity level to be moderate if the challenge requires individuals to change their diet structures (e.g., eating certain types of food) or to moderately restrict the amount of calorie intake (e.g., cutting the serving size). Finally, we determine the intensity level to be high if the challenge requires individuals to significantly restrict

the amount of calorie intake, such as going through a fast. For a physical-activity-oriented challenge, the intensity level can be determined by the metabolic equivalent of task (MET). Specifically, walking and strolling are defined to be at a low-intensity level; activities such as bicycling, yoga, and walking with speed more than 3 mph are determined to be at a moderate intensity level; and activities such as jogging or running, situps, pushups, or jumping rope are considered to be at a high-intensity level. For a weight-loss-oriented challenge, we define the intensity level based on the average weight-loss rate per week. According to the Centers for Disease Control and Prevention (CDC), individuals can generally lose 1-2 lbs. per week by burning 500-1000 calories more than they consume each day. Therefore, challenges that require individuals to lose weight within this range are considered moderate. Challenges that require individuals to lose less than 1 lbs. per week are considered low-intensity, and challenges that require individuals to lose more than 2 lbs. per week are considered high-intensity.

In addition to the above-described SMART-based features, we consider two other features that may affect individuals' engagement in challenge participation: *Motivational* and *Self-Monitoring*. *Motivational* denotes whether a challenge contains encouraging words or sentences that can help individuals to build up the inner motivation for weight-loss. For example, "Just adding a piece of fruit to the diet daily. Nothing crazy here, healthy habits take a while to kick in, so stick with it!" is a diet-oriented challenge that is described in a motivational way. *Self-Monitoring* describes whether a challenge encourages individuals to regularly self-monitor their weight-loss progress (e.g., body weight, daily diet, running mileage, etc.). For instance, "8 weeks of eating a 1,200 calorie diet; keep track of all food and weigh in every Monday" is a diet-oriented challenge that encourages individuals to self-monitor their diet and body weight.

C.4 DATA SUMMARY STATISTICS

We summarize key data statistics in Table C.2.

Table C.2 Summary Statistics

Variable	Mean	Std. Dev.	Min	Max
<i>Challenge Description</i>				
No. of Available Challenges Each Week	51.06	2.30	45	54
Challenge Duration	5.01	2.84	1	20
Total No. of Challenges: 165				
<i>User Statistics</i>				
Gender (Female = 1, Male = 0)	0.52	0.50	0	1
Age	38.07	11.26	19	57
Membership Duration	394.10	157.97	126	672
Initial Weight	210.62	49.10	127	298
Weight Changes Per Week	0.02	3.05	-24	22
No. of Weigh-ins Per Week	1.01	1.14	0	6
No. of New Friends Per Week	0.97	0.88	0	4
Total No. of Users: 1049				
<i>User-Challenge Interactions</i>				
No. of Challenges Selected Per Week	1.94	1.69	0	13

C.5 TOP-K PRECISION, RECALL, AND F-MEASURE

Our recommendation algorithm offers top-K challenges to each user on a weekly basis. In Table C.3, we report the precision, recall, and F-measure for $K = 5, 10, 15$.⁷

⁷ Although we provide these measures for reference purposes, it is worth noting that our recommendation is different from a classification task. In Section A6, we introduce the major evaluation approaches that we adopt in our main analysis.

Table C.3 SMART Challenge Features

Models	Precision			Recall			F-Measure		
	K=5	K=10	K=15	K=5	K=10	K=15	K=5	K=10	K=15
proposed	0.5861	0.6054	0.6186	0.0315	0.0643	0.0998	0.0598	0.1163	0.1718
No constraints	0.5853	0.5974	0.6146	0.0315	0.0640	0.0991	0.0597	0.1156	0.1707
No embeddings	0.4829	0.5519	0.5554	0.0260	0.0589	0.0896	0.0493	0.1064	0.1543
No embs, no cons	0.4713	0.5462	0.5374	0.0253	0.0583	0.0867	0.0481	0.1054	0.1493
No challenge embs	0.4810	0.5625	0.5639	0.0259	0.0604	0.0910	0.0491	0.1091	0.1567
No user embs	0.5802	0.5933	0.6144	0.0312	0.0636	0.0991	0.0592	0.1149	0.1707
UCB	0.5225	0.5266	0.5554	0.0281	0.0563	0.0896	0.0533	0.1017	0.1543
ε -greedy	0.5773	0.5242	0.6058	0.0310	0.0561	0.0977	0.0588	0.1014	0.1683
CACF	0.3486	0.3149	0.3122	0.0187	0.0342	0.0505	0.0355	0.0617	0.0869
SCF	0.3781	0.3970	0.3952	0.0204	0.0414	0.0636	0.0387	0.0750	0.1096
PMF	0.2190	0.2124	0.2122	0.0118	0.0224	0.0342	0.0224	0.0405	0.0589
CAMF	0.2759	0.2653	0.2372	0.0148	0.0286	0.0382	0.0281	0.0516	0.0658
hybrid_pure	0.2403	0.2492	0.2349	0.0132	0.0266	0.0378	0.0250	0.0481	0.0651
hybrid_cacf	0.2893	0.2897	0.2747	0.0157	0.0313	0.0445	0.0298	0.0565	0.0766
CB	0.2350	0.2660	0.2692	0.0126	0.0283	0.0434	0.0239	0.0512	0.0747
pure_exploit	0.5797	0.5760	0.6181	0.0312	0.0612	0.0997	0.0592	0.1106	0.1717
pure_explore	0.5651	0.5690	0.5686	0.0304	0.0608	0.0917	0.0577	0.1099	0.1579

C.6 EVALUATION APPROACH

As noted in the main text body, we assess recommendation performance through three evaluation approaches: doubly-robust estimation (Dudík et al. 2011), offline precision evaluated by preference set (Qin et al. 2014), and simulation. In the following, we will explain each of the approaches.

C.6.1 Doubly-Robust Estimation

Doubly-robust estimation is an offline evaluation approach that enables one to assess the value of a policy using a historical dataset (Dudík et al. 2011). This method combines two popular policy evaluation methods, direct simulation (DS) and inverse propensity score (IPS). The former forms an estimate of the expected reward conditioned on the context and action, and the latter forms an

approximation of the action propensity in data collection to correct for the shift in action proportions between the offline dataset and the focal policy to be evaluated. By combining these two methods, the doubly-robust estimator is able to adjust the potential bias caused by the data collection process.

Formally, let $G = \{(\mathbf{x}, a, r_{a,\mathbf{x}})\}$ denote the offline evaluation dataset, with each row containing information about context $\mathbf{x}$, action a , and reward $r_{a,\mathbf{x}}$. Corresponding to our focal context, $\mathbf{x}$ refers to users' weight-management contexts, such as weight variations and weight-management behavior paths; a refers to a certain challenge that the platform provides; and $r_{a,\mathbf{x}}$ is users' challenge-selection decision, which depends on individuals' weight-management contexts and the challenge attributes. As noted in the manuscript, we use user embeddings and challenge embeddings to capture individuals' weight-management contexts and challenge attributes, respectively. As the platform did not promote any personalized recommendation during our data-collection period, we observe the same set of challenges for each individual. The number of challenges we observe each week may be different, as the platform may provide different challenges across time. Figure C.3 provides an illustration of our offline dataset.




Figure C.3 Offline Dataset for Performance Evaluation

In Figure C.3, C_t denotes the set of challenges that are available at time t . The total number of rows in the dataset G is $I \cdot \sum_{t=1}^T |C_t| \triangleq |G|$, in which I is the number of users and T is the number of time periods. The cumulative recommendation performance during the entire periods, which is also known as the “policy value,” can be estimated based on G as follows:

$$\bar{R}_T = \hat{V}_G^\pi = \frac{1}{|G|} \sum_{i,t,a \in C_t} \left\{ \frac{1}{K} \sum_{k \in S_{it}} \left[\hat{\phi}_{v_{itk}} + \frac{(r_{a,\mathbf{x}_{it}} - \hat{\phi}_{v_{ita}}) I_{[k=a]}}{\hat{p}(a|\mathbf{x}_{it})} \right] \right\}, \quad (\text{C.26})$$

where $\hat{\phi}$ is an estimate of the expected reward conditioned on individuals’ decision-making context, which can be obtained through estimating a reward function using the offline data, $\hat{p}(a|\mathbf{x})$ is the probability of observing the platform providing challenge a under individuals’ weight-management context $\mathbf{x}$, and S_{it} is the set of challenges provided to user i at time t according to the recommendation policy.

A high-level rationale behind this method is that, when data are not available, we directly apply $\hat{\phi}$ to predict the reward; otherwise, we use the actual data to adjust the predictor $\hat{\phi}$, taking

into account the propensity of observations in the data. Following the convention, we randomly split users into a training set and a test set, and we use the training set to build the reward predictor $\hat{\phi}$ and use the test set to evaluate recommendation performance. In particular, we adopt a logistic reward predictor: $r_{a,v} = (1 + \exp(-\mathbf{v}^T \boldsymbol{\zeta}))^{-1}$, where $\boldsymbol{\zeta}$ is fitted using the training data only. Because the platform does not prioritize weight-loss challenges for users, the probability of observing each challenge conditioned on context, i.e., $\hat{p}(a | \mathbf{x})$, is uniform across the available challenge set. That is, $\hat{p}(a | \mathbf{v}, a \in C_t) \equiv |C_t|^{-1}$.

For completeness, we define the reward for providing challenge k to user i at time t as

$$r_t(i, k) = \frac{1}{|C_t|} \sum_{a \in C_t} \left[\hat{\phi}_{v_{ik}} + \frac{(r_{a, \mathbf{x}_{it}} - \hat{\phi}_{v_{ia}}) I_{[k=a]}}{\hat{p}(a | \mathbf{x}_{it})} \right], \quad (\text{C.27})$$

and the performance (or reward) for all the recommendations provided to users at time t is

$$\bar{r}_t = \frac{1}{I \cdot |C_t|} \sum_{i, a \in C_t} \left\{ \frac{1}{K} \sum_{k \in S_{it}} \left[\hat{\phi}_{v_{ik}} + \frac{(r_{a, \mathbf{x}_{it}} - \hat{\phi}_{v_{ia}}) I_{[k=a]}}{\hat{p}(a | \mathbf{x}_{it})} \right] \right\}. \quad (\text{C.28})$$

C.6.2 Offline Precision

Following prior recommendation work (Qin et al. 2014; Qin and Zhu 2013), we measure the precision of recommendation by users' offline preference set. A user's preference set is the set of challenges that the user has chosen in the offline data. We use L_i to denote the preference set for user i . The recommendation precision is the overlap ratio between the recommendation set and the preference set. Formally, we define the precision of the recommendation provided to user i at week t as

$$p_{i,t} = \frac{|S_{it} \cap L_i|}{|S_{it}|}. \quad (\text{C.29})$$

The average recommendation precision across all users at week t is $\bar{r}_t = I^{-1} \sum_i p_{i,t}$, and the average recommendation precision across all users up to week t is $\bar{R}_t = (t \cdot I)^{-1} \sum_{s=1}^t \sum_i p_{i,s}$.

C.6.3 Simulation

In light of previous theoretical MAB studies (Hertz et al. 2018; Sani et al. 2012), we also consider a simulation approach to evaluate our model. We use logistic regression to train a predictor for users' binary challenge-selection decisions. Specifically, let $r_t(i, k)$ denote user i 's decision for challenge k at week t , and let $\mathbf{v}$ be a concatenated vector of user embeddings and challenge embeddings. Our logistic predictor can be expressed as $r_t(i, k) = (1 + \exp(-\mathbf{v}_{ik}^T \zeta))^{-1}$. Again, we build our predictor on a training dataset, and the performance evaluation is conducted on a test set. We simulate users' challenge selection decisions based on the estimated choice-making propensity; that is, $\hat{r}_t(i, k) = (1 + \exp(-\mathbf{v}_{ik}^T \hat{\zeta}))^{-1}$. The average challenge-selection rate among users at week t is estimated by

$$\bar{r}_t = \frac{1}{I \cdot K} \sum_i \sum_{k \in S_{it}} \hat{r}_t(i, k), \quad (\text{C.30})$$

and the average challenge-selection rate among users up to week t is $\hat{R}_t = t^{-1} \sum_{s=1}^t \bar{r}_s$.

C.7 BENCHMARK MODELS

C.7.1 Baseline MABs

As noted, we compare our model with a set of baseline MABs to test each of our proposed model components. These baseline MABs are those without user embeddings or challenge embeddings or those without diversity constraint. When user embeddings or challenge embeddings are not incorporated, we use ordinary user features or challenge features to capture the contextual

information, and the MAB algorithm is the same as the one we use in the main analysis. When the diversity constraint is not incorporated, we use an ordinary TS algorithm, which does not contain the step of solving the optimization problem and instead directly selects the challenge choices with the highest simulated rewards.

C.7.2 Benchmark State-of-the-Art Recommendation Systems

Our benchmark state-of-the-art recommendation systems include (1) bandit models, such as UCB and ε -greedy, (2) batch-learning-based models, such as context-aware collaborative filtering (CACF), social collaborative filtering (SCF), probabilistic matrix factorization (PMF), context-aware matrix factorization (CAMF), content-based filtering (CB), and hybrid recommendation systems, (3) pure exploitation and pure exploration. In the following, we introduce the implementation details of these benchmark models.

UCB algorithms incorporate the opportunity for exploration by choosing the alternative with the highest upper confidence bound of the reward (Auer et al. 2002). That is, a UCB chooses an alternative a_t on round t according to $a_t = \arg \max_k (\mu_t(k) + U_t(k))$, where k indexes alternatives and $U_t(k)$ is the upper confidence bound that plays the role of uncertainty bonus. ε -greedy deploys a random exploration strategy. Specifically, it chooses the arm with the seemingly highest average payoff with probability $1 - \varepsilon$ and explores a random arm with probability ε . We present these two algorithms as follows.

Algorithm: UCB

Input: Prior mean m_j and prior variance v_j for parameters $\theta_j, j=1,2,\dots,d$.

Let $\boldsymbol{\theta} = (\theta_1, \dots, \theta_d)$, $\mathbf{m} = (m_1, \dots, m_d)$, $\mathbf{A} = \text{diag}(v_1, \dots, v_d)$.

For $t=1,2,\dots,T$ **do**

For $i=1,2,\dots,I$ **do**

For arm $k \in C_t$ (C_t is the available challenge set at week t) **do**

 Compute $r_t(i,k) = (1 + \exp(-\mathbf{v}_{ik}^T \mathbf{m}))^{-1} + \sqrt{\boldsymbol{\zeta}^T \mathbf{A}^{-1} \boldsymbol{\zeta}}$,

where $\zeta = -\mathbf{v}_{itk} \exp(-\mathbf{m}^T \mathbf{v}_{itk}) / (1 + \exp(-\mathbf{m}^T \mathbf{v}_{itk}))^2$, $\mathbf{v}_{itk}$ is the context vector.

End for

Select the K arms with the highest $r_t(i, k)$ value, denote this set as S_{it} , $|S_{it}| = K$.

Observe a new batch of data $(\mathbf{v}_{itk}, r_t(i, k))$, $i \in [I]$, $k \in S_{it}$.

Update the posterior mean by:

$$\mathbf{m} = \arg \min_{\mathbf{0}} \frac{1}{2} \sum_{j=1}^d \nu_j^{-1} (\theta_j - m_j)^2 + \sum_{i=1}^I \sum_{k \in S_{it}} \log(1 + \exp(-r_t(i, k) \mathbf{\theta}^T \mathbf{v}_{itk})).$$

Update the posterior variance by:

$$\nu_j^{-1} = \nu_j^{-1} + \sum_{i=1}^I \sum_{k \in S_{it}} \nu_{jikt}^2 p_{ikt} (1 - p_{ikt}),$$

where $p_{ikt} = (1 + \exp(-\mathbf{m}^T \mathbf{v}_{itk}))^{-1}$, ν_{jikt} is the j -th element of $\mathbf{v}_{itk}$.

End for

End for

Algorithm: ε -greedy

Input: Prior mean m_j and prior variance ν_j for parameters θ_j , $j = 1, 2, \dots, d$. Random exploration rate ε .

For $t = 1, 2, \dots, T$ **do**

For $i = 1, 2, \dots, I$ **do**

For arm $k \in C_t$ (C_t is the available challenge set at week t) **do**

 Compute $r_t(i, k) = (1 + \exp(-\mathbf{v}_{itk}^T \mathbf{m}))^{-1}$, where $\mathbf{\theta} = (\theta_1, \dots, \theta_d)$, $\mathbf{v}_{itk}$ is the context vector.

End for

 Draw $y \sim \text{Uniform}(0, 1)$

 If $y < \varepsilon$, randomly select K arms; Otherwise, select the K arms with the highest $r_t(i, k)$ value.

 Denote the set of the selected arms as S_{it} , $|S_{it}| = K$.

 Observe a new batch of data $(\mathbf{v}_{itk}, r_t(i, k))$, $i \in [I]$, $k \in S_{it}$.

 Update the posterior mean by:

$$\mathbf{m} = \arg \min_{\mathbf{0}} \frac{1}{2} \sum_{j=1}^d \nu_j^{-1} (\theta_j - m_j)^2 + \sum_{i=1}^I \sum_{k \in S_{it}} \log(1 + \exp(-r_t(i, k) \mathbf{\theta}^T \mathbf{v}_{itk})).$$

 Update the posterior variance by:

$$\nu_j^{-1} = \nu_j^{-1} + \sum_{i=1}^I \sum_{k \in S_{it}} \nu_{jikt}^2 p_{ikt} (1 - p_{ikt}),$$

 where $p_{ikt} = (1 + \exp(-\mathbf{m}^T \mathbf{v}_{itk}))^{-1}$, ν_{jikt} is the j -th element of $\mathbf{v}_{itk}$.

End for

End for

Context-Aware Collaborative Filtering (CACF) incorporates the contexts of users' item selections as weights into a normal collaborative filtering procedure (Chen 2005). Formally, we define R_{itk} as the decision of user i on item k weighted by contexts; that is,

$R_{itk} = \eta_1 \sum_{\tau < t} r_{\tau}(i, k) \cdot \text{sim}(\mathbf{v}_{i\tau k}, \mathbf{v}_{itk})$, where η_1 is a normalizing factor such that the weights sum to unity, and $\text{sim}(\mathbf{v}_{i\tau k}, \mathbf{v}_{itk})$ denotes the similarity between contexts $\mathbf{v}_{i\tau k}$ and $\mathbf{v}_{itk}$. Define R_t^{new} as the weighted challenge selection matrix at time t ; that is, $R_t^{new} = \{R_{itk}\}_{i \in [I], k \in C_t}$. The predicted decision of user i on item k is derived based on the weighted challenge selection data; that is, $\hat{r}_t(i, k) = \bar{r}_i + \eta_2 \sum_{j \neq i} (R_{jtk} - \bar{r}_j) \cdot \text{sim}(R_t^{new}[i, :], R_t^{new}[j, :])$, where $\bar{r}_i$ denotes user i 's overall challenge selection tendency, $\text{sim}(R_t^{new}[i, :], R_t^{new}[j, :])$ denotes the similarity between user i and user j 's challenge selection histories, and η_2 is a normalizing factor.

Social Collaborative Filtering (SCF) formulates a neighborhood-based method for cold-start collaborative filtering in a generalized matrix algebra framework (Sedhain et al. 2014). Specifically, let $X_{I \times d}$ be users' context matrix, with each row representing a set of context features (e.g., gender, age, social activities, etc.) of user i . Let $R_{I \times L}$ be users' challenge selection matrix. I is the number of users, L is the total number of challenges, and d is the feature dimension. SCF calibrates the pseudo similarity between users' context features and challenge selection histories; that is, $w_{jk} = \text{sim}(X[:, j], R[:, k])$. The method then utilizes the pseudo similarity matrix W to make predictions on users' challenge selection decisions; that is, $R_{ik} \propto \text{sim}(X[i, :], W[:, k])$, where R_{ik} is user i 's decision on item k , $W = \{w_{jk}\}_{j \in [d], k \in [L]}$. For both CACF and SCF, we use cosine similarity to measure the affinity between two vectors.

Probabilistic Matrix Factorization (PMF) is a model-based collaborative filtering method that assumes users' preferences to be determined by a set of latent factors. It models the probabilistic distribution of users' preferences based on a latent user matrix U and a latent item matrix V . We define the conditional distribution of users' challenge selection tendency as

$p(R|U, V, \sigma^2) = \prod_{i \in [I]} \prod_{k \in [L]} N(R_{ik} | U_i^T V_k, \sigma^2)$, where U_i and V_k denote user-specific and item-specific latent feature vectors, respectively. The priors of user and item latent features are assumed to be normally distributed; that is, $p(U | \sigma_U^2) = \prod_i N(U_i | 0, \sigma_U^2 I)$, $p(V | \sigma_V^2) = \prod_k N(V_k | 0, \sigma_V^2 I)$. The objective of PMF is to estimate U and V such that the posterior of R is maximized, or, equivalently, the following sum-of-squared-errors function with quadratic regularizers is minimized: $F = \sum_{i \in [I]} \sum_{k \in [L]} (R_{ik} - U_i^T V_k)^2 + \lambda_U \sum_{i \in [I]} \|U_i\|_2^2 + \lambda_V \sum_{k \in [L]} \|V_k\|_2^2$, where $\lambda_U = \sigma^2 / \sigma_U^2$, $\lambda_V = \sigma^2 / \sigma_V^2$. User i 's decision on item k can then be predicted by $\hat{R}_{ik} \sim N(\hat{U}_i^T \hat{V}_k, \hat{\sigma}^2)$.

Context-Aware Matrix factorization (CAMF) is an extension of the classic matrix factorization approach for incorporating contextual information (Baltrunas et al. 2011). Specifically, it minimizes the following objective function:

$$\min_{U, V, B, b} \sum_{i, k} [(R_{ik} - U_i^T V_k - \bar{r}_k - b_i - \sum_j B_{jc_j})^2 + \lambda(b_i^2 + \|U_i\|_2^2 + \|V_k\|_2^2 + \sum_j \sum_{c_j} B_{jc_j}^2)], \quad (C.31)$$

where R_{ik} denotes user i 's decision on item k ; U_i is the latent user feature vector, and V_k is the latent item feature vector; $\bar{r}_k$ is the average selection rate of item k in the data; b_i is the baseline parameter for user i ; B_{jc_j} are parameters capturing the effects of contextual conditions, in which j indexes the contextual factors and c_j indexes the possible values of contextual factor j ; λ is a regularization factor. Users' challenge-selection decisions can be predicted by $\hat{R}_{ik} = \hat{U}_i^T \hat{V}_k + \bar{r}_k + \hat{b}_i + \sum_j \hat{B}_{jc_j}$.

Our content-based filtering method is similar to the Prod2Vec or Item2Vec method in the prior studies (Barkan and Koenigstein 2016; Vasile et al. 2016). We use challenge embeddings to evaluate the similarity between two challenges, and we recommend the top K challenges that are most similar to users' challenge-selection profile. Again, our similarity measure is cosine. Our

hybrid models combine content-based filtering and collaborative filtering by a mixed hybridization approach (Burke 2002). In particular, the first hybrid model combines content-based filtering with a pure collaborative filtering model, which entirely leverages users' challenge selection histories. Our second hybrid model combines content-based filtering with CACF, which further incorporates contextual information in collaborative filtering. For the above described batch-learning-based models, we use 4-week data to train the predictors. As the recommendation goes on, additional user profile data can be acquired and the predictors can be updated.

Pure exploitation is a Bayesian learning model that selects the choice with the highest posterior payoff expectation. In stark contrast to pure exploitation, pure exploration does not form any posterior and fully randomizes its choice. We present the details of these two algorithms as follows.

Algorithm: Pure Exploitation

Input: Prior mean m_j and prior variance v_j for parameters $\theta_j, j=1,2,\dots,d$.

Let $\boldsymbol{\theta}=(\theta_1,\dots,\theta_d)$, $\mathbf{m}=(m_1,\dots,m_d)$, $\mathbf{A}=\text{diag}(v_1,\dots,v_d)$.

For $t=1,2,\dots,T$ **do**

For $i=1,2,\dots,I$ **do**

For arm $k \in C_t$ (C_t is the available challenge set at week t) **do**

 Compute $r_t(i,k)=(1+\exp(-\mathbf{v}_{ik}^T \mathbf{m}))^{-1}$, where $\mathbf{v}_{ik}$ is the context vector.

End for

 Select the K arms with the highest $r_t(i,k)$ value, denote this set as S_{it} , $|S_{it}|=K$.

 Observe a new batch of data $(\mathbf{v}_{ik}, r_t(i,k)), i \in [I], k \in S_{it}$.

 Update the posterior mean by:

$$\mathbf{m} = \arg \min_{\boldsymbol{\theta}} \frac{1}{2} \sum_{j=1}^d v_j^{-1} (\theta_j - m_j)^2 + \sum_{i=1}^I \sum_{k \in S_{it}} \log(1 + \exp(-r_t(i,k) \boldsymbol{\theta}^T \mathbf{v}_{ik})) .$$

 Update the posterior variance by:

$$v_j^{-1} = v_j^{-1} + \sum_{i=1}^I \sum_{k \in S_{it}} v_{jikt}^2 p_{ikt} (1 - p_{ikt}) ,$$

 where $p_{ikt} = (1 + \exp(-\mathbf{m}^T \mathbf{v}_{ik}))^{-1}$, v_{jikt} is the j -th element of $\mathbf{v}_{ik}$.

End for

End for

Algorithm: Pure Exploration

For $t=1,2,\dots,T$ **do**

```

For  $i=1,2,\dots,I$  do
  Randomly select  $K$  arms from choice set  $C_t$ . Denote the selected arm set as  $S_{it}$ ,  $|S_{it}| = K$ .
End for
End for

```

C.8 SUPPLEMENTARY RESULTS OF EXPERIMENT 4

As noted, in experiment 4, we evaluate recommendation performance on a subset of test users whose choices of challenge types tend to change frequently, and we compare the new recommendation results with the original results evaluated on the full test set. In the main paper, we provide the results measured by doubly-robust estimation. For completeness, in this section, we present the results measured by offline precision and simulation. As can be seen in Figure C.4, the performance variation patterns remain consistent with the ones shown in the main paper. That is, the bandit-based models tend to have a performance increase, whereas the batch-learning-based models generally have a performance decrease. These results indicate that the performance variation patterns are robust under different evaluation approaches, which again amplifies the capability of online learning frameworks to adapt to users' changing preferences.




Figure C.4 Performance Variation for Dynamic Users

References

- Auer P, Cesa-Bianchi N, and Fischer P. (2002) Finite-Time Analysis of the Multiarmed Bandit Problem. *Machine learning* 47(2-3): 235-256.
- Baltrunas L, Ludwig B, and Ricci F (2011) Matrix Factorization Techniques for Context Aware Recommendation. *Proceedings of the fifth ACM conference on Recommender systems*, 301-304.
- Barkan O, and Koenigstein N (2016) Item2vec: Neural Item Embedding for Collaborative Filtering. *2016 IEEE 26th International Workshop on Machine Learning for Signal Processing (MLSP)* (IEEE), 1-6.
- Burke R. (2002) Hybrid Recommender Systems: Survey and Experiments. *User modeling and user-adapted interaction* 12(4): 331-370.
- Chen A (2005) Context-Aware Collaborative Filtering System: Predicting the User's Preference in the Ubiquitous Computing Environment. *International Symposium on Location-and Context-Awareness* (Springer), 244-253.
- Dudík M, Langford J, and Li L. (2011) Doubly Robust Policy Evaluation and Learning. *arXiv preprint arXiv:1103.4601*.
- Hertz U, Bahrami B, and Keramati M. (2018) Stochastic Satisficing Account of Confidence in Uncertain Value-Based Decisions. *PloS one* 13(4): e0195399.
- Johnson P E, Veazie P J, Kochevar L, O'connor P J, Potthoff S J, Verma D, and Dutta P. (2002) Understanding Variation in Chronic Disease Outcomes. *Health Care Management Science* 5(3): 175-189.
- Qin L, Chen S, and Zhu X (2014) Contextual Combinatorial Bandit and Its Application on Diversified Online Recommendation. *Proceedings of the 2014 SIAM International Conference on Data Mining* (SIAM), 461-469.
- Qin L, and Zhu X (2013) Promoting Diversity in Recommendation by Entropy Regularizer. *Twenty-Third International Joint Conference on Artificial Intelligence*.

Sani A, Lazaric A, and Munos R (2012) Risk-Aversion in Multi-Armed Bandits. *Advances in Neural Information Processing Systems*), 3275-3283.

Sedhain S, Sanner S, Braziunas D, Xie L, and Christensen J (2014) Social Collaborative Filtering for Cold-Start Recommendations. *Proceedings of the 8th ACM Conference on Recommender systems*), 345-348.

Vasile F, Smirnova E, and Conneau A (2016) Meta-Prod2vec: Product Embeddings Using Side-Information for Recommendation. *Proceedings of the 10th ACM Conference on Recommender Systems*), 225-232.

VITA

Tingxin Zhou is a researcher in Information Systems. Her research interests broadly concern the transitional and transformational impact of information technology (IT) on business forms and individual behaviors. Her current research mostly focuses on the integration of IT and healthcare. She leverages data-driven techniques such as machine learning and econometrics to obtain insights regarding individuals' health-behavior patterns, IT-enabled care delivery, patient empowerment, and operating mechanism of health IT artifacts. She receives her Ph.D. degree at the Michael G. Foster School of Business, University of Washington. Prior to joining Foster, she received her Bachelor's degree in School of Statistics at Shanghai University of Finance and Economics (SHUFE). As of August 2021, she will be joining W. P. Carey School of Business at Arizona State University as an Assistant Professor.