

Original Paper

New Approach to Equitable Intervention Planning to Improve Engagement and Outcomes in a Digital Health Program: Simulation Study

Jackson A Killian^{1,2,3}, PhD; Manish Jain³, PhD; Yugang Jia², MPH, PhD; Jonathan Amar², PhD; Erich Huang², MD, PhD; Milind Tambe¹, PhD

¹Harvard University, Cambridge, MA, United States

²Verily Life Sciences, South San Francisco, CA, United States

³Google Research, Palo Alto, CA, United States

Corresponding Author:

Yugang Jia, MPH, PhD
Verily Life Sciences
269 East Grand Avenue
South San Francisco, CA, 94080
United States
Phone: 1 9143744981
Email: yugang@verily.com

Abstract

Background: Digital health programs provide individualized support to patients with chronic diseases and their effectiveness is measured by the extent to which patients achieve target individual clinical outcomes and the program's ability to sustain patient engagement. However, patient dropout and inequitable intervention delivery strategies, which may unintentionally penalize certain patient subgroups, represent challenges to maximizing effectiveness. Therefore, methodologies that optimize the balance between success factors (achievement of target clinical outcomes and sustained engagement) equitably would be desirable, particularly when there are resource constraints.

Objective: Our objectives were to propose a model for digital health program resource management that accounts jointly for the interaction between individual clinical outcomes and patient engagement, ensures equitable allocation as well as allows for capacity planning, and conducts extensive simulations using publicly available data on type 2 diabetes, a chronic disease.

Methods: We propose a restless multiarmed bandit (RMAB) model to plan interventions that jointly optimize long-term engagement and individual clinical outcomes (in this case measured as the achievement of target healthy glucose levels). To mitigate the tendency of RMAB to achieve good aggregate performance by exacerbating disparities between groups, we propose new equitable objectives for RMAB and apply bilevel optimization algorithms to solve them. We formulated a model for the joint evolution of patient engagement and individual clinical outcome trajectory to capture the key dynamics of interest in digital chronic disease management programs.

Results: In simulation exercises, our optimized intervention policies lead to up to 10% more patients reaching healthy glucose levels after 12 months, with a 10% reduction in dropout compared to standard-of-care baselines. Further, our new equitable policies reduce the mean absolute difference of engagement and health outcomes across 6 demographic groups by up to 85% compared to the state-of-the-art.

Conclusions: Planning digital health interventions with individual clinical outcome objectives and long-term engagement dynamics as considerations can be both feasible and effective. We propose using an RMAB sequential decision-making framework, which may offer additional capabilities in capacity planning as well. The integration of an equitable RMAB algorithm further enhances the potential for reaching equitable solutions. This approach provides program designers with the flexibility to switch between different priorities and balance trade-offs across various objectives according to their preferences.

(*JMIR Diabetes* 2024;9:e52688) doi: [10.2196/52688](https://doi.org/10.2196/52688)

KEYWORDS

chronic disease; type-2 diabetes; T2D; restless multiarmed bandits; multi-armed bandit; multi-armed bandits; machine learning; resource allocation; digital health; equity

Introduction

Chronic diseases, while obviously heterogeneous in their physiology, pose a series of common management challenges. One of them is that, by the very nature of these conditions, interventions have to impact multiple aspects of the patient's daily living to be effective. This scenario is propitious for the implementation of digital health programs (via wearables, mobile apps, or virtual care), such as vida (Vida) and welldoc (Welldoc), that provide patient-centric support between in-clinic visits. These digital health programs may lead to improved clinical outcomes [1-3].

The success of digital health programs, however, hinges on the dynamic balance of several factors. The ultimate metric of success of any program is always the improvement of the individual health outcomes of participants in the program. However, these programs need to sustain participant engagement to be effective [4]. The importance of patients engaging with specific intervention points is clear since only the interventions that patients receive can have an effect. However, sustained engagement over time is a critical success factor in itself, as it can mediate enduring and (potentially) disease-modifying long-term shifts in patients' attitudes and perceptions about the management of their own health and lifestyle [5]. Yet, attrition and dropout across programs are estimated to be as high as ~50% [6], representing a major barrier to optimal effectiveness. Moreover, these programs may have capacity limitations (eg, the volume of coaches or health counselors) and need to allocate intervention resources proactively. The options for resource management in digital health programs can vary widely, depending on the metrics and time horizon on which success is measured. In that context, it can be challenging to estimate which approach will be the most effective for a given set of goals.

Our focus is on type 2 diabetes (T2D), which is a representative chronic disease condition. T2D is a high-prevalence, high-burden disease. In the United States, 30 million people are estimated to live with diabetes, the 8th leading cause of mortality [7], and it is estimated to account for over US \$300 billion of economic cost [7-9]. The physiologic hallmark of the disease is elevated blood glucose, and success in clinical management is monitored by testing the levels of hemoglobin A_{1c} (HbA_{1c}) tests [10]. T2D can lead to organ damage, but it is manageable through medication and lifestyle changes. Our work is based on a digital program that supports patients through a mobile app, virtual coaching (web-based and app-based), and integration of sensor-collected information.

Our digital program of interest contacts patients to maintain engagement and direct and support specific patient actions. Resource investment into those outreach interventions often relies on an intuitive strategy guided by a present clinical state (eg, in this case, giving preference to patients with the highest HbA_{1c}). However, such engagement-agnostic strategies may

not lead to the best possible health outcomes at the population level, since reactive strategies that only prioritize immediate clinical improvements may do so at the expense of future engagement, reducing the ability to deliver interventions to patients who have dropped out.

Digital programs that consider the joint dynamics of engagement and clinical status may arrive at better determinations about intervention strategies. This is a problem that entails long-term planning usually in resource-constrained settings, therefore it can naturally be cast as a restless multiarmed bandit (RMAB) framework, of the type used for studying resource allocation in the context of stochastic scheduling problems [11]. Recent examples of applying RMABs to health-related problems include computing optimal cancer screening regimens [12], improving maternal health through telehealth [13], and planning hepatitis-C treatment delivery [14].

RMAB frameworks generate sequential resource allocation strategies in pursuit of desired outcomes (in our case it would be optimal health status and engagement) but may be prone to maximizing system-level rewards by sacrificing certain groups to favor the "most promising" ones, hence leading to inequities [15]. In a disease-management context, these (potentially) inequitable policies would translate into disparate outcomes across demographic groups, potentially exacerbating existing systemic inequities in health care [15]. To mitigate this issue, there have been recent studies of fairness in RMAB, in the sense of generating resource allocation strategies with a degree of distributive fairness, where all arms have an opportunity to receive the intervention of interest (in this case resources). Specifically, some works view fairness from the lens of equality, guaranteeing a lower bound of receiving an intervention for all groups [16,17]. Fairness has also been set by modulating risk sensitivity, encoding risk-averseness or risk-prevalence levels to shape the reward functions [18].

In this work, we aimed to develop a resource allocation strategy for a digital health app to support patients with T2D applying an RMAB framework. We intentionally sought to incorporate equity as a desirable feature of our approach, aiming to leverage recent innovations in health care, such as the emergence of digital health, without perpetuating systemic flaws in care delivery, such as societal inequities [15]; moreover, T2D represents an unfortunate example where the presence of systemic inequities continues to have a negative impact in care [19]. We introduce a new solution, equitable RMAB (ERMAB), which requires that allocation policies take affirmative steps to distribute resources in a way that equalizes outcomes across prespecified groups. That is, we focus on fairness through the lens of achieving equitable outcomes in resource allocation. We applied this paradigm to the resource allocation of outreach interventions in our program, evaluating an engagement-health dynamics model and an equitable intervention planning approach via an extensive simulation study using publicly available statistics about digital T2D management. Subsequently, we

carried out a Pareto analysis to further study the interplay of engagement-clinical outcome dynamics under different intervention strategies, and perform sensitivity analyses to demonstrate our framework's robustness across RMAB parameter settings.

Methods

Model

Overview

Our model needed to simultaneously address the following facets essential to digital health programs: (1) evolution of clinical outcomes per patient, (2) joint engagement-health dynamics per patient, (3) limited observability of clinical outcomes, and (4) limited resource availability.

We model the problem as a restless bandit with $n \in 1, \dots, N$ arms representing each patient, discrete per-arm state space S_n , per-arm action space $A_n = \{\text{User self-care, Intervention}\}$ (equivalently $\{U, I\}$), per-arm transition functions P_n defining the probability of arm n transitioning from state s to state s' given action a , per-arm reward function $R_n(s)$ defining the reward for an arm being in state s , time horizon H , and action budget B . For ease of exposition, S_n , A_n , and $R_n(s)$ are the same for all arms, so we drop the subscript n from these, but our methods apply to the general setting where arms have different state, action, and reward functions. Let s^t be the N -length vector of arm states at time t , indexed as s^t_n , and let a^t be an N -length 1-hot encoding of the arms that receive interventions from the program in time period t . The planner must take actions to maximize their objective, subject to a per-round budget constraint, $|a^t|_1 \leq B \forall t \in 1, \dots, H$.

To capture the joint dynamics of engagement and health in digital health programs, we included a dimension for each factor

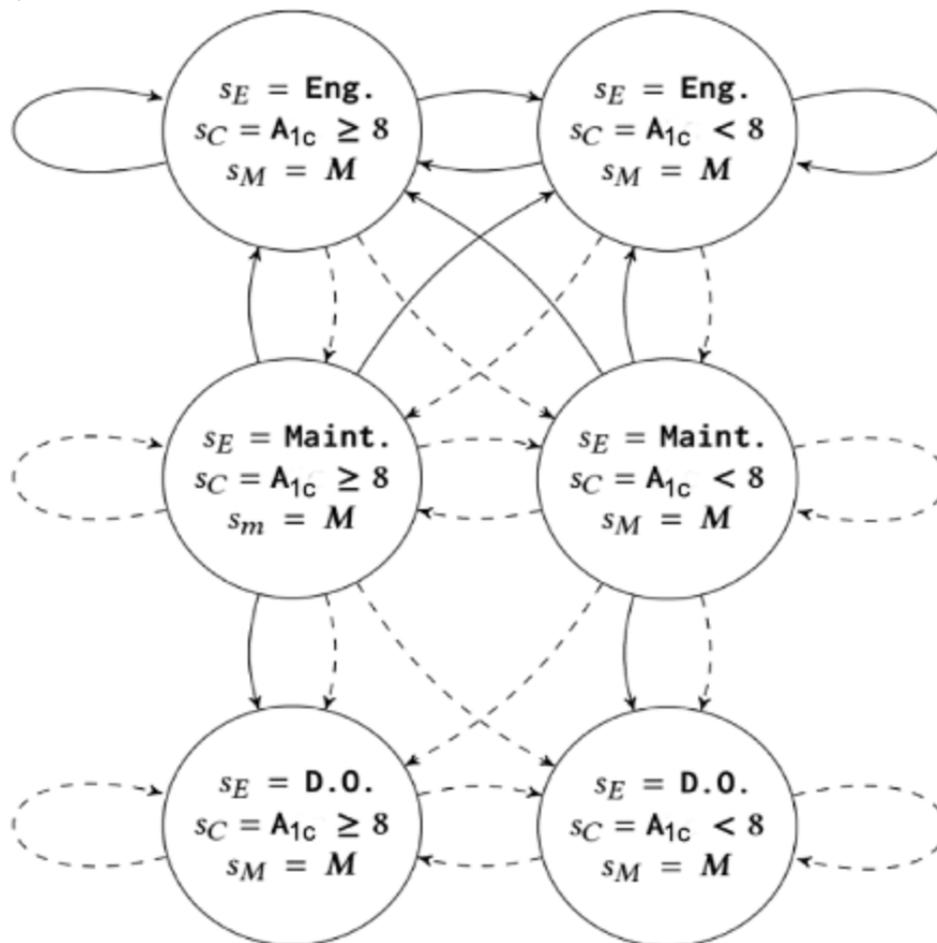
in our state space. For the T2D domain, we also include a dimension for memory, since intervention effects have a delayed impact on clinical outcomes. We represent this 3D state space S by a 3-tuple $(s_E, s_C, \text{ and } s_M)$, where s_E captures the arm's engagement, s_C captures the arm's clinical (ie, health) state, and s_M is a 2-length memory vector. All dimensions of the state space are modeled as discrete, where continuous spaces are discretized via threshold rules, described next.

The engagement dimension, s_E , has 3 states: $\{\text{Engaged, Maintenance, and Dropout}\}$. A patient is *Engaged* if they received an intervention from the care team and they responded to the team within the app in the current time period. A patient is in the *Maintenance* state if they have produced any interactions within the app, but did not respond to an intervention if it was attempted in the current time period. A patient is in the *Dropout* state if they have not produced any interactions in the app in the current time period and will no longer do so in any future time period (eg, they have deleted the app). These states are chosen to capture the primary high-level engagement dynamics seen in our digital program.

The clinical dimension, s_C , captures a user's HbA_{1c} value (via 2 states: $\{\text{HbA}_{1c} < 8, \text{HbA}_{1c} \geq 8\}$). This threshold was chosen to model the clinical outcome target for app users in publicly available data, that is, reducing their HbA_{1c} below 8. Finally, the memory dimension, s_M , is a 2-length vector for recording previous values of s_E , so its entries can take the same values as the s_E dimension. The memory serves to implement a 3-month delay between an intervention and its impact on the clinical state. This effect is observed in data and is due to the biological nature of HbA_{1c} progression, that is, it is a summary measure of the body's blood sugar over the previous 3 months. Let s_{M_i} reference the i th entry of the 0-indexed, 2-length memory vector.

Transition dynamics are summarized below (Figures 1 and 2).

Figure 1. State transition diagram for 1 arm. Bold arrows are transitions when a = intervention and dotted arrows represent transitions when a = user self-care. Eng: engaged; Maint: maintenance.



Engagement Dynamics

The engagement model is made up of 4 main effects. First, each patient has their own independent probability of responding to an intervention and transitioning to the *Engaged* state from either the *Engaged* or *Maintenance* states. Second, the probability of a patient responding to an intervention if they were previously in the *Engaged* state is higher than if they were previously in the *Maintenance* state. Third, the probability of a patient transitioning to a *Dropout* state is lower if the patient receives an intervention, than if they do not. Lastly, patients in the *Dropout* state will never respond to an intervention. In summary, this corresponds to 4 open parameters for the engagement dynamics, p^I_{MtoE} , p^I_{EtoE} , p^I_{MtoD} , and p^U_{MtoD} , where superscripts, *I* or *U*, denote the action.

Clinical Dynamics

There are 2 meaningful clinical dynamics, corresponding to the clinical evolution of patients who did and did not respond to an

intervention. Specifically, we assume that patients who received and responded to an intervention (ie, were in the *Engaged* state) will have a higher probability of transitioning to a healthy clinical state than a patient who did not receive or respond to an intervention. In addition, all effects are delayed by 3 months via the memory states as described in the equations below (Figure 2). Note that we assume that HbA_{1c} progression is the same for users who were in the *Maintenance* and *Dropout* states. We show the evolution of the clinical state s'_C , given the memory state s_{Ml} (ie, clinical state 3 months ago), and the current clinical state s_C , in Table 1. Row 1 of Table 1 represents users who received and responded to an intervention 3 months ago, whereas row 2 represents users who did not receive or respond to an intervention 3 months ago. Note that this requires estimating only 4 parameters for clinical progression, that is, $p^E_{A_{1c} \geq 8}$, $p^E_{A_{1c} < 8}$, $p^E_{A_{1c} \geq 8}$, and $p^E_{A_{1c} < 8}$, all of which encode the probability of having an HbA_{1c} level less than 8 in 3 months.

Figure 2. Construction of the delayed intervention effect on clinical state, s_C , via zoomed in view of Figure 1. Each transition (arrow) in Figure 1 encodes 2 transitions with different probabilities (the dashed and dotted arrows in this figure), each of which depend on the engagement state of the user 3 months ago, that is, the last entry of the memory state M . Specifically, the probability of transitioning to a better clinical state will be larger if the user was in the engaged state 3 months ago. Eng: engaged.

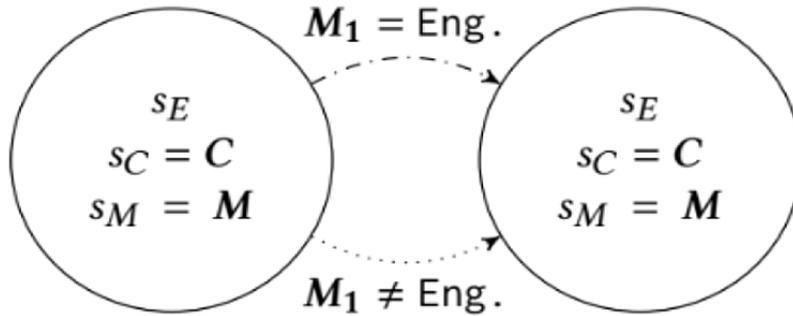


Table 1. The table shows the evolution of clinical state $P(s'_C = \text{HbA}_{1c} < 8 | r, c)$ where r represents the memory state s_{M1} and c represents the current clinical state s_C .

Evolution of clinical state $P(s'_C = \text{HbA}_{1c} < 8 r, c)$	$s'_C = \text{HbA}_{1c} \geq 8$	$s'_C = \text{HbA}_{1c} < 8$
$s_{M1} = \text{Eng}^a$	$P^E_{A1c \geq 8}$	$P^E_{A1c < 8}$
$s_{M1} \neq \text{Eng}$	$P^{\bar{E}}_{A1c \geq 8}$	$P^{\bar{E}}_{A1c < 8}$

^aEng: Engaged.

Memory Dynamics

The memory dimension is a sliding window to record the engagement state of the previous 3 months:

$$P(s'_M0 = s_E, s'_M1 = s_M0 | s_E, s_M0) = 1$$

Finally, note that the arrows in Figures 1 and 2 represent joint engagement-clinical-memory transition probabilities. These are obtained by multiplying the engagement, clinical, and memory transition rules.

Observability

By definition, the engagement state s_E , and thus memory state s_M , are fully observable. However, the clinical state s_C relies on a patient collecting a measurement of their HbA_{1c} in a given time period. We assume that users in the *Engaged* state have fully observable s_C , for example, they will measure their HbA_{1c} upon request from the program, patients in the *Maintenance* state have a partially observable HbA_{1c} , for example, they will measure their HbA_{1c} in a given round with probability q^{Obs}_{Maint} , and users in the *Dropout* state have an unobservable HbA_{1c} . To handle this partial observability in a computationally scalable way, we convert the partially observable system via a belief-state conversion which allows us to treat the converted system as fully observable [20]. The main benefit is that it allows us to use more efficient optimization tools, at the cost of having a slightly larger state space in the converted system.

Rewards

We assign rewards based on the current state of each patient and represent them as $R(s)$. In general, our objective is to jointly boost engagement and clinical state. To capture that objective, we define rewards for each state dimension independently as:

$$r_E(s_E = \text{D.O.}) = 0, r_E(s_E = \text{Maint}) = 0.5, r_E(s_E = \text{Eng}) = 1 \quad (1)$$

$$r_C(s_C = \text{A1c} \geq 8) = 0, r_C(s_C = \text{A1c} < 8) = 1 \quad (2)$$

The reward for a patient's full state is then computed as $R([s_E, s_C, s_M]) = \alpha r_E(s_E) + (1 - \alpha) r_C(s_C)$.

Thus the parameter α represents the relative weight on the engagement reward and it can be tuned based on the planner's desired objective.

Equitable Restless Bandit Problem

Overview

We model the problem as an RMAB, a framework for finding optimal allocations of constrained resources across many Markov Decision Processes and across time. In this work, we enforce that solutions must also be equitable across groups of arms, introducing a new class of ERMAB. Here, we give a brief overview of the ERMAB framework and the equitable objectives considered for our simulation analysis. For full technical background on restless bandits and full derivations of ERMABs and their solutions, please see Killian et al [21].

Preliminaries

We consider predefined groups of arms (patients) G , indexed by g . Let $M-1(g)$ be the set of arms in group g . Given a time horizon H , a start state s^0_g , and per-round budget b_g , a reward-maximizing allocation policy for a group of arms can be found by computing the value function $V^0_g(s^0_g, b_g)$, where:

$$V_g(s^t_g, b_g) = \max_{a^t} \{ \sum_{n=t}^H R(s^n) + E[V^{t+1}(s^{t+1}) | s^t, a^t] \} \quad (3)$$

$$\text{s.t. } \sum_{n \in g} a^k_n \leq b_g \quad \forall k \in t, \dots, H \quad (4)$$

and $V_g^H(\cdot) = 0$. However, solving this exactly is PSPACE-Hard [22], due to the coupling between arms imposed by the budget constraint. Thus, it is more common to work with objectives that relax the budget constraint equation 4 in a Lagrangian fashion, trading some solution quality for computational tractability. Solutions to the relaxed value functions are denoted $L_g^t(s_g^t, b_g)$, rather than $V_g^t(s_g^t, b_g)$.

Equitable Objectives

In ERMABs, our objective is to both maximize reward and ensure that rewards are distributed equitably across groups of arms. Below, we give 2 objectives for planning such policies.

Maximin Reward

Maximin reward (MMR) is a robust objective that maximizes the minimum prospective total reward of any group.

$$\begin{aligned} & \max_{b_g, r^*} r^* \\ & \text{s.t. } L_g^0(s_g^0, b_g) \geq r^* \quad \forall g \in G \quad (5) \\ & \sum_{g \in G} b_g^t = B \quad \forall t \in [0, \dots, H-1] \end{aligned}$$

where B is the total per-round budget constraint over all groups. This objective takes a bottom-up approach to equity, ensuring that the groups that are the worst-off are prioritized for resources. However, since the objective focuses only on maximizing the worst case, on some data distributions, it may over-commit resources to a subset of groups with very low potential for improved outcomes, at the expense of potential gains to other groups, which may be undesirable. To account for this, we also consider a second equitable objective that is

sensitive to gains across the distribution of groups, while still prioritizing the worst-off.

Maximum Nash Welfare

$$\begin{aligned} & \text{Max}_{b_g} \prod_{g \in G} L_g^0(s_g^0, b_g) \\ & \sum_{g \in G} b_g^t = B \quad \forall t \in [0, \dots, H-1] \quad (6) \end{aligned}$$

The maximum Nash welfare (MNW) objective gives diminishing returns as the prospective total reward of a group becomes larger. This leads to prioritizing allocations that improve the rewards of all groups more equitably. However, if 1 or a subset of groups have little potential for gains, the allocations will go to the next-worst-off groups which may see some meaningful utility increase from the allocation.

Both objectives represent a natural bilevel optimization problem, where the inner problem solves for the value function *within* 1 group, and the outer problem solves for the equitable distribution of resources *across* groups. To solve equation 5, we use algorithm 1 (Figure 3 [21]), an efficient water filling procedure that incrementally assigns a budget to the group with the smallest long-term value L , until the total budget B is exhausted. To solve equation 6, we use algorithm 2 (Figure 3 [21]), an efficient greedy approach that incrementally assigns a budget to the group that will see the largest marginal (log) increase in its long-term value L , until the total budget B is exhausted. The algorithm also includes nuance which corrects for computational biases that occur in the presence of unequal group sizes. To take actions (assign resources) in the simulations, we follow the actions implied by the value functions $L_g^t(s_g^t, b_g)$ output by algorithms 1 or 2 (for complete algorithm derivation, with additional proofs and technical detail, see Killian et al [21]).

Figure 3. Algorithms. ERMAB: equitable restless multiarmed bandit.**Algorithm 1:** ERMAB water filling: maximin reward

Data: $\mathcal{G}, B, \mathbf{s}, h$
 $\mathbf{b} = 0$ // $|\mathcal{G}|$ -length vector, of budgets
for $g \in \mathcal{G}$ **do** // Initialize
 $L(\mathbf{s}_g, b_g) = \text{INNEROPT}(g, b_g, \mathbf{s}_g, h)$
 $/|M^{-1}(g)|$
for $b \in [1, \dots, B]$ **do**
 $g^* = \text{ARGMIN}(L(\mathbf{s}_g, b_g))$
 $b_{g^*} += 1$
 $L(\mathbf{s}_{g^*}, b_{g^*}) = \text{INNEROPT}(g^*, b_{g^*}, \mathbf{s}_{g^*}, h)$
 $/|M^{-1}(g)|$
return L, \mathbf{b}

Algorithm 2: ERMAB greedy: max Nash welfare

Data: $\mathcal{G}, B, \mathbf{s}, h$
 $\mathbf{b} = 0$ // $|\mathcal{G}|$ -length vector, of budgets
 $\theta = \max_g \{|M^{-1}(g)|\}$
for $g \in \mathcal{G}$ **do** // Initialize
 $\text{UPSAMPLE}(g, \theta)$ // Resample arms,
 until g has size θ
 $L_0(\mathbf{s}_g, b_g) = \text{INNEROPT}(g, b_g, \mathbf{s}_g, h)$
 $L_1(\mathbf{s}_g, b_g) = \text{INNEROPT}(g, b_g + 1, \mathbf{s}_g, h)$
 $L_\Delta(\mathbf{s}_g, b_g) = \log(L_1(\mathbf{s}_g, b_g)) - \log(L_0(\mathbf{s}_g, b_g))$
for $b \in [1, \dots, B]$ **do**
 $g^* = \text{ARGMAX}(L_\Delta(\mathbf{s}_g, b_g))$
 $b_{g^*} += 1$
 $L_0(\mathbf{s}_{g^*}, b_{g^*}) = L_1(\mathbf{s}_{g^*}, b_{g^*})$
 $L_1(\mathbf{s}_{g^*}, b_{g^*}) = \text{INNEROPT}(g^*, b_{g^*} + 1, \mathbf{s}_{g^*}, h)$
 $L_\Delta(\mathbf{s}_{g^*}, b_{g^*}) =$
 $\log(L_1(\mathbf{s}_{g^*}, b_{g^*})) - \log(L_0(\mathbf{s}_{g^*}, b_{g^*}))$
 $\text{RESCALE}(\mathbf{b}, \mathcal{G}, \theta)$ // Rescale budgets
 proportional to original group size
return L, \mathbf{b}

Simulation

MarketScan Datasource

To derive baseline statistics on clinical evolution, we relied on the widely used Truven Health MarketScan Commercial Database [23], a convenience sample of medical insurance claims from patients who are privately insured in the United States over the years 2018 to 2020, which includes measurements of HbA_{1c}. We consider users enrolled for more than 6 months that have T2D only, that is, excluding those with hypertension, depression, heart failure, or cancer. We then group users by age, gender, and starting HbA_{1c} to derive statistics per group on monthly HbA_{1c} change (full details in [Multimedia Appendix 1](#)). These provide values of $p^{\text{E}}_{A_{1c} \geq 8}$ and $p^{\text{E}}_{A_{1c} < 8}$ of approximately 7.5% and 0.5%, respectively, with about 1% variation across groups. The MarketScan data set is publicly

accessible and provides a reasonable estimate for the background rate of HbA_{1c} change for users not in a specific digital health program, but receiving standard care. It provides a conservative baseline for our experiments.

For the engagement dynamics, statistics on monthly dropout rates by demographic groups from digital health programs are not readily available. Therefore, we use age and gender-based monthly dropout statistics published by the National Diabetes Prevention Program (NDPP) lifestyle change program, primarily made up of in-person meetings [24]. With monthly dropout rates near 10%, this again forms a reasonable conservative baseline for experiments, serving as a proxy for patients' willingness to engage with T2D-related ongoing behavior change coaching. These statistics populate $p^{\text{U}}_{\text{MtoD}}$ in our model, with about 4% variation between groups.

The remaining parameters require estimates from digital health program data which are not readily available publicly. Thus we make the following assumptions to instantiate their values. For $p_{A_{1c} \geq 8}^E$ and $p_{A_{1c} < 8}^E$, that is, the clinical probabilities of patients who received and responded to intervention, the patients in age ranges of aged 30-44, 45-54, and 55-64 years receive 25%, 50%, and 75% boost in their clinical probability of transitioning to $HbA_{1c} < 8$, respectively. We found that this leads to clinical trajectories in line with 1 published observational study of a digital diabetes management program [25], and included age-based variation to align with variation observed in NDPP's monthly dropout statistics. For p_{EtoE}^I and p_{MtoD}^I , we assign values of 99% and 3%, respectively, encoding an assumption that patients are more likely to stay in the program if intervened or if already engaged. For p_{EtoE}^I , we assign values with a mean of 75%, but with the same group variation as was present in the data for NDPP's dropout statistics.

Finally, we set the probability of observing the clinical state of a patient in the maintenance state, that is, q_{Maint}^{Obs} to 30%, in line with statistics from MarketScan.

MMR Counterexample Data

Since MMR objectives are prone to “getting stuck” on unmovable targets, we include a domain to serve as a counterexample that induces this effect. To achieve this, we adopt the probabilities of the MarketScan data, but change the probabilities of 1 group such that interventions are barely effective. Full details are given in the [Multimedia Appendix 1](#).

Analyses

Our simulation analyses quantify the extent to which target clinical outcomes are achieved by calculating the numbers and proportions of patients reaching target HbA_{1c} levels (< 8). For all simulation experiments, we started with all patients in the *Engaged* state, with $HbA_{1c} \geq 8$, and a memory state of [M, M]. We divided data sets by 3 age ranges (aged 30-44, 45-54, and 55-64 years) and 2 genders (man and woman), creating 6 groups in total. The 6 groups had relative sizes of 0.175, 0.15, 0.2, 0.15, 0.125, and 0.2. To ensure each patient followed a unique behavior profile in simulation, for each patient in a group, we instantiated their transition probabilities by sampling each parameter from a normal distribution using the group value as the mean and $\sigma = 0.05$ SD.

Policies were optimized with $\alpha = .0$ unless otherwise noted.

We generated simulation results based on our 2 new equitable policies, MMR and MNW-EG which implemented the MMR and max Nash welfare (with equalized groups) policies, respectively.

We compared simulation results against 2 baselines that served as proxies for how our digital health program of interest, and similar ones, assign intervention resources, that is, based only on the current clinical state. Specifically, allocating interventions randomly each round on patients who are “High Risk,” that is,

patients with $s_C = HbA_{1c} \geq 8$ (termed high-risk random allocation), and a round robin approach which prioritized acting on patients with both $s_C = HbA_{1c} \geq 8$ and with the longest time period without an intervention (termed high-risk round robin allocation).

Additionally, we included a No Action baseline which simulated without assigning any intervention resources, to generate a lower bound of expected outcomes, that is the outcomes observed if individuals were not enrolled in a digital health program, but solely passively seeking care from the traditional primary care system.

We also compared against a state-of-the-art baseline (termed Opt), which assigns resources according to the asymptotically optimal utility-maximizing Whittle index policy [11,26].

Ethical Considerations

This is a simulation study, without human subject participation. World Medical Association Helsinki Declaration and informed consent guidelines are not applicable.

Results

Overview

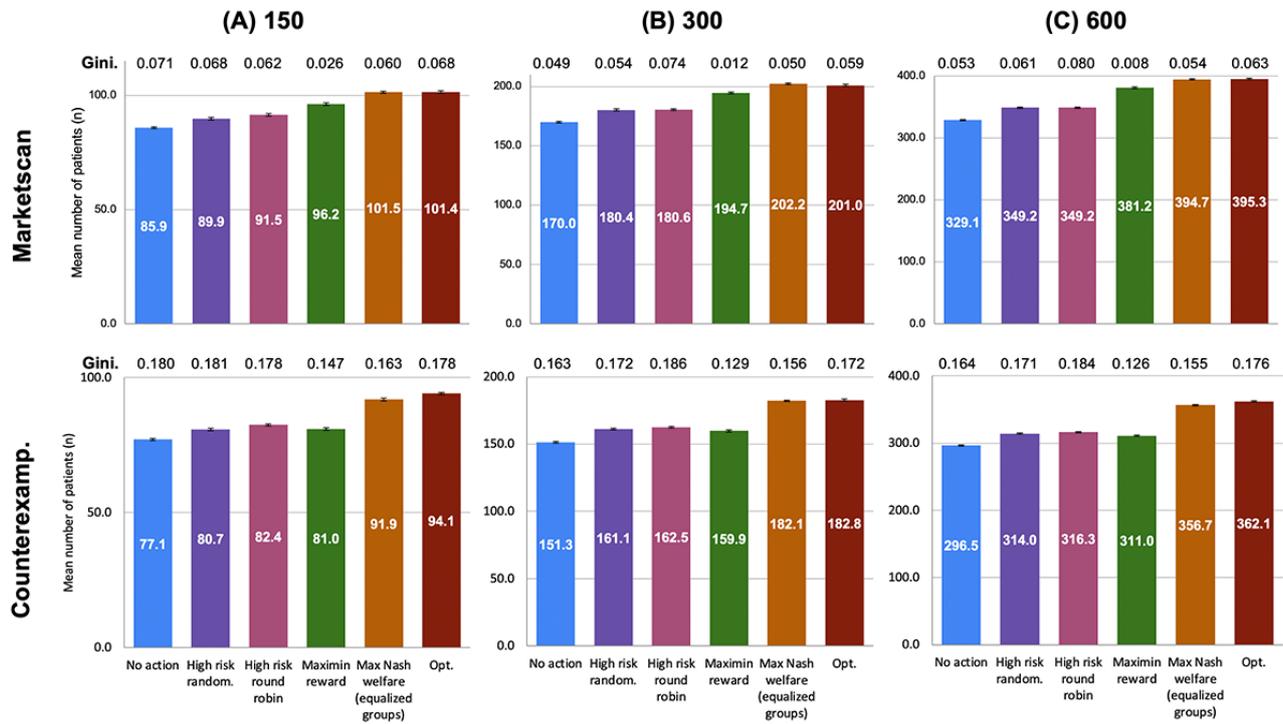
We evaluated our modeling and algorithmic contributions in simulation environments with data derived from publicly available sources on diabetes progression and health program engagement.

We ran experiments for $N \in \{150, 300, 600\}$ patients, horizon of $H = 18$ months, for budget values $B \in \{30, 60, 75, 100, 150\}$ and $\alpha \in \{0, .25, .50, .75, 1.0\}$. To simulate gradual patient enrollment over time, a real-world consideration raised by our digital program, 20% of patients are randomly added to the simulation in each of the first 5 months. Final statistics are all reported based on the health state of each patient after their 12th month in the simulation. We use the Gini coefficient [27] concerning each group's average final reward to measure the equity of each policy applied to each data distribution. Each combination of parameters was run for 50 random seeds, and the results show the average and SE over the seeds.

Achievement of Target Individual Health Outcomes

After 12 months, the Opt, MMR, and MNW-EG policies produced better individual clinical outcomes (measured by number of patients reaching healthy HbA_{1c} levels) and engagement than the baselines (Figure 4). The baselines increased the number of users with healthy HbA_{1c} after 12 months by roughly 5%, whereas at the same budget level, assignment policies considering joint clinical-engagement dynamics, that is, the Opt, MMR, and MNW-EG RMAB policies, could double this improvement, up to a further 10% on the MarketScan data set simulation analysis. MMR finds policies nearly 4-times more equitable, for little system-level cost. On the counterexample, MNW-EG avoids the pitfalls of maximin approaches, achieving more equity for little system-level cost.

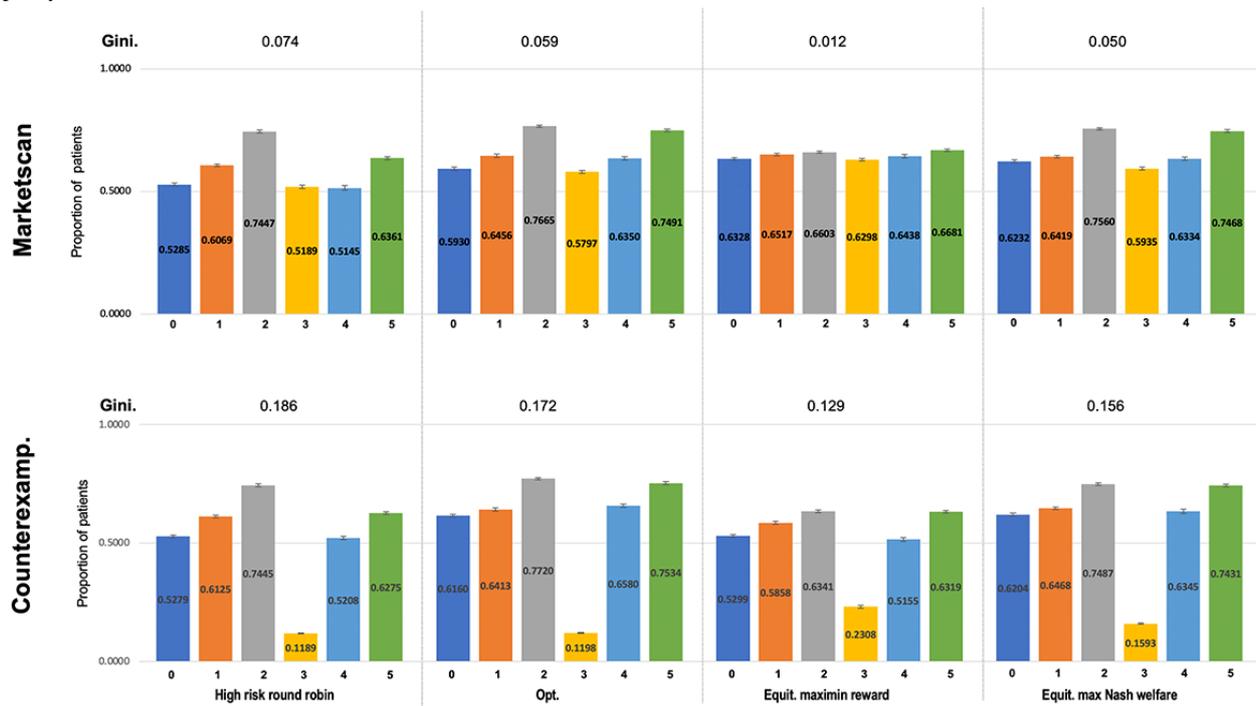
Figure 4. Individual clinical outcomes (average number of patients reaching healthy HbA_{1c} level) with each policy after 12 months, with a monthly intervention budget of B = N 10. Bars show average proportions by policy. Gini coefficient is displayed atop each policy (lower is better). Top: MarketScan. Bottom: MMR-Counterexample. Panels A, B, and C: analyses with N ∈ [150, 300, 600] patients, respectively. Counterexamp: counterexample; Max: maximum; Opt: baseline policy that assigns, based on the optimal utility-maximizing Whittle index policy; random: randomization.



MMR and MNW-EG achieved their lift in the proportion of patients with healthy HbA_{1c} while ensuring greater equity of outcomes across the groups (Figure 5). Specifically, MMR reduced inequity by nearly a factor of 4, at only a slight performance cost. In the counterexample domain (bottom row in the figure), we found that the overly conservative (by design) MMR over-committed resources to improving outcomes of the unmovable group, at the expense of the performance of all other

groups. However, in this case, MNW-EG maintained performance as good as Opt, while achieving the most equitable outcomes of any non-MMR policy. We included additional results in the Multimedia Appendix 1 that show analogous results when policies optimize strictly for engagement (ie, α = 1.0), conclusions held similarly, although the fair policies were able to achieve even greater improvements to equity over baselines.

Figure 5. Individual clinical outcomes (proportion of patients reaching healthy HbA_{1c} level) across demographic subgroups. Bars show average proportions by group (0-5) and policy. Gini coefficient is displayed atop each policy. N = 300, B = 30. Top: MarketScan. Bottom: MMR-Counterexample. Counterexamp: counterexample; Equit: equity; Max: maximum; Opt: baseline policy that assigns, based on the optimal utility-maximizing Whittle index policy.

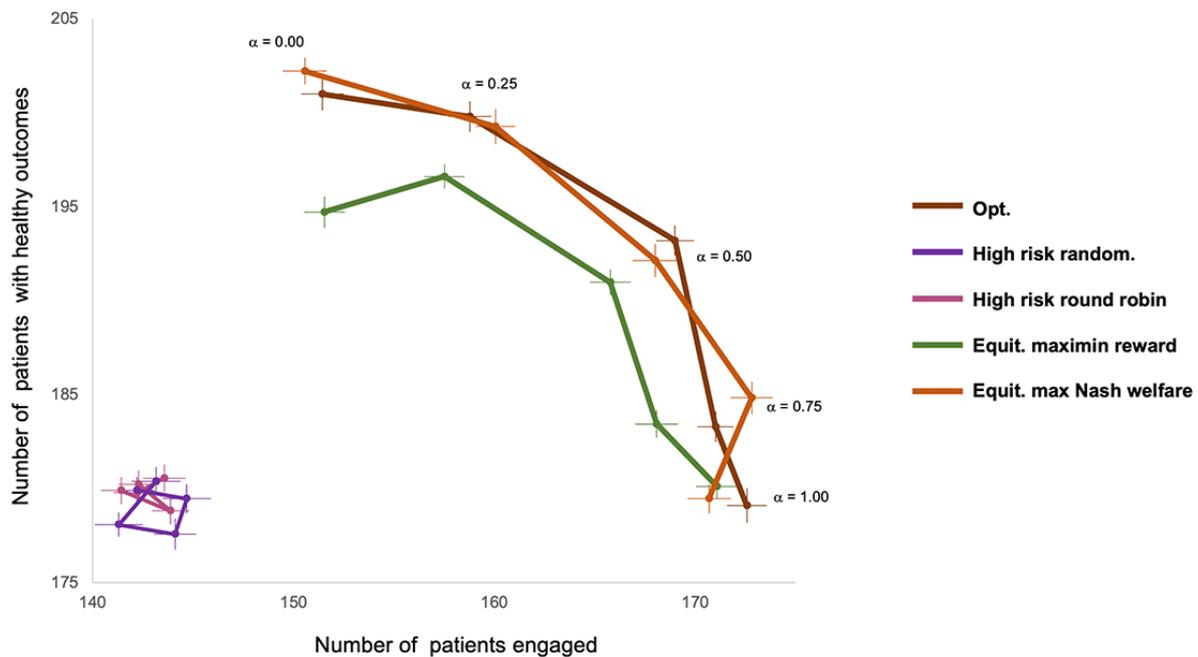


Policy Performance Under Different Preferred Specifications (Pareto Analysis)

Pareto analyses (Figure 6) showed that, even with the choice of $\alpha = 0$ (ie, optimizing only for health), MNW-EG and MMR approaches could achieve both improved health and improved engagement compared to clinical-only baselines. Interestingly, for the MarketScan data set, optimizing with $\alpha = 0.25$, that is, a quarter of reward weighted by engagement, could lead to

roughly a 10% total reduction in 12-month dropout compared to baselines, while maintaining the 10% boost in 12-month HbA_{1c} targets. We hypothesize that this is due to the “sticky” nature of healthy HbA_{1c} in the MarketScan data set, that is, patients with HbA_{1c} < 8 have a <1% chance of flipping back to HbA_{1c} > 8 in the next month. We give additional results in the Multimedia Appendix 1 for more values of the monthly budget B, and with the Gini index as an axis—the equitable policies remained fairer than Opt across choices of α and B.

Figure 6. Pareto curve for each policy as α varies from 0 to 1, with number of engaged patients after 12 months (ie, in E or M states) on the x-axis and number of patients with healthy clinical outcomes ($HbA_{1c} < 8$) on the y-axis. The results are shown for the MarketScan data set with $N = 300$ and $B = N/10$. Equit: equity; Max: maximum; Opt: baseline policy that assigns, based on the optimal utility-maximizing Whittle index policy; random: randomization.

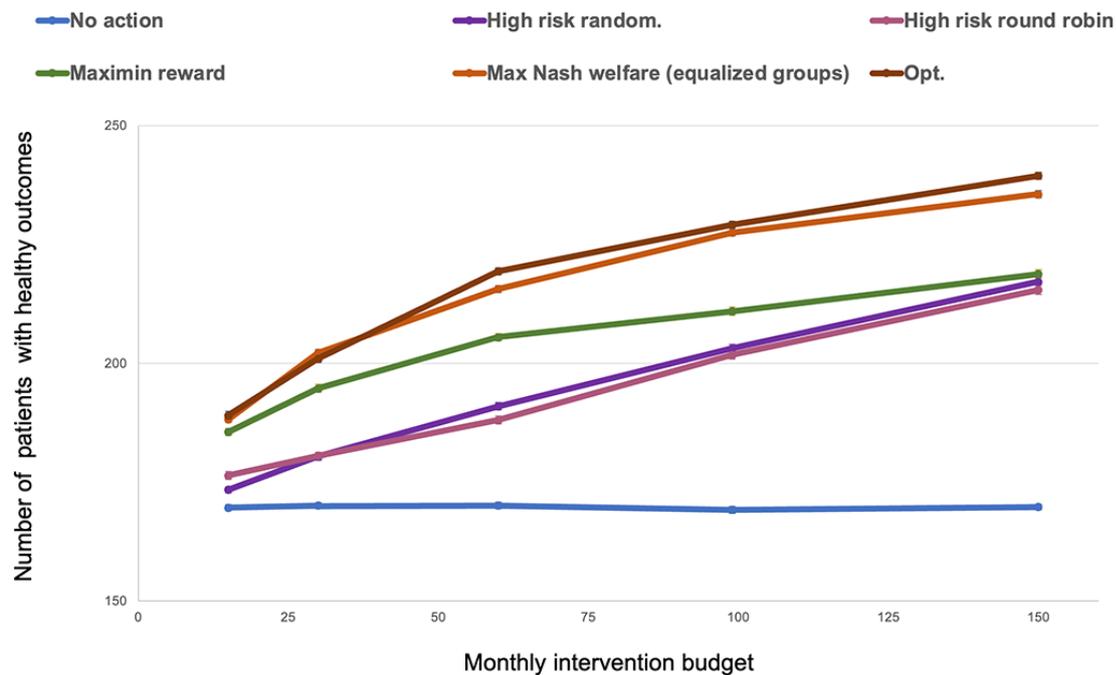


Clinical Outcomes According to Resource Allocations: Capacity Planning

Using the MarketScan data set, we performed analyses to estimate the clinical outcomes resulting from different levels of intervention resource allocations. These analyses demonstrated the capability to perform resource capacity planning for prospective cohorts using our MNW-EG and MMR approaches (Figure 7). For example, if the 12-month target was to reach 200 users with $HbA_{1c} < 8$, this analysis suggested that roughly 30 intervention resources would be needed if following the Opt policy or MNW-EG policies and 45 resources if

following the MMR policies. In contrast, the use of our baseline approaches to reach comparable goals would nearly double the budget, up to 100 monthly intervention resources. Additional results for the counterexample domain, and for α -weighted targets, found similar conclusions (Multimedia Appendix 1). These capacity planning plots allowed us to compute the “cost of fairness,” that is, the additional monthly intervention resources required for a more equitable policy to achieve the same total system-level return as the unfair optimal one, by estimating the horizontal difference between where each policy’s line intersects with the target dashed line. In our analysis, the cost of fairness for MNW-EG was negligible, but it was roughly 15 monthly resources for MMR.

Figure 7. Analysis of individual clinical outcomes according to resource allocation, MarketScan data set, N = 300. In this case, clinical outcomes are measured as the number of patients with healthy outcomes (ie, with $HbA_{1c} < 8$). Max: maximum; Opt: baseline policy that assigns, based on the optimal utility-maximizing Whittle index policy; random: randomization.



Discussion

Principal Findings

In this study of a digital health program in T2D, we used a simulation exercise to present a methodological approach to allocate resources in a digital health program with the potential to balance optimization of clinical outcomes, engagement of participants, and distribution of resources in an equitable fashion across participant subgroups. As an example of that potential, in our simplified simulation exercise, optimized intervention policies based on our proposed ERMAB framework led to 10% more patients reaching a healthy clinical outcome (defined by target HbA_{1c} levels) after 12 months, with a 10% reduction in program engagement dropout compared to standard-of-care baselines. Further, these new equitable policies reduced the mean absolute difference (a common equity measure) of engagement and health outcomes across 6 demographic groups by up to 85% compared to the state-of-the-art. We also demonstrated a new capability for a principled capacity planning system. That is, our system allows planners to estimate the required number of intervention resources needed for this digital health program to support a prospective cohort of patients, each with unique support needs and starting state, in reaching target HbA_{1c} levels. While this study was performed in a T2D setting, we believe that the general tenets of our observations may have applications across a spectrum of chronic diseases. Note that, for simulation, we streamlined our modeling approach, with simplified health goals and demographic groups based on age. Therefore, our quantitative results are merely illustrative, but the principles of this approach could be applied and enriched with more sophisticated modeling and other criteria, such as race or ethnicity, geographic location, or other salient sources

of existing inequity (as documented in diabetes care [19]), when information about those factors is available.

Comparison to Other Work

This work is related to a wide literature on using machine learning to make predictions in support of the delivery of digital health. Examples include predicting mood and depression [31], predicting medication adherence [32], ranking the efficacy of smoking cessation messages [33], and predicting heart arrhythmias from smart watches [34]. There are, however, several elements that contrast this study from others. While other works make predictions about the current or future states of a patient's health, they do not offer tools for planning the allocation of resources. Our work focuses on building up the algorithmic tools required for the long-term planning of allocations of limited resources in ways that will benefit the digital health system as a whole.

This study is also the first to formulate an RMAB model of digital health which has the novel characteristic of a multidimensional state space that encodes the joint dynamics of engagement and clinical health, giving the problem a relevant new structure, but increasing the computational complexity over previous domains.

Furthermore, we had equity-focused objectives, which viewed fairness through the lens of taking affirmative steps toward equitable outcomes. Overall considerations of equity in digital health are an underdeveloped area of study; prior or ongoing studies are still trying to measure the inequality problem in digital health in terms of usability, access, or feedback opportunities [35-39]. Most results show that societal inequalities at large have a reflection in the field of digital health, compounded by the issue of uneven technical access. These findings lend more urgency to the development of

optimizing strategies that tackle the problem of inequality intentionally and proactively. Our work is novel in that it proposes to formulate digital health programs to achieve outcome-based fairness. To our knowledge, this is the first study of its kind leveraging restless bandits and the first to give a principled framework for solving the problem of equitable outcomes with guarantees, in contrast to previous work on probabilistic fairness, which merely guaranteed each arm a lower bound of being considered for an intervention [16,17].

Specific Strengths

In addition, we demonstrated a key new capability of interest to digital health program administrators, namely the ability to perform resource capacity planning for prospective cohorts. This feature allows, for instance, to answer the question of whether the digital health program needs the same number of intervention resources to support a cohort of people aged 55-64 years from a particular region as it does to support a cohort of people aged 35-44 years from a different location. Given estimates of each cohort's clinical and engagement behaviors derived from historical data, one can simulate their preferred intervention policies to understand how many resources are needed to reasonably expect each cohort to reach their clinical goals. Capacity planning analysis, coupled with group-level evaluations of policy equity should allow planners to make principled decisions about resource needs for different populations.

Limitations

We acknowledge that this study also has limitations. As reported in this paper, we have only conducted simulation exercises with the analytical framework that we are proposing. We found the simulated results encouraging regarding the potential of our approach to achieve the objectives of allocating digital health program resources in a manner that is effective for reaching individual target clinical outcomes, and for maintaining patient engagement and population-level equitable care delivery throughout the process. However, further research applying this ERMAB framework in a real-world context is warranted to confirm the upside potential shown in simulations. In addition, our T2D model is simplified and we used claims data for our

simulations; claims have limitations as sources when inquiries go beyond information directly related to medical procedures, thus we opted for a simplified strategy accordingly. First, we are modeling a binary distinction for HbA_{1c} outcomes (< 8 or ≥ 8); while there is precedent for this approach, this simplification is still a limitation of the model. This cutoff point may not be optimal for all patients [40]. Second, the model does not consider comorbidities, which are highly relevant in diabetes, and chronic conditions in general, and could have meaningful effects on outputs, particularly the individual health outcomes. However, this model can be expanded with more granularity, as long as it can learn additional parameters from more sophisticated real-world data sets. These considerations (more individualized HbA_{1c} outcomes, comorbidities, and relevant subcohorts to the investigation of inequity) will all be important for future research based on other sources (such as electronic health records or clinical registries), to determine to which extent increasing complexity in the desired outcomes may affect the model's performance, and the practical implementation of the results.

Conclusions

In conclusion, our work showed the potential feasibility of planning interventions in digital health attending to several important factors in today's societal environment and resource-constrained systems. Our approach to intervention planning accounts not only for individual clinical outcome objectives but also for long-term participant engagement dynamics, using an RMAB sequential decision-making framework. We were able to simulate more equitable policies that could jointly improve engagement as well as clinical outcomes and demonstrated how the RMAB simulation framework could also provide key new capabilities in capacity planning, and objectively analyze how to trade-off between different program outcomes. Finally, we make a key new algorithmic contribution by introducing ERMABs and designing an efficient and fair approach for reaching population-level equitable solutions. We hope that ERMABs will add to the arsenal of tools available to practitioners addressing resource allocation problems in ethically sensitive domains.

Acknowledgments

The authors wish to acknowledge writing and editing support from Julia Saiz, from Verily Life Sciences. The code will be made available for publication at GitHub (GitHub, Inc).

Authors' Contributions

Study concept and design, draft writing and review, and draft approval for submission were done by all authors. Data collection was by JA. Data analysis and interpretation were performed by JAK, YJ, EH, and MJ.

Conflicts of Interest

This study was sponsored by Verily Life Sciences and Google Health. YJ, EH, and JA report employment and equity ownership in Verily Life Sciences. JAK was a student researcher at Google LLC and Verily Life Sciences. MJ and MT are employees of Google LLC and own Alphabet stock.

Multimedia Appendix 1

Additional description of methods.

[PDF File (Adobe PDF File), 1237 KB-Multimedia Appendix 1]

References

1. Widmer RJ, Collins NM, Collins CS, West CP, Lerman LO, Lerman A. Digital health interventions for the prevention of cardiovascular disease: a systematic review and meta-analysis. *Mayo Clin Proc.* 2015;90(4):469-480. [FREE Full text] [doi: [10.1016/j.mayocp.2014.12.026](https://doi.org/10.1016/j.mayocp.2014.12.026)] [Medline: [25841251](https://pubmed.ncbi.nlm.nih.gov/25841251/)]
2. Yardley L, Morton K, Greenwell K, Stuart B, Rice C, Bradbury K, et al. Digital interventions for hypertension and asthma to support patient self-management in primary care: the DIPSS research programme including two RCTs. In: Programme Grants for Applied Research. Southampton (UK). National Institute for Health and Care Research; 2022.
3. Ridho A, Alfian SD, van Boven JFM, Levita J, Yalcin EA, Le L, et al. Digital health technologies to improve medication adherence and treatment outcomes in patients with tuberculosis: systematic review of randomized controlled trials. *J Med Internet Res.* 2022;24(2):e33062. [FREE Full text] [doi: [10.2196/33062](https://doi.org/10.2196/33062)] [Medline: [35195534](https://pubmed.ncbi.nlm.nih.gov/35195534/)]
4. Torous J, Nicholas J, Larsen ME, Firth J, Christensen H. Clinical review of user engagement with mental health smartphone apps: evidence, theory and improvements. *Evid Based Ment Health.* 2018;21(3):116-119. [FREE Full text] [doi: [10.1136/eb-2018-102891](https://doi.org/10.1136/eb-2018-102891)] [Medline: [29871870](https://pubmed.ncbi.nlm.nih.gov/29871870/)]
5. Rosenlund M, Kinnunen UM, Saranto K. The use of digital health services among patients and citizens living at home: scoping review. *J Med Internet Res.* 2023;25:e44711. [FREE Full text] [doi: [10.2196/44711](https://doi.org/10.2196/44711)] [Medline: [36972122](https://pubmed.ncbi.nlm.nih.gov/36972122/)]
6. Meyerowitz-Katz G, Ravi S, Arnolda L, Feng X, Maberly G, Astell-Burt T. Rates of attrition and dropout in app-based interventions for chronic disease: systematic review and meta-analysis. *J Med Internet Res.* 2020;22(9):e20283. [FREE Full text] [doi: [10.2196/20283](https://doi.org/10.2196/20283)] [Medline: [32990635](https://pubmed.ncbi.nlm.nih.gov/32990635/)]
7. Ahmad FB, Cisewski JA, Anderson RN. Provisional mortality data—United States, 2021. *MMWR Morb Mortal Wkly Rep.* 2022;71(17):597-600. [FREE Full text] [doi: [10.15585/mmwr.mm7117e1](https://doi.org/10.15585/mmwr.mm7117e1)] [Medline: [35482572](https://pubmed.ncbi.nlm.nih.gov/35482572/)]
8. American Diabetes Association. Economic costs of diabetes in the U.S. in 2017. *Diabetes Care.* 2018;41(5):917-928. [FREE Full text] [doi: [10.2337/dci18-0007](https://doi.org/10.2337/dci18-0007)] [Medline: [29567642](https://pubmed.ncbi.nlm.nih.gov/29567642/)]
9. National Diabetes Statistics Report website. Centers for Disease Control and Prevention. 2022. URL: <https://www.cdc.gov/diabetes/data/statistics-report/index.html> [accessed 2023-08-08]
10. American Diabetes Association. 2. Classification and diagnosis of diabetes: standards of medical care in diabetes-2020. *Diabetes Care.* 2020;43(Suppl 1):S14-S31. [FREE Full text] [doi: [10.2337/dc20-S002](https://doi.org/10.2337/dc20-S002)] [Medline: [31862745](https://pubmed.ncbi.nlm.nih.gov/31862745/)]
11. Whittle P. Restless bandits: activity allocation in a changing world. *J Appl Probab.* 1988;25(A):287-298. [doi: [10.1017/s0021900200040420](https://doi.org/10.1017/s0021900200040420)]
12. Lee E, Lavieri MS, Volk M. Optimal screening for hepatocellular carcinoma: a restless bandit model. *Manuf Serv Oper Manag.* 2019;21(1):198-212. [FREE Full text] [doi: [10.1287/msom.2017.0697](https://doi.org/10.1287/msom.2017.0697)]
13. Mate A, Madaan L, Taneja A, Madhiwalla N, Verma S, Singh G, et al. Field study in deploying restless multi-armed bandits: assisting non-profits in improving maternal and child health. *Proc AAAI Conf Artif Intell.* 2022;36(11):12017-12025. [doi: [10.1609/aaai.v36i11.21460](https://doi.org/10.1609/aaai.v36i11.21460)]
14. Ayer T, Zhang C, Bonifonte A, Spaulding AC, Chhatwal J. Prioritizing hepatitis C treatment in U.S. prisons. *Oper Res.* 2019;67(3):853-873. [FREE Full text] [doi: [10.1287/opre.2018.1812](https://doi.org/10.1287/opre.2018.1812)]
15. 2019 National Healthcare Quality and Disparities Report. Rockville (MD). Agency for Healthcare Research and Quality (US); 2020.
16. Herlihy C, Prins A, Srinivasan A, Dickerson JP. Planning to fairly allocate: probabilistic fairness in the restless bandit setting. *ArXiv Preprint posted online on 19 Jul 2023.* 2021. [FREE Full text] [doi: [10.1145/3580305.3599467](https://doi.org/10.1145/3580305.3599467)]
17. Li D, Varakantham P. Towards soft fairness in restless multi-armed bandits. *ArXiv Preprint posted online on 27 Jul 2022.* 2022. [FREE Full text]
18. Mate A, Perrault A, Tambe M. Risk-Aware Interventions in Public Health: Planning with Restless Multi-Armed Bandits. Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems; Presented at: AAMAS '21: Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems; May 3-7, 2021, 2021;880-888; Virtual Event United Kingdom. URL: <https://dl.acm.org/doi/proceedings/10.5555/3463952> [doi: [10.1145/860620.860623](https://doi.org/10.1145/860620.860623)]
19. Agarwal S, Wade AN, Mbanya JC, Yajnik C, Thomas N, Egede LE, et al. The role of structural racism and geographical inequity in diabetes outcomes. *Lancet.* 2023;402(10397):235-249. [doi: [10.1016/S0140-6736\(23\)00909-1](https://doi.org/10.1016/S0140-6736(23)00909-1)] [Medline: [37356447](https://pubmed.ncbi.nlm.nih.gov/37356447/)]
20. Kaelbling LP, Littman ML, Cassandra AR. Planning and acting in partially observable stochastic domains. *Artif Intell.* 1998;101(1-2):99-134. [FREE Full text] [doi: [10.1016/s0004-3702\(98\)00023-x](https://doi.org/10.1016/s0004-3702(98)00023-x)]
21. Killian JA, Jain M, Jia Y, Amar J, Huang E, Tambe M. Equitable restless multi-armed bandits: a general framework inspired by digital health. *ArXiv Preprint posted online on 17 Aug 2023.* 2023. [FREE Full text]
22. Papadimitriou CH, Tsitsiklis JN. The complexity of optimal queuing network control. *Math Oper Res.* 1999;24(2):293-305. [doi: [10.1287/moor.24.2.293](https://doi.org/10.1287/moor.24.2.293)]

23. IBM MarketScan Research Databases help organizations take advantage of in-depth, patient-level information for easier, more accurate analysis. Watson Health. 2021. URL: https://www.ibm.com/docs/en/announcements/marketscan-research-databases?mhsrc=ibmsearch_a&mhq=marketscan [accessed 2024-02-24]
24. Cannon MJ, Masalovich S, Ng BP, Soler RE, Jabrah R, Ely EK, et al. Retention among participants in the national diabetes prevention program lifestyle change program, 2012-2017. *Diabetes Care*. 2020;43(9):2042-2049. [FREE Full text] [doi: [10.2337/dc19-2366](https://doi.org/10.2337/dc19-2366)] [Medline: [32616617](https://pubmed.ncbi.nlm.nih.gov/32616617/)]
25. Bergenstal RM, Layne JE, Zisser H, Gabbay RA, Barleen NA, Lee AA, et al. Remote application and use of real-time continuous glucose monitoring by adults with type 2 diabetes in a virtual diabetes clinic. *Diabetes Technol Ther*. 2021;23(2):128-132. [FREE Full text] [doi: [10.1089/dia.2020.0396](https://doi.org/10.1089/dia.2020.0396)] [Medline: [33026839](https://pubmed.ncbi.nlm.nih.gov/33026839/)]
26. Weber RR, Weiss G. On an index policy for restless bandits. *J Appl Probab*. 1990;27(3):637-648. [doi: [10.1017/s0021900200039176](https://doi.org/10.1017/s0021900200039176)]
27. Gini index. In: *The Concise Encyclopedia of Statistics*. New York, NY. Springer; 2008.
31. Burns MN, Begale M, Duffecy J, Gergle D, Karr CJ, Giangrande E, et al. Harnessing context sensing to develop a mobile intervention for depression. *J Med Internet Res*. 2011;13(3):e55. [FREE Full text] [doi: [10.2196/jmir.1838](https://doi.org/10.2196/jmir.1838)] [Medline: [21840837](https://pubmed.ncbi.nlm.nih.gov/21840837/)]
32. Killian JA, Wilder B, Sharma A, Choudhary V, Dilkina B, Tambe MS. Learning to Prescribe Interventions for Tuberculosis Patients Using Digital Adherence Data. New York, NY, United States. Association for Computing Machinery; Presented at: KDD '19: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining; August 4-8, 2019, 2019;2430-2438; Anchorage AK USA. URL: <https://dl.acm.org/doi/proceedings/10.1145/3292500> [doi: [10.1145/3292500.3330777](https://doi.org/10.1145/3292500.3330777)]
33. Sadasivam RS, Borglund EM, Adams R, Marlin BM, Houston TK. Impact of a collective intelligence tailored messaging system on smoking cessation: the perspect randomized experiment. *J Med Internet Res*. 2016;18(11):e285. [FREE Full text] [doi: [10.2196/jmir.6465](https://doi.org/10.2196/jmir.6465)] [Medline: [27826134](https://pubmed.ncbi.nlm.nih.gov/27826134/)]
34. Aschbacher K, Yilmaz D, Kerem Y, Crawford S, Benaron D, Liu J, et al. Atrial fibrillation detection from raw photoplethysmography waveforms: a deep learning application. *Heart Rhythm O2*. 2020;1(1):3-9. [FREE Full text] [doi: [10.1016/j.hroo.2020.02.002](https://doi.org/10.1016/j.hroo.2020.02.002)] [Medline: [34113853](https://pubmed.ncbi.nlm.nih.gov/34113853/)]
35. Areias AC, Molinos M, Moulder RG, Janela D, Scheer JK, Bento V, et al. The potential of a multimodal digital care program in addressing healthcare inequities in musculoskeletal pain management. *NPJ Digit Med*. 2023;6(1):188. [FREE Full text] [doi: [10.1038/s41746-023-00936-2](https://doi.org/10.1038/s41746-023-00936-2)] [Medline: [37816899](https://pubmed.ncbi.nlm.nih.gov/37816899/)]
36. Blount MA, Douglas MD, Li C, Walston DT, Nelms PL, Hughes CL, et al. Opportunities and challenges to advance health equity using digital health tools in underserved communities in Southeast US: a mixed methods study. *J Prim Care Community Health*. 2023;14:21501319231184789. [FREE Full text] [doi: [10.1177/21501319231184789](https://doi.org/10.1177/21501319231184789)] [Medline: [37401631](https://pubmed.ncbi.nlm.nih.gov/37401631/)]
37. Estevez M, Domecq S, Montagni I, Ramel V. Evaluating a public health information service according to users' socioeconomic position and health status: protocol for a cross-sectional study. *JMIR Res Protoc*. 2023;12:e51123. [FREE Full text] [doi: [10.2196/51123](https://doi.org/10.2196/51123)] [Medline: [37999943](https://pubmed.ncbi.nlm.nih.gov/37999943/)]
38. van de Vijver S, Tensen P, Asiki G, Requena-Méndez A, Heidenrijk M, Stronks K, et al. Digital health for all: how digital health could reduce inequality and increase universal health coverage. *Digit Health*. 2023;9:20552076231185434. [FREE Full text] [doi: [10.1177/20552076231185434](https://doi.org/10.1177/20552076231185434)] [Medline: [37434727](https://pubmed.ncbi.nlm.nih.gov/37434727/)]
39. Lyles CR, Nguyen OK, Khoong EC, Aguilera A, Sarkar U. Multilevel determinants of digital health equity: a literature synthesis to advance the field. *Annu Rev Public Health*. 2023;44:383-405. [FREE Full text] [doi: [10.1146/annurev-publhealth-071521-023913](https://doi.org/10.1146/annurev-publhealth-071521-023913)] [Medline: [36525960](https://pubmed.ncbi.nlm.nih.gov/36525960/)]
40. ElSayed NA, Aleppo G, Aroda VR, Bannuru RR, Brown FM, Bruemmer D, et al. 6. Glycemic targets: standards of care in diabetes-2023. *Diabetes Care*. 2023;46(Suppl 1):S97-S110. [FREE Full text] [doi: [10.2337/dc23-S006](https://doi.org/10.2337/dc23-S006)] [Medline: [36507646](https://pubmed.ncbi.nlm.nih.gov/36507646/)]

Abbreviations

ERMAB: equitable restless multiarmed bandit

HbA_{1c}: hemoglobin A_{1c}

MMR: maximin reward

MNW: maximum Nash welfare

NDPP: National Diabetes Prevention Program

Opt: baseline policy that assigns, based on the optimal utility-maximizing Whittle index policy

RMAB: restless multiarmed bandit

T2D: type 2 diabetes

Edited by T Leung; submitted 12.09.23; peer-reviewed by J Liu; comments to author 11.10.23; revised version received 18.10.23; accepted 15.02.24; published 15.03.24

Please cite as:

Killian JA, Jain M, Jia Y, Amar J, Huang E, Tambe M

New Approach to Equitable Intervention Planning to Improve Engagement and Outcomes in a Digital Health Program: Simulation Study

JMIR Diabetes 2024;9:e52688

URL: <https://diabetes.jmir.org/2024/1/e52688>

doi: [10.2196/52688](https://doi.org/10.2196/52688)

PMID: [38488828](https://pubmed.ncbi.nlm.nih.gov/38488828/)

©Jackson A Killian, Manish Jain, Yugang Jia, Jonathan Amar, Erich Huang, Milind Tambe. Originally published in JMIR Diabetes (<https://diabetes.jmir.org>), 15.03.2024. This is an open-access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work, first published in JMIR Diabetes, is properly cited. The complete bibliographic information, a link to the original publication on <https://diabetes.jmir.org/>, as well as this copyright and license information must be included.