

Field Study in Deploying Restless Multi-Armed Bandits: Assisting Non-Profits in Improving Maternal and Child Health

Aditya Mate^{12*}, Lovish Madaan^{1†}, Aparna Taneja¹, Neha Madhiwalla³, Shresth Verma¹,
Gargi Singh¹, Aparna Hegde³, Pradeep Varakantham⁴, Milind Tambe¹²

¹ Google Research

² Harvard University

³ ARMMAN

⁴ Singapore Management University

aditya_mate@g.harvard.edu, {lovishm, aparnataneja}@google.com, neha@armman.org, {vermashresth, gargisingh}@google.com, aparnahegde@armman.org, pradeepv@smu.edu.sg, milindtambe@google.com

Abstract

The widespread availability of cell phones has enabled non-profits to deliver critical health information to their beneficiaries in a timely manner. This paper describes our work to assist non-profits that employ automated messaging programs to deliver timely preventive care information to beneficiaries (new and expecting mothers) during pregnancy and after delivery. Unfortunately, a key challenge in such information delivery programs is that a significant fraction of beneficiaries drop out of the program. Yet, non-profits often have limited health-worker resources (time) to place crucial service calls for live interaction with beneficiaries to prevent such engagement drops. To assist non-profits in optimizing this limited resource, we developed a Restless Multi-Armed Bandits (RMABs) system. One key technical contribution in this system is a novel clustering method of offline historical data to infer unknown RMAB parameters. Our second major contribution is evaluation of our RMAB system in collaboration with an NGO, via a real-world service quality improvement study. The study compared strategies for optimizing service calls to 23003 participants over a period of 7 weeks to reduce engagement drops. We show that the RMAB group provides statistically significant improvement over other comparison groups, reducing $\sim 30\%$ engagement drops. To the best of our knowledge, this is the first study demonstrating the utility of RMABs in real world public health settings. We are transitioning our RMAB system to the NGO for real-world use.

1 Introduction

The wide-spread availability of cell phones has allowed non-profits to deliver targeted health information via voice or text messages to beneficiaries in underserved communities, often with significant demonstrated benefits to those communities (Pfammatter et al. 2016; Kaur et al. 2020). We focus in particular on non-profits that target improving maternal and infant health in low-resource communities in the global south.

*Work done during an internship at Google Research USA

†These authors contributed equally.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

These non-profits deliver ante- and post-natal care information via voice and text to prevent adverse health outcomes (Johnson 2017; ARMMAN 2020; HelpMum 2021).

Unfortunately, such information delivery programs are often faced with a key shortcoming: a large fraction of beneficiaries who enroll may drop out or reduce engagement with the information program. Yet non-profits often have limited health-worker time available on a periodic (weekly) basis to help prevent engagement drops. More specifically, there is limited availability of health-worker time where they can place crucial service calls (phone calls) to a limited number of beneficiaries, to encourage beneficiaries’ participation, address complaints and thus prevent engagement drops.

Optimizing limited health worker resources to prevent engagement drops requires that we prioritize beneficiaries who would benefit most from service calls on a periodic (e.g., weekly) basis. We model this resource optimization problem using Restless Multi-Armed Bandits (RMABs), with each beneficiary modeled as an RMAB arm. RMABs have been well studied for allocation of limited resources motivated by a myriad of application domains including preventive interventions for healthcare (Mate et al. 2020), planning anti-poaching patrols (Qian et al. 2016), machine repair and sensor maintenance (Glazebrook, Ruiz-Hernandez, and Kirkbride 2006) and communication systems (Sombabu et al. 2020). However, RMABs have rarely seen real world deployment, and to the best of our knowledge, never been deployed in the context of large-scale public health applications.

This paper presents first results of an RMAB system in real world public health settings. Based on available health worker time, RMABs choose m out of N total beneficiaries on a periodic (e.g., weekly) basis for service calls, where the m are chosen to optimize prevention of engagement drops. The paper presents two main contributions. First, previous work often assumes RMAB parameters as either known or easily learned over long periods of deployment. We show that both assumptions do not hold in our real-world contexts; instead, we present clustering of offline historical data as a novel approach to infer unknown RMAB parameters.

Our second contribution is a real world evaluation showing the benefit of our RMAB system, conducted in partnership with ARMMAN¹, an NGO in India focused on maternal and child care. ARMMAN conducts a large-scale health information program, with concrete evidence of health benefits, which has so far served over a million mothers. As part of this program, an automated voice message is delivered to an expecting or new mother (beneficiary) over her cell phone on a weekly basis throughout pregnancy and for a year post birth in a language and time slot of her preference.

Unfortunately, ARMMAN’s information delivery program also suffers from engagement drops. Therefore, in collaboration with ARMMAN we conducted a service quality improvement study to maximize the effectiveness of their service calls to ensure beneficiaries do not drop off from the program or stop listening to weekly voice messages. More specifically, the current standard of care in ARMMAN’s program is that any beneficiary may initiate a service call by placing a so called “missed call”. This beneficiary-initiated service call is intended to help address beneficiaries’ complaints and requests, thus encouraging engagement. However, given the overall decreasing engagement numbers in the current setup, key questions for our study are to investigate an approach for effectively conducting additional ARMMAN-initiated service calls (these are limited in number) to reduce engagement drops. To that end, our service quality improvement study comprised of 23,003 real-world beneficiaries spanning 7 weeks. Beneficiaries were divided into 3 groups, each adding to the current standard of care. The first group exercised ARMMAN’s current standard of care (CSOC) without additional ARMMAN-initiated calls. In the second, the RMAB group, ARMMAN staff added to the CSOC by initiating service calls to 225 beneficiaries on average per week chosen by RMAB. The third was the Round-Robin group, where the exact same number of beneficiaries as the RMAB group were called every week based on a systematic sequential basis.

Results from our study demonstrate that RMAB provides statistically significant improvement over CSOC and round-robin groups. This improvement is also practically significant — the RMAB group achieves a $\sim 30\%$ reduction in engagement drops over the other groups. Moreover, the round-robin group does not achieve statistically significant improvement over the CSOC group, i.e., RMAB’s optimization of service calls is crucial. To the best of our knowledge, this is the first large-scale empirical validation of use of RMABs in a public health context. Based on these results, the RMAB system is currently being transitioned to ARMMAN to optimize service calls to their ever growing set of beneficiaries. Additionally, this methodology can be useful in assisting engagement in many other awareness or adherence programs, e.g., Thirumurthy and Lester (2012); Chen et al. (2021). *Our RMAB code would be released upon acceptance.*

¹<https://armman.org/>

2 Related Work

Patient adherence monitoring in healthcare has been shown to be an important problem (Martin et al. 2005), and is closely related to the churn prediction problem, studied extensively in the context of industries like telecom (Dahiya and Bhatia 2015), finance (Xie et al. 2009; Shaaban et al. 2012), etc. The healthcare domain has seen several studies on patient adherence for diseases like HIV (Tuldrà et al. 1999), cardiac problems (Son et al. 2010; Corotto et al. 2013), Tuberculosis (Killian et al. 2019; Pilote et al. 1996), etc. These studies use a combination of patient background information and past adherence data, and build machine learning models to predict future adherence to prescribed medication². However, such models treat adherence monitoring as a single-shot problem and are unable to appropriately handle the sequential resource allocation problem at hand. Additionally, the pool of beneficiaries flagged as high risk can itself be large, and the model can’t be used to prioritize calls on a periodic basis, as required in our settings.

Campaign optimization (via phone outreach) has also been studied previously. Most existing works (Leskovec, Adamic, and H. 2007; Eagle, Macy, and Claxton 2010) however, rely on the availability of a customer social network based on preferences, behavior or demographics, to help identify the set of key customers who will increase the reach of the campaign. In our domains of interest, there is no evidence of a social network among the beneficiaries, so such campaign optimization techniques are inapplicable. Furthermore, campaign optimization relies on single-shot interventions for optimization, whereas, our problem requires tracking progress of beneficiaries over multiple timesteps.

The Restless Multi-Armed Bandit (RMAB) framework has been popularly adopted to tackle such sequential resource allocation problems (Whittle 1988; Jung and Tewari 2019). Computing the optimal solution for RMAB problems is shown to be PSPACE-hard. Whittle proposed an index-based heuristic (Whittle 1988), that can be solved in polynomial time and is now the dominant technique used for solving RMABs. It has been shown to be asymptotically optimal for the time average reward problem (Weber and Weiss 1990), and other families of RMABs arising from stochastic scheduling problems (Glazebrook, Ruiz-Hernandez, and Kirkbride 2006). Several works as listed in Section 1, show applicability of RMABs in different domains but these unrealistically assume perfect knowledge of the RMAB parameters, and have not been tested in real-world contexts. Biswas et al. (2021); Avrachenkov and Borkar (2020), present a Whittle Index-based Q-learning approach for unknown RMAB parameters. However, their techniques either assume identical arms or rely on receiving thousands of samples from each arm, which is unrealistic in our setting, given limited overall stay of a beneficiary in an information program — a beneficiary may drop out or stop engaging with the program few weeks post enrolment un-

²Similarly, in our previous preliminary study (anonymous 2020) published in a non-archival setting, we used demographic and message features to build models for predicting beneficiaries likely to drop-off from ARMMAN’s information program.

less a service call convinces them to do otherwise. Instead, we present a novel approach that applies clustering to the available historical data to infer model parameters.

Clustering in the context of Multi-Armed Bandit and Contextual Bandits has received significant attention in the past (Gentile, Li, and Zappella 2014; Li, Chen, and Leung 2019; Yang et al. 2020; Li, Wu, and Wang 2021), but these settings do not consider restless bandit problems. (Mintz et al. 2020) tackles a non-stationary setup with stochastic rewards, while (Ayer et al. 2019) infers model parameters from independent studies in absence of historic data. In contrast, we focus on learning RMAB parameters using clustered historic beneficiary data. (Zhou et al. 2018; Liao et al. 2020) propose building predictive models per beneficiary in an online fashion, which is infeasible in our setup given the short stay of the beneficiaries.

3 Preliminaries

Background: Restless Multi-Armed Bandits

An RMAB instance consists of N independent 2-action Markov Decision Processes (MDP) (Puterman 1994), where each MDP is defined by the tuple $\{\mathcal{S}, \mathcal{A}, R, \mathcal{P}\}$. \mathcal{S} denotes the state space, \mathcal{A} is the set of possible actions, R is the reward function $R : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ and \mathcal{P} represents the transition function. We use $P_{s,s'}^\alpha$ to denote the probability of transitioning from state s to state s' under the action α . The policy π , is a mapping $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that selects the action to be taken at a given state. The total reward accrued can be measured using either the discounted or average reward criteria to sum up the immediate rewards accrued by the MDP at each time step. Our formulation is amenable to both, although we use the discounted reward criterion in our study.

The expected *discounted reward* starting from state s_0 is defined as $V_\beta^\pi(s_0) = \mathbb{E}[\sum_{t=0}^{\infty} \beta^t R(s_t, \pi(s_t), s_{t+1} | \pi, s_0)]$ where the next state is drawn according to $s_{t+1} \sim P_{s_t, s_{t+1}}^{\pi(s_t)}$, $\beta \in [0, 1]$ is the discount factor and actions are selected according to the policy mapping π . The planner’s goal is to maximize the total reward.

We model the engagement behavior of each beneficiary by an MDP corresponding to an arm of the RMAB. Pulling an arm corresponds to an active action, i.e., making a service call (denoted by $\alpha = a$), while $\alpha = p$ denotes the passive action of abstaining from a call. The state space \mathcal{S} consists of binary valued states, s , that account for the recent engagement behavior of the beneficiary; $s \in [NE, E]$ (or equivalently, $s \in [0, 1]$) where E and NE denote the ‘Engaging’ and ‘Not Engaging’ states respectively. For example, in our domain, ARMMAN considers that if a beneficiary stays on the automated voice message for more than 30 seconds (average message length is 1 minute), then the beneficiary has engaged. If a beneficiary engages at least once with the automated voice messages sent during a week, they are assigned the engaging (E) state for that time step and non-engaging (NE) state otherwise. For each action $\alpha \in \mathcal{A}$, the beneficiary states follow a Markov chain represented by the 2-state Gilbert-Elliot model (Gilbert 1960) with transition parameters given by $P_{ss'}^\alpha$, as shown in Figure 1. With slight abuse

of notation, the reward function $R(\cdot)$ of n^{th} MDP is simply given by $R_n(s) = s$ for $s \in \{0, 1\}$.

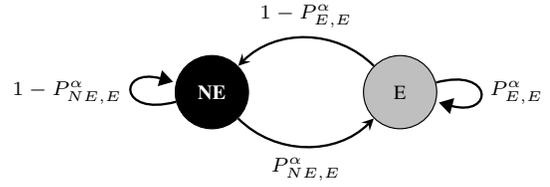


Figure 1: The beneficiary transitions from a current state s to a next state s' under action α , with probability $P_{ss'}^\alpha$.

We adopt the Whittle solution approach described previously for solving the RMAB. It hinges around the key idea of a ‘passive subsidy’, which is a hypothetical reward offered to the planner, in addition to the original reward function for choosing the passive action. The Whittle Index is then defined as the infimum subsidy that makes the planner indifferent between the ‘active’ and the ‘passive’ actions, i.e.,:

$$W(s) = \inf_{\lambda} \{ \lambda : Q_{\lambda}(s, p) = Q_{\lambda}(s, a) \} \quad (1)$$

Data Collected by ARMMAN

Beneficiaries enroll into ARMMAN’s information program with the help of health workers, who collect the beneficiary’s demographic data such as age, education level, income bracket, phone owner in the family, gestation age, number of children, preferred language and preferred slots for the automated voice messages during enrolment. These features are referred to as Beneficiary Registration Features in rest of the paper. Beneficiaries provided both written and digital consent for receiving automated voice messages and service calls. ARMMAN also stores listenership information regarding the automated voice messages together with the registration data in an anonymized fashion.

4 Problem Statement

We assume the planner has access to an offline historical data set of beneficiaries, \mathcal{D}_{train} . Each beneficiary data point $\mathcal{D}_{train}[i]$ consists of a tuple, $\langle f, \mathcal{E} \rangle$, where f is beneficiary i ’s feature vector of static features, and \mathcal{E} is an episode storing the trajectory of (s, α, s') pairs for that beneficiary, where s denotes the start state, α denotes the action taken (passive v/s active), and s' denotes the next state that the beneficiary lands in after executing α in state s . We assume that these (s, α, s') samples are drawn according to fixed, latent transition matrices $P_{ss'}^a[i]$ and $P_{ss'}^p[i]$ (corresponding to the active and passive actions respectively), unknown to the planner, and potentially unique to each beneficiary.

Given \mathcal{D}_{train} , we now consider a new beneficiary cohort \mathcal{D}_{test} , consisting of N beneficiaries, marked $\{1, 2, \dots, N\}$, that the planner must plan service calls for. The MDP transition parameters corresponding to beneficiaries in \mathcal{D}_{test} are unknown to the planner, but assumed to be drawn at random from a distribution similar to the joint distribution of features and transition parameters of beneficiaries in the historical data distribution. We assume the planner has access to the feature vector f for each beneficiary in \mathcal{D}_{test} .

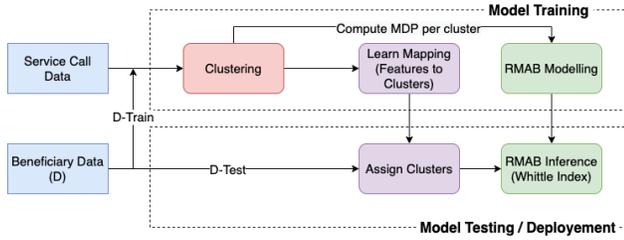


Figure 2: RMAB Training and Testing pipelines proposed

We now define the service call planning problem as follows. The planner has upto m resources available per round, which the planner may spend towards delivering service calls to beneficiaries. Beneficiaries are represented by N arms of the RMAB, of which the planner may pull upto m arms (i.e., m service calls) at each time step. We consider a round or timestep of one week which allows planning based on the most recent engagement patterns of the beneficiaries.

5 Methodology

Figure 2 shows our overall solution methodology. We use clustering techniques that exploit historical data D_{train} to estimate an offline RMAB problem instance relying solely on the beneficiaries’ static features and state transition data. This enables overcoming the challenge of limited samples (time-steps) per beneficiary. Based on this estimation, we use the Whittle Index approach to prioritize service calls.

Clustering Methods

We use historical data D_{train} to learn the impact of service calls on transition probabilities. While there is limited service call data (active transition samples) for any single beneficiary, clustering on the beneficiaries allows us to combine their data to infer transition probabilities for the entire group. Clustering offers the added advantage of reducing computational cost for resource limited NGOs; since all beneficiaries within a cluster share identical transition probability values we can compute their Whittle index all at once. We present four such clustering techniques below:

1. Features-only Clustering (FO): This method relies on the correlation between the beneficiary feature vector f and their corresponding engagement behavior. We employ k-means clustering on the feature vector f of all beneficiaries in the historic dataset D_{train} , and then derive the representative transition probabilities for each cluster by pooling all the (s, α, s') tuples of beneficiaries assigned to that cluster. At test time, the features f of a new, previously unseen beneficiary in D_{test} map the beneficiary to their corresponding cluster and estimated transition probabilities.

2. Feature + All Probabilities (FAP) In this 2-level hierarchical clustering technique, the first level uses a rule-based method, using features to divide beneficiaries into a large number of pre-defined buckets, B . Transition probabilities are then computed by pooling the (s, α, s') samples from all the beneficiaries in each bucket. Finally, we perform a

k-means clustering on the transition probabilities of these B buckets to reduce them to k clusters ($k \ll B$). However, this method suffers from several smaller buckets missing or having very few active transition samples.

3. Feature + Passive Probabilities (FPP): This method builds on the FAP method, but only considers the passive action probabilities to preclude the issue of missing active transition samples.

4. Passive Transition-Probability based Clustering (PPF): The key motivation here is to group together beneficiaries with similar transition behaviors, irrespective of their features. To this end, we use k-means clustering on passive transition probabilities (to avoid issues with missing active data) of beneficiaries in D_{train} and identify cluster centers. We then learn a map ϕ from the feature vector f to the cluster assignment of the beneficiaries that can be used to infer the cluster assignments of new beneficiaries at test-time solely from f . We use a random forest model as ϕ .

The rule-based clustering on features involved in *FPP* and *FAP* methods can be thought of as using one specific, hand-tuned mapping function ϕ . In contrast, the *PPF* method *learns* such a map ϕ from data, eliminating the need to manually define accurate and reliable feature buckets.

Evaluation of Clustering Methods

We use a historical dataset, D_{train} from ARMMAN consisting of 4238 beneficiaries in total, who enrolled into the program between May-July 2020. We compare the clustering methods empirically, based on the criteria described below.

1. Representation: Cluster centers that are representative of the underlying data distribution better resemble the ground truth transition probabilities. This is of prime importance to the planner, who must rely on these values to plan actions. Fig 3 plots the ground truth transition probabilities and the resulting cluster centers determined using the proposed methods. Visual inspection reveals that the *PPF* method represents the ground truth well, as is corroborated by the quantitative metrics of Table 1 that compares the RMSE error across different clustering methods.

2. Balanced cluster sizes: A low imbalance across cluster sizes is desirable to preclude the possibility of arriving at few, gigantic clusters which will assign identical whittle indices to a large groups of beneficiaries. Working with smaller clusters also aggravates the missing data problem in estimation of active transition probabilities. Considering the variance in cluster sizes and RMSE error for the different clustering methods with $k = \{20, 40\}$ as shown in Table 1, *PPF* outperforms the other clustering methods and was chosen for the pilot study.

Next we turn to choosing k , the number of clusters: as k grows, the clusters become sparse in number of active samples aggravating the missing data problem while a smaller k suffers from a higher RMSE. We found $k = 40$ to be optimal and chose it for the pilot study .

Finally, we adopt the Whittle solution approach for RMABs to plan actions and pre-compute all of the possible $2 * k$ index values that beneficiaries can take (corresponding to combinations of k possible clusters and 2 states). The

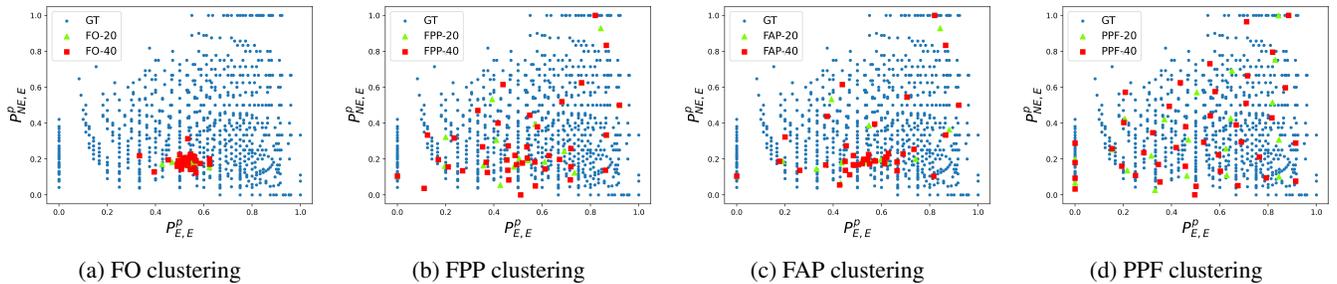


Figure 3: Comparison of passive transition probabilities obtained from different clustering methods with cluster sizes $k = \{20, 40\}$ with the ground truth transition probabilities. Blue dots represent the true passive transition probabilities for every beneficiary while red or green dots represent estimated cluster centres.

Table 1: Average RMSE and cluster size variance over all beneficiaries for different methods. Total Beneficiaries = 4238, $\mu_{20} = 211.9$, $\mu_{40} = 105.95$ (μ = average beneficiaries per cluster)

Clustering Method	Average RMSE		Standard Deviation	
	$k = 20$	$k = 40$	$k = 20$	$k = 40$
FO	0.229	0.228	143.30	74.22
FPP	0.223	0.222	596.19	295.01
FAP	0.224	0.223	318.46	218.37
PPF	0.041	0.027	145.59	77.50

indices can then be looked up at all future time steps in constant time, making this an optimal solution for large scale deployment with limited compute resources.

As we got this RMAB system ready for real-world use, there was as an important observation for social impact settings: real-world use also required us to carefully handle several domain specific challenges, which were time consuming. For example, despite careful clustering, a few clusters may still be missing active probability values, which required employing a data imputation heuristic. Moreover, there were other constraints specific to ARMMAN, such as a beneficiary should receive only one service call every η weeks, which was addressed by introducing “sleeping states” for beneficiaries who receive a service call.

6 Experimental Study

In this section, we discuss a real-world quality improvement study. We also simulate the expected outcome in other synthetically constructed situations and demonstrate good performance of our approach across the board.

Service Quality Improvement Study

Setup This cohort of beneficiaries registered in the program between Feb 16, 2021 and March 15, 2021 as D_{test} and started receiving automated voice messages few days post enrolment as per their gestational age. Additionally, as per the current standard of care, any of these beneficiaries could initiate a service call by placing a “missed call”. The 23003 beneficiaries are randomly distributed across 3 groups, each group adding to the CSOC as follows:

- **Current-Standard-of-Care (CSOC) Group:** The beneficiaries in this group follow the original standard of care, where there are no ARMMAN initiated service calls. The listenership behavior of beneficiaries in this group is used as a benchmark for the RR and RMAB groups.
- **RMAB group:** In this group, beneficiaries are selected for ARMMAN-initiated service call per week via the Whittle Index policy described in Section 3. Even though all beneficiaries within a cluster are modeled by identical MDP parameters, their states may evolve independently, and so the Whittle indices are tracked for each beneficiary separately, leading to an RMAB with 7668 arms.
- **Round Robin (RR) group:** By default, NGOs including ARMMAN often conduct service calls using some systematic set order – the idea here is to have an easily executable policy, that services enough of a cross-section of beneficiaries and can be scaled up or down per week based on available resources. To recreate this setting, we generate service calls to beneficiaries based on the ascending order of their date of enrollment for this RR group, as recommended by ARMMAN. If this method succeeds compared to CSOC, then a simple manual strategy is enough; RMAB style optimization may not be needed.

Table 2 shows the absolute number of beneficiaries in states E or NE, where the state is computed using one week of engagement data between April 19 - April 26, 2021.

Beneficiaries across all three groups receive the same automated voice messages regarding pregnancy and post-birth care throughout the program, and no health related information is withheld from any beneficiary. The study only aims to evaluate the effectiveness of ARMMAN-initiated outbound service calls with respect to improving engagement with the program across the three groups. No interviews or research data or feedback was collected from the beneficiaries.

The study started on April 26, 2021, with m beneficiaries selected from the RMAB and RR group each ($m \ll N$) per week for ARMMAN-initiated service calls. ARMMAN staff performing service calls were blind to the experimental groups that the beneficiaries belonged to. Recall, the goal of the service calls is to encourage the beneficiaries to engage with the health information message program in the future. For this study, number of service calls m was on average 225

Table 2: Beneficiary distribution in the three groups and their start states during week 0 of the study.

Group	Engaging (E)	Non-Engaging (NE)	Total
RMAB	3571	4097	7668
RR	3647	4021	7668
CSOC	3661	4006	7667

per week for each of RMAB and RR groups to reflect real-world constraints on service calls. The study was scheduled for a total of 7 weeks, during which 20% of the RMAB (and RR) group had received a service call, which is closer to the percentage of population that may be reached in service calls by ARMMAN.³

Results We present our key results from the study in Figure 4. The results are computed at the end of 7 weeks from the start of the quality improvement study on April 26, 2021.

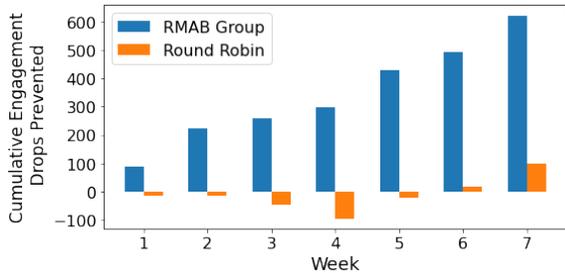


Figure 4: Cumulative number of weekly engagement drops prevented (in comparison to the CSOC group) by RMAB far exceed those prevented by RR.

Figure 4 measures the impact of service calls by the RMAB and RR policies in comparison to the CSOC Group. Beneficiaries’ engagement with the program typically starts to dwindle with time. In Figure 4, we measure the impact of a service call policy as the cumulative drop in engagement prevented compared to the CSOC Group. We consider drop in engagement instead of the raw engagement numbers themselves, because of the slight difference in the numbers of beneficiaries in engaging (E) state at the start of the study. The drop in engagement under a policy π at time t can be measured as the change in engagement:

$$\Delta_{current}^{\pi}(t) := \sum_{n \in N} (R_n(s_0) - R_n(s_t)) \quad (2)$$

where $R_n(s_t)$ represents the reward for n^{th} beneficiary in state s_t at time step t and cumulative drop in engagement is:

$$\Delta_{cumulative}^{\pi}(t) := \sum_{n \in N} \sum_{\zeta=0}^{\zeta=t} (R_n(s_0) - R_n(s_{\zeta})) \quad (3)$$

³Each beneficiary group also received very similar beneficiary-initiated calls, but these were less than 10% of the ARMMAN-initiated calls in RMAB or RR groups over 7 weeks.

The cumulative drop in engagement prevented by a policy π , in comparison to the CSOC Group is thus simply:

$$\Delta_{cumulative}^{\pi}(t) - \Delta_{cumulative}^{CSOC}(t) \quad (4)$$

and is plotted on the y-axis of Figure 4.

Figure 4 shows that the RMAB policy prevents a total 622 instances of a drop in automated health message engagement, at the end of 7 weeks, as compared to CSOC. RR group, on the other hand, only prevents 101 engagement drops by the end of week 7. Given that there are a total of 1944 engagement drops in the CSOC group, we show in the first row of Table 3, that the RMAB group has 32.0% and 28.3% less cumulative engagement drops as compared to the CSOC and RR groups respectively by the end of the study.

Table 3: Statistical significance for service call policy impact at week 7 is tested using a linear regression model. We use: * $p < 0.05$; † $p < 0.1$

	RMAB vs CSOC	RR vs CSOC	RMAB vs RR
% reduction in cumulative engagement drops	32.0%	5.2%	28.3%
p-value	0.044*	0.740	0.098†
Coefficient β	-0.0819	-0.0137	-0.0068

Statistical Analysis To investigate the benefit from use of RMAB policy over policies in the RR and CSOC groups, we use regression analysis (Angrist and Pischke 2008)⁴. Specifically, we fit a linear regression model to predict number of cumulative engagement drops at week 7 while controlling for treatment assignment and covariates specified by beneficiary registration features. The model is given by:

$$Y_i = k + \beta T_i + \sum_{j=1}^J \gamma_j x_{ij} + \epsilon_i$$

where for the i_{th} beneficiary, Y_i is the outcome variable defined as number of cumulative engagement drops at week 7, k is the constant term, β is the treatment effect, T_i is the treatment indicator variable, x_i is a vector of length J representing the i_{th} beneficiary’s registration features, γ_j represents the impact of the j_{th} feature on the outcome variable and ϵ_i is the error term. For evaluating the effect of RMAB service calls as compared to CSOC group, we fit the regression model only for the subset of beneficiaries assigned to either of these two groups. T_i is set to 1 for beneficiaries belonging to the RMAB group and 0 for those in CSOC group. We repeat the same experiment to compare RR vs CSOC group and RMAB vs RR group.

The results are summarized in Table 3. We find that RMAB has a statistically significant treatment effect in reducing cumulative engagement drop (negative β , $p < 0.05$) as compared to CSOC group. However, the treatment effect is not statistically significant when comparing RR with CSOC group ($p = 0.740$). Additionally, comparing RMAB

⁴See Appendix A for erratum

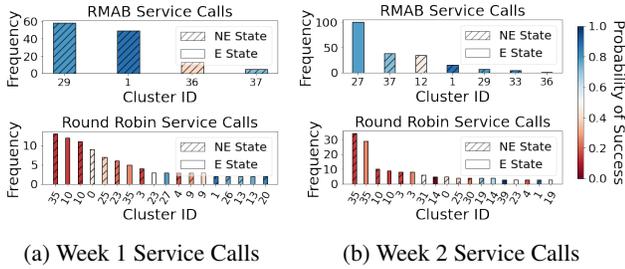


Figure 5: Distributions of clusters picked for service calls by RMAB and RR are significantly different. RMAB is very strategic in picking only a few clusters with a promising probability of success, RR displays no such selection.

group with RR, we find β , the RMAB treatment effect, to be significant ($p < 0.1$). This shows that RMAB policy has a statistically significant effect on reducing cumulative engagement drop as compared to both the RR policy and CSOC. RR fails to achieve statistical significance against CSOC. Together these results illustrate the importance of RMAB’s optimization of service calls, and that without such optimization, service calls may not yield any benefits.

RMAB Strategies We analyse RMAB’s strategic selection of beneficiaries in comparison to RR using Figure 5, where we group beneficiaries according to their whittle indices, equivalently their $\langle \text{cluster}, \text{state} \rangle$. Figure 5 plots the frequency distribution of beneficiaries (shown via corresponding clusters) who were selected by RMAB and RR in the first two weeks. For example, the top plot in Figure 5a shows that RMAB selected 60 beneficiaries from cluster 29 (NE state). First, we observe that RMAB was clearly more selective, choosing beneficiaries from just four (Figure 5a) or seven (Figure 5b) clusters, rather than RR that chose from 20. Further, we assign each cluster a hue based on their probability of transitioning to engaging state from their current state given a service call. Figure 5 reveals that RMAB consistently prioritizes clusters with high probability of success (blue hues) while RR deploys no such selection; its distribution emulates the overall distribution of beneficiaries across clusters (mixed blue and red hues).

Furthermore, Figure 6a further highlights the situation in week 1, where RMAB spent 100% of its service calls on beneficiaries in the non-engaging state while RR spent the same on only 64%. Figure 6b shows that RMAB converts 31.2% of the beneficiaries shown in Figure 6a from non-engaging to engaging state by week 7, while RR does so for only 13.7%. This further illustrates the need for optimizing service calls for them to be effective, as done by RMAB.

Synthetic Results

We run additional simulations to test other service call policies beyond those included in the quality improvement study and confirm the superior performance of RMAB. Specifically, we compare to the following baselines: (1) RANDOM is a naive baseline that selects m arms at random. (2) MYOPIC is a greedy algorithm that pulls arms optimizing for the

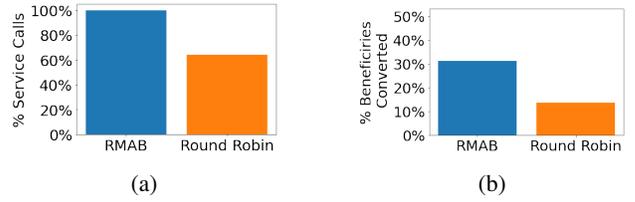


Figure 6: (a) % of week 1 service calls on non-engaging beneficiaries (b) % of non-engaging beneficiaries of week 1 receiving service calls that converted to engaging by week 7

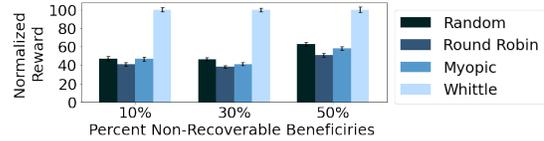


Figure 7: Performance of MYOPIC can be arbitrarily bad and even worse than RANDOM, unlike the Whittle policy.

reward in the immediate next time step. WHITTLE is our algorithm. We compute a normalized reward of an algorithm ALG as: $\frac{100 \times (\bar{R}^{\text{ALG}} - \bar{R}^{\text{CSOC}})}{\bar{R}^{\text{WHITTLE}} - \bar{R}^{\text{CSOC}}}$ where \bar{R} is the total discounted reward. Simulation results are averaged over 30 independent trials and run over 40 weeks.

Figure 7 presents simulation of an adversarial example (Mate et al. 2020) consisting of $x\%$ of non-recoverable and $100 - x\%$ of self-correcting beneficiaries for different values of x . Self-correcting beneficiaries tend to miss automated voice messages sporadically, but revert to engaging ways without needing a service call. Non-recoverable beneficiaries are those who may drop out for good, if they stop engaging. We find that in such situations, MYOPIC proves brittle, as it performs even worse than RANDOM while WHITTLE performs well consistently. The actual quality improvement study cohort consists of 48.12% non-recoverable beneficiaries (defined by $P_{01}^p < 0.2$) and the remaining comprised of self-correcting and other types of beneficiaries.

7 Conclusions and Lessons Learned

The widespread use of cell-phones, particularly in the global south, has enabled non-profits to launch massive programs delivering key health messages to a broad population of beneficiaries in a cost-effective manner. We present an RMAB based system to assist these non-profits in optimizing their limited service resources. To the best of our knowledge, ours is the first study to demonstrate the effectiveness of such RMAB-based resource optimization in real-world public health contexts. These encouraging results have initiated the transition of our RMAB software to ARMMAN for real-world deployment. We hope this work paves the way for use of RMABs in many other health service applications.

Some key lessons learned from this research, which complement some of the lessons outlined in (Wilder et al. 2021; Floridi et al. 2020; Tomašev et al. 2020) include the following. First, social-impact driven engagement and design

iterations with the NGOs on the ground is crucial to understanding the right AI model for use and appropriate research challenges. As discussed in footnote 1, our initial effort used a one-shot prediction model, and only after some design iterations we arrived at the current RMAB model. Next, given the missing parameters in RMAB, we found that the assumptions made in literature for learning such parameters did not apply in our domain, exposing new research challenges in RMABs. In short, *domain partnerships with NGOs to achieve real social impact automatically revealed requirements for use of novel application of an AI model (RMAB) and new research problems in this model.*

Second, *data and compute limitations of non-profits are a real world constraint, and must be seen as genuine research challenges in AI for social impact, rather than limitations.* In our domain, one key technical contribution in our RMAB system is deploying clustering methods on offline historical data to infer unknown RMAB parameters. Data is limited as not enough samples are available for any given beneficiary, who may stay in the program for a limited time. Non-profit partners also cannot bear the burden of massive compute requirements. Our clustering approach allows efficient offline mapping to Whittle indices, addressing both data and compute limits, enabling scale-up to service 10s if not 100s of thousands of beneficiaries. Third, *in deploying AI systems for social impact, there are many technical challenges that may not need innovative solutions, but they are critical to deploying solutions at scale.* Indeed, deploying any system in the real world is challenging, but even more so in domains where NGOs may be interacting with low-resource communities. We hope this work serves as a useful example of deploying an AI based system for social impact in partnership with non-profits in the real world and will pave the way for more such solutions with real world impact.

Finally, there are also some important topics for future work in improving the RMAB system, which include handling fairness (Mate, Perrault, and Tambe 2021), changing the current RMAB model with two actions to incorporate multiple actions (Killian et al. 2021), and improving the RMAB model from interactions with beneficiaries (Biswas et al. 2021).

Acknowledgement

We thank Bryan Wilder and Aakriti Kumar for valuable feedback throughout the project and Divy Thakkar and Manoj Sandur Karnik for program management help. We also thank Suresh Chaudhary and Sonali Nandlaskar from the ARMMAN team for helping set up the deployment pipeline. Additionally, we are grateful for support from the ARMMAN staff who made this field study possible. Pradeep Varakantham is supported by the National Research Foundation, Singapore under the AI Singapore Programme (AISG Award No: AISG2-RP-2020-017)

References

- Angrist, J. D.; and Pischke, J.-S. 2008. *Mostly harmless econometrics*. Princeton university press.
- ARMMAN. 2020. mMitra. <https://armman.org/mmitra/>.
- Avrachenkov, K.; and Borkar, V. S. 2020. Whittle index based Q-learning for restless bandits with average reward. *arXiv preprint arXiv:2004.14427*.
- Ayer, T.; Zhang, C.; Bonifante, A.; Spaulding, A. C.; and Chhatwal, J. 2019. Prioritizing hepatitis C treatment in US prisons. *Operations Research*, 67(3): 853–873.
- Biswas, A.; Aggarwal, G.; Varakantham, P.; and Tambe, M. 2021. Learn to Intervene: An Adaptive Learning Policy for Restless Bandits in Application to Preventive Healthcare. In Zhou, Z., ed., *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, 4039–4046. ijcai.org.
- Chen, R.; Santo, K.; Wong, G.; Sohn, W.; Spallek, H.; Chow, C.; and Irving, M. 2021. Mobile Apps for Dental Caries Prevention: Systematic Search and Quality Evaluation. *JMIR mHealth and uHealth*, 9.
- Corotto, P. S.; McCarey, M. M.; Adams, S.; Khazanie, P.; and Whellan, D. J. 2013. Heart failure patient adherence: epidemiology, cause, and treatment. *Heart failure clinics*, 9(1): 49–58.
- Dahiya, K.; and Bhatia, S. 2015. Customer churn analysis in telecom industry. In *2015 4th International Conference on Reliability, Infocom Technologies and Optimization (ICRITO)(Trends and Future Directions)*, 1–6. IEEE.
- Eagle, N.; Macy, M.; and Claxton, R. 2010. Network Diversity and Economic Development. *Science*, 328(5981): 1029–1031.
- Floridi, L.; Cows, J.; King, T.; and Taddeo, M. 2020. How to Design AI for Social Good: Seven Essential Factors. *Science and Engineering Ethics*, 26.
- Gentile, C.; Li, S.; and Zappella, G. 2014. Online clustering of bandits. In *International Conference on Machine Learning*, 757–765. PMLR.
- Gilbert, E. N. 1960. Capacity of a burst-noise channel. *Bell system technical journal*, 39(5): 1253–1265.
- Glazebrook, K. D.; Ruiz-Hernandez, D.; and Kirkbride, C. 2006. Some indexable families of restless bandit problems. *Advances in Applied Probability*, 38(3): 643–672.
- Hadad, V.; Hirshberg, D. A.; Zhan, R.; Wager, S.; and Athey, S. 2021. Confidence intervals for policy evaluation in adaptive experiments. *Proceedings of the National Academy of Sciences*, 118(15): e2014602118.
- HelpMum. 2021. PREVENTING MATERNAL AND INFANT MORTALITY IN NIGERIA. <https://helpmum.org/>.
- Janson, L. 2022. Private Communication.
- Johnson, J. . 2017. MomConnect: Connecting Women to Care, One Text at a Time. <https://www.jnj.com/our-giving/momconnect-connecting-women-to-care-one-text-at-a-time>.
- Jung, Y. H.; and Tewari, A. 2019. Regret bounds for thompson sampling in episodic restless bandit problems. *Advances in Neural Information Processing Systems*.

- Kaur, J.; Kaur, M.; Chakrapani, V.; Webster, J.; Santos, J.; and Kumar, R. 2020. Effectiveness of information technology-enabled ‘SMART Eating’ health promotion intervention: A cluster randomized controlled trial. *PLOS ONE*, 15: e0225892.
- Killian, J. A.; Biswas, A.; Shah, S.; and Tambe, M. 2021. Q-Learning Lagrange Policies for Multi-Action Restless Bandits. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 871–881.
- Killian, J. A.; Wilder, B.; Sharma, A.; Choudhary, V.; Dilkina, B.; and Tambe, M. 2019. Learning to Prescribe Interventions for Tuberculosis Patients Using Digital Adherence Data. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- Leskovec, J.; Adamic, L. A.; and H., B. A. 2007. The Dynamics of Viral Marketing. *ACM Trans. Web*, 1(1): 5–es.
- Li, C.; Wu, Q.; and Wang, H. 2021. Unifying Clustered and Non-stationary Bandits. In *International Conference on Artificial Intelligence and Statistics*, 1063–1071. PMLR.
- Li, S.; Chen, W.; and Leung, K.-S. 2019. Improved algorithm on online clustering of bandits. *arXiv preprint arXiv:1902.09162*.
- Liao, P.; Greenewald, K.; Klasnja, P.; and Murphy, S. 2020. Personalized heartsteps: A reinforcement learning algorithm for optimizing physical activity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(1): 1–22.
- Martin, L. R.; Williams, S. L.; Haskard, K. B.; and DiMatteo, M. R. 2005. The challenge of patient adherence. *Therapeutics and clinical risk management*, 1(3): 189.
- Mate, A.; Killian, J.; Xu, H.; Perrault, A.; and Tambe, M. 2020. Collapsing Bandits and Their Application to Public Health Intervention. *Advances in Neural Information Processing Systems*, 34.
- Mate, A.; Perrault, A.; and Tambe, M. 2021. Risk-Aware Interventions in Public Health: Planning with Restless Multi-Armed Bandits. *Autonomous Agents and Multi-Agent Systems (AAMAS)*.
- Mintz, Y.; Aswani, A.; K., P.; Flowers, E.; and Fukuoka, Y. 2020. Nonstationary bandits with habituation and recovery dynamics. *Operations Research*, 68(5): 1493–1516.
- Pfammatter, A.; Spring, B.; Saligram, N.; Davé, R.; Gowda, A.; Blais, L.; Arora, M.; Ranjani, H.; Ganda, O.; Hedeker, D.; Reddy, S.; and Ramalingam, S. 2016. mHealth Intervention to Improve Diabetes Risk Behaviors in India: A Prospective, Parallel Group Cohort Study. *Journal of Medical Internet Research*, 18: e207.
- Pilote, L.; Tulskey, J. P.; Zolopa, A. R.; Hahn, J. A.; Schechter, G. F.; and Moss, A. R. 1996. Tuberculosis Prophylaxis in the Homeless: A Trial to Improve Adherence to Referral. *Archives of Internal Medicine*, 156(2): 161–165.
- Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. Wiley. ISBN 978-0-47161977-2.
- Qian, Y.; Zhang, C.; Krishnamachari, B.; and Tambe, M. 2016. Restless Poachers: Handling Exploration-Exploitation Tradeoffs in Security Domains. In Jonker, C. M.; Marsella, S.; Thangarajah, J.; and Tuyls, K., eds., *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016*, 123–131. ACM.
- Shaaban, E.; Helmy, Y.; Khedr, A.; and Nasr, M. 2012. A proposed churn prediction model. *International Journal of Engineering Research and Applications*, 2(4): 693–697.
- Sombabu, B.; Mate, A.; Manjunath, D.; and Moharir, S. 2020. Whittle Index for AoI-Aware Scheduling. In *2020 12th International Conference on Communication Systems & Networks*. IEEE.
- Son, Y.-J.; Kim, H.-G.; Kim, E.-H.; Choi, S.; and Lee, S.-K. 2010. Application of support vector machine for prediction of medication adherence in heart failure patients. *Healthcare informatics research*, 16(4): 253–259.
- Thirumurthy, H.; and Lester, R. T. 2012. M-health for health behaviour change in resource-limited settings: applications to HIV care and beyond. *Bulletin of the World Health Organization*, 90: 390–392.
- Tomašev, N.; Cornebise, J.; Hutter, F.; Mohamed, S.; Picciariello, A.; Connelly, B.; Belgrave, D.; Ezer, D.; Haert, F.; Mugisha, F.; Abila, G.; Arai, H.; Almiraat, H.; Proskurnia, J.; Snyder, K.; Otake, M.; Othman, M.; Glasmachers, T.; Wever, W.; and Clopath, C. 2020. AI for social good: unlocking the opportunity for positive impact. *Nature Communications*, 11: 2468.
- Tuldrà, A.; Ferrer, M. J.; Fumaz, C. R.; Bayés, R.; Paredes, R.; Burger, D. M.; and Clotet, B. 1999. Monitoring Adherence to HIV Therapy. *Archives of Internal Medicine*, 159(12): 1376–1377.
- Weber, R. R.; and Weiss, G. 1990. On an index policy for restless bandits. *Journal of applied probability*, 637–648.
- Whittle, P. 1988. Restless bandits: Activity allocation in a changing world. *Journal of applied probability*, 287–298.
- Wilder, B.; Onasch-Vera, L.; Diguseppi, G.; Petering, R.; Hill, C.; Yadav, A.; Rice, E.; and Tambe, M. 2021. Clinical Trial of an AI-Augmented Intervention for HIV Prevention in Youth Experiencing Homelessness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 14948–14956.
- Xie, Y.; Li, X.; Ngai, E.; and Ying, W. 2009. Customer churn prediction using improved balanced random forests. *Expert Systems with Applications*, 36(3, Part 1): 5445 – 5449.
- Yang, L.; Liu, B.; Lin, L.; Xia, F.; Chen, K.; and Yang, Q. 2020. Exploring Clustering of Bandits for Online Recommendation System. In *Fourteenth ACM Conference on Recommender Systems*, 120–129.
- Zhang, K. W.; Janson, L.; and Murphy, S. A. 2022. Statistical Inference After Adaptive Sampling in Non-Markovian Environments. *arXiv preprint arXiv:2202.07098*.
- Zhou, M.; Mintz, Y.; Fukuoka, Y.; Goldberg, K.; Flowers, E.; Kaminsky, P.; Castillejo, A.; and Aswani, A. 2018. Personalizing mobile fitness apps using reinforcement learning. In *CEUR workshop proceedings*, volume 2068. NIH Public Access.

A Clarification on Statistical Analysis

It has been brought to our attention that the standard statistical test employed in the paper, while being popular, is still imperfect for the unique nature of the adaptively collected data analyzed in our study. Designing a valid statistical test for concluding from adaptively collected data from trials such as ours, is still an area of modern research (Hadad et al. 2021; Zhang, Janson, and Murphy 2022).

The unique challenge is that the outcomes of individual beneficiaries within an experimental arm, are not independent owing to the budget constraint. In fact, we expect the outcomes to be negatively correlated under the assumption that interventions are never detrimental — reason being that while allocating an intervention to a beneficiary tends to improve their outcomes, it robs some other beneficiary of the intervention resource, leading to a decline in their expected outcome. This negative correlation is expected to translate to a lower sample variation in the measured outcomes than what is accounted for under the independent-outcomes assumption. Our discussion (Janson 2022) suggests that this is likely to result in a less heavy tail than expected, and consequently, our computation should likely yield conservative p-values. We justify this intuition further via an illustrative example involving a simpler setting below.

Because the computed p-values are conservative, the conclusions drawn in our study should still remain valid.

Justification of why negative correlation leads to p-value inflation:

We provide intuition through an example, for why negative correlation leads to inflated p-values.

Consider a simplified linear model, similar to the one employed in the paper, without contributions from intermediate terms such as k, T_i, x_{ij} , retaining only β, ϵ . Consider the null hypothesis $H_0 := \beta = 0$. To reject this hypothesis, consider employing the z-test that would construct a z-statistic as: $z\text{-stat} = \frac{\sqrt{n} \cdot \bar{Y}}{\sigma}$. If the Y_i 's were assumed to be independent, we would obtain the p-value as α in terms of the tail distribution Z of a standard normal distribution as:

$$\left| \frac{\sqrt{n} \cdot \bar{Y}}{\sigma} \right| = Z_{1-\frac{\alpha}{2}}$$

However, in reality, since the Y_i 's are not independent, but in fact, negatively correlated, the actual variance in \bar{Y} is smaller. This is because the second term in the expansion of the variance of the mean outcomes, contributes negatively as in:

$$\text{var}(\bar{Y}) = \frac{1}{n^2} \left[\sum \text{var}(Y_i) + \sum \text{cov}(Y_i, Y_j) \right]$$

This leads to σ^2 being an inflated version of the actual variance and equivalently, this leads to an under-inflated value of $\left| \frac{\sqrt{n} \cdot \bar{Y}}{\sigma} \right|$. This translates to a larger mass on the right tail of the Z distribution, leading to a conservative (i.e. larger) p-value.