Learning to Call: A Field Trial of a Collaborative Bandit Algorithm for Improved Message Delivery in Mobile Maternal Health

Arpan Dasgupta*
Google Deepmind
India
arpandg@google.com

Mizhaan Maniyar*
Google Deepmind
India
mizhaan@google.com

Awadhesh Srivastava
ARMMAN
India
awadhesh@armman.org

Sanat Kumar
ARMMAN
India
sanat@armman.org

Amrita Mahale
ARMMAN
India
amrita@armman.org

Aparna Hedge
ARMMAN
India
aparnahedge@armman.org

Arun Suggala
Google Deepmind
India
arunss@google.com

Karthikeyan Shanmugam
Google Deepmind
India
karthikeyanvs@google.com

Aparna Taneja
Google Deepmind
India
arparnataneja@google.com

Milind Tambe
Google Deepmind
India
milindtambe@google.com

ABSTRACT

Mobile health (mHealth) programs ([23], [10], [17]) utilize automated voice messages to deliver health information, particularly targeting underserved communities, demonstrating the effectiveness of using mobile technology to disseminate crucial health information to these populations, improving health outcomes through increased awareness and behavioral change. India's Kilkari program delivers vital maternal health information via weekly voice calls to millions of mothers. However, the current random call scheduling often results in missed calls and reduced message delivery. This study presents a field trial of a collaborative bandit algorithm designed to optimize call timing by learning individual mothers' preferred call times. We deployed the algorithm with $\sim 8,700$ Kilkari participants as a pilot study, comparing its performance to the baseline random calling approach. Our results demonstrate a statistically significant improvement in call pickup rates with the bandit algorithm, indicating its potential to enhance message delivery and impact millions of mothers across India. This research highlights the efficacy of personalized scheduling in mobile health interventions and underscores the potential of machine learning to improve maternal health outreach at scale.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

KEYWORDS

mHealth, AI for social good, Collaborative bandits

ACM Reference Format:

Arpan Dasgupta*, Mizhaan Maniyar*, Awadhesh Srivastava, Sanat Kumar, Amrita Mahale, Aparna Hedge, Arun Suggala, Karthikeyan Shanmugam, Aparna Taneja, and Milind Tambe. 2025. Learning to Call: A Field Trial of a Collaborative Bandit Algorithm for Improved Message Delivery in Mobile Maternal Health. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 7 pages.

1 INTRODUCTION

Maternal health remains a critical public health concern in India, with millions of women with limited access to timely and accurate information during pregnancy and postpartum. Recognizing this need, the Government of India launched the Kilkari program, a nationwide mobile health initiative that delivers weekly voice messages that contain essential maternal health information to more than 10 million registered mothers [5]. mHealth programs such as these play a vital role in reducing maternal mortality rates - a key target within the WHO's Sustainable Development Goals [2]. These messages cover vital topics such as iron and calcium supplementation, antenatal care, and postnatal practices, aiming to improve maternal health outcomes throughout the country.

However, the effectiveness of this large-scale program is contingent upon successful message delivery. Currently, Kilkari employs a random call scheduling strategy, attempting to reach mothers, with up to nine re-attempts (until the call is picked up), but without considering individual preferences for call timing. This approach often results in missed calls, crucial bandwidth spent on re-attempts and most importantly limiting the reach and impact of crucial health information [16, 18, 22]. To address this challenge, this paper presents a field trial of the collaborative bandit algorithm [28] designed to optimize call scheduling by learning mothers' preferred call times.

Collaborative bandit algorithms offer a promising approach for personalized intervention delivery in mobile health. By iteratively learning from user responses and interactions, these algorithms can adapt to individual preferences and maximize engagement. In this study, we implemented a collaborative bandit algorithm within the Kilkari platform and conducted a field trial involving approximately 8,700 mothers. Our goal was to evaluate the algorithm's ability to improve call pickup rates compared to the baseline random calling strategy.

This research contributes to the growing body of literature in the application of machine learning in mobile health interventions [21, 25, 31]. By demonstrating the effectiveness of a collaborative bandit algorithm in a real-world setting, we highlight the potential for personalized call scheduling to enhance the reach and impact of maternal health programs at scale. Given the national scope of Kilkari and the potential for improved message delivery to millions of mothers, our findings have significant implications for public health policy and practice in India and beyond.

Key note on the experiments reported in this paper: This work was conducted as a joint effort between a research team from ARMMAN [1], a non-profit organization in India, and Google DeepMind, as reflected in the co-authorship of this paper. It is crucial to highlight that the beneficiary data utilized in this research is fully anonymized, and no socio-demographic features were available to the research team. To ensure data privacy and security, the experimental infrastructure was managed exclusively by the ARMMAN team, who were the only individuals with access to the raw beneficiary data. The collaborative bandit algorithm, previously developed by Google researchers, was implemented under their guidance by the ARMMAN team. The Google researchers contributed by advising the ARMMAN team on the algorithm's implementation and subsequently collaborating on the analysis of the resulting study.

2 RELATED WORK

mHealth programs provide essential health information though automated voice messages to a large number of beneficiaries [14, 24], which implies any improvement to the program positively affects a lot of mothers. Previously, AI has been applied to schedule interventions [21, 31] and showed positive behavioral outcomes [12]. LeFevre et al. [20] talk about the protocol for an individually controlled randomized control trial in an attempt to show the effectiveness of Kilkari.

Scheduling the time of the day the beneficiaries are called using collaborative bandits [28] showed promise in simulation and this work aims to test it out in a pilot study. The analysis in [7] shows that there are preferred slots for calling beneficiaries. They further show that most calls which are picked are done so by the third attempt. These factors point towards the advantages of scheduling these calls in a non-random manner.

Multi-armed bandits represent a well-researched and potent approach for tackling diverse resource allocation challenges. Numerous methodologies, including phased elimination [19, 29], Upper Confidence Bound (UCB) [6], Thompson Sampling [3, 30], and Best-arm Identification [4, 13], have undergone thorough investigation. The collaborative bandit problem has witnessed a surge in interest recently, driven by the widespread adoption of recommender systems [8, 11]. Under specific conditions, several algorithms with robust theoretical guarantees have been developed [15, 26]. An algorithm suited for scenarios with approximate low-rank structure, was introduced by [28] and is evaluated in this field study.

3 BACKGROUND

3.1 Kilkari

Kilkari [5] is the world's largest mobile health program focused on maternal and child health. It is conducted by India's Ministry of Health and Family Welfare in partnership with the NGO ARMMAN. Kilkari uses pre-recorded voice calls to deliver vital preventive care information on maternal and infant health to pregnant women and new mothers. The program aims to improve access to healthcare information for pregnant women, mothers of infants, and their families, particularly in underserved communities.

However, these programs face challenges, including limited beneficiary phone access and unknown time preferences, which hinder timely outreach and lead to poor engagement. Specifically, low listenership of the automated voice messages is a major challenge. Even with multiple call attempts, approximately 23% of beneficiaries are not reached. Consistent low listenership can even lead to beneficiaries being dropped from the program, which can occur if beneficiaries listen to less than 25% of the messages for six weeks in a row.

To address the challenge of low listenership, the use of a stochastic bandit approach could be very useful to learn the favored time slot of individual mothers/beneficiaries. This is important because factors such as limited phone access, working hours, and household responsibilities significantly affect the likelihood of answering a call at a given time slot. By quickly identifying good time slots for each beneficiary, engagement with the calls can be improved, and beneficiaries can be retained in the program. Furthermore, optimizing the time slot to send automated voice messages can help to reduce automatic dropouts and save bandwidth.

3.2 Collaborative Bandits

Pal et al. [28] address the challenge of optimizing time slot selection in mobile health programs like Kilkari, where the goal is to deliver automated voice messages to beneficiaries at times they are most likely to engage. To tackle this, [28] formulate the problem as a multi-agent multi-armed bandit problem [29]. Here, each beneficiary is modeled as an agent, and each possible time slot for delivering the message is considered an arm. The key idea is to learn the preferences of these agents (beneficiaries) for different arms (time slots) through repeated interactions (i.e., attempting to deliver messages).

To efficiently solve this multi-agent bandit problem, the underlying algorithm in [26, 27] is applied to the current problem. This framework leverages the assumption that the preferences of beneficiaries are not entirely independent but rather share some underlying structure. Specifically, it assumes that the matrix representing beneficiaries' preferences (e.g., the probability of a beneficiary picking up a call at a given time slot) is approximately low-rank. This low-rank assumption implies that there are a few latent factors that explain a significant portion of the variability in beneficiaries' time slot preferences. The collaborative bandit algorithm exploits this low-rank structure to learn more efficiently by sharing information across beneficiaries, rather than learning each beneficiary's preferences in isolation.

Pal et al. [28] introduces two novel algorithms: Greedy Matrix Completion (MC) and Phased MC. Phased MC is the key algorithm we use in this work. Greedy MC first has a long random phase where arms are picked randomly, followed by prediction which is followed thereafter. Phased MC operates by updating the estimates in "phases" which implies that no lengthy exploration phase is required to obtain an initial estimate. To allow for exploration during prediction, a Boltzmann noise [9] is added. It also uses variance reduction techniques to improve robustness to noise. Since the exploration phase required in this algorithm is not as large, it hence prevents chances of dropout early in the trial due to random calls.

4 EXPERIMENTAL DESIGN

This field trial was conducted in the Kalahandi and Puri districts of Odisha, India (upon the guidance of the NGO), to evaluate the effectiveness of the collaborative bandit algorithm in improving call pickup rates within the Kilkari maternal health program.

4.1 Randomization

Beneficiaries were randomly assigned to either the Random or Treatment group to minimize bias and ensure comparability between the groups.

Random Group (Control): This group received calls using the current Kilkari random call scheduling approach. This group initially consisted of 8694 beneficiaries.

Treatment Group (Collaborative Bandit): This group received calls scheduled using a collaborative bandit algorithm designed to learn and adapt to individual beneficiaries' preferred call times. This group initially consisted of 8724 beneficiaries.

4.2 Trial Phases

The trial consisted of two distinct phases:

(i) **Baseline Phase (Weeks 1-3) [7th January - 26th January, 2025]**: Both the Random and Treatment groups received calls using the standard Kilkari random calling strategy. This phase served to establish a baseline for call pickup rates and to collect data for the collaborative bandit algorithm in the Treatment group to initiate learning beneficiaries' preferences. The exact number of calls attempted for each beneficiary during this phase will be detailed in Section 6.

(ii) **Intervention Phase (Weeks 4-5) [27th January - 9th February 2025]**: The Random group continued to receive calls using the random calling strategy. The Treatment group, however, received calls scheduled based on the preferences learned by the collaborative bandit algorithm during the Baseline Phase in an iterative manner. The exact number of calls attempted for each beneficiary during this phase will be detailed in Section 6.

4.3 Data Collection

Call logs were collected for all beneficiaries in both groups, recording the date, time, and outcome (answered/missed) of each call attempt. The total number of beneficiaries in each group will be noted in the Section 6.

4.4 Outcome Measure

The primary outcome measure was the call pickup rate, defined as the proportion of successful call pickups out of the total number of call attempts, for each group during the Intervention Phase.

Ethical Considerations: No ethical approvals were required for this study as it was deployed on an existing program and counts as a program improvement.

Statistical Analysis: Statistical analysis was conducted to compare the call pickup rates between the Random and Treatment groups during the Intervention Phase. We used a simple two-sample t-statistic to verify the statistical significance between the two groups, across the baseline and intervention phase. More information can be found in the next section.

5 PRELIMINARIES

In this section we mathematically define the call pickup-rates and its variants that are used for the analysis in the later section. The index i represents a beneficiary (or user), $j \in [1, 7]$ being one of the seven time slot IDs chosen, t being the day, and $r \in [0, 2]$ being the re-attempt number for that slot, i.e. $r = 0$ being the first call, and $r = 1$ being the second call made, if the first call wasn't picked up. Let $call$ be mapped uniquely to the tuple (i, j, t, r). Let $A_{call} \equiv A_{i,j,t,r} \in \{0, 1\}$ denote whether a call attempt was made for user i during time-slot j on day t and whether it was the r-th re-attempt. Similarly, let $p_{call} \equiv p_{i,j,t,r} \in \{0, 1\}$ denote whether an attempted call was picked or not. We assume that the set of calls $\{call | p_{call} = 1\} \subseteq \{call | A_{call} = 1\}$. We now define the pooled pickup-rate as,

$$PR_{pooled} = \frac{\sum_{\forall i,j,t,r} p_{i,j,t,r}}{\sum_{\forall i,j,t,r} A_{i,j,t,r}} \equiv \frac{\sum_{\forall call} p_{call}}{\sum_{\forall call} A_{call}}. \tag{1}$$

Alternatively, we can define a user-specific pickup-rate (PR) and its average as such

$$PR_i = \frac{\sum_{\forall j,t,r} p_{i,j,t,r}}{\sum_{\forall j,t,r} A_{i,j,t,r}}, \qquad PR_{user} = \frac{\sum_{\forall i} PR_i}{\sum_{\forall i} 1}. \tag{2}$$

Table 1: Pooled call pickup-rates across calls made to only those users who didn't drop out in the intervention phase, i.e. active users.

Group	PR^{active}_{pooled} (baseline)	PR^{active}_{pooled} (intervention)	% reduction	p-value
Treatment	0.470	0.463	**-1.52**%	0.0849
Control	0.465	0.448	-3.62 %	4.65e-05
p-value	0.1345	0.0006	-	-

The metric PR_i can be seen as an estimate of the probability of user i picking up a call.

Two-sample t-test

We use this method in our analysis in Section 6 to determine whether there is a significant difference between the means of two independent groups. Given two samples $X_1, X_2, \ldots, X_{n_1}$ and $Y_1, Y_2, \ldots, Y_{n_2}$, drawn from normal distributions with means μ_X and μ_Y, we test the null hypothesis, $H_0 : \mu_X = \mu_Y$ against the alternative hypothesis, $H_A : \mu_X \neq \mu_Y$.

Test Statistic. The test statistic for the two-sample t-test is given by:

$$t = \frac{\bar{X} - \bar{Y}}{S_p\sqrt{\frac{1}{n_1} + \frac{1}{n_2}}}, \tag{3}$$

where $\bar{X}$ and $\bar{Y}$ are the sample means, and S_p is the pooled standard deviation, computed as:

$$S_p = \sqrt{\frac{(n_1 - 1)S_X^2 + (n_2 - 1)S_Y^2}{n_1 + n_2 - 2}}, \tag{4}$$

where S_X^2 and S_Y^2 are the sample variances.

P-Value Calculation. The p-value is obtained by comparing the observed t-statistic to the critical values of the Student's t-distribution with n_1+n_2-2 degrees of freedom, $p = 2P(T > |t|)$ where T follows a t-distribution under the null hypothesis. Therefore, a small p-value (typically $p < 0.05$) suggests strong evidence against the null hypothesis, leading to the conclusion that the two population means differ significantly. This test is widely used in various scientific fields, including medicine, psychology, and economics.

6 RESULTS

In this section we analyse the pickup-rates obtained for the two groups, i.e. treatment and control during the two phases. We present the pooled pickup-rate values obtain from 1 in Table 1. We notice that during the baseline phase the performance was similar across the two groups using a t-test. Furthermore, the reduction in performance between the two phases for the treatment group is of lesser significance (p-value > 0.05) as compared to the control group ((p-value << 0.05). Furthermore, the performance difference in the treatment group is much better and significant as compared to the control group in the intervention phase.

We now dive deeper in analysing the difference that arose in the intervention phase, i.e. from the 27th of January to 9th February, comparing the pickup-rates in the treatment group with the control group. In order to remove outliers, we segregate beneficiaries with very high PR_i, i.e. those who always pickup their calls and very low PR_i, i.e. those who never pickup their calls. For the treatment group, we have 40.59% users with a $PR_i = 1$ and 6.56% users with $PR_i = 0$. While the control group has values of 38.46% and 6.99% respectively. In order to maintain a fair comparison, we remove the same fraction of users from both these groups, i.e. removing the top 40.59% ($\max\{40.59\%, 38.46\%\}$) and bottom 6.99% ($\max\{6.56\%, 6.99\%\}$) from both the groups according to their pickup-rate probability, i.e. PR_i obtained via 2. We call these tiers, High Tier, Mid Tier and Low Tier, respectively, emphasising on the Mid Tier for most of the analysis results.

6.1 Call Volumes

Table 2 summarizes the number of calls made to each arm within each tier for the intervention phase only.

Table 2: Call Volumes by Tier and Arm

Tier	Treatment	Control
High	5077	5222
Mid	16775	17345
Low	2789	2542

6.2 Call Pickup Rates

Table 3 presents the call pickup rates for each arm within each tier, along with the corresponding p-values for statistical significance. We use a 2-sample t-test for the two arms in each of the 3 tiers and obtain the p-value according to the methodology mentioned in the appendix (add ref).

Table 3: Call Pickup Rates by Tier and Arm

Tier	Treatment	Control	% improvement	p-value
High	1.0000	0.9732	2.75%	4.66e-32
Mid	0.3763	0.3555	5.83%	7.07e-05
Low	0.0100	0.0000	NaN	3.98e-07

High Tier. By construction, the treatment group will have all its users with a $PR_i = 1, \forall i$, while the control group having an average, i.e $PR_{user} < 1$ representing the mean for this tier only.

Mid Tier. The treatment group achieved a call pickup rate of 0.3763 (37.63%), while the control group achieved a rate of 0.3555 (35.55%). This difference amounts to a 5.83% improvement, and it was statistically significant (p = 7.07e-05), demonstrating that the collaborative bandit algorithm significantly improved call pickup rates for beneficiaries in the middle tier.

Low Tier. For the Low Tier, the treatment group had a call pickup rate of 0.01 (1.00%), and the control group had a rate of 0.0000, as expected by construction, similar with the High Tier.

6.3 Time slot wise analysis

In order to see which time slots saw the most improvement, we analyse the calls made for the mid tier group (Tier 2) in Table 4. However for each time slot, we use the following formulae for Tables 4,

$$PR^j_{pooled} = \frac{\sum_{\forall i,t,r} p_{i,j,t,r}}{\sum_{\forall i,t,r} A_{i,j,t,r}}. \tag{5}$$

Table 4: Pooled call pickup-rates PR^j_{pooled} across all calls made in the respective time slot j given by 5.

Time Slot ID	Treatment	Control	% pickup-rate	p-value
1	0.3584	0.3337	7.4104	**0.0563**
2	0.3510	0.3365	4.2896	0.2695
3	0.3908	0.3625	7.8111	**0.0438**
4	0.3841	0.3683	4.2734	0.2609
5	0.3753	0.3686	1.8228	0.6385
6	0.3598	0.3223	11.6131	**0.0060**
7	0.4197	0.4121	1.8303	0.6115

We see a positive improvement in pickup-rate across all time slots and a huge improvement of 11.61% in time slot id 6 especially. The p-values are significant for the slots with $> 5\%$ improvement at the 0.10 level.

6.4 Summary

The tiered analysis reveals that the collaborative bandit algorithm (the treatment group) significantly improved call pickup rates compared to the random calling strategy (random control group), particularly for beneficiaries in the middle and bottom tiers. This demonstrates the algorithm's effectiveness in optimizing call scheduling and enhancing message delivery within the Kilkari program.

REFERENCES

[1] [n.d.]. Armman Home - ARMMAN - Helping Mothers and Children — armman.org. https://armman.org. [Accessed 31-05-2024].

[2] [n.d.]. SDG Target 3.1 Maternal mortality — who.int. https://www.who.int/data/gho/data/themes/topics/sdg-target-3-1-maternal-mortality. [Accessed 31-05-2024].

[3] Shipra Agrawal and Navin Goyal. 2012. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on learning theory*. JMLR Workshop and Conference Proceedings, 39–1.

[4] Shubhada Agrawal, Sandeep Juneja, and Peter Glynn. 2020. Optimal δ-Correct Best-Arm Selection for Heavy-Tailed Distributions. In *Algorithmic Learning Theory*. PMLR, 61–110.

[5] ARMMAN. 2023. Kilkari. https://armman.org/kilkari/

[6] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47 (2002), 235–256.

[7] Jean Juste Harrisson Bashingwa, Diwakar Mohan, Sara Chamberlain, Salil Arora, Jai Mendiratta, Sai Rahul, Vinod Chauhan, Kerry Scott, Neha Shah, Osama Ummer, et al. 2021. Assessing exposure to Kilkari: a big data analysis of a large maternal mobile messaging service across 13 states in India. *BMJ global health* 6, Suppl 5 (2021), e005213.

[8] Guy Bresler, Devavrat Shah, and Luis Filipe Voloch. 2016. Collaborative filtering with low regret. In *Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science*. 207–220.

[9] Nicolò Cesa-Bianchi, Claudio Gentile, Gábor Lugosi, and Gergely Neu. 2017. Boltzmann exploration done right. *Advances in neural information processing systems* 30 (2017).

[10] Mahbub Elahi Chowdhury, Shafayatul Islam Shiblee, and Heidi E. Jones. 2019. Does mHealth voice messaging work for improving knowledge and practice of maternal and newborn healthcare? *BMC Medical Informatics and Decision Making* 19, 1 (05 Sep 2019), 179. https://doi.org/10.1186/s12911-019-0903-z

[11] Hamid Dadkhahi and Sahand Negahban. 2018. Alternating linear bandits for online matrix-factorization recommendation. *arXiv preprint arXiv:1810.09401* (2018).

[12] Arpan Dasgupta, Niclas Boehmer, Neha Madhiwalla, Aparna Hedge, Bryan Wilder, Milind Tambe, and Aparna Taneja. 2024. Preliminary Study of the Impact of AI-Based Interventions on Health and Behavioral Outcomes in Maternal Health Programs. *arXiv preprint arXiv:2407.11973* (2024).

[13] Aurélien Garivier and Emilie Kaufmann. 2016. Optimal best arm identification with fixed confidence. In *Conference on Learning Theory*. PMLR, 998–1027.

[14] Aparna Hegde and Riddhi Doshi. 2016. Assessing the Impact of Mobile-based Intervention on Health Literacy among Pregnant Women in Urban India. In *American Medical Informatics Association Annual Symposium*. 1423.

[15] Prateek Jain and Soumyabrata Pal. 2022. Online low rank matrix completion. *arXiv preprint arXiv:2209.03997* (2022).

[16] Bashingwa JJH, Mohan D, Chamberlain S, Arora S, Mendiratta J, Rahul S, Chauhan V, Scott K, Shah N, Ummer O, Ved R, Mulder N, and LeFevre AE. 2021. Assessing exposure to Kilkari: a big data analysis of a large maternal mobile messaging service across 13 states in India. *BMJ Glob Health* (2021).

[17] Johnblack K Kabukye, Onaedo Ilozumba, Jacqueline E W Broerse, Nicolette de Keizer, and Ronald Cornet. 2021. Implementation of an interactive voice response system for cancer awareness in Uganda: Mixed methods study. *JMIR MHealth UHealth* 9, 1 (Jan. 2021), e22061.

[18] Arshika Lalan, Shresth Verma, Kumar Madhu Sudan, Amrita Mahale, Aparna Hegde, Milind Tambe, and Aparna Taneja. 2023. Analyzing and Predicting Low-Listenership Trends in a Large-Scale Mobile Health Program: A Preliminary Investigation. *arXiv preprint arXiv:2311.07139* (2023).

[19] Tor Lattimore and Csaba Szepesvári. 2020. *Bandit algorithms*. Cambridge University Press.

[20] Amnesty LeFevre, Smisha Agarwal, Sara Chamberlain, Kerry Scott, Anna Godfrey, Rakesh Chandra, Aditya Singh, Neha Shah, Diva Dhar, Alain Labrique, et al. 2019. Are stage-based health information messages effective and good value for money in improving maternal newborn and child health outcomes in India? protocol for an individually randomized controlled trial. *Trials* 20 (2019), 1–12.

[21] Aditya Mate, Lovish Madaan, Aparna Taneja, Neha Madhiwalla, Shresth Verma, Gargi Singh, Aparna Hegde, Pradeep Varakantham, and Milind Tambe. 2022. Field study in deploying restless multi-armed bandits: Assisting non-profits in improving maternal and child health. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 12017–12025.

[22] Diwakar Mohan, Kerry Scott, Neha Shah, Jean Juste Harrisson Bashingwa, Arpita Chakraborty, Osama Ummer, Anna Godfrey, Priyanka Dutt, Sara Chamberlain, and Amnesty Elizabeth LeFevre. 2021. Can health information through mobile phones close the divide in health behaviours among the marginalised? An equity analysis of Kilkari in Madhya Pradesh, India. *BMJ Global Health* 6, Suppl 5 (2021), e005512.

[23] Nirmala Murthy, Subhashini Chandrasekharan, Muthu Perumal Prakash, Aakash Ganju, Joanne Peter, Nadi Kaonga, and Patricia Mechael. 2020. Effects of an mHealth voice message service (mMitra) on maternal health knowledge and practices of low-income women in India: findings from a pseudo-randomized controlled trial. *BMC Public Health* 20, 1 (01 Jun 2020), 820. https://doi.org/10.1186/s12889-020-08965-2

[24] Nirmala Murthy, Subhashini Chandrasekharan, Muthu Perumal Prakash, Aakash Ganju, Joanne Peter, Nadi Kaonga, and Patricia Mechael. 2020. Effects of an mHealth voice message service (mMitra) on maternal health knowledge and practices of low-income women in India: findings from a pseudo-randomized controlled trial. *BMC Public Health* 20, 1 (2020), 1–10.

[25] Vineet Nair, Kritika Prakash, Michael Wilbur, Aparna Taneja, Corinne Namblard, Oyindamola Adeyemo, Abhishek Dubey, Abiodun Adereni, Milind Tambe, and Ayan Mukhopadhyay. 2022. Adviser: Ai-driven vaccination intervention optimiser for increasing vaccine uptake in nigeria. *arXiv preprint arXiv:2204.13663* (2022).

[26] Soumyabrata Pal, Arun Sai Suggala, Karthikeyan Shanmugam, and Prateek Jain. 2023. Optimal algorithms for latent bandits with cluster structure. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 7540–7577.

[27] Soumyabrata Pal, Arun Sai Suggala, Karthikeyan Shanmugam, and Prateek Jain. 2023. Optimal Algorithms for Latent Bandits with Cluster Structure. arXiv:2301.07040 [cs.LG] https://arxiv.org/abs/2301.07040

[28] Soumyabrata Pal, Milind Tambe, Arun Suggala, and Aparna Taneja. 2024. Improving Mobile Maternal and Child Health Care Programs: Collaborative Bandits for Time slot selection. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*. 1540–1548.

[29] Aleksandrs Slivkins et al. 2019. Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning* 12, 1-2 (2019), 1–286.
[30] William R Thompson. 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25, 3-4 (1933), 285–294.
[31] Shresth Verma, Gargi Singh, Aditya S Mate, Paritosh Verma, Sruthi Gorantala, Neha Madhiwalla, Aparna Hegde, Divy Hasmukhbhai Thakkar, Manish Jain, Milind Shashikant Tambe, et al. 2023. Deployed SAHELI: Field Optimization of Intelligent RMAB for Maternal and Child Care. In *Innovative Applications of Artificial Intelligence (IAAI)*.

APPENDIX

A CALL DISTRIBUTIONS

In this section, we will see the call distributions and how they changed during the different phases especially for the treatment group (collaborative bandits algorithm). The calls were also made according to a 3-2-2-2 pattern, i.e. if a user doesn't pickup the first time we call them 2 more times with an interval of 10-20 minutes that day. If they still don't pickup up, we call them twice the next day in the slot recommended by the algorithm for that day – the slot with the highest pickup-rate probability – and twice the day after etc. until a call has been picked. If a call was picked in any one of these 9 attempts, the next call is made a calendar week after the first call. We use the following method to obtain the call recommendation distribution $\pi(j)$ for time-slot j for both the groups in the intervention phase,

$$\texttt{count}(j) = \sum_{\forall i,t} A_{i,j,t,r=0}, \qquad \pi(j) = \frac{\texttt{count}(j)}{\sum_{j=1}^{7} \texttt{count}(j)}.$$

In Figure 1, we can see how the call recommendation distribution of the treatment group deviates from that of the control group indicating that our algorithm is potentially finding the right time slots to call at.

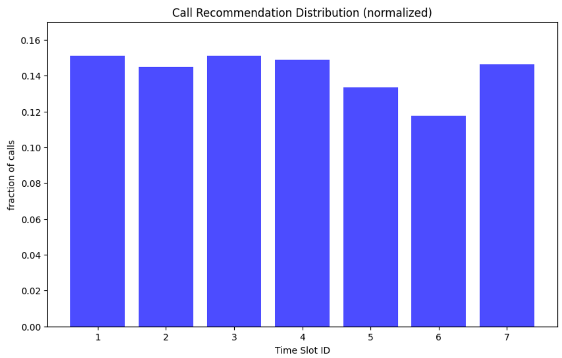

(a) Treatment Group

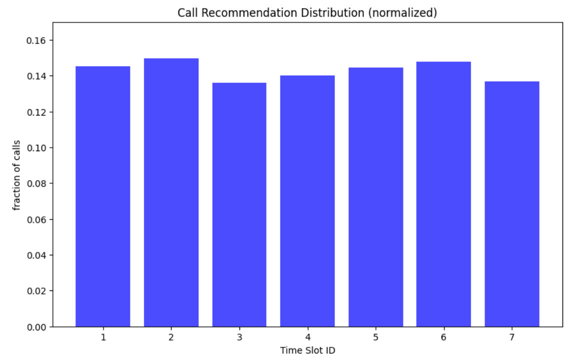

(b) Control Group

Figure 1: The above bars represent the fraction of unique calls, i.e. the first call recommended by the algorithm without considering re-attempts. These give a truer representation of the call recommendation distribution or policy. (a) despite the distribution having a notable dip in slot 5 and 6, we still observe good pickup success rates in slot 6 from Table 4, and for (b) the distribution is almost uniform as expected.