Measuring the Impact of AI Interventions in mHealth Programs

Arpan Dasgupta*
Google DeepMind
Bengaluru, India
arpandg@google.com

Sarvesh Gharat*
IIT Bombay
Mumbai, India
sarveshgharat19@gmail.com

Neha Madhiwalla
ARMMAN
Mumbai, India
neha@armman.org

Milind Tambe
Harvard University, Google DeepMind
Boston, USA
milindtambe@google.com

Aparna Taneja
Google DeepMind
Bengaluru, India
aparnataneja@google.com

ABSTRACT

Automated voice calls are a proven method for disseminating maternal and child health information. In our previous work, we showed that AI could identify beneficiaries who benefit most from live service call interventions, increasing listenership. We had also shown the positive trend in health and behavioral outcomes of AI scheduled interventions in the mHealth programme mMitra. This study directly links these findings by analyzing mothers' post-natal health knowledge via surveys. We demonstrate that AI-scheduled interventions, which enhance listenership, lead to statistically significant improvements in mothers' understanding of critical health topics during pregnancy and infancy. This underscores the potential of AI to drive meaningful improvements in maternal and child health.

KEYWORDS

AI for Maternal Health, Decision Focused Learning, Intervention Allocation

ACM Reference Format:
Arpan Dasgupta*, Sarvesh Gharat*, Neha Madhiwalla, Milind Tambe, and Aparna Taneja. 2025. Measuring the Impact of AI Interventions in mHealth Programs. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 8 pages.

1 INTRODUCTION

Timely access to reliable health information plays a crucial role in improving maternal and infant health outcomes, particularly in underserved communities where traditional healthcare services may be limited. Mobile health (mHealth) programs have emerged as an effective way to bridge this gap, leveraging the widespread use of mobile phones to deliver essential health education at scale. However, despite their potential reach, a decline in beneficiary engagement due to loss of interest remains a challenge.

To address this issue, interventions such as calls or visits from community health workers can help keep beneficiaries engaged by providing personalized support and encouragement. This increases the likelihood that they will remain active participants. However, given staffing and resource constraints, it is not feasible to reach every beneficiary, making it essential to identify those who would benefit the most.

This presents a complex prediction and resource allocation problem, but recent advancements in AI have proven effective in optimizing call schedules to maximize listenership. In this paper, in collaboration with ARMMAN [1], we focus on their mHealth program mMitra [2] which is the second-largest maternal mHealth program in the world.

In this work, our primary objective was to investigate the potential of AI-scheduled interventions to enable improvements in both behavioral and health outcomes in the context of the mMitra program. To demonstrate this, we administer a comprehensive survey aimed at the program's beneficiaries. Learning from our previous work, this survey was meticulously designed to assess the beneficiaries' understanding of crucial health practices while making sure the questions are simple to answer and to evaluate.

We deployed a randomized controlled trial (RCT) that comprised two distinct arms: an intervention arm, where beneficiaries received interventions precisely as scheduled by the AI algorithm, and a control arm, where beneficiaries did not receive such interventions. By contrasting the knowledge and behavioral outcomes of the beneficiaries across these two arms, we establish the outcome of the AI interventions.

Subsequent analysis yielded compelling findings. Firstly, we successfully established a statistically significant increase in listenership among beneficiaries within the intervention arm, thereby validating the efficacy of AI-scheduled interventions. Secondly, we observed a general trend of improved behavior and health practices among beneficiaries in the intervention arm in some questions, as shown by survey responses. In our previous study [5], we struggled to establish a statistically significant difference between the intervention and control arms with a high degree of certainty. This limitation stemmed primarily from the relatively small sample size inherent in our study, coupled with the presence of considerable noise and variability within the survey responses. In this study, by targeting a more focused group, we are able to counteract the problem of having a small sample size, and with more focused questions, we are able to reduce the noise in the responses. Hence, we are able to establish a statistically significant gain for the intervention arm in some questions pertaining to the knowledge and behavior of the participants.

This work is licensed under a Creative Commons Attribution International 4.0 License.

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).

In conclusion, this work provides evidence demonstrating the potential impact of AI-scheduled interventions on the health and behavior of mothers enrolled in the mMitra program, translating the previously established listenership difference into difference in knowledge among participants.

The paper is organized as follows. In Section 2, we briefly discuss previous work in this domain. In Section 3, we discuss the setup of the study, including intervention scheduling and the conducted survey. In Section 4, we describe our method of analysis for the survey and the results obtained.

2 RELATED WORK

The allocation of finite resources is a recurring challenge across various domains necessitating strategic planning. Restless multi-armed bandits (RMABs) serve as a prevalent instrument for addressing such sequential resource allocation problems within uncertain environments. Specifically, RMABs have demonstrated their utility in applications such as anti-poaching surveillance [15], multi-channel communication optimization [9], task scheduling [3, 22], unmanned aerial vehicle routing [23], and equipment maintenance alongside sensor monitoring [6]. These resource-constrained allocation problems also naturally manifest during the planning of interventions within mHealth programs [10].

Prior research has validated that the health information disseminated through mHealth programs results in enhanced infant care practices and maternal knowledge acquisition among mothers [7, 12]. Notably, Hegde and Doshi [7] employed a randomized controlled trial to assess the impact of personalized voice calls on mothers participating in mMitra. Hegde and Doshi [7] established statistically significant findings concerning improved infant care knowledge among mothers, as well as a direct influence on infant health, as measured by birth weight.

These findings motivated ARMMAN to enhance beneficiary listenership through health worker service calls.

In collaboration with ARMMAN, Mate et al. [11] presented an AI-driven methodology for scheduling intervention calls. This approach determines the allocation of service calls utilizing the RMAB framework, where each beneficiary is modeled as a Markov decision process. Their methodology underwent initial testing via simulations, followed by a field study, culminating in its large-scale deployment in practical settings [18]. A core challenge in SAHELI has been the learning of transition probabilities for Markov decision processes modeling beneficiaries. After iterative refinements, Wang et al. [19] and Verma et al. [17] implemented decision-focused learning (DFL) [16] for RMABs to facilitate the learning of transition probabilities, thereby improving program performance in deployment.

To date, the primary observable objective optimized by SAHELI and similar intervention scheduling programs for ARMMAN has been the mother's listenership of automated voice calls; consequently, program performance has consistently been evaluated through improvements in listenership metrics. However, a correlation between AI-scheduled interventions and behavioral outcomes had not been demonstrated, until [5] which showed positive trends in some questions in a health study. However, it failed to establish concrete statistical significance due to several limitations. This work aims to do so by improving upon the methodology of conducting the survey as well as asking questions in a more targeted manner.

3 STUDY SETUP

Our initial methodological step involved segmenting registered beneficiaries into distinct cohorts, categorized by their respective program enrollment dates. Within each cohort, we then implemented a randomized allocation procedure, creating intervention (DFL) and control (Dummy) groups. It is important to note that all participants, regardless of group assignment, consistently received automated voice calls disseminating essential health information throughout their program participation. However, only those assigned to the intervention group were considered for supplementary interventions. The AI algorithm played a crucial role in determining, on a weekly basis, which intervention group members would receive personalized live service calls from health workers. Finally, to evaluate the impact of these interventions on behavioral and health knowledge, we conducted a comprehensive survey, administered to representative subsets of both the intervention and control groups.

3.1 Experiment Arms

3.1.1 Cohorts. The study was conducted in three cohorts with a combined number of 34453 beneficiaries.

- Cohort 1 : 12749 beneficiaries. Registered between 1st of October 2023 to 31st of October 2023.
- Cohort 2 : 9122 beneficiaries. Registered between 1st November 2023 to 31st November 2023.
- Cohort 3 : 12582 beneficiaries. Registered between 1st December 2022 to 31st December 2023.

As detailed in Section Section 3.2.1, the program does not treat these cohorts as entirely distinct groups; rather, they serve primarily as a mechanism for establishing intervention eligibility timelines.

3.1.2 Division Into Arms. Within each cohort, beneficiaries were randomly assigned to either the intervention or control/dummy arm, ensuring a similar distribution of key attributes between the two groups. This procedure mirrors covariate adaptive randomization, a technique designed to balance the distribution of relevant covariates across experimental groups, as described by Lachin et al. [8]. We specifically balanced the following attributes:

Engagement States
- For each beneficiary and a given automated voice call, we denote the engagement state $E@T$ at a threshold T as $E@T = 1$ if the beneficiary listened to the call for at least T seconds, and $E@T = 0$ otherwise.
- We compute $E@T\_w$ for each beneficiary, representing the engagement state over w weeks leading up to the cohort's anticipated intervention start date.
- To ensure comparable listenership patterns between the two arms, we aim to achieve approximately equivalent values of $E@T\_w$ for thresholds $T \in \{1, 5, 10, 30, 100\}$ and time windows $w \in \{1, 2, 3\}$.

Demographic Features We incorporate the beneficiaries' gestational age, categorized into trimesters, as a feature for cohort formation. This categorization is achieved by dividing the gestational age in weeks by 14, resulting in four equally sized bins. We ensure that both arms have a balanced representation of beneficiaries across each trimester, which inherently also equates to a balanced distribution of beneficiaries who have completed delivery.

To ensure a balanced distribution of these attributes between arms, we first construct a feature vector Y for each beneficiary by concatenating the attributes. Subsequently, we partition the beneficiaries into two equally sized groups, employing Y as the stratification criterion [19]. This is accomplished by treating Y as a categorical label and utilizing a stratified splitting mechanism to produce two balanced subsets. Specifically, we leverage the stratified option within the *train_test_split* function from the sklearn library [14]. Given the sufficient number of beneficiaries within each cohort, we successfully achieve an exact split, resulting in perfectly balanced groups.

3.2 Conducting Interventions

Interventions in mMitra are service calls made by healthcare workers that aim to boost the future listenership of automated messages of the called beneficiary.

3.2.1 Number of Interventions per week. Interventions began on February 2024. We only intervene on beneficiaries that have been present for at least 3 weeks in the program. We perform interventions on beneficiaries of cohort 1, 2 and 3 one after the other, intervening on a total of 35% of beneficiaries in each cohort, this is to simulate the actual scenario where it is impractical to intervene on a lot of beneficiaries.

- 5th February 2024 to 3rd March 2024 (4 weeks) - consider only Cohort 1 for interventions.
- 4th March 2024 to 24th March 2024 (3 weeks) consider Cohort 2 for interventions.
- 25th March 2024 to 21st April 2024 (4 weeks) consider Cohort 3 for interventions.

We conduct approximately 330 interventions per week for cohorts 1 and 3 and about 285 for Cohort 2, while ensuring that each beneficiary can be intervened on only once. This is because we wanted to keep the approximate number of interventions similar for each week, and since Cohort 2 had lesser registrations, the number of weeks is lesser. We end up conducting interventions on about 12000 beneficiaries.

3.2.2 Eligibility for Interventions. Beneficiaries are deemed eligible for interventions when the following criteria are met:

(1) Active Program Participation: They maintain an active enrollment status and continue to receive automated voice messages.
(2) Recent Engagement: They have listened to at least one automated voice message within the four weeks preceding the intervention period for their respective cohort.
(3) No Prior Intervention: They have not previously received a live service call intervention.

These eligibility requirements serve to ensure the efficient and equitable allocation of the program's constrained intervention resources.

3.2.3 Conducting Interventions. Each week, the DFL-RMAB [19] algorithm, our chosen AI algorithm, determines the set of beneficiaries from the intervention arm who will receive an intervention. We then store all beneficiaries from the intervention arm that have received an intervention in some week into a list I_D. We also simulate the AI algorithm on the control (dummy) arm to determine the set of beneficiaries that would have been selected for an intervention (assuming we conducted the same number of interventions as in the intervention arm). As in the intervention arm, beneficiaries in the control arm cannot be selected multiple times. We collect the beneficiaries from the control arm that have been selected by the algorithm in some week into a list I_C. We create an intervention list I that combines I_D and I_C.

The idea behind this is that we later compare the behavior of beneficiaries from I_D and I_C, as we can think of beneficiaries from I_C as the counterfactual counterparts of those from I_D.

3.3 Health Study

3.3.1 Conducting the Survey. The health survey is conducted on the beneficiaries from the intervention list I between 20th June to $30th$ October in 2024. We also exclude from I_C and I_D, the beneficiaries who have delivered before the interventions (since this is a post-natal survey). Out of all the people who are scheduled to receive interventions in I_D, only a fraction pick up the intervention call (say this set is $I_{D'}$). This combined set of I_C and $I_{D'}$ is then surveyed. This is a major difference from the previous study [5] where we surveyed the entire set I_D, leading to fewer people in the survey who were actually intervened on.

This subset of beneficiaries are then called by a health worker and asked to answer the questions from the survey. However, the survey calls are only picked up by a fraction of beneficiaries. This makes it difficult to evaluate the outcome of the study as we know the survey results for only a subset of the beneficiaries that are willing and available to answer to the survey questions (in particular, this group of beneficiaries is not chosen uniformly at random). Hence, we have to re-balance the control and intervention group for the final comparison.

3.3.2 Survey Questions. Each participant was presented with 23 questions, designed to evaluate their program engagement and knowledge across various health-related domains. This assessment is aligned with the content of the automated voice messages, aiming to gauge the beneficiaries' comprehension of the delivered information. Specifically, the survey encompassed categories such as program engagement, general knowledge of health practices, breastfeeding practices, communication with family members, and health supplement usage. Figure 1 provides a comprehensive list of these questions. For each question, beneficiaries received a score based on their responses. Further, to ensure fairness, the interviewers were blinded and had no idea whether the interviewee belonged to Control or DFL group.

Among the 23 questions, 13 were structured as single-choice (Yes/No) questions. The remaining 8 were questions where scores were determined based on multiple possible correct answers, with varying weights assigned to different answers. Beneficiaries could select one or more of these answers, and their final score was the

	Registered	After removing inactive	Intervened On [About 35%] (for DFL)	Picked up Intervention (for DFL)	After imposing delivery-date constraint (delivery after intervention : check section 3.3.1)	Picked up Survey
DFL	12749 + 9122 + 12582 = 34453	19970	4495	3469	1496	701
Control			4495	4495	1901	850

Table 1: Beneficiary Counts at Different Stages

1. Did you know your baby's birth weight ?
2. Iron tablets ANC
3. Calcium tablets ANC
4. Scientific name iron pill
5. Scientific name calcium pill
6. Are you currently taking iron tablets after delivery?
7. Are you still taking calcium pills after delivery?
8. Did you share this with your husband?
9. Should the baby be fed the first solid yellow milk that comes to the mother after childbirth?
10. Should the baby be fed at night?
11. Social smile
12. Have you heard calls from mMitra regularly?
13. Have you ever discussed with your husband/family about the information you heard/told in the call?

1. If , yes what was the baby's weight at birth ?
2. If yes, what did the husband do?
3. When should the baby be fed the first milk after birth?
4. How many times a day do you feed your baby?
5. What to do to dry baby's umbilical cord after birth?
6. If yes, what made you want to listen to the call?
7. Does anyone in your family know that you are getting informative calls from mMitra?
8. Do they hear these calls too?
9. Do you pick up mMitra calls even when other people are around?
10. Do you pick up mMitra call even when you are busy at work?

Figure 1: Single Correct (left) and Multiple Correct Questions (right) asked in the study.

cumulative score of all correctly identified answers within that question.

4 RESULTS

4.1 Comparison Methodology

A direct comparison of the people surveyed (control vs DFL) does not suffice for comparison. The reason is that the set of people who pick up the intervention call and then the survey call are non-random. This implies that some other metric has to be used to ensure a fair comparison between the two sets.

In the paper [5], the method used was to choose a set of people from intervention group and then create an equivalent set in the control by performing one-to-one matching by using the Mahalonobis distance metric between the feature vectors of beneficiaries. These features included pre-intervention listenership, gestational age and number of children previously conceived.

In this study, we perform matching of beneficiaries using the Whittle Index [21], which is the metric used to decide who gets an intervention. The Whittle Index is a priority score assigned to each "arm" in a restless multi-armed bandit problem, guiding resource allocation by prioritizing arms that offer the most expected marginal reward for a given threshold. If the problem is "indexable" [13, 20], a simple policy of always choosing arms with the highest index values can be optimal for maximizing overall reward under resource constraints.

For each beneficiary, when the DFL algorithm decides to perform an intervention, the Whittle indices are stored. Whittle indices are computed both for intervention and control group, but the actual interventions only go out to the the former group. For the latter the whittle indices help us identify beneficiaries who would have been intervened on by the model had we conducted interventions. The advantage of using this index is that it takes demographic information and listenership into account automatically, and represents the preference of the algorithm to conduct intervention on a beneficiary. The people in the intervention set $I_{D'}$ are matched to a beneficiary in control I_C with a similar whittle index from the same cohort, after which the scores are compared.

4.2 Key Positive Results

4.2.1 Establishing Improved Listenership. A key objective of our study is to determine whether the intervention successfully enhances listenership among beneficiaries. To quantify the gain in listenership, we define it as the difference between post-intervention listenership and pre-intervention listenership, where post-intervention listenership is calculated by taking the mean of the average listenership per user over the two weeks after intervention and pre-intervention listenership is calculated by taking the mean of the average listenership per user over the two weeks before intervention.

In Figure 2, the x-axis represents the week number, indicating the point in time at which the listenership gain is calculated. The y-axis represents the average listenership gain in seconds, computed as the mean listenership gain for all beneficiaries who were intervened (or would have been intervened) in a given week. This approach ensures that for each week, we only consider the listenership gain of the relevant beneficiaries, providing a clear week-wise trend of the intervention's impact.

As seen in the figure, across all weeks, we observe a general trend where the intervened group performs better than the control. Additionally, for the first two cohorts, listenership gains remain high throughout all weeks. For Cohort 3, control also sees a significant improvement in listenership compared to Cohorts 1 and 2. However, the gain still remains lower than that of the intervened group, showing that the interventions have had a positive impact on listenership among beneficiaries. Overall, the listenership gain for intervention is significant compared to control.

4.2.2 Establishing Health Benefits. The survey responses were compared across the single and multi correct questions and some topics showed significant improvement. This subsection presents key findings from the study, focusing on the improvement in knowledge regarding iron and calcium tablet intake among beneficiaries in Cohorts 1 and 2. Additionally, we highlight the improved results for birth weight in the intervened arm across these two cohorts.

As seen in Figures 3, 4, and 5, the x-axis represents cumulative weeks, while the y-axis denotes scores. Here, cumulative weeks refer to the inclusion of all beneficiaries who received interventions up to a particular week. For example, in Cumulative Week 3, we consider all beneficiaries who were intervened in Weeks 1, 2, and 3. Similarly, for the control (dummy) arm, although no actual intervention takes place, we already know which beneficiaries would have been intervened. This allows us to easily identify the relevant list of beneficiaries from I_C to compare against the intervened arm beneficiaries from $I_{D'}$.

The Y-axis represents the average score for a particular question corresponding to the respective groups. This score is derived from beneficiary responses to the specific knowledge-based questions asked. Although we see a general improvement in several questions, we highlight the results for three questions for which we see statistically significant improvement:

(1) "Are you still taking iron pills after delivery?" (Single correct Q6)
(2) "Are you still taking calcium pills after delivery?" (Single correct Q7)
(3) "What was the baby's weight at birth?" (Multiple correct Q1)

A higher score on the first two questions indicates **improved knowledge regarding supplementation**, which is essential for the health of the mother [4]. In the case of birth weight knowledge, the responses indicate whether whether the mother kept track of the baby's weight after birth, which is essential knowledge to **understand whether the baby was born healthy**. The mother knowing the exact weight also helps in future checkups by health workers, helping them **keep track of the baby's health**. Hence, the higher score obtained by the DFL group indicates improvement in knowledge of critical health information.

As seen in Figures 3 and 4, we observe that towards the end of Cohort 2, the score differences between the intervened and control groups increase, with the intervened group performing better. These results become more evident as we also observe a significant decrease in the p-values (refer Table **??**) for Cumulative Weeks 6 and 7.

In contrast, for Cohort 3, although the actual scores are higher than those of Cohorts 1 and 2, we cannot clearly distinguish the differences between the intervened and control groups. This is likely because the beneficiaries in Cohort 3 already have high listenership (refer fig. 6), suggesting that any further increase in listenership does not significantly impact the knowledge gained by the beneficiaries.

Similarly, for the question "What was the baby's weight at birth?", the scores follow a similar pattern, with statistical significance observed across all weeks (except Week 1) for Cohorts 1 and 2. While Cohort 3 shows scores comparable to those of Cohorts 1 and 2, the high p-values prevent us from drawing any meaningful conclusions about performance.

5 CONCLUSION

In this paper, we show that AI scheduled interventions can lead to significant difference in knowledge and behaviour among mothers when deployed at scale in a mHealth program. We improve upon the previous work by conducting a focused and better designed study, and establish statistically significant difference in questions based on knowledge and behaviours in particular importance of taking iron and calcium supplements and keeping track of baby's weight.

REFERENCES

[1] Armman Foundation. 2024. Armman Foundation. https://www.armman.org/ Accessed: March 6, 2024.
[2] Armman Foundation. 2024. mMitra Program. https://armman.org/mmitra/ Accessed: March 6, 2024.
[3] Saeed Bagheri and Anna Scaglione. 2015. The restless multi-armed bandit formulation of the cognitive compressive sensing problem. *IEEE Transactions on Signal Processing* 63, 5 (2015), 1183–1198.
[4] LaVerne L Brown, Barbara E Cohen, Emmeline Edwards, Courtney E Gustin, and Zara Noreen. 2021. Physiological need for calcium, iron, and folic acid for women of various subpopulations during pregnancy and beyond. *Journal of Women's Health* 30, 2 (2021), 207–211.
[5] 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

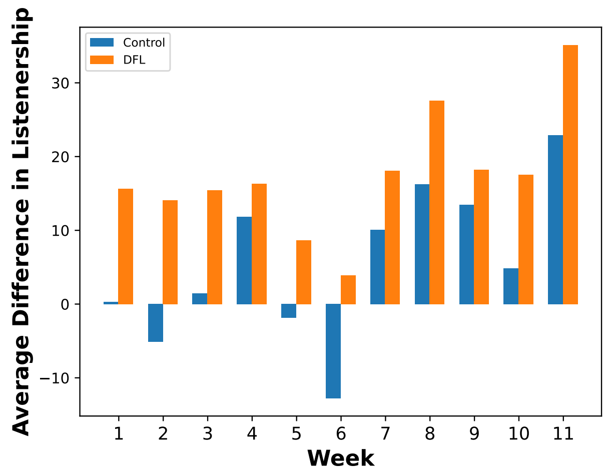

Figure 2: Comparison of Listenership Gain for DFL and Control for all the Cohorts

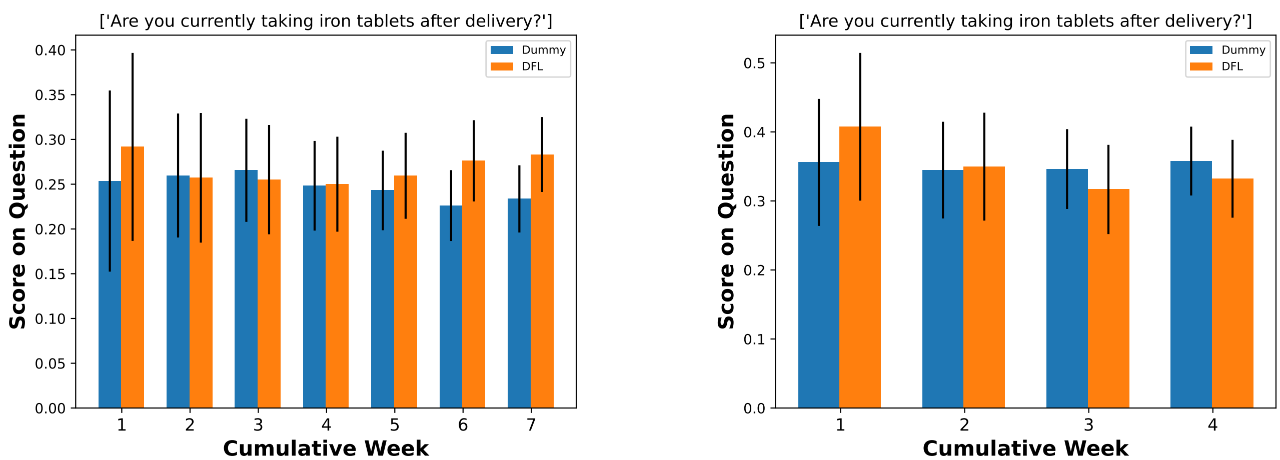

(a) Cohort 1 and 2

(b) Cohort 3

Figure 3: Are you currently taking iron tablets after delivery

	Cohort 1 and 2							Cohort 3			
	Week 1	Week 2	Week 3	Week 4	Week 5	Week 6	Week 7	Week 1	Week 2	Week 3	Week 4
Q1	0.7457	1.0000	0.8956	1.0000	0.6878	0.1207	**0.0981**	0.5716	1.0000	0.5713	0.5611
Q2	1.0000	0.9851	0.9971	0.8596	0.6517	**0.0801**	**0.0413**	0.4416	0.7949	0.6102	0.6818
Q3	0.1003	**0.0204**	**0.0163**	**0.0114**	**0.0173**	**0.0186**	**0.0080**	0.2792	0.5707	0.3896	0.7840

Table 2: p-values for cumulative weeks in Cohort 1,2 and Cohort 3. Q1 represents the response to 'Are you currently taking iron tablets after delivery?', Q2 corresponds to 'Are you taking calcium tablets after delivery?', and Q3 relates to 'If yes, what was the baby's weight at birth?'

Health Programs. *arXiv preprint arXiv:2407.11973* (2024).
[6] Kevin D Glazebrook, Diego Ruiz-Hernandez, and Christopher Kirkbride. 2006. Some indexable families of restless bandit problems. *Advances in Applied Probability* 38, 3 (2006), 643–672.
[7] 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.
[8] John M Lachin, John P Matts, and LJ Wei. 1988. Randomization in clinical trials: conclusions and recommendations. *Controlled clinical trials* 9, 4 (1988), 365–374.

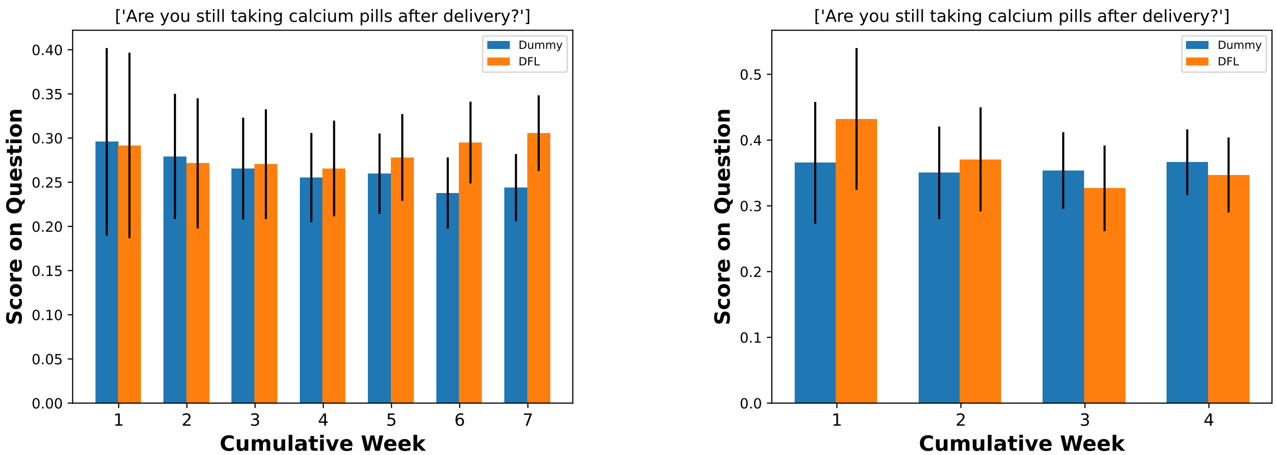

(a) Cohort 1 and 2

(b) Cohort 3

Figure 4: Are you currently taking iron tablets after delivery

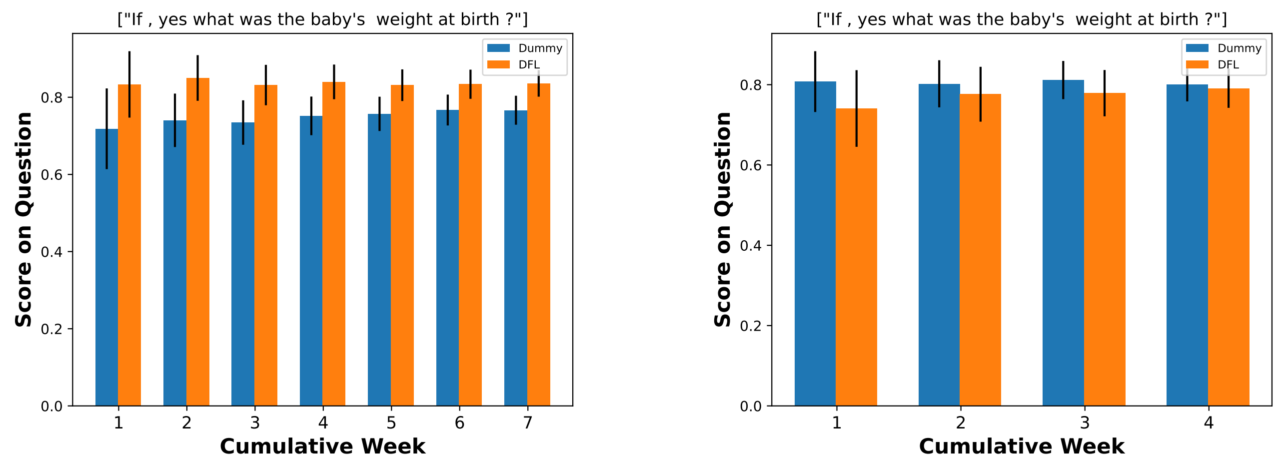

(a) Cohort 1 and 2

(b) Cohort 3

Figure 5: What was the baby's weight at birth?

[9] Keqin Liu and Qing Zhao. 2010. Indexability of restless bandit problems and optimality of whittle index for dynamic multichannel access. *IEEE Transactions on Information Theory* 56, 11 (2010), 5547–5567.

[10] Aditya Mate, Jackson Killian, Haifeng Xu, Andrew Perrault, and Milind Tambe. 2020. Collapsing bandits and their application to public health intervention. *Advances in Neural Information Processing Systems* 33 (2020), 15639–15650.

[11] 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 Thirty-Sixth AAAI Conference on Artificial Intelligence.* AAAI Press, 12017–12025.

[12] 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.

[13] Jose Nino-Mora. 2001. Restless bandits, partial conservation laws and indexability. *Advances in Applied Probability* 33, 1 (2001), 76–98.

[14] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research* 12 (2011), 2825–2830.

[15] Yundi Qian, Chao Zhang, Bhaskar Krishnamachari, and Milind Tambe. 2016. Restless poachers: Handling exploration-exploitation tradeoffs in security domains. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems.* 123–131.

[16] Sanket Shah, Kai Wang, Bryan Wilder, Andrew Perrault, and Milind Tambe. 2022. Decision-focused learning without decision-making: Learning locally optimized decision losses. *Advances in Neural Information Processing Systems* 35 (2022), 1320–1332.

[17] Shresth Verma, Aditya Mate, Kai Wang, Neha Madhiwalla, Aparna Hegde, Aparna Taneja, and Milind Tambe. 2023. Restless Multi-Armed Bandits for Maternal and Child Health: Results from Decision-Focused Learning. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems.* ACM, 1312–1320.

[18] Shresth Verma, Gargi Singh, Aditya Mate, Paritosh Verma, Sruthi Gorantla, Neha Madhiwalla, Aparna Hegde, Divy Thakkar, Manish Jain, Milind Tambe, and

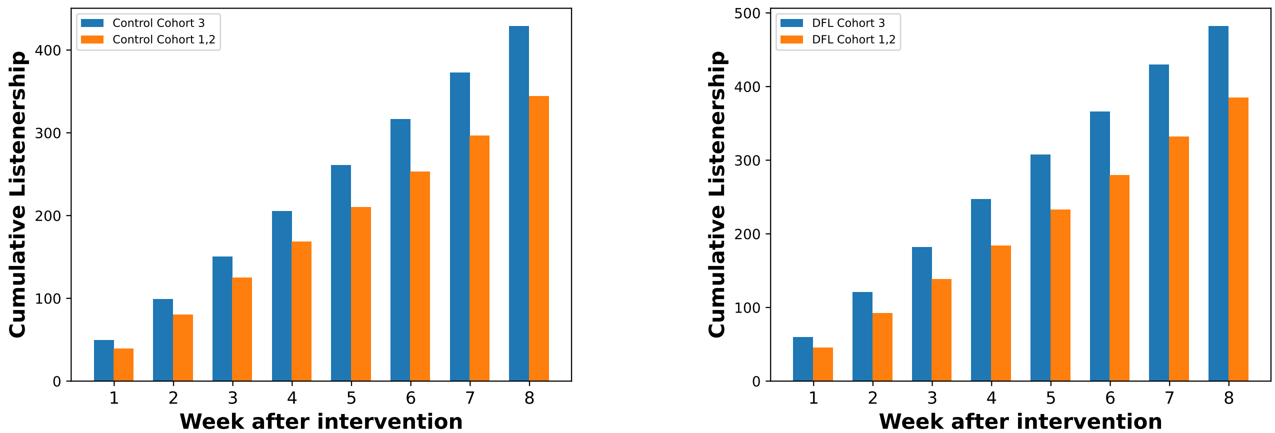

(a) Control **(b) DFL**

Figure 6: Comparison of the cumulative listenership after "intervention" for DFL and Control. Both the plots indicate that the base listenership for Cohort 3 was higher.

Aparna Taneja. 2023. Expanding impact of mobile health programs: SAHELI for maternal and child care. *AI Magazine* 44, 4 (2023), 363–376.

[19] Kai Wang, Shresth Verma, Aditya S. Mate, Sanket Shah, Aparna Taneja, Neha Madhiwalla, Aparna Hegde, and Milind Shashikant Tambe. 2023. Scalable Decision-Focused Learning in Restless Multi-Armed Bandits with Application to Maternal and Child Care. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence*. AAAI Press, 12138–12146.

[20] Richard R Weber and Gideon Weiss. 1990. On an index policy for restless bandits. *Journal of applied probability* 27, 3 (1990), 637–648.

[21] Peter Whittle. 1988. Restless bandits: Activity allocation in a changing world. *Journal of applied probability* 25, A (1988), 287–298.

[22] Zhe Yu, Yunjian Xu, and Lang Tong. 2018. Deadline scheduling as restless bandits. *IEEE Transactions on Automatic Control* 63, 8 (2018), 2343–2358.

[23] Qing Zhao, Bhaskar Krishnamachari, and Keqin Liu. 2008. On myopic sensing for multi-channel opportunistic access: structure, optimality, and performance. *IEEE Transactions on Wireless Communications* 7, 12 (2008), 5431–5440.