

---

# Prioritization Strategies for LLM-Designed Restless Bandit Rewards in Public Health

---

**Shresth Verma**  
Harvard University

**Niclas Boehmer**  
Harvard University

**Lingkai Kong**  
Harvard University

**Milind Tambe**  
Harvard University

## Abstract

LLMs are increasingly used to design reward functions based on human preferences in Reinforcement Learning (RL). We focus on LLM-designed rewards for Restless Multi-Armed Bandits, a framework for allocating limited resources among agents. In applications such as public health, this approach empowers grassroots health workers to tailor automated allocation decisions to community needs. In the presence of multiple agents, altering the reward function based on human preferences can impact subpopulations very differently, leading to complex tradeoffs and a multi-objective resource allocation problem. We are the first to present a principled method termed *Social Choice Language Model* for dealing with these tradeoffs for LLM-designed rewards for multiagent planners in general and restless bandits in particular. The novel part of our model is a transparent and configurable selection component, called an *adjudicator*, external to the LLM that controls complex tradeoffs via a user-selected social welfare function. Our experiments demonstrate that our model reliably selects more effective, aligned, and balanced reward functions compared to purely LLM-based approaches.

## 1 Introduction

Reward functions play a fundamental role in the generation of optimal policies for sequential decision-making via reinforcement learning. Previous work has shown that LLMs are an effective tool for designing reward functions that can be guided and customized via human language prompts [27, 10, 23, 49, 28, 52, 19]. We study the problem of designing high-quality reward functions aligned with human preference prompts in the context of *multiagent* optimization and planning problems. We present a transparent framework around LLMs that constructs effective, aligned, and balanced reward functions for complex human prompts.

We study the reward design problem for restless multi-armed bandits (RMABs), a popular model in multiagent systems for sequentially allocating a limited number of resources to a set of agents [48, 35]. In RMABs, each agent is represented by an individual Markov Decision Process including a reward function. By choosing these reward functions, one can control which agents are more or less likely to receive a resource. RMABs have been used in various domains such as machine maintenance [1], anti-poaching [37], and healthcare [5, 46]. In many of them, system organizers have evolving allocation priorities based on agents' features that need to be incorporated into the resource allocation process [15, 47]. For instance, in a healthcare program, a healthcare worker might want to change the allocation policy to prioritize low-income beneficiaries who are at higher risk or older beneficiaries who have transportation barriers for healthcare access [33, 43] via the preference prompt: *Prioritize low-income beneficiaries and older beneficiaries*.

Translating such human language prompts to effective and aligned reward functions is a general, non-trivial challenge in RL [19, 10, 27], e.g., as it is unclear how changes in the reward influence the used policy downstream. However, the multiagent nature of the RMAB problem adds a new twist to the problem of reward design in RL: It becomes fundamentally multi-objective. Consider the above example prompt asking for the prioritization of two subpopulations. As these subpopulations may contain different agents, selecting a reward function will most likely involve trading off the interests of the low-income vs. older beneficiaries, making this a multi-objective problem. If this multi-objective nature is ignored, a selected reward function might heavily favor one of the two groups (e.g., leading to the allocation of many resources to low-income beneficiaries, and no resources to older ones). Another tradeoff that emerges because of the multiagent nature of the problem is between aligning with the preference prompt and impacting agents not referenced in the prompt: Changing the reward functions to align with a prompt can lead to drastic changes in the number of resources distributed to subpopulations not referenced in the user prompt (e.g., in the example prompt, the selected reward might lead to taking away all resources from well-educated beneficiaries). If not adequately balanced, this can cause strong side effects unintended by the user.

To our knowledge, we are the first to address the multi-objective nature of LLM-powered reward design in RMABs in particular and multiagent planners in general. Closest to our paper is the work by Behari et al. [8] who proposed a fully LLM-based Decision-Language Model for RMABs to generate and select reward functions (as code) from human language prompts. However, as argued in Sections 2 and 4, the DLM model is not properly equipped to handle the multi-objective nature of the problem, as the LLM selects functions in an unpredictable, hard-to-control and sometimes (clearly) suboptimal way that does not adequately take into account and balance the different objectives.

We present a Social Choice Language Model (SCLM) that designs reward functions (as Python code) aligned with complex, multi-objective human language preferences in the context of a public health program (see Figure 2). Our pipeline separates the generation of candidate reward functions in the *generator* from the selection of one function in the *adjudicator*. For the generator, we use LLM-powered evolutionary search to generate a pool of reward functions [8]. In the transparent and customizable adjudicator, we take a new social choice perspective to address the multi-objective nature of our problem: We create a scorer component that evaluates the quality of generated reward functions according to the different objectives (e.g., different prioritization requests). Subsequently, a social welfare function aggregates “alignment scores” to select the best reward function. The user can select social welfare function and thereby has additional control over the preferred trade-off between objectives, e.g., maximizing the summed vs. minimum alignment of all objectives. We demonstrate in our experiments that our SCLM leads to the selection of reward functions significantly better aligned with complex, multi-objective prompts. Moreover, we show how it can be used to mitigate risks of using rewards designed from human prompts: unintended effects for other agents and ineffective allocation of resources. *SCLM combines the generative power of LLMs to design reward functions with the capabilities of social choice to handle multi-objective decision-making scenarios.*

## 2 Related Works

LLMs have emerged as a powerful tool to enhance RL. Recent work has used LLMs to generate reward functions based on natural language descriptions [27, 50, 52]. For instance, Goyal et al. [18], Carta et al. [11], Mirchandani et al. [31], Hazra et al. [19] shape rewards by training an RL agent to learn and complete intermediate tasks guided by language, yet focusing on very different (non-multiagent) environments.

The work of Behari et al. [8] is the first to present a Decision-Language Model for generating reward functions for RMABs from human prompts. The model performs a form of evolutionary search to find reward functions aligned with the given prompt in two interleaving phases: generation and reflection. In the generation phase, an LLM generates a set of reward functions. Based on reward function’s performances, in the reflection phase [27, 41], the LLM selects the function best aligned with the prompt. This function is then included in the prompt for the next generation phase or returned. In contrast to our work, DLM mixes generation with selection and does not explicitly account for the multi-objective nature of the reward design problem. Furthermore, in contrast to our work, they focus on small RMAB instances ( $\sim 20$  arms). Throughout the paper, we will use a slightly modified variant of DLM adjusted to our setting (see Section C.2) as a baseline.

LLMs can also be integrated into model-based and explainable RL, serving as world model simulators [26] or policy interpreters [14]. They can also act as information processors, effectively managing environment data and task instructions [7, 42] and as direct decision makers [25] to improve the sample efficiency of offline RL. While this area of research is rapidly evolving, there remains a significant gap in addressing multi-objective, multi-agent scenarios within LLM-enhanced RL, which we begin to address in this paper.

### 3 Preliminaries

An instance of Restless Multi-Armed Bandits (RMAB) is defined by a set of  $N$  arms, a time horizon  $T$ , and a budget  $K$ . We also refer to arms as agents. Each arm  $i \in [N]$  is an independently evolving MDP with state space  $\mathcal{S}_i$ , actions  $\mathcal{A}_i = \{0, 1\}$ , transition function  $P_i : \mathcal{S}_i \times \mathcal{A}_i \times \mathcal{S}_i \rightarrow \mathbb{R}_{\geq 0}$ , and reward function  $R_i : \mathcal{S}_i \rightarrow \mathbb{R}$ . We refer to 1 as the *active* action corresponding to pulling the arm (i.e., allocating a resource) and 0 as the *passive* action corresponding to not pulling the arm. We focus on the popular case where each MDP consists of two states, i.e.,  $\mathcal{S}_i = \{0, 1\}$  for all  $i \in [N]$ , yet our methodology applies to MDPs with arbitrary state spaces. We refer to 0 as the *bad* and 1 as the *good* state. For each step in which an agent is in the good state, they derive a *utility* of 1, while they derive a utility of 0 in the bad state. Accordingly, agents’ *default reward function*  $R^*$  is  $R^*(s) = s$ . We assume that there is a set of categorical features. Each arm is associated with a value of each feature. A *global reward function* is a reward function defined over features, which induces a reward function for each arm by plugging in its feature values (see Example 4.1).

In each step within the time horizon  $T$ , the planner observes the state of all arms and decides to pull a subset of at most  $K$  arm. As solving the RMAB problem optimally is computationally intractable [36], we make use of the very popular state-dependent Whittle index [48, 35], which tries to quantify for each state of each arm the reward gain achieved from applying the active action to the arm in this state. In the Whittle index policy  $\Pi$ , in each step, we compute the Whittle index for each arm (based on its current state) and pull the arms with the  $K$  highest Whittle indices. We will use it as the solution strategy in the following.

For a global reward function  $R$ , we write  $\Pi(R)$  to denote the Whittle index policy for  $R$ , i.e., the Whittle index policy for the instance where each agent uses the function  $R$  after plugging in their feature values as their reward. We refer to  $\Pi(R^*)$  as the *default policy*. To assess the quality of a global reward function  $R$ , we often consider the *utility feature distribution* for some feature  $X$ . This distribution shows for each value of the feature, the expected utility generated by arms with this feature value under the policy  $\Pi(R)$  (see Figure 4a in Section B for an example).

### 4 Problem Statement & Challenges

We assume that we are given a human-language preference prompt, concatenating one or multiple *preference clauses*. Each preference clause specifies a single optimization goal. We explicitly consider two types of preference clauses (yet our methodology extends to arbitrary ones):

- Give priority to agents with certain feature values, i.e., increase utility generated by agents,
- do not shift the utility distribution for some feature

We mostly focus on the first type and refer to them as *prioritization clauses and prompts*. A preference prompt is a set  $P = \{p_1, p_2, \dots\}$  of the involved preference clauses. We call a prompt  $P$  *singular* if  $|P| = 1$  and *composite* otherwise; our focus is on the latter. We can influence the utility agents generate by selecting a single global reward function (inducing reward functions  $(R_i)_{i \in [n]}$  for all agents).

**Example 4.1.** Consider an RMAB instance with three binary features  $A$ ,  $B$ , and  $C$ . A preference prompt  $P$  could be “Prioritize agents with  $A = 0$  and prioritize agents with  $B = 1$ ”, i.e.,  $P = \{\text{“prioritize agents with } A = 0\text{”}, \text{“prioritize agents with } B = 1\text{”}\}$ . Two possible global reward functions for the prompt are  $R'(s) = s \cdot (1 - A) \cdot B$  and  $R''(s) = s \cdot (1 - A) + s \cdot B$ . For function  $R''$ , the reward of an agent  $i$  with  $A = 0$  and  $B = 1$  is  $R_i(s) = 2s$ , while the reward of an agent  $j$  with  $A = 1$  and  $B = 1$  is  $R_j(s) = s$ . Selecting  $R''$ , agent  $i$  is more likely to receive a resource than agent  $j$ , as the good state contributes more reward for  $i$ .

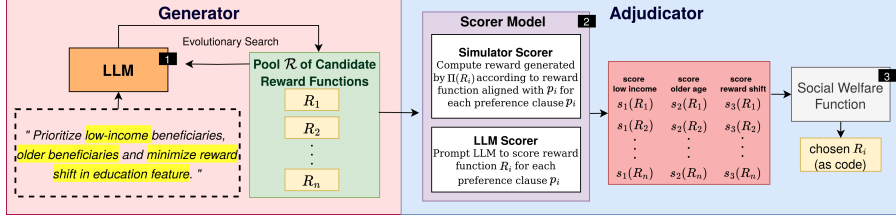
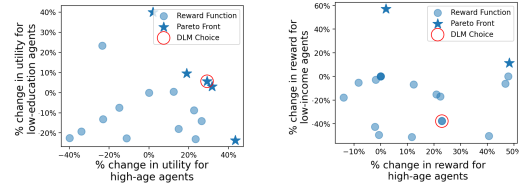


Figure 2: Overview of SCLM. In step 1, the preference prompt is passed to the generator, which performs an evolutionary search to create a pool  $\mathcal{R}$  of candidate reward functions. In step 2, these functions are passed to the adjudicator where a scorer model (e.g., the simulator or LLM scorer) computes the alignment scores. In step 3, a user-defined social welfare function selects a reward function based on the alignment scores.

We want to design a single global reward function that is “well-aligned” with all clauses of the given human-language preference prompt. However, as clauses can contradict each other, perfect alignment with all clauses becomes impossible. For instance, if a prompt requests the prioritization of two fully disjoint subpopulations, each resource will only benefit one of the two. When picking the reward function, we need to carefully balance the interests of the two groups of agents against each other. Generally, in the presence of multiple agents and limited resources, each clause can be viewed as a separate independent objective that we want to optimize, rendering this a multi-objective problem.

To illustrate the tradeoff decisions we face between different clauses when selecting a reward function, in Figure 1, we show two instances from our experiments for a prompt consisting of two prioritization clauses. Every point represents a LLM-designed reward function. The  $x$  and  $y$  axes represent the quality of the reward function from the perspective of the two prioritized subgroups where higher percentage values indicate more benefits (see Section 6.4 for details). Reward functions marked with stars lie on the Pareto fronts (no other available function dominates them).



(a) Prompt: “Prioritize agents with old age and agents with low education” (b) Prompt: “Prioritize agents with high age and agents with low income”

Figure 1: Tradeoffs between prioritization clauses.

In our experiments, we observe that DLM picks functions from very different parts of the Pareto frontier, potentially clearly prioritizing one subgroup over another (see Figure 1a for an example). In many other instances, it also picks suboptimal functions, i.e., functions that do not lie on the frontier, that may even harm one of the subgroups while strongly benefiting the other (see Figure 1b). This highlights the risks (and shortcomings of DLM) in not accounting for the multi-objective nature of the problem, as it picks reward functions that are inefficient (i.e., dominated) and unfair (i.e., heavily favoring one clause over the other).

*Our goal is to create a model that handles the tradeoffs posed by composite “multi-objective” prompts in a principled, transparent, and customizable fashion and outputs a single effective and fairly aligned global reward function.*

## 5 Social Choice Language Model (SCLM)

We propose a Social Choice Language Model to generate rewards from human language composite preference prompts (see Figure 2 for a visualization). Separating the generation and selection of reward functions, the model consists of two sequential components. The LLM-powered *generator* generates a set of candidate reward functions. Subsequently, taking a social-choice-inspired viewpoint, the *adjudicator* selects a reward function from the pool to be returned to the user in two steps: First, a scorer model computes an alignment score for each reward function with each prioritization clause (i.e., we judge each reward function from the perspective of all relevant “objectives”). Second, a user-defined social welfare function aggregates these scores into a “winning” candidate reward function. By selecting the social welfare function, the user can control the area of the Pareto frontier from which reward functions get selected. While we remark that our model can also be used to

tackle multi-objective issues arising when designing rewards in single-agent RL, the details of our components (e.g., the reflection in the generator and the computation of alignment scores) are specific to the multiagent nature of the RMAB problem.

## 5.1 Generator

Given a prompt, our generator creates a set of candidate reward functions (as Python code) via a variant of evolutionary search [8]: We proceed in multiple steps. First, inputting the problem description, feature descriptions and the preference prompt, we ask an LLM to generate code for a reward function. We repeat this query  $n_p$  times to obtain a set  $\mathcal{R}$  of  $n_p$  candidate reward functions. Afterwards, for each function  $R \in \mathcal{R}$  we compute the utility feature distributions of the policy  $\Pi(R)$  induced by the reward function  $R$  on the given RMAB instance (via repeatedly simulating the policy on the instance). Then, the prompt and the set of candidate reward functions together with the associated utility feature distributions are passed to an LLM, which is asked to select the reward function  $R'$  from  $\mathcal{R}$  best aligned with the prompt [27, 41]. Now, we repeat the whole process, this time including the selected policy  $R'$  as a seed in the reward function generation prompts. Once we have executed the process  $n_r$  times, we add all generated  $n_p \cdot n_r$  candidate reward functions  $R$  to the pool  $\mathcal{R}$  (see Section C.2 for details).

## 5.2 Adjudicator

The adjudicator selects a reward function from a given pool of candidate reward functions returned by the generator. To handle the complex tradeoffs arising within composite prompts and the resulting multi-objective optimization problem, the adjudicator follows a social choice approach. Social choice is a discipline at the intersection of economics, philosophy, and mathematics and concerned with aggregating the potentially contradicting preferences of a set of voters into a fair compromise alternative from a given candidate set [4, 32]. It thus provides a theoretically grounded and well-studied methodology for balancing competing interests. In our problem, we can interpret the reward functions as the candidates and the preference clauses in the prompt as the voters with their preferences over the candidates reflecting the reward function’s alignment with the clause. This view gives rise to the following strategy: Given a prompt  $P = \{p_1, p_2, \dots, p_\ell\}$ , we evaluate each reward function  $R \in \mathcal{R}$  from the perspective of each preference clause  $p_i$  by computing an (alignment) score  $s_i(R)$ .  $s_i(R)$  measures the alignment of  $\Pi(R)$  with preference clause  $p_i$ , i.e., how much the voter representing  $p_i$  “likes” the candidate  $R$ . We describe at the end of this section how these scores are computed.

**Social Welfare Functions** Social welfare functions select an alternative based on input preferences of voters. The pros and cons of individual social welfare functions have been extensively researched and debated in social choice [4, 39]. We make use of cardinal social welfare functions [22] which take as input our alignment scores  $(s_i(R))_{i \in [\ell], R \in \mathcal{R}}$  and output the winning reward function. We consider the three arguably most popular social welfare functions:

**Utilitarian** Return the reward function maximizing the sum of its scores, i.e.,  

$$\arg \max_{R \in \mathcal{R}} \sum_{i \in [\ell]} s_i(R).$$

**Nash** Return the reward function maximizing the product of its scores, i.e.,  

$$\arg \max_{R \in \mathcal{R}} \prod_{i \in [\ell]} s_i(R).$$

**Egalitarian** Return the reward function maximizing its minimum score, i.e.,  

$$\arg \max_{R \in \mathcal{R}} \min_{i \in [\ell]} s_i(R).$$

Social welfare functions also allow for assigning a different importance to clauses: The user could submit an importance score  $w_i$  for each clause  $p_i$ , which can be easily incorporated in the social welfare function, e.g., the Utilitarian welfare function becomes  $\arg \max_{R \in \mathcal{R}} \sum_{i \in [\ell]} w_i \cdot s_i(R)$ .

Selecting the social welfare function gives us control over the tradeoffs between objectives: By picking the egalitarian function, we ensure that one clause will not get prioritized over another. In contrast, the Utilitarian function prioritizes the summed alignment, allowing for mismatches between clauses; the Nash function strikes a balance between the two functions. Further, the adjudicator makes the selection process more transparent, as the different objectives, the selection criterion, and the performance of the candidate reward functions regarding the objectives become explicit.

It remains to describe how the alignment scores  $s_i(R)$  are computed. We present two general methods to compute alignment scores, which we will use for prioritization clauses. Subsequently, we discuss two more customized methods for the prevention of unintended utility shifts or drops in total generated utility.

**Simulator Scorer Model (SCLM-SIM)** For each preference clause  $p_i \in P$ , we compute a reward function  $R_i$  aligned with  $p_i$  by casting it as a singular prompt to the DLM pipeline (see Section C.2). For each  $R \in \mathcal{R}$ , we compute as  $s_i(R)$  the expected reward according to reward function  $R_i$  produced by policy  $\Pi(R)$  (again, we approximate this quantity by running multiple simulations).  $s_i(R)$  quantifies the quality of the policy induced by the candidate reward function  $R$  from the perspective of  $p_i$  (as captured by  $R_i$ ). As the scale of the reward functions can vary significantly among preference clauses, we normalize the scores by the performance of the default policy, i.e., we compute  $\frac{s_i(R) - s_i(R^*)}{s_i(R^*)}$ .

**LLM Scorer Model (SCLM-LLM)** The Simulator Scorer Model assumes access to reward functions capturing individual preference clauses well. If no well-aligned reward functions can be obtained, the performance of SCLM-SIM can deteriorate because it can become very noisy. Another disadvantage of SCLM-SIM is that the scores in SCLM-SIM are all computed via simulation, which can become computationally quite costly. Motivated by this, we propose a quicker and more flexible LLM-based approach, where we prompt an LLM to rate the alignment of a candidate reward function with a preference clause. In particular, for each  $R \in \mathcal{R}$  and  $p_i \in P$ , we use a prompt that includes  $R$ ,  $p_i$ , and the utility feature distributions produced by policy  $\Pi(R)$ . We ask the LLM to rank how well  $R$  aligns with the preference clause  $p_i$  on a scale from 1 to 5 (see Section E for prompt texts).

**Preventing Unintended Utility Shifts and Utility Drop** Aligning reward functions to a prioritization prompt might cause (unintended) utility shifts in other features (e.g., due to feature correlations, shifting utility to low-income beneficiaries might shift it away from more educated ones). See Section B for a concrete example. SCLM offers users the option to explicitly prevent these shifts by adding additional clauses (“objectives”) to the prompt: Given a prompt  $P$  (e.g., the prompt from Example 4.1), for each feature not referenced in the prompt, the user can add a new preference clause requesting a minimum shift in the utility distribution of this feature (e.g., for Example 4.1 they could add “do not change the utility distribution for feature  $C$ ”). To compute the alignment score  $s_i(R)$  between a reward function  $R$  and a clause  $p_i$ =“minimize utility shift for feature  $X$ ”, we compare feature  $X$ ’s utility distribution under the default policy with its utility distribution under the policy  $\Pi(R)$ . Specifically, we quantify the difference using the Earth mover’s distance (EMD) between the two distributions. Afterward, we apply 0-1 normalization to all scores  $s_i(R)_{R \in \mathcal{R}}$  for prompt  $p_i$ , which are input to the social welfare function (along with the alignment scores for the other clauses).

Another potential risk of aligning a reward function with a prioritization prompt is that it can sharply decrease the summed utility generated by all agents: The user might request the prioritization of a subpopulation that does not benefit much from receiving a resource, leading to severe drops in the summed utility generated by all agents. Users can address this issue in our model by adding a clause  $p_i$ =“maximize the total generated utility” to the prompt. As the alignment score  $s_i(R)$  of  $p_i$  with some reward function  $R$  we compute the summed utility generated by all agents under the policy  $\Pi(R)$  (computed via multiple simulations of the policy on the given instance). We again apply 0-1 normalization to all scores  $s_i(R)_{R \in \mathcal{R}}$  for prompt  $p_i$ .

## 6 Experiments

We describe our testset (Section 6.1), the compared methods (Section 6.2), and our experimental results both for dealing with composite prioritization prompts (Section 6.4) and additionally minimizing unintended side effects (Section 6.5). Following the work of Behari et al. [8], which constitutes our most important baseline, we use Gemini Pro [2] as the LLM in our experiments.

### 6.1 Dataset Description

**Synthetic Domain** We create three synthetic datasets, each with three features ( $A$ ,  $B$ , and  $C$ ). For some agent and feature, we randomly sample the agent’s feature value between 0 and 1. Arm’s

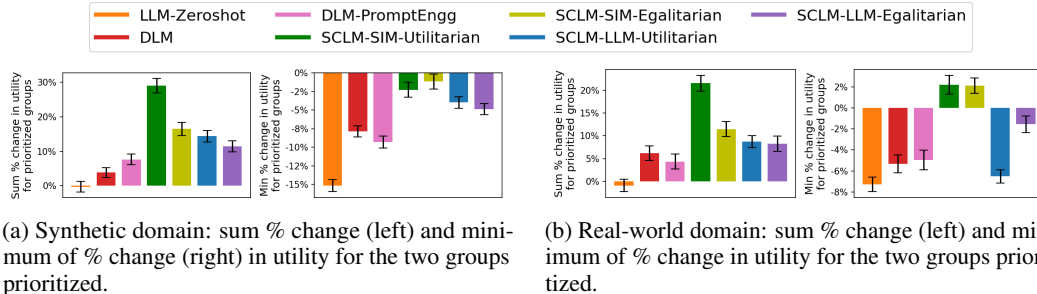


Figure 3: Results comparing the quality of reward design methods for composite prioritization prompts. Results are averaged across  $180 = 12 \cdot 15$  values: 12 composite prompts on 15 RMAB instances (from 3 datasets). Error bars represent std-error.

passive transition probabilities ( $P(s, a = 0, s')$  for  $s, s' \in [0, 1]$ ) are sampled uniformly between 0 and 1. Our three datasets differ in how active transition probabilities are sampled. For each dataset, we set a three-dimensional weight vector  $W \in [0, 1]^3$  specifying how much each feature impacts the effect of applying an active action (see Section C.3 for details). For each agent, let  $f$  denote their feature values; we sample their active transition probabilities for  $s, s' \in [0, 1]$  as  $P(s, a = 1, s') = P(s, a = 0, s') + \delta$ ,  $\delta \sim \mathcal{N}(\Delta, \sigma)$ , where  $\Delta = W \cdot f^T$  and  $\sigma$  is a hyperparameter. Subsequently, we discretize the values of all features into 5 different equal-sized buckets to be consistent with the real-world domain. This also makes our experiments robust to small changes in feature values. For each dataset, we sample 5 instances according to the above procedure with  $N = 2100$  arms, a budget of  $B = 210$ , and a time horizon of  $T = 12$  (to replicate the setup of the real-world domain described below where roughly 3% – 10% of beneficiaries receive a resource every week for up to three months).

**Real-World Domain** ARMMAN [3] is a non-profit in India that operates large-scale Maternal and Child Care Mobile Health programs for underserved communities. One of their programs disseminates critical health information via weekly automated voice messages. The goal of the NGO is to maximize beneficiaries’ engagement, i.e., the number of messages they listen to. A limited number of beneficiaries are called by health workers every week to boost engagement. The problem of planning which beneficiaries to call has been modeled and solved as an RMAB, where the good/bad state corresponds to a high/low weekly engagement of the beneficiary. We use anonymized data from a quality improvement study conducted in January 2022 [46]. For each beneficiary, we have access to their income, education, and age level, which we use as our three features. These features are available as categorical variables, hence beneficiaries with small changes in the exact value of income, for instance, would still have same categorical value of income group. Beneficiaries’ historic listenership values are used to estimate their transition probabilities under the passive action [30]. One problem in estimating transition probabilities under the active action is that due to the limited number of service calls made, such transitions are rare. Thus, active transition probability estimates are noisy. To alleviate this issue, we use the features and passive transition probabilities from ARMMAN together with the procedure to sample active transitions described in the synthetic domain. Again, we create datasets for three different weight vectors  $W$ . Each dataset consists of five sampled RMAB instances with  $N = 2100$  arms, a budget of  $B = 210$  and a time horizon of  $T = 12$ .

**Problem Instances** Instances of our problem consist of two parts: A preference prompt and an RMAB instance. We initially focus on prioritization prompts. Specifically, for each feature  $X$ , we consider two different prioritization clauses “Prioritize agents with low/high value of feature  $X$ ”. This gives rise to 6 singular prompts consisting of one prioritization clause, two for each feature. For composite prompts, we take all combinations of two features and the two prioritization clauses for each feature (e.g. “Prioritize agents with high value of feature  $A$  and also prioritize agents with low value of feature  $B$ ”). This results in  $3 \cdot 4 = 12$  composite prompts. For each domain, we run each prompt on the 15 RMAB instances from the three datasets.

## 6.2 Models & Baselines

We analyze four different variants of SCLM differing in the used social welfare function (Utilitarian or Egalitarian) and scorer model (Simulator or LLM), e.g., we denote as *SCLM-SIM-Egalitarian* SCLM with the Simulator Scorer Model and the Egalitarian social welfare function.<sup>1</sup> In our generator, we generate 4 candidate reward functions in each step and run 5 iterations to generate a total of 20 candidate reward functions. In addition, we consider several LLM-focused baselines (see Section E for detailed descriptions):

**LLM-Zeroshot:** This baseline only queries the LLM once. It asks to return a reward function aligned with the given preference prompt and provides the problem and feature description as additional context in the prompt.

**DLM** This baseline implements the Decision-Language Model by Behari et al. [8] (see Section C.2).

**DLM-PromptEngg** This is a modified version of DLM where within the reflection prompt, we include examples for singular queries of how the LLM should reason over the different reward function choices (see Section E).

## 6.3 Performance Evaluation Metrics

As our goal is to fulfill the preferences specified by the user, we need to quantify the alignment of the returned reward function with the given prompt  $P$  to evaluate our models. Due to the composite, multi-objective nature of our prompts, we start by measuring the alignment of the returned reward function  $R$  with each prioritization clause  $p_i \in P$  in a separate evaluation score  $e_i(R)$ . However, as our prompts are written in human language, the subpopulations specified in the prompt are not precisely defined (as the prompts only speak of agents with “high”/“low” value of some feature  $X$ ). To address this issue, we assume that the terms “low” and “high” in the input prompts refer to the most extreme feature values. We compute the evaluation score  $e_i(R)$  as follows:

1. Let  $p_i$  be some prompt prioritizing agents with a high/low value of some feature  $X$ .
2. Calculate  $e_i(R)$ , as the summed utility generated by the agents with highest/lowest value of  $X$  under the policy  $\Pi(R)$  normalized by the utility generated by these agents under the default policy  $\Pi(R^*)$ .

Further, reflecting the multi-objective nature of our problem, we consider two metrics for measuring the alignment of a reward  $R$  with a full composite prompt: the sum and the minimum of evaluation scores  $e_i(R)$  across all prioritization clauses  $p_i \in P$ .

## 6.4 Results: Composite Prioritization Prompts

We analyze the performance on the 12 composite prompts which request the prioritization of two subpopulations (see Section D.1 for additional results).

In Figure 3, we show the averaged results from the synthetic and real-world domain. We depict the average summed and minimum alignment with the two clauses of the composite prompt, i.e., the minimum/summed change in the utility generated by the prioritized group of agents.

We start by focusing on SCLM with Simulator Scorer *SCLM-SIM* (green-shaded bars), our strongest method. On both domains, *SCLM-SIM* significantly outperforms all baselines for both minimum and summed % change independent of whether the Utilitarian or Egalitarian social welfare function is chosen. *SCLM-SIM-Utilitarian* outperforming the baselines for the minimum change and *SCLM-SIM-Egalitarian* outperforming them for the summed change highlights the advantages of the SCLM, as these objectives are not explicitly optimized by the respective models, e.g., *SCLM-SIM-Utilitarian* aims at maximizing the summed change and not the minimum one, but still performs well regarding the minimum change. This indicates that SCLM independent of the chosen welfare function does a better job at picking effective and aligned reward functions (on the Pareto front). Comparing *SCLM-SIM-Utilitarian* and *SCLM-SIM-Egalitarian*, the two methods exhibit a big difference under the summed change criterion, while the difference regarding the minimum change is much smaller. Examining individual instances we find in Section D.1 that both functions lead to very different

<sup>1</sup>We also ran our experiments for the Nash welfare function. Its performance was very similar, yet slightly inferior to the Utilitarian welfare function in all relevant evaluation dimensions.



Table 1: Results comparing different methods for aligning with prioritization prompts while minimizing unintended utility shifts. A higher summed utility change implies a better alignment, whereas less unintended shift is better.

|        | Synthetic Domain                              |                  | Real-World Domain                             |                   |
|--------|---|------------------|---|-------------------|
|        | Summed % Utility Change in Desired Feature(s) | Unintended Shift | Summed % Utility Change in Desired Feature(s) | Unintended Shift  |
| DLM-OP | -0.6±0.58                                     | 2.315±0.05       | 2.016±1.39                                    | 0.276±0.03        |
| DLM-EP | -1.882±0.87                                   | 2.346±0.08       | -0.254±1.45                                   | 0.276±0.04        |
| SCLM   | <b>4.854±0.55</b>                             | <b>0.12±0.01</b> | <b>6.01±1.8</b>                               | <b>0.062±0.02</b> |

selections on the instance level; unsurprisingly, the Egalitarian method creates rewards that benefit both groups in a more balanced fashion.

If we replace the Simulation Scorer with the LLM Scorer, the performance of the SCLM decreases, but is still better than all of our three baselines. The difference between the LLM and Simulation Scorer highlights the advantage of the additional information acquired through the more complex and computationally expensive simulation method. Regarding the performance of the baselines, our DLM baseline with prompt engineering *DLM-PromptEngg* improves slightly upon the results of DLM in the synthetic domain, while in the real-world domain, their performance is similar. This suggests that prompt engineering itself is not sufficient to adequately deal with the multi-objective nature of composite prompts; an external component (like our adjudicator) is needed. Finally, LLM-zeroshot consistently performs the worst, which highlights the non-triviality of the reward design problem and the need for a guided extensive search within the reward function search space.

### 6.5 Results: Utility Shifts & Cumulative Utility

As discussed in Section 5.2, we can also use our pipeline to prevent unintended side-effects of aligning reward functions with prioritization clauses, i.e., shifts in the utility feature distribution of features not included in the prompt. We analyze all 6 singular and 12 composite prompts, where we add additional clauses to prevent shifts in the utility distribution of all features not referenced in the prompt. We use *SCLM-SIM-Utilitarian* as our method. As baselines, we consider DLM only prompted with the prioritization clause(s) (called DLM-OP) and DLM prompted with the prioritization clause(s) and clause(s) asking for the minimizing of utility shifts (called DLM-EP). To compute the alignment with the prioritization clauses, similar to Section 6.4, we compute the average change in the utility generated by the prioritized subpopulations. To quantify the unintended utility shifts, we compute the average Earth mover’s distance between the utility feature distribution under the candidate and default reward function for each feature not included in one of the prioritization clauses. Table 1 shows the results. Comparing *DLM-OP* and *DLM-EP*, we find that adding the additional objective to the prompt does not result in a better performance for the synthetic and real-world domains. In contrast, SCLM chooses reward functions resulting in significantly higher utility increases for the prioritized subpopulations and significantly fewer unintended utility shifts. The fact that SCLM performs advantageously for both (conflicting) objectives highlights the quality of the pipeline and its capabilities to effectively address multiple objectives (of different types).

## 7 Discussion

We present a customizable Social Choice Language Model to handle the multi-objective nature of preference prompts in reward design for RMABs in a public health setting. We showcase how methods from social choice can be used to improve the quality and transparency of decision-making of LLM-based frameworks, as we present an adjudicator component that makes the final decision from options generated by the LLM. SCLM significantly improves the quality of the chosen reward functions. We demonstrate that SCLM can not only handle composite prioritization prompts but arbitrary prompts containing multiple objectives, e.g., balancing the prioritization of subpopulations with the total utility generated by all agents. For future work, SCLM can be applied to other problems from multiagent planning and reinforcement learning. Further, SCLM can easily be extended to handle multiple preference prompts specified by multiple health workers or end users.

## References

- [1] Abderrahmane Abbou and Viliam Makis. Group maintenance: A restless bandits approach. *INFORMS J. Comput.*, 31(4):719–731, 2019. doi: 10.1287/IJOC.2018.0863. URL <https://doi.org/10.1287/ijoc.2018.0863>.
- [2] Rohan Anil et al. Gemini: A family of highly capable multimodal models. *CoRR*, abs/2312.11805, 2023. doi: 10.48550/ARXIV.2312.11805. URL <https://doi.org/10.48550/arXiv.2312.11805>.
- [3] ARMMAN. ARMMAN: Advancing Reduction in Mortality and Morbidity of Mothers, Children, and Neonates, 2024. <https://armman.org/>.
- [4] Kenneth J Arrow, Amartya Sen, and Kotaro Suzumura. *Handbook of social choice and welfare*. Elsevier, 2010.
- [5] Turgay Ayer, Can Zhang, Anthony Bonifonte, Anne C Spaulding, and Jagpreet Chhatwal. Prioritizing hepatitis c treatment in us prisons. *Operations Research*, 67(3):853–873, 2019.
- [6] Michiel A. Bakker, Martin J. Chadwick, Hannah Sheahan, Michael Henry Tessler, Lucy Campbell-Gillingham, Jan Balaguer, Nat McAleese, Amelia Glaese, John Aslanides, Matt M. Botvinick, and Christopher Summerfield. Fine-tuning language models to find agreement among humans with diverse preferences. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Annual Conference on Neural Information Processing Systems 2022*, 2022. URL [http://papers.nips.cc/paper\\_files/paper/2022/hash/f978c8f3b5f399cae464e85f72e28503-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/f978c8f3b5f399cae464e85f72e28503-Abstract-Conference.html).
- [7] Shreyas Basavatia, Keerthiram Murugesan, and Shivam Ratnakar. Starling: Self-supervised training of text-based reinforcement learning agent with large language models. *arXiv preprint arXiv:2406.05872*, 2024.
- [8] Nikhil Behari, Edwin Zhang, Yunfan Zhao, Aparna Taneja, Dheeraj Nagaraj, and Milind Tambe. A decision-language model (DLM) for dynamic restless multi-armed bandit tasks in public health. *CoRR*, abs/2402.14807, 2024. doi: 10.48550/ARXIV.2402.14807. URL <https://doi.org/10.48550/arXiv.2402.14807>.
- [9] Daan Bloembergen, Karl Tuyls, Daniel Hennes, and Michael Kaisers. Evolutionary dynamics of multi-agent learning: A survey. *Journal of Artificial Intelligence Research*, 53:659–697, 2015.
- [10] Yuji Cao, Huan Zhao, Yuheng Cheng, Ting Shu, Guolong Liu, Gaoqi Liang, Junhua Zhao, and Yun Li. Survey on large language model-enhanced reinforcement learning: Concept, taxonomy, and methods. *CoRR*, abs/2404.00282, 2024. doi: 10.48550/ARXIV.2404.00282. URL <https://doi.org/10.48550/arXiv.2404.00282>.
- [11] Thomas Carta, Pierre-Yves Oudeyer, Olivier Sigaud, and Sylvain Lamprier. Eager: Asking and answering questions for automatic reward shaping in language-guided rl. *Advances in Neural Information Processing Systems*, 35:12478–12490, 2022.
- [12] Vincent Conitzer, Rachel Freedman, Jobst Heitzig, Wesley H. Holliday, Bob M. Jacobs, Nathan Lambert, Milan Mossé, Eric Pacuit, Stuart Russell, Hailey Schoelkopf, Emanuel Tewolde, and William S. Zwicker. Social choice for AI alignment: Dealing with diverse human feedback. *CoRR*, abs/2404.10271, 2024. doi: 10.48550/ARXIV.2404.10271. URL <https://doi.org/10.48550/arXiv.2404.10271>.
- [13] Jessica Dai and Eve Fleisig. Mapping social choice theory to RLHF. *CoRR*, abs/2404.13038, 2024. doi: 10.48550/ARXIV.2404.13038. URL <https://doi.org/10.48550/arXiv.2404.13038>.
- [14] Devleena Das, Sonia Chernova, and Been Kim. State2explanation: Concept-based explanations to benefit agent learning and user understanding. *Advances in Neural Information Processing Systems*, 36:67156–67182, 2023.
- [15] Katrina V Deardorff, Arianna Rubin Means, Kristjana H Ásbjörnsdóttir, and Judd Walson. Strategies to improve treatment coverage in community-based public health programs: a systematic review of the literature. *PLoS neglected tropical diseases*, 12(2):e0006211, 2018.

- [16] Sara Fish, Paul Gözl, David C. Parkes, Ariel D. Procaccia, Gili Rusak, Itai Shapira, and Manuel Wüthrich. Generative social choice. *CoRR*, abs/2309.01291, 2023. doi: 10.48550/ARXIV.2309.01291. URL <https://doi.org/10.48550/arXiv.2309.01291>.
- [17] Luise Ge, Daniel Halpern, Evi Micha, Ariel D. Procaccia, Itai Shapira, Yevgeniy Vorobeychik, and Junlin Wu. Axioms for AI alignment from human feedback. *CoRR*, abs/2405.14758, 2024. doi: 10.48550/ARXIV.2405.14758. URL <https://doi.org/10.48550/arXiv.2405.14758>.
- [18] Prasoon Goyal, Scott Niekum, and Raymond J Mooney. Using natural language for reward shaping in reinforcement learning. *arXiv preprint arXiv:1903.02020*, 2019.
- [19] Rishi Hazra, Alkis Sygkounas, Andreas Persson, Amy Loutfi, and Pedro Zuidberg Dos Martires. Revolve: Reward evolution with large language models for autonomous driving. *CoRR*, abs/2406.01309, 2024. doi: 10.48550/ARXIV.2406.01309. URL <https://doi.org/10.48550/arXiv.2406.01309>.
- [20] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29, 2016.
- [21] Jiechuan Jiang and Zongqing Lu. Learning fairness in multi-agent systems. *Advances in Neural Information Processing Systems*, 32, 2019.
- [22] Ralph L Keeney and Craig W Kirkwood. Group decision making using cardinal social welfare functions. *Management Science*, 22(4):430–437, 1975.
- [23] Minae Kwon, Sang Michael Xie, Kalesha Bullard, and Dorsa Sadigh. Reward design with language models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL <https://openreview.net/forum?id=10uNUgI5Kl>.
- [24] Pawel Ladosz, Lilian Weng, Minwoo Kim, and Hyondong Oh. Exploration in deep reinforcement learning: A survey. *Information Fusion*, 85:1–22, 2022.
- [25] Shuang Li, Xavier Puig, Chris Paxton, Yilun Du, Clinton Wang, Linxi Fan, Tao Chen, De-An Huang, Ekin Akyürek, Anima Anandkumar, et al. Pre-trained language models for interactive decision-making. *Advances in Neural Information Processing Systems*, 35:31199–31212, 2022.
- [26] Jessy Lin, Yuqing Du, Olivia Watkins, Danijar Hafner, Pieter Abbeel, Dan Klein, and Anca Dragan. Learning to model the world with language. *arXiv preprint arXiv:2308.01399*, 2023.
- [27] Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Eureka: Human-level reward design via coding large language models. *CoRR*, abs/2310.12931, 2023. doi: 10.48550/ARXIV.2310.12931. URL <https://doi.org/10.48550/arXiv.2310.12931>.
- [28] Yecheng Jason Ma, William Liang, Hung-Ju Wang, Sam Wang, Yuke Zhu, Linxi Fan, Osbert Bastani, and Dinesh Jayaraman. Dreureka: Language model guided sim-to-real transfer. *CoRR*, abs/2406.01967, 2024. doi: 10.48550/ARXIV.2406.01967. URL <https://doi.org/10.48550/arXiv.2406.01967>.
- [29] Gabriel B Margolis and Pulkit Agrawal. Walk these ways: Tuning robot control for generalization with multiplicity of behavior. In *Conference on Robot Learning*, pages 22–31. PMLR, 2023.
- [30] Aditya Mate, Lovish Madaan, Aparna Taneja, Neha Madhiwalla, Shresth Verma, Gargi Singh, Aparna Hegde, Pradeep Varakantham, and Milind Tambe. 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*, volume 36, pages 12017–12025, 2022.
- [31] Suvir Mirchandani, Siddharth Karamcheti, and Dorsa Sadigh. Ella: Exploration through learned language abstraction. *Advances in neural information processing systems*, 34:29529–29540, 2021.
- [32] Hervé Moulin. *Fair division and collective welfare*. MIT press, 2004.

- [33] Lyndsay A Nelson, Shelagh A Mulvaney, Tebeb Gebretsadik, Yun-Xian Ho, Kevin B Johnson, and Chandra Y Osborn. Disparities in the use of a mhealth medication adherence promotion intervention for low-income adults with type 2 diabetes. *Journal of the American Medical Informatics Association*, 23(1):12–18, 2016.
- [34] Andrew Y Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *Icml*, volume 99, pages 278–287, 1999.
- [35] José Niño-Mora. Markovian restless bandits and index policies: A review. *Mathematics*, 11(7): 1639, 2023.
- [36] Christos H Papadimitriou and John N Tsitsiklis. The complexity of optimal queueing network control. In *Proceedings of IEEE 9th annual conference on structure in complexity Theory*, pages 318–322. IEEE, 1994.
- [37] Yundi Qian, Chao Zhang, Bhaskar Krishnamachari, and Milind Tambe. Restless poachers: Handling exploration-exploitation tradeoffs in security domains. In Catholijn M. Jonker, Stacy Marsella, John Thangarajah, and Karl Tuyls, editors, *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016*, pages 123–131. ACM, 2016. URL <http://dl.acm.org/citation.cfm?id=2936946>.
- [38] Yundi Qian, Chao Zhang, Bhaskar Krishnamachari, and Milind Tambe. Restless poachers: Handling exploration-exploitation tradeoffs in security domains. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 123–131, 2016.
- [39] John Rawls. A theory of justice. In *Applied ethics*, pages 21–29. Routledge, 2017.
- [40] Diederik M Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48: 67–113, 2013.
- [41] Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: language agents with verbal reinforcement learning. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Annual Conference on Neural Information Processing Systems 2023*, 2023. URL [http://papers.nips.cc/paper\\_files/paper/2023/hash/1b44b878bb782e6954cd888628510e90-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2023/hash/1b44b878bb782e6954cd888628510e90-Abstract-Conference.html).
- [42] Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su. Llm-planner: Few-shot grounded planning for embodied agents with large language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2998–3009, 2023.
- [43] Samina T Syed, Ben S Gerber, and Lisa K Sharp. Traveling towards disease: transportation barriers to health care access. *Journal of community health*, 38:976–993, 2013.
- [44] Peter Vamplew, Richard Dazeley, Adam Berry, Rustam Issabekov, and Evan Dekker. Empirical evaluation methods for multiobjective reinforcement learning algorithms. *Machine learning*, 84:51–80, 2011.
- [45] Kristof Van Moffaert and Ann Nowé. Multi-objective reinforcement learning using sets of pareto dominating policies. *The Journal of Machine Learning Research*, 15(1):3483–3512, 2014.
- [46] Shresth Verma, Gargi Singh, Aditya Mate, Paritosh Verma, Sruthi Gorantla, Neha Madhiwalla, Aparna Hegde, Divy Thakkar, Manish Jain, Milind Tambe, and Aparna Taneja. Expanding impact of mobile health programs: SAHELI for maternal and child care. *AI Mag.*, 44(4): 363–376, 2023. doi: 10.1002/AAAI.12126. URL <https://doi.org/10.1002/aaai.12126>.
- [47] Shresth Verma, Yunfan Zhao, Niclas Boehmer Sanket Shah, Aparna Taneja, and Milind Tambe. Group fairness in predict-then-optimize settings for restless bandits. [openreview.net/pdf?id=GJIZbpLWX3](https://openreview.net/pdf?id=GJIZbpLWX3), 2024.

- [48] Peter Whittle. Restless bandits: Activity allocation in a changing world. *Journal of applied probability*, 25(A):287–298, 1988.
- [49] Tianbao Xie, Siheng Zhao, Chen Henry Wu, Yitao Liu, Qian Luo, Victor Zhong, Yanchao Yang, and Tao Yu. Text2reward: Reward shaping with language models for reinforcement learning. In *The Twelfth International Conference on Learning Representations, ICLR 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=tUM39YTRxH>.
- [50] Tianbao Xie, Siheng Zhao, Chen Henry Wu, Yitao Liu, Qian Luo, Victor Zhong, Yanchao Yang, and Tao Yu. Text2reward: Reward shaping with language models for reinforcement learning. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=tUM39YTRxH>.
- [51] Runzhe Yang, Xingyuan Sun, and Karthik Narasimhan. A generalized algorithm for multi-objective reinforcement learning and policy adaptation. *Advances in neural information processing systems*, 32, 2019.
- [52] Wenhao Yu, Nimrod Gileadi, Chuyuan Fu, Sean Kirmani, Kuang-Huei Lee, Montserrat Gonzalez Arenas, Hao-Tien Lewis Chiang, Tom Erez, Leonard Hasenclever, Jan Humprik, Brian Ichter, Ted Xiao, Peng Xu, Andy Zeng, Tingnan Zhang, Nicolas Heess, Dorsa Sadigh, Jie Tan, Yuval Tassa, and Fei Xia. Language to rewards for robotic skill synthesis. In Jie Tan, Marc Toussaint, and Kourosh Darvish, editors, *Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pages 374–404. PMLR, 2023. URL <https://proceedings.mlr.press/v229/yu23a.html>.
- [53] Chen Bo Calvin Zhang and Giorgia Ramponi. Hip-rl: Hallucinated inputs for preference-based reinforcement learning in continuous domains. In *ICML 2023 Workshop The Many Facets of Preference-Based Learning*, 2023.
- [54] Matthieu Zimmer, Claire Glanois, Umer Siddique, and Paul Weng. Learning fair policies in decentralized cooperative multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 12967–12978. PMLR, 2021.

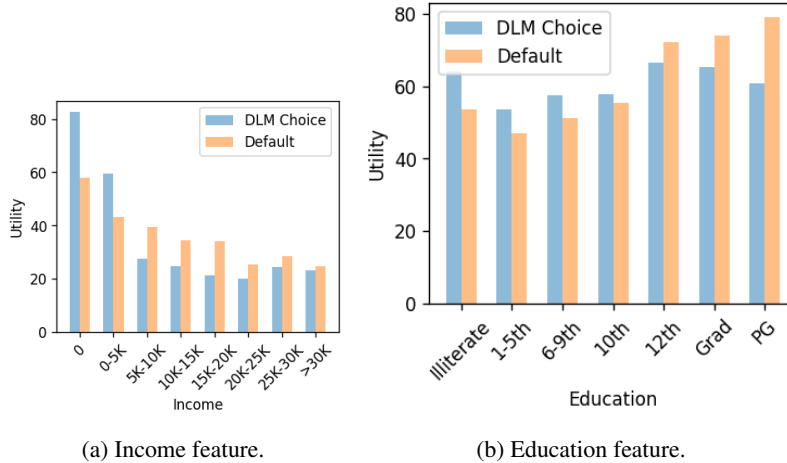


Figure 4: Utility feature distributions for the default reward function (orange) and reward function returned for prompt “Prioritize agents with low income” (blue). The  $x$ -axis depicts the values of the feature and the  $y$ -axis the total utility generated by all agents with this feature value.

## Appendix

### A Additional Related Work

**Multi-Objective Reinforcement Learning** Real-world decision-making often involves trade-offs between multiple, conflicting objectives, making multi-objective reinforcement learning (MORL) essential. A common approach in MORL is to scalarize multiple objectives into a single reward function, but this can oversimplify the problem. To address this, Pareto-based methods identify Pareto-optimal solutions that reflect different trade-offs between objectives, providing a more comprehensive decision-making framework [40, 45]. Recent advancements in MORL have introduced algorithms that enhance scalability, handle non-linear objectives, and improve exploration strategies. For example, [51] proposed an algorithm that adapts policies based on changing objective importance. Additionally, empirical methods have been developed to evaluate MORL performance across scenarios [44].

In multi-agent systems, MORL adds complexity as agents must balance their objectives while considering others’. Specialized approaches, such as extending Pareto-based methods to multi-agent contexts, aim to find Pareto-efficient solutions for all agents involved [9]. Moreover, the concept of fairness becomes important in multi-agent MORL, as solutions must be equitable to all agents while still optimizing overall system performance. Researchers have begun to investigate fairness-aware MORL algorithms that explicitly consider the distribution of rewards among agents, ensuring that no single agent disproportionately benefits at the expense of others [21, 54].

**(Pre-LLM) Reward Function Design** Designing effective reward functions for RL has long been a challenge due to the lack of robust learning signals. [24]. Previous approaches for reward design in RL include manually specifying reward components and tuning their coefficients [34, 29], and learning reward models from demonstrations [20] or preferences [53]. While promising, these methods often demand substantial domain expertise, large amounts of data, or computationally expensive processes.

**LLMs and Social Choice** In our adjudicator, we take a social choice perspective to model the problem of selecting a final reward function in a way that balances different objectives against each other. Principles from social choice have informed LLM-based algorithm in previous works, albeit in the very different context of creating a consensus statements for a group of users with diverse views [6, 16]. Moreover, social choice principles are also increasingly used to inform the fine-tuning of language models, including its alignment with human feedback [6, 12, 17, 13].

## B Additional Material for Section 4

**Unintended Utility Shifts** Moving from the default reward function to a reward function aligned with a given (prioritization) prompt causes shifts within the distribution of resources and utility. Due to correlations between features, this might change might lead to unintended utility shifts for features not specified in the prompt. In Figure 4, we show an instance of our experiments. We present the utility feature distribution for the two features income and education for two reward functions: The reward function selected by DLM for the prompt “Prioritize agents with low income” (orange) and the default reward (blue). While the utility generated by low-income agents increases, the utility generated by highly educated agents decreases, a side-effect a user is likely unaware of and that might conflict with their allocation goals.

## C Implementation Details

### C.1 Whittle Index Policy

To formulate the Whittle Index Policy, first we define the value function for an arm  $i \in [N]$  under the subsidy  $\lambda$  as

$$V_i^\lambda(s) = \max_{a \in \{0,1\}} Q_i(s, a_i, \lambda). \tag{1}$$

Here,  $Q_i(s, a_i, \lambda)$  measures the expected discounted cumulative future reward where a subsidy  $\lambda$  is added to the reward when the passive action is taken. The Whittle index associated to the state  $s_i$  is then defined as:

$$W_i(s_i) := \inf_m \{Q_i(s_i, a_i = 0, m) = Q_i(s_i, a_i = 1, m)\}.$$

The whittle index is thus the value of subsidy such that the passive ( $a = 0$ ) and active ( $a = 1$ ) actions have the same value. Intuitively, this captures the value of taking active action on an arm.

To implement the Whittle Index computation, we use the method in [38] based on binary search. Additionally, for all the experiments, we cache computed whittle index for a given set of transition probabilities and reward functions to reduce computation time.

### C.2 DLM Pipeline

In our work, we use a modified version of the DLM pipeline [8], which employs the Whittle Index policy as the planner for RMAB. Our approach differs from Behari et al. [8] in that we use the Whittle Index policy specifically for simulating RMAB, whereas Behari et al. [8] use the PreFeRMAB policy (Zhao and Behari et al. 2023). This modification allows for faster and more stable simulations and effectively decouples the learning problem from the planning problem.

Specifically, the DLM consists of three components. First, a user provides a natural language preference  $P$ . We then create a prompt for LLM which includes the context of the RMAB problem, the preference  $P$ , a description of features available, and the index of those features in the feature array. Finally, the LLM is asked to generate the Python code of the reward function in text format. We describe the prompt used in Figures 6 and 7.

The LLM is queried  $n_p$  times to obtain  $n_p$  reward functions. For each reward function, we also compute the reward distribution over all the features. Next, given all the generated reward functions and their corresponding reward distributions, we query the LLM to select the best reward function. We describe the prompt used in Figures 8 and 9. This is called the reflection step. The best reward function is then used inside the prompt for next step of generating  $n_p$  reward functions. This process is repeated  $n_r$  times to obtain  $n_p \cdot n_r$  reward functions. In all our reward function generation experiments, we query the LLM  $n_p = 4$  times and run the reflection loop  $n_r = 5$  times resulting in 20 candidate reward functions.

As an LLM, we use the Gemini Pro model by Google. We query the LLM using python based API from the *generative-ai-python* library.

Table 2: Weight vector parameters for different synthetic datasets.

|                  | Feature A | Feature B | Feature C | $\sigma$ |
|------------------|-----------|-----------|-----------|----------|
| <b>Dataset 1</b> | 0.8       | -1.5      | 1         | 0.1      |
| <b>Dataset 2</b> | 10        | -1.5      | 1         | 0.1      |
| <b>Dataset 3</b> | 1         | -1.5      | 10        | 0.1      |

### C.3 Synthetic Dataset Generation

An RMAB problem is defined by the transition probabilities of Markov Decision Process governing every arm and the reward functions. We consider a 2-state, 2-action Markov decision process in all our experiments. This results in 8 transition probability parameters  $P_i(s, a, s') \forall s \in \{0, 1\}, a \in \{0, 1\}, s' \in \{0, 1\}$  for every arm  $i$ . Out of these 8 parameters, we only need to define 4 parameters  $P_i(s, a, s' = 1) \forall s \in \{0, 1\}, a \in \{0, 1\}$  independently, and the rest 4 parameters can be calculated as the compliment values  $P_i(s, a, s' = 0) = 1 - P_i(s, a, s' = 1) \forall s \in \{0, 1\}, a \in \{0, 1\}$ .

To simulate the effects of conflicting preferences and trade-offs in a controlled setup, we configure the four transition probability parameters to depend on the features describing each arm. We consider each arm to be characterized by a vector of continuous features  $f$ , with three values ranging between 0 and 1 ( $f \in [0, 1]^3$ ). We use the following setup to stochastically generate the 4 transition probability parameters.

1. For every arm  $i$ , uniformly sample the two passive probability parameters  $P_i(s, a = 0, s' = 1) \sim \text{Uniform}(0, 1), \forall s \in \{0, 1\}$ .
2. For every arm  $i$ , uniformly sample the three features  $f \sim \text{Uniform}([0, 1]^3)$
3. Define a three-dimensional weight vector  $W \in [0, 1]^3$  and a standard deviation parameter  $\sigma$ .
4. Sample the effect of intervention from the normal distribution as  $\delta \sim \mathcal{N}(\Delta, \sigma)$ , where  $\Delta = W \cdot f^T$ .
5. Calculate active transition probabilities as  $P(s, a = 1, s') = P(s, a = 0, s') + \delta$
6. Calculate the complimentary probabilities as  $P_i(s, a, s' = 0) = 1 - P_i(s, a, s' = 1), \forall s \in \{0, 1\}, a \in \{0, 1\}$

The magnitude of the weight value determines the extent to which a feature influences the effect of the intervention, while the direction indicates whether a low or high feature value amplifies this effect. The standard deviation parameter,  $\sigma$ , controls the spread of sampled probabilities around the mean effect of the intervention. Table 2 shows the Weight vector used for the three synthetic datasets generated in our experiments. We consider weight values that describe the following scenarios : i) all weights are roughly equal (dataset 1) ii) one of the features has much higher weight than the other two (dataset 2) and we switch which feature has maximum weight (dataset 3).

### C.4 Real World Dataset

As described in Section 6.1, one challenge in estimating transition probabilities under active actions is the rarity of such transitions due to the limited number of service calls. Consequently, the estimates for active transition probabilities tend to be noisy. To mitigate this issue, we follow a procedure similar to that in the synthetic domain to construct the datasets, but with some modifications. Specifically, instead of uniformly sampling passive probabilities in step 1, we use the transition probabilities estimated from real-world data [46]. Next, in step 2, we use the attributes describing the real-world beneficiaries to define features. Thus, instead of features A, B and C, we use features Age, Income and Education. Finally, we use the same weight vector values as in Table 2 to generate three real world datasets. This allows us to generate multiple datasets with varying levels of effect of intervention based on features while still using the realistic passive transition probabilities and feature distributions.



Table 3: Results summary comparing different reward function choice strategies aggregated across the three synthetic Datasets. Cells in bold indicate the top 2 best values (higher is better). DLM: Decision-Language Model, SCLM: Social Choice Language Model, SCLM-SIM: Simulation based Scorer Model, SCLM-LLM: LLM based Scorer Model

| Method         | Social Welfare Function | Minimum of % change in utility |                      |                      | Minimum of % change in utility |                      |                      |
|----------------|-------------------------|--------------------------------|----------------------|----------------------|--------------------------------|----------------------|----------------------|
|                |                         | top/bottom 1 bucket            | top/bottom 2 buckets | top/bottom 3 buckets | top/bottom 1 bucket            | top/bottom 2 buckets | top/bottom 3 buckets |
| LLM-zeroshot   |                         | -0.266%±1.51%                  | -0.228%±1.3%         | -3.269%±1.14%        | -15.16%±0.81%                  | -13.249%±0.71%       | -12.485%±0.61%       |
| DLM            |                         | 3.879%±1.42%                   | 5.379%±1.27%         | 1.89%±0.99%          | -7.844%±0.74%                  | -6.249%±0.62%        | -6.761%±0.49%        |
| DLM-PromptEngg |                         | 7.607%±1.53%                   | 8.665%±1.44%         | 4.087%±1.24%         | -9.301%±0.82%                  | -7.17%±0.81%         | -7.767%±0.71%        |
| SCLM-SIM       | Utilitarian             | <b>28.944%±2.12%</b>           | <b>21.936%±1.55%</b> | <b>13.654%±1.07%</b> | <b>-2.278%±1.02%</b>           | <b>-3.579%±0.83%</b> | <b>-3.559%±0.55%</b> |
| SCLM-LLM       |                         | 14.348%±1.66%                  | 10.304%±1.22%        | 6.333%±0.99%         | -3.973%±0.76%                  | -4.428%±0.6%         | -4.817%±0.54%        |
| SCLM-SIM       | Egalitarian             | 16.448%±1.95%                  | 11.425%±1.34%        | 8.6%±0.94%           | <b>-1.176%±1.01%</b>           | <b>-2.028%±0.71%</b> | <b>-1.833%±0.49%</b> |
| SCLM-LLM       |                         | 11.421%±1.62%                  | 7.845%±1.21%         | 4.141%±0.92%         | -4.877%±0.73%                  | -4.373%±0.59%        | -4.68%±0.5%          |
| SCLM-SIM       | Nash                    | <b>28.262%±2.42%</b>           | <b>20.416%±1.73%</b> | <b>11.102%±1.19%</b> | -4.053%±1.18%                  | -5.408%±0.89%        | -5.261%±0.62%        |
| SCLM-LLM       |                         | 9.478%±1.68%                   | 6.608%±1.23%         | 2.957%±1.02%         | -5.973%±0.77%                  | -5.782%±0.6%         | -6.117%±0.54%        |

Table 4: Results summary comparing different reward function choice strategies aggregated across the three Real World Datasets. Cells in bold indicate the top 2 best values (higher is better). DLM: Decision-Language Model, SCLM: Social Choice Language Model, SCLM-SIM: Simulation based Scorer Model, SCLM-LLM: LLM based Scorer Model

| Method         | Social Welfare Function | Minimum of % change in utility |                      |                      | Minimum of % change in utility |                      |                      |
|----------------|-------------------------|--------------------------------|----------------------|----------------------|--------------------------------|----------------------|----------------------|
|                |                         | top/bottom 1 bucket            | top/bottom 2 buckets | top/bottom 3 buckets | top/bottom 1 bucket            | top/bottom 2 buckets | top/bottom 3 buckets |
| LLM-zeroshot   |                         | -0.893%±1.38%                  | -5.541%±0.79%        | -10.436%±0.46%       | -7.285%±0.69%                  | -7.003%±0.45%        | -6.971%±0.26%        |
| DLM            |                         | 6.219%±1.59%                   | -0.485%±0.95%        | -8.733%±0.43%        | -5.317%±0.84%                  | -6.454%±0.39%        | -7.674%±0.19%        |
| DLM-PromptEngg |                         | 4.341%±1.65%                   | -3.57%±0.91%         | -9.417%±0.57%        | -4.957%±0.94%                  | -7.366%±0.5%         | -7.586%±0.37%        |
| SCLM-SIM       | Utilitarian             | <b>21.502%±1.76%</b>           | <b>10.643%±1.06%</b> | <b>-3.434%±0.53%</b> | <b>2.206%±0.89%</b>            | <b>-1.203%±0.45%</b> | <b>-4.423%±0.29%</b> |
| SCLM-LLM       |                         | 8.711%±1.3%                    | -0.574%±1.08%        | -8.285%±0.82%        | -6.512%±0.63%                  | -6.457%±0.55%        | -6.756%±0.43%        |
| SCLM-SIM       | Egalitarian             | 11.481%±1.61%                  | 7.92%±1.01%          | <b>-0.416%±0.37%</b> | <b>2.109%±0.75%</b>            | <b>0.287%±0.4%</b>   | <b>-1.974%±0.22%</b> |
| SCLM-LLM       |                         | 8.239%±1.67%                   | -1.919%±0.98%        | -9.488%±0.48%        | -1.561%±0.8%                   | -4.67%±0.39%         | -7.059%±0.2%         |
| SCLM-SIM       | Nash                    | <b>20.579%±1.8%</b>            | <b>10.195%±1.1%</b>  | -4.376%±0.54%        | 2.09%±0.9%                     | -1.568%±0.47%        | -4.886%±0.29%        |
| SCLM-LLM       |                         | -0.732%±1.25%                  | -4.854%±0.95%        | -10.455%±0.68%       | -9.877%±0.65%                  | -8.373%±0.53%        | -8.107%±0.36%        |

## C.5 Hyperparameters

We run all experiments with  $N = 2100$  arms,  $B = 210$  as budget,  $T = 12$  weeks and  $\gamma = 0.9$  as discount factor. For every dataset, 5 RMAB instances are generated based on the different weight vectors as described in the section above. Additionally, we run each RMAB simulation with 10 different seeds. We estimate the cumulative rewards and feature-level utility by calculating the mean of the discounted sum of rewards over  $T$  timesteps across all 10 seeds.

## C.6 Computing Resources

All experiments are run on a machine with 16GB RAM and M1 chip CPU. For every RMAB instance, it took roughly 3 hours to generate 360 candidate reward functions (20 reward functions for each of the 18 prompts). The primary bottleneck in speed was the API call limits for using LLM (Gemini Pro).

## D Additional Results

### D.1 Composite Prioritization Prompts

In Tables 3 and 4, we show results aggregated for the synthetic and real world datasets. Specifically, we show the Sum of % change in utility and Minimum of % change in utility not just for the highest/lowest value of a feature but for top/bottom two and three buckets. We also include results from the scorer model that optimizes for Nash Social Welfare function.

Overall, we observe that the Simulator-based scorer (SCLM-SIM) performs best in all scenarios, both when optimizing for the Utilitarian objective or the Egalitarian objective. It is also worth noting that when optimizing for one of the objectives (for instance, Utilitarian objective maximizes the sum of % change in utility), SCLM outperforms baselines in the objective it is not explicitly optimizing for (for instance, Minimum of % change in utility). Lastly, we observe that optimizing for the Nash objective yields very similar performance as optimizing for the Utilitarian objective.

Table 5: Results comparing different reward function choice selection strategies for deciding tradeoff between utility maximization and preference alignment, aggregated across the three synthetic datasets (left) and three real world datasets (right). Social Choice Language Model has the highest percent change in reward for the desired attribute while also ensuring the low drop in utility as compared to default.

| Method                 | Synthetic                     |                      |                          |                    | Real World                    |                      |                          |                    |
|------------------------|-------------------------------|----------------------|--------------------------|--------------------|-------------------------------|----------------------|--------------------------|--------------------|
|                        | % Change in Desired Attribute |                      | Utility Metrics          |                    | % Change in Desired Attribute |                      | Utility Metrics          |                    |
|                        | top/bottom 1 bucket           | top/bottom 2 buckets | Normalized Utility Score | % drop in Utility  | top/bottom 1 bucket           | top/bottom 2 buckets | Normalized Utility Score | % drop in Utility  |
| DLM-PrioritizationOnly | -0.812±0.85                   | 0.072±0.83           | 0.507±0.01               | -4.867±0.21        | 5.678±0.67                    | 1.177±0.43           | 0.327±0.01               | -6.051±0.14        |
| DLM-ExtendedPrompt     | -2.508±1.2                    | -1.089±1.26          | 0.548±0.02               | -3.825±0.29        | 3.416±1.04                    | 0.461±0.64           | 0.347±0.02               | -5.689±0.25        |
| SCLM                   | 12.389±1.01                   | <b>10.276±0.8</b>    | <b>0.71±0.01</b>         | <b>-1.141±0.16</b> | <b>9.058±0.61</b>             | <b>4.868±0.4</b>     | <b>0.532±0.01</b>        | <b>-3.662±0.16</b> |

## D.2 Effect of Social Choice Function

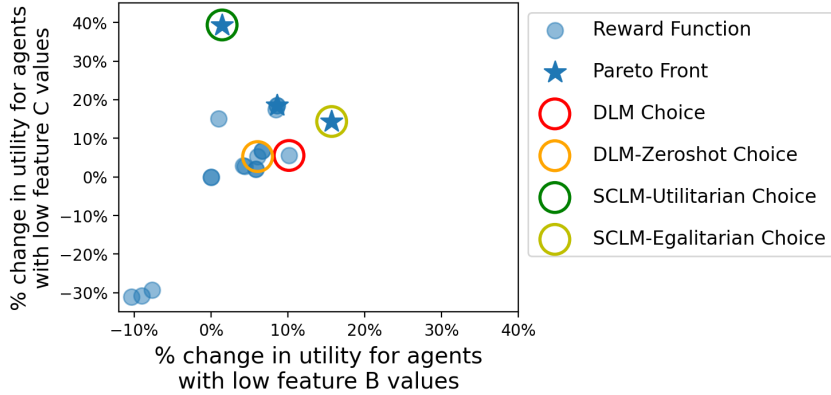


Figure 5: Comparison of different methods using the prompt “Prioritize agents with a low value for feature  $B$  and agents with a low value for feature  $C$ ”.

While both the Utilitarian and Egalitarian social welfare function leads to good overall results, the choice between them significantly influences which reward function is selected at the instance level. Figure 5 shows one example instance for a prompt with two prioritization clauses. As in Figure 1, each point corresponds to a candidate reward function with the axes measuring alignment with the two respective clauses. (as described in the beginning of the section 6.3). The SCLM-Utilitarian model chooses a reward function from the Pareto frontier that shows a much stronger effect for the second clause (i.e., the utility increase for agents with a low value of feature  $C$  is more pronounced). In contrast, the reward function selected by SCLM-Egalitarian is much more balanced. DLM-based baselines fail to pick a reward function from the Pareto frontier.

## D.3 Tradeoffs in preferences and cumulative utility

To show the general applicability of our framework, we consider prompts where the user can specify preferences for multiple attributes and also specify a preference clause that prefers maximizing utility. For all 6 singular and 12 composite prompts, we add additional clauses to maximize utility (or minimize drop in utility). As a baseline, we consider two models: DLM prompted only by the prioritization clause (DLM-Prioritization Only) and DLM prompted by both the prioritization clause(s) and a clause requesting the maximization of overall utility (DLM-ExtendedPrompt).

To assess alignment with prioritization clauses, we calculate the average change in utility for the prioritized subgroup, following the method outlined in Section 6.3. To quantify the score for maximizing utility, we use the utility generated by the default reward function, normalizing this value between 0 and 1 to ensure all score components are on the same scale. Finally, with scores for all prompt components, we can apply a user-defined Social Welfare Function.

Table 5 shows the results. Comparing *DLM-Prioritization Only* and *DLM-ExtendedPrompt*, we find that adding the additional objective to the prompt does not result in a better performance on both the synthetic and real-world domains. In contrast, SCLM chooses reward functions resulting in significantly higher utility increases for the prioritized subpopulations and significantly higher utility

(measured through Normalized Utility Score) and lower drop in utility as compared to default reward function. The fact that SCLM performs advantageously for both (conflicting) objectives highlights the quality of the pipeline and its capabilities to effectively address multiple objectives (of different types).

## E Prompt Texts

In Figures 6-15, we show the prompts passed to LLM for various experiments described in the paper. Specifically, prompt in Figures 6,7 show how to generate a reward function in code form based on the problem description, the indices of features, and preference goals for synthetic and real-world problem domains, respectively. Figures 8, 9 show how the LLM is prompted to select the best reward function from those generated in the previous step for synthetic and real-world domains, respectively. Together, these prompts establish the DLM baseline in the paper. LLM-Zeroshot baseline has the exact same prompt as Figures 6, 7 for synthetic and real world domains. The only difference is that there is no reflection step, and the first reward generated by LLM is chosen as the best reward.

In Figure 10 and 11, we show how the prompt is enhanced with additional information to assist in selecting a reward function. The additional information is highlighted in bold. Section 6.4 explains how SCLM is used to specify additional objectives, such as minimizing utility shifts in features not included in the preference prompt and maximizing overall utility. In Figures 12, 13 we show the prompts for the DLM-EP baselines, which explicitly include instructions to minimize utility shifts in unintended features. In Figures 14, 15, we show the prompts that explicitly include instructions for maximizing overall utility.

## F Ethics Statement

### ARMMAN domain

**Data Usage, Collection and Consent** The data used for the realworld domain contains fully-anonymized datasets. Our experiments constitute secondary analysis of the data and are approved by ARMMAN’s ethics board. The paper does not involve any realworld deployment of the proposed algorithms for ARMMAN. For data collection consent is taken from each beneficiary in ARMMAN’s automated voice call program. Data exchange and use was regulated through clearly defined exchange protocols including anonymization, read-access only to researchers, restricted use of the data for research purposes only, and approval by ARMMAN’s ethics review committee.

**Accessibility of Health Information** In our simulation experiments, we change the reward functions for every agent in mobile health program. But this only improves the quality of service calls and no health information is withheld from any agent. Participants in the program can still request service calls through free missed call service.

## Generator: Synthetic Domain

### Prompt

Create a Python reward function for RL in resource allocation problem for agents, with the objective of prioritizing higher states and: Focus on the agents with low value of feature A.. The function should use 'state' (value is either 0,1) and features 'agent\_feats' (length 15 array) to direct the RL agent. Here is a description of the features you may use along with the index in 'agent\_feats' array:

Index Name DataType

0. Feature A bucket 1 - Binary
1. Feature A bucket 2 - Binary
2. Feature A bucket 3 - Binary
3. Feature A bucket 4 - Binary
4. Feature A bucket 5 - Binary
5. Feature B bucket 1 - Binary
6. Feature B bucket 2 - Binary
7. Feature B bucket 3 - Binary
8. Feature B bucket 4 - Binary
9. Feature B bucket 5 - Binary
10. Feature C bucket 1 - Binary
11. Feature C bucket 2 - Binary
12. Feature C bucket 3 - Binary
13. Feature C bucket 4 - Binary
14. Feature C bucket 5 - Binary

All buckets are in increasing order of the feature values. For example, 'Feature A bucket 1' would consist of bucket of lowest values of feature A while 'Feature A bucket 5' would consist of highest value of feature A. This is true for Feature A, B and C.

Your task: 1. Write a simple, single-line Python reward function. Exclude the word 'return' and exclude non-standard libraries. Format your code with triple \$ signs: \$\$\$YOUR FUNCTION\$\$\$ Note that HIGHER states are always preferred, so ensure reward increases as state increases. Make sure reward is always positive and increasing with state.

Example Prompt: Prioritize agents that have low feature A and high feature C

Example Response: Python Code: '\$\$\$ state+state \* ((agent\_feats[0] or agent\_feats[1] and (agent\_feats[17] or agent\_feats[18] or agent\_feats[19])) \$\$\$' or '\$\$\$ state \* (agent\_feats[0] or 3\*agent\_feats[19]) \$\$\$' or '\$\$\$ state + 2\*state \* ((5\*agent\_feats[0]+agent\_feats[1] and agent\_feats[19]) \$\$\$' In these example, agent\_feats[0] and agent\_feats[1] represent agents with low values for feature A and agent\_feats[17], agent\_feats[18], agent\_feats[19] represent agents with high values for feature C

It is upto you to decide which features will represent a preference. For example low values could be the lowest feature bucket, or lower three feature buckets or so on. Come up with a unique new reward for the specified goal: Focus on the agents with low value of feature A.. Goal: Focus on the agents with low value of feature A.

### Output

```
$$$ 2*state + state * (1*agent_feats[0]+ 0.5*agent_feats[1]) $$$
```

Figure 6: Prompt passed to the LLM to generate reward function based on the problem scenarios in the Synthetic Domain.

### Generator: Real World Domain

#### Prompt

Create a Python reward function for RL in phone call resource allocation to mothers in India, with the objective of prioritizing higher states and: Focus on the young mothers by age. The function should use 'state' (value is either 0,1) and features agent feats (length 43 array) to direct the RL agent. Here is a description of the features you may use along with the index in agent feats array:

- | Index | Name  | Data Type |
|-------|---|-----------|
| 0.    | Ages 10-20                                  | - Binary  |
| 1.    | Ages 21-30                                  | - Binary  |
| 2.    | Ages 31-40                                  | - Binary  |
| 3.    | Ages 41-50                                  | - Binary  |
| 4.    | Ages 51-60                                  | - Binary  |
| 5.    | Education level 1/7 – illiterate            | - Binary  |
| 6.    | Education level 2/7 – 1-5th Grade Completed | - Binary  |
| 7.    | Education level 3/7 – 6-9th Grade Completed | - Binary  |
| 8.    | Education level 4/7 – 10th Grade Passed     | - Binary  |
| 9.    | Education level 5/7 – 12th Grade Passed     | - Binary  |
| 10.   | Education level 6/7 – Graduate              | - Binary  |
| 11.   | Education level 7/7 – Post graduate         | - Binary  |
| 12.   | Income bracket 1 (e.g., 0-5000)             | - Binary  |
| 13.   | Income bracket 2 (e.g., 5001-10000)         | - Binary  |
| 14.   | Income bracket 3 (e.g., 10001-15000)        | - Binary  |
| 15.   | Income bracket 4 (e.g., 15001-20000)        | - Binary  |
| 16.   | Income bracket 5 (e.g., 20001-25000)        | - Binary  |
| 17.   | Income bracket 6 (e.g., 25001-30000)        | - Binary  |
| 18.   | Income bracket 7 (e.g., 30000-999999)       | - Binary  |

#### Your task:

Write a simple, single-line Python reward function. Exclude the word 'return' and exclude non-standard libraries. Format your code with triple \$ signs: \$\$\$[YOUR FUNCTION]\$\$\$. Note that HIGHER states are always preferred, so ensure reward increases as state increases. Make sure reward is always positive and increasing with state. Example Prompt: Prioritize agents that are older and rich

Example Response:

Python Code: '\$\$\$ state \* (agent\_feats[4] and agent\_feats[18]) \$\$\$'

Come up with a unique new reward for the specified goal: Focus on the young mothers by age.

Goal: Focus on the young mothers by age

#### Output

\$\$\$ state \* agent\_feats[7] \$\$\$

Figure 7: Prompt passed to the LLM to generate reward function based on the problem scenarios in the Real World Domain.

## DLM Choice: Synthetic Domain

### Prompt

My goal was to create a Python reward function for RL in resource allocation, with the objective of: Focus on the agents with low value of feature A. I tried several reward functions for this task.

Below are the reward functions I used and their corresponding reward distributions:

**Function Number 0:** Reward Function:  $2*state + state * (1*agent\_feats[0]+0.5*agent\_feats[1])$

Reflection: '

Category: A Feature A bucket 1: 113.20

Feature A bucket 2: 137.07

Feature A bucket 3: 56.51

Feature A bucket 4: 56.82

Feature A bucket 5: 54.60

Category: B Feature B bucket 1: 82.89

Feature B bucket 2: 65.36

Feature B bucket 3: 60.33

Feature B bucket 4: 50.13

Feature B bucket 5: 46.00

Category: C Feature C Bucket 1: 46.78

Feature C Bucket 2: 49.96

Feature C Bucket 3: 46.64

Feature C Bucket 4: 66.58

Feature C Bucket 5: 62.24

**Function Number 1:** Reward Function:  $state * (agent\_feats[0] \text{ and not } (agent\_feats[1] \text{ or } agent\_feats[2]))$

Reflection: ' Category: A Feature A bucket 1: 177.84

Feature A bucket 2: 54.76

Feature A bucket 3: 55.93

Feature A bucket 4: 57.00

Feature A bucket 5: 55.29

Category: B Feature B bucket 1: 64.47

Feature B bucket 2: 58.76

Feature B bucket 3: 56.96

Feature B bucket 4: 51.11

Feature B bucket 5: 50.13

Category: C Feature C Bucket 1: 49.80

Feature C Bucket 2: 51.69

Feature C Bucket 3: 48.40

Feature C Bucket 4: 65.89

Feature C Bucket 5: 60.00

Based on the above reward distributions and the given goal: Focus on those with high education., please identify the FUNCTION NUMBER of the most effective reward function. Provide your answer EXACTLY IN the following format: 'The best reward function is at number: [FUNCTION NUMBER]'.

### Output:

The best reward function is at number: 1

Figure 8: Prompt passed to the LLM to choose a reward function based on the context of problem scenario in the Synthetic Domain, the generated reward functions and the reward distribution corresponding to every reward function.

## DLM Choice: Real World Domain

### Prompt

My goal was to create a Python reward function for RL in resource allocation, with the objective of: Focus on those with high education. I tried several reward functions for this task. Below, I have the given reward function, and the corresponding distribution of reward achieved across 20 agent features.

Below are the reward functions I used and their corresponding reward distributions:

**Function Number 0:** Reward Function: `-agent_feats[5] -agent_feats[6]-agent_feats[7]-agent_feats[8]-agent_feats[9]-agent_feats[10]-agent_feats[11]`

Reflection: '

Category: Age Ages 10-20: 121.73

Ages 21-30: 421.04

Ages 31-40: 244.49

Ages 41-50: 64.11

Ages 51-60: 10.58

Category: Income Income bracket 1 (e.g., 0-5000): 126.82

Income bracket 2 (e.g., 5001-10000): 373.62

Income bracket 3 (e.g., 10001-15000): 234.87

Income bracket 4 (e.g., 15001-20000): 77.40

Income bracket 5 (e.g., 20001-25000): 35.58

Income bracket 6 (e.g., 25001-30000): 2.58

Income bracket 7 (e.g., 30000-999999): 11.09

Category: Education Illiterate: 39.91

1-5th Grade Completed: 157.84

6-9th Grade Completed: 281.36

10th Grade Passed: 197.64

12th Grade Passed: 103.18

Graduate: 21.13

Post graduate: 60.89

**Function Number 1:** Reward Function: `state * agent_feats[10]`

Reflection: ' Category: Age Ages 10-20: 134.22

Ages 21-30: 469.16

Ages 31-40: 270.44

Ages 41-50: 72.80

Ages 51-60: 11.96

Category: Income Income bracket 1 (e.g., 0-5000): 138.40

Income bracket 2 (e.g., 5001-10000): 414.44

Income bracket 3 (e.g., 10001-15000): 266.44

Income bracket 4 (e.g., 15001-20000): 85.33

Income bracket 5 (e.g., 20001-25000): 40.20

Income bracket 6 (e.g., 25001-30000): 2.80

Income bracket 7 (e.g., 30000-999999): 10.96

Category: Education Illiterate: 45.07

1-5th Grade Completed: 173.82

6-9th Grade Completed: 314.07

10th Grade Passed: 217.31

12th Grade Passed: 113.02

Graduate: 29.36

Post graduate: 65.93

Based on the above reward distributions and the given goal: Focus on those with high education., please identify the FUNCTION NUMBER of the most effective reward function. Provide your answer EXACTLY IN the following format: 'The best reward function is at number: [FUNCTION NUMBER]'.

### Output:

The best reward function is at number: 1

Figure 9: Prompt passed to the LLM to choose a reward function based on the context of problem scenario in Real World Domain, the generated reward functions and the reward distribution corresponding to every reward function.

## DLM Choice with Prompt Engineering (DLM-PromptEngg): Real World Domain

### Prompt

My goal was to create a Python reward function for RL in resource allocation, with the objective of: Focus on those with high education. I tried several reward functions for this task. Below, I have the given reward function, and the corresponding distribution of reward achieved across 20 agent features.

Below are the reward functions I used and their corresponding reward distributions:

**Function Number 0:** Reward Function: `-agent_feats[5] -agent_feats[6]-agent_feats[7]-agent_feats[8]-agent_feats[9]-agent_feats[10]-agent_feats[11]`

Reflection: '

Category: Age Ages 10-20: 121.73

Ages 21-30: 421.04

Ages 31-40: 244.49

Ages 41-50: 64.11

Ages 51-60: 10.58

Category: Income Income bracket 1 (e.g., 0-5000): 126.82

Income bracket 2 (e.g., 5001-10000): 373.62

Income bracket 3 (e.g., 10001-15000): 234.87

Income bracket 4 (e.g., 15001-20000): 77.40

Income bracket 5 (e.g., 20001-25000): 35.58

Income bracket 6 (e.g., 25001-30000): 2.58

Income bracket 7 (e.g., 30000-999999): 11.09

Category: Education Illiterate: 39.91

1-5th Grade Completed: 157.84

6-9th Grade Completed: 281.36

10th Grade Passed: 197.64

12th Grade Passed: 103.18

Graduate: 21.13

Post graduate: 60.89

**Function Number 1:** Reward Function: `state * agent_feats[10]`

Reflection: ' Category: Age Ages 10-20: 134.22

Ages 21-30: 469.16

Ages 31-40: 270.44

Ages 41-50: 72.80

Ages 51-60: 11.96

Category: Income Income bracket 1 (e.g., 0-5000): 138.40

Income bracket 2 (e.g., 5001-10000): 414.44

Income bracket 3 (e.g., 10001-15000): 266.44

Income bracket 4 (e.g., 15001-20000): 85.33

Income bracket 5 (e.g., 20001-25000): 40.20

Income bracket 6 (e.g., 25001-30000): 2.80

Income bracket 7 (e.g., 30000-999999): 10.96

Category: Education Illiterate: 45.07

1-5th Grade Completed: 173.82

6-9th Grade Completed: 314.07

10th Grade Passed: 217.31

12th Grade Passed: 113.02

Graduate: 29.36

Post graduate: 65.93

Figure 10: Enhanced Prompt passed to the LLM to choose a reward function based on the context of the problem scenario in the Real World Domain, the generated reward functions, the reward distributions corresponding to every reward function and additional examples on what to look at when choosing a reward function aligned with the preference.



Based on the above reward distributions and the given goal: Focus on the those with low education., please identify the FUNCTION NUMBER of the most effective reward function. **You can look at the reward distributions for different features and based on them, judge the effectiveness of the correponding reward function. For instance, if the query wants to prioritize low income agents, you should look if the rewards are indeed high for low income features. it is upto you to decide which features describe low income preference.** Provide your answer EXACTLY IN the following format: 'The best reward function is at number: [FUNCTION NUMBER]'..

**Output:**

The best reward function is at number: 1

Figure 11: Enhanced Prompt passed to the LLM to choose a reward function based on the context of the problem scenario in the Real World Domain, the generated reward functions, the reward distributions corresponding to every reward function and additional examples on what to look at when choosing a reward function aligned with the preference.

## DLM Choice with Extended Prompt for Minimizing Utility Shifts (DLM-EP): Real World Domain

### Prompt

My goal was to create a Python reward function for RL in resource allocation, with the objective of: **Focus on the young mothers by age and also focus on those with low education.** I tried several reward functions for this task. Below, I have the given reward function, and the corresponding distribution of reward achieved across 20 agent features.

Below are the reward functions I used and their corresponding reward distributions:

**Function Number 0:** Reward Function: `state * (agent_feats[0] or agent_feats[1]) and (agent_feats[5] or agent_feats[6])`

Reflection: ' "

Category: Age Ages 10-20: 163.24

Ages 21-30: 547.98

Ages 31-40: 269.78

Ages 41-50: 72.11

Ages 51-60: 10.91

Category: Income Income bracket 1 (e.g., 0-5000): 154.40

Income bracket 2 (e.g., 5001-10000): 472.98

Income bracket 3 (e.g., 10001-15000): 293.53

Income bracket 4 (e.g., 15001-20000): 89.82

Income bracket 5 (e.g., 20001-25000): 40.84

Income bracket 6 (e.g., 25001-30000): 2.91

Income bracket 7 (e.g., 30000-99999): 9.53

Category: Education Illiterate: 66.47

1-5th Grade Completed: 257.87

6-9th Grade Completed: 312.69

10th Grade Passed: 224.22

12th Grade Passed: 113.53

Graduate: 23.42

Post graduate: 65.82

**Function Number 1:** Reward Function: `state * (agent_feats[0] or agent_feats[1]) * (agent_feats[5] or agent_feats[6])`

Reflection: ' " Category: Age Ages 10-20: 163.24

Ages 21-30: 547.98

Ages 31-40: 269.78

Ages 41-50: 72.11

Ages 51-60: 10.91

Category: Income Income bracket 1 (e.g., 0-5000): 154.40

Income bracket 2 (e.g., 5001-10000): 472.98

Income bracket 3 (e.g., 10001-15000): 293.53

Income bracket 4 (e.g., 15001-20000): 89.82

Income bracket 5 (e.g., 20001-25000): 40.84

Income bracket 6 (e.g., 25001-30000): 2.91

Income bracket 7 (e.g., 30000-99999): 9.53

Category: Education Illiterate: 66.47

1-5th Grade Completed: 257.87

6-9th Grade Completed: 312.69

10th Grade Passed: 224.22

12th Grade Passed: 113.53

Graduate: 23.42

Post graduate: 65.82

**Additional Information - Rewards from Default reward function** (*Reward distribution from Default reward function. Truncated for brevity.*)

Figure 12: Enhanced Prompt passed to the LLM to choose a reward function based on the context of the problem scenario in the Real World Domain, the generated reward functions, the reward distributions corresponding to every reward function and additional information to minimize the unintended utility shifts in dimensions not specified in the preference.

Based on the above reward distributions and the given goal: **Focus on the young mothers by age and also focus on those with low education**, please identify the FUNCTION NUMBER of the most effective reward function. **Also make sure that that you choose a reward function that does not cause unintended shifts in reward. Unintended shifts in reward here means that the chosen reward function shouldn't drastically change the distribution in reward with respect to features not specified in the prompt For example, if the prompt is to prefer agents with low education, then the chosen reward function shouldn't change the distribution in reward w.r.t the default reward distribution too much in the income feature buckets .** Provide your answer EXACTLY IN the following format: 'The best reward function is at number: [FUNCTION NUMBER]'.

**Output:**

The best reward function is at number: 1

Figure 13: Continued: Enhanced Prompt passed to the LLM to choose a reward function based on the context of the problem scenario in the Real World Domain, the generated reward functions, the reward distributions corresponding to every reward function and additional information to minimize the unintended utility shifts in dimensions not specified in the preference.

## DLM Choice with Extended Prompt for Maximizing Overall Utility (DLM-EP): Real World Domain

### Prompt

My goal was to create a Python reward function for RL in resource allocation, with the objective of: **Focus on the young mothers by age and also focus on those with low education.** I tried several reward functions for this task. Below, I have the given reward function, and the corresponding distribution of reward achieved across 20 agent features.

Below are the reward functions I used and their corresponding reward distributions:

**Function Number 0:** Reward Function: `state * (agent_feats[0] or agent_feats[1]) and (agent_feats[5] or agent_feats[6])`

Reflection: ' "

Category: Age Ages 10-20: 163.24

Ages 21-30: 547.98

Ages 31-40: 269.78

Ages 41-50: 72.11

Ages 51-60: 10.91

Category: Income Income bracket 1 (e.g., 0-5000): 154.40

Income bracket 2 (e.g., 5001-10000): 472.98

Income bracket 3 (e.g., 10001-15000): 293.53

Income bracket 4 (e.g., 15001-20000): 89.82

Income bracket 5 (e.g., 20001-25000): 40.84

Income bracket 6 (e.g., 25001-30000): 2.91

Income bracket 7 (e.g., 30000-99999): 9.53

Category: Education Illiterate: 66.47

1-5th Grade Completed: 257.87

6-9th Grade Completed: 312.69

10th Grade Passed: 224.22

12th Grade Passed: 113.53

Graduate: 23.42

Post graduate: 65.82

**Function Number 1:** Reward Function: `state * (agent_feats[0] or agent_feats[1]) * (agent_feats[5] or agent_feats[6])`

Reflection: ' " Category: Age Ages 10-20: 163.24

Ages 21-30: 547.98

Ages 31-40: 269.78

Ages 41-50: 72.11

Ages 51-60: 10.91

Category: Income Income bracket 1 (e.g., 0-5000): 154.40

Income bracket 2 (e.g., 5001-10000): 472.98

Income bracket 3 (e.g., 10001-15000): 293.53

Income bracket 4 (e.g., 15001-20000): 89.82

Income bracket 5 (e.g., 20001-25000): 40.84

Income bracket 6 (e.g., 25001-30000): 2.91

Income bracket 7 (e.g., 30000-99999): 9.53

Category: Education Illiterate: 66.47

1-5th Grade Completed: 257.87

6-9th Grade Completed: 312.69

10th Grade Passed: 224.22

12th Grade Passed: 113.53

Graduate: 23.42

Post graduate: 65.82

Figure 14: Enhanced Prompt passed to the LLM to choose a reward function based on the context of the problem scenario in the Real World Domain, the generated reward functions, the reward distributions corresponding to every reward function and additional information to maximize the overall utility.

Based on the above reward distributions and the given goal: **Focus on the young mothers by age and also focus on those with low education**, please identify the FUNCTION NUMBER of the most effective reward function. **Also make sure that that you choose a reward function which also maximizes the total reward. You can calculate this by adding up rewards in each feature bucket..** Provide your answer EXACTLY IN the following format: 'The best reward function is at number: [FUNCTION NUMBER]'.

**Output:**

The best reward function is at number: 1

Figure 15: Continued: Enhanced Prompt passed to the LLM to choose a reward function based on the context of the problem scenario in the Real World Domain, the generated reward functions, the reward distributions corresponding to every reward function and additional information to maximize the overall utility.