Optimizing Message-Driven Recruitment on Networks

Anonymous Authors[1]

Abstract

We study adaptive message-driven recruitment on social networks, motivated by digital public-health deployments in which growth relies on non-monetary behavioral nudges rather than paid referral incentives. We introduce a sequential influence model in which a decision maker repeatedly broadcasts one of several message types (or no message) to the currently active population; each broadcast has instantaneous, label-dependent effects that temporarily boost recruitment probabilities along social ties whose relationship contexts match the message. Policies are evaluated using a discounted active-mass objective that rewards both scale and timeliness, capturing the value of early engagement and equity-weighted reach. We show that optimizing such policies is computationally intractable even under strong structural simplifications, and that the problem is hard to approximate along two natural dimensions: maximizing recruitment within a fixed time horizon, and minimizing time to reach a target recruitment level. On the positive side, we analyze a natural adaptive policy that greedily maximizes expected one-step gain, and establish a parameterized approximation guarantee in terms of an instance-dependent amplification parameter that quantifies how strongly recruitment can compound over time. Finally, we illustrate our theoretical findings through experiments on synthetic networks and recruitment scenarios derived from real-world data.

1. Introduction

Our work draws from two complementary lines of public-health research. The first shows that non-monetary digital nudges, such as text messages, reminders, and message framing, can meaningfully affect health-related behaviour (Lester et al., 2010; Milkman et al., 2021; Dai et al., 2021; Chiam et al., 2024). The second shows that public-health interventions can be amplified by leveraging social networks, peer-to-peer influence, and trusted community leaders (Valente, 2012).

Recent algorithmic work in public health has also highlighted that such interventions often give rise to sequential decision-making problems: platforms must decide, over time, how to allocate limited outreach effort as the target population evolves (Verma et al., 2023). Motivated by this perspective, we study message-driven recruitment on social networks. Rather than paying users or directly selecting new participants, a platform repeatedly broadcasts non-monetary nudges to currently active users, with the goal of encouraging themto recruit others through existing social ties.

This setting departs from classical influence maximization (Kempe et al., 2003), where the decision-maker typically selects seed. In our setting, interventions are often global (broadcast to all currently active users), transient (affecting only the current interaction window), and context-dependent (effective only for certain relationships or behavioral contexts). For example, a reminder or notification may successfully induce one user to engage or recruit peers in a family or workplace setting, while having little effect in other contexts. This combination of indirect control, heterogeneous responses, and evolving network effects makes principled intervention design fundamentally challenging.

Case study: HIV care in South Africa. South Africa operates one of the world's largest HIV treatment programs, serving millions of individuals monthly. Despite this scale, prevalence remains critically high, and key populations, such as men and adolescents, may be systematically less likely to engage with traditional clinic-based services. Structural barriers such as limited operating hours, discrimination, and the perception of clinics as non-friendly spaces hinder or delay access to care (UNAIDS, n.d.; Cornell et al., 2011; Colvin, 2019; Sileo et al., 2019; Cornell et al., 2021). To bridge this gap, additional efforts are needed, including the use of digital tools such as chatbots, to serve as a complementary layer of health infrastructure capable of reaching underserved communities (Van Heerden et al., 2023; Chen et al., 2025; Mbewe et al., 2026).

---

[1] Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

Our work is motivated by an ongoing field trial in KwaZulu-Natal (Bill & Melinda Gates Foundation, 2024) to improve delivery of HIV care, with a central objective to extend access to underserved populations. As resources are limited, network growth must be driven by non-monetary behavioral nudges — by optimizing the content, framing, and timing of messages — to encourage existing users to recruit peers within their social graph.

Our approach. We formalize this setting as a sequential decision process on a labeled social network; see Section 3 for details and Fig. 1 for an illustration. At each round, a decision-maker selects a broadcast message that temporarily modifies recruitment probabilities along edges whose relationship context matches the message. Recruitment then proceeds stochastically, expanding the active population over time. The objective is to maximize a discounted notion of active mass, capturing both scale and timeliness of engagement, and reflecting the importance of early and sustained reach.

This formulation departs from classical influence maximization (Kempe et al., 2003). Rather than selecting seed nodes, the decision-maker must adaptively steer an ongoing diffusion process through global interventions. Because each action simultaneously affects many edges and interacts with the evolving active set, the resulting objective is highly non-submodular, and standard approximation guarantees no longer apply.

A central difficulty is the presence of amplification effects: activating one individual may unlock substantial downstream recruitment opportunities. While desirable in practice, this creates strong intertemporal dependencies and makes the optimization problem both computationally intractable and hard to approximate. From a deployment perspective, this raises a key question: when can simple, interpretable policies be trusted, and when do such systems require more complex long-horizon planning?

Our contributions.

- **Hardness:** We show that optimizing message-driven recruitment is NP-hard even under strong structural restrictions, and difficult to approximate under natural objectives such as maximizing recruitment within a fixed horizon or minimizing time to reach a target.
- **Parameterized guarantees:** We analyze a simple adaptive greedy policy and show that its performance is governed by an interpretable amplification parameter. When amplification is limited — a regime consistent with high-friction, real-world settings — the greedy policy achieves a constant-factor approximation.
- **Empirical validation:** Experiments on synthetic and real-world inspired networks demonstrate that amplification governs both problem difficulty and the effectiveness of simple policies.

Collectively, our work reframes influence maximization around the optimization of transient, non-monetary interventions applied to dynamic sets of users. It provides a principled theoretical foundation for understanding when such message-driven recruitment can be efficiently optimized and when worst-case network structures render the problem inherently hard.

2. Related work

Influence maximization and diffusion on networks. A large body of work studies the problem of influence maximization, where the goal is to select an initial set of seed nodes so as to maximize the expected spread of influence under a diffusion model; see the seminal work of (Kempe et al., 2003) and the extensive literature that follow. The two canonical diffusion models in this literature are the *Independent Cascade* (IC) and *Linear Threshold* (LT) models. In the IC model (Goldenberg et al., 2001a;b), influence propagates through independent probabilistic trials along directed edges: when a node u first becomes active, it gets a *single chance* to activate each currently inactive neighbor v with probability $p_{u,v}$. In the LT model (Granovetter, 1978; Schelling, 2006), each node v draws a latent threshold and becomes active once the cumulative influence from its active neighbors exceeds this threshold. A key reason these models became standard is that, under either IC or LT, the expected spread as a function of the seed set is monotone submodular, yielding a $(1 - 1/e)$-approximation via a greedy algorithm (Kempe et al., 2003). This structural property has enabled a rich algorithmic literature, including scalable estimators and near-linear-time algorithms based on sampling and reverse-reachable sets (Borgs et al., 2014; Guo et al., 2022). Subsequent works also studied multi-round variant of IC model (Yan et al., 2011; Yadav et al., 2016), allowing influence attempts to recur over time.

Many variants of influence maximisation have been studied in the literature. The two models most closely related to ours are those of (Yu et al., 2020), where the action space consists of adding edges to the graph, and (Lin et al., 2017), where actions boost the activation probabilities of selected nodes. Both frameworks capture the intuition that the decision-maker cannot directly activate nodes; instead, actions increase the likelihood of activation in a manner contingent on neighbouring nodes also being active—otherwise, the introduction of new edges or boosted probabilities has no effect. The key distinction from our model is that their actions are inherently local, affecting specific nodes or edges, whereas our actions correspond to sending messages that exert a global effect across the entire network.

Much like (Yan et al., 2011; Yadav et al., 2016), our setting considers a multi-round influence dynamics where influence attempts are not exhausted after a single exposure: active users may repeatedly attempt to recruit inactive neighbors across rounds, with success probabilities that depend on the message broadcast in that round. However, our setting is fundamentally different from classic influence maximization: instead of selecting a seed set (or repeatedly selecting additional seeds), our action is a *global broadcast message* that instantaneously modifies recruitment probabilities on many edges simultaneously in a label-dependent manner. This repeated-attempt, message-modulated influence departs from the classical IC/LT framework and destroys the submodularity structure that underpins the classical $(1 - 1/e)$ guarantees.

Adaptive and sequential influence under non-submodular actions. Our problem is inherently sequential: decisions are taken over time, the active set evolves stochastically, and future actions adapt to observed activations. This places our work in the broad space of adaptive influence maximization and sequential decision-making on graphs, including adaptive seeding and multi-round interventions. A unifying framework for many of these problems is *adaptive submodularity* (Golovin & Krause, 2011), which extends submodular guarantees to certain stochastic, adaptive decision problems. When adaptive submodularity holds, greedy policies retain approximation guarantees analogous to the non-adaptive case.

Unfortunately, while our problem is inherently sequential and adaptive, it crucially does *not* satisfy adaptive submodularity: broadcasting a message affects many edges simultaneously through their labels, and the marginal value of an action can increase as the active set grows, due to amplification effects whereby newly activated users unlock additional high-impact recruitment opportunities in future rounds. As a result, the diminishing-returns property fails even in expectation. This rules out standard adaptive greedy analyses and motivates both our hardness results and our alternative, instance-dependent approximation guarantees based on an amplification parameter.

3. Model

Notation. For any integer n, we write $[n]$ to denote the set $\{1, 2, \dots, n\}$. The set of natural numbers is denoted by $\mathbb{N}$, and $\mathbb{N}_{\geq 2} = \{2, 3, \dots\}$. Sets are denoted in bold (e.g., $\mathbf{V}, \mathbf{E}$), and graphs in calligraphic font (e.g., $\mathcal{G}$). Expectations over the full stochastic evolution induced by a policy π are denoted by $\mathbb{E}_\pi[\cdot]$, while $\mathbb{E}_{\mathbf{A}}[\cdot \mid \mathbf{S}, \ell]$ denotes expectation over the one-step activation randomness given state $\mathbf{S}$ and broadcast message ℓ. $\mathbb{1}[\cdot]$ denotes the indicator function.

3.1. The Message-Driven Recruitment (MDR) Model

We are given a directed social network $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ and a finite set of message types $\mathbf{L}$. Each node $v \in \mathbf{V}$ represents an individual and has a nonnegative weight $w_v \geq 0$. For any subset $\mathbf{X} \subseteq \mathbf{V}$, we write $w(\mathbf{X}) = \sum_{v \in \mathbf{X}} w_v$. Node weights allow us to encode heterogeneous value or priority (e.g., emphasizing under-served populations); the special case $w_v \equiv 1$ corresponds to maximizing the expected number of recruited individuals.

Each directed edge $(u, v) \in \mathbf{E}$ represents a potential recruitment relationship and is associated with (i) a baseline recruitment probability $p_{u,v} \in [0, 1]$, and (ii) a (possibly empty) set of labels $\mathbf{L}(u, v) \subseteq \mathbf{L}$. Labels represent relationship contexts under which a broadcast message may be persuasive (e.g., `work`, `school`, `family`, `partner`, `community`, `religion`). Allowing multiple labels per edge accommodates multi-context relationships, e.g., a coworker who is also a friend.

Time is discrete and indexed by rounds $t \in \mathbb{N}_{\geq 1}$. Let $\mathbf{S}^{(t)} \subseteq \mathbf{V}$ denote the set of active (recruited) individuals at the *start* of round t. The initial seed set is $\mathbf{S}^{(1)}$. Once a node becomes active, it remains active permanently.

At each round t, the decision maker broadcasts a *single global message* $\ell^{(t)} \in \mathbf{L} \cup \{\emptyset\}$, where $\emptyset$ denotes sending no message. As illustrated in Fig. 1, a broadcast instantaneously modifies recruitment probabilities only on those *directed edges* whose labels match the broadcasted message; all other edges use their baseline probabilities. Boosts apply only for the current round and do not persist. See Fig. 1 for an illustration.

Instantaneous boosting. At round t, broadcasting a message $\ell \in \mathbf{L} \cup \{\emptyset\}$ induces a round-specific edge probability $p_{u,v}^{(t)} \in [0, 1]$, which depends only on the baseline probability $p_{u,v}$, the edge label set $\mathbf{L}(u, v)$, and the message ℓ. Boosts are instantaneous and do not persist across rounds. Two canonical boosting rules are possible.

Additive boosting. Given boosting parameters $\mathcal{B} = \{\beta_\ell \geq 0 : \ell \in \mathbf{L}\}$, define

$$p_{u,v}^{(t)} = \min\{1, p_{u,v} + \beta_\ell \cdot \mathbb{1}\{\ell \in \mathbf{L}(u, v)\}, p_{u,v}^{\ell}\}$$

Multiplicative boosting. Given boosting parameters $\mathcal{B} = \{\kappa_\ell \geq 1 : \ell \in \mathbf{L}\}$, define

$$p_{u,v}^{(t)} = \min\{1, p_{u,v} \cdot \kappa_\ell^{\mathbb{1}\{\ell \in \mathbf{L}(u,v)\}}\}$$

Recruitment dynamics. Given an active set $\mathbf{S} \subseteq \mathbf{V}$ and a broadcast message ℓ, each inactive node $v \notin \mathbf{S}$ becomes newly active independently with probability $q_v(\mathbf{S}, \ell) \in [0, 1]$. We assume that recruitment attempts are *repeated*

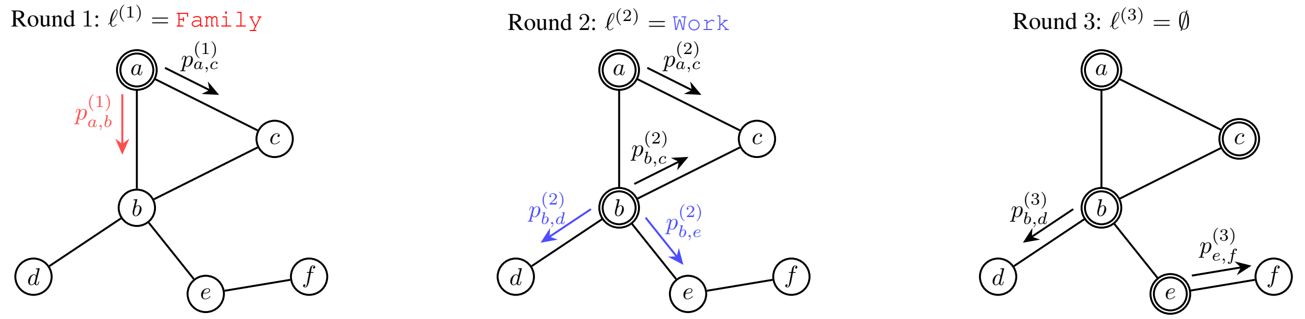

Figure 1. Double circle are active nodes. At each round t, a single global message $\ell^{(t)}$ is broadcast (shown at the top of each figure). Directed edges whose label matches $\ell^{(t)}$ receive a probability boost (shown as colored arrows); all other edges use baseline probabilities. Boosts are instantaneous and do not persist across rounds. Activations are stochastic and permanent.

across rounds: as long as $u \in \mathbf{S}$ and $v \notin \mathbf{S}$, node u may attempt to recruit v in every round, with success probability determined by the message broadcast in that round. A canonical instantiation of this process is a repeated-attempt variant of the Independent Cascade (IC) model. Under this model, $q_v(\mathbf{S}, \ell) = 1 - \prod_{u \in \mathbf{S}}(1 - p_{u,v}(\ell))$, where $p_{u,v}(\ell)$ is the message-dependent recruitment probability on edge (u, v) in the current round. This dynamics reflects message-driven recruitment settings in which referrals are episodic and non-exhaustive, and where exposure in one round does not preclude future influence.

Let $\mathbf{A}^{(t)} \subseteq \mathbf{V}$ denote the random set of nodes newly activated in round t. For a fixed state-message pair $(\mathbf{S}, \ell)$, the expected newly activated weight in one round is

$$\mathop{\mathbb{E}}_{\mathbf{A}}[w(\mathbf{A}) \mid \mathbf{S}, \ell] = \sum_{v \notin \mathbf{S}} q_v(\mathbf{S}, \ell) \cdot w_v \tag{1}$$

Under a policy π which picks ℓ given $\mathbf{S}$, the active set evolves according to $\mathbf{S}^{(t+1)} = \mathbf{S}^{(t)} \cup \mathbf{A}^{(t)}$.

Objective. Fix a horizon $T \in \mathbb{N}_{\geq 2}$ and a discount factor $\gamma \in (0, 1)$. The discounted active-mass objective is

$$R_T(\pi) = \mathop{\mathbb{E}}_{\pi}\left[\sum_{t=1}^{T} \gamma^{t-1} \cdot w(\mathbf{S}^{(t)})\right] \tag{2}$$

Since the contribution of the initial seed set $\mathbf{S}^{(1)}$ is policy-independent, we work with the normalized objective

$$\widetilde{R}_T(\pi) = R_T(\pi) - w(\mathbf{S}^{(1)}) \sum_{t=1}^{T} \gamma^{t-1} \tag{3}$$

For algorithmic analysis, it is convenient to define the value-to-go

$$V_T(\mathbf{S}) = \max_{\pi} \mathop{\mathbb{E}}_{\pi}\left[\sum_{t=1}^{T} \gamma^{t-1} \cdot \mathop{\mathbb{E}}_{\mathbf{A}^{(t)}}\left[w(\mathbf{A}^{(t)}) \mid \mathbf{S}^{(t)}, \ell^{(t)}\right] \middle| \mathbf{S}^{(1)} = \mathbf{S}\right]$$

which captures the maximum expected discounted new activation achievable over T rounds starting from state $\mathbf{S}$.

Definition 3.1 (MDR)**.** Given a labeled directed graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, baseline recruitment probabilities $\{p_{u,v}\}$, boost parameters $\mathcal{B}$, an initial seed set $\mathbf{S}^{(1)}$, horizon T, and discount factor γ, the goal is to compute a policy π that selects a broadcast message $\ell^{(t)} \in \mathbf{L} \cup \{\emptyset\}$ at each round t so as to maximize the expected normalized objective $\widetilde{R}_T(\pi)$.

4. Hardness

We establish computational lower bounds showing that MDR is difficult even under substantial simplifications. In contrast to classical influence maximization under IC/LT, where submodularity yields a $(1 - 1/e)$ greedy guarantee, our broadcast control acts globally and can create amplification effects that destroy (adaptive) submodularity. In particular, we prove NP-hardness under strong structural restrictions and show strong inapproximability results for both (i) minimizing time to reach a target recruitment level and (ii) maximizing recruitment within a fixed horizon.

To do so, we first define the decision version of MDR. Recall from Definition 3.1 that an instance of MDR can be parameterized by $(\mathcal{G}, \mathbf{S}, \mathbf{L}, \mathcal{B}, T, \kappa)$ — a directed graph $\mathcal{G}$, a set of initial seed nodes $\mathbf{S}$, a set of message labels $\mathbf{L}$, a set of boosting parameters $\mathcal{B}$, and number of rounds T. Throughout this section, we focus on the unit-weight case $w_v \equiv 1$ and measure performance by the expected number of *new* activations over a horizon. Given an initial active set $\mathbf{S} \subseteq \mathbf{V}$ and a policy π, define

$O_T(\mathbf{S}, \pi) = \mathbb{E}_\pi\left[\sum_{t=1}^{T} |\mathbf{A}^{(t)}| \mid \mathbf{S}^{(1)} = \mathbf{S}\right]$, where $\mathbf{A}^{(t)}$ is the set of nodes newly activated in round t.

Definition 4.1 (MDR decision version)**.** Given an instance $(\mathcal{G}, \mathbf{S}, \mathbf{L}, \mathcal{B}, T, \kappa)$, decide whether there exists a policy π such that $O_T(\mathbf{S}, \pi) \geq \kappa$.

Our NP-hardness result follows from reducing two known NP-hard problems (shortest common supersequence and minimum set cover) to Definition 4.1. Both reductions hold even with unit node weights. The former reduction holds even when the input graph is a tree with message set size $|\mathbf{L}| = 2$ while the latter reduction holds even on a star graph.

Theorem 4.2 (NP-hardness on trees with two messages)**.** *Suppose all nodes are unit weight, i.e.,* $w_v = 1$ *for all* $v \in \mathbf{V}$*. The decision problem in Definition 4.1 is NP-hard even under either structural restriction: (i)* $\mathcal{G}$ *is a tree and the message set has size two (*$|\mathbf{L}| = 2$*), (ii)* $\mathcal{G}$ *is a star graph.*

Theorem 4.2 shows that hardness is not driven by dense graphs, large action spaces, or heterogeneous weights. Even when the diffusion is restricted to a tree and the decision maker chooses between two broadcast types, deciding whether one can recruit a target number of individuals in expectation is NP-hard.

Given the NP-hardness of MDR, a natural question is whether approximation algorithms exist. The problem is inherently bi-criteria: one may either (A) fix a desired recruitment level and seek to minimize the number of rounds required, or (B) fix a time horizon and seek to maximize the expected number of recruited individuals. We show that both formulations are hard to approximate in a strong sense by giving approximation-preserving reductions from *minimum set cover* and *k-densest subgraph*, thereby inheriting their known inapproximability bounds in each of the bi-criteria dimensions.

(A) Fixed target, minimize time. For a target $\kappa \in \mathbb{N}$, define the optimal number of rounds as $T^\star(\kappa) = \min\{T \in \mathbb{N}_{\geq 1} : \exists \pi \text{ with } O_T(\mathbf{S}, \pi) \geq \kappa\}$.

Theorem 4.3 ($(1-\varepsilon)\ln n$-inapproximability for time-to–target)**.** *For every constant* $\varepsilon > 0$*, unless* P $=$ NP*, no polynomial-time algorithm can approximate* $T^\star(\kappa)$ *within a factor better than* $(1-\varepsilon)\ln n$*. This hardness holds even with unit node weights.*

(B) Fixed horizon, maximize recruitment. For a fixed horizon $T \in \mathbb{N}$, define the optimal expected number of new activations as $\kappa^\star(T) = \max_\pi O_T(\mathbf{S}, \pi)$.

Theorem 4.4 (Strong inapproximability for recruitment-at-horizon)**.** *Assuming the Exponential Time Hypothesis (ETH), there exists a universal constant* $c > 0$ *such that no polynomial-time algorithm can approximate* $\kappa^\star(T)$ *within a factor better than* $n^{-1/(\log\log n)^c}$*. Equivalently, under ETH, there is no polynomial-time algorithm that always returns a policy achieving expected recruitment at least* $n^{-1/(\log\log n)^c} \cdot \kappa^\star(T)$*.*

Theorem 4.3 and Theorem 4.4 rule out efficient worst-case guarantees that are independent of instance structure: even if one is willing to approximate, optimizing recruitment over time is hard both when (i) the objective is to minimize the number of rounds to reach a target and when (ii) the objective is to maximize recruitment within a fixed horizon. These results motivate our parameterized analysis in Section 5, which identifies regimes where simple adaptive policies admit provable performance bounds.

5. Parameterized approximation for a natural greedy algorithm

Our hardness results in Section 4 motivate approximation guarantees under additional structural control. Here, we analyze a natural myopic policy and show that its performance can be bounded in terms of a single *amplification parameter* that quantifies how much activating one individual can increase future recruitment opportunities.

Greedy policy. Fix tie-breaking (e.g., lexicographic) so the policy is deterministic. At each round $t \geq 1$, given the current active set $\mathbf{S}^{(t)}$, the greedy policy broadcasts $\ell^{(t)}_{\text{greedy}} \in \arg\max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \mathbb{E}_{\mathbf{A}^{(t)}}\left[w(\mathbf{A}^{(t)}) \mid \mathbf{S}^{(t)}, \ell\right]$. Equivalently, at each step, greedy always chooses an action attaining the one-step expected gain $g(\mathbf{S})$ defined as:

$$g(\mathbf{S}) = \max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \mathop{\mathbb{E}}_{\mathbf{A}}\left[w(\mathbf{A}) \mid \mathbf{S}, \ell\right] \tag{4}$$

Amplification parameter. We quantify how sensitive one-step recruitment is to enlarging the active set. Define the global problem amplification parameter $\rho \in [0, \infty)$ as:

$$\rho = \max_{\substack{\mathbf{S} \subseteq \mathbf{V}, v \notin \mathbf{S} \\ \ell \in \mathbf{L} \cup \{\emptyset\}, w_v > 0}} \frac{1}{w_v} \cdot \left(\mathop{\mathbb{E}}_{\mathbf{A}}[w(\mathbf{A}) \mid \mathbf{S} \cup \{v\}, \ell] - \mathop{\mathbb{E}}_{\mathbf{A}}[w(\mathbf{A}) \mid \mathbf{S}, \ell]\right) \tag{5}$$

Intuitively, the parameter ρ measures the maximum additional *immediate* recruitment potential unlocked by activating one unit of weight, and serves as a global Lipschitz constant controlling how sensitive next-round recruitment is to changes in the active set. Small ρ corresponds to diminishing or limited amplification, whereas large ρ captures worst-case instances, such as those in our hardness constructions (Section 4), where activating a single node can unlock a large number of otherwise unreachable recruits in the next round.

Given the definitions above, we can show the following parameterized approximation guarantee for π_{greedy}. Note that no assumption is made on the value of the amplification parameter ρ; if ρ is large, the approximation ratio will degrade accordingly.

Theorem 5.1 (Parameterized approximation for π_{greedy}). *Fix an* MDR *instance with horizon $T \in \mathbb{N}_{\geq 2}$ and discount factor $\gamma \in (0,1)$. Then, the greedy policy π_{greedy} achieves*

$$\widetilde{R}_T(\pi_{\text{greedy}}) \geq \alpha(T,\gamma,\rho) \cdot \max_{\pi} \widetilde{R}_T(\pi)$$

where the approximation factor is

$$\alpha(T,\gamma,\rho) = \frac{1}{\sum_{k=0}^{T-1} \left(\gamma \cdot (1+\rho)\right)^k} \cdot \frac{1-\gamma}{1-\gamma^{T-1}}$$

Interpreting $\alpha(T,\gamma,\rho)$. The factor $\sum_{k=0}^{T-1}(\gamma(1+\rho))^k$ arises from how a one-step gain at time t can be amplified (by at most a factor $1+\rho$ per additional step) and discounted (by γ per step). When $\gamma \cdot (1+\rho) < 1$, amplification decays faster than discounting, yielding a constant-factor guarantee. As $\gamma \cdot (1+\rho)$ approaches or exceeds 1, early activations can unlock increasingly large future gains, and the guarantee degrades with horizon. Meanwhile, the ratio $\frac{1-\gamma}{1-\gamma^{T-1}} = \frac{1}{\sum_{t=1}^{T-1} \gamma^{t-1}}$ can be treated as a finite-horizon normalization or horizon-truncation penalty; this ratio approaches $1-\gamma$ as the horizon T increases.

Proof sketch (see Section B.1). The proof has two steps. First, we introduce an auxiliary discounted new-activation objective $H_T(\pi) := \mathbb{E}_{\pi}[\sum_{t=1}^{T-1} \gamma^{t-1} w(\mathbf{A}^{(t)})]$ and show that greedy attains a $1/\sum_{k=0}^{T-1}(\gamma(1+\rho))^k$ fraction of the optimal H_T. This uses a Bellman decomposition plus a Lipschitz bound implied by ρ. Second, we relate $\widetilde{R}_T(\pi)$ to $H_T(\pi)$ via simple geometric bounds on the coefficient $\Gamma_t = \sum_{i=t+1}^{T} \gamma^{i-1}$ multiplying $w(\mathbf{A}^{(t)})$ in $\widetilde{R}_T$. Combining the two steps yields the stated approximation factor.

We complement our parameterized upper bound with a hardness guarantee for the greedy policy π_{greedy}.

Proposition 5.2 (Lower bound for π_{greedy}). *Fix horizon $T \in \mathbb{N}_{\geq 2}$ and discount factor $\gamma \in (0,1)$. Then, for every $\epsilon > 0$ and $\rho > 1+\epsilon$, there exist an instance of* MDR *where greedy policy π_{greedy} achieves*

$$\widetilde{R}_T(\pi_{\text{greedy}}) \leq \alpha'(T,\gamma,\rho) \cdot \max_{\pi} \widetilde{R}_T(\pi)$$

where the approximation factor is

$$\begin{aligned} &\alpha'(T,\gamma,\rho) \\ &= \alpha(T,\gamma,\rho-\varepsilon) \cdot \frac{(1-\gamma^{T-1}) \cdot (1+2\varepsilon) \cdot (2-\gamma-\gamma^T)}{(1-\gamma)^3} \end{aligned}$$

In other words, Proposition 5.2 tells us that the parameterized bound of Theorem 5.1 is optimal up to constant factors for any fixed discount factor γ.

Proof sketch (see Section B.2). We will construct an MDR instance defined over a graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ with $n = T + 2^T$ nodes and $T+1$ labels such that the greedy policy π_{greedy} achieves

$$\widetilde{R}_T(\pi_{\text{greedy}}) \leq \frac{(1+2\varepsilon) \cdot (2-\gamma-\gamma^T)}{(1-\gamma)^2}$$

while an optimal policy π_{OPT} achieves

$$\widetilde{R}_T(\pi_{\text{OPT}}) \geq \frac{1}{\sum_{k=0}^{T-1} \left(\gamma \cdot (1+\rho-\varepsilon)\right)^k}$$

Thus, establishing our result.

6. Experiments

We benchmark our proposed GREEDY policy against several natural baselines—EMPTY, RANDOM, and GREEDY-OVER-RANDOM—on both synthetic and real-world graphs for MDR under discounted rewards. The policies are defined as follows.

- GREEDY: The greedy policy defined in Section 5, which at each round selects the action that maximises the one-step expected gain.
- EMPTY: Selects no labels, (i.e. $\emptyset$), at every round.
- RANDOM: Selects a label uniformly at random at each round.
- GREEDY-OVER-RANDOM: For each candidate action, estimates the expected discounted reward by fixing that label in the current round and simulating RANDOM for all remaining rounds; the action with the highest estimated reward is selected.

Average Amplification Parameter. Motivated by the amplification parameter ρ in Section 5, we consider the notion of average unboosted amplification factor. For $u \in \mathbf{V}$, let $N(u) = \{v \in \mathbf{V} \mid (u,v) \in \mathbf{E}\}$ be the out neighbours of u. Then, define the unboosted amplification parameter of u $\rho'(u) \in [0,\infty)$ as:

$$\rho'(u) = \sum_{v \in N(u)} \frac{p_{u,v} \cdot w(v)}{w(u)}$$

Then, the average unboosted amplification parameter of an instance $\mathbf{G}$ is:

$$\rho' = \frac{1}{|\mathbf{V}|} \sum_{u \in \mathbf{V}} \rho'(u) = \frac{1}{|\mathbf{V}|} \sum_{u \in \mathbf{V}} \sum_{v \in N(u)} \frac{p_{u,v} \cdot w(v)}{w(u)}$$

Synthetic Dataset We extend the Stochastic Block Model (Holland et al., 1983) and its weighted variant (Ng & Murphy, 2021) to allow edges to carry labels. We assume that each node is assigned independently to one of Q latent clusters. Conditioned on these cluster assignments, all remaining aspects of the network are generated independently across ordered node pairs: the presence of a directed edge, its associated baseline probability, and its set of labels are each drawn independently from distributions that depend only on the clusters of the two endpoints. Further details are provided in Section C.1.

Email-EU Dataset We use the Email-EU dataset (Lin et al., 2017; Leskovec et al., 2007), which consists of approximately 1,000 nodes and 25,600 directed edges, where an edge represents an email exchange between members of a large European research institution. Each node belongs to exactly one of 42 departments within the institution. For each edge, we draw the baseline recruitment probability from a uniform distribution, and assign the edge label as the target node's department index modulo 10; this aggregation ensures that each action affects a sufficiently large fraction of nodes.

In all experiments, we fix the number of labels to $|\mathbf{L}| = 10$ and use multiplicative boosting with $\kappa_\ell = 2$ for all $\ell \in \mathbf{L}$. The initial seed set consists of 0.5% of the nodes, sampled uniformly at random.

6.1. Results

Effect of the Average Unboosted Amplification ρ'. In Fig. 2, we can see that for both the synthetic and Email-EU dataset, as ρ' increase, the performance gains of GREEDY over RANDOM decreases. This can be explain by the fact that at high ρ', even without sending any message in the EMPTY policy, there is a good chance that an unactivated node will become activated, hence there is less need for sending effective messages.

We also observe that under the EMPTY policy, for the same value of ρ', node activations occur at a much faster rate in the synthetic experiments than in the Email-EU experiments. We attribute this difference to structural properties of the synthetic graphs: nodes are more easily reachable and there are effectively no hard-to-reach nodes.

Specifically, across 5,000 synthetically instances from the stochastic block model, all nodes are always reachable from the initial seeds, even when the average degree is set to 20 (which is lower than the average degree of approximately 25 in the Email-EU dataset). Moreover, the maximum shortest-path distance from the initial seed nodes is at most 3. In contrast, in instances generated from the Email-EU network, the expected number of unreachable nodes is 39.8, and the maximum distance from the initial seed nodes to a reachable node is 5.

Performance of GREEDY-OVER-RANDOM We run GREEDY-OVER-RANDOM with the number of simulated trajectories set to $s \in \{5, 10, 15, 20\}$. We evaluate the relative improvement of both GREEDY and GREEDY-OVER-RANDOM against RANDOM in terms of the discounted reward R_T, using a discount factor $\gamma = 0.95$.

As shown in Fig. 3, while GREEDY-OVER-RANDOM consistently outperforms RANDOM, it also consistently underperforms GREEDY. Moreover, increasing the number of simulated trajectories does not appear to yield systematic performance improvements. A plausible explanation is that RANDOM is highly suboptimal and therefore provides a poor baseline for guiding greedy action selection.

Finally, we note that GREEDY-OVER-RANDOM incurs a substantially higher computational cost than GREEDY. The per-round runtime of GREEDY-OVER-RANDOM is $O(\ell \cdot s \cdot T \cdot (|\mathbf{V}| + |\mathbf{E}|))$, where T denotes the number of remaining timesteps, compared to $O(\ell(|\mathbf{V}| + |\mathbf{E}|))$ for GREEDY.

7. Conclusion and Discussions

We introduced a model for adaptive message-driven recruitment on labeled social networks with instantaneous broadcast effects, and studied its algorithmic foundations. Our hardness results show that even when the underlying network structure and recruitment parameters are fully known, optimizing adaptive broadcast policies is computationally intractable in general. At the same time, our parameterized analysis identifies a regime in which a simple adaptive policy admits provable guarantees: when the amplification parameter ρ is small, myopic maximization of one-step expected gain achieves a constant-factor approximation over long horizons, whereas large ρ captures instances where long-horizon planning can be dramatically more valuable.

An important limitation of our formulation is that it assumes full knowledge of the social network and recruitment probabilities. In practice, deployments rarely know $\mathcal{G}$ upfront, and the most relevant edges — who talks to whom, in which contexts, and with what responsiveness to which messages — are only partially observed and may shift over time. Our results suggest that this is not merely a technical nuisance: the problem is already hard even with complete information, so partial-information settings will require new algorithmic ideas beyond exact optimization.

Several directions appear particularly promising. First, one can study learning-augmented and online variants in which the decision maker must simultaneously learn influence parameters (or even the network) while optimizing broadcasts, potentially under bandit feedback. Second, realistic deployments may impose additional constraints (fairness or equity

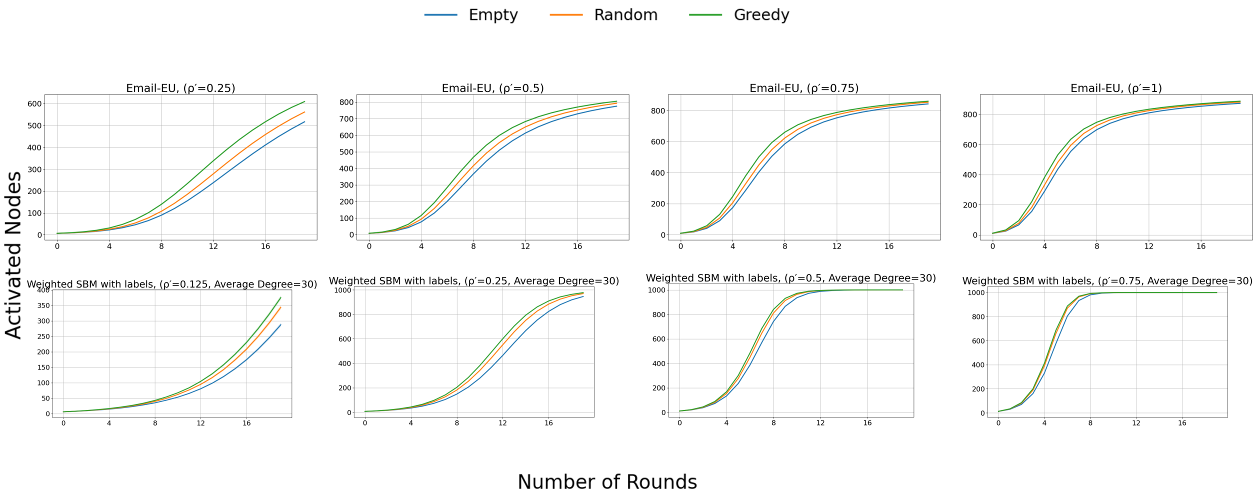

Figure 2. **Synthetic dataset, for $\rho' \in \{0.125, 0.25, 0.5, 0.75\}$ (top) and Email-EU dataset, for $\rho' \in \{0.25, 0.5, 0.75, 1\}$ (bottom).** Expected number of activated nodes over time under three policies: EMPTY, RANDOM, and GREEDY.

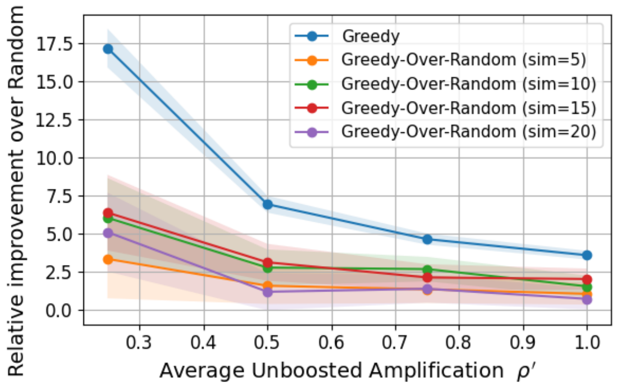

Figure 3. Relative improvement over RANDOM as a function of the target unboosted amplification ρ'.

constraints across subpopulations, limits on message fatigue, or budgets on the number of broadcasts), which interact non-trivially with amplification. Finally, the label structure itself could be learned or refined from data, connecting message design and causal modeling to sequential control. We view the present work as a first step toward a principled theory of adaptive broadcast optimization for real-world recruitment systems.

References

Bill & Melinda Gates Foundation. Hiv lift: Leveraging innovation and friendships to improve hiv testing, 2024. IATI Identifier: DAC-1601-INV-070192. Available at https://datastore.iatistandard.org/activity/DAC-1601-INV-070192.

Borgs, C., Brautbar, M., Chayes, J., and Lucier, B. Maximizing social influence in nearly optimal time. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pp. 946–957. SIAM, 2014.

Chen, S., Yu, F.-y., Fang, Y., Yau, J., Ng, N., Zhang, Q., Ni, Z., Duong, M. C., Sun, F., Mo, P. K., et al. Chatbot-delivered real-time support to improve hiv self-testing rates: A randomized clinical trial. *JAMA Network Open*, 8(11):e2544821–e2544821, 2025.

Chiam, J., Lim, A., and Teredesai, A. Nudgerank: Digital algorithmic nudging for personalized health. In *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining*, pp. 4873–4884, 2024.

Colvin, C. J. Strategies for engaging men in hiv services. *The Lancet HIV*, 6(3):e191–e200, 2019.

Cornell, M., McIntyre, J., and Myer, L. Men and antiretroviral therapy in africa: our blind spot. *Tropical Medicine & International Health*, 16(7):828–829, 2011.

Cornell, M., Majola, M., Johnson, L. F., and Dubula-Majola, V. Hiv services in sub-saharan africa: the greatest gap is men. *The Lancet*, 397(10290):2130–2132, 2021.

Dai, H., Saccardo, S., Han, M. A., Roh, L., Raja, N., Vangala, S., Modi, H., Pandya, S., Sloyan, M., and Croymans, D. M. Behavioural nudges increase COVID-19 vaccinations. *Nature*, 597(7876):404–409, 2021.

Dinur, I. and Steurer, D. Analytical approach to parallel repetition. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pp. 624–633, 2014.

Goldenberg, J., Libai, B., and Muller, E. Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing letters*, 12(3):211–223, 2001a.

Goldenberg, J., Libai, B., and Muller, E. Using complex systems analysis to advance marketing theory development: Modeling heterogeneity effects on new product growth through stochastic cellular automata. *Academy of Marketing Science Review*, 9(3):1–18, 2001b.

Golovin, D. and Krause, A. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42:427–486, 2011.

Granovetter, M. Threshold models of collective behavior. *American journal of sociology*, 83(6):1420–1443, 1978.

Guo, Q., Wang, S., Wei, Z., Lin, W., and Tang, J. Influence maximization revisited: efficient sampling with bound tightened. *ACM Transactions on Database Systems (TODS)*, 47(3):1–45, 2022.

Holland, P. W., Laskey, K. B., and Leinhardt, S. Stochastic blockmodels: First steps. *Social networks*, 5(2):109–137, 1983.

Kempe, D., Kleinberg, J., and Tardos, É. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 137–146, 2003.

Leskovec, J., Kleinberg, J., and Faloutsos, C. Graph evolution: Densification and shrinking diameters. *ACM transactions on Knowledge Discovery from Data (TKDD)*, 1 (1):2–es, 2007.

Lester, R. T., Ritvo, P., Mills, E. J., Kariri, A., Karanja, S., Chung, M. H., Jack, W., Habyarimana, J., Sadatsafavi, M., Najafzadeh, M., et al. Effects of a mobile phone short message service on antiretroviral treatment adherence in kenya (weltel kenya1): a randomised trial. *The Lancet*, 376(9755):1838–1845, 2010.

Lin, Y., Chen, W., and Lui, J. C. Boosting information spread: An algorithmic approach. In *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*, pp. 883–894. IEEE, 2017.

Manurangsi, P. Almost-polynomial ratio eth-hardness of approximating densest k-subgraph. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 954–961, 2017.

Mbewe, A., Sarkwah, H. N., Ijaodola, O., Ratevosian, J., Simba, S., Smith, T. R., Atiwoto, W. K., Cain, K., Konopka, S. N., Ujam, C., et al. The role of digital health and artificial intelligence in improving the reach and effectiveness of hiv prevention in africa. *The Lancet Global Health*, 14(1):e131–e142, 2026.

Milkman, K. L., Patel, M. S., Gandhi, L., Graci, H. N., Gromet, D. M., Ho, H., Kay, J. S., Lee, T. W., Akinola, M., Beshears, J., et al. A megastudy of text-based nudges encouraging patients to get vaccinated at an upcoming doctor's appointment. *Proceedings of the National Academy of Sciences*, 118(20):e2101165118, 2021.

Ng, T. L. J. and Murphy, T. B. Weighted stochastic block model. *Statistical Methods & Applications*, 30(5):1365–1398, 2021.

Räihä, K.-J. and Ukkonen, E. The shortest common supersequence problem over binary alphabet is np-complete. *Theoretical Computer Science*, 16(2):187–198, 1981. ISSN 0304-3975. doi: https://doi.org/10.1016/0304-3975(81)90075-X. URL `https://www.sciencedirect.com/science/article/pii/030439758190075X`.

Schelling, T. C. *Micromotives and macrobehavior*. WW Norton & Company, 2006.

Sileo, K. M., Fielding-Miller, R., Dworkin, S. L., and Fleming, P. J. A scoping review on the role of masculine norms in men's engagement in the hiv care continuum in sub-saharan africa. *AIDS care*, 2019.

UNAIDS. Global HIV & AIDS statistics - Fact sheet, n.d. Available at `https://www.unaids.org/en/resources/fact-sheet`.

Valente, T. W. Network interventions. *science*, 337(6090): 49–53, 2012.

Van Heerden, A., Bosman, S., Swendeman, D., and Comulada, W. S. Chatbots for hiv prevention and care: a narrative review. *Current HIV/AIDS Reports*, 20(6):481–486, 2023.

Verma, S., Singh, G., Mate, A., Verma, P., Gorantla, S., Madhiwalla, N., Hegde, A., Thakkar, D., Jain, M., Tambe, M., et al. Increasing Impact of Mobile Health Programs: SAHELI for Maternal and Child Care. In *Proceedings of the aaai conference on artificial intelligence*, volume 37, pp. 15594–15602, 2023.

Yadav, A., Chan, H., Jiang, A. X., Xu, H., Rice, E., and Tambe, M. Using social networks to aid homeless shelters: Dynamic influence maximization under uncertainty. In *AAMAS*, volume 16, pp. 740–748, 2016.

Yan, Q., Guo, S., and Yang, D. Influence maximizing and local influenced community detection based on multiple spread model. In *International Conference on Advanced Data Mining and Applications*, pp. 82–95. Springer, 2011.

Yu, L., Li, G., and Yuan, L. Maximizing boosted influence spread with edge addition in online social networks. *ACM Transactions on Data Science*, 1(2):1–21, 2020.

A. Deferred proofs

We complement our algorithmic guarantees with hardness results showing that MDR is computationally difficult even under substantial simplifications.

We first state the decision variant of MDR.

MDR

Input: An instance $\mathcal{I} = (\mathcal{G}, \mathbf{S}, \mathbf{L}, \mathcal{B}, T, \kappa)$, consisting of a graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, an initial seed set $\mathbf{S} \subseteq \mathbf{V}$, a message set $\mathbf{L}$ with boosting parameters $\mathcal{B}$, number of rounds T and a target number of newly activated nodes $\kappa \in \mathbb{N}$.

Question: Does there exist a policy π such that the expected number of newly activated nodes after T rounds is at least κ? In other words, does there exists π such that

$$O_T(\mathbf{S}, \pi) \geq \kappa?$$

A.1. NP-hardness under structural restrictions

We first define the notion of an *boosted and unboosted activation*

Definition A.1 (Boosted and Unboosted Activation)**.** Consider an activation event in which node u activates node v at round r. We call this event a *boosted activation* if the message ℓ_r broadcast in round r satisfies $\ell_r \in \mathbf{L}(u, v)$, and an *unboosted activation* otherwise, i.e.,

$$\ell_r \notin \mathbf{L}(u, v).$$

We note by the linearity of expectation, the expected number of new nodes activated is the sum of the expected number of boosted and unboosted activation. Now, we bound the expected number of such activations over T rounds for any policy, as a function of the maximum unboosted propagation probability.

Lemma A.2. *Suppose that for every ordered pair (u, v), the unboosted propagation probability satisfies*

$$p_{(u,v)} \leq \frac{\varepsilon}{n(n-1)T}.$$

Then, for any policy π, the expected total number of unboosted activations over T rounds is at most ε.

Proof. Fix an arbitrary policy π. Consider a single round r. For any ordered pair (u, v), the probability that (u, v) forms an unboosted activation in round r is at most $p_{(u,v)} \leq \frac{\varepsilon}{n(n-1)T}$. Since there are at most $n(n-1)$ ordered pairs of nodes, the expected number of unboosted activations in round r is at most

$$n(n-1) \cdot \frac{\varepsilon}{n(n-1)T} = \frac{\varepsilon}{T}.$$

By linearity of expectation, the expected total number of unboosted activations over T rounds is at most

$$T \cdot \frac{\varepsilon}{T} = \varepsilon.$$

This bound holds for any policy π, completing the proof. □

We note that by Markov's inequality, the probability that no unboosted activations occur is at least $1 - \varepsilon$.

Next, we show that MDR remains NP-hard even in a highly restricted setting: the underlying graph $\mathcal{G}$ is a tree, all node weights are unit ($w_v \equiv 1$), and the message set has size two ($|\mathbf{L}| = 2$).

Theorem A.3. MDR *is NP-hard even when the underlying graph $\mathcal{G}$ is a tree, all node weights are unit ($w_v \equiv 1$), and the message set has size two ($|\mathbf{L}| = 2$)*

Proof. For clarity, in our following reduction, the graph $\mathcal{G}$ is a forest. This can be easily extended to a tree by adding a new root node and connecting it to the root of each subtree in $\mathcal{G}$ via edges with propagation probability 1, and allowing one additional round.

We reduce from the shortest common supersequence problem over binary alphabet (SCSB). The shortest common supersequence is known to be NP-hard even over a binary alphabet.(Räihä & Ukkonen, 1981).

SHORTEST COMMON SUPERSEQUENCE OVER BINARY ALPHABET (SCSB)

Input: An instance $\mathcal{I} = (\mathcal{S}, \kappa)$ consisting of a binary alphabet $\Sigma = \{0, 1\}$, a collection of strings $\mathcal{S} = \{s_1, s_2, \ldots, s_m\}$ with $s_i \in \Sigma^*$, and a target length $\kappa \in \mathbb{N}$.

Question: Does there exist a string $w \in \Sigma^*$ of length κ such that every string $s_i \in \mathcal{S}$ is a (not necessarily contiguous) subsequence of w?

Let I_1 be an instance of $SCSB$. Then, let $n = \sum_{i \in [m]} 1 + s_i$.

We construct instance I_2 of MDR with a graph of n nodes and $|\mathcal{S}|$ starting node. Let $p_e = \frac{1}{2n(n-1)T}$ for all $e \in E$, and let $\beta_\ell = 1 - \frac{1}{2n(n-1)T}$ under additive boosting and $\beta_\ell = 2n(n-1)T$ under multiplicative boosting for all messages $\ell \in \mathbf{L}$. In other words, each edge has a propogation probability of $\frac{1}{2n(n-1)T}$ if unboosted and a propogation probability of 1 if boosted.

For each string $s \in \mathcal{S}$, write

$$s = p_s(1) \circ p_s(2) \circ \cdots \circ p_s(|s|),$$

where $p_s(i)$ denotes the ith (from the left) character of s. Next, for each string $s \in \mathcal{S}$, we construct a directed path subgraph $\mathcal{G}_s$ consisting of $|s| + 1$ nodes

$$n_{(s,0)}, n_{(s,1)}, \ldots, n_{(s,|s|)}.$$

For each $i \in \{1, \ldots, |s|\}$, we add a directed edge $(n_{(s,i-1)}, n_{(s,i)})$ and assign it the label

$$\mathbf{L}(n_{(s,i-1)}, n_{(s,i)}) = \{p_s(i)\}.$$

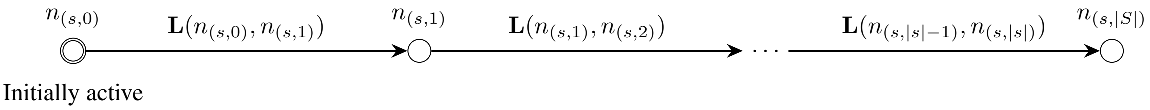

Figure 4. SCSB chain graph for $\mathcal{G}_s$, where $\mathbf{L}(n_{(s,i-1)}, n_{(s,i)}) = \{p_s(i)\}$ for $i \in \{1, \ldots, |s|\}$.

Let $\mathcal{G} = \biguplus_{s \in \mathcal{S}} \mathcal{G}_s$ be a collection of all subgraphs $\mathcal{G}_s$. We define the set of seed nodes to be the first node of the path graph corresponding to each string $s \in \mathcal{S}$. That is,

$$\mathbf{S} = \{n_{(s,0)} \mid s \in \mathcal{S}\}.$$

Note that the constructed graph $\mathcal{G}$ is a forest, since it is formed by a collection of path subgraphs. In addition, all node weights are unit, and the message set $L = \{0, 1\}$ has size two.

We now show that $I_1 = (\mathcal{S}, \kappa)$ is a `YES` instance of SCSB if and only if $I_2 = (\mathcal{G}, \mathbf{S}, \mathbf{L}, \mathcal{B}, \kappa, n - |\mathcal{S}|)$ is a `YES` instance of MDR.

($\Rightarrow$) If I_1 is a `YES` instance of SCSB, then there exists a string w of length of length κ such that every string $s_i \in \mathcal{S}$ is a subsequence of w. Let π_w be the policy that relay messages in the order of w. We now show that under this policy, all nodes in $\mathcal{G}$ are activated and hence $O_T(\mathbf{S}, \pi_w) = n - |S|$.

Fix any $s \in \mathcal{S}$. As s is a subsequence of w, there exists a strictly increasing sequence of indices

$$(r_1, r_2, \ldots, r_{|s|}) \subseteq [|w|]$$

such that $p_s(i) = w_{r_i}$ for every $i \in [|s|]$. We now show by induction that, , the node $n_{s,i}$ in the subgraph $\mathcal{G}_s$ is active after round r_i.

Base case. At round r_1, the node $n_{(s,0)}$ is active since $n_{(s,0)} \in \mathbf{S}$ is a seed node. Moreover, because

$$w_{r_1} = p_s(1) \in \mathbf{L}(n_{(s,0)}, n_{(s,1)}),$$

the node $n_{(s,1)}$ becomes active after round r_1.

Inductive step. Assume that node $n_{(s),i)}$ is active after round r_i for some $i \in \{1, \ldots, |s| - 1\}$. Since $r_i < r_{i+1}$, the node $n_{(s,i)}$ remains active during round $r_{(i+1)}$. Furthermore, because

$$w_{r_{i+1}} = p_s(i+1) \in \mathbf{L}(n_{(s,i)}, n_{(s,i+1)}),$$

the node $n_{(s,i+1)}$ is activated after round r_{i+1}.

By induction, all nodes $n_{s,i}$ for $i \in \{0, \ldots, |s|\}$ are active after T rounds. Therefore, the constructed instance I_2 is a YES-instance.

(⇐). Assume I_2 is a YES-instance of MDR. Then there exists a policy π such that

$$O_T(\mathbf{S}, \pi) \geq n - |\mathcal{S}|.$$

Since $p_e = \frac{1}{2n(n-1)T}$ for all $e \in E$, Lemma A.7 implies that for any policy (and in particular for π), the expected number of *unboosted* activations over T rounds is at most 0.5. Since, $O_T(\mathbf{S}, \pi) \geq n - |\mathcal{S}|$, the expected number of *boosted* activations is at least $n - |S| - 0.5$. For each trajectory, as the number of *boosted* activations is integral for each realisation, there must be a trajectory with $n - |S|$ boosted activations.

Let $w = w_1 w_2 \cdots w_\kappa$ be the sequence of messages broadcast by π in one such trajectory with $n - |S|$ boosted activation, where w_t is the message sent in round $t \in [\kappa]$. Fix any $s \in \mathcal{S}$ and consider its path $n_{(s,0)} \to n_{(s,1)} \to \cdots \to n_{(s,|s|)}$. For each $i \in \{1, \ldots, |s|\}$, the node $n_{(s,i)}$ must be activated at some round $r_i \leq T$ via the incoming edge $(n_{(s,i-1)}, n_{(s,i)})$, and this activation must be boosted. Since $\mathbf{L}(n_{(s,i-1)}, n_{(s,i)}) = \{p_s(i)\}$, boosting this edge in round r_i implies

$$w_{r_i} = p_s(i).$$

Moreover, because $n_{(s,i)}$ cannot be activated before $n_{(s,i-1)}$, we can choose the rounds so that $1 \leq r_1 < r_2 < \cdots < r_{|s|} \leq T$. Hence s is a subsequence of w. Since s was arbitrary, every string in $\mathcal{S}$ is a subsequence of w, and $|w| = T = \kappa$. Therefore I_1 is a YES-instance of SCSB.

□

A.2. Hardness of Approximation

We now turn to the inapproximability of MDR. The problem is inherently bi-criteria, involving two competing parameters: the expected number of activated nodes κ and the number of rounds T. Consequently, approximation hardness arises along two natural dimensions, depending on which parameter is treated as fixed.

In the first direction, we fix a target number of expected activated nodes κ and seek to minimize the number of rounds T required to achieve this coverage. We show that this variant is hard to approximate within any factor better than $\log(n)$ of the optimum.

In the second direction, we fix the time horizon T and aim to maximize the number of nodes that can be activated within T rounds. Under the Exponential Time Hypothesis (ETH), we show that no algorithm can achieve an approximation guarantee of $n^{-1/(\log \log n)^c}$ of the optimum for some universal constant $c > 0$. This rules out polylogarithmic approximation guarentees under ETH.

To establish these inapproximability results, we formulate and analyze appropriate *gap versions* of the decision problem associated with MDR.

GAP-α-MDR (FIXED k, MINIMIZE T)

Input: A parameter $\alpha \in \mathbb{R}_{>1}$ and an instance $\mathcal{I} = (\mathcal{G}, \mathbf{S}, \mathbf{L}, \mathcal{B}, T, \kappa)$, consisting of a graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, an initial seed set $\mathbf{S} \subseteq \mathbf{V}$, a message set $\mathbf{L}$ with boosting parameters $\mathcal{B}$, and a target number of newly activated nodes $\kappa \in \mathbb{N}$.

Promise : Let $T^\star(\mathcal{I})$ denote the minimum number of rounds T for which there exists a policy π with

$$O_T(\mathbf{S}, \pi) \geq \kappa.$$

The input instance satisfies either

$$\text{YES: } T^\star(\mathcal{I}) \leq T \qquad \text{or} \qquad \text{NO: } T^\star(\mathcal{I}) \geq \alpha \cdot T,$$

Question: Decide whether the instance is a YES-instance or a NO-instance.

Theorem A.4 ((1-ε)ln n-Inapproximability for Fixed k)**.** *For every constant $\varepsilon > 0$, there is no polynomial time algorithm for* GAP-$(1-\varepsilon)\ln n$-MDR (FIXED k, MINIMIZE T)*. Equivalently, unless* $\mathrm{P} = \mathrm{NP}$*, no polynomial-time algorithm can approximate the minimum number of rounds $T^\star(\mathcal{I})$ within a factor better than $(1-\varepsilon)\ln n$, where $n := |\mathbf{V}|$.*

Proof. We reduce from the Set Cover problem. Dinur and Steurer (Dinur & Steurer, 2014) show that for every constant $\varepsilon > 0$, SET COVER is inapproximable in polynomial time within a factor $(1-\varepsilon)\ln |U|$, where $|U|$ denotes the universe size.

GAP-α-SET COVER

Input: An instance $\mathcal{I} = (U, \mathcal{F}, k)$ consisting of a universe U with $|U| = n$, a family of subsets $\mathcal{F} = \{F_1, \ldots, F_m\} \subseteq 2^U$, and an integer $k \in \mathbb{N}$.

Promise : Let $\mathrm{OPT}(\mathcal{I})$ denote the minimum number of sets whose union equals U. The instance satisfies either

$$\text{YES: } \mathrm{OPT}(\mathcal{I}) \leq k \qquad \text{or} \qquad \text{NO: } \mathrm{OPT}(\mathcal{I}) \geq \alpha \cdot k.$$

Question: Decide whether the instance is a YES-instance or a NO-instance.

Let $I_1 = (U, \mathcal{F}, k)$ be an instance of Gap-α-SET COVER. We construct an instance I_2 of Gap-α-MDR in the following ways:

We construct instance I_2 of MDR with a graph of $n = |U| + 1$ nodes and 1 starting seed node n_0. Add a message type for each subset in $\mathcal{F}$. In other words, let $\mathbf{L} = \mathcal{F}$. Let $p_e = \frac{1}{2n(n-1)T}$ for all $e \in E$, and let $\beta_\ell = 1 - \frac{1}{2n(n-1)T}$ under additive boosting and $\beta_\ell = 2n(n-1)T$ under multiplicative boosting for all messages $\ell \in \mathbf{L}$. In other words, each edge has a propogation probability of $\frac{1}{2n(n-1)T}$ if unboosted and a propogation probability of 1 if boosted.

For each $u \in U$, we add a node n_e to $\mathcal{G}$. Then, we add an edge (n_0, n_u) such that the label of (n_0, n_u) is the all the subsets $F \in \mathcal{F}$ such that $u \in F$.

$$\mathbf{L}(n_0, n_u) = \{F \in \mathcal{F} \mid u \in F\}$$

We now show that $I_1 = (U, \mathcal{F}, k)$ is a `YES` instance of Gap-α-SET COVER. if and only if $I_2 = (\mathcal{G}, \mathbf{S}, \mathbf{L}, \mathcal{B}, k, n-1)$ is a `YES` instance of Gap-α-MDR.

($\Rightarrow$) If I_1 is a `YES` instance of Gap-α-SET COVER, then there exists a collection of subsets $C \subseteq \mathcal{F}$ such that $|C| \leq k$ and $\bigcup_{F \in C} F = U$. Then, let π be the policy that relay messages that correspond to the subsets in C (in any order). We now show that under this policy, all nodes in $\mathcal{G}$ are activated and hence $O_k(I_2) = n$.

Fix an element $u \in U$. By assumption, there exists a $F \in C$ such that $u \in F$. Let r be the round where the message corresponding to F was sent. Hence, in round r, n_0 was active since n_0 was a seed node and $F \in \mathbf{L}(n_0, n_u) = \{F \in \mathcal{F} \mid u \in F\}$. Hence, node n_u is active after round k. As all nodes are activated after round k, I_2 is a `YES` instance.

($\Leftarrow$) If I_1 is a `NO` instance of Gap-α-SET COVER, then for all collections of subsets $C \subseteq \mathcal{F}$ such that $|C| \leq \alpha k - 1$, $\bigcup_{F \in C} F \neq U$.

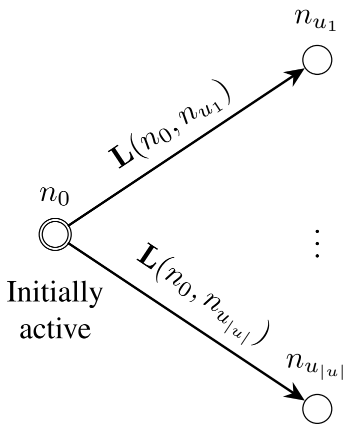

Figure 5. Set cover reduction where $U = \{u_1, \ldots, u_{|u|}\}$, where $\mathbf{L}(n_0, n_{u_j}) = \{F_i \in \mathcal{F} \mid u_j \in \mathcal{F}_i\}$ for $j \in \{1, \ldots, |u|\}$

Suppose, I_2 is not a NO instance of Gap-α-MDR. Then there exists a policy π such that

$$O_{\alpha k-1}(\mathbf{S}, \pi) \geq \kappa' = n - 1.$$

Lemma A.7 implies that for any policy (and in particular for π), the expected number of *unboosted* activations over T rounds is at most 0.5. Since, $O_T(\mathbf{S}, \pi) \geq n - 1$, the expected number of *boosted* activations is at least $n - 1.5$. For each trajectory, as the number of *boosted* activations be integral for each realisation, there must be a trajectory with $n - 1$ boosted activation.

Let $w = w_1 w_2 \cdots w_{\alpha k-1}$ be the sequence of messages broadcast by π in one such trajectory with $n - 1$ boosted activation, where w_t is the message sent in round t. Fix $u \in U$. As n_u was activated through a boosted activation, then there must be a $w_i \in w$ such that $w_i \in \mathbf{L}(n_0, n_u)$. Hence, as $\mathbf{L}(n_0, n_u) = \{F \in \mathcal{F} \mid u \in F\}$, this implies that $u \in w_i$. Since, u was arbitrary, every element $u \in U$ is contained by some subset $w_i \in w$ and there is a collection of subset w such that $|w| \leq \alpha k - 1$ and $\bigcup_{F \in w} F \neq U$ and we have arrived at a contradiction.

Observe that the reduction preserves the approximation gap: a cover of size at most k corresponds to a policy achieving activation within k rounds, while any solution requiring at least αk sets translates to requiring at least αk rounds. Moreover, since the constructed graph has $n = |U| + 1$ nodes, the universe size and the instance size differ only by a constant factor. Consequently, the $(1 - \varepsilon) \ln n$ inapproximability bound for SET COVER is preserved under the reduction.

□

GAP-α-MDR (FIXED T, MAXIMIZE κ)

Input: An instance $\mathcal{I} = (\mathcal{G}, \mathbf{S}, \mathbf{L}, \mathcal{B}, T, \kappa)$, consisting of a graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, an initial seed set $\mathbf{S} \subseteq \mathbf{V}$, a message set $\mathbf{L}$ with boosting parameters $\mathcal{B}$, a fixed number of rounds $T \in \mathbb{N}$, and a target number of newly activated nodes $\kappa \in \mathbb{N}$.

Promise : Let $\kappa^\star(\mathcal{I})$ denote the maximum expected number of newly activated nodes achievable within T rounds, i.e.,

$$\kappa^\star(\mathcal{I}) := \max_\pi O_T(\mathbf{S}, \pi).$$

The input instance satisfies either

$$\text{YES: } \kappa^\star(\mathcal{I}) \geq \kappa \qquad \text{or} \qquad \text{NO: } \kappa^\star(\mathcal{I}) \leq \frac{\kappa}{\alpha}.$$

Question: Decide whether the instance is a YES-instance or a NO-instance.

Theorem A.5 (Strong Inapproximability for Fixed T)**.** *Assuming the Exponential Time Hypothesis, there exists a universal constant $c > 0$ such that there is no polynomial time algorithm for* GAP-α-MDR (FIXED T, MAXIMIZE κ) *for*

$$\alpha = n^{1/(\log \log n)^c},$$

where $n := |\mathbf{V}|$. Equivalently, under ETH, no polynomial-time algorithm can approximate $\kappa^\star(\mathcal{I})$ within a factor better than $n^{-1/(\log\log n)^c}$.

Proof. We reduce from the densest k subgraph (DkS). Manurangsi (Manurangsi, 2017) show that for that under the epxonential time hypothesis, there is no polynomial time algorithm that approximates DkS to within a $n^{-1/(\log\log n)^c}$ factor of the optimum, where $c > 0$ is a universal constant independent of n.

GAP-α-DKS

Input: An instance $\mathcal{I} = (G, k, M)$ consisting of an undirected graph $G = (V, E)$ with $|V| = n$, an integer $k \in \mathbb{N}$, and a threshold $M \in \mathbb{N}$.

Promise : Let

$$\mathrm{OPT}(\mathcal{I}) := \max_{S \subseteq V,\ |S|=k} |E(S)|$$

denote the maximum number of edges induced by any k-vertex subgraph. The instance satisfies either

$$\text{YES: } \mathrm{OPT}(\mathcal{I}) \geq M \qquad \text{or} \qquad \text{NO: } \mathrm{OPT}(\mathcal{I}) \leq \frac{M}{\alpha}.$$

Question: Decide whether the instance is a YES-instance or a NO-instance.

Let $I_1 = (G_1 = (V_1, E_1), k, M)$ be an instance of Gap-α-DKS. Without loss of generality, assume that $k \geq 2$ and hence $\mathrm{OPT}(\mathcal{I}) \geq 1$. We construct an instance of I_2 of Gap-$\alpha/2$-MDR in the following ways:

We construct instance I_2 of MDR with a graph of $n = |E_1| \cdot (k^3 - k^2 + k^3|E_1|)$ nodes. Add a message type for each node $u \in v_1$. In other words, $\mathbf{L} = V_1$. In our reduction, the first k message sent correspond to the selection of the k-vertex subgraph S.

For each edge $e \in E$, we create a gadget of 2 components, C_1^e with $k^3 - k^2$ node, and C_2^e with $k^3|E_1|$ nodes.

Let $e = (u, v)$. C_1^e consists of $k(k-1)$ disjoint directed path subgraphs each of length k. Path graph $g_{(e,i,j)}$ represents the case where u is the ith message selected and v is the jth message selected for a total of $k(k-1)$ path subgraphs. Path subgraph $g_{(e,i,j)}$ consisting of $k+1$ nodes

$$n_{((e,i,j),0)}, \ldots, n_{((e,i,j),k)}.$$

For each $x \in \{1, \ldots, k\}$, we add a directed edge $(n_{((e,i,j),x-1)}, n_{((e,i,j),x)})$ and assign it the label

$$\mathbf{L}\big(n_{((e,i,j),x-1)}, n_{((e,i,j),x)}\big) = \begin{cases} \{u\}, & \text{if } i = x, \\ \{v\}, & \text{if } j = x, \\ V_1, & \text{otherwise.} \end{cases}$$

.

C_2^e consists of $k^3|E_1|$ nodes that is activated if any of the path subgraph is activated. For all combinations (e, i, j) and all $n' \in C_2^e$, add an edge between $n_{((e,i,j),k)}$ and n' such that $\mathbf{L}(n_{((e,i,j),k)}, n') = V_1$.

The active nodes are $n_{((e,i,j),0)}$ for all combinations of (e, i, j) for all edges $e \in E_1$.

Let $p_e = \frac{1}{n^2(n-1)T}$ for all $e \in E$, and let $\beta_\ell = 1 - \frac{1}{n^2(n-1)T}$ under additive boosting and $\beta_\ell = n^2(2n-1)T$ under multiplicative boosting for all messages $\ell \in \mathbf{L}$. In other words, each edge has a propogation probability of $\frac{1}{n^2(n-1)T}$ if unboosted and a propogation probability of 1 if boosted.

We now show that $I_1 = (G_1, k, M)$ is a `YES` instance of Gap-α-DKS if and only if $I_2 = (\mathcal{G}, \mathbf{S}, \mathbf{L}, \mathcal{B}, k+1, Mk^3|E_1|)$ is a `YES` instance of Gap-$\alpha/2$-MDR.

($\Rightarrow$) If I_1 is a `YES` instance of Gap-α-DKS, then there exists a subgraph $S \subseteq V_1$ such that $|S| \leq k$ and $|E(S)| \geq M$. Then, let π be the policy that relay messages that correspond to the nodes in S and send any non-empty message at round $k+1$. We now show that under this policy, for every edge in $e \in E(s)$, all edges in C_2^e are activated for a minimum of $Mk^3|E_1|$ nodes.

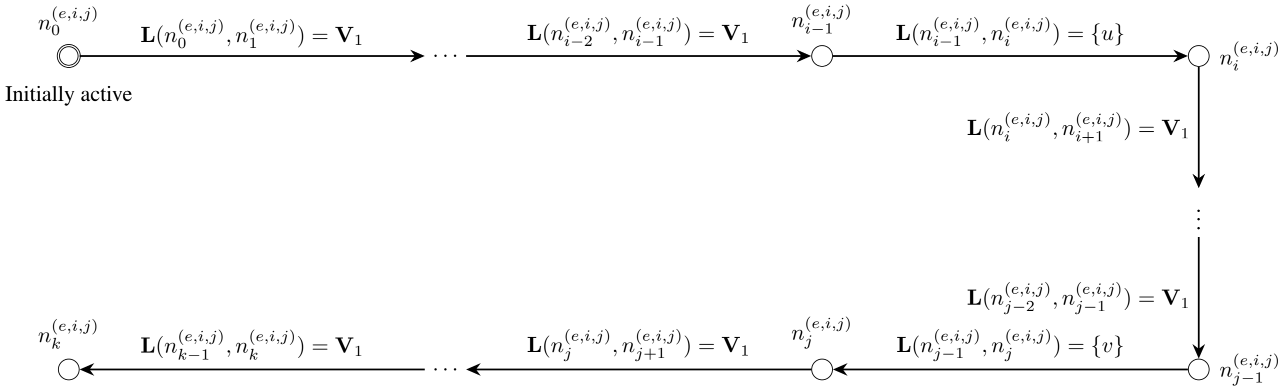

Figure 6. Path graph $g_{(e,i,j)}$. All messages are $\mathbf{V}_1$ except for the messages right before $n_i^{(e,i,j)}$ and $n_j^{(e,i,j)}$, which are $\{u\}$ and $\{v\}$ respectively. In this manner, $n_k^{(e,i,j)}$ will be activated only if messages u and v are broadcasted at timesteps i and j respectively.

Let $e = (u, v)$. As $e \in E(S)$, both the nodes u and v is part of the subgraph S. Hence, there is a round i where message corresponding to u is chosen and a round j where node v is chosen. After k timesteps, all nodes in the path subgraph $g_{(e,i,j)}$ is activated. As $n_{((e,i,j),k)}$ is active after timestep k, and $\mathbf{L}(n_{((e,i,j),k)}, n') = V_1$ for all $n' \in C_2^e$. As long as the message chosen in timestep $k+1$ is not the empty message, all nodes in C_2^e will be activated.

Hence, as there are $k^3|E_1|$ nodes in C_2^e for all $e \in E(S)$ and $|E(S)| \geq M$, the policy π activates at least $Mk^3|E_1|$ nodes after $k+1$ timesteps and I_2 is a `YES` instance.

$(\Leftarrow)$ If I_1 is a `NO` instance of Gap-α-DKS, then for all subgraph $S \subseteq V_1$ such that $|S| \leq k$ and $|E(S)| \leq \alpha M$. Suppose, I_2 is not a `NO` instance of Gap-α-MDR. Then there exists a policy π such that

$$O_{k+1}(\mathbf{S}, \pi) \geq \frac{2Mk^3|E_1|}{\alpha}.$$

We now split the cases based on whether in the realisation there was an unboosted activation or not. Let $\mathcal{E}$ be the event when under π there is no unboosted activation and $\bar{\mathcal{E}}$ be the event when under π there is at least 1 unboosted activation. By the linearity of expectation:

$$O_{k+1}(\mathbf{S}, \pi) = \Pr(\mathcal{E}) \cdot \mathbb{E}\left[O_{k+1}(\mathbf{S}, \pi) \mid \mathcal{E}\right] + \Pr(\bar{\mathcal{E}}) \cdot \mathbb{E}\left[O_{k+1}(\mathbf{S}, \pi) \mid \bar{\mathcal{E}}\right]$$

Since unboosted probabilities are at most $\frac{1}{n^2(n-1)T}$, Lemma A.7 implies that for any policy (and in particular for π), the expected number of *unboosted* activations over T rounds is at most $\frac{1}{n}$. Hence, by Markov's Inequality, $\mathcal{P}(\bar{\mathcal{E}}) \leq \frac{1}{n}$. As the number of newly activated nodes is upper bounded by the number of nodes, $\mathbb{E}\left[O_{k+1}(\mathbf{S}, \pi) \mid \bar{\mathcal{E}}\right] \leq n$ and hence the term

$$\Pr(\bar{\mathcal{E}}) \cdot \mathbb{E}\left[O_{k+1}(\mathbf{S}, \pi) \mid \bar{\mathcal{E}}\right] \leq 1$$

Furthermore, as $\Pr(\mathcal{E}) \leq 1$, after rearranging the terms, we have that

$$\mathbb{E}\left[O_{k+1}(\mathbf{S}, \pi) \mid \mathcal{E}\right] \geq O_{k+1}(\mathbf{S}, \pi) - 1 \geq \frac{2Mk^3|E_1|}{\alpha} - 1 > \frac{Mk^3|E_1|}{\alpha} + |E_1|(k^3 - k))$$

where the last inequality holds as we only consider the case when $k \geq 2$ and hence $M \geq 1$.

As $|\bigcup_{e \in E_1} C_1^e| = |E_1|(k^3 - k))$ and $E\left[O_{k+1}(\mathbf{S}, \pi) \mid UB\right] > \frac{Mk^3|E_1|}{\alpha} + |E_1|(k^3 - k))$, this implies that the number of activated nodes in $\bigcup_{e \in E_1} C_2^e$ is more than $\frac{Mk^3|E_1|}{\alpha}$. Let $E' \subseteq E_1$ be the set of edges such that at least 1 node in C_e^2 is activated. By Pigeonhole Principle, as $|C_2^e| = k^3|E_1|$ for all $e \in E_1$ we establish that $|E'| > M/\alpha$.

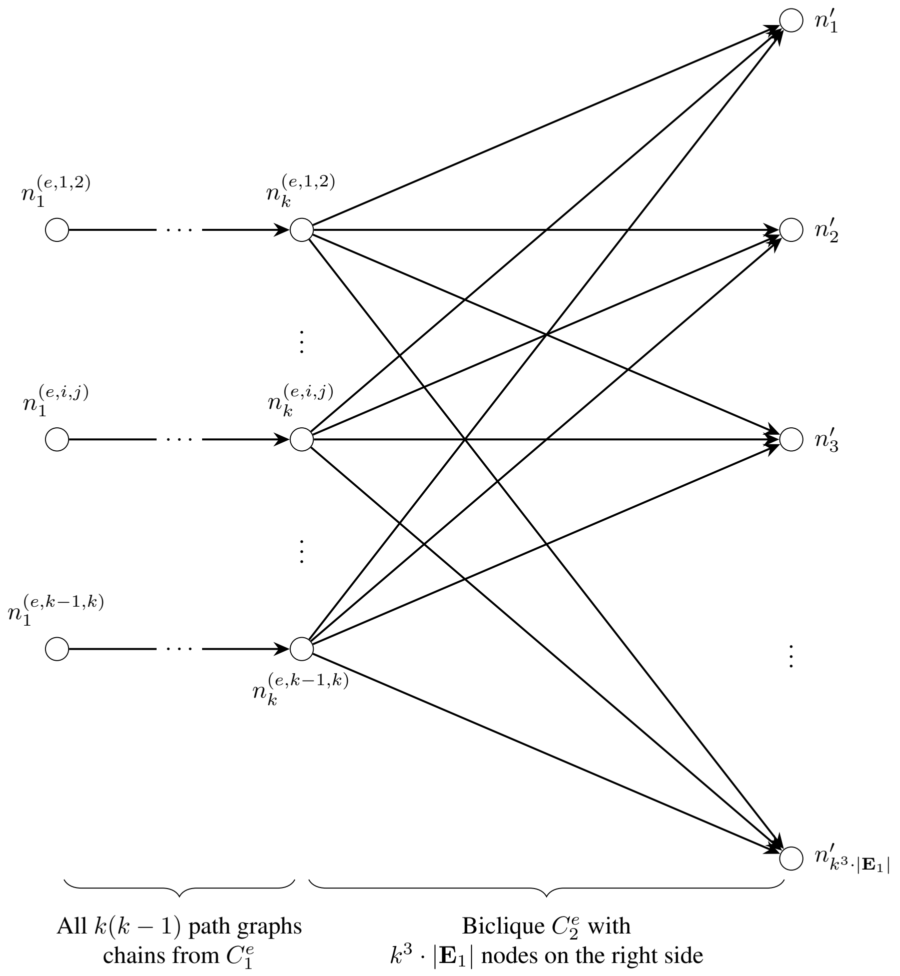

Figure 7. DKS blow up.

Fix $e \in E_1$ and let $e = (u, v)$. We note that if $n' \in C_e^2$ is only active after timestep $k+1$ if and only there exists $i, j \in [k]$ such that $n_{((e,i,j),k)}$ is active after timestep k. Furthermore, as we are in the case where there are no unboosted activations, if $n_{((e,i,j),k)}$ is active after timestep k, then the ith message must be u and the jth message must be v. Let $S \subseteq V_1$ be the subgraph that contains all the nodes sent as a message in the first k timestep of policy π in the realistation when there are no unboosted activation. We note that $u, v \in S$ and $(u, v) \in E(S)$.

Hence, as E' is the set of edges such that at least 1 node in C_2^e is activated, $E' \subseteq E(S)$. Thus, $|S| \leq k$ and $|E(S)| \geq |E'| > M/\alpha$ contradicting our premise that I_1 is a `NO` instance.

One subtlety is that the number of nodes in the constructed instance I_2 is larger than the number of vertices in the original DKS instance I_1. Let $n_0 := |V_1|$ and let $n := |\mathcal{V}|$ denote the number of nodes in the constructed graph $\mathcal{G}$. Suppose there were a polynomial-time algorithm that approximates MDR within a factor

$$N^{-1/(\log\log N)^{c_0}}$$

for some constant $c_0 > 0$. Since $n = \mathrm{poly}(n_0)$, we have $\log\log n = \Theta(\log\log n_0)$ and hence there exists a constant $c_1 > 0$ such that

$$n^{-1/(\log\log n)^{c_0}} \;\geq\; n_0^{-1/(\log\log n_0)^{c_1}}.$$

Therefore, such an algorithm would yield an approximation for DKS that contradicts the inapproximability of Manurangsi (Manurangsi, 2017). Consequently, MDR inherits the same $n^{-1/(\log\log n)^c}$- inapproximability, when stated in terms of the input size of MDR.

□

A.3. A technical lemma: unboosted activations are negligible when base probabilities are tiny

Definition A.6 (Boosted vs. unboosted activation)**.** Consider an activation event in which u activates v in round t via the edge (u, v). This activation is *boosted* if the broadcast $\ell^{(t)}$ satisfies $\ell^{(t)} \in \mathbf{L}(u, v)$, and *unboosted* otherwise.

Lemma A.7 (Bounding unboosted activations)**.** *Let* $n := |\mathbf{V}|$. *Suppose every directed edge* (u, v) *has* unboosted *propagation probability at most*

$$p_{u,v} \leq \frac{\varepsilon}{n(n-1)T}.$$

Then for any policy π*, the expected total number of unboosted activations over* T *rounds is at most* ε*.*

Proof. Fix any policy π and any round t. For each ordered pair (u, v), the probability that an unboosted activation occurs along (u, v) in round t is at most $p_{u,v}$. There are at most $n(n-1)$ ordered pairs, so the expected number of unboosted activations in round t is at most $n(n-1) \cdot \frac{\varepsilon}{n(n-1)T} = \frac{\varepsilon}{T}$. Summing over $t = 1, \ldots, T$ and using linearity of expectation yields the claim. □

A useful corollary is that if $\varepsilon \leq 1/2$, then with constant probability *no* unboosted activations occur. Formally, let U be the (random) number of unboosted activations over T rounds. Then $\mathbb{E}[U] \leq 1/2$, so by Markov, $\Pr[U \geq 1] \leq 1/2$ and hence $\Pr[U = 0] \geq 1/2$.

B. Greedy proof

B.1. Proof of Theorem 5.1

To prove Theorem 5.1, we rely on the following helper lemmas.

Lemma B.1. *Fix an* MDR *instance with horizon* $T \in \mathbb{N}_{\geq 2}$ *and discount factor* $\gamma \in (0, 1)$*. For any policy* π*, we have*

$$\widetilde{R}_T(\pi) = \mathop{\mathbb{E}}_{\pi}\left[\sum_{t=1}^{T-1} w(\mathbf{A}^{(t)}) \cdot \Gamma_t\right] \tag{6}$$

where $\Gamma_t = \sum_{i=t+1}^{T} \gamma^{i-1} = \gamma^t \cdot \sum_{j=0}^{T-t-1} \gamma^j = \gamma^t \cdot \frac{1-\gamma^{T-t}}{1-\gamma}$*, for* $1 \leq t \leq T-1$*.*

Proof. For $t \in [T-1]$, recall that the random set of newly activated nodes $\mathbf{A}^{(t)} = \mathbf{S}^{(t+1)} \setminus \mathbf{S}^{(t)}$ and $\mathbf{S}^{(t)} = \mathbf{S}^{(1)} \cup \bigcup_{i=1}^{t-1} \mathbf{A}^{(i)}$. By construction, the sets $\mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(t)}$ are mutually disjoint, so $w(\mathbf{S}^{(t)}) = w(\mathbf{S}^{(1)}) + \sum_{i=1}^{t-1} w(\mathbf{A}^{(i)})$. Multiplying γ^{t-1} and summing over $t \in [T]$, we get

$$\begin{aligned}
\sum_{t=1}^{T} \gamma^{t-1} \cdot w(\mathbf{S}^{(t)}) &= \sum_{t=1}^{T} \gamma^{t-1} \cdot \left(w(\mathbf{S}^{(1)}) + \sum_{i=1}^{t-1} w(\mathbf{A}^{(i)})\right) && \text{(From above)} \\
&= w(\mathbf{S}^{(1)}) \cdot \left(\sum_{t=1}^{T} \gamma^{t-1}\right) + \sum_{t=1}^{T-1} w(\mathbf{A}^{(t)}) \cdot \left(\sum_{i=t+1}^{T} \gamma^{i-1}\right) \\
& && \text{(Rearranging and interchanging order of summation)} \\
&= w(\mathbf{S}^{(1)}) \cdot \left(\sum_{t=1}^{T} \gamma^{t-1}\right) + \sum_{t=1}^{T-1} w(\mathbf{A}^{(t)}) \cdot \Gamma_t && \text{(Definition of } \Gamma_t\text{)}
\end{aligned}$$

Taking expectations on both sides, we see that

$$\begin{aligned}
\mathbb{E}_{\pi}\left[\sum_{t=1}^{T}\gamma^{t-1}w(\mathbf{S}^{(t)})\right] &= \mathbb{E}_{\pi}\left[w(\mathbf{S}^{(1)})\cdot\left(\sum_{t=1}^{T}\gamma^{t-1}\right)+\sum_{t=1}^{T-1}w(\mathbf{A}^{(i)})\cdot\Gamma_t\right]\\
&= w(\mathbf{S}^{(1)})\cdot\left(\sum_{t=1}^{T}\gamma^{t-1}\right)+\mathbb{E}_{\pi}\left[\sum_{t=1}^{T-1}w(\mathbf{A}^{(i)})\cdot\Gamma_t\right]\\
&\qquad\text{(Linearity of expectation and since the first term is independent of } \pi\text{)}
\end{aligned}$$

Eq. (6) then follows by recalling the definitions of $\widetilde{R}_T(\pi)$ and $R_T(\pi)$ from Eq. (3) and Eq. (2). ☐

Lemma B.2. *Fix any horizon $T \in \mathbb{N}_{\geq 2}$ and initial active set $\mathbf{S} \subseteq \mathbf{V}$. We have*

$$V_T(\mathbf{S}) = \max_{\ell\in\mathbf{L}\cup\{\emptyset\}}\left\{\mathbb{E}_{\mathbf{A}^{(1)}}\left[w(\mathbf{A}^{(1)})+\gamma\cdot V_{T-1}(\mathbf{S}^{(1)}\cup\mathbf{A}^{(1)})\,\middle|\,\mathbf{S}^{(1)}=\mathbf{S},\ell\right]\right\}$$

Proof. We will prove it by showing that inequality holds in both directly separately, thus establishing equality.

Fix an arbitrary message $\ell \in \mathbf{L}\cup\{\emptyset\}$ and consider any policy π whose first action given $\mathbf{S}^{(1)} = \mathbf{S}$ is $\ell^{(1)} = \ell$. Observe that

$$\begin{aligned}
&\mathbb{E}_{\pi}\left[\sum_{t=1}^{T}\gamma^{t-1}\cdot\mathbb{E}_{\mathbf{A}^{(t)}}[w(\mathbf{A}^{(t)})\mid\mathbf{S}^{(t)},\ell^{(t)}]\,\middle|\,\mathbf{S}^{(1)}=\mathbf{S}\right]\\
&=\mathbb{E}_{\pi}\left[\mathbb{E}_{\mathbf{A}^{(1)}}\left[w(\mathbf{A}^{(1)})\mid\mathbf{S}^{(1)},\ell^{(1)}\right]+\gamma\cdot\sum_{s=1}^{T-1}\gamma^{s-1}\cdot\mathbb{E}_{\mathbf{A}^{(s+1)}}\left[w(\mathbf{A}^{(s+1)})\mid\mathbf{S}^{(s+1)},\ell^{(s+1)}\right]\right]\\
&\qquad\qquad\text{(Pulling out first term and re-indexing)}\\
&\leq\mathbb{E}_{\pi}\left[\mathbb{E}_{\mathbf{A}^{(1)}}\left[w(\mathbf{A}^{(1)})\mid\mathbf{S}^{(1)},\ell^{(1)}\right]+\gamma\cdot V_{T-1}(\mathbf{S}^{(1)}\cup\mathbf{A})\,\middle|\,\mathbf{S}^{(1)}=\mathbf{S},\ell^{(1)}\right]\qquad\text{(Definition of }V_{T-1}(\mathbf{S}^{(1)}\cup\mathbf{A})\text{)}\\
&=\mathbb{E}_{\mathbf{A}^{(1)}}\left[w(\mathbf{A}^{(1)})+\gamma\cdot V_{T-1}(\mathbf{S}^{(1)}\cup\mathbf{A})\,\middle|\,\mathbf{S}^{(1)}=\mathbf{S},\ell^{(1)}\right]\qquad\text{(Since policy }\pi\text{ picks }\ell^{(1)}\text{)}
\end{aligned}$$

Thus,

$$\begin{aligned}
V_T(\mathbf{S}) &= \max_{\pi}\mathbb{E}_{\pi}\left[\sum_{t=1}^{T}\gamma^{t-1}\cdot\mathbb{E}_{\mathbf{A}^{(t)}}[w(\mathbf{A}^{(t)})\mid\mathbf{S}^{(t)},\ell^{(t)}]\,\middle|\,\mathbf{S}^{(1)}=\mathbf{S}\right]\\
&\leq\mathbb{E}_{\mathbf{A}^{(1)}}\left[w(\mathbf{A}^{(1)})+\gamma\cdot V_{T-1}(\mathbf{S}^{(1)}\cup\mathbf{A})\,\middle|\,\mathbf{S}^{(1)}=\mathbf{S},\ell^{(1)}\right]\qquad\text{(From above)}\\
&\leq\max_{\ell\in\mathbf{L}\cup\{\emptyset\}}\left\{\mathbb{E}_{\mathbf{A}^{(1)}}\left[w(\mathbf{A}^{(1)})+\gamma\cdot V_{T-1}(\mathbf{S}^{(1)}\cup\mathbf{A}^{(1)})\,\middle|\,\mathbf{S}^{(1)}=\mathbf{S},\ell\right]\right\}
\end{aligned}$$

To establish equality, we now show that $V_T(\mathbf{S}) \geq \max_{\ell\in\mathbf{L}\cup\{\emptyset\}}\left\{\mathbb{E}_{\mathbf{A}^{(1)}}\left[w(\mathbf{A}^{(1)})+\gamma\cdot V_{T-1}(\mathbf{S}^{(1)}\cup\mathbf{A}^{(1)})\,\middle|\,\mathbf{S}^{(1)}=\mathbf{S},\ell\right]\right\}$.

Fix an arbitrary message $\ell \in \mathbf{L}\cup\{\emptyset\}$. Consider the policy that broadcasts ℓ in the first round, then after observing $\mathbf{S}^{(2)}\cup\mathbf{S}^{(1)}\cup\mathbf{A}^{(1)}$, follows an optimal length $T-1$ policy from state $\mathbf{S}^{(2)}$. By definition, this continuation achieves value exactly $V_{T-1}(\mathbf{S}^{(1)}\cup\mathbf{A}^{(1)})$, for each realized $\mathbf{A}^{(1)}$. Therefore, the expected value of this composite policy is exactly

$$\mathbb{E}_{\mathbf{A}^{(1)}}\left[w(\mathbf{A}^{(1)})+\gamma\cdot V_{T-1}(\mathbf{S}^{(1)}\cup\mathbf{A}^{(1)})\,\middle|\,\mathbf{S}^{(1)},\ell\right]$$

Since this is achievable for every fixed $\ell \in \mathbf{L}\cup\{\emptyset\}$, it is also achievable for the maximizing ℓ, implying that

$$V_T(\mathbf{S}^{(1)}) \geq \max_{\ell\in\mathbf{L}\cup\{\emptyset\}}\mathbb{E}_{\mathbf{A}^{(1)}}\left[w(\mathbf{A}^{(1)})+\gamma\cdot V_{T-1}(\mathbf{S}^{(1)}\cup\mathbf{A}^{(1)})\,\middle|\,\mathbf{S}^{(1)},\ell\right]$$

In other words, $V_T(\mathbf{S}) = \max_{\ell\in\mathbf{L}\cup\{\emptyset\}}\left\{\mathbb{E}_{\mathbf{A}^{(1)}}\left[w(\mathbf{A}^{(1)})+\gamma\cdot V_{T-1}(\mathbf{S}^{(1)}\cup\mathbf{A}^{(1)})\,\middle|\,\mathbf{S}^{(1)}=\mathbf{S},\ell\right]\right\}$, as desired. ☐

Lemma B.3 (Lipschitzness of the one-step gain under ρ)**.** *For any active set* $\mathbf{S} \subseteq \mathbf{V}$*, subset* $\mathbf{A} \subseteq \mathbf{V} \setminus \mathbf{S}$*, and any message* $\ell \in \mathbf{L} \cup \{\emptyset\}$*, we have*

$$\mathop{\mathbb{E}}_{\mathbf{A}'}\left[g(\mathbf{S} \cup \mathbf{A}) \mid \mathbf{S}, \ell\right] \leq g(\mathbf{S}) + \rho \cdot w(\mathbf{A}) \tag{7}$$

Proof. The claim holds trivially when $\mathbf{A} = \emptyset$ since $\mathbb{E}_{\mathbf{A}'}\left[g(\mathbf{S}) \mid \mathbf{S}, \ell\right] = g(\mathbf{S})$ and $\rho \geq 0$. Now, suppose $\mathbf{A} \neq \emptyset$ and enumerate its elements as $\mathbf{A} = \{v_1, \ldots, v_r\}$. Let us define the following sequence of active sets $\mathbf{S}^{(j)} = \mathbf{S} \cup \{v_1, \ldots, v_j\}$ for $j \in \{0, 1, \ldots, r\}$ with $\mathbf{S}^{(r)} = \mathbf{S} \cup \mathbf{A}$. By definition of ρ (see Appendix B.2), we have that

$$\mathop{\mathbb{E}}_{\mathbf{A}'}\left[w(\mathbf{A}') \mid \mathbf{S}^{(j)}, \ell\right] - \mathop{\mathbb{E}}_{\mathbf{A}'}\left[w(\mathbf{A}') \mid \mathbf{S}^{(j-1)}, \ell\right] \leq \rho \cdot w_{v_j} \tag{8}$$

for all $j \in \{1, \ldots, r\}$. Therefore, telescoping yields

$$\begin{aligned}
\mathop{\mathbb{E}}_{\mathbf{A}'}\left[w(\mathbf{A}') \mid \mathbf{S} \cup \mathbf{A}, \ell\right] - \mathop{\mathbb{E}}_{\mathbf{A}'}\left[w(\mathbf{A}') \mid \mathbf{S}, \ell\right] &= \mathop{\mathbb{E}}_{\mathbf{A}'}\left[w(\mathbf{A}') \mid \mathbf{S}^{(r)}, \ell\right] - \mathop{\mathbb{E}}_{\mathbf{A}'}\left[w(\mathbf{A}') \mid \mathbf{S}^{(0)}, \ell\right] \\
&= \sum_{j=1}^{r} \left(\mathop{\mathbb{E}}_{\mathbf{A}'}\left[w(\mathbf{A}') \mid \mathbf{S}^{(j)}, \ell\right] - \mathop{\mathbb{E}}_{\mathbf{A}'}\left[w(\mathbf{A}') \mid \mathbf{S}^{(j-1)}, \ell\right] \right) \\
&\leq \sum_{j=1}^{r} \rho \cdot w_{v_j} && \text{(By Eq. (8))} \\
&= \rho \cdot w(\mathbf{A})
\end{aligned}$$

In other words, $\mathbb{E}_{\mathbf{A}'}\left[w(\mathbf{A}') \mid \mathbf{S} \cup \mathbf{A}, \ell\right] - \mathbb{E}_{\mathbf{A}'}\left[w(\mathbf{A}') \mid \mathbf{S}, \ell\right] \leq \rho \cdot w(\mathbf{A})$.

Recall from Eq. (4) that $g(\mathbf{S}) = \max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \mathbb{E}_{\mathbf{A}}[w(\mathbf{A}) \mid \mathbf{S}, \ell]$ for any active set $\mathbf{S}$ and message ℓ. Let $\ell^* \in \arg\max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \mathbb{E}_{\mathbf{A}}[w(\mathbf{A}) \mid \mathbf{S}, \ell]$ so that $g(\mathbf{S} \cup \mathbf{A}) = \mathbb{E}_{\mathbf{A}}[w(\mathbf{A}) \mid \mathbf{S} \cup \mathbf{A}, \ell^*]$. Then,

$$\begin{aligned}
g(\mathbf{S} \cup \mathbf{A}) &= \mathop{\mathbb{E}}_{\mathbf{A}'}\left[w(\mathbf{A}') \mid \mathbf{S} \cup \mathbf{A}, \ell^*\right] \\
&\leq \mathop{\mathbb{E}}_{\mathbf{A}'}\left[w(\mathbf{A}') \mid \mathbf{S}, \ell^*\right] + \rho \cdot w(\mathbf{A}) \\
&\leq \max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \mathop{\mathbb{E}}_{\mathbf{A}'}\left[w(\mathbf{A}') \mid \mathbf{S}, \ell\right] + \rho \cdot w(\mathbf{A}) \\
&= g(\mathbf{S}) + \rho \cdot w(\mathbf{A}) && \text{(By definition of } g(\mathbf{S})\text{)}
\end{aligned}$$

This concludes the proof. ☐

Lemma B.4. *Fix any arbitrary horizon* $t \in \mathbb{N}_{\geq 1}$ *and initial active set* $\mathbf{S}$*. We have* $V_t(\mathbf{S}) \leq c_t \cdot g(\mathbf{S})$ *where* $c_t = \sum_{k=0}^{t-1} \left(\gamma \cdot (1 + \rho)\right)^k$.

Proof. We prove by applying induction on t.

Base case ($t = 1$**).** Fix an arbitrary active set $\mathbf{S} \subseteq \mathbf{V}$. By definition of $\mathbf{V}_1(\mathbf{S})$ and $g(\mathbf{S})$ (see Eq. (4)), we have

$$V_1(\mathbf{S}) = \max_{\pi} \mathop{\mathbb{E}}_{\pi}\left[\mathop{\mathbb{E}}_{\mathbf{A}^{(1)}}\left[w(\mathbf{A}^{(1)}) \mid \mathbf{S}^{(1)}, \ell^{(1)}\right] \,\middle|\, \mathbf{S}^{(1)} = \mathbf{S} \right] = \max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \mathop{\mathbb{E}}_{\mathbf{A}}[w(\mathbf{A}) \mid \mathbf{S}, \ell] = g(\mathbf{S})$$

Furthermore, since $c_1 = \sum_{k=0}^{0} \left(\gamma \cdot (1 + \rho)\right)^k = 1$, we have $V_1(\mathbf{S}) = c_1 \cdot g(\mathbf{S})$.

Inductive case ($t \geq 2$**).** Suppose the claim holds for $t - 1$, i.e., $V_{t-1}(\mathbf{S}') \leq c_{t-1} \cdot g(\mathbf{S}')$ for *any* active set $\mathbf{S}' \subseteq \mathbf{V}$. Fix an arbitrary active set $\mathbf{S} \subseteq \mathbf{V}$. Then,

$$\begin{aligned}
V_t(\mathbf{S}^{(1)}) &= \max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \left\{ \mathop{\mathbb{E}}_{\mathbf{A}^{(1)}} \left[w(\mathbf{A}^{(1)}) + \gamma \cdot V_{t-1}(\mathbf{S}^{(1)} \cup \mathbf{A}^{(1)}) \;\middle|\; \mathbf{S}^{(1)} = \mathbf{S}, \ell \right] \right\} && \text{(From Lemma B.2)} \\
&\leq \max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \left\{ \mathop{\mathbb{E}}_{\mathbf{A}^{(1)}} \left[w(\mathbf{A}^{(1)}) + \gamma \cdot c_{t-1} \cdot g(\mathbf{S}^{(1)} \cup \mathbf{A}^{(1)}) \;\middle|\; \mathbf{S}^{(1)} = \mathbf{S}, \ell \right] \right\} && \text{(By induction hypothesis)} \\
&\leq \max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \left\{ \mathop{\mathbb{E}}_{\mathbf{A}^{(1)}} \left[w(\mathbf{A}^{(1)}) + \gamma \cdot c_{t-1} \cdot \left(g(\mathbf{S}^{(1)}) + \rho \cdot w(\mathbf{A}^{(1)}) \right) \;\middle|\; \mathbf{S}, \ell \right] \right\} && \text{(From Eq. (7))} \\
&= \gamma \cdot c_{t-1} \cdot g(\mathbf{S}) + \max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \left\{ \mathop{\mathbb{E}}_{\mathbf{A}^{(1)}} \left[(1 + \gamma \cdot c_{t-1} \cdot \rho) \cdot w(\mathbf{A}^{(1)}) \;\middle|\; \mathbf{S}, \ell \right] \right\} \\
& && \text{(Since } \gamma \cdot c_{t-1} \cdot g(\mathbf{S}) \text{ is independent of } \ell\text{)} \\
&= \gamma \cdot c_{t-1} \cdot g(\mathbf{S}) + (1 + \gamma \cdot c_{t-1} \cdot \rho) \cdot \max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \mathop{\mathbb{E}}_{\mathbf{A}^{(1)}} \left[w(\mathbf{A}^{(1)}) \;\middle|\; \mathbf{S}, \ell \right] && \text{(Since } (1 + \gamma \cdot c_{t-1} \cdot \rho) \geq 0\text{)} \\
&= \gamma \cdot c_{t-1} \cdot g(\mathbf{S}) + (1 + \gamma \cdot c_{t-1} \cdot \rho) \cdot g(\mathbf{S}) && \text{(Definition of } g(\mathbf{S})\text{)} \\
&= (1 + \gamma \cdot c_{t-1} \cdot (1 + \rho)) \cdot g(\mathbf{S})
\end{aligned}$$

It remains to show that $c_t = (1 + \gamma \cdot c_{t-1} \cdot (1 + \rho))$. Observe that

$$c_t = \sum_{k=0}^{t-1} (\gamma \cdot (1+\rho))^k = 1 + \gamma \cdot (1+\rho) \cdot \sum_{k=0}^{t-2} (\gamma \cdot (1+\rho))^k = 1 + \gamma \cdot (1+\rho) \cdot c_{t-1}$$

This completes the proof. ☐

We are now ready to prove Theorem 5.1.

Theorem 5.1 (Parameterized approximation for π_{greedy})**.** *Fix an* MDR *instance with horizon* $T \in \mathbb{N}_{\geq 2}$ *and discount factor* $\gamma \in (0, 1)$. *Then, the greedy policy* π_{greedy} *achieves*

$$\widetilde{R}_T(\pi_{\text{greedy}}) \geq \alpha(T, \gamma, \rho) \cdot \max_{\pi} \widetilde{R}_T(\pi)$$

where the approximation factor is

$$\alpha(T, \gamma, \rho) = \frac{1}{\sum_{k=0}^{T-1} (\gamma \cdot (1+\rho))^k} \cdot \frac{1 - \gamma}{1 - \gamma^{T-1}}$$

Proof. Let us define the discounted new-activation objective as

$$H_T(\pi) = \mathop{\mathbb{E}}_{\pi} \left[\sum_{t=1}^{T-1} \gamma^{t-1} \cdot w(\mathbf{A}^{(t)}) \right] = \mathop{\mathbb{E}}_{\pi} \left[\sum_{t=1}^{T-1} \gamma^{t-1} \cdot \mathop{\mathbb{E}}_{\mathbf{A}^{(t)}} \left[w(\mathbf{A}^{(t)}) \mid \mathbf{S}^{(t)}, \ell^{(t)} \right] \right]$$

where the second equality is because we have $\mathbb{E}_{\pi}[w(\mathbf{A}^{(t)})] = \mathbb{E}_{\pi}\left[\mathbb{E}_{\mathbf{A}^{(t)}}\left[w(\mathbf{A}^{(t)}) \mid \mathbf{S}^{(t)}, \ell^{(t)}\right]\right]$ for all $t \in [T]$, via the law of iterated expectations. By definition, we see that $V_T(\mathbf{S}_1) = \max_{\pi} H_T(\pi)$.

By definition of π_{greedy}, we choose $\ell^{(t)}_{\text{greedy}} \in \arg\max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \mathbb{E}_{\mathbf{A}^{(t)}}[w(\mathbf{A}^{(t)}) \mid \mathbf{S}^{(t)}, \ell^{(t)}]$. In other words, by definition of $g(\mathbf{S}^{(t)})$ (see Eq. (4)), we have $\mathbb{E}_{\mathbf{A}^{(t)}}[w(\mathbf{A}^{(t)}) \mid \mathbf{S}^{(t)}, \ell^{(t)}] = g(\mathbf{S}^{(t)})$ for all t along the execution trajectory of π_{greedy}. So,

$$\begin{aligned}
H_T(\pi_{\text{greedy}}) &= \mathop{\mathbb{E}}_{\pi_{\text{greedy}}} \left[\sum_{t=1}^{T-1} \gamma^{t-1} \cdot g(\mathbf{S}^{(t)}) \right] && \text{(Definition of } H_T(\pi_{\text{greedy}})\text{)} \\
&\geq \mathop{\mathbb{E}}_{\pi_{\text{greedy}}} \left[g(\mathbf{S}^{(1)}) \right] && \text{(Keep only the } t = 1 \text{ term)} \\
&\geq g(\mathbf{S}^{(1)}) && \text{(} g(\mathbf{S}^{(1)}) \text{ is a deterministic term)}
\end{aligned}$$

Meanwhile, Lemma B.4 tells us that $V_t(\mathbf{S}) \leq c_t \cdot g(\mathbf{S})$, where $c_t = \sum_{k=0}^{t-1} (\gamma \cdot (1+\rho))^k$ for any horizon $t \in \mathbb{N}_{\geq 1}$. Putting together, we see that

$$\max_{\pi} H_T(\pi) = V_T(\mathbf{S}_1) \leq c_T \cdot g(\mathbf{S}_1) \leq c_T \cdot H_T(\pi_{\text{greedy}}) \tag{9}$$

Let us now relate H_T with $\widetilde{R}_T$ using Lemma B.1. Recall that $\widetilde{R}_T(\pi) = \mathbb{E}_{\pi}\left[\sum_{t=1}^{T-1} w(\mathbf{A}^{(t)}) \cdot \Gamma_t\right]$, where $\Gamma_t = \gamma^t \cdot \frac{1-\gamma^{T-t}}{1-\gamma}$. For $\gamma \in (0,1)$, observe that $\frac{\Gamma_t}{\gamma^{t-1}}$ decreases in t, so $\gamma = \frac{\Gamma_{T-1}}{\gamma^{t-1}} \leq \frac{\Gamma_t}{\gamma^{t-1}} \leq \frac{\Gamma_1}{\gamma^{t-1}} = \gamma \cdot \frac{1-\gamma^{T-1}}{1-\gamma}$. So, we see that

$$\widetilde{R}_T(\pi) = \mathop{\mathbb{E}}_{\pi}\left[\sum_{t=1}^{T-1} w(\mathbf{A}^{(t)}) \cdot \Gamma_t\right] \geq \mathop{\mathbb{E}}_{\pi}\left[\sum_{t=1}^{T-1} w(\mathbf{A}^{(t)}) \cdot \gamma^{t-1} \cdot \frac{\Gamma_{T-1}}{\gamma^{t-1}}\right] = \gamma \cdot H_T(\pi) \tag{10}$$

$$\widetilde{R}_T(\pi) = \mathop{\mathbb{E}}_{\pi}\left[\sum_{t=1}^{T} w(\mathbf{A}^{(t)}) \cdot \Gamma_t\right] \leq \mathop{\mathbb{E}}_{\pi}\left[\sum_{t=1}^{T-1} w(\mathbf{A}^{(t)}) \cdot \gamma^{t-1} \cdot \frac{\Gamma_1}{\gamma^{t-1}}\right] = \gamma \cdot \frac{1-\gamma^{T-1}}{1-\gamma} \cdot H_T(\pi) \tag{11}$$

Let π^* be any optimal policy attaining $\max_{\pi} H_T(\pi)$. Then,

$$\begin{aligned}
\widetilde{R}_T(\pi_{\text{greedy}}) &\geq \gamma \cdot H_T(\pi_{\text{greedy}}) && \text{(From Eq. (10))}\\
&\geq \gamma \cdot \frac{1}{c_T} \cdot \max_{\pi} H_T(\pi) && \text{(From Eq. (9))}\\
&\geq \gamma \cdot \frac{1}{c_T} \cdot H_T(\pi^*) && \text{(Definition of } \pi^*\text{)}\\
&\geq \gamma \cdot \frac{1}{c_T} \cdot \frac{1}{\gamma \cdot \frac{1-\gamma^{T-1}}{1-\gamma}} \cdot \widetilde{R}_T(\pi^*) && \text{(From Eq. (11))}\\
&= \frac{1}{c_T} \cdot \frac{1-\gamma}{1-\gamma^{T-1}} \cdot \widetilde{R}_T(\pi^*)
\end{aligned}$$

Therefore, we have $\widetilde{R}_T(\pi_{\text{greedy}}) \geq \frac{1}{c_T} \cdot \frac{1-\gamma}{1-\gamma^{T-1}} \cdot \widetilde{R}_T(\pi^*) = \alpha(T, \gamma, \rho) \cdot \max_{\pi} \widetilde{R}_T(\pi)$ as claimed. □

B.2. Proof of Proposition 5.2

To prove Proposition 5.2, we rely on the following helper lemma.

Lemma B.5. *For any integer $T \geq 1$ and any $\gamma \neq 1$, we have*

$$\sum_{k=1}^{T} (k+1)\gamma^{k-1} = \frac{2 - \gamma - (T+2)\gamma^T + (T+1)\gamma^{T+1}}{(1-\gamma)^2}.$$

Proof. Fix an arbitrary $\gamma \neq 1$. Define

$$G(\gamma) := \sum_{k=0}^{T} \gamma^k = \frac{1-\gamma^{T+1}}{1-\gamma} \qquad (\gamma \neq 1).$$

Differentiating both sides with respect to γ gives

$$G'(\gamma) = \sum_{k=1}^{T} k\,\gamma^{k-1}.$$

Using the quotient rule on $\frac{1-\gamma^{T+1}}{1-\gamma}$,

$$\begin{aligned}
G'(\gamma) &= \frac{(-(T+1)\gamma^T)(1-\gamma) - (1-\gamma^{T+1})(-1)}{(1-\gamma)^2}\\
&= \frac{-(T+1)\gamma^T + (T+1)\gamma^{T+1} + 1 - \gamma^{T+1}}{(1-\gamma)^2}\\
&= \frac{1 - (T+1)\gamma^T + T\gamma^{T+1}}{(1-\gamma)^2}
\end{aligned}$$

Since $\sum_{k=1}^{T}(k+1)\gamma^{k-1} = \sum_{k=1}^{T} k \cdot \gamma^{k-1} + \sum_{k=1}^{T} \gamma^{k-1} = G'(\gamma) + \sum_{k=1}^{T} \gamma^{k-1}$, it remains to add $\sum_{k=1}^{T} \gamma^{k-1}$ to $G'(\gamma)$ above. Since $\sum_{k=1}^{T} \gamma^{k-1} = \frac{1-\gamma^T}{1-\gamma}$, we have

$$\sum_{k=1}^{T}(k+1)\gamma^{k-1} = G'(\gamma) + \sum_{k=1}^{T}\gamma^{k-1} = \frac{1-(T+1)\gamma^T + T\gamma^{T+1}}{(1-\gamma)^2} + \frac{1-\gamma^T}{1-\gamma} = \frac{2-\gamma-(T+2)\gamma^T + (T+1)\gamma^{T+1}}{(1-\gamma)^2}$$

establishing Lemma B.5. □

We now prove Proposition 5.2

Proposition 5.2 (Lower bound for π_{greedy})**.** *Fix horizon $T \in \mathbb{N}_{\geq 2}$ and discount factor $\gamma \in (0,1)$. Then, for every $\epsilon > 0$ and $\rho > 1+\epsilon$, there exist an instance of* MDR *where greedy policy π_{greedy} achieves*

$$\widetilde{R}_T(\pi_{\text{greedy}}) \leq \alpha'(T,\gamma,\rho) \cdot \max_{\pi} \widetilde{R}_T(\pi)$$

where the approximation factor is

$$\begin{aligned} &\alpha'(T,\gamma,\rho) \\ &= \alpha(T,\gamma,\rho-\varepsilon) \cdot \frac{(1-\gamma^{T-1}) \cdot (1+2\varepsilon) \cdot (2-\gamma-\gamma^T)}{(1-\gamma)^3} \end{aligned}$$

Proof. We will construct an MDR instance defined over a graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ with $n = T + 2^T$ nodes and $T+1$ labels such that the greedy policy π_{greedy} achieves

$$\widetilde{R}_T(\pi_{\text{greedy}}) \leq \frac{(1+2\varepsilon)\cdot(2-\gamma-\gamma^T)}{(1-\gamma)^2}$$

while an optimal policy π_{OPT} achieves

$$\widetilde{R}_T(\pi_{\text{OPT}}) \geq \frac{1}{\sum_{k=0}^{T-1}\left(\gamma\cdot(1+\rho-\varepsilon)\right)^k}$$

Thus, establishing our result.

Let $\varepsilon' = \frac{\varepsilon}{n(n-1)T^2\Delta}$ where $\Delta = w(\mathbf{G})$. For every edge $e \in \mathbf{E}$, let $w(e) = \varepsilon'$. Choose boosting parameters so that if an edge is boosted, its probability becomes 1. For example, under additive boosting set $\beta_\ell = 1 - \frac{1}{\varepsilon'}$ for both $\ell \in \{0,1\}$; under multiplicative boosting set $\kappa_\ell = \frac{1}{\varepsilon'}$.

Then, we construct $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ as follows:

- Initially, there is one active node n_0 with $w(n_0) = 1$.
- Let $G = \{g_1, \ldots, g_n\}$. Treat n_0 as g_0. For $i \in [T]$, $w(g_i) = 1 + \varepsilon - \varepsilon'$ and we add an edge between g_{i-1} and g_i with label $\mathbf{L}(g_{i-1}, g_i) = \{0\}$
- Let $S_1 = \{s\}$ contains 1 node with weight 1 and we add an edge between n_0 and s with label $\mathbf{L}(n_0, s) = \{1\}$
- For $i \in \{2, \ldots, T\}$, we create S_i recurssively through the following process. For $u \in \bigcup_{j\in[i-1]} S_j$, we add a node u' to S_j such that $w(u') = \rho' \cdot w(u)$. Add an edge between u and u' with label $\mathbf{L}(u, u') = \{i\}$

Justifying the parameter ρ For $u \in \mathbf{V}$, let $N(u) = \{v \in \mathbf{V} \mid (u,v) \in \mathbf{E}\}$ be the list of outgoing neighbours of u. Let, $p_{u,v^{(\ell)}}$ be the edge weight of (u,v) when messege ℓ is sent. We note that

$$\max_{\substack{u\in\mathbf{V},\\ \ell\in\mathbf{L}\cup\{\emptyset\}}} \frac{\sum_{v\in N(u)} p_{u,v^{(\ell)}} \cdot w(v)}{w(u)} \geq \max_{\substack{\mathbf{S}\subseteq\mathbf{V}, v\notin\mathbf{S}\\ \ell\in\mathbf{L}\cup\{\emptyset\}, w_v>0}} \frac{1}{w_v} \cdot \left(\mathop{\mathbb{E}}_{\mathbf{A}}[w(\mathbf{A}) \mid \mathbf{S}\cup\{v\}, \ell] - \mathop{\mathbb{E}}_{\mathbf{A}}[w(\mathbf{A}) \mid \mathbf{S}, \ell] \right)$$

We now show that for all $u \in \mathbf{V}$, $\max_{\ell\in\mathbf{L}\cup\{\emptyset\}} \frac{\sum_{v\in N(u)} p_{u,v^{(\ell)}} \cdot w(v)}{w(u)} \leq \rho$

- n_0 has 2 neighbours, 1 with weight $1+\varepsilon-\varepsilon'$ and 1 with weight ρ'. As the edge label between any 2 outgoing edges from n_0 are disjoint, and $1+\varepsilon-\varepsilon' \le \rho'$ (as $1+\varepsilon-\varepsilon' \le \rho$):

$$\max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \frac{\sum_{v \in N(n_0)} p_{n_0, v^{(\ell)}} \cdot w(v)}{w(n_0)} \le \varepsilon' + \rho' < \rho$$

- For $g_i \in G$, g_i has 1 outgoing neighbour with weight $1+\varepsilon-\varepsilon'$.

$$\max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \frac{\sum_{v \in N(g_i)} p_{g_i, v^{(\ell)}} \cdot w(v)}{w(g_i)} = 1 < \rho$$

- For all nodes $u \in \bigcup_{i \in [T]} S_i$, node u has at most T outgoing neighbours all of weight $\rho' w(u)$ and the edge label any 2 outgoing edges from e are disjoint

$$\max_{\ell \in \mathbf{L} \cup \{\emptyset\}} \frac{\sum_{v \in N(u)} p_{u, v^{(\ell)}} \cdot w(v)}{w(u)} \le (T-1)\varepsilon' + \rho' < \rho$$

as $\varepsilon' < \frac{\varepsilon}{T}$

Upper Bound on $\widetilde{R}_T(\pi_{\text{greedy}})$: Let $\mathcal{E}$ be the event when under π there is no unboosted activation and $\bar{\mathcal{E}}$ be the event when under π there is at least 1 unboosted activation. By the linearity of expeectation

$$\begin{aligned}
\widetilde{R}_T(\pi_{\text{greedy}}) &= \widetilde{R}_T(\pi_{\text{greedy}} \mid \mathcal{E}) \Pr(\mathcal{E}) + \widetilde{R}_T(\pi_{\text{greedy}} \mid \bar{\mathcal{E}}) \Pr(\bar{\mathcal{E}}) \\
&\ge \widetilde{R}_T(\pi_{\text{greedy}} \mid \mathcal{E}) \Pr(\mathcal{E}) + \widetilde{R}_T(\pi_{\text{greedy}} \mid \bar{\mathcal{E}})
\end{aligned}$$

By Lemma A.7, $\Pr(\bar{\mathcal{E}}) \le \frac{\varepsilon}{T\delta}$. Furthermore, as the maximum possible reward is less than $T\delta$, that means that

$$\widetilde{R}_T(\pi_{\text{greedy}} \mid \mathcal{E}) \Pr(\mathcal{E}) \le \varepsilon$$

To evluate $\widetilde{R}_T(\pi_{\text{greedy}} \mid \bar{\mathcal{E}})$, we note that the greedy policy would send the label 0 at everytime step. Consequently, the active nodes after timestep $t \in [T]$ is

$$\mathbf{S}^{(t)} = \{n_0\} \cup \{g_i \mid i \in [t]\}$$

Thus, the discounted reward at timestep t is

$$\gamma^{t-1} \cdot w(\mathbf{S}^{(t)}) = \gamma^{t-1} \cdot (1 + t(1+\varepsilon-\varepsilon')) \le \gamma^{t-1}(t+1)(1+\varepsilon)$$

$$\begin{aligned}
\widetilde{R}_T(\pi_{\text{greedy}} \mid \bar{\mathcal{E}}) &\le \sum_{k=1}^{T} (1+\varepsilon)(k+1)\gamma^{k-1} \\
&= (1+\varepsilon)\frac{2-\gamma-(T+2)\gamma^T + (T+1)\gamma^{T+1}}{(1-\gamma)^2} && \text{(by Lemma B.5)} \\
&= (1+\varepsilon)\frac{2-\gamma-\gamma^T-(T+1)\gamma^T + (T+1)\gamma^{T+1}}{(1-\gamma)^2} \\
&\le (1+\varepsilon)\frac{2-\gamma-\gamma^T}{(1-\gamma)^2} && \text{(as } \gamma < 1)
\end{aligned}$$

Hence, as $\frac{2-\gamma-\gamma^T}{(1-\gamma)^2} > 1$,

$$\begin{aligned}\widetilde{R}_T(\pi_{\text{greedy}} \mid \mathcal{E}) \Pr(\mathcal{E}) + \widetilde{R}_T(\pi_{\text{greedy}} \mid \bar{\mathcal{E}}) &\leq \varepsilon + (1+\varepsilon)\frac{2-\gamma-\gamma^T}{(1-\gamma)^2} \\ &\leq \frac{(1+2\varepsilon)(2-\gamma-\gamma^T)}{(1-\gamma)^2}\end{aligned}$$

Lower Bound on $\widetilde{R}_T(\pi_{\text{OPT}})$: Consider the fixed policy of sending the sequence of messages $(1, \dots, T)$. We first note that for all $t \in [T]$, $\bigcup_{i\in[t]} S_i \subseteq \mathbf{S}^{(t)}$. We now prove by induction that $w(\mathbf{S}^{(t)}) \geq (1+\rho')^{t-1}$ for all $t \in [T]$. We show this by proving that $w(\bigcup_{i\in[t]} S_i) = (1+\rho')^{t-1}$ for all $t \in [T]$

- Base Case: When $t = 1$, $w(S_1) = 1 = (1+\rho')^0$.
- Inductive Step: Assume $w(\bigcup_{i\in[t]} S_i) = (1+\rho')^{t-1}$ for all $t \in [k]$ for some $1 \leq k < T$. Then, we note that $w(S_{k+1}) = \rho' \cdot w(\bigcup_{i\in[t]} S_i) = \rho' \cdot (1+\rho')^{t-1}$. Hence,

$$w\left(\bigcup_{i\in[k+1]} S_i\right) = w\left(\bigcup_{i\in[k]} S_i\right) + w(S_{k+1}) = (1+\rho')^{k-1} + \rho' \cdot (1+\rho')^{k-1} = (1+\rho')^k$$

Thus,

$$\begin{aligned}\widetilde{R}_T(\pi_{\text{OPT}}) &= \sum_{t\in[T]} \gamma^{t-1} \cdot w(\mathbf{S}^{(t)}) \\ &\geq \sum_{t\in[T]} \gamma^{t-1} \cdot (1+\rho')^{t-1} \\ &= \sum_{t\in[T]} (\gamma \cdot (1+\rho'))^{t-1} \\ &= \sum_{t\in[T]} (\gamma \cdot (1+\rho-\varepsilon'))^{t-1} \\ &\geq \sum_{t\in[T]} (\gamma \cdot (1+\rho-\varepsilon))^{t-1}\end{aligned}$$

□

C. Experimental Details

C.1. Weighted Stochastic Block Model with Labels.

We extend the Stochastic Block Model (Holland et al., 1983) and Weighted Stochastic Block Model (Ng & Murphy, 2021) to include labels. Let $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ be a directed social network on $n := |\mathbf{V}|$ nodes and let $\mathbf{L}$ be a finite set of message types. Fix a number of clusters $Q \in \mathbb{N}$.

- **Cluster assignments and node weights.** Let $\boldsymbol{\pi} = (\pi_1, \dots, \pi_Q)$ be a probability vector, i.e., $\pi_q \geq 0$ for all $q \in [Q]$ and $\sum_{q=1}^{Q} \pi_q = 1$. Each node $v \in \mathbf{V}$ is assigned a cluster

$$z_v \sim \text{Categorical}(\boldsymbol{\pi}).$$

- **Edge existence.** For each ordered pair of distinct nodes $u, v \in \mathbf{V}$ with $u \neq v$, a directed edge from u to v is generated according to
$$A_{u,v} \sim \text{Bernoulli}\big(\rho_{z_u,z_v}\big),$$
where $\rho \in [0,1]^{Q\times Q}$ is a block connectivity matrix. If $A_{u,v} = 0$, then $(u,v) \notin \mathbf{E}$.
- **Baseline recruitment probabilities.** Conditioned on the event that $A_{u,v} = 1$, the baseline recruitment probability associated with edge (u,v) is sampled as
$$p_{u,v} \sim F^{(p)}_{z_u,z_v},$$
where $F^{(p)}_{q,r}$ is a distribution supported on $[0,1]$ that depends only on the clusters of u and v. If $A_{u,v} = 0$, we set $p_{u,v} = 0$.
- **Edge labels.** Conditioned on $A_{u,v} = 1$, the edge (u,v) is assigned a (possibly empty) set of labels
$$\mathbf{L}(u,v) \subseteq \mathbf{L},$$
drawn from a block-dependent distribution
$$\mathbf{L}(u,v) \sim F^{(L)}_{z_u,z_v}, \qquad F^{(L)}_{q,r} \text{ is a distribution over } 2^{\mathbf{L}}.$$
If $A_{u,v} = 0$, we set $\mathbf{L}(u,v) = \emptyset$.
- A convenient special case assumes independent label inclusion: for each label $\ell \in \mathbf{L}$,
$$\mathbb{P}\big(\ell \in \mathbf{L}(u,v) \mid z_u = q, z_v = r\big) = \theta_{q,r,\ell},$$
where $\theta_{q,r,\ell} \in [0,1]$. Under this parameterization, labels are drawn independently across $\ell \in \mathbf{L}$.

C.2. Graph Generation Details

Synthetic graph generation. We generate directed graphs with $n = 1000$ nodes using the WSBM with label model.

- **Clusters:** $Q = 4$ clusters with independent assignments $z_v \sim \text{Categorical}(1/4, 1/4, 1/4, 1/4)$.
- **Target degree:** $d \in \{10, 20, 30\}$. Block connectivity $\rho_{i,i} = x$, $\rho_{i,j} = x/3$ for $i \neq j$, with $x = \frac{2d}{|\mathbf{V}|}$ ensuring $\mathbb{E}[|N(u)|] \approx d$.
- **Baseline recruitment:** $p_{u,v} \sim U(0,\beta)$ with $\beta = \frac{2\rho}{d}$, yielding $\mathbb{E}[\rho'(u)] = \rho'$.
- **Edge labels:** $|\mathbf{L}| = 10$ labels, included independently with probabilities $\theta_{z_u,z_v} \sim U(0, 0.15)$.

We now show that setting $x = \frac{2d}{|\mathbf{V}|}$ ensures $\mathbb{E}[|N(u)|] \approx d$ and setting $\beta = \frac{2\rho}{d}$ ensures $\mathbb{E}[\rho'(u)] = \rho'$.

Proposition C.1. *If* $x = \frac{2d}{|\mathbf{V}|}$*, then* $\mathbb{E}[|N(u)|] \approx d$

Proof. As nodes are assigned independently and uniformly to one of the $Q = 4$ clusters, for a fixed node u, by linearity of expectation,
$$\mathbb{E}[|N(u)|] = \sum_{v \in \mathbf{V}\setminus\{u\}} \mathbb{E}[\mathbb{1}_{\mathbf{E}}(u,v)].$$

Conditioning on whether u and v belong to the same cluster,
$$\begin{aligned}\mathbb{E}[\mathbb{1}_{\mathbf{E}}(u,v)] &= \Pr(z_u = z_v)\ \mathbb{E}[\mathbb{1}_{\mathbf{E}}(u,v) \mid z_u = z_v] + \Pr(z_u \neq z_v)\ \mathbb{E}[\mathbb{1}_{\mathbf{E}}(u,v) \mid z_u \neq z_v] \\ &= \frac{1}{4}x + \frac{3}{4}\frac{x}{3} = \frac{x}{2}.\end{aligned}$$

Thus,

$$\mathbb{E}[|N(u)|] = (|\mathbf{V}| - 1)\frac{x}{2} \approx \frac{|\mathbf{V}|x}{2}.$$

Setting $x = \frac{2d}{|\mathbf{V}|}$, gives

$$\mathbb{E}[|N(u)|] \approx d$$

$\square$

Proposition C.2. *If the average degree is* d *and* $\beta = \frac{2\rho}{d}$*, then average unboosted amplification parameters will* ρ'*.*

Proof. Suppose we have a target ρ'. Recall that

$$\rho' = \frac{1}{|\mathbf{V}|}\sum_{u\in\mathbf{V}} \rho'(u).$$

Since $\mathbb{E}[\rho'(u)]$ is identical for all $u \in \mathbf{V}$, it suffices to ensure $\mathbb{E}[\rho'(u)] = \rho'$. Indeed,

$$\begin{aligned}\mathbb{E}[\rho'(u)] &= \sum_{v\in\mathbf{V}\setminus\{u\}} \mathbb{E}[\mathbb{1}_{\mathbf{E}}(u,v)]\cdot\mathbb{E}[p_{u,v}] \\ &= d\cdot\frac{\beta}{2},\end{aligned}$$

where the final equality follows from the definition of the expected degree d and the fact that $\mathbb{E}[p_{u,v}] = \beta/2$. Hence, after rearranging the terms, it suffices that $\beta = \frac{2\rho}{d}$

$\square$

C.3. Experimental Results Data

Trajectories at different values of ρ' and average degree d

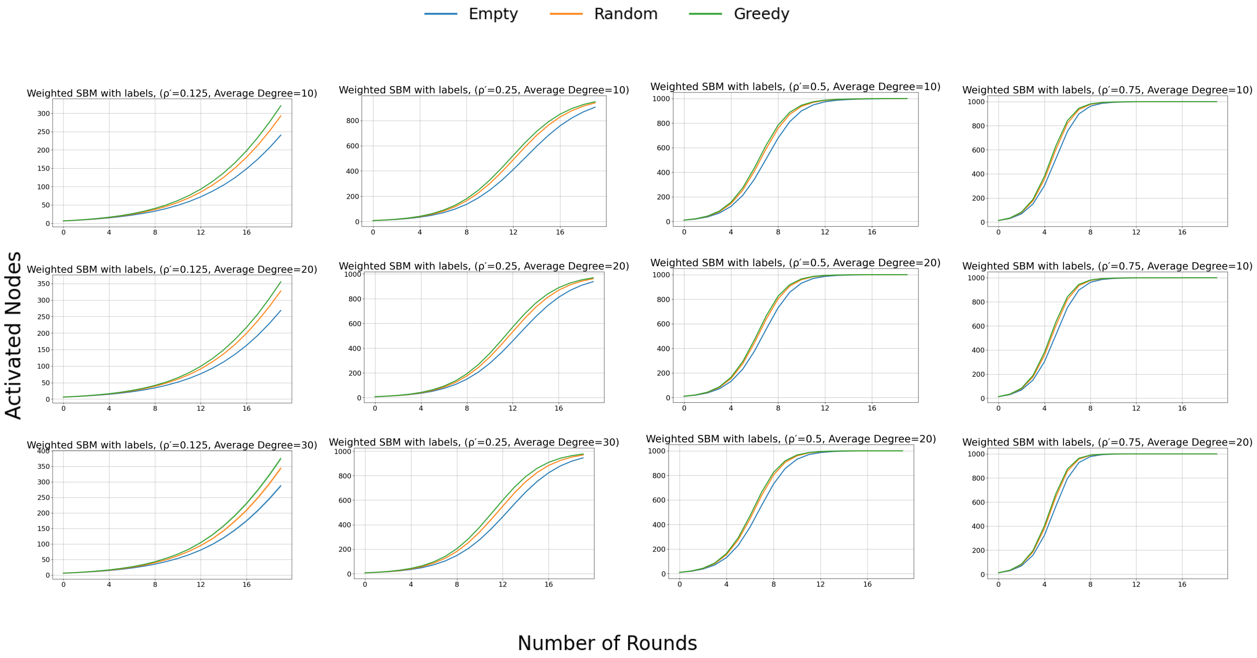

Figure 8. $d = 10$ **(top),** $d = 20$ **(middle), and** $d = 30$**(bottom). Synthetic dataset, for** $\rho' \in \{0.125, 0.25, 0.5, 0.75\}$ Expected number of activated nodes over time under three policies: EMPTY, RANDOM, and GREEDY.