

# Steering Social Activity: A Stochastic Optimal Control Point Of View\*

**Ali Zarezade**

*Sharif University of Technology  
Teheran, Iran*

ZAREZADE@CE.SHARIF.EDU

**Abir De**

*Max Planck Institute for Software Systems  
Kaiserslautern, Germany*

ADE@MPI-SWS.ORG

**Utkarsh Upadhyay**

*Max Planck Institute for Software Systems  
Kaiserslautern, Germany*

UTKARSHU@MPI-SWS.ORG

**Hamid R. Rabiee**

*Sharif University of Technology  
Teheran, Iran*

RABIEE@SHARIF.EDU

**Manuel Gomez-Rodriguez**

*Max Planck Institute for Software Systems  
Kaiserslautern, Germany*

MANUELGR@MPI-SWS.ORG

**Editor:** Nikos Vlassis

## Abstract

User engagement in online social networking depends critically on the level of *social activity* in the corresponding platform—the number of online *actions*, such as posts, shares or replies, taken by their users. Can we design data-driven algorithms to increase social activity? At a user level, such algorithms may increase activity by helping users decide when to take an action to be more likely to be *noticed* by their peers. At a network level, they may increase activity by *incentivizing* a few influential users to take more actions, which in turn will trigger additional actions by other users.

In this paper, we model social activity using the framework of marked temporal point processes, derive an alternate representation of these processes using stochastic differential equations (SDEs) with jumps and, exploiting this alternate representation, develop two efficient online algorithms with provable guarantees to steer social activity both at a user and at a network level. In doing so, we establish a previously unexplored connection between optimal control of jump SDEs and doubly stochastic marked temporal point processes, which is of independent interest. Finally, we experiment both with synthetic and real data gathered from Twitter and show that our algorithms consistently steer social activity more effectively than the state of the art.

**Keywords:** marked temporal point processes, stochastic optimal control, stochastic differential equations with jumps, social networks, information networks

---

\*Preliminary version of this work appeared in Zarezade et al. (2017).

## 1. Introduction

People perform a wide variety of online actions in a growing number of online social networking sites. In most cases, these online actions can be broadly categorized into *exogenous* actions, taken by users at their own initiative, and *endogenous* actions, taken by users as a *response* to previous actions by other users. For example, users often post small pieces of information at their own initiative, which can then trigger additional posts, shares or replies by other users. In this paper, our goal is designing online algorithms that, by steering exogenous actions, are able to increase the number of endogenous actions and thus the overall user activity.

We first address the above problem from the perspective of an individual user whose posts compete for *attention* with dozens, if not hundreds, of posts *simultaneously* shared by other users that her *followers* follow (Backstrom et al., 2011; Gomez-Rodriguez et al., 2014). In this context, recent empirical studies have shown that stories at the top of their followers’ feed are more likely to be *noticed* and consequently liked, shared or replied to (Hodas and Lerman, 2012; Kang and Lerman, 2015; Lerman and Hogg, 2014). Therefore, we design an online algorithm, REDQUEEN, to help a user decide when-to-post to increase her chances to stay at the top, increasing her *visibility*. Then, we tackle the problem from the perspective of an entire online social networking site, where an endless stream of stories posted by its users are constantly eliciting a variable number of likes, shares or replies by other users (Goel et al., 2012; Farajtabar et al., 2014; Rizoiu et al., 2017). Here, we design another online algorithm, CHESHIRE, to find how much should we incentivize a small number of influential users to post more over time to increase the overall number of additional posts, shares or replies in the site.

More specifically, we represent users’ actions using the framework of marked temporal point processes, which characterizes the continuous time interval between actions using conditional intensity functions (Aalen et al., 2008), and model endogenous and exogenous actions using multidimensional Hawkes processes (Hawkes, 1971). Then, we derive an alternate representation of these processes using stochastic differential equations (SDEs) with jumps and, exploiting this alternate representation, address the design of both algorithms from the perspective of optimal control for SDEs with jumps (Hanson, 2007). Our approach differs from previous literature in two key technical aspects, which are of independent interest:

- I. The control signals are conditional (posting) intensities, which are used to sample stochastic events (*i.e.*, stories to post). In contrast, previous work considered the control signal to be a time-varying real vector. As a consequence, our approach requires another layer of stochasticity.
- II. The (posting) intensities are stochastic Markov processes and thus the dynamics are doubly stochastic. This requires us to redefine the cost-to-go to incorporate the instantaneous value of these intensities as additional arguments. Previous work has typically considered constant intensities until very recently (Wang et al., 2016, 2017).

These technical aspects have implications beyond steering social activity since they enable us to establish a previously unexplored connection between optimal control of jump SDEs

and doubly stochastic temporal point processes (*e.g.*, Hawkes processes), which have been increasingly used to model social activity (Farajtabar et al., 2014, 2015b; Zhao et al., 2015).

In both cases, we find that the solution to the corresponding optimal control problem is surprisingly simple. At a user level, the optimal posting intensity for a user to achieve high visibility depends linearly on the position of her most recent post on each of her follower’s feed and, at a network level, the optimal level of incentivized actions depends linearly on the current level of overall actions. Moreover, at a user level, the coefficients of the linear relationship only depend on tunable parameters, which are predefined, and, at a network level, the coefficients can be found by solving a matrix Riccati differential equation (Garrett, 2013) and a first order differential equation, which has a closed form solution. As a consequence, both REDQUEEN and CHESHIRE are simple and highly efficient online procedures, which can be implemented in a few lines of code (refer to Algorithms 1 and 3). Finally, we perform experiments on both synthetic and real-data gathered from Twitter and show that our algorithms can consistently steer social activity at a user and at a network level more effectively than the state of the art.

## 1.1 Related Work

Our work relates to previous work on the when-to-post problem, the activity shaping problem, stochastic optimal control, and temporal point processes.

The “when-to-post” problem was first studied by Spasojevic et al. (2015), who performed a large empirical study on the best times for a user to post in Twitter and Facebook, measuring attention a user elicits by means of the number of responses to her posts. Moreover, they designed several heuristics to pinpoint at the times that elicit the greatest attention in a training set and showed that these times also lead to more responses in a held-out set. Since then, algorithmic approaches to the “when-to-post” problem with provable guarantees have been largely lacking. Only very recently, Karimi et al. (2016) introduced a convex optimization framework to find optimal broadcasting strategies, measuring attention a user elicits as the time that at least one of her posts is among the  $k$  most recent stories received in her followers’ feed. However, their algorithm requires expensive data pre-processing, it does not adapt to changes in the users’ feeds dynamics, and in practice, it is less effective than our algorithm REDQUEEN, as shown in Section 5.

The “activity shaping” problem was first studied by Farajtabar et al. (2015a), who derived a time dependent linear relation between the *intensity* of exogenous actions and the overall intensity of actions in a social network under a model of actions based on multidimensional Hawkes processes and, exploiting this connection, developed a convex optimization framework for activity shaping. One of the main shortcomings of their framework is that it provides deterministic exogenous intensities that do not adapt to changes in the users’ intensities and, as a consequence, it is less effective than our algorithm CHESHIRE, as shown in Section 5. More recently, Farajtabar et al. (2016) developed a heuristic method that splits the time window of interest into stages and adapts to changes in the users’ intensities at the beginning of each stage. However, their method is suboptimal, it does not have *provable* guarantees, and it achieves lower performance than our method.

In the traditional control literature (Hanson, 2007), two key aspects of our approach—intensities as control signals and stochastic intensities—have been largely understudied.

Only very recently, Wang et al. (2016) and Wang et al. (2017) have considered these aspects. However, in Wang et al. (2016), the intensities are not stochastic and the resulting algorithm needs to know the future actions, hindering its applicability, and, in Wang et al. (2017), the solution is open-loop and the control policy depends on the expectation of the uncontrolled dynamics, which needs to be calculated approximately by a time consuming sampling process. In contrast, our framework consider double stochastic intensities, our solution is closed-loop, our control policies only depend on the current state of the dynamics, and the feedback coefficients only need to be calculated once off-line.

Temporal point processes have been increasingly used for representation and modeling in a wide range of applications in social and information systems, *e.g.*, information propagation (Gomez-Rodriguez et al., 2011; Du et al., 2013; Zhao et al., 2015), opinion dynamics (De et al., 2016), product competition (Valera and Gomez-Rodriguez, 2015; Zarezade et al., 2016b), spatiotemporal social activity (Jankowiak and Gomez-Rodriguez, 2017; Zarezade et al., 2016a), information reliability (Tabibian et al., 2017), or human learning (Mavroforakis et al., 2017). However, in such context, algorithms based on stochastic optimal control of temporal point processes have been lacking.

## 2. Preliminaries

In this section, we first revisit the framework of temporal point processes (Aalen et al., 2008) and then describe how to use such framework to represent actions and feeds (Karimi et al., 2016) as well as to model endogenous and exogenous actions in social networks (Farajtabar et al., 2014).

### 2.1 Temporal Point Processes

A univariate temporal point process is a stochastic process whose realization consists of a sequence of discrete events localized in time,  $\mathcal{H} = \{t_i \in \mathbb{R}^+ \mid i \in \mathbb{N}^+, t_i < t_{i+1}\}$ . It can also be represented as a counting process  $N(t)$ , which counts the number of events before time  $t$ , *i.e.*,

$$N(t) = \sum_{t_i \in \mathcal{H}} \mathbb{I}(t - t_i \geq 0),$$

where  $\mathbb{I}(\cdot)$  is the indicator function. Then, we can characterize the counting process using the conditional intensity function  $\lambda^*(t)$ , which is the conditional probability of observing an event in an infinitesimal window  $[t, t + dt)$  given the history of event times up to time  $t$ ,  $\mathcal{H}(t) = \{t_i \in \mathcal{H} \mid t_i < t\}$ , *i.e.*,

$$\lambda^*(t) dt = \mathbb{P} \{ \text{event in } [t, t + dt) \mid \mathcal{H}(t) \} = \mathbb{E}[dN(t) \mid \mathcal{H}(t)],$$

where  $dN(t) := N(t+dt) - N(t) \in \{0, 1\}$ , the sign \* means that the intensity may depend on the history  $\mathcal{H}(t)$ . Moreover, given a function  $f(t)$ , it will be useful to define the convolution with respect to  $dN(t)$  as

$$f(t) \star dN(t) := \int_0^t f(t-s) dN(s) = \sum_{t_i \in \mathcal{H}(t)} f(t-t_i).$$

One can readily extend the above definitions to multivariate (or multidimensional) temporal point processes, which have been recently used to represent many different types of event data produced in social networks, such as the times of tweets (Farajtabar et al., 2014), retweets (Zhao et al., 2015), or links (Farajtabar et al., 2015b). More specifically, a realization of an  $m$ -dimensional temporal point process,  $\mathcal{H} = \{(t_i, u_i) \mid i \in \mathbb{N}^+, t_i \in \mathbb{R}_+, u_i \in [m], t_i < t_{i+1}\}$ , consists of  $m$  sequences of discrete events localized in time,  $\mathcal{H} = \cup_{u \in [m]} \mathcal{H}_u$ , where  $\mathcal{H}_u = \{(t_i, u_i) \in \mathcal{H} \mid u_i = u\}$ . Equivalently, it can be represented by an  $m$ -dimensional counting process  $\mathbf{N}(t)$ , where  $N_u(t)$  counts the number of events in the  $u$ -th sequence before time  $t$ , and this counting process can be characterized by  $m$  intensity functions, *i.e.*,

$$\boldsymbol{\lambda}^*(t)dt = \mathbb{E}[d\mathbf{N}(t)|\mathcal{H}(t)],$$

where  $\mathcal{H}(t) = \{(t_i, u_i) \in \mathcal{H} \mid t_i < t\}$ ,  $d\mathbf{N}(t) := (dN_u(t))_{u \in [m]} := (N_u(t+dt) - N_u(t))_{u \in [m]}$ , and  $\boldsymbol{\lambda}^*(t) := (\lambda_u^*(t))_{u \in [m]}$  denotes the associated intensities, which may depend on history  $\mathcal{H}(t)$ . Finally, given a function  $f(t)$ , one can define the convolution with respect to  $d\mathbf{N}(t)$  as

$$f(t) \star d\mathbf{N}(t) := \int_0^t f(t-s)d\mathbf{N}(s) = \left( \sum_{t_i \in \mathcal{H}_u(t)} f(t-t_i) \right)_{u \in [m]}.$$

In the remainder of the paper, to simplify the notation, we drop the sign  $*$  from the intensities.

## 2.2 Representation of Actions and Feeds

Given a (directed) social network  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with  $|\mathcal{V}| = n$  users, we assume that each user  $i$  can take a variety of online actions, *e.g.*, posting, sharing or replying, and she will be exposed to the online actions taken by the users she *follows* through her feed. Here, we assume that user  $i$  follows user  $j$  if and only if  $(i, j) \in \mathcal{E}$ .

Then, we represent the times when users take online actions as a multidimensional counting process  $\mathbf{N}(t)$ , where the  $i$ -th dimension,  $N_i(t)$ , counts the number of actions taken by user  $i$  up to time  $t$ . Here, we denote the history of times of the actions taken by user  $i$  by time  $t$  as  $\mathcal{H}_i(t)$ , the entire history of times as  $\mathcal{H}(t) = \cup_{i \in \mathcal{V}} \mathcal{H}_i(t)$ , and characterize this multidimensional process using  $n$  intensity functions, *i.e.*,  $\mathbb{E}[d\mathbf{N}(t)|\mathcal{H}(t)] = \boldsymbol{\lambda}(t) dt$ .

Given the adjacency matrix  $\mathbf{A} \in \{0, 1\}^{n \times n}$ , where  $\mathbf{A}_{ij} = 1$  indicates that user  $j$  follows user  $i$ , we can represent the times of the actions users are exposed to through their feeds as a sum of counting processes,  $\mathbf{A}^T \mathbf{N}(t)$ , and calculate the corresponding conditional intensities as  $\boldsymbol{\gamma}(t) = \mathbf{A}^T \boldsymbol{\lambda}(t)$ . Here, we denote the history of times of the actions user  $j$  is exposed to by time  $t$  as  $\mathcal{F}_j(t) := \cup_{i \in \mathcal{N}(j)} \mathcal{H}_i(t)$ , where  $\mathcal{N}(j)$  is the set of users that  $j$  follows.

Finally, from the perspective of a user  $i$ , it is useful to define the multidimensional counting process  $\mathbf{M}_{\setminus i}(t) = \mathbf{A}^T \mathbf{N}(t) - \mathbf{A}_i N_i(t)$ , in which the  $j$ -th dimension,  $M_{j \setminus i}(t)$ , counts the number of actions taken by other users that user  $j$  follows up to time  $i$ , and  $\mathbf{A}_i$  is the  $i$ -th row of the adjacency matrix  $\mathbf{A}$ . Moreover, for each dimension, the conditional intensity is given by  $\gamma_{j \setminus i}(t) = \gamma_j(t) - \lambda_i(t)$  and the history is given by  $\mathcal{F}_{j \setminus i}(t) := \mathcal{F}_j(t) \setminus \mathcal{H}_i(t)$ . When there is no ambiguity, we will omit the subscript  $i$  and write  $\mathbf{M}_{\setminus i}(t) = \mathbf{M}(t)$  and  $\gamma_{\setminus i}(t) = \boldsymbol{\gamma}(t)$ .

### 2.3 Modeling Endogenous and Exogenous Actions

Following previous work (Farajtabar et al., 2014, 2015b; Zhao et al., 2015; De et al., 2016; Rizoïu et al., 2017), we model endogenous and exogenous actions using multidimensional Hawkes processes (Hawkes, 1971). More specifically, we proceed as follows.

From the perspective of an individual user  $i$ , we assume we observe the online actions her followers are exposed to through their feeds. Then, consider the following functional form for the intensities  $\gamma_i(t) = \gamma(t)$  of the followers' feeds from Section 2.2:

$$\gamma(t) = \gamma_0(t) + \mathbf{D} \kappa_\omega(t) \star d\mathbf{M}(t) = \gamma_0(t) + \mathbf{D} \int_0^t \kappa_\omega(t-s) d\mathbf{M}(s) \quad (1)$$

where  $\gamma_0(t)$  are time-varying functions that model exogenous actions, *i.e.*, actions users take at their own initiative, the second term, with  $\mathbf{D} = \text{diag}(\alpha_j)$ ,  $\alpha_j \geq 0$ , models endogenous actions, *i.e.*, actions users take as *response* to previous actions (their own as well as the ones taken by others), and  $\kappa_\omega(t) = e^{-\omega t} \mathbb{I}(t \geq 0)$  is an exponential kernel modeling the decay on influence over time<sup>1</sup>.

From the perspective of an entire online social networking site, we assume we observe the online actions taken by all the users<sup>2</sup>. Thus, we consider the following functional form for the users' intensities  $\lambda(t)$  from Section 2.2:

$$\lambda(t) = \lambda_0(t) + \mathbf{B} \kappa_\omega(t) \star d\mathbf{N}(t) = \lambda_0(t) + \mathbf{B} \int_0^t \kappa_\omega(t-s) d\mathbf{N}(s), \quad (2)$$

where  $\lambda_0(t) \geq 0$  are time-varying functions that model exogenous actions and the second term models endogenous actions—actions users take as response to the actions taken by the users they follow. Here, we parameterize the strength of *influence* between users using a sparse nonnegative *influence matrix*  $\mathbf{B} = (\beta_{uv}) \in \mathbb{R}_+^{m \times m}$ , where  $\beta_{uv}$  means user  $u$ 's actions directly triggers *follow-ups* from user  $v$  and we assume  $\beta_{uv} = 0$  if  $(u, v) \notin \mathcal{E}$ .

In both cases, the second term makes the intensity dependent on the history and a stochastic process itself. Moreover, the following alternative representation of multidimensional Hawkes processes will be useful to design our stochastic optimal control algorithms (proven in Appendix A):

**Proposition 1** *Let  $\mathbf{N}(t)$  be an  $m$ -dimensional Hawkes process with an associated intensity  $\lambda(t) = \lambda_0(t) + \mathbf{C} \int_0^t \kappa_\omega(t-s) d\mathbf{N}(s)$ , where  $\lambda_0(t) \geq 0$  is a differentiable time-varying function,  $\kappa_\omega(t) = e^{-\omega t} \mathbb{I}(t \geq 0)$  is an exponential triggering kernel, and  $\omega \geq 0$  and  $\mathbf{C} \in \mathbb{R}_+^{m \times m}$  are given parameters. Then, the tuple  $(\mathbf{N}(t), \lambda(t))$  is a doubly stochastic Markov process, whose dynamics can be defined by the following jump SDEs:*

$$d\lambda(t) = [\lambda'_0(t) + \omega \lambda_0(t) - \omega \lambda(t)] dt + \mathbf{C} d\mathbf{N}(t), \quad (3)$$

- 
1. Exponential kernels have been shown to provide relatively good predictive performance (Farajtabar et al., 2014, 2015b; Zhao et al., 2015; De et al., 2016; Rizoïu et al., 2017), however, we acknowledge that power-law may have higher predictive performance in some scenarios, as recently shown by Mishra et al. (2016). Here, we opt for exponential kernels since, in such case, the dynamics of the corresponding intensity can be expressed by a linear SDE with jumps and this will be helpful in the derivation of our stochastic optimal control algorithms.
  2. We acknowledge that there may be scenarios in which one only has access to a subset of the online actions taken by the users. However, in those cases, one could resort to an increasing number of methods to fit multidimensional Hawkes processes from incomplete observations, *e.g.*, refer to Xu et al. (2017)

with the initial condition  $\lambda(0) = \lambda_0(0)$ .

**Remark.** From the perspective of an individual user, we assume the user has a local view, *i.e.*, she only observes the online actions her followers are exposed to through their feeds. Alternatively, if the user would have a global view, one could naturally opt for the following functional form for the intensities  $\gamma(t)$  of the followers’ feeds:

$$\gamma(t) = \mathbf{A}^T \lambda(t) = \mathbf{A}^T \lambda_0(t) + \mathbf{A}^T \mathbf{B} \int_0^t \kappa_\omega(t-s) d\mathbf{N}(s). \quad (4)$$

However, for simplicity, from the perspective of an individual user, we will assume the user has a local view and thus consider the functional form defined by Eq. 1. Considering an individual user has a global view is an interesting, albeit challenging venue for future work.

### 3. A Local Problem: When-to-post

In this section, we take the perspective of an individual user and design an online algorithm to help her decide when-to-post to be *noticed*—achieve high *visibility*. To this aim, we first define our visibility measure and then derive a jump stochastic differential equation that links our measure to the user’s posting intensity and the feeds’ intensities due to other users her followers follow. Finally, we formally state the when-to-post problem for this visibility measure and address the problem from the perspective of stochastic optimal control of SDEs with jumps. For ease of exposition, we assume users can only take one type of exogenous and endogenous action—posting stories. Moreover, throughout the section, we refer to users who posts stories as *broadcasters*.

#### 3.1 Visibility Definition and Dynamics

Given a broadcaster  $i$  and one of her followers  $j$ , we define the visibility function  $r_{ij}(t)$  as the position or *rank* of the most recent story posted by  $i$  in  $j$ ’s feed by time  $t$ , which clearly depends on the feed ranking mechanism in the corresponding social network. Here, for simplicity, we assume each user’s feed ranks stories in reverse chronological order<sup>3</sup>. However, our framework can be easily extended to any feed ranking mechanisms, as long as its rank dynamics can be expressed as a jump SDE<sup>4</sup>.

Under the reverse chronological ordering assumption, position at time  $t$  is simply the number of stories that other broadcasters posted in  $j$ ’s feed from the time of the most recent story posted by  $i$  until  $t$ . Then, when a new story arrives to a user’s feed, it appears at the top of the feed and the other stories are shifted down by one. If we denote the time of the most recent story posted by  $i$  by time  $t$  as  $\tau_i(t) = \max\{t_k \in \mathcal{H}_i(t)\}$ , then the visibility is formally defined as

$$r_{ij}(t) = M_{j \setminus i}(t) - M_{j \setminus i}(\tau_i(t)). \quad (5)$$

Note that, if the last story posted by  $i$  is at the top of  $j$ ’s feed at time  $t$ , then  $r_{ij}(t) = 0$ .

3. At the time of writing, Weibo rank stories in reverse chronological order by default and Twitter and Facebook allow choosing such an ordering.

4. This would require either having access to the corresponding feed ranking mechanism or reverse engineering it, which is out of the scope of this work.

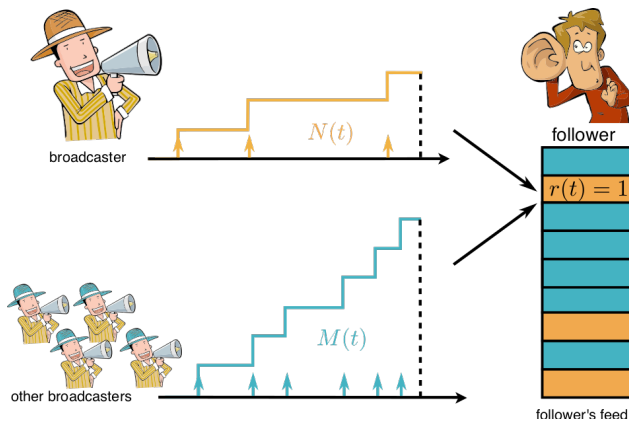


Figure 1: The dynamics of visibility. A broadcaster  $i$  posts  $N_i(t) = N(t)$  stories with intensity  $\lambda_i(t) = u(t)$ . Her stories accumulate in her follower  $j$ 's feed, competing for attention with  $M_{j \setminus i}(t) = M(t)$  other stories posted by other broadcasters user  $j$  follows with intensity  $\gamma_{j \setminus i}(t) = \gamma(t)$ . The visibility function  $r_{ij}(t) = r(t)$  is the position or rank of the most recent story posted by broadcaster  $i$  in the follower  $j$ 's feed by time  $t$ .

**Dynamics of visibility.** Given a broadcaster  $i$  with broadcasting counting process  $N_i(t)$  and one of her followers  $j$  with feed counting process due to other broadcasters  $M_{j \setminus i}(t)$ , the rank of  $i$  in  $j$ 's feed  $r_{ij}(t)$  satisfies the following equation:

$$r_{ij}(t+dt) = \underbrace{(r_{ij}(t) + 1)dM_{j \setminus i}(t)(1 - dN_i(t))}_{1. \text{ Increases by one}} + \underbrace{0}_{2. \text{ Becomes zero}} + \underbrace{r_{ij}(t)(1 - dM_{j \setminus i}(t))(1 - dN_i(t))}_{3. \text{ Remains the same}},$$

where each term models one of the three possible situations:

1. The other broadcasters post a story in  $(t, t + dt]$ ,  $dM_{j \setminus i}(t) = 1$ , and broadcaster  $i$  does not post,  $dN_i(t) = 0$ . The position of the last story posted by  $i$  in  $j$ 's feed steps down by one, *i.e.*,  $r_{ij}(t + dt) = r_{ij}(t) + 1$ .
2. Broadcaster  $i$  posts a story in  $(t, t + dt]$ ,  $dN_i(t) = 1$ , and the other broadcasters do not,  $dM_{j \setminus i}(t) = 0$ . No matter what the previous rank was, the new rank is  $r_{ij}(t + dt) = 0$  since the newly posted story appears at the top of  $j$ 's feed.
3. No one posts any story in  $(t, t + dt]$ ,  $dN_i(t) = 0$  and  $dM_{j \setminus i}(t) = 0$ . The rank remains the same, *i.e.*,  $r_{ij}(t + dt) = r_{ij}(t)$ .

We skip the case in which  $M_{j \setminus i}(t) = 1$  and  $dN_i(t) = 1$  in the same time interval  $(t, t + dt]$  because, by the Blumenthal zero-one law (Blumenthal, 1957), it has zero probability. Now, by rearranging terms and using that  $dN_i(t)dM_{j \setminus i}(t) = 0$ , we uncover the following jump SDE for the visibility (or rank) dynamics:

$$dr_{ij}(t) = -r_{ij}(t) dN_i(t) + dM_{j \setminus i}(t), \quad (6)$$

where  $dr_{ij}(t) = r_{ij}(t + dt) - r_{ij}(t)$ . Figure 1 illustrates the concept of visibility for one broadcaster and one follower.

### 3.2 Problem Formulation

Given a broadcaster  $i$  and her followers  $\mathcal{N}(i)$ , our goal is to find the optimal posting intensity  $\lambda_i(t) = u(t)$  that minimizes the expected value of a particular nondecreasing convex loss



function  $\ell(\mathbf{r}(t), u(t))$  of the broadcaster's visibility on each of her follower's feed,  $\mathbf{r}(t) = [r_{ij}(t)]_{j \in \mathcal{N}(i)}$ , and the intensity itself,  $u(t)$ , over a time window  $(t_0, t_f]$ , *i.e.*,

$$\begin{aligned} & \underset{u(t_0, t_f]}{\text{minimize}} && \mathbb{E}_{(N_i, \mathbf{M}_{\setminus i})(t_0, t_f]} \left[ \phi(\mathbf{r}(t_f)) + \int_{t_0}^{t_f} \ell(\mathbf{r}(s), u(s)) ds \right] \\ & \text{subject to} && u(t) \geq 0 \quad \forall t \in (t_0, t_f], \end{aligned} \quad (7)$$

where  $u(t_0, t_f]$  denotes the broadcaster  $i$ 's intensity from  $t_0$  to  $t_f$ , the expectation is taken over all possible realizations of the counting processes associated to the broadcaster  $i$  and all other broadcasters from  $t_0$  to  $t_f$ , denoted as  $(N_i, \mathbf{M}_{\setminus i})(t_0, t_f]$ , and  $\phi(\mathbf{r}(t_f))$  is an arbitrary penalty function<sup>5</sup>. Here, by considering a loss which is nondecreasing in the rank  $r(t)$  and the posting intensity  $u(t)$ , we penalize times when the position of the most recent story on each of the follower's feeds is high (*i.e.*, the most recent story does not stay *at the top*) and we limit the posting intensity, which in turn limits the number of stories the broadcaster can post. Finally, note that the optimal intensity  $u(t)$  for broadcaster  $i$  at time  $t$  may depend on the visibility  $\mathbf{r}(t)$  with respect to each of her followers and, thus, the associated counting process  $N_i(t)$  may be doubly stochastic.

### 3.3 Stochastic Optimal Control Algorithm

In this section, we tackle the when-to-post problem defined by Eq. 7 from the perspective of stochastic optimal control of jump SDEs (Hanson, 2007). More specifically, we first derive a solution to the problem considering only one follower, provide an efficient practical implementation of the solution and then generalize it to the case of multiple followers. We conclude this section by deriving a solution to the problem given an (idealized) oracle that knows the times of all stories in the followers' feeds *a priori*, which we will use as baseline in our experiments (refer to Section 5.1).

**Optimizing for one follower.** Given a broadcaster  $i$  with  $N_i(t) = N(t)$  and  $\lambda_i(t) = u(t)$  and only one of her followers  $j$  with  $M_{j \setminus i}(t) = M(t)$  and  $\gamma_{j \setminus i}(t) = \gamma(t)$ , we can rewrite the when-to-post problem defined by Eq. 7 as

$$\begin{aligned} & \underset{u(t_0, t_f]}{\text{minimize}} && \mathbb{E}_{(N, M)(t_0, t_f]} \left[ \phi(r(t_f)) + \int_{t_0}^{t_f} \ell(r(s), u(s)) ds \right] \\ & \text{subject to} && u(t) \geq 0 \quad \forall t \in (t_0, t_f], \end{aligned} \quad (8)$$

where, using Eqs. 1, 3 and 6, the dynamics of  $M(t)$  and  $r(t)$  are given by the following two coupled jump SDEs:

$$\begin{aligned} d\gamma(t) &= [\gamma'_0(t) + \omega\gamma_0(t) - \omega\gamma(t)] dt + \alpha dM(t), \\ dr(t) &= -r(t) dN(t) + dM(t), \end{aligned} \quad (9)$$

with initial conditions  $r(t_0) = r_0$  and  $\gamma(t_0) = \gamma_0$ , and the dynamics of  $N(t)$  are given by the intensity  $u(t)$  that we aim to optimize.

---

5. The final penalty function  $\phi(\mathbf{r}(t_f))$  is necessary to derive the optimal intensity  $u^*(t)$  in Section 3.3. However, the actual optimal intensity  $u^*(t)$  will not depend on the particular choice of terminal condition.

Next, we will define a novel optimal cost-to-go function that accounts for the above unique aspects of our problem, showing that the Bellman's principle of optimality still follows, and finally find the optimal solution using the corresponding Hamilton-Jacobi-Bellman (HJB) equation.

**Definition 2** *The optimal cost-to-go  $J(r(t), \gamma(t), t)$  is defined as the minimum of the expected value of the cost of going from state  $r(t)$  with intensity  $\gamma(t)$  at time  $t$  to a final state at time  $t_f$ , i.e.,*

$$J(r(t), \gamma(t), t) = \min_{u(t, t_f]} \mathbb{E}_{(N, M)(t, t_f]} \left[ \phi(r(t_f)) + \int_t^{t_f} \ell(r(s), u(s)) ds \right], \quad (10)$$

where the expectation is taken over all trajectories of the control and noise jump process,  $N$  and  $M$ , in the  $(t, t_f]$  interval, given the initial values of  $r(t)$ ,  $\gamma(t)$  and  $u(t)$ .

To find the optimal control  $u(t, t_f]$  and cost-to-go  $J$ , we break the problem into smaller subproblems, using the Bellman's principle of optimality (Bertsekas, 1995), which the above definition allows (proven in Appendix B):

**Lemma 3 (Bellman's Principle of Optimality)** *The optimal cost satisfies the following recursive equation:*

$$J(r(t), \gamma(t), t) = \min_{u(t, t+dt]} \{ \mathbb{E} [J(r(t+dt), \gamma(t+dt), t+dt)] + \ell(r(t), u(t)) dt \}, \quad (11)$$

where the expectation is taken over all trajectories of the control and noise jump processes,  $N$  and  $M$ , in  $(t, t+dt]$ . Then, we use the Bellman's principle of optimality to derive a partial differential equation on  $J$ , often called the Hamilton-Jacobi-Bellman (HJB) equation (Hanson, 2007). To do so, we first assume  $J$  is continuous and then rewrite Eq. 11 as

$$\begin{aligned} J(r(t), \gamma(t), t) &= \min_{u(t, t+dt]} \{ \mathbb{E} [J(r(t), \gamma(t), t) + dJ(r(t), \gamma(t), t)] + \ell(r(t), u(t)) dt \} \\ 0 &= \min_{u(t, t+dt]} \{ \mathbb{E} [dJ(r(t), \gamma(t), t)] + \ell(r(t), u(t)) dt \}, \end{aligned} \quad (12)$$

where  $dJ(r(t), \gamma(t), t) = J(r(t+dt), \gamma(t+dt), t+dt) - J(r(t), \gamma(t), t)$ . Then, we differentiate  $J$  with respect to time  $t$ ,  $r(t)$  and  $\gamma(t)$  using to following Lemma (proven in the Appendix C):

**Lemma 4** *The differential  $dJ(r(t), \gamma(t), t)$  of the cost-to-go function  $J(r(t), \gamma(t), t)$ , as defined by Eq. 10, is given by*

$$\begin{aligned} dJ(r(t), \gamma(t), t) &= J_t(r(t), \gamma(t), t)dt + [\gamma'_0(t) + \omega\gamma_0(t) - \omega\gamma(t)] J_\gamma(r(t), \gamma(t), t)dt \\ &\quad + [J(0, \gamma(t), t) - J(r(t), \gamma(t), t)]dN(t) \\ &\quad + [J(r(t) + 1, \gamma(t) + \alpha, t) - J(r(t), \gamma(t), t)]dM(t), \end{aligned}$$

where  $J_t$  and  $J_\gamma$  are derivatives of  $J$  with respect to  $t$  and  $\gamma$ , respectively.

---

**Algorithm 1:** REDQUEEN for fixed  $s$ ,  $q$  and one follower.
 

---

```

1: Input: parameters  $q$  and  $s$  ;
2: Output: Returns time for the next post ;
3:  $t \leftarrow \infty$ ;  $\tau \leftarrow \text{othersNextPost}()$  ;
4: while  $\tau < t$  do
5:    $\Delta \sim \text{Sample}(\sqrt{s/q})$ ;
6:    $t \leftarrow \min(t, \tau + \Delta)$ ;
7:    $\tau \leftarrow \text{othersNextPost}()$ ;
8: return  $t$ ;
    
```

---

Next, if we plug in the above equation in Eq. 12, it follows that

$$\begin{aligned}
 0 = \min_{u(t,t+dt)} & \left\{ J_t(r(t), \gamma(t), t)dt + [\gamma'_0(t) + \omega\gamma_0(t) - \omega\gamma(t)] J_\gamma(r(t), \gamma(t), t)dt \right. \\
 & + [J(0, \gamma(t), t) - J(r(t), \gamma(t), t)]\mathbb{E}[dN(t)] \\
 & \left. + [J(r(t) + 1, \gamma(t) + \alpha, t) - J(r(t), \gamma(t), t)]\mathbb{E}[dM(t)] + \ell(r(t), u(t)) dt \right\}.
 \end{aligned}$$

Now, using  $\mathbb{E}[dN(t)] = u(t)dt$  and  $\mathbb{E}[dM(t)] = \gamma(t)dt$ , and rearranging terms, the HJB equation follows:

$$\begin{aligned}
 0 = & J_t(r(t), \gamma(t), t) + [\gamma'_0(t) + \omega\gamma_0(t) - \omega\gamma(t)] J_\gamma(r(t), \gamma(t), t) \\
 & + [J(r(t) + 1, \gamma(t) + \alpha, t) - J(r(t), \gamma(t), t)]\gamma(t) \\
 & + \min_{u(t,t+dt)} \ell(r(t), u(t)) + [J(0, \gamma(t), t) - J(r(t), \gamma(t), t)]u(t). \tag{13}
 \end{aligned}$$

To be able to continue further, we need to define the loss  $\ell$  and the penalty  $\phi$ . Following the literature on stochastic optimal control (Hanson, 2007), we consider the following quadratic forms, which will turn out to be a tractable choice<sup>6</sup>:

$$\phi(r(t_f)) = \frac{1}{2}r^2(t_f) \quad \text{and} \quad \ell(r(t), u(t)) = \frac{1}{2}s(t)r^2(t) + \frac{1}{2}qu^2(t),$$

where  $s(t)$  is a time significance function  $s(t) \geq 0$ , which favors some periods of times (*e.g.*, times in which the follower is online<sup>7</sup>), and  $q$  is a given parameter, which trade-offs visibility and number of broadcasted posts.

Under these definitions, we take the derivative with respect to  $u(t)$  of Eq. 13 and uncover the relationship between the optimal intensity and the optimal cost:

$$u^*(t) = q^{-1} [J(r(t), \gamma(t), t) - J(0, \gamma(t), t)]. \tag{14}$$

Finally, we substitute the above expression in Eq. 13 and find that the optimal cost  $J$  needs to satisfy the following nonlinear differential equation:

$$\begin{aligned}
 0 = & J_t(r(t), \gamma(t), t) + \frac{1}{2}s(t)r^2(t) - \frac{1}{2}q^{-1} [J(r(t), \gamma(t), t) - J(0, \gamma(t), t)]^2 \\
 & + [\gamma'_0(t) + \omega\gamma_0(t) - \omega\gamma(t)] J_\gamma(r(t), \gamma(t), t) + [J(r(t) + 1, \gamma(t) + \alpha, t) - J(r(t), \gamma(t), t)]\gamma(t)
 \end{aligned} \tag{15}$$

---

6. Considering other losses with a specific semantic meaning (*e.g.*,  $\mathbb{I}(r(t) \leq k)$ ) is a challenging direction for future work.

7. Such information may be hidden but one can use the followers' posting activity or geographic location as a proxy (Karimi et al., 2016).

with  $J(r(t_f), \gamma(t_f), t_f) = \phi(r(t_f))$  as the terminal condition. The following technical lemma provides us with a solution to the above equation (proven in Appendix D):

**Lemma 5** *In the space of  $m$ -degree polynomials, the following polynomial is the only solution to Eq. 15:*

$$J(r(t), \gamma(t), t) = f(t) + \sqrt{s(t)q}r(t) + \sum_{j=1}^m g_j(t)\gamma^j(t),$$

where  $f(t)$  and  $g_j(t)$  are time-varying functions which can be found by solving a linear system of differential equations.

Given the above Lemma and Eq. 14, the optimal intensity is readily given by following theorem:

**Theorem 6** *The optimal intensity for the when-to-post problem defined by Eq. 8 with quadratic loss and penalty function is given by  $u^*(t) = \sqrt{s(t)/q}r(t)$ .*

The optimal intensity only depends on the position of the most recent post by broadcaster  $i$  in her follower's feed and, thus, allows for a very efficient procedure to sample posting times, which exploits the superposition theorem (Kingman, 1992). The key idea is as follows: at any given time  $t$ , we can view the process defined by the optimal intensity as a superposition of  $r(t)$  inhomogeneous Poisson processes with intensity  $\sqrt{s(t)/q}r(t)$  which starts at jumps of the rank  $r(t)$ , and find the next sample by computing the minimum across all samples from these processes. Algorithm 1 summarizes our (sampling) method, which we name REDQUEEN (Carroll, 1917). Within the algorithm,  $Sample(\sqrt{s/q})$  samples from a Poisson process with intensity  $\sqrt{s/q}$  and  $othersNextPost()$  returns the time of the next event by other broadcasters in the followers' feeds, once the event happens. In practice, we only need to know if the event happens before we post. Remarkably, it only needs to sample  $M(t_f)$  times from a (in)homogeneous Poisson process (if significance is time varying) and requires  $O(1)$  space.

**Optimizing for multiple followers.** Given a broadcaster  $i$  with  $N_i(t) = N(t)$  and  $\lambda_i(t) = u(t)$  and her followers  $\mathcal{N}(i)$  with  $M_{\setminus i}(t) = M(t)$  and  $\gamma_{\setminus i}(t) = \gamma(t)$ , we can write the dynamics of  $M(t)$  and  $r(t)$ , which we need to solve Eq. 7, using Eqs. 1, 3 and 6:

$$\begin{aligned} d\gamma(t) &= [\gamma'_0(t) + w\gamma_0(t) - w\gamma(t)] dt + \mathbf{D} dM(t) \\ d\mathbf{r}(t) &= -\mathbf{r}(t) dN(t) + dM(t). \end{aligned}$$

Then, consider the following quadratic forms for the loss  $\ell$  and the penalty  $\phi$ :

$$\begin{aligned} \phi(\mathbf{r}(t_f)) &= \sum_{i=1}^n \frac{1}{2} r_i^2(t_f) \\ \ell(\mathbf{r}(t), u(t), t) &= \sum_{i=1}^n \frac{1}{2} s_i(t) r_i^2(t) + \frac{1}{2} q u^2(t). \end{aligned}$$

---

**Algorithm 2:** Optimal posting times with an oracle.
 

---

```

1: Input: Initial state  $r_0$ , interval widths  $w_1, \dots, w_{m+1}$ , parameter  $q$  and significance  $s(t) = s$ ;
2: Output: Overall cost  $J(r_0, 0)$ , optimal control  $u_0^*, \dots, u_m^*$ ;
3: for  $r \leftarrow r_0 + m$  to 0 do
4:    $J(r, m+1) \leftarrow \frac{1}{2}r^2$ ;
5: for  $k \leftarrow m$  to 0 do
6:   for  $r \leftarrow r_0 + k - 1$  to 0 do
7:      $J(r, k) = \min\{\frac{1}{2}q + J(0, k+1), \frac{1}{2}sw_{k+1}(r+1)^2 + J(r+1, k+1)\}$ ;
8: for  $k \leftarrow 0$  to  $m$  do
9:   if  $\frac{1}{2}q + J(0, k+1) < \frac{1}{2}sw_{k+1}(r_k+1)^2 + J(r_k+1, k+1)$  then
10:     $u_k^* \leftarrow 1$ ;  $r_{k+1} \leftarrow 0$ ;
11:   else
12:     $u_k^* \leftarrow 0$ ;  $r_{k+1} \leftarrow r_k + 1$ ;
return  $J(r_0, 0)$ ,  $u_0^*, \dots, u_m^*$ 
    
```

---

where  $s_i(t)$  is the time significance function for follower  $i$ , as defined above, and  $q$  is a given parameter. Proceeding similarly as in the case of one follower, we can show that:

$$u^*(t) = \sum_{i=1}^n \sqrt{s_i(t)/q} r_i(t), \quad (16)$$

which only depends on the position of the most recent post by broadcaster  $i$  in her followers' feeds. Finally, we can readily adapt REDQUEEN (Algorithm 1) to efficiently sample the posting times using the above intensity—it only needs to sample  $|\cup_{j \in \mathcal{N}(i)} \mathcal{F}_{j \setminus i}(t_f)|$  values from an (in)homogeneous Poisson process (if significance is time varying) and requires  $O(|\mathcal{N}(i)|)$  space.

**Optimizing with an oracle.** In this section, we consider a broadcaster  $i$  with  $N_i(t) = N(t)$  and  $\lambda_i(t) = u(t)$ , only one of her followers  $j$  with  $M_{j \setminus i}(t) = M(t)$ , and a constant significance  $s(t) = s$ . The derivation can be easily adapted to the case of multiple followers and time-varying significance.

Suppose there is an (idealized) oracle that reveals  $M(t)$  from  $t_0$  to  $t_f$ , *i.e.*, the history  $\mathcal{F}_{j \setminus i}(t_f) = \mathcal{F}(t_f)$  is given, and  $M(t_f) = |\mathcal{F}(t_f)| = m$ . Then, we can rewrite Eq. 7 as

$$\begin{aligned} & \underset{u(t_0, t_f)}{\text{minimize}} && \mathbb{E}_{N(t_0, t_f)} \left[ \phi(r(t_f)) + \int_{t_0}^{t_f} \ell(r(s), u(s)) ds \right] \\ & \text{subject to} && u(t) \geq 0 \quad \forall t \in (t_0, t_f], \end{aligned}$$

where the expectation is only taken over all possible realizations of the counting process  $N(t_0, t_f]$  since  $M(t_0, t_f]$  is revealed by the oracle and, thus, deterministic.

Similar to the previous sections, assume the loss  $\ell$  and penalty  $\phi$  are quadratic. The best times for user  $i$  to post will always coincide with one of the times in  $\mathcal{F}(t_f)$  since, given a posting time  $\tau_i \in (t_k, t_{k+1})$ , where  $t_k, t_{k+1} \in \mathcal{F}(t_f)$ , one can reduce the cost by  $(1/2)q(\tau_i - t_k)r^2(t_k)$  by choosing instead to post at  $t_k$ . As a consequence, we can discretize the dynamics of  $r(t)$  in times  $\mathcal{F}(t_f)$ , and write  $r_{k+1} = r_k + 1 - (r_k + 1)u_k$ , where  $r_k = r(t_k^-)$ ,

$u_k = u(t_k^+) \in \{0, 1\}$ ,  $t_k \in \mathcal{F}(t_f)$ . We can easily see that  $r_k$  is bounded by  $0 \leq r_k < r_0 + m$ . Similarly, we can derive the optimal cost-to-go in discrete-time as

$$J(r_k, k) = \min_{u_k, \dots, u_m} \frac{1}{2} r_{m+1}^2 + \sum_{i=k}^m \frac{1}{2} q w_{i+1} r_{i+1}^2 + \frac{1}{2} s u_i^2,$$

where  $w_i = t_i - t_{i-1}$ . Next, we can break the minimization and use Bellman's principle of optimality,

$$J(r_k, k) = \min_{u_k} \frac{1}{2} q w_{k+1} r_{k+1}^2 + \frac{1}{2} s u_k^2 + J(r_{k+1}, k+1),$$

and, since  $u_k \in \{0, 1\}$ , the above recursive equation can be written as

$$J(r_k, k) = \min \left\{ \frac{1}{2} s + J(0, k+1), \frac{1}{2} q w_{k+1} (r_k + 1)^2 + J(r_k + 1, k+1) \right\}.$$

Finally, we can find the optimal control  $u_k^*$ ,  $k = 0, \dots, m$  and cost  $J(r_0, 0)$  by backtracking from the terminal condition  $J(r_{m+1}, m+1) = r_{m+1}^2/2$  to the initial state  $r_0$ , as summarized in Algorithm 2, which can be adapted to multiple followers. Note that, in this case, the optimal strategy is not stochastic and consists of a set of optimal posting times, as one could have guessed. However, for multiple followers, the complexity of the algorithm is  $O(m^2)$ , where  $m = |\cup_{j \in \mathcal{N}(i)} \mathcal{F}_{j \setminus i}(t_f)|$ .

## 4. A Global Problem: Activity Maximization

In this section, we take the perspective of an entire online social networking site and design an online algorithm to find how much should we incentivize a small number of influential users to post more over time to increase the overall number of additional posts, shares or replies in the site. To this aim, we first describe how to formally model such incentive mechanism in social networks. Then, we state the activity shaping problem and address the problem from the perspective of stochastic optimal control of SDEs with jumps, similarly as in Section 3.

### 4.1 Triggering Additional Endogenous Actions

Given a social network  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with  $|\mathcal{V}| = n$  users, we trigger additional endogenous user actions by directly incentivizing (*e.g.*, paying) for  $\mathbf{P}(t)$  actions, where  $P_i(t)$  counts the number of directly incentivized actions taken by user  $i \in \mathcal{V}$  before time  $t$  and we can characterize  $\mathbf{P}(t)$  by means of  $n$  intensity functions  $\mathbf{u}(t)$ , *i.e.*,  $\mathbb{E}[d\mathbf{P}(t)] = \mathbf{u}(t)dt$ . Then, if we assume the strength of influence  $\mathbf{B}$  between users is the same both for organic and incentivized actions<sup>8</sup>, as previous work (Farajtabar et al., 2014, 2016), we can rewrite the

---

8. This assumption seems reasonable in some scenarios, *e.g.*, it is often difficult to notice whether an influential user is being paid for posting a message. However, one could relax this assumption by considering a different influence matrix  $C \neq B$  for the additional endogenous actions in Eq. 17, change the SDE, Eq. 18 and HJB, Eq. 22 accordingly, and derived the new optimal control in Theorem 11.

users' intensities  $\boldsymbol{\lambda}(t)$ , given by Eq. 2, as

$$\boldsymbol{\lambda}(t) = \boldsymbol{\lambda}_0(t) + \mathbf{B} \int_0^t \kappa_\omega(t-s) d\mathbf{N}(s) + \underbrace{\mathbf{B} \int_0^t \kappa_\omega(t-s) d\mathbf{P}(s)}_{\text{Additional endogenous actions}} . \quad (17)$$

Here, note that treating directly incentivized actions as a different counting process  $\mathbf{P}(t)$  prevents the intensities  $\boldsymbol{\lambda}(t)$  from including their intensity but instead including only the intensity of their follow-ups (*i.e.*, additional endogenous actions), which we aim to maximize. Then, it is easy to derive the following alternative representation, similarly as in Proposition 1, which we will use in our stochastic optimal control algorithm:

**Proposition 7** *Let  $\mathbf{N}(t)$  and  $\mathbf{P}(t)$  be two multidimensional counting processes with associated intensities  $\boldsymbol{\lambda}(t)$ , given by Eq. 17, and  $\mathbf{u}(t)$ , respectively. Then, the dynamics of the intensity  $\boldsymbol{\lambda}(t)$  can be expressed using the following jump SDEs:*

$$d\boldsymbol{\lambda}(t) = [\boldsymbol{\lambda}'_0(t) + \omega\boldsymbol{\lambda}_0(t) - \omega\boldsymbol{\lambda}(t)] dt + \mathbf{B} d\mathbf{N}(t) + \mathbf{B} d\mathbf{P}(t) \quad (18)$$

with the initial condition  $\boldsymbol{\lambda}(0) = \boldsymbol{\lambda}_0(0)$ .

In the remainder of the section, for ease of exposition, we assume  $\boldsymbol{\lambda}_0(t) = \boldsymbol{\lambda}_0$  and thus  $\boldsymbol{\lambda}'_0(t) = 0$ .

## 4.2 Problem Formulation

Given a social network  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with  $|\mathcal{V}| = n$  users, our goal is to find the optimal users' intensities for directly incentivized actions  $\mathbf{u}(t)$  (for short, control intensities) that minimize the expected value of a particular loss function  $\ell(\mathbf{u}(t), \boldsymbol{\lambda}(t))$  of the control intensities and the users' intensities for not directly incentivized actions over a time window  $(t_0, t_f]$ , *i.e.*,

$$\begin{aligned} & \underset{\mathbf{u}(t_0, t_f]}{\text{minimize}} && \mathbb{E}_{(\mathbf{N}, \mathbf{P})(t_0, t_f]} \left[ \phi(\boldsymbol{\lambda}(t_f)) + \int_{t_0}^{t_f} \ell(\boldsymbol{\lambda}(s), \mathbf{u}(s)) ds \right] \\ & \text{subject to} && u_i(t) \geq 0, \forall t \in (t_0, t_f], i = 1, \dots, n \end{aligned} \quad (19)$$

where  $\mathbf{u}(t_0, t_f]$  denotes the control intensities from  $t_0$  to  $t_f$ , the dynamics of  $\mathbf{N}(t)$  are given by Eq. 18, and the expectation is taken over all possible realizations of the two counting processes  $\mathbf{N}(t)$  and  $\mathbf{P}(t)$  during interval  $(t_0, t_f]$ . Here, by considering a loss that is nonincreasing (nondecreasing) with respect to the intensities  $\boldsymbol{\lambda}(t)$  ( $\mathbf{u}(t)$ ), we will trade-off number of directly incentivized actions and number of additional endogenous actions. Finally, note that the optimal intensities  $\mathbf{u}(t)$  at time  $t$  may depend on the intensities  $\boldsymbol{\lambda}(t)$  and thus the associated counting process  $\mathbf{P}(t)$  may be doubly stochastic.

## 4.3 Stochastic Optimal Control Algorithm

In this section, we proceed similarly as in Section 3.3 and tackle the activity maximization problem defined by Eq. 19 from the perspective of stochastic optimal control of SDEs with jumps (Hanson, 2007). More specifically, we first define a novel optimal cost-to-go function specially designed for this global problem and then derive and solve the corresponding HJB equation to find the optimal control intensities.

**Definition 8** The optimal cost-to-go  $J(\boldsymbol{\lambda}(t), t)$  is defined as the minimum of the expected value of the cost of going from the state with intensity  $\boldsymbol{\lambda}(t)$  at time  $t$  to the final state at time  $t_f$ , i.e.,

$$J(\boldsymbol{\lambda}(t), t) = \min_{\mathbf{u}(t, t_f)} \mathbb{E}_{(\mathbf{N}, \mathbf{P})(t, t_f)} \left[ \phi(\boldsymbol{\lambda}(t_f)) + \int_t^{t_f} \ell(\boldsymbol{\lambda}(s), \mathbf{u}(s)) ds \right], \quad (20)$$

where the expectation is taken over all trajectories of the counting processes  $\mathbf{N}$  and  $\mathbf{P}$  in the  $(t, t_f]$  interval, given the initial values of  $\boldsymbol{\lambda}(t)$  and  $\mathbf{u}(t)$ .

Next, we use the Bellman's principle of optimality (Bertsekas, 1995), which the above cost-to-go  $J(\boldsymbol{\lambda}(t), t)$  also allows for (proof is similar to that of Lemma 3), to break the problem into smaller subproblems and rewrite Eq. 20 as

$$\begin{aligned} J(\boldsymbol{\lambda}(t), t) &= \min_{\mathbf{u}(t, t+dt)} \left\{ \mathbb{E} [J(\boldsymbol{\lambda}(t), t) + dJ(\boldsymbol{\lambda}(t), t)] + \ell(\mathbf{u}(t), \boldsymbol{\lambda}(t)) dt \right\} \\ 0 &= \min_{\mathbf{u}(t, t+dt)} \left\{ \mathbb{E} [dJ(\boldsymbol{\lambda}(t), t)] + \ell(\mathbf{u}(t), \boldsymbol{\lambda}(t)) dt \right\}, \end{aligned} \quad (21)$$

where  $dJ(\boldsymbol{\lambda}(t), t) = J(\boldsymbol{\lambda}(t+dt), t+dt) - J(\boldsymbol{\lambda}(t), t)$ . Then, we explicitly differentiate  $J$  with respect to time  $t$  and  $\boldsymbol{\lambda}(t)$  using the following Lemma (proven in the Appendix E):

**Lemma 9** The differential  $dJ(\boldsymbol{\lambda}(t), t)$  of the cost-to-go function  $J(\boldsymbol{\lambda}(t), t)$ , as defined by Eq. 20, is given by

$$\begin{aligned} J_t(\boldsymbol{\lambda}(t), t) dt + [w\boldsymbol{\lambda}_0 - w\boldsymbol{\lambda}(t)]^T \nabla_{\boldsymbol{\lambda}} J(\boldsymbol{\lambda}(t), t) dt + \sum_i [J(\boldsymbol{\lambda}(t) + \mathbf{b}_i, t) - J(\boldsymbol{\lambda}(t), t)] dN_i(t) \\ + \sum_i [J(\boldsymbol{\lambda}(t) + \mathbf{b}_i, t) - J(\boldsymbol{\lambda}(t), t)] dP_i(t), \end{aligned}$$

where  $\mathbf{b}_i$  is the  $i$ 'th column of  $\mathbf{B}$ ,  $J_t$  is derivative of  $J$  with respect to  $t$  and  $\nabla_{\boldsymbol{\lambda}} J$  is the gradient of  $J$  with respect to  $\boldsymbol{\lambda}(t)$ .

Next, if we plug the above equation in Eq. 21 and use that  $\mathbb{E}[N_i(t)] = \lambda_i(t) dt$  and  $\mathbb{E}[P_i(t)] = u_i(t) dt$ , the HJB equation follows:

$$0 = J_t(\boldsymbol{\lambda}(t), t) + [w\boldsymbol{\lambda}_0 - w\boldsymbol{\lambda}(t)]^T \nabla_{\boldsymbol{\lambda}} J(\boldsymbol{\lambda}(t), t) + \boldsymbol{\lambda}^T(t) \Delta_B J + \min_{\mathbf{u}(t)} \ell(\mathbf{u}(t), \boldsymbol{\lambda}(t)) + \mathbf{u}^T(t) \Delta_B J, \quad (22)$$

where  $\Delta_B J$  denotes a vector whose  $i$ 'th element is given by  $(\Delta_B J)_i = J(\boldsymbol{\lambda}(t) + \mathbf{b}_i, t) - J(\boldsymbol{\lambda}(t), t)$ .

To solve the above equation, we need to define the loss and penalty functions,  $\ell$  and  $\phi$ . Similarly as in Section 3, we consider the following quadratic forms, which will turn out to be a tractable choice:

$$\ell(\mathbf{u}(t), \boldsymbol{\lambda}(t)) = -\frac{1}{2} \boldsymbol{\lambda}^T(t) \mathbf{Q} \boldsymbol{\lambda}(t) + \frac{1}{2} \mathbf{u}^T(t) \mathbf{S} \mathbf{u}(t) \quad \text{and} \quad \phi(\boldsymbol{\lambda}(t_f)) = -\frac{1}{2} \boldsymbol{\lambda}^T(t_f) \mathbf{F} \boldsymbol{\lambda}(t_f),$$

where  $\mathbf{Q}$ ,  $\mathbf{F}$  and  $\mathbf{S}$  are given symmetric matrices<sup>9</sup> with  $q_{ij} \geq 0$ ,  $f_{ij} \geq 0$  and  $s_{ij} \geq 0$  for all  $i, j \in [n]$ . These matrices allow us to trade-off the number of directly incentivized actions

9. In practice, we will consider diagonal matrices.



---

**Algorithm 3:** CHESHIRE: it returns user  $i$  and time  $\tau$  for the next incentivized action

---

```

1: Initialization:
2: Compute  $\mathbf{H}(t)$  and  $\mathbf{g}(t)$  ;
3:  $\mathbf{u}(t) \leftarrow -\mathbf{S}^{-1}[\mathbf{B}^T(\mathbf{g}(t) + \mathbf{H}(t)\boldsymbol{\lambda}_0) + \frac{1}{2} \text{diag}(\mathbf{B}^T \mathbf{H}(t) \mathbf{B})]$  ;
4: General subroutine:
5:  $(i, \tau) \leftarrow \text{Sample}(\mathbf{u}(t))$  ;
6:  $(j, s) \leftarrow \text{NextAction}()$  ;
7: while  $s < \tau$  do
8:    $\boldsymbol{\lambda}_N(t) \leftarrow \mathbf{B}\mathbf{e}_j\kappa_\omega(t-s)$  ;
9:    $\mathbf{u}_N(t) \leftarrow -\mathbf{S}^{-1}\mathbf{B}^T \mathbf{H}(t)\boldsymbol{\lambda}_N(t)$  ;
10:   $(k, r) \leftarrow \text{Sample}(\mathbf{u}_N(t))$  ;
11:  if  $r < \tau$  then
12:     $\tau \leftarrow r$  ;
13:     $i \leftarrow k$  ;
14:   $\mathbf{u}(t) \leftarrow \mathbf{u}(t) + \mathbf{u}_N(t)$  ;
15:   $(j, s) \leftarrow \text{NextAction}()$  ;
16:   $\boldsymbol{\lambda}_P(t) \leftarrow \mathbf{B}\mathbf{e}_i\kappa_\omega(t-\tau)$  ;
17:   $\mathbf{u}_P(t) \leftarrow -\mathbf{S}^{-1}\mathbf{B}^T \mathbf{H}(t)\boldsymbol{\lambda}_P(t)$  ;
18:   $\mathbf{u}(t) \leftarrow \mathbf{u}(t) + \mathbf{u}_P(t)$  ;
19: return  $(i, \tau)$ 

```

---

over time and the number of additional endogenous actions, both over time and at time  $t_f$ . Under these definitions, we can find the relationship between the optimal intensity and the optimal cost by solving the minimization in the HJB equation:

$$\underset{\mathbf{u}(t)}{\text{minimize}} \quad \mathbf{u}^T(t) \Delta_B J + \frac{1}{2} \mathbf{u}^T(t) \mathbf{S} \mathbf{u}(t) \quad \text{subject to} \quad u_i(t) \geq 0, \quad i = 1, \dots, n.$$

By taking the differentiation with respect to  $\mathbf{u}(t)$ , the solution of the unconstrained minimization is given by

$$\mathbf{u}^*(t) = -\mathbf{S}^{-1} \Delta_B J, \tag{23}$$

which is the same as the solution to the constrained problem given that  $s_{ij} \geq 0$ , by definition, and  $(\Delta_B J)_i \leq 0$ , as proved in the Appendix F. Then, we substitute Eq. 23 into Eq. 22 and find that the optimal cost  $J$  needs to satisfy the following partial differential equation:

$$0 = J_t(\boldsymbol{\lambda}(t), t) + [w\boldsymbol{\lambda}_0 - w\boldsymbol{\lambda}(t)]^T \nabla_{\boldsymbol{\lambda}} J(\boldsymbol{\lambda}(t), t) + \boldsymbol{\lambda}^T(t) \Delta_B J - \frac{1}{2} \boldsymbol{\lambda}^T(t) \mathbf{Q} \boldsymbol{\lambda}(t) - \frac{1}{2} \Delta_B J^T \mathbf{S}^{-1} \Delta_B J, \tag{24}$$

with  $J(\boldsymbol{\lambda}(t_f), t_f) = \phi(\boldsymbol{\lambda}(t_f))$  as the terminal condition. The following lemma provides us with a solution to the above equation (proven in Appendix G):

**Lemma 10** *In the space of  $m$ -degree degree polynomials, the following quadratic form is the only solution to Eq. 24:*

$$J(\boldsymbol{\lambda}(t), t) = f(t) + \mathbf{g}(t)^T \boldsymbol{\lambda}(t) + \frac{1}{2} \boldsymbol{\lambda}(t)^T \mathbf{H}(t) \boldsymbol{\lambda}(t).$$

where  $\mathbf{g}(t)$  and  $\mathbf{H}(t)$  can be found by solving the following differential equations:

$$\begin{aligned}\dot{\mathbf{H}}(t) &= (\omega\mathbf{I} - \mathbf{B})^T \mathbf{H}(t) + \mathbf{H}(t)(\omega\mathbf{I} - \mathbf{B}) + \mathbf{H}(t)\mathbf{B}\mathbf{S}^{-1}\mathbf{B}^T \mathbf{H}(t) + \mathbf{Q} \\ \dot{\mathbf{g}}(t) &= [\omega\mathbf{I} - \mathbf{B}^T + \mathbf{H}(t)\mathbf{B}\mathbf{S}^{-1}\mathbf{B}^T]\mathbf{g}(t) - \omega\mathbf{H}(t)\boldsymbol{\lambda}_0 + \frac{1}{2} [\mathbf{H}(t)\mathbf{B}\mathbf{S}^{-1} - \mathbf{I}] \text{diag}(\mathbf{B}^T \mathbf{H}(t)\mathbf{B}).\end{aligned}$$

In the above lemma, note that the first differential equation is a matrix Riccati differential equation, which can be solved using many well-known efficient numerical solvers (Garrett, 2013), and the second one is a first order differential equation which has closed form solution. Both equations are solved backward in time with final conditions  $\mathbf{g}(t_f) = \mathbf{0}$  and  $\mathbf{H}(t_f) = -\mathbf{F}$ .

Finally, given the above cost-to-go function, the optimal intensity is given by the following theorem:

**Theorem 11** *The optimal intensity for the activity maximization problem defined by Eq. 19 with quadratic loss and penalty function is given by*

$$\mathbf{u}^*(t) = -\mathbf{S}^{-1}[\mathbf{B}^T \mathbf{g}(t) + \mathbf{B}^T \mathbf{H}(t)\boldsymbol{\lambda}(t) + \frac{1}{2} \text{diag}(\mathbf{B}^T \mathbf{H}(t)\mathbf{B})].$$

Since the optimal intensity is linear in  $\boldsymbol{\lambda}(t)$ , it allows for an efficient procedure to sample the times of the users' directly incentivized actions: at any given time  $t$ , we can view the multidimensional control signal  $\mathbf{u}(t)$  as a superposition of inhomogeneous multidimensional Poisson processes, one per non incentivized action, which start when the actions take place. Algorithm 3 summarizes our method, which we name CHESHIRE (Carroll, 1917).

Within the algorithm, *NextAction()* returns the time of the next (non directly incentivized) action in the network as well as the identity of the user who takes the action, once the action happens,  $\mathbf{e}_j$  is an indicator vector where the entry corresponding to user  $j$  is 1, and *Sample*( $\mathbf{u}(t)$ ) samples from a multidimensional inhomogeneous Poisson process with intensity  $\mathbf{u}(t)$  and it returns both the sampled time and dimension (*i.e.*, user). Moreover, note that the algorithm initially plans a user  $i$  and time  $\tau$  for the next directly incentivized action, however, if before  $\tau$ , a new organic action takes place at  $s < \tau$  and the intensity  $\boldsymbol{\lambda}(t)$  changes, then the algorithm updates the user and time for the next directly incentivized action using the superposition principle. To sample from a multidimensional inhomogeneous Poisson process, there exist multiple methods *e.g.*, refer to Lewis and Shedler (1979). Finally, note that one can precompute most of the quantities the algorithm needs, *e.g.*, lines 2-3,  $\mathbf{B}\mathbf{e}_j$  in line 8, and  $\mathbf{S}^{-1}\mathbf{B}^T \mathbf{H}(t)$  in line 9. Given these precomputations, the algorithm only needs to perform  $O(n)$  operations and sample  $\mathbf{1}^T \mathbf{N}(t_f)$  times from an inhomogeneous Poisson process.

## 5. Experiments

In this section, we validate REDQUEEN (Algorithm 1) and CHESHIRE (Algorithm 3) using both synthetic and real data gathered from Twitter and compare their performance with several state of the art methods and competitive baselines He et al. (2012).




Figure 2: Optimizing for one follower. Performance of REDQUEEN in comparison with the oracle and the method by Karimi et al. (2016) against number of broadcasted events. The feeds counting processes  $M(t)$  due to other broadcasters are Hawkes processes with  $\gamma_0 = 10$ ,  $\alpha = 1$  and  $w = 10$ . In all cases, the time horizon  $t_f - t_0$  is chosen such that the number of stories posted by other broadcasters is  $\sim 1000$ . Error bars are too small to be seen.

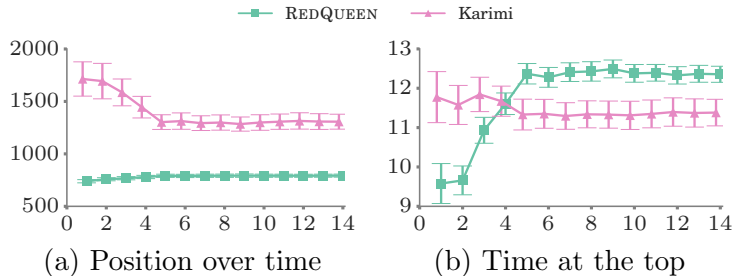


Figure 3: Optimizing for multiple followers. Performance of REDQUEEN in comparison with the method by Karimi et al. (2016) against number of followers. The feeds counting processes  $M(t)$  due to other broadcasters follow piecewise constant intensities, where the intensity of each follower remains constant within each piece, it varies as a half-sinusoid across pieces and it starts with a random initial phase. The performance of both methods stays constant upon addition of more followers.

## 5.1 When-to-post

### 5.1.1 EXPERIMENTS ON SYNTHETIC DATA

**Experimental setup.** We evaluate the performance via two quality measures: position over time,  $\int_0^T r(t)dt$ , and time at the top,  $\int_0^T \mathbb{I}(r(t) < 1)dt$  and compare the performance of REDQUEEN against the oracle, described in Section 3.3, and the method by Karimi et al. (2016), which, to the best of our knowledge, is the state of the art. Unless otherwise stated, we set the significance  $s_i(t) = 1, \forall t, i$  and use the parameter  $q$  to control the number of posts by REDQUEEN<sup>10</sup>.

**Optimizing for one follower.** We first experiment with one broadcaster and one follower against an increasing number of events (or budget). We generate the counting processes

10. The expected number of posts by REDQUEEN are a decreasing function of  $q$ . Hence, we can use binary search to guess  $q$  and then use averaging over multiple simulation runs to estimate the number of posts made.




Figure 4: Performance of REDQUEEN and the method by Karimi et al. (2016) for 2000 Twitter users, picked at random. The solid horizontal line (square) shows the median (mean) quality measure, normalized with respect to the value achieved by the users’ actual *true* posts, and the box limits correspond to the 25%-75% percentiles.

$M(t)$  due to other broadcasters using Hawkes processes, as defined by Eq. 1. We perform 10 independent simulation runs and compute the average and standard error (or standard deviation) of the quality measures. Fig. 2 summarizes the results, which show that our method: (i) consistently outperforms the method by Karimi et al. by large margins; (ii) achieves at most  $3\times$  higher position over time than the oracle as long as the budget is  $<30\%$  of the posted events by all other broadcasters; and, (iii) achieves  $>40\%$  of the value of time at the top that the oracle achieves.

**Optimizing for multiple followers.** Next, we experiment with one broadcaster and multiple followers. In this case, we generate the counting processes  $M(t)$  due to other broadcasters using piece-wise constant intensity functions. More specifically, we simulate the feeds of each follower for 1 day, using 24 1-hour long segments, where the rate of posts remains constant per follower in each segment and the rate itself varies as a half-sinusoid (*i.e.*, from  $\sin 0$  to  $\sin \pi$ ), with each follower starting with a random initial phase. This experimental setup reproduces volume changes throughout the day across followers’ feeds in different time-zones and closely resembles the settings in previous work (Karimi et al., 2016). The total number of posts by the REDQUEEN broadcaster is kept nearly constant and is used as the budget for the other baselines. Additionally, for Karimi’s method, we provide as input the true empirical rate of tweets per hour for each user. Here, we do not compare with the oracle since, due to its quadratic complexity, it does not scale.

Figure 3 summarizes the results. In terms of position over time, REDQUEEN outperforms Karimi’s method by a factor of 2. In terms of time at the top, REDQUEEN achieves  $\sim 18\%$  lower values than Karimi’s method for 1-4 followers but  $\sim 10\%$  higher values for  $>5$  followers. A potential reason Karimi’s method performs best in terms of time at the top for a low number of followers and piecewise constant intensities is that, while the number of followers is low, there are segments which are clearly favorable and thus Karimi’s method concentrates posts on those. However, as the number of followers increases, there are no clear favorable segments and advance planning does not give Karimi’s method any advantage. On the other hand, REDQUEEN, due to its online nature, is able to adapt to transient variations in the feeds.

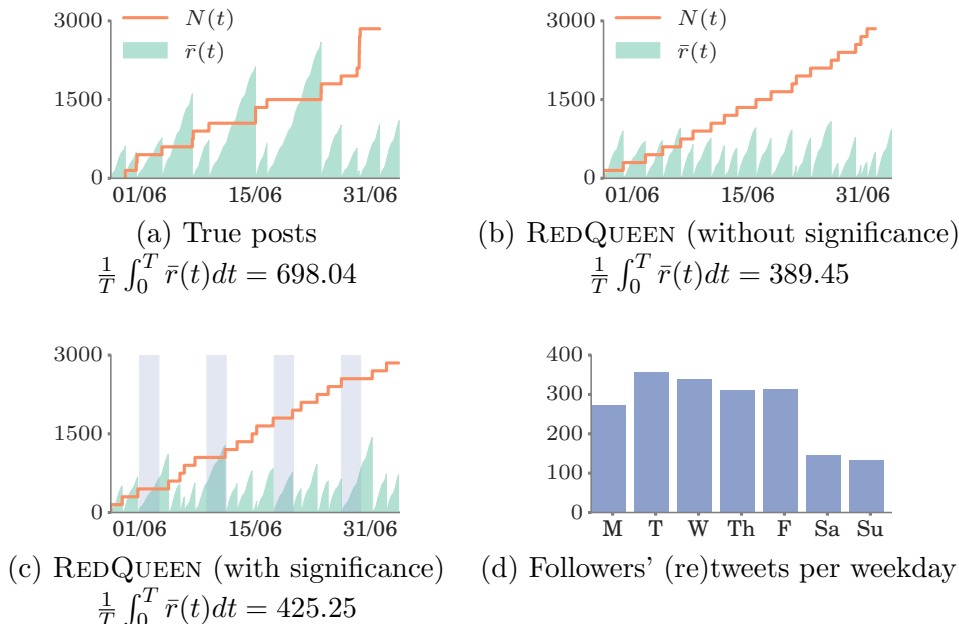


Figure 5: A broadcaster chosen from real data. Panels compares the position over time  $\bar{r}(t) = \sum_{i=0}^{\mathcal{N}} r(t)/\mathcal{N}$  (in green; lower is better) for the most recent tweet posted by a real user against the most recent one *posted* by a simulation run of REDQUEEN without and with significance. Here, the orange staircases represent the counts  $N(t)$  of the tweets posted by the real user and REDQUEEN over time. The shaded area in panel (c) highlights weekends. We can see that REDQUEEN avoided tweeting on weekends, when the followers are less likely to be active/logged-in, as seen in panel (d).

### 5.1.2 EXPERIMENTS ON REAL DATA

**Dataset description and experimental setup.** We use data gathered from Twitter as reported in previous work (Cha et al., 2010), which comprises profiles of 52 million users, 1.9 billion directed follow links among these users, and 1.7 billion public tweets posted by the collected users. The follow link information is based on a snapshot taken at the time of data collection, in September 2009. Here, we focus on the tweets published during a two month period, from July 1, 2009 to September 1, 2009, in order to be able to consider the social graph to be approximately static, and sample 2000 users uniformly at random as broadcasters and record all the tweets they posted. Then, for each of these broadcasters, we track down their followers and record all the (re)tweets they posted as well as reconstruct their timelines by collecting all the (re)tweets published by the people they follow. We assign equal significance to each follower but filter out those who follow more than 500 people since, otherwise, they would dominate the optimal strategy. Finally, we tune  $q$  such that the total number of tweets posted by our method is equal to the number of tweets the broadcasters tweeted during the two month period (with a tolerance of 10%).

**Solution quality.** We only compare the performance of our method against the method by Karimi et al. (2016) since the oracle does not scale to the size of real data. Moreover,

for the method by Karimi et al., we divide the two month period into ten segments of approximately one week to fit the piecewise constant intensities of the followers’ timelines, which the method requires. Fig. 4 summarizes the results by means of box plots, where position over time and time at the top are normalized with respect to the value achieved by the broadcasters’ actual *true* posts during the two month period. That means, if  $y = 1$ , the optimized intensity achieves the same position over time or time at the top as the broadcaster’s true posts. In terms of position over time and time at the top, REDQUEEN consistently outperforms competing methods by large margins and achieves  $0.28\times$  lower average position and  $3.5\times$  higher time at the top, in average, than the broadcasters’ true posts—in fact, it achieves lower position over time (higher time at the top) for 100% (99.1%) of the users.

**Time significance.** We look at the actual broadcasting strategies for one real user and investigate the effect of a time varying significance. We define  $s_i(t)$  to be the probability that follower  $i$  is online on that weekday, estimated empirically using the (re)tweets the follower posted as in Karimi et al. (2016). Fig. 5 compares the position over time for the most recent tweet posted by a real user against the most recent one *posted* by a simulation run of REDQUEEN with and without time varying significance. We can see that without significance information, REDQUEEN posts at nearly an even pace. However, when we supply empirically estimated significance, REDQUEEN avoids tweeting at times the followers are unlikely to be active, *i.e.*, the weekends, denoted by the shaded areas in panel (c) of Fig. 5. Due to this, the average position (maximum position) falls from 389.45 (1085.17) to 425.25 (1431.0), but is still lower than 698.04 (2597.9) obtained by the user’s original posting schedule.

## 5.2 Activity Maximization

### 5.2.1 EXPERIMENTS ON SYNTHETIC DATA

In this section, we first shed light on CHESHIRE’s sampling strategy in two small Kronecker networks (Leskovec et al., 2010) by recording, on the one hand, the number of directly incentivized actions per node and, on the other hand, the number of not directly incentivized actions per node in comparison with an uncontrolled setup. Then, we compare the performance of our method against several baselines and state of the art methods (Farajtabar et al., 2014, 2016) on a variety of large Kronecker networks and provides a scalability analysis.

**Sampling strategy.** We generate two small Kronecker networks with 64 nodes, a small core-periphery (parameter matrix  $[0.96, 0.3; 0.3, 0.96]$ ) and a dissortative network (param. matrix  $[0.3, 0.96; 0.96, 0.3]$ ), shown in Figure 6. For each network, we draw  $\mathbf{B}$  from a uniform distribution  $U(0, 10)$ ,  $\lambda_0$  also from a uniform distribution  $U(0, 10)$  for 20% of the nodes and  $\lambda_0 = 0$  for the remaining 80%, and set  $\omega = 16$ ,  $t_0 = 0$  and  $t_f = 5.5$ . Then, we compare the number of actions  $\mathbf{N}(t)$  over time under uncontrolled dynamics (Eq. 2; without CHESHIRE) and controlled dynamics (Eq. 17; with CHESHIRE). In both cases, we perform 20 simulation runs and sample non directly incentivized actions using Ogata’s thinning algorithm (Ogata, 1981). In the case of controlled dynamics, we sample directly incentivized actions using Algorithm 3.

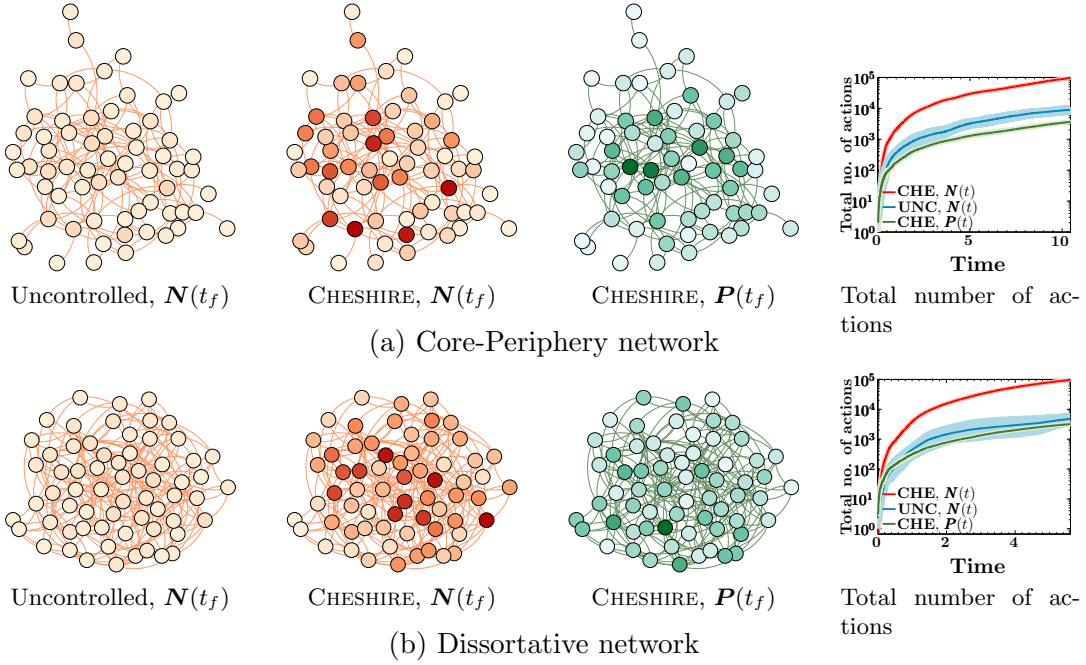


Figure 6: Activity on two 64-node networks,  $G_1$  (top) and  $G_2$  (bottom) under uncontrolled (Eq. 2; without CHESHIRE) and controlled (Eq. 17; with CHESHIRE) dynamics. The first and second columns visualize the final number of non incentivized actions  $N(t_f)$  under uncontrolled and controlled dynamics, where darker red corresponds to higher number of actions. The third column visualizes the final number of incentivized actions  $P(t_f)$  under controlled dynamics, where darker green corresponds to higher number of actions. The fourth column shows the temporal evolution of the number of incentivized and non incentivized actions across the whole networks for controlled and uncontrolled dynamics, where the solid line is the average across simulation runs and the shadowed region represents the standard error. By incentivizing a relatively small number of actions ( $\sim 3,600$  actions), CHESHIRE is able to increase the overall number of (non incentivized) actions dramatically ( $\sim 96,000$  vs  $\sim 4,800$  actions).

Figure 6 summarizes the results in terms of the number of non directly incentivized and incentivized actions, which show that: (i) a relatively small number of incentivized actions (fourth column;  $\sim 3,600$  actions) result in a dramatic increase in the overall number of (non directly incentivized) actions with respect the uncontrolled setup (fourth column;  $\sim 96,000$  vs  $\sim 4,800$  actions); (ii) the majority of the incentivized actions concentrate in a small set of influential nodes (third column; nodes in dark green); and, (iii) the variance of the overall number of actions (fourth column; shadowed regions) is consistently reduced when using CHESHIRE, in other words, the networks become more robust.

**Performance.** We experiment with five different types of large Kronecker networks (Leskovec et al., 2010) with 512 nodes: (i) homophily (assortative) networks (parameter matrix  $[0.96, 0.3; 0.3, 0.96]$ ); (ii) heterophily (dissortative) networks ( $[0.3, 0.96; 0.96, 0.3]$ ); (iii) random networks ( $[0.7, 0.7; 0.7, 0.7]$ ); (iv) hierarchical networks ( $[0.9, 0.1; 0.1, 0.9]$ ); and, (v)

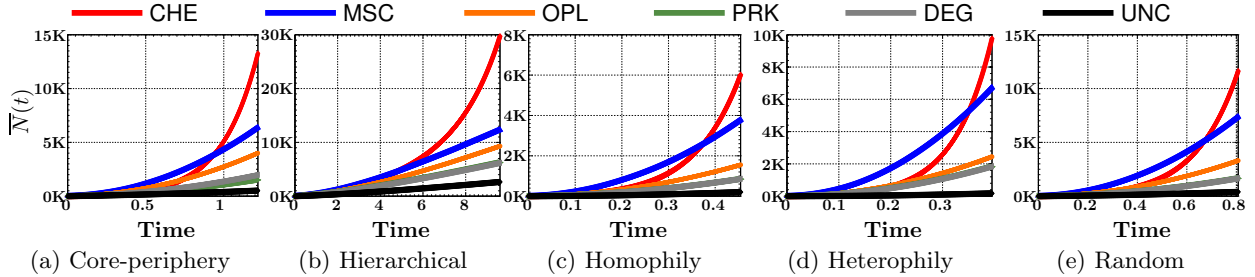


Figure 7: Performance over time of CHESHIRE against several competitors for several types of Kronecker networks. Performance is measured in terms of overall number of tweets  $\bar{N}(t) = \sum_{u \in \mathcal{V}} \mathbb{E}[N_u(t)]$ . In all cases, we tune the parameters  $Q$ ,  $S$  and  $F$  such that the total number of incentivized tweets *posted* by our method is equal to the budget used in the competing methods and baselines.




Figure 8: Scalability of CHESHIRE against several competitors. Panel (a) shows the running time against the cut-off time  $t_f$  for a 1,000 node Kronecker network. Panel (b) shows the running time for several Kronecker networks of increasing size with  $t_f = 0.5$ . In both panels, the average degree per node is 10. The experiments are carried out in a single machine with 24 cores and 64 GB of main memory.

core-periphery networks  $([0.9, 0.5; 0.5, 0.3])$ . For each network, we draw  $\lambda_0$  and  $\mathbf{B}$  from a uniform distribution  $U(0, 1)$  and set  $\omega = 100$ . We compare the performance of our algorithm, CHESHIRE, with two state of the art methods, which we denote as OPL (Farajtabar et al., 2014) and MSC (Farajtabar et al., 2016), and three baselines, which we denote as PRK, DEG, and UNC. OPL is a provably optimal algorithm for the offline setting and MSC is a suboptimal algorithm that discretizes the time window of interest in several rounds and, at any given round, it computes the control signal using the feedback from previous rounds to maximize the activity. PRK and DEG distribute the users’ control intensities  $\mathbf{u}(t)$  proportionally to the user’s page rank and outgoing degree in the network, respectively. UNC simply considers all control intensities to be zero, *i.e.*, it represents the uncontrolled dynamics.

For the large Kronecker networks, Figure 7 compares the performance of our algorithm against others in terms of overall average number of tweets  $\bar{N}(t) = \sum_{u \in \mathcal{V}} \mathbb{E}[N_u(t)]$  for a fixed budget  $\bar{P}(t_f) = \sum_{u \in \mathcal{V}} \mathbb{E}[P_u(t_f)] \approx 6.1\text{K}$ . We find that: (i) our algorithm consistently outperforms competing methods by large margins at time  $t_f$ ; (ii) it triggers up to 50%–100% more posts than the second best performer by time  $t_f$ ; (iii) MSC tends to use the



Dataset	$ \mathcal{V} $	$ \mathcal{E} $	$ \mathcal{H}(T_{\text{Data}}) $	$T = T_{\text{simulation}}$
Elections	231	1108	1584	120.2
Verdict	1059	10691	17452	22.11
Club	703	4154	9409	19.23
Sports	703	4154	7431	21.53
TV Show	947	10253	13203	12.11

Table 1: Real datasets statistics

budget too early, as a consequence, although it initially beats our method, it eventually gets outperformed by time  $t_f$ ; and, (iv) the baselines PRK and DEG have an underwhelming performance, suggesting that the network structure alone is not an accurate measure of influence.

**Scalability.** Figure 8 shows that our algorithm scales to large networks and is almost an order of magnitude faster than the second best performer, MSC (Farajtabar et al., 2016). For example, our algorithm takes  $\sim 30$  seconds to steer a network with 1,000 nodes and average degree of 10 while MSC takes  $\sim 4$  minutes.

### 5.2.2 EXPERIMENTS ON REAL DATA

In this section, we experiment with data gathered from Twitter and show that our model can maximize the number of online actions more effectively than several baselines and state of the art methods (Farajtabar et al., 2014, 2016).

**Experimental setup.** We experiment with five Twitter data sets about current real-world events, where actions are tweets (and retweets). To create each data set, we used the Twitter search API<sup>11</sup> to collect all the tweets (corresponding to a 2-3 weeks period around the event date) that contain hashtags related to:

- **Elections:** British election, from May 7 to May 15, 2015.
- **Verdict:** Verdict for the corruption-case against Jayalalitha, an Indian politician, from May 6 to May 17, 2015.
- **Club:** Barcelona getting the first place in La-liga, from May 8 to May 16, 2016.
- **Sports:** Champions League final in 2015, between Juventus and Real Madrid, from May 8 to May 16, 2015.
- **TV Show:** The promotion on the TV show “Games of Thrones”, from May 4 to May 12, 2015.

We then built the follower-followee network for the users that posted the collected tweets using the Twitter rest API<sup>12</sup>. Finally, we filtered out users that posted less than 200 tweets during the account lifetime, follow less than 100 users, or have less than 50 followers. An account of the data set statistics is given in Table 1.

Similarly as in the synthetic experiments, we compare the performance of our algorithm with two state of the art methods, OPL (Farajtabar et al., 2014) and MSC (Farajtabar

11. <https://dev.twitter.com/rest/public/search>

12. <https://dev.twitter.com/rest/public>

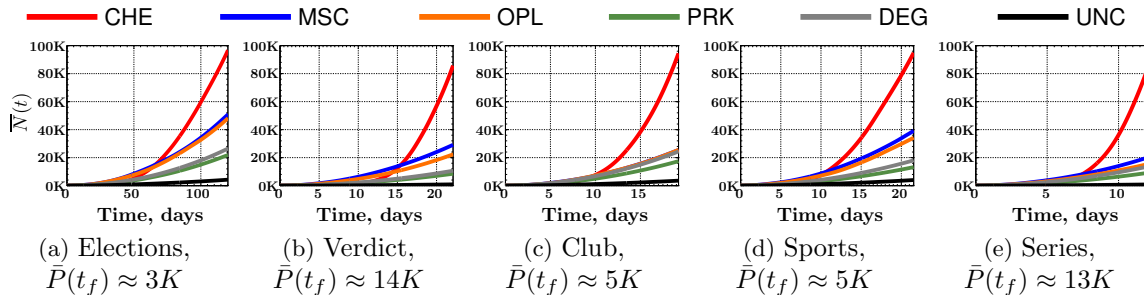


Figure 9: Performance over time of CHESHIRE against several competitors for each Twitter data set. Performance is measured in terms of overall number of tweets  $\bar{N}(t) = \sum_{u \in \mathcal{V}} \mathbb{E}[N_u(t)]$ . In all cases, we tune the parameters  $\mathbf{Q}$ ,  $\mathbf{S}$  and  $\mathbf{F}$  to be diagonal matrices such that the total number of incentivized tweets *posted* by our method is equal to the budget used in the competing methods and baselines. CHESHIRE (in red) consistently outperforms competing methods over time and it triggers up to 100%–400% more posts than the second best performer (in blue) as time goes by.

et al., 2016), and three baselines, PRK, DEG and UNC. More in detail, we proceed as follows.

For each data set, we estimate the influence matrix  $\mathbf{B}$  of the multidimensional Hawkes process defined by Eq. 2 using maximum likelihood, as elsewhere (Farajtabar et al., 2014; Valera and Gomez-Rodriguez, 2015). Moreover, we set the decay parameter  $\omega$  of the corresponding exponential triggering kernel  $\kappa(t)$  by cross-validation. Then, we perform 20 simulation runs for each method and baseline, where we sample non directly incentivized actions from the multidimensional Hawkes process learned from the corresponding Twitter data set using Ogata’s thinning algorithm (Ogata, 1981). For the competing methods and baselines, the control intensities  $\mathbf{u}(t)$  are deterministic and thus we only need to sample incentivized actions from inhomogeneous Poisson processes (Lewis and Shedler, 1979). For our method, the control intensities are stochastic and thus we sample incentivized actions using Algorithm 3. Here, we compare their performance in terms of the (empirical) average number of tweets  $\mathbb{E}[\mathbf{N}(t)]$ . In the above procedure, for a fair comparison, we tune the parameters  $\mathbf{Q}$ ,  $\mathbf{S}$  and  $\mathbf{F}$  to be diagonal matrices such that the total number of incentivized tweets *posted* by our method is equal to the budget used in the state of the art methods and baselines.

**Results.** We first compare the performance of our algorithm against others in terms of overall average number of tweets  $\bar{N}(t) = \sum_{u \in \mathcal{V}} \mathbb{E}[N_u(t)]$  for a fixed budget  $\bar{P}(t_f) = \sum_{u \in \mathcal{V}} P_u(t_f)$ . Figure 9 summarizes the results, which show that: (i) our algorithm consistently outperforms competing methods by large margins; (ii) it triggers up to 100%–400% more posts than the second best performer as time goes by; and, (iii) the baselines PRK and DEG have an underwhelming performance, suggesting that the network structure alone is not an accurate measure of influence.

Next, we evaluate the performance of our algorithm against others with respect to the available budget. To this aim, we compute the average time  $\bar{t}_{30K}$  required by each method to reach a milestone of 30,000 tweets against the number of directly incentivized tweets

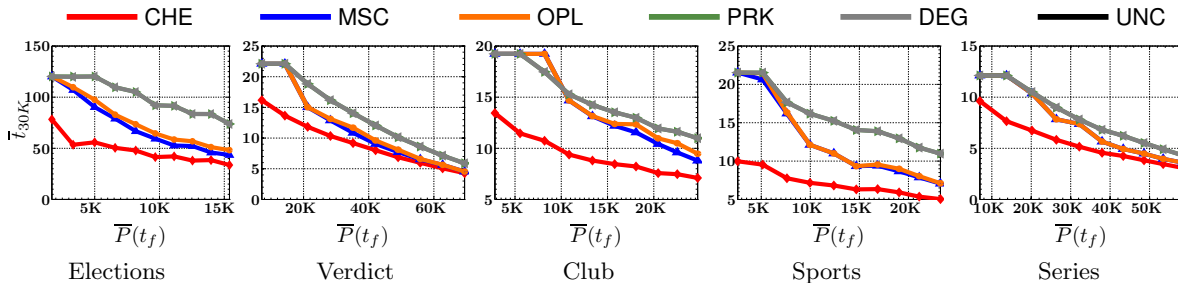


Figure 10: Performance vs. number of directly incentivized tweets for each Twitter data set. Performance is measured in terms of the average time  $\bar{t}_{30K}$  required by each method to reach a milestone of 30,000 tweets. CHESHIRE (in red) consistently reaches the milestone faster than the competing methods, *e.g.*, 20%–50% faster than the second best performer (in blue) for low budgets.

$\bar{P}(t_f)$  (*i.e.*, the budget). Here, we do not report the results for the uncontrolled dynamics (UNC) since it did not reach the milestone after  $10\times$  the time the slowest competitor took to reach it. Figure 10 summarizes the results, which show that: (i) our algorithm consistently reaches the milestone faster than the competing methods; (ii) it exhibits a greater competitive advantage when the budget is low; and, (iii) it reaches the milestone 20%–50% faster than the second best performer for low budgets.

## 6. Conclusions

In this paper, we developed efficient online algorithms for steering social activity, both at a user (local) and network (global) level, based on stochastic optimal control of stochastic differential equations (SDEs) with jumps. In doing so, we established a previously unexplored connection between optimal control of jump SDEs and doubly stochastic marked temporal point processes, which is of independent interest. We experimented with synthetic and real-world data gathered from Twitter and showed that our algorithms can consistently steer social activity more effectively than the state of the art.

Our work also opens many venues for future work. For example, from the perspective of an individual user, we considered social networks that sort stories in the users’ feeds in reverse chronological order (*e.g.*, Twitter, Weibo). Extending our methodology to social networks that sort stories algorithmically (*e.g.*, Facebook) is a natural next step. Moreover, we assume that only one broadcaster is using REDQUEEN. A very interesting follow-up would be augmenting our framework to consider multiple broadcasters under cooperative, competitive and adversarial environments. From the perspective of an entire social networking site, our experimental evaluation is based on simulation, using models whose parameters  $(\mathbf{B}, \omega)$  are learned from data. It would be very interesting to evaluate our method using actual interventions in a social network. From both perspectives, our algorithms optimize quadratic losses and assume the model parameters do not change over time, however, it would be useful to derive optimal broadcasting intensities for other losses with well-defined semantics and time-varying model parameters. Finally, optimal control of jump SDEs with doubly stochastic temporal point processes can be potentially applied to design online al-

gorithms for a wide variety of control problems in social and information systems, such as human learning (Reddy et al., 2016) or rumor control (Friggeri et al., 2014).

## Appendix A. Proof of Proposition 1

Using the left continuity of Poisson processes and the definition of derivative  $d\boldsymbol{\lambda}(t) = \boldsymbol{\lambda}(t + dt) - \boldsymbol{\lambda}(t)$  we can find the dynamics of the process using Ito's calculus (Hanson, 2007) as follows:

$$\begin{aligned}
 d\boldsymbol{\lambda}(t) &= \boldsymbol{\lambda}'_0(t) dt + \mathbf{C} \int_0^{t+dt} \kappa_\omega(t + dt - s) d\mathbf{N}(s) - \mathbf{C} \int_0^t \kappa_\omega(t - s) d\mathbf{N}(s) \\
 &= \boldsymbol{\lambda}'_0(t) dt + \mathbf{C} \int_0^{t+dt} (\kappa_\omega(t - s) + \kappa'_\omega(t - s) dt) d\mathbf{N}(s) - \mathbf{C} \int_0^t \kappa_\omega(t - s) d\mathbf{N}(s) \\
 &= \boldsymbol{\lambda}'_0(t) dt + \mathbf{C} \int_t^{t+dt} \kappa_\omega(t - s) d\mathbf{N}(s) + dt \mathbf{C} \int_0^{t+dt} \kappa'_\omega(t - s) d\mathbf{N}(s) \\
 &= \boldsymbol{\lambda}'_0(t) dt + \mathbf{C} \kappa_\omega(0) d\mathbf{N}(t) - \omega dt \mathbf{C} \int_0^{t+dt} \kappa_\omega(t - s) d\mathbf{N}(s) \\
 &= \boldsymbol{\lambda}'_0(t) dt + \mathbf{C} d\mathbf{N}(t) - \omega dt \mathbf{C} \int_0^t \kappa_\omega(t - s) d\mathbf{N}(s) \\
 &= [\boldsymbol{\lambda}'_0(t) + \omega \boldsymbol{\lambda}_0(t) - \omega \boldsymbol{\lambda}(t)] dt + \mathbf{C} d\mathbf{N}(t).
 \end{aligned}$$

## Appendix B. Proof of Lemma 3

$$\begin{aligned}
 J(\gamma(t), r(t), t) &= \min_{u(t, t_f)} \mathbb{E}_{(N, M)(t, t_f)} \left[ \phi(r(t_f)) + \int_t^{t_f} \ell(r(\tau), u(\tau)) d\tau \right] \\
 &= \min_{u(t, t_f)} \mathbb{E}_{(N, M)(t, t_f)} \left[ \phi(r(t_f)) + \int_t^{t+dt} \ell(r(\tau), u(\tau)) d\tau + \int_{t+dt}^{t_f} \ell(r(\tau), u(\tau)) d\tau \right] \\
 &= \min_{u(t, t_f)} \mathbb{E}_{(N, M)(t, t+dt)} \left[ \mathbb{E}_{(N, M)(t+dt, t_f)} \left[ \phi(r(t_f)) + \ell(t, r, u) dt + \int_{t+dt}^{t_f} \ell(r(\tau), u(\tau)) d\tau \right] \right] \\
 &= \min_{u(t, t+dt)} \min_{u(t+dt, t_f)} \mathbb{E}_{(N, M)(t, t+dt)} \left[ \ell(r(t), \gamma(t), t) dt + \mathbb{E}_{(N, M)(t+dt, t_f)} \left[ \phi(r(t_f)) + \int_{t+dt}^{t_f} \ell(r(\tau), u(\tau)) d\tau \right] \right] \\
 &= \min_{u(t, t+dt)} \mathbb{E}_{(N, M)(t, t+dt)} [J(\gamma(t + dt), r(t + dt), t + dt)] + \ell(r(t), u(t)) dt.
 \end{aligned}$$

## Appendix C. Proof of Lemma 4

According to the definition of differential,

$$\begin{aligned}
 dJ(r(t), \gamma(t), t) &:= J(r(t + dt), \gamma(t + dt), t + dt) - J(r(t), \gamma(t), t) \\
 &= J(r(t) + dr(t), \gamma(t) + d\gamma(t), t + dt) - J(r(t), \gamma(t), t).
 \end{aligned}$$

To evaluate the first term in the right hand side of the above equality we substitute  $dr(t)$  and  $d\gamma(t)$  using Eq. 9. Then, using the zero-one jump law (Kingman, 1992) we can write:

$$dJ = J(0, \gamma(t) + m(t)dt, t + dt) dN(t) + J(r(t) + g(t), \gamma(t) + m(t)dt + \alpha, t + dt) dP(t) \\ + J(r, \gamma(t) + m(t)dt, t + dt) (1 - dN(t))(1 - dP(t)) - J(r(t), \gamma(t), t),$$

where  $m(t) = \gamma'_0(t) + \omega\gamma_0(t) - \omega\gamma(t)$ . Then, we can expand the first three terms in the right hand sides:

$$J(0, \gamma + m(t)dt, t + dt) = J(0, \gamma(t), t) + J_\gamma(0, \gamma(t), t)m(t)dt + J_t(0, \gamma(t), t)dt \\ J(r(t) + 1, \gamma(t) + m(t)dt + \alpha, t + dt) = J(r(t) + 1, \gamma(t) + \alpha, t) + J_\gamma(r + 1, \gamma(t) + \alpha, t)m(t)dt \\ J(r(t), \gamma(t) + m(t)dt, t + dt) = J(r(t), \gamma(t), t) + J_\gamma(r(t), \gamma(t), t)m(t)dt + J_t(r(t), \gamma(t), t)dt,$$

using that the bilinear differential form  $dt dN(t) = 0$  (Hanson, 2007) and  $dN(t)dP(t) = 0$  by the zero-one jump law. Finally,

$$dJ(r(t), \gamma(t), t) = J_t(r(t), \gamma(t), t)dt + [\gamma'_0(t) + \omega\gamma_0(t) - \omega\gamma(t)] J_\gamma(r(t), \gamma(t), t)dt \\ + [J(0, \gamma(t), t) - J(r(t), \gamma(t), t)]dN(t) \\ + [J(r(t) + 1, \gamma(t) + \alpha, t) - J(r(t), \gamma(t), t)]dP(t),$$

which concludes the proof.

## Appendix D. Proof of Lemma 5

Consider the following proposal for the cost-to-go:

$$J(r(t), \gamma(t), t) = \sum_{i=0}^n \sum_{j=0}^m f_{ij}(t) r^i(t) \gamma^j(t),$$

where  $m$  and  $n$  are arbitrary large numbers and  $f_{ij}(t)$  are time-varying functions. Now, substitute this proposal in to Eq. 15:

$$0 = \sum_{i=1}^n f'_{i0} r^i + f_{i0} (r + 1)^i \gamma - f_{i0} r^i \gamma + \sum_{j=1}^m f'_{0j} \gamma^j + j(\gamma'_0 + w\gamma_0 - w\gamma) f_{0j} \gamma^{j-1} + f_{0j} (\gamma + \alpha)^j \gamma - f_{0j} \gamma^{j+1} \\ + \sum_{i=1}^n \sum_{j=1}^m j(\gamma'_0 + w\gamma_0 - w\gamma) f_{ij} r^i \gamma^{j-1} + \sum_{i=1}^n \sum_{j=1}^m f_{ij} (r + 1)^i (\gamma + \alpha)^j \gamma - f_{ij} r^i \gamma^{j+1} \\ - \frac{1}{2} q^{-1} \left[ \sum_{i=1}^n f_{i0} r^i + \sum_{i=1}^n \sum_{j=1}^m f_{ij} r^i \gamma^j \right]^2 + \frac{1}{2} s r^2 + f'_{00}$$

where for notational simplicity we omitted the time argument of functions. To find the unknown functions  $f_{ij}(t)$ , we equate the coefficient of different variables. If we consider the coefficient of  $r^{2n}$ , we have  $f_{n0}(t) = 0$ . We can continue this argument for  $n - 1, n - 2, \dots, 2$  to show that  $\forall i \geq 2; f_{i0}(t) = 0$ . Similar reasoning for coefficients of  $r^{2i} \gamma^{2j}$  shows that  $\forall j, i \geq$

2;  $f_{ij}(t) = 0$ . Finally, the coefficient of  $r^2$  is  $s(t)/2 - q^{-1}f_{10}^2(t)/2 = 0$ , so  $f_{10}(t) = \sqrt{s(t)q}$ . If we rename  $f_{0j}(t)$  to  $g_j(t)$  and  $f_{00}(t)$  to  $f(t)$ , then it follows that

$$J(r(t), \gamma(t), t) = f(t) + \sqrt{s(t)q}r(t) + \sum_{j=1}^m g_j(t)\gamma^j(t).$$

We can continue the previous method to find the remaining coefficients and completely define the cost-to-go function. If we equate the coefficient of  $\gamma^j$  to zero we would have a system of first order differential equation which its  $j$ 'th row for  $j > 1$  is

$$g'_j(t) - g_{j-1}(t) + j(\alpha - w)g_j(t) + (j+1)(\gamma'_0(t) + w\gamma_0(t) + \frac{j}{2}\alpha^2)g_{j+1}(t) + \sum_{k=2}^{m-j} \binom{j+k}{k+1} \alpha^{k+1} g_{j+k}(t) = 0$$

We can also evaluate the corresponding differential equation for  $j = 1$ . When  $\gamma_0(t) = \gamma_0$ , we can express these differential equations using a matrix differential equation  $\mathbf{g}'(t) = \mathbf{A}\mathbf{g}(t)$ . and its solution is  $\mathbf{g}(t) = c_1 e^{\zeta_1 t} \mathbf{u}_1 + c_2 e^{\zeta_2 t} \mathbf{u}_2 + \dots + c_n e^{\zeta_n t} \mathbf{u}_n$  where  $\zeta_i$  and  $\mathbf{u}_i$  are eigenvalue and eigenvector of matrix  $\mathbf{A}$  and  $c_i$  is a constant found using the terminal conditions. Since in triangular matrices diagonal entries are eigenvalues, we have  $\mathbf{g}(t) = \sum_{i=1}^m c_i e^{j(w-\alpha)t} \mathbf{u}_i$ . We can approximate general time varying  $\gamma_0(t)$  using piecewise function and repeat the above procedure for each piece.

## Appendix E. Proof of Lemma 9

According to the definition of differential,

$$dJ(\boldsymbol{\lambda}(t), t) := J(\boldsymbol{\lambda}(t+dt), t+dt) - J(\boldsymbol{\lambda}(t), t) = J(\boldsymbol{\lambda}(t) + d\boldsymbol{\lambda}(t), t+dt) - J(\boldsymbol{\lambda}(t), t).$$

To evaluate the first term in the right hand side of the above equality we substitute  $d\boldsymbol{\lambda}(t)$  using Eq. 18. Then, using the zero-one jump law (Kingman, 1992) we can write:

$$\begin{aligned} J(\boldsymbol{\lambda}(t) + d\boldsymbol{\lambda}(t), t+dt) &= J(\boldsymbol{\lambda}(t) + m(t)dt + \mathbf{B}d\mathbf{N}(t) + \mathbf{B}d\mathbf{P}(t), t+dt) \\ &= \sum_i J(\boldsymbol{\lambda}(t) + m(t)dt + \mathbf{b}_i, t+dt)dN_i(t) + \sum_i J(\boldsymbol{\lambda}(t) + m(t)dt + \mathbf{b}_i, t+dt)dP_i(t) \\ &\quad + J(\boldsymbol{\lambda}(t) + m(t)dt, t+dt) \prod_i [1 - dN_i(t)][1 - dP_i(t)] \\ &= J(\boldsymbol{\lambda}(t) + m(t)dt, t+dt) [1 - \sum_i dN_i(t) + dP_i(t)] + \sum_i J(\boldsymbol{\lambda}(t) + m(t)dt + \mathbf{b}_i, t+dt) [dN_i(t) + dP_i(t)] \\ &= J(\boldsymbol{\lambda}(t) + m(t)dt, t+dt) + \sum_i [J(\boldsymbol{\lambda}(t) + m(t)dt + \mathbf{b}_i, t+dt) - J(\boldsymbol{\lambda}(t) + m(t)dt, t+dt)] [dN_i(t) + dP_i(t)] \end{aligned}$$

where  $m(t) := \omega\boldsymbol{\lambda}_0 - \omega\boldsymbol{\lambda}(t)$  and we used that the bilinear differential form  $dt d\mathbf{N}(t) = 0$  (Hanson, 2007). By total derivative rule, it follows that

$$\begin{aligned} J(\boldsymbol{\lambda}(t) + m(t)dt + \mathbf{b}_i, t+dt) &= J(\boldsymbol{\lambda}(t) + \mathbf{b}_i, t) + \nabla_{\boldsymbol{\lambda}} J(\boldsymbol{\lambda}(t) + \mathbf{b}_i, t)m(t)dt + J_t(\boldsymbol{\lambda}(t) + \mathbf{b}_i, t)dt \\ J(\boldsymbol{\lambda}(t) + m(t)dt, t+dt) &= J(\boldsymbol{\lambda}(t), t) + \nabla_{\boldsymbol{\lambda}} J(\boldsymbol{\lambda}(t), t)m(t)dt + J_t(\boldsymbol{\lambda}(t), t)dt. \end{aligned}$$

Then, the differential is given by:

$$dJ(\boldsymbol{\lambda}(t) + d\boldsymbol{\lambda}(t), t + dt) = J(\boldsymbol{\lambda}(t), t) + (\omega\boldsymbol{\lambda}_0 - \omega\boldsymbol{\lambda}(t))^T \nabla_{\boldsymbol{\lambda}} J(\boldsymbol{\lambda}(t), t) dt + J_t(\boldsymbol{\lambda}(t), t) dt + \sum_i [J(\boldsymbol{\lambda}(t) + \mathbf{b}_i, t) - J(\boldsymbol{\lambda}(t), t)] [dN_i(t) + dP_i(t)],$$

which completes the proof.

### Appendix F. Proof of $(\Delta_B J)_i \leq 0$

Lets  $t < s$ , then according to the definition we can write,

$$\boldsymbol{\lambda}(s) = \boldsymbol{\lambda}_0 + \int_0^s \kappa_\omega(s - \tau) \mathbf{B} d\mathbf{N}(\tau) = \boldsymbol{\lambda}_0 + \int_0^t \kappa_\omega(s - \tau) \mathbf{B} d\mathbf{N}(\tau) + \int_t^s \kappa_\omega(s - \tau) \mathbf{B} d\mathbf{N}(\tau).$$

For the exponential kernel  $\kappa_\omega(t) = e^{-wt}$  we have,

$$\int_0^t \kappa_\omega(s - \tau) \mathbf{B} d\mathbf{N}(\tau) = \int_0^t e^{-w(s-\tau)} \mathbf{B} d\mathbf{N}(\tau) = e^{-w(s-t)} \int_0^t e^{-w(t-\tau)} \mathbf{B} d\mathbf{N}(\tau) = e^{-w(s-t)} (\boldsymbol{\lambda}(t) - \boldsymbol{\lambda}_0)$$

so given the value of  $\boldsymbol{\lambda}(t)$  at time  $t$  then we can write  $\boldsymbol{\lambda}(s)$  for later times as

$$\boldsymbol{\lambda}(s) = \boldsymbol{\lambda}_0 + e^{-w(s-t)} (\boldsymbol{\lambda}(t) - \boldsymbol{\lambda}_0) + \int_t^s \kappa_\omega(s - \tau) \mathbf{B} d\mathbf{N}(\tau)$$

Lets consider a process  $\boldsymbol{\xi}(s)$  with intensity value at time  $t$  equal to  $\boldsymbol{\lambda}(t) + \mathbf{b}_i$  as,

$$\boldsymbol{\xi}(s) = \boldsymbol{\lambda}_0 + e^{-w(s-t)} (\boldsymbol{\lambda}(t) + \mathbf{b}_i - \boldsymbol{\lambda}_0) + \int_t^s \kappa_\omega(s - \tau) \mathbf{B} d\mathbf{N}(\tau).$$

Since  $\mathbf{b}_i \succeq 0$ , then given the same history in interval  $(t, s)$ , we have  $\boldsymbol{\xi}(s) \succeq \boldsymbol{\lambda}(s)$ . Then, we have:

$$\ell(\boldsymbol{\xi}(s), \mathbf{u}(s)) \leq \ell(\boldsymbol{\lambda}(s), \mathbf{u}(s)).$$

Now, taking the integration, then expectation (over all histories) and finally the minimization from the above inequality does not change the direction of inequality. So it readily follows the required result.

$$J(\boldsymbol{\lambda}(t) + \mathbf{b}_i, t) \leq J(\boldsymbol{\lambda}(t), t).$$

### Appendix G. Proof of Lemma 10

Consider the following proposal of degree three for the cost-to-go function:

$$J(\boldsymbol{\lambda}(t), t) = f(t) + \sum_i g_i(t) \lambda_i(t) + \sum_i \sum_j \lambda_i(t) \lambda_j(t) H_{ij}(t) + \sum_i \sum_j \sum_k \lambda_i(t) \lambda_j(t) \lambda_k(t) H_{ijk}(t)$$

If we plug this proposal in Eq. 22, and evaluate the coefficient of fourth degree terms like  $\lambda_i^2(t)\lambda_j^2(t)$  and equate them to zero, then we can find the unknown coefficients  $H_{ijk}(t)$ 's as follows:

$$\forall i, j, t : \sum_k \left( \sum_\ell B_{\ell k}^2 \right) H_{ijk}^2(t) = 0$$

Since the sum of positives terms is zero if and only if they all be zero, then  $H_{ijk}(t)$  and consequently the terms with degree three in the proposal are all zero. So the proposal reduces to a quadratic proposal.

It is quite straightforward to extend this argument for proposals with order  $m > 3$  and by equating the degree  $2m - 1$  terms, similarly conclude that coefficients of degree  $m$  terms in the proposal are zero. If we repeat this argument for  $m - 1, \dots, 3$ , we deduce that any proposal with arbitrary degree  $m \geq 2$ , would result in a quadratic optimal cost.

## References

- O. Aalen, O. Borgan, and H. K. Gjessing. *Survival and event history analysis: a process point of view*. Springer, 2008.
- L. Backstrom, E. Bakshy, J. M. Kleinberg, T. M. Lento, and I. Rosenn. Center of attention: How facebook users allocate attention across friends. *Proceedings of the 5th International AAAI Conference on Weblogs and Social Media*, 2011.
- D. P. Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena Scientific Belmont, MA, 1995.
- R. M. Blumenthal. An extended markov property. *Transactions of the American Mathematical Society*, 85(1):52–72, 1957.
- L. Carroll. *Through the looking glass: And what Alice found there*. Rand, McNally, 1917.
- M. Cha, H. Haddadi, F. Benevenuto, and P. K. Gummadi. Measuring User Influence in Twitter: The Million Follower Fallacy. *Proceedings of the 4th International AAAI Conference on Weblogs and Social Media*, 2010.
- A. De, I. Valera, N. Ganguly, S. Bhattacharya, and M. Gomez-Rodriguez. Learning and forecasting opinion dynamics in social networks. In *Advances in Neural Information Processing Systems*, 2016.
- N. Du, L. Song, M. Gomez-Rodriguez, and H. Zha. Scalable influence estimation in continuous-time diffusion networks. In *Advances in Neural Information Processing Systems*, 2013.
- M. Farajtabar, N. Du, M. Gomez-Rodriguez, I. Valera, H. Zha, and L. Song. Shaping social activity by incentivizing users. In *Advances in Neural Information Processing Systems*, 2014.



- M. Farajtabar, M. Gomez-Rodriguez, N. Du, M. Zamani, H. Zha, and L. Song. Back to the past: Source identification in diffusion networks from partially observed cascades. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, 2015a.
- M. Farajtabar, Y. Wang, M. Gomez-Rodriguez, S. Li, H. Zha, and L. Song. COEVOLVE: A joint point process model for information diffusion and network co-evolution. In *Advances in Neural Information Processing Systems*, 2015b.
- M. Farajtabar, X. Ye, S. Harati, L. Song, and H. Zha. Multistage campaigning in social networks. In *Advances in Neural Information Processing Systems*, 2016.
- A. Friggeri, L. Adamic, D. Eckles, and J. Cheng. Rumor cascades. In *Proceedings of the 8th International AAAI Conference on Weblogs and Social Media*, 2014.
- C. K. Garrett. *Numerical integration of matrix Riccati differential equations with solution singularities*. PhD thesis, The University of Texas at Arlington, 2013.
- S. Goel, D. J. Watts, and D. G. Goldstein. The structure of online diffusion networks. In *Proceedings of the 13th ACM conference on electronic commerce*, 2012.
- M. Gomez-Rodriguez, D. Balduzzi, and B. Schölkopf. Uncovering the temporal dynamics of diffusion networks. In *Proceedings of the 28th International Conference on Machine Learning*, 2011.
- M. Gomez-Rodriguez, K. Gummadi, and B. Schölkopf. Quantifying information overload in social media and its impact on social contagions. In *Proceedings of the 8th International AAAI Conference on Weblogs and Social Media*, 2014.
- F. B. Hanson. *Applied stochastic processes and control for Jump-diffusions: modeling, analysis, and computation*, volume 13. SIAM, 2007.
- A. G. Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83–90, 1971.
- X. He, G. Song, W. Chen, and Q. Jiang. Influence blocking maximization in social networks under the competitive linear threshold model. In *SDM*, pages 463–474. SIAM, 2012.
- N. Hodas and K. Lerman. How visibility and divided attention constrain social contagion. *SocialCom*, 2012.
- M. Jankowiak and M. Gomez-Rodriguez. Uncovering the spatiotemporal patterns of collective social activity. In *Proceedings of the 2017 SIAM International Conference on Data Mining*, pages 822–830. SIAM, 2017.
- J. H. Kang and K. Lerman. Vip: Incorporating human cognitive biases in a probabilistic model of retweeting. In *Social Computing, Behavioral-Cultural Modeling, and Prediction*, pages 101–110. Springer, 2015.

- M. Karimi, E. Tavakoli, M. Farajtabar, L. Song, and M. Gomez-Rodriguez. Smart Broadcasting: Do you want to be seen? In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016.
- J. F. C. Kingman. *Poisson processes*. Oxford University Press, 1992.
- K. Lerman and T. Hogg. Leveraging position bias to improve peer recommendation. *PLoS one*, 9(6), 2014.
- J. Leskovec, D. Chakrabarti, J. M. Kleinberg, C. Faloutsos, and Z. Ghahramani. Kronecker graphs: An approach to modeling networks. *Journal of Machine Learning Research*, 11(2):985–1042, 2010.
- P. A. Lewis and G. S. Shedler. Simulation of nonhomogeneous poisson processes by thinning. *Naval research logistics quarterly*, 26(3):403–413, 1979.
- C. Mavroforakis, I. Valera, and M. Gomez-Rodriguez. Modeling the dynamics of online learning activity. In *Proceedings of the 26th International World Wide Web Conference*, 2017.
- S. Mishra, M. RizoIU, and L. Xie. Feature driven and point process approaches for popularity prediction. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management, CIKM '16*, pages 1069–1078, 2016.
- Y. Ogata. On lewis' simulation method for point processes. *Information Theory, IEEE Transactions on*, 27(1):23–31, 1981.
- S. Reddy, I. Labutov, S. Banerjee, and T. Joachims. Unbounded human learning: Optimal scheduling for spaced repetition. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016.
- M. A. RizoIU, L. Xie, S. Sanner, M. Cebrian, H. Yu, and P. V. Hentenryck. Expecting to be hip: Hawkes intensity processes for social media popularity. In *Proceedings of the 26th International Conference on World Wide Web*, 2017.
- N. Spasojevic, Z. Li, A. Rao, and P. Bhattacharyya. When-to-post on social networks. In *Proceedings of the 21st ACM SIGKDD international conference on Knowledge discovery and data mining*, 2015.
- B. Tabibian, I. Valera, M. Farajtabar, L. Song, B. Schoelkopf, and M. Gomez-Rodriguez. Distilling information reliability and source trustworthiness from digital traces. In *Proceedings of the 26th International World Wide Web Conference*, 2017.
- I. Valera and M. Gomez-Rodriguez. Modeling diffusion of competing products and conventions in social media. In *Proceedings of the IEEE International Conference on Data Mining*, 2015.
- Y. Wang, E. Theodorou, A. Verma, and L. Song. A stochastic differential equation framework for guiding online user activities in closed loop. *arXiv preprint arXiv:1603.09021*, 2016.

- Y. Wang, G. Williams, E. Theodorou, and L. Song. Variational policy for guiding point processes. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.
- H. Xu, D. Luo, and H. Zha. Learning Hawkes processes from short doubly-censored event sequences. In *Proceedings of the 34th International Conference in Machine Learning*, 2017.
- A. Zarezade, S. Jafarzadeh, and H. R. Rabiee. Spatio-temporal modeling of check-ins in location-based social networks. *arXiv preprint arXiv:1611.07710*, 2016a.
- A. Zarezade, A. Khodadadi, M. Farajtabar, H. R. Rabiee, and H. Zha. Correlated Cascades: Compete or Cooperate. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 2016b.
- A. Zarezade, U. Upadhyay, H. R. Rabiee, and M. Gomez-Rodriguez. Redqueen: An online algorithm for smart broadcasting in social networks. In *Proceedings of the 10th ACM International Conference on Web Search and Data Mining*, 2017.
- Q. Zhao, M. Erdogdu, H. He, A. Rajaraman, and J. Leskovec. Seismic: A self-exciting point process model for predicting tweet popularity. In *Proceedings of the 21st ACM SIGKDD international conference on Knowledge discovery and data mining*, 2015.