

Dynamic Influence Maximization Under Increasing Returns to Scale

Haifeng Zhang
Vanderbilt University
haifeng.zhang@vanderbilt.edu

Ariel D. Procaccia
Carnegie Mellon University
arielpro@cs.cmu.edu

Yevgeniy Vorobeychik
Vanderbilt University
yevgeniy.vorobeychik@vanderbilt.edu

ABSTRACT

Influence maximization is a problem of maximizing the aggregate adoption of products, technologies, or even beliefs. Most past algorithms leveraged an assumption of submodularity that captures diminishing returns to scale. While submodularity is natural in many domains, early stages of innovation adoption are often better characterized by convexity, which is evident for renewable technologies, such as rooftop solar. We formulate a dynamic influence maximization problem under increasing returns to scale over a finite time horizon, in which the decision maker faces a budget constraint. We propose a simple algorithm in this model which chooses the best time period to use up the entire budget (called *Best-Stage*), and prove that this policy is optimal in a very general setting. We also propose a heuristic algorithm for this problem of which *Best-Stage* decision is a special case. Additionally, we experimentally verify that the proposed “best-time” algorithm remains quite effective even as we relax the assumptions under which optimality can be proved. However, we find that when we add a “learning-by-doing” effect, in which the adoption costs decrease not as a function of time, but as a function of aggregate adoption, the “best-time” policy becomes suboptimal, and is significantly outperformed by our more general heuristic.

Categories and Subject Descriptors: G.1.6 [Numerical Analysis]: Optimization—Dynamic Influence Maximization, Constrained optimization

Keywords: Dynamic Influence Maximization; Convexity; Backward Induction; Heuristic Algorithm

1. INTRODUCTION

One of the important algorithmic questions in marketing on social networks is how one should leverage network effects in promoting products so as to maximize long-term product uptake. Indeed, a similar question arises in political science, if framed in terms of maximizing uptake of beliefs and attitudes, leading to particular voting behavior. Crucial to such problems is a model of social influence on networks, and a number of such models have been proposed [15, 9, 19, 10, 17, 5, 11, 13]. Some of the prominent models give rise to global influence functions (capturing the expected number

of adopters as a function of a subset of initially targeted individuals) that are *submodular*, that is, possess a natural diminishing returns property: targeting a greater number of individuals yields a lower marginal increase in the outcome. Submodularity is a powerful property for algorithm design; in particular, a simple greedy heuristic has a provable approximation guarantee, and is typically very close to optimal in experiments [10, 5, 13, 16]. Submodularity as typically defined is only helpful in a static setting, that is, when we choose individuals to target in a single shot. But an extension, termed *adaptive submodularity*, was proposed to make use of greedy heuristics in dynamic environments where individuals can be targeted over time [8].

The diminishing returns feature naturally arises in many settings. However, early adoption trends can exhibit the opposite property of *increasing returns to scale*. For example, in the classic Bass model [3] the early portion of the famous “S-curve” is convex, and if one uses logistic regression to model individual adoption choice—a natural model that was used in recent work by the authors that learned individual rooftop solar adoption behavior [20]—the model is convex when probabilities are small. Arguably, early adoption settings are most significant for the development of effective product promotion strategies, since overall uptake is quite critical to the ultimate success of the product line. This is an especially acute concern in the “green energy” sector, where renewable energy technologies, such as rooftop solar, are only at a very early stage of overall adoption—indeed, adoption has been negligible except in a few states, such as California and Arizona.

We consider the problem of influence maximization with network effects by aggregating a social network into an “aggregate” adoption function which takes as input the number of adopters at a given time t and outputs the number of adopters at time $t + 1$. Our main theoretical result is that when this function is convex and marketing budget can be reinvested at a fixed interest rate δ (equivalently, marketing or “seeding” costs decrease exponentially over time), influence over a finite time period can be maximized by using up the entire targeted marketing budget at a single time point (rather than splitting up the budget among multiple time periods); we refer to the resulting simple algorithm as the *Best-Stage* algorithm.

We study the degree to which the theoretical optimality of the best-time algorithm holds in practice, using real data of rooftop solar adoption in San Diego county [20]. As a baseline, we develop a more general heuristic algorithm that splits the budget equally over a set of consecutive months,

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.
Copyright © 2015, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

with the size of this set and the starting month allowed to vary. We find that investing the whole budget in a single month is indeed (almost) optimal under a variety of simplified seeding cost functions (exponential, polynomial, or even linear with time), despite a number of gaps between the theory and the experimental setup — suggesting that the theoretical model is somewhat robust.

In contrast, the best-time algorithm becomes suboptimal in an “ideal” model that was previously validated to confirm its predictive efficacy [20]. Through careful analysis, we find that this is the result of “learning-by-doing” effects, where marketing costs (for example, when actual products are given away for free) depend on the total product uptake in the marketplace.

2. RELATED WORK

We build on the extensive literature of economics of diffusion with network effects [7, 4, 2, 6]. Largely, however, this literature is concerned with equilibria that arise, rather than algorithmic considerations. The latter are extensively studied in the literature on influence maximization on social networks. A number of models have been proposed to quantify influence diffusion on a network, perhaps the most prominent of which are linear threshold and independent cascade models [15, 9, 19, 10, 12, 21]. These and many related models give rise to a submodular “expected adoption” function. Many past approaches to “one-shot” influence maximization take advantage of this property in algorithmic development; in particular, a simple greedy algorithm is highly effective in practice, and offers a constant-factor approximation guarantee relative to a globally optimal solution to the associated combinatorial optimization problem [10, 13, 12].

While one-shot influence maximization on social networks has received much attention, significantly less work addresses the problem of dynamic influence maximization, where individuals can be seeded in sequence. In an important effort in this vein, Golovin and Krause show that when dynamics (and uncertainty) satisfy the property of *adaptive submodularity*, a natural *dynamic* greedy heuristic is guaranteed to be only a constant factor worse than the optimal clairvoyant algorithm [8]. Our problem is distinct from this effort in several ways. The first, and key, distinction is that we are concerned with increasing returns to scale. The second is that we capture network effects simply in terms of total numbers of all past adoptions (thus, the social network is completely connected for our purposes). The third is that we introduce another key element of tension into the problem by supposing that there is a fixed total budget allocated for seeding, and either this budget (or any portion of it) can be set aside to collect interest, or the costs of seeding fall over time (commonly, costs of products are expected to fall over time as a result of learning-by-doing, or supply-side network effects where better processes and technology reduce production costs with increasing product uptake and experience in the marketplace). To our knowledge, we are the first to consider the algorithmic problem of dynamic influence maximization in such a context.

3. THE DYNAMIC INFLUENCE MAXIMIZATION MODEL

We consider a problem of adoption diffusion with network externalities. Our model of network externalities is simplified to consider only the aggregate adoption, which we denote by D . Specifically, we assume that the diffusion takes place in discrete time, and the aggregate adoption at time t , D_t , is a function only of adoption by the previous time step, D^{t-1} : $D_t = f(D^{t-1})$. In other words, adoption dynamics are deterministic and first-order Markovian. We make three assumptions on the nature of the diffusion function f :

1. f is strictly convex (i.e., increasing returns to scale),
2. f is strictly monotonically increasing, and
3. $\forall D > 0, f(D) > D$.

Assumptions (2) and (3) ensure that aggregate adoption increases over time (i.e., with every application of f); note that they are not redundant. Assumption (1) of increasing returns to scale captures a common model of early adoption dynamics (the convex portion of the “S-curve” [3], or the logistic regression model of adoption over a low-probability region discussed in our experiments below). We suppose that at time 0 (that is, initial decision time), there is some initial number of adopters in the population, $D^0 \geq 0$.

As stated, the problem poses no algorithmic tension: if one had a fixed budget for stimulating adoption, the entire budget should be used up immediately, as it would then take maximal advantage of network effects. We now introduce the principal source of such tension: an exponentially increasing budget. Specifically, suppose an agent is initially given a budget B^0 and any remaining budget will accrue by a factor of δ . For example, we can decide to invest residual budget at an interest rate δ . Alternatively, if we are giving away a product, its cost may decrease over time as technology matures, a process often referred to as “learning-by-doing” [1, 14, 18]. Such learning-by-doing effects are paramount when we consider technology evolution in its early stages, which is arguably the setting where we would be most interested in promoting the product by giving it away to a subset of the individuals in a population. Notice that either saving some of the budget at a fixed interest rate, or costs of seeding decreasing at a constant rate, both give rise to exponentially increasing purchasing power of the budget over time. This gives rise to a non-trivial tension between seeding early so as to maximally leverage network effects and seeding later so as to maximize the number of individuals seeded (as well as subsequent network effects). The algorithmic question we pose is: **how should we use a given budget B over a fixed time horizon T so as to maximize the total number of adopters at time T ?** For simplicity, we assume unit cost for every seed; in other words, any budget B will create exactly B new adopters. This assumption will be relaxed in our experiments below.

Given the deterministic T -stage diffusion model and the budget accrual process described above, as well as an initial budget B^0 and aggregate adoption D^0 , we can define a “seeding” policy π as a sequence of fractions of the budget allocated in each stage, that is

$$\pi = (\alpha_0, \alpha_1, \dots, \alpha_{T-1}), \quad \alpha_t \geq 0.$$

The dynamic influence maximization problem aspires to compute a policy π^* which leads to the maximum number of

adopters based on the diffusion model f starting with initial adoption D^0 , an initial budget B^0 , and a budget growth rate δ . If we define the total number of adopters at time T as the *utility* of a policy, the problem can be written as

$$\pi^* = \arg \max_{\pi} U(f, \pi, D^0, B^0, \delta).$$

4. ALGORITHMS FOR DYNAMIC INFLUENCE MAXIMIZATION

To solve the dynamic maximization problem, it is convenient to view it as a T -stage decision process illustrated in Figure 1.




Figure 1: T Stage Decision Process

At any stage t , we consider two types of actions for an agent: spending all of the remaining budget ($\alpha = 1$) and fraction (or none) of budget ($0 \leq \alpha_t < 1$). Particularly, as long as one chooses $\alpha_t = 1$, the decision process is terminated and the utility can be obtained in terms of the number of final users. Otherwise, one should proceed until the budget is exhausted. Note that an agent will always spend whatever budget is left at $t = T - 1$ (since there is no utility from keeping a fraction of it intact thereafter).

In principle, one can solve the dynamic influence maximization problem using backward induction. For a T -stage decision problem, backward induction starts with considering optimal decisions at the last stage (i.e. $t = T - 1$) for all possible realizations of D and B , then considers optimal choices at the second-to-last stage (i.e. $t = T - 2$) using the optimal decisions in the final stage, and so on. The process proceeds until the very first decision node (i.e. $t = 0$) is reached, and finally an optimal policy can be returned. However, such an approach quickly becomes intractable when the evaluation of f is very time consuming (for example, as in the instance described below, f corresponds to an agent-based simulation of individual adoption decisions). For example, if we suppose that each simulation takes only 1 second, our population involves 10,000 individuals, the budget can seed only 20 individuals, and our time horizon is 20 stages (e.g., 20 months), dynamic programming would take over 1000 hours, or about 1.5 months. Our primary goal is to develop algorithmic approaches which are orders of magnitude faster.

4.1 Optimal Algorithm

To start, consider Algorithm 1 (called *Best-Stage*), which simply finds a single best stage t at which to use up all of the budget accrued by this stage. Our main result in this section is that the *Best-Stage* algorithm is optimal.

THEOREM 4.1. *Let t be the stage returned by the Best-Stage algorithm. Then the policy $\pi = (0, \dots, 0, B^0 \delta^t, \dots)$ where the only non-zero entry is at time t , is optimal.¹*

The proof of this result is rather involved, and we develop it in a series of steps.

¹Note that the nature of the policy after time t is irrelevant.

Data: D^0, B^0

Result: t

$u^* = 0;$

$t^* = 0;$

for $t=0..T-1$ **do**

$u = \text{FindAdoptions}(D^0, B^0 \delta^t, t);$

if $u > u^*$ **then**

$u^* = u;$

$t^* = t;$

end

end

Algorithm 1: The *Best-Stage* algorithm. The function $\text{FindAdoptions}(D, B, t)$ computes the total number of adoptions at time T when there are D adopters at time 0 and the entire budget B is used for seeding at time t .

Let us consider a general backward induction step from the descendant nodes at stage $t + 1$ to their parent node at stage t . Suppose at stage t there are D^t users and B^t budget. Particularly, if the entire budget is used at t , an agent will have utility

$$U(\alpha_t = 1) = f^{(T-t)}(D^t + B^t) \quad (1)$$

where notation $f^{(m)}(\cdot)$ stands for applying f repeatedly for m stages. If some fraction $0 \leq \alpha_t < 1$ of the budget is used at stage t , there will be

$$D^{t+1} = f(D^t + \alpha_t B^t) \quad (2)$$

adopters and

$$B^{t+1} = (1 + \delta)(1 - \alpha_t)B^t \quad (3)$$

available budget in the next stage $t + 1$. Let π^* be an optimal policy, and α_t^* the fraction of budget allocated to seeding under π^* in stage t . Our next series of results characterize the properties that π^* (and fractions α_t^*) must have, by virtue of its optimality.

LEMMA 4.2. *Suppose that at some stage t the optimal policy has $\alpha_t^* = 0$. Further, suppose that there is some $t^* > t$ with $\alpha_{t^*}^* = 1$ and $\alpha_{t'}^* = 0$ for all $t < t' < t^*$. Then $\alpha_{t-1}^* = 1$ cannot be optimal.*

PROOF. Let $\alpha_{t-1}^* = a$, and let D^t and B^t be the number of adopters and remaining budget at time t if the optimal policy π^* is followed. Notice that there are $T - t$ stages remaining when we start at stage t . Thus, the utility in stage $t - 1$ is given by

$$U(\alpha_{t-1}^* = a, \alpha_t^* = 0) = f^{(T-t^*)}(f^{(t^*-t)}(D^t) + (1 + \delta)^{t^*-t} B^t),$$

where $D^t = f(D^{t-1} + aB^{t-1})$, $B^t = (1 + \delta)(1 - a)B^{t-1}$ and $t^* > t$ is the stage of the optimal policy π^* after t at which the entire budget is used up. The lemma states that if $\alpha_t^* = 0$, then $\alpha_{t-1}^* = 1$ must not be. In formal notation, we can state this as $U(\alpha_{t-1}^* = a, \alpha_t^* = 0) > U(\alpha_{t-1}^* = 1)$ for $a < 1$. By monotonicity of f this is equivalent to

$$f^{(t^*-t)}(D^t) + (1 + \delta)^{t^*-t} B^t > f^{(t^*-t+1)}(D^{t-1} + B^{t-1}). \quad (4)$$

By optimality of $\alpha_t^* = 0$, we further have

$$f^{(T-t^*)}(f^{(t^*-t)}(D^t) + (1 + \delta)^{t^*-t} B^t) \geq f^{(T-t)}(D^t + B^t),$$

which by monotonicity of f is equivalent to

$$f^{(t^*-t)}(D^t) + (1+\delta)^{t^*-t} B^t \geq f^{(t^*-t)}(D^t + B^t). \quad (5)$$

If we could show that

$$f^{(t^*-t)}(D^t + B^t) > f^{(t^*-t+1)}(D^t + B^t)$$

or, equivalently,

$$D^t + B^t > f(D^{t-1} + B^{t-1}),$$

then combining this with (5) will imply that (4) must hold. Let us rewrite (5) as follows

$$\frac{f^{(t^*-t)}(D^t + B^t) - f^{(t^*-t)}(D^t)}{B^t} \leq (1+\delta)^{t^*-t}, \quad (6)$$

or, equivalently, as

$$\begin{aligned} & \frac{f^{(t^*-t)}(D^t + B^t) - f^{(t^*-t)}(D^t)}{f^{(t^*-t-1)}(D^t + B^t) - f^{(t^*-t-1)}(D^t)} \\ & \times \dots \times \frac{f(D^t + B^t) - f(D^t)}{B^t} \leq (1+\delta)^{t^*-t} \end{aligned}$$

Because $f^{(t^*-t)}(D^t) > \dots > f(D^t)$ and by strict convexity, it follows that

$$\begin{aligned} & \frac{f^{(t^*-t)}(D^t + B^t) - f^{(t^*-t)}(D^t)}{f^{(t^*-t-1)}(D^t + B^t) - f^{(t^*-t-1)}(D^t)} \\ & > \dots > \frac{f(D^t + B^t) - f(D^t)}{B^t}, \end{aligned}$$

which in turn implies that

$$\frac{f(D^t + B^t) - f(D^t)}{B^t} < 1 + \delta. \quad (7)$$

Additionally, observe that

$$D^t = f(D^t + aB^{t-1}) > D^t + aB^{t-1}$$

and

$$B^t = (1+\delta)(1-a)B^{t-1} > (1-a)B^{t-1}$$

By strict convexity, it then follows that

$$\frac{f(D^{t-1} + B^{t-1}) - f(D^{t-1} + aB^{t-1})}{(1-a)B^{t-1}} < \frac{f(D^t + B^t) - f(D^t)}{B^t}$$

which together with (7) implies that

$$\frac{f(D^{t-1} + B^{t-1}) - f(D^{t-1} + aB^{t-1})}{(1-a)B^{t-1}} < 1 + \delta,$$

which is equivalent to

$$\begin{aligned} f(D^{t-1} + B^{t-1}) & < f(D^{t-1} + aB^{t-1}) + (1+\delta)(1-a)B^{t-1} \\ & = D^t + B^t, \end{aligned}$$

completing the proof. $\square$

The next lemma builds on Lemma 4.2 to significantly strengthen its result.

LEMMA 4.3. *Suppose that at some stage t the optimal policy has $\alpha_t^* = 0$. Further, suppose that there is some $t^* > t$ with $\alpha_{t^*}^* = 1$ and $\alpha_{t'}^* = 0$ for all $t < t' < t^*$. Then $\alpha_{t-1}^* = 0$.*

PROOF. In this lemma, we will show that it cannot be the case that $\alpha_{t-1}^* \in (0, 1)$. Together with Lemma 4.2, it will imply the desired result.

We prove this lemma by contradiction. Suppose that the optimal choice is $\alpha_{t-1}^* = a$ with $0 < a < 1$. The optimal utility is then given by

$$\begin{aligned} U(\alpha_{t-1}^* = a, \alpha_t^* = 0) \\ = f^{(T-t^*)}(f^{(t^*-t)}(D^t) + (1+\delta)^{t^*-t} B^t), \end{aligned}$$

where t^* is the stage after t at which the entire budget is spent and

$$D^t = f(D^{t-1} + aB^{t-1}), B^t = (1+\delta)(1-a)B^{t-1}.$$

By optimality of π^* , this policy should be (weakly) better than a policy which spends nothing at stage $t-1$ and spends the entire budget in stage t^* (with the same t^* as above). The utility of such a policy is given by

$$U(\alpha_{t-1}^* = 0, \alpha_t^* = 0) = f^{(T-t^*)}(f^{(t^*-t)}(\bar{D}^t) + (1+\delta)^{t^*-t} \bar{B}^t),$$

where

$$\bar{D}^t = f(D^{t-1}), \bar{B}^t = (1+\delta)B^{t-1}.$$

Since we assume $U(\alpha_{t-1}^* = 0, \alpha_t^* = 0) \leq U(\alpha_{t-1}^* = a, \alpha_t^* = 0)$,

$$\begin{aligned} & f^{(t^*-t)}(f(D^{t-1})) + (1+\delta)^{t^*-t+1} B^{t-1} \leq \\ & f^{(t^*-t)}(f(D^{t-1} + aB^{t-1})) + (1+\delta)^{t^*-t+1} (1-a)B^{t-1} \end{aligned}$$

or

$$\begin{aligned} & f^{(t^*-t+1)}(D^{t-1} + aB^{t-1}) - f^{(t^*-t+1)}(D^{t-1}) \\ & \geq (1+\delta)^{t^*-t+1} aB^{t-1}, \end{aligned}$$

which we can rewrite as

$$\frac{f^{(t^*-t+1)}(D^{t-1} + aB^{t-1}) - f^{(t^*-t+1)}(D^{t-1})}{aB^{t-1}} \geq (1+\delta)^{t^*-t+1}.$$

On the other hand, by strict convexity of f

$$\begin{aligned} & \frac{f^{(t^*-t+1)}(D^{t-1} + aB^{t-1}) - f^{(t^*-t+1)}(D^{t-1})}{aB^{t-1}} \\ & < \frac{f^{(t^*-t+1)}(D^{t-1} + B^{t-1}) - f^{(t^*-t+1)}(D^{t-1} + aB^{t-1})}{(1-a)B^{t-1}}. \end{aligned}$$

Moreover, by the same argument as used in the proof of Lemma 4.2 to arrive at (6), the optimal choice a must satisfy

$$\begin{aligned} & \frac{f^{(t^*-t+1)}(D^{t-1} + B^{t-1}) - f^{(t^*-t+1)}(D^{t-1} + aB^{t-1})}{(1-a)B^{t-1}} \\ & \leq (1+\delta)^{t^*-t+1}, \end{aligned}$$

which implies that

$$\frac{f^{(t^*-t+1)}(D^{t-1} + aB^{t-1}) - f^{(t^*-t+1)}(D^{t-1})}{aB^{t-1}} < (1+\delta)^{t^*-t+1},$$

a contradiction. $\square$

LEMMA 4.4. *Suppose that $\alpha_t^* = 1$. Then $\alpha_{t-1}^* \in (0, 1)$ cannot be optimal.*

PROOF. The utility of choosing $\alpha_{t-1}^* = a$, where $a \in [0, 1)$ at stage $t-1$ is given by $U(\alpha_{t-1}^* = a, \alpha_t^* = 1) = f^{(T-t)}(D^t + B^t)$, where $D^t = f(D^{t-1} + aB^{t-1})$ and $B^t = (1+\delta)(1-a)B^{t-1}$.

$a)B^{t-1}$. Suppose that $a \neq 1$. Then $U(a = 1) \leq U(0 \leq a < 1, \alpha_t^* = 1)$, which means that

$$\begin{aligned} f^{(T-t+1)}(D^{t-1} + B^{t-1}) \\ \leq f^{(T-t)}(f(D^{t-1} + aB^{t-1}) + (1+\delta)(1-a)B^{t-1}), \end{aligned}$$

or, by monotonicity,

$$f(D^{t-1} + B^{t-1}) \leq f(D^{t-1} + aB^{t-1}) + (1+\delta)(1-a)B^{t-1}.$$

Rearranging, we obtain

$$\frac{f(D^{t-1} + B^{t-1}) - f(D^{t-1} + aB^{t-1})}{(1-a)B^{t-1}} \leq 1 + \delta. \quad (8)$$

On the other hand, the presumption of optimality of $\alpha_t^* = 1$ implies that $\alpha_t = 1$ is a (weakly) better choice than $\alpha_t = 0$. In other words, spending all budget at stage t is nevertheless better than saving it and spending at any other stages after t . This implies that for any $\tau > t$,

$$f^{(T-\tau)}(f^{(\tau-t)}(D^t) + (1+\delta)^{\tau-t}B^t) \leq f^{(T-t)}(D^t + B^t).$$

By monotonicity this is equivalent to

$$f^{(\tau-t)}(D^t) + (1+\delta)^{\tau-t}B^t \leq f^{(\tau-t)}(D^t + B^t),$$

which we can rewrite as

$$\frac{f^{(\tau-t)}(D^t + B^t) - f^{(\tau-t)}(D^t)}{B^t} \geq (1+\delta)^{\tau-t}.$$

In particular, if we let $\tau = t + 1$, we have

$$\frac{f(D^t + B^t) - f(D^t)}{B^t} \geq 1 + \delta.$$

By strict convexity and assumption (3) on f , it follows that

$$\frac{f(D^t + B^t) - f(D^t)}{B^t} > \frac{f(D^{t-1} + B^{t-1}) - f(D^{t-1} + aB^{t-1})}{(1-a)B^{t-1}}.$$

We now show that if $a \neq 1$, it must be the case that $a = 0$. Notice that if $U(a = 0, \alpha_t = 1) \leq U(0 < a < 1, \alpha_t^* = 1)$, it must be that

$$\begin{aligned} f^{(T-t)}(f(D^{t-1}) + (1+\delta)B^{t-1}) \\ \leq f^{(T-t)}(f(D^{t-1} + aB^{t-1}) + (1+\delta)(1-a)B^{t-1}), \end{aligned}$$

or, equivalently,

$$f(D^{t-1}) + (1+\delta)B^{t-1} \leq f(D^{t-1} + aB^{t-1}) + (1+\delta)(1-a)B^{t-1},$$

which can be written as

$$\frac{f(D^{t-1} + aB^{t-1}) - f(D^{t-1})}{aB^{t-1}} \geq 1 + \delta. \quad (9)$$

By strict convexity, we have

$$\begin{aligned} \frac{f(D^{t-1} + aB^{t-1}) - f(D^{t-1})}{aB^{t-1}} \\ < \frac{f(D^{t-1} + B^{t-1}) - f(D^{t-1} + aB^{t-1})}{(1-a)B^{t-1}}, \end{aligned}$$

which, together with (8) and (9) implies that $1 + \delta < 1 + \delta$, a contradiction. Consequently, either $a = 0$ or $a = 1$. $\square$

Armed with Lemmas 4.2-4.4, we are now ready to prove our main result.

PROOF OF THEOREM 4.1. We prove the theorem by induction.

Base Case: $T = 1$ For an 1-stage decision problem, it is clearly optimal to spend the entire budget budget, so the optimal policy is $\alpha_0^* = 1$.

Inductive Step: Suppose the argument holds for $T = K \geq 1$. That is for any K -stage decision problem, the optimal policy is using all budget at some stage t^* , s.t., $t^* \in 0, 1, \dots, K-1$. Let us now consider a $T = K+1$ -stage decision problem. Assume we are at the final backward induction step from nodes in stage $t = 1$ the root node (i.e. $t = 0$). Note that those different decision nodes at stage $t = 1$ are results from different values of α_0 picked at $t = 0$. Generally, for any node (D^1, B^1) in stage $t = 1$, by the inductive assumption, an optimal policy must spend the entire budget at some stage between stage 1 and $T-1$; in other words, starting in stage $t = 1$, *Best-Stage* is optimal. Such a policy can only be one of the following three types:

1. $\alpha_0^* = 1$: spend the entire budget at $t = 0$;
2. $\alpha_0^* = 0$: spend all budget at a single stage $t^* > 0$;
3. $0 < \alpha_0^* < 1$: spend only a fraction of budget at stage $t = 0$, and the rest at a single stage $t^* > 0$.

Clearly, we only need to rule out the third type.

Case 1: $\alpha_1^* = 0$ By Lemma 4.3, it must be the case that $\alpha_0^* = 0$.

Case 2: $\alpha_1^* = 1$ By Lemma 4.4, it must be the case that either $\alpha_0^* = 0$ or $\alpha_0^* = 1$. In either case, the optimal policy has the form of *Best-Stage*. $\square$

The *Best-Stage* algorithm is clearly far faster than dynamic programming, with running time $O(T)$, compared to $O(DBT)$ for the former, where D is the size of the population, B the maximum budget, and T the time horizon. In our example above, it will take only 20 seconds, as compared to 1.5 months!

4.2 A Heuristic Search Algorithm

The model we described above is clearly stylized. As with any stylized model, a natural question is how far its assumptions can be relaxed without losing the guarantees that come with them—in our case, a very simple and provably optimal algorithm for dynamic influence maximization. If we are to undertake such an analysis experimentally, it is simply not feasible to use dynamic programming as a baseline, for reasons made clear earlier. As an alternative, we propose a heuristic algorithm which generalizes the *Best-Stage* algorithm, but does not incur too prohibitive a cost in terms of running time. Rather than choosing a single best stage, this algorithm, which we term *Best-K-Stages* (see Algorithm 2) iterates over integers K between 1 and $T-1$, splitting the budget equally among K consecutive time steps. This algorithm only takes $O(T^2)$, which is still quite manageable (about 200 seconds in our example above).

5. EXPERIMENTS

To verify the efficacy of our proposed algorithm, we utilize an agent-based simulation introduced by [20] which forecasts rooftop solar adoption in a representative zip code in San Diego county. This simulation has several features of relevance to us. First, it utilizes agent models which were learned from actual adoption, as well as auxiliary data, using

```

 $U \leftarrow -\infty;$ 
 $\pi^* \leftarrow \emptyset;$ 
for  $i=1, \dots, T-1$  do
  for  $j=0, \dots, T-i-1$  do
     $\pi \leftarrow (\alpha^j, \dots, \alpha^{j+i-1});$ 
    if  $U(M, \pi, B^0, \delta) > U$  then
       $U \leftarrow \max\{U, U(M, \pi, B^0, \delta)\};$ 
       $\pi^* \leftarrow \pi;$ 
    end
  end
end
return  $\pi^*$ 

```

Algorithm 2: The *Best-K-Stages* Algorithm.

a logistic regression model. Second, the model includes *network effects*: aggregate adoption is one of the variables that has a significant impact on individual adoption decisions. Third, the solar market is still very much in its developing stages, even in San Diego county, and consequently the relevant region of the logistic regression model is convex in the network effects variable. Fourth, while in the policy context of this model the budget remains fixed over the T -stage time horizon, costs decrease over time. In particular, in the previously validated version of the model [20], costs actually decrease as a function of “learning-by-doing”, modeled as aggregate adoption in San Diego county (which in turn increases over time). Moreover, the costs decrease linearly, rather than exponentially. As a result, the cost in this simulation does not satisfy the assumptions that guarantee the optimality of the *Best-Stage* algorithm.

An additional deviation from the idealized model we consider above is that the system cost enters the consideration in two ways: first, decreasing cost translates into an effectively increasing budget, and second, cost enters the adoption model directly as a part of *economic* considerations in rooftop solar adoption. A final important imperfection is that the budget need not be perfectly divisible by cost. In this section, we systematically investigate the extent to which these deviations from the “ideal” modeled above impact the optimality of the *Best-Stage* algorithm.

Throughout, time is discretized in months, and we fix our time horizon at 24 months (that is, 24 stages). We use the *Best-K-Stages* as our benchmark, observing in each case below expected adoption at horizon T as a function of K . Thus, when we find $K = 1$ yields an optimal or a near-optimal solution, we conclude that the deviation from model assumptions is not significant, while optimal $K > 1$ suggests the converse. We used the budget allocated by the California Solar Initiative (CSI) to the San Diego county incentive program as our baseline, and consider amplifications of this budget, denoted by “BX”. For example, 50X budget means that 50 times the CSI budget was used in an experiment. Moreover, we explored the impact of the magnitude of the peer effect variable in the model by considering as a baseline the network effect coefficient produced by learning the adoption model from data, as well as its amplifications, denoted by “PX”; for example “2X” network effect means that we doubled the network effect coefficient.


5.1 Exponential Cost

Our hypothesis is that the primary consideration is the nature of the cost model, with the remaining “imperfections” introduced by the complex considerations of the agent-based model in question being significantly less important. To in-

vestigate, we replace the cost model in the original model with a much simpler cost function which decreases exponentially with time t (equivalently, the budget is exponentially increasing over time): $C(t) = C_0 e^{\omega t}$, or, equivalently, $\log(C) = \log(C_0) + \omega t$. We used the rooftop solar cost data for San Diego county to estimate the parameters $\log(C_0)$ and ω in this model (see Table 1). The new cost model was then “injected” into the otherwise unmodified agent-based model.


Table 1: Exponential Cost Model ($R^2 = 0.020$)

Parameter	Coefficient
$\log(C_0)$	10.55
ω	-0.0059



(a) Different Lengths (K values). (b) Different Seeding Month.


Figure 2: Exponential Cost: 50X Budget, 1X network effects.



(a) Different Lengths (K values). (b) Different Seeding Month

Figure 3: Exponential Cost: 50X Budget, 1.75X network effects.

The maximum expected adoption of seeding policies over different lengths (values of K in *Best-K-Stages*) is illustrated in Figure 2a. The results suggest that seeding in a single stage tends to be a better policy than splitting the budget over multiple stages. Moreover, for length-1 policies we find that in this case seeding at very end (i.e., in stage $T - 1$) is optimal, as shown in Figure 2b, where the horizontal axis is the seeding month. This is likely because benefit of the exponential cost decay in this variation of the model exceeds the gains from seeding early due to network effects. As the importance of network effects increases, we expect that seeding earlier would become more beneficial. To investigate, we manually varied the coefficient of network effects in the adoption model, multiplying it by a factor of 1.75 and 2, and comparing the outcomes. Seeding in a single stage is still quite effective (Figure 3a and 4a; the jagged nature of the



(a) Different Lengths (K values). (b) Different Seeding Month
Figure 4: Exponential Cost: 50X Budget, 2X network effects.

plots is likely due to the indivisibilities discussed above), but the optimal month to seed shifts earlier, as shown in Figure 3b and 4b.


5.2 Original Agent-Based Model

As shown in the previous section, the *Best-Stage* algorithm performs impressively despite a number of imperfections in the agent-based model of aggregate rooftop solar adoption. Next, we test its effectiveness in the context of an “ideal” model that was previously thoroughly validated in forecasting solar adoption. Significantly, the cost function used in this model includes a number of relevant parameters (such as system size), and in place of explicit dependence on time, it is a decreasing function of the overall uptake of solar systems in San Diego county (see Table 2). Finally, the cost function is modeled as linear in its parameters (with a fixed lower bound at zero).

Table 2: Cost function in the original agent-based model ($R^2 = 0.8399$).

Predictor	Coefficient
(Intercept)	11,400
Home Value	7.38e-04
Living Square Feet	0.15
System Capacity	6,210
San Diego County Adoption	-1.06

The results in Figure 5a are surprising: the *Best-Stage* algorithm performs *significantly* worse than policies that split budget over a relatively long contiguous series of months ($K > 5$). Figure 5b further reveals that we optimally wish to push the initial month of this contiguous sequence rather late in the period; in other words, network effects are relatively weak. The key question is: what goes wrong? The observations in the previous section strongly suggest that it is the form of the cost function that is at the heart of the issue. We therefore proceed to carefully investigate what, precisely, about the nature of the cost function in this model is the cause of this qualitative change relative to our stylized model in Section 3. Specifically, we start with a simplified model of solar adoption that conforms to the assumptions of our main result (i.e., using only time as a variable), and incrementally relax it to bring it closer to the cost function actually used in the original simulation environment. In particular, we begin with a polynomial cost function, proceed to investigate a linear cost model (in time only), and finally consider a linear cost function that depends on aggregate product uptake (learning-by-doing) rather than time.



(a) Different Lengths (K values). (b) Different Start Month
Figure 5: Actual Cost Learning-by-doing: 10X Budget, 1X network effects.


5.3 Polynomial Cost

The exponential cost function, or exponentially increasing budget, seems quite dramatic. What if we slow this growth in budget buying power to be polynomial?

We formulate the following simple model of polynomial cost: $C(t) = C_0 t^\theta$, where t is time variable, and C_0 and θ are parameters. This function is equivalent to $\log(C(t)) = \log(C_0) + \theta \log(t)$, which we can fit using linear regression. The resulting coefficients are given in Table 3.

Table 3: Polynomial Cost Model ($R^2 = 0.014$)

Parameter	Coefficient
$\log(C_0)$	10.70
$\log(t)$	-0.098




(a) 0.5X network effects. (b) 1X network effects.

Figure 6: Polynomial Cost: 50X Budget.

Figures 6a, 6b, and 7a suggest that *Best-Stage* is still a very good policy here, albeit the jagged nature of the plots (likely due to indivisibilities) makes this observation somewhat equivocal when network effects are very weak. However, as the magnitude of network effects increases, the advantage of *Best-Stage* over alternatives becomes more pronounced. On balance, it seems clear that the polynomial vs. exponential nature of the cost function does not give rise to a qualitative difference in the effectiveness of our underlying model.

5.4 Linear Cost

Given that the polynomial cost model did not appear to bring about a substantial difference, we proceed to relax to a linear cost model, inching even closer to the “ideal” model used in the simulation. Our linear cost model has linear dependence on time, implying a slower decay rate than the polynomial cost function: $C(t) = a + bt$. As before, parameters a and b are estimated using solar system cost




(a) Different Lengths (K values). (b) Different Seeding Month.
Figure 7: Polynomial Cost: 50X Budget, 2X network effects.

data for San Diego county, with the results given in Table 4. As before, we ran the experiments by “plugging in” this

Table 4: Linear Cost Model ($R^2 = 0.012$)

Parameter	Coefficient
Intercept	42,053
t	-201

cost model into the simulation environment (retraining the individual adoption propensities). Our experiments show that seeding in a single month is, again, more effective than seeding in multiple consecutive stages (See Figure 8a and 8b), even as we vary the importance of network effects.



(a) 1X network effects. (b) 2X network effects.


Figure 8: Linear Cost: 50X Budget.

5.5 Linear Cost with Learning-by-Doing

Both “ideal” cost model and linear cost model are linear in their features. The key difference is that the ideal cost function depends on cumulative adoption, whereas the latter only depends on time. We now move yet closer to the ideal model, replacing the linear dependence on time with linear dependence on aggregate solar system uptake in San Diego county: $C(t) = c + dy(t)$, where $y(t)$ is number of solar adoption up to time t . The parameters c and d are estimated via linear regression, and are given in Table 5. Remarkably, the results now echo our observations in the “ideal” model (Figures 9a and 9b): *Best-Stage* is decidedly suboptimal, and policies that split the budget among $K = 5$ or more consecutive stages perform significantly better. Additionally, we can see that the “optimal” number of stages to seed (at least in our heuristic algorithm) increases as the magnitude of network effects increases (Figure 9a and 9b). The key distinction from the idealized model is that learning-by-doing makes the temporal benefit of waiting endogenous: now seeding earlier will directly reduce future seeding costs and, consequently, the effectiveness of residual budget. As a

Table 5: Linear Cost Model with Learning-by-Doing ($R^2 = 0.013$)

Parameter	Coefficient
(Intercept)	40,356
y	-1.74



(a) 1X network effects.

(b) 2X network effects.

Figure 9: Linear Cost with Learning-by-Doing: 10X Budget

result, we observe what amounts to an “interior” optimum in budget allocation, with some of the budget used in seeding in order to make residual budget more valuable later.

6. CONCLUSION

We formulate a novel dynamic influence maximization problem under increasing returns to scale and prove that the optimal policy must use up budget at a single stage, giving rise to a simple *Best-Stage* search algorithm. In addition, we propose a heuristic algorithm, *Best-K-Stages*, that includes *Best-Stage* as a special case. We experimentally verify that the proposed *Best-Stage* algorithm remains quite effective even as we relax the assumptions to different time-involved cost dynamics, i.e., polynomial and linear cost. On the other hand, we find that when we replace the time dependency of the cost function by cumulative adoption (learning-by-doing), *Best-K-Stages* significantly outperforms *Best-Stage*.

Looking forward, there are several possible directions we would like to pursue in the future. First, it is clear that there must exist even better policies remaining unexplored for the “ideal” cost model with learning-by-doing. Our heuristic search algorithm only covers a special subset of policies. Design of more efficient algorithms to find optimal solutions in such a realistic setting can be a meaningful extension to our current work. Second, the dynamic influence maximization problem proposed in this paper assumes seeding does not discriminate among individuals. This is a very strong assumption, but enables us to make significant progress. Seeding in a social network with heterogeneous individuals has been shown to be *NP*-hard even for “one-shot” decisions and a simple submodular diffusion model [10]. A relaxation to individual heterogeneity is sure to create further algorithmic challenges.

7. ACKNOWLEDGMENTS

This work was partially supported by the U.S. Department of Energy (DOE) office of Energy Efficiency and Renewable Energy, under the Solar Energy Evolution and Diffusion Studies (SEEDS) program. It was also partially supported by the National Science Foundation (NSF) under grants CCF-1215883 and IIS-1350598.

REFERENCES

- [1] Kenneth J. Arrow. The economic implications of learning by doing. *Review of Economic Studies*, 29(3):155–173, 1962.
- [2] Emmanuelle Auriol and Michel Benaim. Standardization in decentralized economies. *American Economic Review*, 90(3):550–570, 2000.
- [3] Frank M. Bass. A new product growth for model consumer durables. *Management Science*, 15(5):215–227, 1969.
- [4] L.M.B. Cabral. On the adoption of innovation with network externalities. *Mathematical Social Sciences*, 19:299–308, 1990.
- [5] Wei Chen, Yajun Wang, and Siyu Yang. Efficient influence maximization in social networks. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 199–208, 2009.
- [6] Jacomo Corbo and Yevgeniy Vorobeychik. The effects of quality and price on adoption dynamics of competing technologies. In *International Conference on Information Systems*, 2009.
- [7] Joseph Farrell and Garth Saloner. Standardization, compatibility, and innovation. *Rand Journal of Economics*, 16(1):70–83, 1985.
- [8] Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42:427–486, 2011.
- [9] M. Granovetter. Threshold models of collective behavior. *American Journal of Sociology*, 83(6):1420–1443, 1978.
- [10] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 137–146, 2003.
- [11] Jon Kleinberg. Cascading behavior in networks: Algorithmic and economic issues. In Noam Nissan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 24, pages 613–632. Cambridge University Press, 2007.
- [12] Dong Li, Zhi-Ming Xu, Nilanjan Chakraborty, Anika Gupta, Katia Sycara, and Sheng Li. Polarity related influence maximization in signed social networks. *PLoS ONE*, 9(7):e102199, 2014.
- [13] Bo Liu, Gao Cong, Dong Xu, and Yifeng Zhang. Time constrained influence maximization in social networks. In *International Conference on Data Mining*, pages 439–448, 2012.
- [14] A. McDonald and L. Schrattenholzer. Learning rates for energy technologies. *Energy Policy*, 29(4):255–261, 2001.
- [15] S. Morris. Contagion. *Review of Economic Studies*, 67, 2000.
- [16] G. Nemhauser, L. Wolsey, and M. Fisher. An analysis of the approximations for maximizing submodular set functions. *Mathematical Programming*, 14:265–294, 1978.
- [17] M. Richardson and P. Domingos. Mining knowledge-sharing sites for viral marketing. In *International Conference on Knowledge Discovery and Data Mining*, 2002.
- [18] Arthur van Benthem, Kenneth Gillingham, and James Sweeney. Learning-by-doing and the optimal solar policy in california. *Energy Journal*, 29(3):131–151, 2008.
- [19] Duncan Watts. A simple model of global cascades in random networks. *Proceedings of the National Academy of Sciences*, 99:5766–5771, 2002.
- [20] H. Zhang, Y. Vorobeychik, J. Letchford, and K. Lakkaraju. Predicting rooftop solar adoption using agent-based modeling. In *AAAI Fall Symposium on Energy Market Prediction*, 2014.
- [21] Huiyuan Zhang, Thang N. Dinh, and My T. Thai. Maximizing the spread of positive influence in online social networks. In *International Conference on Distributed Computing Systems*, pages 317–326, 2013.