

Understanding Cyclic Trends in Social Choices

Anish Das Sarma
Google Research
anish.dassarma@gmail.com

Sreenivas Gollapudi
Microsoft Research
sreenig@microsoft.com

Rina Panigrahy
Microsoft Research
rina@microsoft.com

Li Zhang
Microsoft Research
lzha@microsoft.com

ABSTRACT

Motivated by trends in popularity of products, we present a formal model for studying trends in our choice of products in terms of three parameters: (1) their innate utility; (2) individual boredom associated with repeated usage of an item; and (3) social influences associated with the preferences from other people. Different from previous work, in this paper we introduce boredom to explain the cyclic pattern in individual and social choices. We formally model boredom and show that a rational individual would make cyclic choices when considering the boredom factor. Furthermore, we extend the model to social choices by showing that a society that votes for a particular style or product can be viewed as a single individual cycling through different choices.

We adopt a natural model of utility an individual derives from using an item, i.e., the utility of an item gets discounted by its repeated use and increases when the item is not used. We address the problem of optimally choosing items for usage, so as to maximize overall user satisfaction over a period of time. First we show that the simple greedy heuristic of always choosing the item with the maximum current composite utility can be arbitrarily worse than the optimal. Second, we prove that even with just a single individual, determining the optimal strategy for choosing items is NP-hard. Third, we show that a simple modification to the greedy algorithm is a provably close approximation to the optimal strategy. Finally, we present an experimental study over real-world data collected from query logs to compare our algorithms.

Categories and Subject Descriptors

F.0 [Theory of Computation]: General

General Terms

Algorithms, Theory, Human Factors, Economics

Keywords

fashion trends, cyclic trends, social choice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WSDM'12, February 8–12, 2012, Seattle, Washington, USA.
Copyright 2012 ACM 978-1-4503-0747-5/12/02 ...\$10.00.

1. INTRODUCTION

When an individual or a society is repeatedly presented with multiple substitutable choices, we often observe a recurring shift of user preferences over time. Many studies have observed that the choices of a society tend to be cyclic [11, 12, 5, 4]. While some trends are relatively easy to explain (e.g., sweater sales increasing in the winter), some other trends may result from a variety of factors, including social influence, novelty of an item, and its innate value to a user.

In this paper, we describe a utility model that combines different factors that may result in cyclic choices. We study the computational issues under the model and provide simple mechanisms by which consumers may make *close to optimal* decisions on which products to consume and when, in order to maximize their *overall utility*. In contrast, we give examples in which the (natural) greedy algorithm that maximizes short term utility can be much worse in long run. Further, we show that under a linear social influence model, the society can be viewed as an individual, and we can apply the algorithm designed for an individual to achieve nearly optimal social welfare. Therefore, our algorithm can be useful for a company or website in deciding which items to show to the users in order to maximize the total user satisfaction in long term. We also conduct experiments to show how various parameters in our model can be estimated as well as validate our algorithms.

In our model, a user lives in discrete time periods and at each time period chooses to consume one from a set of *items*, representing any product, good, or concept the user consumes repeatedly. We consider three broad categories of factors that may impact the user's utility of choosing an item. The first, and perhaps the most basic, cause of a product becoming trendy is its innate utility, intuitively capturing the value it gives an individual. Second, psychologically, a person's utility of consuming a product may be discounted by repeated consumption of the same item. This results in the user choosing new and different products as she gets tired of existing products. Third, while at an individual level, we have certain inclinations based on our tastes, these are influenced by social phenomena, such as friends' or celebrities' preferences. Our utility model is based on the aforementioned factors.

Innate utility: The utility of an item captures the innate value the item provides to an individual. We assume it is fixed, independent of other influences.

Individual boredom: If we use any item for too long, we

get bored of it, and our appreciation for it goes down. This is modeled as a negative component added to the utility. This factor grows if one repeatedly uses the same item and fades away when one stops consuming the item.

To model boredom on any item at any given time t , we associate with each past usage of the item, say at time t' , a factor in the form of $(1 - r)^{t-t'}$ for some $r < 1$. Then the total boredom on the item takes the sum of this factor from all the past usage of the item. This definition captures the intuition that the boredom grows if an item is repeatedly used.

Utility maximization with boredom turns out to be hard — as we show in the paper, it is NP-hard to determine the optimal consumptions under the boredom model. One natural heuristic is the greedy algorithm in which the user, at each time t , always chooses the item with the maximum $u_i(t)$, based on the consumption history of the user. We construct an example to show that the natural greedy algorithm may result in overall utility far less than the optimum. On the other hand, we show a slightly adapted “greedy” algorithm can lead to close to the optimal utility. In the proposed algorithm, in addition to the past history, the user also factors in the boredom effect the item might have on the future consumption and calculates a “shadow” utility of each item by double the boredom. We show that by picking the item that maximize the “shadow” utility every time, the user can achieve close to optimal utility in long run when the boredom decay rate is small.

Social influence: Our valuation of an item can change significantly by the valuation of our friends or influencing people. For example, when we see that many people around us like something we may start liking it; or we may consciously want to differ from some other people around us. We model such influences as a weighted linear combination from other people.¹

The effect of social influence can be formalized using a linear model. For example, to model social influence, consider one item and a society consisting of m people. Let G denote the influence graph on these people that is directed where each edge is labeled with a weight that indicates the strength of this influence. A high value on an edge, such as outgoing edges from celebrities indicates a strong outgoing influence; on the other hand a negative value indicates a desire to distance oneself or be different from the source node. Let A denote the corresponding influence matrix. let $u_i(t)$ denote the utility of the item to the i^{th} person; let $\mathbf{u}(t)$ denote the vector of utilities. If we assume that for each time step the influence from all friends of a person add linearly then we may write $\mathbf{u}(t + 1) = A\mathbf{u}(t)$, which is similar to [11]. Note that for stability of this iterative powers, we should assume that its top eigenvector has magnitude 1. We will show that under this influence model, we may treat the society as an individual making choices under the effect of boredom.

From a broader standpoint, one of the main points our model makes is that trends in social choices can be viewed as not just the effect of the influence of a privileged few but more as a democratic process that churns the social boredom

¹Additional influence may come from the association of a product to things/concepts we like or dislike. For instance, someone may be very fond of green technologies or dislike things that are scary. We may simply model such concepts as individuals.




Figure 1: Cycles resulting from negative influences: Consider a system with two types of people: (1) followers, who always imitate the trendsetters, and (2) trendsetters who always attempt to distinguish themselves from the followers. Such a system will result in a cyclic preferences.

and channels the innate instinct for change. Boredom is the innate psychological force that dulls the effect of a constant stimulus over a period of time and make us look for newer stimuli. Thus the *coolness* of an item drops over time and things that we haven’t seen or used in a long time begin to appear more *cool*. Another point the paper makes is that several scenarios involving individual and social choices are essentially driven by the same underlying principles. Boredom is thus the single most and simplest explanation for oscillations in individual and social choices. This is not at all surprising; indeed boredom is perhaps a strong influence when we make choices on items such as food, clothes, and songs. Social influence no doubt plays a large part in individual choices. But when we look at the social system as a whole the influences across individuals are forces within the system and in the net effect it simply gives a larger voice to the more influential individuals. We also note that influence by itself is not sufficient to create cyclic behavior. In fact, if all influences are positive then without any boredom the system $u(t + 1) = Au(t)$ converges to a fixed value resulting in a fixed social choice.

Finally, we recognize that the factors we consider are by no means comprehensive; several other factors may change the values of nodes. Our decay model for boredom and linear model for influence may be too simplistic. Nonetheless, we believe boredom and social influence capture several important aspects of the underlying psychological processes that people use to value items.

Related work. Understanding social choices particularly in the context of fashion trends are of significant academic interest as well as commercial importance in various fields, including brand advertising and market economics. There’s a large body of work in multiple disciplines – sociology (e.g. [5, 4]), economics (e.g. [7]), and marketing (e.g. [11, 12]), on theories for evolution of fashion. Our work is inspired by [11] in which the authors propose a mathematical model for unifying many existing explanations for fashion cycles. We will describe [11] in more details and contrast the difference of our work.

Reference [11] proposed a formal model of fashion based on association rules. Intuitively, an individual’s utility for

an item is impacted by how similar it is to items he likes, and how dissimilar it is to items he dislikes. Further, he is influenced by the society through other individuals' preferences for various items. Consider a single item, whose consumption vector is given by $c(t)$ at time t . Considering the recurrence $c(t+1) = Wc(t)$, where W is the weight influence matrix, [11] observed that if the matrix W has a complex top eigenvalue (corresponding to negative influences), then the item's consumption pattern may be periodic, producing cycles in preferences (as shown in Figure 1). Our model of utility is similar to the consumption model in [11]. However, we consider an additional parameter of boredom that is essential to explain fashion cycles in a society with non-negative influences as such a matrix W always has a real top eigenvalue. The following argument shows that a matrix with non-negative influences cannot result in cycles.

OBSERVATION 1.1. *The process $c(t + 1) = Wc(t)$ converges to a fixed point if W has no non-negative entries and is eigen-decomposable where the magnitude of any eigenvalue is at most 1. That is, it cannot result in cyclic behavior where c takes different values repeatedly as $t \rightarrow \infty$*

PROOF. The basic idea is that if W is eigen decomposable, it can be written as $W = QDQ^{-1}$ where D is the diagonal matrix of eigenvalues and the columns of Q are the eigenvectors. Now $W^t = QD^tQ^{-1}$ which is dominated by the largest (in magnitude) eigenvalue. We already assumed the magnitudes are at most 1 for stability. If the top eigenvalue in magnitude is -1 (or complex) then we will see a behavior like $(-1)^t$ which cycles through the values $1, -1, 1, -1, \dots$. If the top eigen values are the k th roots of 1 then we will get cyclic behavior with periodicity k .

But, if W is non-negative, the Perron-Frobenius theorem [2] states that there is a unique dominant positive eigenvalue. Then $W^t = QD^tQ^{-1}$ converges to 0 if eigenvalue of top magnitude is less than 1 and to $q_1'q_1$ if the unique top eigenvalue is 1, where q_1 is the eigenvector corresponding to the top eigenvalue. \square

Some other recent work (e.g., [6]) study behavioral influences in social networks, such as in terms of information propagation. For instance, [8] studies how two competing products spread in society, [9] provides techniques for tracking and representing "memes", which may be used to analyze news cycles, and [10] studies how recommendations propagate in a network through social influence.

The focus of our paper is on formalizing a theory for social choices in the presence of boredom, combined with social influence. Therefore, for a large part of the paper we consider only a single individual and study cyclic trends based on boredom. In our extension to multiple individuals, we assume a linear weighting of influences from friends' preferences for particular items.

2. CONTRIBUTIONS OF OUR STUDY

2.1 Modeling individual boredom

We consider a user living in discrete time periods $0, 1, \dots$ and consuming one item among n substitutable items at each time; for example, a person needs to decide which restaurant to go to every night or which political party to vote for every four years. We assume that each item i brings

a base utility v_i to the user. Now if we assume that the utilities are fixed then the user would always choose the item with the maximum v_i . This would be inconsistent with the observed common behavior of cycling among multiple items. In order to explain cyclic choices, it is necessary to model the utility dependence of the consumptions across different time periods.

We propose a simple model in which the utility of an item at any time t is the base utility discounted by a boredom factor proportional to the "memory" the person has developed by using this item in the past. The more often or more recent the user has used the item, the more memory and boredom is developed for the item, and consequently the less utility the item has to the user.

We naturally assume that the memory drops geometrically over time, and the total memory of a person is bounded. This leads to the following definition of memory. Let $0 < r < 1$ be a memory decay rate, i.e., the rate at which a person "forgets" about things. Let $x_i(t) \in \{0, 1\}$ indicate if the user uses the item i at time t . Then the memory of i at time t is $M_i(t) = r \sum_{\tau=0}^{t-1} x_i(\tau)(1-r)^{t-\tau}$. We add the factor r so that $M_i(t) \leq 1$. This is equivalent to the recurrence $M_i(t+1) = (1-r)M_i(t) + rx_i(t)$. Note that r can be viewed as the decay rate because if an item is used once in the beginning and then never used then its memory after time t is $(1-r)^t \approx e^{-rt}$ (for small r).

Another important point to note is that a small decay rate is very different from a decay rate of 0. Look at the equation $M_i(t+1) = (1-r)M_i(t) + rx_i(t)$. If $r = 0$ then $M_i(t)$ is always 0 - there is no memory. If $r \rightarrow 0$, $M_i(t) = (1-r)M_i(t) + rx_i(t) = r \sum_{\tau=0}^{t-1} x_i(\tau)(1-r)^{t-\tau} = \frac{\sum_{\tau=0}^{t-1} x_i(\tau)(1-r)^{t-\tau}}{\sum_{\tau=0}^{t-1} (1-r)^{t-\tau}}$ which is like the weighted-average fraction of time the item was used where most of the weight is in the last $O(1/r)$ time steps. Thus if $r = 0$ memory is always 0, but if r is tiny then on continuous use of an item its memory will reach close to 1 in $O(1/r)$ time steps.

The boredom $b_i(t) = \alpha_i M_i(t)$ is proportional to the memory and depends on the item. The utility of item i is defined as $u_i(t) = v_i - b_i(t) = v_i - \alpha_i M_i(t)$. Henceforth, we will refer to v as the base utility and α as the boredom coefficient.

2.2 Utility optimization with boredom

With the above model, one natural question is to compute the choices of the items to maximize the user's overall utility. If we allow the user to choose at continuous time, the maximization problem becomes relatively easy as the best way to consume an item is to do it cyclically at regular time intervals. However, such regular placement may not be realizable in the discrete case or is hard to find. As we will show below, it is NP-hard to compute the best consumption sequence.

We also consider the natural greedy strategy and show that under the greedy strategy, the utility of each item is always bounded in a narrow band and so each item is consumed approximately cyclically. The greedy strategy, however, may have produce a sequence giving poor overall utility. We provide a simple heuristics, called double-greedy strategy, and show that it emulates the cyclic pattern of the optimal solution on the real line and yields utility close to the optimal when r is small.

2.2.1 Greedy algorithm

In the greedy strategy, at each time t , the user consumes the item with the maximum utility $u_i(t)$. This strategy is intuitive and probably consistent with how we make our daily decisions. We show that the utility gap between any two items is small all the time. We provide an example to show it has poor performance in terms of utility maximization, even when $r \rightarrow 0$. Denote by $\alpha = \max_i \alpha_i$.

THEOREM 2.1. *There exists a time T such that for any $t \geq T$, $\mu \leq \max_i u_i(t) \leq \mu + O(r\alpha \log n)$ where μ is the unique solution to the following system:*

There is an $f_i \geq 0$ for each item such that (1) $\sum f_i = 1$; (2) if $f_i > 0$, $v_i - f_i \alpha_i = \mu$, and $v_i < \mu$ otherwise.

Intuitively, f_i is the long run frequency of the item i gets consumed. While the greedy algorithm has the nice property of keeping the utility gap between any items small, it may produce a sequence with poor overall utility.

OBSERVATION 2.1. *The Greedy strategy of always picking the highest utility item each day is not optimal.*

To see the non-optimality of greedy, simply consider two items for beverage, say “water” and “soda”. Assume water has low innate utility say 1 and zero boredom coefficient. Soda on the other hand has high utility say 10 but also a high boredom coefficient say 10. So if one drank soda every day its utility would drop to below that of water. Observe that the greedy strategy will choose soda till its utility drops to that of water and then switch back to soda whenever its utility rises even slightly over 1. So the average utility of the greedy strategy is close to 1, as soda’s utility fluctuates around 1. A smarter strategy is to hold off on the soda even if it is a better choice today so as to enjoy it even more on a later day. Thus it is possible to derive an average utility that is much higher than 1. For example, we can get average utility of at least 3 by alternating between water and soda in the above example. Since the memory on an item is at most 1, the memory on an item is at most $1/2$ if it is consumed every other day. Hence the utility on soda is $10 - 10/2 = 5$ in the above strategy. It can be easily seen that the parameters in the above example can be picked to create arbitrarily large gap between the greedy algorithm and the optimal solution. We emphasize that this gap exists for any value of r .

This naturally raises the question: what is the optimal strategy? More importantly is there an optimal strategy that is a simple ‘rule of thumb’ that is easy to remember and employ as we make the daily choices. Unfortunately it turns out that computing the optimal strategy is NP-hard.

THEOREM 2.2. *Given a period T , target utility U^* , and n items, it is NP-hard to determine whether there exists a selection of items with period T such that the total utility of the selection is at least U^* .*

The proof is detailed in the appendix.

2.2.2 Double-greedy algorithm

On the positive side we show that there is indeed a simple “rule of thumb” that gives an almost optimal solution when r is small. The strategy “double-greedy” waits longer for items that we get bored of too quickly. It is a simple twist on the greedy strategy: instead of picking the item that maximizes the utility $u_i(t) = u_i - b_i(t)$, it picks the one which maximizes

$w_i(t) = u_i - 2b_i(t)$. Thus it doubles the boredom of all items and then runs the greedy strategy. We arrive at the double-greedy strategy by considering a continuous version of the problem. Intuitively, the double-greedy strategy tries to equalize the “marginal utility”, rather than the utility as in the greedy algorithm, each item brings. Indeed, we will bound the gap between marginal utilities and show that it leads to close to optimal utility when r is small.

THEOREM 2.3. *Let \bar{U} denote the average utility obtained by the double greedy algorithm and U^* the optimal utility. Then $\bar{U} \geq U^* - O(r\alpha \log n)$ where $\alpha = \max_i \alpha_i$.*

2.3 Fashion as a Social Choice

A choice is a fashion, if it is the choice of a large fraction of the society. Thus a society only supports a small number of fashions. Industries often target one type of fashion for each market segment. Consider a situation where the entire society consists of one fashion market segment. We will see how in this case such a society can be compared to an individual making choices to maximize utility under the effect of boredom. Each individuals utilities depend not only on his base utility and boredom but also on the influence from other individuals.

Consider a society of n people and m possible item choices. The society needs to choose one item out of these at every time step. We will study the problem of the making the optimal choice so as to maximize welfare. This is applicable in the following scenarios: A business is launching the next fashion style for its market segment, or a radio channel is broadcasting songs in a sequence to maximize the welfare to its audience. Let $u_{ij}(t+1)$ denote the utility of item i to person j at time t ; let $b_{ij}(t)$ denote the boredom value; let $\mathbf{u}_i(t)$ denote the vector of utilities to the n people for item i , $\mathbf{v}_i(t)$ denote the vector of base utilities, and $\mathbf{b}_i(t)$ denote the vector of boredom values. In the absence of boredom we will say $\mathbf{u}_i(t+1) = A\mathbf{u}_i(t)$ where A is the influence matrix. Accounting for boredom we will say, $\mathbf{u}_i(t+1) = A\mathbf{u}_i(t) - [\mathbf{b}_i(t+1) - \mathbf{b}_i(t)]$. Note that this is consistent with the case when there is only one individual where $u_i(t+1) = u_i(t) - [b(t+1) - b(t)]$. Observe that ignoring the effect of boredom we simply get the recurrence $\mathbf{u}_i(t+1) = A\mathbf{u}_i(t)$ or $\mathbf{u}_i(t) = A^t \mathbf{v}_i$. This recurrence reflects the diffusion of influence through the social network. Note that if the largest eigenvalue of A has magnitude more than 1 then the process will diverge and if all eigenvalues are < 1 it will eventually converge to 0. So we will assume the maximum eigenvalue of A is has magnitude 1. If the gap between the magnitude of the largest and the second largest eigenvalue is at least ϵ then this diffusion process converges quickly in about $\tilde{O}(1/\epsilon)$ steps. We will focus on the case when rate of boredom r is much slower than the diffusion rate (this corresponds to the case where influences spread fast and the boredom grows slowly). We then study the problem of making social choices of items over time so as to maximize welfare.

We will assume that A is diagonalizable and has a real top eigenvalue of 1 and all the other eigenvalues are smaller in magnitude. In that case it is well known that for any vector x $A^t x$ converges to to a fixed point and the speed of convergence depends on the gap between the largest and second largest eigenvalue. We show that under certain conditions if r/ϵ is small then the choices made by the society is comparable to the choices made by an individual with appropriate base utilities and boredom coefficients. Let $W_i(t)$ denote

the welfare of the society at time t by choosing item i ; then $W_i(t)/n$

THEOREM 2.4. *Consider a society with influence matrix A that has largest eigenvalue 1 and second largest eigenvalue of magnitude at most $1 - \epsilon$. For computing the welfare over a sequence of social choices approximately, such a society can be modelled as a single individual with base utilities \tilde{v}_i and boredom coefficients $\tilde{\alpha}_i$, where $\tilde{v}_i = \mathbf{c}'\mathbf{v}_i$ and $\tilde{\alpha}_i = \mathbf{c}'\alpha_i$ for some vector \mathbf{c} . Let $\tilde{u}_i(t)$ denote the utility of item i to such an individual at time t .*

More precisely, differences in the average utility of the society for the same sequence of choices until any time $|W_i(t)/n - \tilde{u}_i(t)| \leq \frac{\epsilon}{2} O(|\alpha_i|_\infty)$ for any $t > T$ for some fixed T . The O notation hides factors that depends on A . For a real, symmetric matrix the constant is 1

3. TECHNICAL DETAILS

In this section, we include the proofs for Theorem 2.1, 2.3, and 2.4 and refer to Appendix B for the proof of Theorem 2.2.

3.1 Utility optimization with boredom

We will first show that the total memory $\sum_i M_i(t)$ approaches a constant over time. The following is straightforward from the definition.

LEMMA 3.1. *$\sum_i M_i(t) \leq 1$, and $\sum_i M_i(t) \rightarrow 1$ for large t . When $t = \Omega(1/r)$, $\sum_i M_i(t) = 1 - O(\exp(-tr))$.*

Recall that $u_i(t) = u_i - \alpha_i M_i(t)$. The above lemma implies that $\sum_i (u_i - u_i(t))/\alpha_i = \sum_i M_i(t)$ is about constant for sufficiently large t . Together with the fact that the change to $u_i(t)$ is bounded by $r\alpha_i$, we are able to show that under the greedy strategy, the gap between $u_i(t)$'s cannot grow too large.

Before we prove Theorem 2.1, we prove the following Lemma that gives us an upper-bound on $\max_i u_i(t)$.

LEMMA 3.2. *For $t \geq T$ and any $i \in S_g$, $\max_j u_j(t) - u_i(t) = O(\alpha r \log n)$, where $\alpha = \max_i \alpha_i$.*

PROOF. The proof is by looking back at the times when each item is last picked and bound the gap iteratively. The details can be found in the appendix. \square

We are now ready to prove Theorem 2.1.

PROOF. (Theorem 2.1) To see that the solution to the given system is unique, note that $f_i = (\frac{v_i - \mu}{\alpha_i})^+$ (where x^+ denotes $\max(x, 0)$), and so $\sum_i (\frac{v_i - \mu}{\alpha_i})^+ = 1$. This must have a unique solution as $\sum_i (\frac{v_i - \mu}{\alpha_i})^+$ is decreasing function of μ and strictly decreasing as long as the sum is positive. Let μ denote the solution to the above system.

We now show $\max_i u_i(t) \geq \mu$ for any t . This is done by contradiction. Suppose that for all i $u_i(t) < \mu$. We have that $\sum_i (\frac{v_i - \mu}{\alpha_i})^+ < \sum_i \frac{v_i - u_i(t)}{\alpha_i}$. But $\sum_i \frac{v_i - u_i(t)}{\alpha_i} = \sum_i M_i(t) \leq 1$. We have that $\sum_i (\frac{v_i - \mu}{\alpha_i})^+ < 1$, a contradiction.

Let S_g denote the set of all the items ever picked by the greedy algorithm. Let T be the time by which each item in S_g has been used at least once. By Lemma 3.1, after some steps $\sum_i M_i(t)$ converges to arbitrarily close to 1. Lets assume for simplicity of argument that it is exactly 1 with

sufficiently large T . The upper-bound on $\max_i u_i(t)$ follows from Lemma 3.2.

Furthermore there exists $i \in S_g$ such that $u_i(t) \leq \mu$ because otherwise it would be the case that $\sum_i (\frac{v_i - \mu}{\alpha_i})^+ > \sum_i \frac{v_i - u_i(t)}{\alpha_i} = \sum_i M_i(t) \approx 1$, a contradiction. Hence $\max_j u_j(t) = \mu + O(\alpha r \log n)$. \square

Using the exactly same argument in the proofs of Theorem 2.1, we can prove that

LEMMA 3.3. *There exists a time T such that for any $t \geq T$, $\mu \leq \max_i w_i(t) \leq \mu + O(\alpha \log nr)$ where μ is the unique solution to the following system:*

There is an $f_i \geq 0$ for each item such that (1) $\sum f_i = 1$; (2) if $f_i > 0$, $v_i - 2f_i\alpha_i = \mu$, and $v_i < \mu$ otherwise.

To prove the approximation, we formulate a continuous version of the problem whose solution serves as an upper-bound to the utility in the discrete case. By using Lemma 3.3, we show that the double greedy strategy maintains each item's "marginal utility" close to that of the optimal solution from which we obtain the approximation bound.

PROOF. (Theorem 2.3) For $0 < f \leq 1$, write $\Delta(f) = r \cdot \frac{(1-r)^{1/f}}{1-(1-r)^{1/f}}$.

Let U^* be the optimal value of the following program.

$$\max U = \sum_i f_i (v_i - \alpha_i f_i) \quad \text{s.t.} \quad \sum_i f_i \leq 1 \quad \text{and} \quad f_i \geq 0. \quad (1)$$

Let OPT denote the optimal average utility. We have that $OPT \leq U^* + \alpha r$. This is by observing that for any $0 < f < 1$, placing an item $1/f$ apart gives an upper bound on the utility of consuming the item with frequency f . The bound is $v - \alpha\Delta(f) \leq v - \alpha(r - f)$ by observing that $\Delta(f) > r - f$.

The objective of (1) is maximized when there exists λ such that $\frac{\partial U}{\partial f_i} = \lambda$ for $f_i > 0$ and $\frac{\partial U}{\partial f_i} < \lambda$ for $f_i = 0$, and $\sum_i f_i = 1$. Since $\frac{\partial U}{\partial f_i} = v_i - 2\alpha_i f_i$, λ is exactly the same as μ in the statement of Lemma 3.3. This explains the intuition of the double greedy heuristics — it tries to equalize the marginal utility gain of each item. Denote the optimal solution by f_i^* . Then for $f_i^* > 0$, $v_i - 2\alpha_i f_i^* = \mu$. Hence,

$$U^* = \sum_i f_i^* (v_i - \alpha_i f_i^*) = \sum_i f_i^* (\mu + \alpha_i f_i^*) = \mu + \sum_i \alpha_i f_i^{*2}.$$

Let k_i denote the number of times item i is used in $[0, T]$ by the double-greedy algorithm, and $f_i = k_i/T$. Let \bar{M}_i denote the average memory on i at the times when i is picked. Then we have that

$$\begin{aligned} \bar{U} &= \sum_{t=0}^T u_{x(t)}(t)/T = \sum_i \sum_{x(t)=i} u_i(t)/T \\ &= \sum_i \sum_{x(t)=i} (w_i(t) + \alpha_i M_i(t))/T \\ &\quad \text{(by Lemma 3.3, } w_i(t) \geq \mu) \\ &\geq \sum_i \sum_{x(t)=i} (\mu + \alpha_i M_i(t))/T \geq \mu + \sum_i \alpha_i f_i \bar{M}_i. \end{aligned} \quad (2)$$

Write $\delta = \alpha r \log n$. By Lemma 3.3, $w_i(t) = \mu + O(\delta)$ for each i, t . We will show that

Claim 1. $\alpha_i f_i = \alpha_i f_i^* - O(\delta)$. Observe that for any item i which is picked k_i times in $[0, T]$, $\min_{0 \leq t \leq T} M_i(t) \leq$

$\Delta(k_i/T) \leq k_i/T = f_i$. Hence, $\max_t w_i(t) \geq v_i - 2\alpha_i M_i(t) \geq v_i - 2\alpha_i f_i$. On the other hand, $w_i(t) = \mu + O(\delta)$. We have $v_i - 2\alpha_i f_i = \mu + O(\delta)$. But $\mu = v_i - 2\alpha_i f_i^*$. Therefore $\alpha_i f_i \geq \alpha_i f_i^* - O(\delta)$.

Claim 2. $\alpha_i \bar{M}_i = \alpha_i f_i^* - O(\delta)$. Since $v_i - 2\bar{M}_i \leq \mu + \delta$, we obtain the bound by following the same argument as in the proof of Claim 1. Now, plugging both claims into (3), we have that

$$\begin{aligned} \bar{U} &\geq \mu + \sum_i \alpha_i (f_i^* - \frac{O(\delta)}{\alpha_i})^2 \geq \mu + \sum_i \alpha_i f_i^{*2} - O(\delta) \\ &= U^* - O(\delta) = OPT - O(\delta). \end{aligned}$$

This last equality follows from $U^* = OPT - \alpha r$. This completes the proof. \square

3.2 Social Choice is equivalent to individual choice

Let $\mathbf{1}$ denote the vector with all coordinates set to 1 and α_i denote the vector of boredom coefficients α_{ij} .

OBSERVATION 3.1. *For any diagonalizable matrix A with largest eigenvalue 1 and the second largest eigenvalue is at most $1 - \epsilon$, there is a vector c so that. $\mathbf{1}' A^t x - c' x \leq (1 - \epsilon)^t \sqrt{n} O(|x|_2)$. The O notation hides factors that depends on A . For a real, symmetric matrix the constant is 1.*

PROOF. We will sketch the proof for real symmetric matrices. The same idea holds for non-symmetric matrices. If p_1, \dots, p_n denote the eigenvectors of A and $1 = \lambda_1, \dots, \lambda_n$ denote the eigenvalues then $A^t x = \sum_j \lambda_j^t p_j p_j' x = p_1 p_1' x + \sum_{j>1} \lambda_j^t p_j p_j' x$. Now, $|\sum_{j>1} \lambda_j^t p_j p_j' x| \leq (1 - \epsilon)^t |x|_2$. So $|\mathbf{1}'(A^t x - p_1 p_1' x)| \leq |\mathbf{1}|_2 (1 - \epsilon)^t |x|_2 = \sqrt{n} (1 - \epsilon)^t |x|_2$ Setting $c = \mathbf{1}' v_1 v_1'$ completes the proof. \square

We will now prove theorem 2.4

PROOF. (Theorem 2.4) Let $\Delta \mathbf{b}(t)$ denote $\mathbf{b}(t) - \mathbf{b}(t - 1)$. Now $\mathbf{u}_i(t) = A \mathbf{u}_i(t - 1) + \Delta \mathbf{b}(t)$. This gives, $\mathbf{u}_i(t) = A^t \mathbf{v}_i + \sum_{j=0}^{t-1} A^j \Delta \mathbf{b}(t)$. Note that $W_i(t) = \mathbf{1}' \mathbf{u}_i(t) = \mathbf{1}' A^t \mathbf{v} + \sum_{j=0}^{t-1} \mathbf{1}' A^j \Delta \mathbf{b}(t - j)$

Note $\Delta b_{ij}(t) = \alpha_{ij}((1 - r)M_{ij}(t) + rI_i(t) - M_{ij}(t)) = \alpha_{ij}r(x_i(t) - M_{ij}(t))$. So $|\Delta b_{ij}(t)| \leq r|\alpha_i|_2$.

Now $|\mathbf{1}' A^t \mathbf{v}_i - c' \mathbf{v}_i| \leq (1 - \epsilon)^t \sqrt{n} O(|\mathbf{v}_i|_2)$. For $t > (1/\epsilon) \Omega(\log(nr|\mathbf{v}_i|_2))$, this is at most r . Also $|\mathbf{1}' A^j \Delta \mathbf{b}_i(t - j) - c' \Delta \mathbf{b}_i(t - j)| \leq (1 - \epsilon)^j \sqrt{n} O(|\Delta \mathbf{b}_i(t - j)|_2) \leq (1 - \epsilon)^j \sqrt{nr} O(|\alpha_i|_2)$. So, $|\sum_j \mathbf{1}' A^j \Delta \mathbf{b}_i(t - j) - \sum_j c' \Delta \mathbf{b}_i(t - j)| \leq (r/\epsilon) \sqrt{n} O(|\alpha_i|_2)$. So, $|\sum_j \mathbf{1}' A^j \Delta \mathbf{b}_i(t - j) - c' \mathbf{b}_i(t)| \leq (r/\epsilon) \sqrt{n} O(|\alpha_i|_2)$. Therefore $|W_i(t) - (c' \mathbf{v}_i - c' \mathbf{b}_i(t))| \leq (r/\epsilon) O(\sqrt{n} |\alpha_i|_2)$. Dividing by n completes the proof. \square

4. EXPERIMENTS

We now provide experimental results to study the techniques presented in the paper. Our primary objectives is to evaluate the quality of greedy and double-greedy algorithms for choosing items based on utility and boredom parameters estimated from the real data. In our experiments, we consider three families of fashionable items, namely songs, movies, and popular events. Based on the data from Google Trends [1], we obtain parameters that describe the utility of various items. We then show that double-greedy algorithm outperforms the greedy algorithm consistently.

4.1 Setup

We obtain data on the popularity of songs, movies, and popular events from Google Trends [1]. We collected weekly aggregate counts from query logs for popular songs from the last 5 years. Similar data was collected for popular movies and events too. The event data was gathered by collecting trends on keywords related to major “events” in the last 5 years. While the popularity of each entity depends on additional factors such as awards won by an album or a movie, our goal was to perform a controlled experiment only based on overall utility and boredom. Therefore, for each item we collected weekly aggregate counts starting from the highest peak in logs till there was an “artificial peak” due to an external event such as an award. Further, we compare the utility obtained by our model with a baseline in which the user selects an item simply based on its utility without any discounting from boredom. We describe how we compute the values of α , v , and r in the next section.

Table 1 shows the v , α , and r values for a set of 10 songs used in our experiments while the corresponding data for the movie and event data set is shown in Tables 2 and 3; here we allow different values of r , but we notice that all r -values within the domain of songs and movies are similar. In the case of events, the lambda values are roughly similar (although there is a little more disparity compared to songs and movies).

Song	v	α	r
The Climb	12.3	9.9	0.097
Lucky	2.6	1.58	0.114
Snow Patrol - Chasing Cars	10.7	6.8	0.127
I know you want me	7.95	6.5	0.077
Viva la vida	12.4	9.1	0.16
Stop and stare	10.5	9.4	0.092
Disturbia	8	7.2	0.092
Pocket full of sunshine	7.6	6.3	0.14
Supernatural superserious	24.2	22	0.15
One step at a time	9.35	8.5	0.075

Table 1: v, α, r for the set of songs

4.2 Computing the model parameters

Figure 2 shows the trend observed for a specific song from our dataset, I Know You Want Me, over a 45-week period starting August 2, 2009. The first natural observation we make is that the total number of queries do indeed display a steady decline, which we attribute to boredom. From the data, we use the maximum count as the peak utility, v_{peak} , and let the final count be denoted v_{final} . We set $\alpha = v_{peak} - v_{final}$. Let $X(t)$ denote the aggregate count for the week t , we obtain the boredom parameter r using the following equation: $e^{-rt} = 1 - \frac{v_{peak} - X(t)}{v_{peak} - v_{min}}$.

We plot $rt = -\ln(1 - \frac{v_{peak} - X(t)}{v_{peak} - v_{min}})$, and fit a linear line on the resulting curve and obtain r from the slope. Figure 3 shows the curve for I Know You Want Me, from which we obtain the r value.

4.3 Results

We ran a set of experiments to verify the effectiveness of the greedy and double-greedy heuristics. We ran the experiments over 100000 steps for both the data sets. The




Figure 2: Weekly aggregate query counts for I Know You Want Me for a 45-week period from Google trends

Movie	v	α	r
Godfather	6.15	5.15	0.123
Hancock	9.6	8.8	0.128
The Bucket List	13.1	11.8	0.102
Quantum of Solace	29.8	29	0.111
Tropic Thunder	25.6	24.8	0.082

Table 2: v, α, r for the set of movies

Event	v	α	r
NASA phoenix	4.02	3.96	0.32
Copenhagen Summit	27.8	27.4	0.14
H1N1	34.4	34	0.23
Lehman Brothers	33.8	33.6	0.09
Hudson River	18.9	18.1	0.3
Hurricane Katrina	6.3	6.2	0.3
Mumbai Terrorist	3.68	3.66	0.21
Michael Jackson	39.8	38.8	0.1
Haiti	80.5	79.5	0.25
Tsunami	186	184	0.19

Table 3: v, α, r for the set of popular events

average utility obtained by the user for both the data sets was computed and is shown in Table 4. We also show results for the baseline approach that always picks the same item with the highest base utility. Tables 5, 6, and 7 illustrate the average utility obtained by the user over the selected songs and movies respectively. The corresponding normalized frequencies are shown in parenthesis. As expected, in the baseline case where the user selects an item according to its base utility, the movie *Quantum of Solace* (with a base utility of 29.8) is always selected while in the case of songs, the song *supernatural superserious* (with a utility of 24.2) is selected. In the case of events, people always followed *Tsunami* (with a base utility of 186). Unsurprisingly, the average utility discounting boredom in all the cases is very low (see Table 4).

In another experiment, we measured the change in the average utility with time. Figure 4 illustrates the change in average utility as the user selects different items at each time step for movies. Naturally, the utility is highest at the very beginning as the user picks an item with the highest base utility and decreases subsequently as she picks items with highest discounted utility at each time step.




Figure 3: Approximate linear trend for I Know You Want Me, with slope giving r

Dataset	Greedy	Double-Greedy	Baseline
Songs	11.94	13.53	5.62
Movies	16.12	17.30	4.25
Events	76.6	88.06	37.74

Table 4: Average utility over 100000 time steps

Song	Greedy	Double-Greedy
The Climb	11.17(0.17)	11.11(0.17)
Snow Patrol - Chasing Cars	- (0)	10.24(0.12)
Viva la vida	11.61(0.17)	11.24(0.21)
Stop and stare	- (0)	10.08(0.08)
Supernatural superserious	12.22(0.67)	17.52(0.40)

Table 5: Avg. utilities (frequencies) for selected songs

Movie	Greedy	Double-Greedy
Hancock	- (0)	9.56(0.04)
The Bucket List	- (0)	11.40(0.20)
Quantum of Solace	16.41(0.55)	20.37(0.40)
Tropic Thunder	15.77(0.45)	18.01(0.36)

Table 6: Avg. utilities (frequencies) for selected movies

Event	Greedy	Double-Greedy
H1N1	- (0)	34.4 (0.03)
Michael Jackson	- (0)	37.81 (0.094)
Haiti	71.35 (0.25)	62.80 (0.38)
Tsunami	78.34 (0.75)	119.60 (0.50)

Table 7: Avg. utilities (frequencies) for selected events

5. FUTURE WORK

As we mentioned, our model is by no means comprehensive. For example, boredom may come from consuming similar items, or there may be a cost when switching from item to item. Taking into account these factors raises some interesting algorithmic issues. Fully incorporating these extensions is left as future work.

Acknowledgments. We thank Atish Das Sarma for useful discussions.

6. REFERENCES

- [1] Google Trends. <http://www.google.com/trends>.
- [2] Perron Frobenius Theorem. http://en.wikipedia.org/wiki/Perron-Frobenius_theorem.




Figure 4: Change in average utility over time.

- [3] A. Bar-Noy, R. Bhatia, J. Naor, and B. Schieber. Minimizing service and operation costs of periodic scheduling. *Math. Oper. Res.*, 27(3):518–544, 2002.
- [4] B. Barber and L. S. Lobel. Fashion in womens clothing and the american social system. *Social Forces*, 1952.
- [5] H. Blumer. Fashion: From class differentiation to collective selection. *Sociological Quarterly*, 1969.
- [6] D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, 2010.
- [7] M. P. Grindereing. Fashion diffusion. *Journal of Home Economics*, 1967.
- [8] N. Immorlica, J. Kleinberg, M. Mahdian, and T. Wexler. The role of compatibility in the diffusion of technologies through social networks. In *Electronic Commerce*, 2007.
- [9] J. K. J. Leskovec, L. Backstrom. Meme-tracking and the dynamics of the news cycle. In *Proc. 15th ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining*, 2009.
- [10] J. Leskovec, A. Singh, and J. Kleinberg. Patterns of influence in a recommendation network. In *PAKDD*, 2006.
- [11] C. M. Miller, S. H. McIntyre, and M. K. Mantrala. Toward formalizing fashion theory. *Journal of Market Research*, 1993.
- [12] W. H. Reynolds. Cars and clothing: Understanding fashion trends. *Journal of Marketing*, 1968.

APPENDIX

A. PROOF OF LEMMA 3.2

PROOF. Denote by $x(t)$ the item that has the maximum utility at time t . It suffices to show that $u_{x(t)} \leq u_i(t) + O(\alpha r \log n)$ for $i \in S_j$.

We recursively compute a decreasing sequence of t_j as follows. Let $t_1 = t$. For $j > 1$, suppose we have computed t_{j-1} . Let $S_{j-1} = \{x(t_1), x(t_2), \dots, x(t_{j-1})\}$. Now let $t_j = \max_{t' < t_{j-1}, x(t') \notin S_{j-1}} t'$. We stop when there is k such that $x(t_k) = i$. Since $t > T$, the process is guaranteed to stop. By the above construction, we know only items in S_{j-1} are picked by the greedy algorithm in the interval $[t_j+1, t_1]$. For any $S \subseteq \{1, \dots, n\}$, let $A(S, t) = \sum_{\ell \in S} \frac{u_\ell(t)}{\alpha_\ell}$, and $B(S) =$

$\sum_{\ell \in S} \frac{1}{\alpha_\ell}$. We will show that for $1 < j \leq k$,

$$A(S_{j-1}, t_{j-1}) \leq \frac{B(S_{j-1})}{B(S_j)} A(S_j, t_j) + r. \quad (4)$$

First observe that

$$A(S, t) = \sum_{\ell \in S} \frac{u_\ell(t)}{\alpha_\ell} = \sum_{\ell \in S} \frac{V_\ell - \alpha_\ell M_\ell(t)}{\alpha_\ell} = \sum_{\ell \in S} \frac{V_\ell}{\alpha_\ell} - \sum_{\ell \in S} M_\ell(t). \quad (5)$$

(4) follows from the following claims.

Claim 1. $A(S_{j-1}, t_{j-1}) \leq A(S_{j-1}, t_j + 1)$.

Since any item picked by the greedy algorithm in $[t_j + 1, t_{j-1}]$ is in S_{j-1} , we have that for $t' \in [t_j + 1, t_{j-1}]$,

$$\begin{aligned} \sum_{\ell \in S_{j-1}} M_\ell(t' + 1) &= (1 - r) \sum_{\ell \in S_{j-1}} M_\ell(t') + r \\ &\geq \sum_{\ell \in S_{j-1}} M_\ell(t'). \end{aligned}$$

The last inequality is by $\sum_{\ell} M_\ell(t') \leq 1$. Therefore

$$\sum_{\ell \in S_{j-1}} M_\ell(t_{j-1}) \geq \sum_{\ell \in S_{j-1}} M_\ell(t_j + 1).$$

By (5), we have $A(S_{j-1}, t_{j-1}) \leq A(S_{j-1}, t_j + 1)$.

Claim 2. $A(S_{j-1}, t_j + 1) \leq A(S_j, t_j) + r$.

Since $t_j \notin S_{j-1}$ is the item picked by the greedy algorithm at t_j , $\sum_{\ell \in S_{j-1}} M_\ell(t_j + 1) = (1 - r) \sum_{\ell \in S_{j-1}} M_\ell(t_j)$. Thus $\sum_{\ell \in S_{j-1}} M_\ell(t_j + 1) - \sum_{\ell \in S_{j-1}} M_\ell(t_j) = r \sum_{\ell \in S_{j-1}} M_\ell(t_j) \leq r$. Again by (5), we have $A(S_{j-1}, t_j + 1) \leq A(S_j, t_j) + r$.

Claim 3. $A(S_{j-1}, t_j) \leq \frac{B(S_{j-1})}{B(S_j)} A(S_j, t_j)$.

Immediately follows from $u_{t_j}(t_j) \geq u_\ell(t_j)$ for $\ell \in S_{j-1}$.

Repeating (4), we have that

$$\begin{aligned} &A(S_{j-1}, t_{j-1}) \\ &\leq \frac{B(S_{j-1})}{B(S_j)} A(S_j, t_j) + r \\ &\leq \frac{B(S_{j-1})}{B(S_j)} \left(\frac{B(S_j)}{B(S_{j+1})} A(S_{j+1}, t_{j+1}) + r \right) + r \\ &= \frac{B(S_{j-1})}{B(S_{j+1})} A(S_{j+1}, t_{j+1}) + r \cdot \frac{B(S_{j-1})}{B(S_j)} + r \dots \\ &\leq \frac{B(S_{j-1})}{B(S_k)} A(S_k, t_k) + r \cdot \sum_{\ell=j-1}^{k-1} \frac{B(S_{j-1})}{B(S_\ell)}. \end{aligned}$$

Hence, we have that

$$\begin{aligned} &u_{x(t_1)}(t_1) \\ &= \alpha_1 A(S_1, t_1) \leq \alpha_1 \left(\frac{B(S_1)}{B(S_k)} A(S_k, t_k) + r \cdot \sum_{\ell=1}^{k-1} \frac{B(S_1)}{B(S_\ell)} \right) \\ &= \frac{1}{B(S_k)} A(S_k, t_k) + r \cdot \sum_{\ell=1}^{k-1} \frac{1}{B(S_\ell)}. \end{aligned}$$

Since $i = x(t_k)$, for any i' , $u_{i'}(t_k) \leq u_i(t_k)$. Therefore $A(S_k, t_k) = \sum_{j \in S_k} \frac{u_j(t_k)}{\alpha_j} \leq u_i(t_k) \sum_{j \in S_k} \frac{1}{\alpha_j} = u_i(t_k) B(S_k)$.

By that $\alpha = \max_i \alpha_i$, we have $B(S_\ell) \leq \ell/\alpha$. Hence

$$\begin{aligned} u_{x(t_1)}(t_1) &\leq \frac{1}{B(S_k)} A(S_k, t_k) + r \cdot \sum_{\ell=1}^{k-1} \frac{1}{B(S_\ell)} \\ &\leq u_i(t_k) + \alpha r \sum_{\ell=1}^{k-1} 1/\ell \\ &= u_i(t_k) + O(\alpha r \log n). \end{aligned}$$

Since item i is not used during the interval of $[t_k + 1, t_1]$, we have $u_i(t_k + 1) \leq u_i(t_1)$, and hence $u_i(t_k) \leq u_i(t_k + 1) + \alpha_i r \leq u_i(t_1) + \alpha_i r$. Therefore we have that, $\max_j u_j(t_1) = u_{x(1)}(t_1) \leq u_i(t_1) + O(\alpha r \log n)$.

□

B. PROOF OF THEOREM 2.2

Restatement of Theorem 2.2: It is NP-hard to decide, given T , and U^* , and n items, whether there exist an assignment Y with period T such that $U(Y) \geq U^*$.

PROOF. The reduction is from the following problem.

Regular assignment problem (RAP). Given positive integers p_1, p_2, \dots, p_n , determine if there exists a sequence y_0, y_1, \dots where $y_t \in \{0, 1, \dots, n\}$ such that for any $i \neq 0$, two consecutive appearances of i in the sequence are exactly p_i apart.

It is shown in [3] that the regular assignment problem is NP-complete. Note that for RAP, a regular assignment exists if and only if it does so on a cycle with length $T = \prod_i p_i$. We will now reduce it to the optimal fashion selection problem.

Given p_1, \dots, p_n , we create $n+1$ items such that a regular assignment, if exists, maximizes the utility of any periodic selection with period T . Hence we can reduce RAP to the optimal selection problem. Item 0 is a special item with $v_0 = 1$ and $\alpha_i = 0$. For $1 \leq i \leq n$, we assign $v_i = \frac{2T}{p_i}$ and $\alpha_i = 1$. Further let $r_i = 1/T$ for $1 \leq i \leq n$. We claim that there exists U^* and $\epsilon \geq 1/T^2$ such that for a regular assignment Y , $U(Y) \geq U^*$, and $U(Y) < U^* - \epsilon$ otherwise.

Consider the case when there is only item and when the selections are made on the real line. Given T and an item with parameters v, α, r , let $\mathcal{Y}_k(v, \alpha, r)$ be the set of all the selections which have period T and choose the item exactly k times on the real interval $[0, T)$. Denote by $U_k(v, \alpha, r) = \max_{Y \in \mathcal{Y}_k(v, \alpha, r)} U(Y)$ and $\delta U_k(v, \alpha, r) = U_k(v, \alpha, r) - U_{k-1}(v, \alpha, r)$. The correctness of the reduction follows from the following claims.

Claim 1. $U_k(v, \alpha, r) = kv - k\alpha \frac{(1-r)^{T/k}}{1-(1-r)^{T/k}}$, and the maximum is achieved with the regular assignment.

Claim 2. For $1 \leq v \leq n$, $U_k(v_i, 1, 1/T) = kv_i - (k^2 - \frac{1}{2}k + \frac{1}{12} + o(1/k))$, and $\delta U_k(v_i, 1, 1/T) = v_i - (2k - \frac{3}{2} + o(1/k))$.

Claim 3. For any non-regular **integral** selection $Y \in \mathcal{Y}_k(v_i, 1, 1/T)$, $U(Y) < U_k(v_i, 1, 1/T) - 1/T^2$.

Claim 1 holds because the total memory is minimized when the k assignments are regularly spaced. Claim 2 is a direct consequence of Claim 1 by Taylor expansion on those particular parameters. Claim 3 follows by comparing the memory caused by adjacent items between regular and non-regular assignments.

From Claim 2, we can see that $\delta U_k(v_i, 1, 1/T) \geq 3/2$ for $k \leq T/p_i$ and < 0 for $k \geq T/p_i + 1$ for $1 \leq i \leq n$, and $\delta U_k(v_0, 0, 1/T) = 1$. Combining it with Claim 3, we have that the utility gap between a regular and non-regular assignment is at least $1/T^2$. Therefore the reduction is correct and can be done in polynomial time.

□