The population dynamics of websites* †

[Extended Abstract]

Kartik Ahuja
Electrical Engineering
UCLA
ahujak@ucla.edu

Simpson Zhang
Economics Department
UCLA
simpsonzhang@ucla.edu

Mihaela van der Schaar
Electrical Engineering
UCLA
mihaela@ee.ucla.edu

ABSTRACT

Websites derive revenue by advertising or charging fees for services and so their profit depends on their user base – the number of users visiting the website. But how should websites control their user base? This paper is the first to address and answer this question. It builds a model in which, starting from an initial user base, the website controls the growth of the population by choosing the intensity of referrals and targeted ads to potential users. A larger population provides more profit to the website, but building a larger population through referrals and targeted ads is costly; the optimal policy must therefore balance the marginal benefit of adding users against the marginal cost of referrals and targeted ads. The nature of the optimal policy depends on a number of factors. Most obvious is the initial user base; websites starting with a small initial population should offer many referrals and targeted ads at the beginning, but then decrease referrals and targeted ads over time. Less obvious factors are the type of website and the typical length of time users remain on the site: the optimal policy for a website that generates most of its revenue from a core group of users who remain on the site for a long time – e.g., mobile and online gaming sites – should be more aggressive and protective of its user base than that of a website whose revenue is more uniformly distributed across users who remain on the site only briefly. When arrivals and exits are stochastic, the optimal policy is more aggressive – offering more referrals and targeted ads.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Design, Policy

Keywords

Websites, referrals, population dynamics

1. INTRODUCTION

Many popular websites such as Facebook, Google, and Netflix [1] derive a significant portion of their revenue through advertising or by charging subscription fees to their users[2]. Given such a revenue model it is critical for the websites to obtain and maintain a healthy user base. Hence, a critical question that needs to be answered by such websites is: how can the websites control their user base? The user traffic on a website comes from two channels: sponsored and non-sponsored. Sponsored traffic is steered to the website through targeted ads, referrals, paid keywords, discounts etc. Hence, the website needs to decide how aggressively it should advertise as well as send referrals to control the user base given the incurred costs.

This paper aims to study and design policies which can be deployed by websites to control their user base through referrals and targeted ads in order to maximize their profits. Among the works on user base dynamics in the economics literature, the ones that relate the most to this work are [1] and [2]. In [1] the authors analyze the effects of search frictions in building a customer base on the firm's profits, investment, sales, etc. In [2] the authors study informative advertising and analyze the effect of a decline in the cost of information dissemination on the firm and customer dynamics. The key results in [1][2] are derived by calibrating models based on data, while this work instead provides a theoretical foundation for understanding user base dynamics. Also, the focus of these works [1][2] are on a firm selling a product to homogeneous users; in contrast the users on the website in our model can be heterogeneous and generate varying revenue depending on the number of advertisements they click.

We propose a dynamic continuous time model for the population of users present on a website. We assume that the website starts with a small initial user base and at every moment in time, the website can reach out to potential users via referrals and targeted ads. The website does this by paying a cost to incentivize its current users to send referrals to friends or by paying for targeted ads on other platforms. Thus, the website must adopt a policy that balances the marginal benefit of increasing its user base versus the marginal cost of providing the referrals and targeted ads. Our model accounts for the fact that users are heterogeneous, and therefore provide different benefits to the website [3]. We model this by making the natural assumption that the benefit function of the website is concave in the popula-

†Authors K. Ahuja, S. Zhang and, M. van der Schaar acknolweledge support from the ONR Mathematical Data Science Program

[1]See http://www.pewinternet.org/2015/01/09/social-media-update-2014/ and http://www.pewglobal.org/2012/12/12/social-networking-popular-across-globe/ .

[2]See http://finance.yahoo.com/news/must-know-assessing-facebook-revenue-170009607.html

[3]For instance, in mobile gaming apps half of the revenue coming from only 0.15% of the user base See http://venturebeat.com/2014/02/26/only-0-15-of-mobile-gamers-account-for-50-percent-of-all-in-game-revenue-

tion level: the value from adding new users decreases with an increasing population level.

If the website starts with a low initial population, we show that the optimal policy will be to give many referrals per unit time initially and then decrease the referrals over time. This is applicable for most websites that are new to the market and are unlikely to be able to acquire a high initial population. We also analyze how the website's optimal policy depends on the average time that a user stays on the website before exiting. We show that this depends critically on the revenue distribution of the website across its users. For instance, some websites rely heavily on a small set of users for revenue, e.g. mobile and online gaming websites, which rely on a small group of players who spend a lot on games. For such websites the optimal policy will increase the referrals and targeted ads if the average stay time of the users decreases. On the other hand, some websites obtain revenue from its users more uniformly, e.g. subscription based websites such as Netflix. For such websites the optimal policy will become more aggressive – to send more referrals and ads if the average time that a user stays on the website increases. We also extend our model to allow for stochastic user arrivals/exits, and we show that the above results are robust: if there is high uncertainty in user arrivals, we show that the website's policy will be to send more referrals and ads on average.

2. MODEL

We assume that there is a continuum of potential users and the firm can attract these users by posting ads on search engines (Google) and other websites (Facebook) or by giving referrals. New users can visit the website either through such sponsored media or arrive directly through exogenous methods, such as arriving upon the website through a non-sponsored link. We assume that the rate of such arrivals is a constant θ. This is a simplifying assumption made due to page limitations and the key results of this work continue to hold even under more general arrival rates that depend on the population levels such as $\theta . p^s$ with $0 \leq s \leq 1$. The users are heterogeneous, i.e. the revenue that the users generate for the website (e.g., the number of ads a user clicks) varies across the users. Every user who visits the website stops using the website after a random time. This time is a random variable drawn from an exponential distribution and the average time a user stays is η. Denote the total mass of the users using the website at time t as $p(t)$, where $p : [0, \infty) \to [0, \infty)$. The initial user base is denoted as $p(0)$. This total mass of the users represents the unique user statistic, which is often used as a metric to evaluate a website's popularity [4].

The revenue that the website generates per unit time increases with the mass of the users. It is denoted as $b(p)$, where $b : \mathbb{R} \to [0, \infty)$ is a continuously differentiable increasing function. Moreover, as mentioned in the introduction $b(p)$ is assumed to be a concave function in p. The long-term benefit of the website considering a discount rate of ρ can be computed as $B(p(.)) = \int_0^\infty b(p(t)) e^{-\rho t} dt$. The website chooses the intensity of advertisements (measured in terms of number of sponsored links, referrals, targeted ads) to be posted on other platforms at each time t and as a result it controls the rate of sponsored arrivals $\lambda(t)$, where $\lambda : [0, \infty) \to \mathbb{R}_+$ is a continuous and bounded function. The rate of change of population on the website is $\frac{dp}{dt} = \theta + \lambda(t) - \frac{1}{\eta} p$. The first two terms in the differential equation represent the rate of direct and sponsored arrivals respectively, while the third term represents the users exiting the website. The firm bears a cost per unit time for the advertisements, which increases with the sponsored user arrivals $\lambda(t)$. It is given by $c(\lambda)$, where $c : \mathbb{R} \to [0, \infty)$ is a continuously differentiable increasing function. Given the heterogeneous pool of potential users it is harder to increase the intensity of arrivals when λ is large. Hence, we assume that the cost $c(\lambda)$ is strictly convex in λ (as in [1]). We also assume that there is no cost when there are no sponsored arrivals $c(0) = 0$, and that the cost becomes unbounded as the sponsored arrivals approach ∞, i.e. $\lim_{\lambda \to} c(\lambda) = \infty$. The long-term discounted average cost to the website is $C(\lambda(.)) = \int_0^\infty c(\lambda(t)) e^{-\rho t} dt$.

The firm desires to maximize its total discounted profit, i.e. $B(p(.)) - C(\lambda(.))$. The continuous-time optimization problem of the firm is stated as follows

$$\max_{\lambda(t) \in \mathbb{R}_+, \forall t \geq 0} B(p(.)) - C(\lambda(.))$$

$$\text{subject to } \frac{dp}{dt} = \theta + \lambda(t) - \frac{1}{\eta} p(t), \; p(0) \text{ is given}$$

Next, we analyze the behavior of the optimal policy.

3. RESULTS

We can show that the optimal policy exists and is denoted by $\lambda_{p(0)}(t)$. Denote the corresponding population dynamic as $p_{p(0)}(t)$. If the optimal policy and the corresponding population dynamic converge, the steady state is achieved. The next theorem provides conditions for the existence of the steady state and also establishes that the steady state is unique. This theorem uses the the steady state population level $\hat{p}$, which is given by

$$X(\hat{p}) = b^{'}(\hat{p}) \frac{1}{\rho + \frac{1}{\eta}} - c^{'}(\frac{1}{\eta} \hat{p} - \theta) = 0 \qquad (1)$$

with $b^{'}(p) = \frac{db(p)}{dp}$ and $c^{'}(\lambda) = \frac{dc(\lambda)}{d\lambda}$.

THEOREM 1. *Steady state: Existence and Uniqueness*

i) If $0 < b^{'}(0) < \infty$ *and* $\lim_{p \to \infty} X(p) < 0$ *then there exists a unique solution* $\hat{p}$ *to the steady state equation (1).*

ii) If $p(0) < \hat{p}$ *then the optimal policy* $\lambda_{p(0)}(t)$ *decreases with time and converges to* $\hat{\lambda} = \frac{1}{\eta} \hat{p} - \theta$ *and the corresponding population dynamic* $p_{p(0)}(t)$ *increases and converges to* $\hat{p}$.

Due to space limitations we do not provide the proofs. Theorem 1 proves that if there is a positive marginal benefit from increasing the user base at very low population levels and if there is a negative marginal benefit from increasing the user base at very high populations then there exists a population level (between very low and very high population levels) where the marginal benefit is zero, which corresponds to the unique steady state. In addition we see that if the initial population level is low, which is true for most of the websites when they are launched, then the firm is more aggressive with advertising in the initial stages (closer to the launch) in comparison to the later stages (closer to the steady state). For the rest of the paper it is assumed that the firm starts with a low initial population, i.e. $p(0) < \hat{p}$. Also,

exclusive/

[4]http://www.pcmag.com/encyclopedia/term/53438/unique-visitors

we will use the terms advertisements and referrals, firm and website interchangeably henceforth.

3.1 Policy for Different User Behaviors

There are several interesting questions that one can ask about how the optimal policy depends on the user's behavior: What happens to the policy if the average time that a user stays on the website decreases? Or if the direct arrivals to the website increase? How does the new policy compare with the old policy both in the steady state (i.e. $t \to \infty$) and at a finite time t? We first address these questions at time $t \to \infty$, i.e. in the steady state.

3.1.1 Comparison of the policy in the steady state

If the average stay time of the users decreases then the firm is faced with the following question: Is spending more on advertisements worth it? The answer is not straightforward because of the following dilemma. The average stay time of the users reduces which discourages the firm, but the users leaving at a faster rate will also reduce the population and the total benefit may thus fall sharply. We answer this question next.

Define an operator $\Phi_{\theta,\eta,\rho}(b(p), p) = \frac{d^2b(p)}{dp^2} p\eta + \frac{db(p)}{dp} \frac{1}{\rho + \frac{1}{\eta}}$. The next theorem shows that if η decreases and $\Phi_{\theta,\eta,\rho}(b(\hat{p}), \hat{p}) < 0$ then the intensity of ads increases, otherwise the intensity decreases.

THEOREM 2. *Policy comparison with change in η in steady state $t \to \infty$: Local Behavior*

i) If the average time that a user stays on the website η is decreased by $\epsilon > 0$ and if $\Phi_{\theta,\eta,\rho}(b(\hat{p}), \hat{p}) < 0$ then the advertisements in the steady state increase.

ii) If the average time that a user stays on the website η is decreased by $\epsilon > 0$ and if $\Phi_{\theta,\eta,\rho}(b(\hat{p}), \hat{p}) > 0$ then the advertisements in the steady state decrease.

We assume that the change in η, i.e. ϵ to be small. Theorem 2 is interpreted as follows. If $\Phi_{\theta,\eta,\rho}(b(\hat{p}), \hat{p}) < 0$ then the percentage increase in the marginal benefit resulting from a decrease in the user population is high, i.e. $\frac{\frac{d^2b(p)}{dp^2}|_{p=\hat{p}}}{\frac{db(p)}{dp}|_{p=\hat{p}}} < -\frac{1}{(\rho+\frac{1}{\eta})\hat{p}\eta}$. Therefore the website should send out more ads if the average stay time reduces. Conversely, if $\Phi_{\theta,\eta,\rho}(b(\hat{p}), \hat{p}) > 0$ then the percentage increase in marginal benefit resulting from a decrease in population is not sufficient to compensate for the marginal cost of additional ads. Therefore, the firm reduces the ads.

The above theorem characterizes the local behavior of the firm's policy under the impact of small changes in the average stay time. Next, we characterize how the firm's global behavior under the impact of arbitrary changes in average stay time. We focus on a specific benefit function for better exposition, while the results presented extend to a larger class of functions. Consider a benefit function defined for $p \geq 0$, as $b(p) = p^a$ with $0 < a < 1$. For a fixed a, we define a threshold $\bar{\eta} = \frac{a}{\rho(1-a)}$. Then the following is true for the behavior of the firm.

RESULT 1. *Policy comparison with change in η in steady state $t \to \infty$: Global Behavior*

If the average stay time of the users decreases upto $\bar{\eta}$ then the firm should increase the intensity of advertisements in the steady state, while if the average stay time falls below $\bar{\eta}$ then the firm should decrease the advertisements.

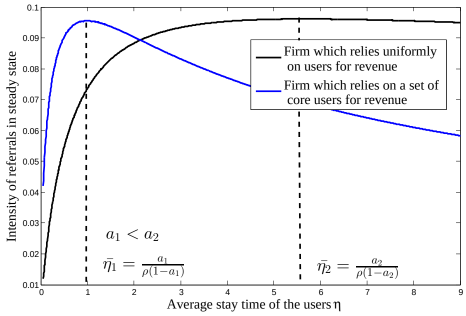

Figure 1: Impact of average stay time on firm's behavior in steady state

Next, we analyze the firm's behavior depending on its revenue distribution across the users, which is reflected by a. If a is small then the firm's benefit saturate very fast. This occurs if the firm relies heavily on a set of core users for revenue. Example of such websites are online and mobile gaming, since these websites rely heavily on their core users. Observe that a small value of a implies that the threshold $\bar{\eta}$ is very small as well. Hence, for such firms a reduction in the average stay time leads to an increase in the advertisements in the steady state. On the other hand, consider the case when a is large and close to 1. If a is close to 1 then the firm's benefit saturates slowly. This occurs if the firm relies uniformly on all the users for revenue. Examples include subscription based websites such as Netflix. Observe that a large value of a implies a high threshold $\bar{\eta}$. Hence, for such firms an increase (decrease) in the average stay time leads to a increase (decrease) in the advertisements in the steady state. As an example, the increase in the content quality and thus user stay time on Netflix led Netflix spending on expanding its user base [5]. Fig. 1 shows the example of the above two types of websites. Next, we understand the effect of an increase in the direct arrivals on the intensity of advertisements in steady state.

THEOREM 3. *Policy comparison with change in θ in steady state $t \to \infty$: Global Behavior*

If the intensity of the direct arrivals θ is increased then the intensity of advertisements in the steady state decreases.

The intuition behind Theorem 3 is as follows. An increase in the direct arrivals makes the firm decrease the advertisements in a controlled manner, such that the total population on the website for the new higher level of direct arrivals is higher than the case with original lower level. Hence, the firm can reduce its cost while simultaneously increasing the total population.

We have analyzed the impact of the revenue distribution of the firm and the average stay time of the users on the firm's policy in steady state. Now we want to understand the impact on the policy while it is on the path to steady state. For this, we will assume quadratic benefit and cost functions from now on due to space limitations. However, the results obtained extend to a larger class of benefit and cost functions.

3.1.2 Comparison of the policy on the path to steady state

[5] http://www.forbes.com/sites/petercohan/2015/01/21/4-reasons-to-invest-in-netflix/

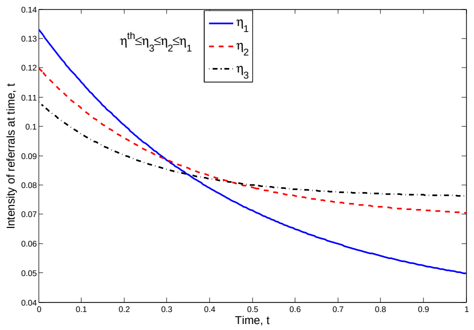

Figure 2: Impact of average stay time of the user η on firm's behavior out of the steady state when $\eta \geq \eta^{th}$

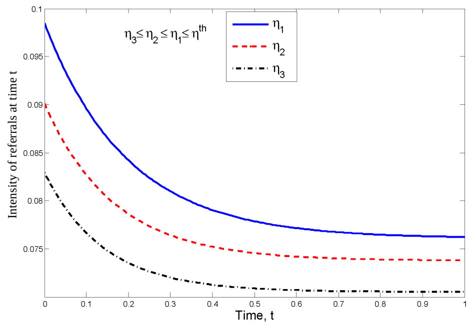

Figure 3: Impact of average stay time of the user η on firm's behavior out of the steady state when $\eta \leq \eta^{th}$

Websites are hosted on servers and the traffic that can be handled by a server is limited. Hence, we assume that the website has a maximum capacity, which corresponds to the maximum mass of users that the website can support with a non-negative marginal benefit. The benefit per unit time is a continuously differentiable concave function $b : \mathbb{R}_+ \to \mathbb{R}$ defined as $b(p) = \gamma^2 - (\gamma - p)^2$, where γ is the capacity of the website. In the previous subsection the benefit function was increasing in p, while this is no longer the case when capacity is incorporated in the model. We will assume that the capacity is sufficiently large, which will ensure that the results of the previous section continue to hold. We also assume that the cost per unit time for sponsored arrivals is quadratic (as in [1]) and is given as $c(\lambda) = c.\lambda^2$.

We will assume through out this section that the capacity is sufficiently high, i.e. $\gamma \geq \theta\eta$. This condition ensures that if there are no advertisements given at any time then the population in the steady state is less than the capacity. Under this assumption we can show that the result in Theorem 1, 2 and, 3 continue to hold. Hence, in the optimal policy the advertisements will decrease with time and achieve the steady state value $\hat{\lambda}$ (from Theorem 1). If the average stay time of the users were to decrease then the firm will continue to increase the advertisements $\hat{\lambda}$ until the stay time falls below a threshold, beyond which the advertisements decrease. Next, we analyze the effect of a decreased stay time on the policy outside the steady state.

THEOREM 4. *Policy comparison with change in η: at finite time t*

i) If the average stay time of the user decreases to a value above a sufficiently high threshold then the intensity of advertisements increases at sufficiently large time instances.

ii) If the average stay time of the user η decreases to a value below a sufficiently low threshold then the intensity of advertisements at every time instance decreases.

This theorem generalizes the result in subsection 2.2.1., where the behavior was analyzed in the steady state. In Fig. 2 we can see that if the average stay time of the user η reduces but stays sufficiently high then the firm spends more on advertisements – this would happen when the website is sufficiently old. This behavior from an incumbent website could serve as a barrier to entry to an entrant website, whose arrival possibly results in the decrease in the average stay time. Fig. 3 shows the second case, i.e. if the average stay time of the user decreases to a sufficiently low value then the website gives out less referrals in the new optimal policy.

Next, we understand the effect of an increase in the direct arrivals on the policy outside the steady state. It can be shown that there will be a decrease in the advertisements at all times. This generalizes Theorem 3. The intuition for this result is the same as that for Theorem 3.

Extension - Uncertain user arrivals/exits: We can show that even under uncertain arrivals/exit of the users the key results presented in the deterministic setting will continue to hold . The population dynamic with uncertain user arrival/exit changes into the following stochastic differential equation $dP = (\theta + \lambda(t) - \frac{1}{\eta}P)dt + \sigma P dW_t$. Here W_t is the Weiner process/Brownian motion. The firm's objective in this case is to maximize the expectation of the long term total discounted profit (benefit-cost). If we consider the quadratic benefit and costs considered in the previous section, Theorem 4 can be extended to the stochastic setting as well . In this case the comparison is done between the expectation of the intensity of sponsored arrivals. An interesting result emerges is when we compare the intensity of the advertisements under varying levels of uncertainty σ. It can be shown that increasing the level of uncertainty makes the optimal policy more aggressive in giving advertisements.

4. CONCLUSION

This paper has been the first to systematically characterize how a website should build its user base. We showed that the optimal policy requires that websites starting with small initial populations send ads aggressively in the beginning and then decrease them with time. For websites that derive revenue from a set of core users, e.g. mobile and online gaming, the optimal policy increases ads when the average stay time of the user on the site decreases, while for websites with a more uniform revenue distribution across users the optimal policy decreases ads. These results extend to a stochastic setting with noise in population dynamics, and it is shown that more uncertainty in user arrivals/exits leads to a more aggressive optimal policy for ads.

5. REFERENCES

[1] F. Gourio and L. Rudanko. Customer capital. *The Review of Economic Studies*, page rdu007, 2014.

[2] E. M. Dinlersoz and M. Yorukoglu. Information and industry dynamics. *American Economic Review*, 102(2):884–913, 2012.