Managing Online Content to Build a Follower Base: Model and Applications

Felipe Caro* Victor Martínez-de-Albéniz†

Forthcoming in INFORMS Journal of Optimization

Abstract

Content providers manage their production by regulating the pace at which content is created and released. They have two types of consumers: existing ones, or *followers*; and new visitors who may become followers. To maximize their effectiveness, providers must consider the direct, short-term effect of content and also its indirect, long-term effect on the retention and expansion of their follower base. We develop a simple model to study the dynamics of building up a follower base and then combine that model with stochastic dynamic programming to optimize the content provider's profit. We find that optimal content policies exhibit trajectories whereby both content and followers increase until reaching a steady state at which revenues equal marginal costs. It is therefore preferable to start slow and then accelerate only after successful content has increased the follower base; thus past success makes the content provider work harder. We apply our model to bloggers by analyzing the traffic from new and returning visitors to several blogs. We find that optimized posting activity yields significantly higher profits than does a steady or seasonal content strategy. Our model contributes to the growing literature on data-driven optimization and prescriptive analytics, specifically for content management.

1. Introduction

The worst part [of success] is that it drags you onto the conveyor belt of working all the time, to maximize the success.

—Phil Collins (in *Rolling Stone*, "The Last Word", November 2016)

The production of cultural content such as press articles, radio or television clips, or online posts is a challenging task faced by media companies and many others. This task is one faced not only by a newspaper's editorial board; even a retailer with a social media strategy must carefully determine what, how, and when to communicate with potential customers. These decisions are sufficiently relevant to be the cause of success or failure of a publication, broadcasting, or social media channel.

There are several types of difficulties in planning the production of new content. First, content is costly and so there are economic factors that a producer must take into account. Second, content

*UCLA Anderson School of Management, Los Angeles, CA 90095, felipe.caro@anderson.ucla.edu
†IESE Business School, University of Navarra, Barcelona, Spain, valbeniz@iese.edu.

strongly affects consumers' interest, and it is seldom easy to know a priori whether a new piece will be appreciated or dismissed by consumers. Third, content is dynamic in the sense that it is not consumed all at once by everybody; rather, its consumption is diffused over time. Consumers are known to be extremely sensitive to recency, which means that the attention given to a piece of content declines quickly over time. Finally, these dimensions (cost, revenue, and dynamics) are complicated by the following phenomenon: content "leaves a mark" on consumers by changing their appreciation of the content provider's platform. Indeed, the more views generated by a piece of content (e.g., from a link on Facebook), the greater the number of individuals who will be aware of the site's existence and the more likely it is that they will "bookmark" it for future visits. The subscribers of a periodical magazine will similarly update their satisfaction every time they receive a new issue: readers who enjoyed it will likely stay with the service; but those who did not may well cancel their subscriptions and thereby reduce its future readership base. These intertemporal effects are of paramount importance, and they enrich yet complicate the content producer's task. It is clearly possible to expand the base of customers – which we refer to as *followers* – by releasing more and/or better content. At the same time, content release decisions become more difficult and the provider needs to carefully balance costs, direct revenues (current and future) from the new content, and also indirect revenues (current and future) derived from enlarging the follower base.

The purpose of this paper is to provide some guidance on such decisions. More, specifically, we develop an advanced dynamic model for content generation that can be used as a data-driven strategic tool to oversee how content should be released over time. The model's main distinctive feature is its consideration of the dynamics of content optimization, which cannot be easily included within a simpler static model capable of yielding closed-form solutions. Our model could be used by managers of content production, and there are several concrete applications for which it offers useful advice.

First, this model supports the process of allocating resources to different content streams, which should reflect the respective "return on followers" that each content type delivers. Although content can certainly produce an immediate traffic spike, that traffic does not always translate into loyal followers who are retained in the long term (see Besbes et al. 2015 for a similar idea). Second, the model can also be used to estimate the market potential of a particular content stream, a feature that aids the manager of a content platform in determining which streams are viable and which are not. Third, our model can help content providers actively manage a follower base by suggesting, at any stage of the platform's life cycle, an appropriate level of content "intensity" that supports the follower base at optimality, i.e., so that traffic revenue minus content costs are maximized.

One can perhaps better appreciate our model's value by considering a manager in charge of digital content for a news platform, such as CNN and its online channel (cnn.com). This manager coordinates a large team of journalists and writers that "feed" the platform. She must know how to allocate these resources among various sections (e.g., international, US politics, finance), which tend to have different follower bases. Allocating resources strongly affects the platform's total traffic, and the model presented here can help the manager boost traffic while containing editorial costs. In addition, our model can be used to estimate the potential reach of each section when content is released optimally – information that is a key component in negotiating the advertisers' placement deals on which the platform relies for funding. Finally, the manager meets periodically with the customer relationship management (CRM) group, which tracks customer retention and often requests new content actions that will expand the customer base. Although the value of active users is widely acknowledged (Mediakix 2017), it is not entirely clear how best to balance the acquisition of new customers with the costs associated with producing new content (O'Neill 2011). Moreover, achieving an appropriate balance is especially difficult when users are not tracked individually.

To achieve this balance, we use dynamic programming (DP) to develop an optimization model that enables the content provider to optimize content release decisions over time. We formalize the concept of "followers" by considering two types of visitors: *new* visitors, who have never before visited the platform; and *returning* visitors that are part of the follower base. We assume that new visitors arrive randomly to the platform and are converted into followers. Returning visitors are also assumed to arrive randomly, but grow with the size of the follower base. In addition, followers can abandon the platform at any time – as captured by the "churn rate" – and we define a "virality factor" that characterizes the feedback loop between past traffic and platform visibility. The random arrival processes are affected in part by the provider's content decisions, so the provider has some control over arrivals albeit at a cost. Finally, the model accommodates time-varying traffic to specific content and thus captures decaying patterns observed in practice. Figure 1 plots this pattern for a single post from one of the blogs analyzed in our case study.

Generic DP formulations typically suffer from a large and growing state space, known as the "curse of dimensionality". Yet our base model assumes that content strength (i.e., attractiveness to visitors) decays exponentially, which allows us to devise a DP formulation with only a four-dimensional state space. We characterize the optimal policy's structure and find that stronger content should be released in a period when (a) the follower base is larger and (b) the past content is stronger. Thus content has a reinforcing effect that leads to increasing trajectories: followers and

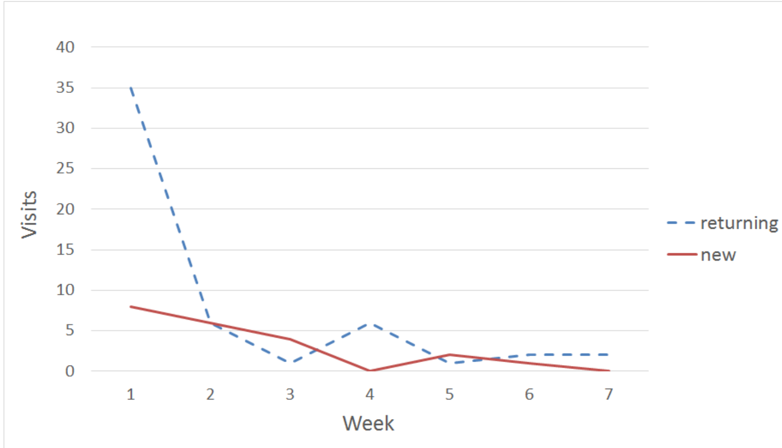

Figure 1: Traffic decay for a single post. The post generates traffic (new and returning) when it is released in week 1, after which the number of visits decays rapidly.

content grow in concert over time. As this growth generates increasing content production costs, followers and content eventually stabilize. Note that this dynamic does not imply that the platform is viable: it only states that the platform has a given potential and our model characterizes what that potential is. Furthermore, when there are fixed costs to sustain the platform, it might be optimal to shut it down despite the increasing trajectories.

In addition to shedding light on the optimal policy's structure, this paper derives lower and upper bounds on the optimal profit based on deterministic DP formulations, which can be solved efficiently when costs are quadratic. Finally, we extend the model to incorporate uncertainty about content quality, general content decay patterns, a broad class of arrival processes, and fixed costs related to managing content.

Following that theoretical development, we use a real-world case study to illustrate our model's value for a manager of blogs. To estimate the model parameters, we use the number of individual visits of new and returning customers to each post, for a period of several years and several blogs. We find that followers are very loyal: the model's weekly retention factor – i.e., the proportion of followers in a given period who are retained in the following period – exceeds 90%. The implication is that current visits do have a "carry-over effect" on future visits. We use estimates derived from our empirical study to compare actual blogs' posting policies with the model's recommendations. The main finding is that a constant or a seasonal blogging effort is suboptimal and may lead to significant losses compared to the optimum in our data. Thus our research contributes to the nascent literature on prescriptive analytics in terms of its application to content management that incorporates follower-base dynamics.

The paper proceeds as follows. Section 2 reviews the streams of literature most relevant to our analysis. Section 3 introduces the model and describes the theoretical results. We present the application to blogs in Section 4 and conclude in Section 5. Extensions and mathematical proofs are given in Appendix A and Appendix B.

2. Literature Review

The work in this paper is related to several streams of literature in computer science, marketing, operations, and economics.

The first relevant field of study is the one addressing social media and blogs in particular, to which we apply our model. Most work in that stream is in the realm of computer science and uses machine learning methods and data mining to detect patterns in the diffusion of content. The main applicable insights from this literature are that content success is heterogeneous and that interest in specific content decays quickly following its release; our model incorporates these features. Leskovec et al. (2007) detect the evolution of power laws in the links established to blogs (which are associated also with traffic), and Goetz et al. (2009) report similar patterns in the time between two consecutive posts. Guo et al. (2009) analyze user posting activities in blogs (and in photos, bookmarks, and responses); these authors find that contributors are seldom homogeneous in their activity and posting intervals. Sun and Zhu (2013) document how posting behavior is affected by commercial incentives due to advertising. Bandari et al. (2012) predict the popularity of tweets based on initial responses and emphasize that interest decays over time. Aggarwal et al. (2012) show that negative posts in a blog can have a positive effect if they attract visitors that become followers. The authors build and estimate an autoregressive model for page views and show that readership typically follows an exponentially increasing curve before stabilizing – a pattern revealed also by our model. Goel et al. (2015) study virality patterns on Twitter and find that viral tweets appear at a rate of a few per million. Finally, Mendelson and Moon (2016) develop a statistical model for app evolution on Facebook; their paper highlights that app growth is driven by customer retention, which is a key element in our model as well.

The research just cited does not directly capture the behavior of returning visitors (i.e., followers); however, the marketing literature has developed models that account for "adoption build-up" (Bass 1969, Bass et al. 1994), and "intertemporal spillover" (Nerlove and Arrow 1962). More recent work has studied how current marketing variables influence present and future customer behavior. Their effect on present behavior is discussed by Ghose and Yang (2009) and Rutz and Trusov (2011), who show how advertising feature can increase click-through behavior and conversion. The effect of marketing variables on future behavior is considered by Rutz and Bucklin (2011), Chan et al. (2011), Joo et al. (2014), and Liaukonyte et al. (2015), who examine the lagged effects of current marketing decisions (e.g., paid online advertising) on future interactions. Our model shares similar modeling features, including the creation of latent variables to account for the size of the future visitor pool. However, our intent is different in that we seek to characterize optimal dynamic policies that include such intertemporal spillovers. Thus we aim to prescribe how content can be used to create new followers who will generate future visits.

Finally, our paper is connected to work in the operations literature that has studied the optimization of product introductions and/or renewals. For instance, Wilson and Norton (1989) identify the optimal time to introduce a line extension; Burns et al. (1997) describe the shortening of life cycles and the necessity of "retail renewal". Krankel et al. (2006) optimize product introductions under a demand diffusion model that incorporates carry-over effects from past sales into the future. More recent dynamic assortment models also incorporate substitution across multiple products: Caro et al. (2014) optimize product introduction times and explicitly include demand decay for products introduced long ago, and Çınar and Martínez-de-Albéniz (2013) analyze the same problem within a DP framework. Martínez-de-Albéniz and Valdivia (2016) apply similar ideas to museum exhibitions. Lobel et al. (2015) and Bernstein and Martínez-de-Albéniz (2017) describe the optimal timing of product introductions in the context of consumers who behave strategically. Optimization models based on service level management, rather than product launches, are developed by Olsen and Parker (2008), Huang et al. (2008), and Afèche et al. (2017), among others.

3. Connecting Content Decisions to Customer Interest

3.1 The Model

We consider the problem of a content provider interested in maximizing the traffic that arrives to the content site. That site could be an online portal such as cnn.com, a newspaper, a YouTube channel, or a blog on a content platform. In these settings, traffic is directly related to revenue – either because it can be monetized through a fee (e.g., the price paid to access specific content) or because advertising income is proportional to traffic (e.g., the revenue collected from showing ads on the online portal).

In seeking to maximize traffic, the content manager must decide how much content to release in each period $t = 1, \ldots, T$; these periods are assumed to be discrete, where t denotes the time period (say, a week) and T could be either finite or infinite. So that our model will be tractable, we assume that the manager's content decision is captured by a univariate metric $z_t \geq 0$ that can be interpreted as the attractiveness or "strength" of the content released in period t. We also assume that z_t is a scalar, although the policy could encompass various content attributes such as the number of posts, the length of those posts, and so forth.[1] The cost associated with producing content of strength z_t is denoted $c(z_t)$, and it is assumed for tractability to be $\mathcal{C}^2$ (i.e., twice differentiable). So if $\mathbf{a}_t$ denotes the multi-dimensional vector of attributes describing the content at time t, then $c(z_t) = \min_{\mathbf{a}_t \geq 0} \tilde{c}(\mathbf{a}_t)$ subject to $\mathbf{e}'\mathbf{a}_t = z_t$, where $\tilde{c}(\mathbf{a}_t)$ is the cost associated with attributes $\mathbf{a}_t$ and $\mathbf{e}$ is a vector of 1s. Thus, by assuming that content attributes $\mathbf{a}_t$ can be mapped to a scalar z_t, in effect we decouple content management (i.e., how much content to release in each period) from the production decisions involved in generating content of a given strength. We shall focus on the problem of content management, though we remark that the production problem is considered by other scholars; for example, Sun and Zhu (2013) identifies the drivers of effective content.

We assume that content is accessed by two types of customers, as described next.

- First, there are returning customers who have visited the content site before. We use B_t to denote the follower base – measured as the number of potential returning visitors – at the start of time t. These customers react to the arrival of new content and access it; for example, the provider might send them a notice (say, via e-mail) that new content is available. Members of the follower base may also want to visit some older content.

  The appeal of new content decreases with time. More precisely, the attractiveness of content released in period s decays exponentially at the rate β_{ret} per period.[2] The traffic of followers will also decay. Namely, the attractiveness at time t of the content released in period s is proportional to $z_s \beta_{\text{ret}}^{t-s}$ and the expected number of follower visits can be written as $r_{st} := B_t z_s \beta_{\text{ret}}^{t-s}$.

- Second, in any given period there is a universe of potential new visitors A_t who may arrive to view new or older content. Such traffic is typically generated from some search source (e.g., results listed by a search engine), but it may also come from referrals. These new visitors become followers in the sense that, if they return to the content site, they will be treated as returning customers. We assume that the content released in period s contributes (on average) $n_{st} := A_t z_s \beta_{\text{new}}^{t-s}$ new visitors in period t, where β_{new} is the exponential decay of content for new visitors and $z_s \beta_{\text{new}}^{t-s}$ is the (residual) attractiveness at time t of the content that was released in period s.

The actual number of returning and new customers may, of course, differ from the average number. To incorporate randomness, our model assumes that the actual number of returning customers

[1] The scalar z_t can be seen as a univariate metric that combines quantity and quality of content released in period t.
[2] General decay rates are considered in Appendix A.1.

in period t who are accessing content posted in period s, denoted R_{st}, is distributed according to a Poisson distribution with mean r_{st}. Similarly, the actual number of new customers in period t who are accessing content posted in period s, denoted N_{st}, is distributed according to a Poisson distribution with mean n_{st}. Thus the arrival processes for returning and new customers each follow a Poisson distribution.[3] Poisson arrivals are often used in contexts such as ours because traffic streams reflect the realization of individual choices and so are typically memoryless. For example, Moe and Fader (2004) use a similar approach to model online page visitors. Moreover, for the application described in Section 4, we find that Poisson arrivals fit the data fairly well. Note, however, that this is not a strong assumption in our setup: all the analytical results extend to any arrival distribution provided that N_{st} and R_{st} are each stochastically increasing in their respective mean parameters; examples include Normal distributions with a fixed standard deviation, Gamma distributions with a fixed scale parameter, and negative binomial distributions with a fixed probability of success (see Appendix A.1 for details). We remark that, in these examples, the sum of traffic from content posted in several periods retains the same distributional form.

The key parameters in this model are A_t and B_t. Their values are updated dynamically and are affected (endogenously) by the strength of current and past content. In particular, the dynamics of the universe of potential new visitors and the follower base are as follows. On the one hand, the former grows over time whereas the latter undergoes a natural attrition (a.k.a. churn) that proceeds at rate $1 - \alpha$ for $\alpha \in [0, 1]$. On the other hand, both streams of customers are reinforced by traffic; that is, they are increasing in the site's number of past visitors. We assume that the potential universe grows with the total traffic experienced in the preceding period and use a "virality" parameter $\kappa \geq 0$ that accounts for possible amplification effects (as might occur, for instance, because of the site's increased visibility to search engines; seeMendelson and Moon 2016). Similarly, the follower base incorporates all new customers that are arriving for the first time.[4] Formally, the dynamics of A_t and B_t are given by

$$A_{t+1} = A_t + \kappa \mathcal{P}\left(\sum_{s=1}^{t} n_{st} + \sum_{s=1}^{t} r_{st} \right) \tag{1}$$

and

$$B_{t+1} = \alpha B_t + \mathcal{P}\left(\sum_{s=1}^{t} n_{st} \right), \tag{2}$$

[3]We can extend the results to include shocks to the Poisson intensity (e.g., uncertainty about content quality), as shown in Appendix A.1.

[4]The growth in A_t and B_t assumes that all past traffic is amplified by κ, regardless of the source (new or returning), and that all first-time visitors become part of the follower base. These simplifying assumptions could be easily extended to the case in which only a fraction of the traffic joins, or when either α or κ depend on t, without any qualitative change in the results. See Appendix A.1 for details.

where $\mathcal{P}(\lambda)$ is a Poisson random variable of rate λ. To initialize the follower base, we let $B_1 = 0$. In contrast, A_1 is a model parameter. The dynamics of A_t and B_t differ: whereas A_t is driven by total visits to the platform and is not diminished by attrition, the value of B_t is dampened over time and any growth is due only to new visitors.

The geometric decay of content strength allows us to summarize the system's state in terms of the following variables:

$$r_{t+1} := \sum_{s=1}^{t} z_s \beta_{\text{ret}}^{t-s} = z_t + \beta_{\text{ret}} r_t; \tag{3}$$

$$n_{t+1} := \sum_{s=1}^{t} z_s \beta_{\text{new}}^{t-s} = z_t + \beta_{\text{new}} n_t. \tag{4}$$

By (3), r_{t+1} is the sum of the return visit probabilities across all the content released until time t. In other words, the state variable r_{t+1} captures the aggregate attractiveness of the content site at the end of period t (or at the start of period $t+1$) for the follower base. The state variable n_{t+1} is likewise a measure of the overall attractiveness of past content for new customers. Hence, we can rewrite the follower base dynamics in Equation (2) as follows:

$$A_{t+1} = A_t + \kappa \mathcal{P}(A_t(z_t + \beta_{\text{new}} n_t) + B_t(z_t + \beta_{\text{ret}} r_t)) \tag{5}$$

and

$$B_{t+1} = \alpha B_t + \mathcal{P}(A_t(z_t + \beta_{\text{new}} n_t)). \tag{6}$$

We are now ready to formulate the content provider's objective function. Let the current monetary value of a repeat visit by a follower be p_{ret} dollars per visit and let the value of a new visit be p_{new} dollars per visit. Revenue and costs are discounted at the rate δ. Then the content management problem can be formulated as the following finite-horizon DP problem, where the value function W_t is equal to the expected discounted future profit from t onward:

$$W_t(A_t, B_t, r_t, n_t) = \max_{z_t \geq 0} \left\{ \begin{array}{c} p_{\text{new}} A_t(z_t + \beta_{\text{new}} n_t) + p_{\text{ret}} B_t(z_t + \beta_{\text{ret}} r_t) \\ -c(z_t) + \delta \mathbb{E}[W_{t+1}(A_{t+1}, B_{t+1}, r_{t+1}, n_{t+1})] \end{array} \right\}; \tag{7}$$

here A_{t+1}, B_{t+1}, r_{t+1}, and n_{t+1} are given by Equations (3)–(6) and $W_{T+1} \equiv 0$. Since c is $\mathcal{C}^2$ and since the Poisson distribution probabilities are C^∞ in the mean, it follows that W_t is both finite and $\mathcal{C}^2$; hence partial derivatives with respect to all variables exist.

For an infinite horizon we define the stationary value function $W(A, B, r, n)$, which must satisfy the following Bellman equation:

$$W(A, B, r, n) = \max_{z \geq 0} \left\{ \begin{array}{l} p_{\text{new}} A(z + \beta_{\text{new}} n) + p_{\text{ret}} B(z + \beta_{\text{ret}} r) - c(z) \\ +\delta \mathbb{E}\left[W \left(\begin{array}{c} A + \kappa \mathcal{P}(A(z + \beta_{\text{new}} n) + B(z + \beta_{\text{ret}} r)), \\ \alpha B + \mathcal{P}(A(z + \beta_{\text{new}} n)),\ z + \beta_{\text{ret}} r,\ z + \beta_{\text{new}} n \end{array} \right) \right] \end{array} \right\}. \tag{8}$$

If the marginal cost of creating content c' is sufficiently high, then the optimal value is finite ($W < \infty$). It can also be shown that value iteration converges; that is, the finite-horizon problem W_t converges to W as t tends to infinity (Bertsekas 2000).

The following theorem characterizes the structure of the optimal policy and value function.

Theorem 1. *For all t:*

(a) *W_t is nondecreasing and supermodular in (A_t, B_t, r_t, n_t) and is unidirectionally convex in A_t and B_t; and*

(b) *z_t^* is component-wise nondecreasing in (A_t, B_t, r_t, n_t).*

These properties carry over to W and z^* for the infinite-horizon case in Equation (8).

This theorem describes a structure of increasing returns in the content management problem. Indeed, convexity and supermodularity imply that an additional follower increases the marginal value of content. The reason is that an additional follower will increase not only the revenue generated per post made in the past but also increases the incentive to create content in the future. Since the system is initialized at $B_1 = 0$ and since the model parameters are stationary, it follows that the content provider's optimal path will have increasing trajectories. In other words, the number of followers can be expected to increase over time. Yet this does not imply that followers are increasing in *every* sample path, because low realizations of the Poisson arrival process might create occasional drops in the size of the follower base. The size of the potential universe similarly increases the marginal value of content.

We can interpret Theorem 1 in the blogging context as follows: past success makes the blogger work harder; thus, when a blog is a hit and many new followers join (or more referral channels are created), it is optimal for the blogger to work harder in the next period to take advantage of the "tailwind". In contrast, if the blog gains little traction then the blogger should not overexert himself in the next period; instead, he should wait for a stronger tailwind. This result implies that, in creative industries such as blogging and music publishing, it is suboptimal to "rest on one's laurels" after a hit – as underscored by our paper's opening quote. It is worth noting that this result continues to hold even when there is a fixed cost associated with blogging (see Appendix A.4). In that event, however, it might be optimal to discontinue the blog at some point if the follower base does not become large enough to recoup the fixed cost.

The formulation in (7) requires that we solve a stochastic dynamic program with a reasonably sized (four-dimensional) state space.[5] Solving this DP numerically would require simulating the

[5] A general formulation requires a state space of $T + 1$ dimensions; see Appendix A.1.

randomness of visitor arrivals and hence would be computationally challenging. We next investigate solution methods that are more tractable by considering the problem's deterministic counterpart.

3.2 Deterministic Counterpart

Suppose that arrivals, instead of following a Poisson process, are deterministic. Then the dynamics of the follower base are described by the recursions

$$A_{t+1} = A_t + \kappa(A_t n_{t+1} + B_t r_{t+1}) \quad \text{and} \quad B_{t+1} = \alpha B_t + A_t n_{t+1}.$$

These approximated dynamics lead to a modified optimization for which the corresponding DP formulation is given by

$$W_t^D(A_t, B_t, r_t, n_t) = \max_{z_t \geq 0} \left\{ \begin{array}{l} p_{\text{new}} A_t(z_t + \beta_{\text{new}} n_t) + p_{\text{ret}} B_t(z_t + \beta_{\text{ret}} r_t) - c(z_t) \\ +\delta \left[W_{t+1}^D \left(\begin{array}{c} A_t + \kappa A_t(z_t + \beta_{\text{new}} n_t) + \kappa B_t(z_t + \beta_{\text{ret}} r_t), \\ \alpha B_t + A_t(z_t + \beta_{\text{new}} n_t),\, z_t + \beta_{\text{ret}} r_t, z_t + \beta_{\text{new}} n_t \end{array} \right) \right] \end{array} \right\} \tag{9}$$

with $W_{T+1}^D \equiv 0$. From Theorem 1 we know that W_t is convex in A_t and B_t. Hence the deterministic counterpart W_t^D constitutes a lower bound on the optimal profit of the stochastic problem, which is formally stated in the following proposition.

Proposition 1. *For all t, we have $W_t \geq W_t^D$.*

In the infinite-horizon case, we are able to provide a stronger characterization when a steady state can be reached. Our next proposition assume no virality ($\kappa = 0$), but it can easily be extended to the case when there is a virality effect for a finite number of periods.[6]

Proposition 2. *Assume that $\kappa = 0$ (no virality), in which case A_t is a constant equal to $A = A_1$. If the deterministic counterpart with state reduction W_t^D admits a finite stationary content strength z_∞, then it is given by the solution to the equation*

$$c'(z_\infty) = m z_\infty + k; \tag{10}$$

here

$$m := p_{\text{ret}} A \left(\frac{E_{\text{new}}^0 E_{\text{ret}}^1}{1-\alpha} + \frac{\delta E_{\text{ret}}^0 E_{\text{new}}^1}{1-\delta\alpha} \right) \quad \text{and} \quad k := p_{\text{new}} A E_{\text{new}}^1,$$

where $E_{\text{new}}^0 := (1-\beta_{\text{new}})^{-1}$, $E_{\text{ret}}^0 := (1-\beta_{\text{ret}})^{-1}$, $E_{\text{ret}}^1 := (1-\delta\beta_{\text{ret}})^{-1}$, and $E_{\text{new}}^1 := (1-\delta\beta_{\text{new}})^{-1}$. Moreover, if $B_1 = r_1 = n_1 = 0$ then the content decisions converge monotonically to z_∞ and the state converges to $B_\infty = AE_{\text{new}}^0/(1-\alpha) z_\infty$, $r_\infty = E_{\text{ret}}^0 z_\infty$, and $n_\infty = E_{\text{new}}^0 z_\infty$.

[6] If $\kappa > 0$ for all periods then $A_\infty = B_\infty = \infty$ and a finite stationary solution does not exist.

Equation (10) shows that the marginal benefit of adding content is equal to a constant k that reflects new traffic plus the multiplicative effect mz_∞ that results from returning visitors. In other words, the optimal policy should internalize the value of stronger (steady state) content by considering the marginal benefit derived from having a larger follower base.

An important particular case occurs when there is no virality ($\kappa = 0$) and the cost $c(z_t)$ is quadratic: $c(z_t) = cz_t^2$. In this case, the deterministic counterpart W_t^D becomes a linear quadratic system (Bertsekas 2000). Hence the optimal policy $z_t^*(x_t)$ is linear in the state $x_t' = (B_t, r_t, n_t)$ and the optimal profit-to-go $W_t^D(x_t)$ is quadratic in that state; these functions can be computed efficiently using a Riccati recursion (see Appendix A.2 for details). Note also that, in the quadratic case, a necessary condition for an interior stationary solution in Proposition 2 is that $c > m/2$, which means that generating content cannot be too cheap. Otherwise the optimal solution is trivial – namely, produce as much content as possible.

3.3 Rule-of-Thumb Heuristic

Here we develop a closed-form heuristic that can also be used whenever the cost is not quadratic. Even in the quadratic case, though, this heuristic has the advantage of not requiring a recursive computation such as the Riccati solution. In what follows we show the derivation of the heuristic when there is no virality ($\kappa = 0$).

Consider the deterministic counterpart W_t^D. The first-order condition for the maximization is

$$c'(z_t) = p_{\text{new}} A + p_{\text{ret}} B_t + \delta \left(A \frac{\partial W_{t+1}^D}{\partial B_{t+1}} + \frac{\partial W_{t+1}^D}{\partial r_{t+1}} + \frac{\partial W_{t+1}^D}{\partial n_{t+1}} \right).$$

If we approximate the partial derivatives by those in the steady state (computed using the envelope theorem), then we obtain

$$c'(z_t) - \frac{\delta p_{\text{ret}} A E_{\text{new}}^1}{1 - \delta\alpha} z_t = k + p_{\text{ret}} E_{\text{ret}}^1 B_t + \frac{\delta p_{\text{ret}} A \beta_{\text{ret}} E_{\text{new}}^1}{1 - \delta\alpha} r_t, \tag{11}$$

where k, E_{ret}^1, and E_{new}^1 are as defined in Proposition 2. Note that if the cost $c(z_t)$ is quadratic, then Equation (11) yields a closed-form heuristic policy z_t that is linear in both B_t and r_t. If the cost is not quadratic, then (11) must be solved in terms of z_t.

In Appendix A.3 we establish an upper bound for the optimal profit-to-go W_t in the case of no virality. This bound can be used to assess the performance of the rule-of-thumb heuristic and of the Riccati solution.

4. Application to Blogs

The objective of this section is to illustrate how our model can, in practice, facilitate better content decisions. Specifically, here we show how the parameters can be estimated from real-world data and how our previous results can yield actionable insights. We apply the model to traffic data from a set of blogs published by faculty at IESE Business School and hosted at blog.iese.edu. For this purpose, we use data from the start of each blog until May 8, 2018.

4.1 Data Description

The site blog.iese.edu hosts 11 blogs with more than 100,000 visits before May 2018. Among these blogs, we decided to study four. Two of them are in English: Expatriatus (412,000 unique visits since March 2011); and Business Ethics (462,000 unique visits since March 2012). The other two are in Spanish: Family Business (255,000 unique visits since October 2010); and Economics, Ethics and CSR (hereafter simply "Economics"; 702,000 unique visits since January 2011).

The data was collected through Google Analytics, which monitors visits to each of the blogs. For every day and every post, we had access to its number of unique views and could further separate these visits into new and returning visitors.[7] In that way, we can distinguish between traffic stemming from the follower base and other traffic. This disaggregated data was obtained by repeatedly running queries through Google Analytics' developer tools website, thereby obtaining exact data rather than sampled data.

The raw data exhibited some traffic seasonality, for which we controlled as follows. First, there were predictable within-week variations (i.e., higher traffic over Monday through Thursday); hence we aggregated traffic by week. Second, there were also monthly variations with lower traffic in August, during Easter week, and around Christmas. Yet since these fluctuations corresponded to lower posting activity, there is no exogenous seasonality beyond that created by the content decisions themselves. In fact, we applied our model to normalized traffic (ratio of blog visits to the total traffic of the www.iese.edu domain, which had qualitatively similar traffic patterns – peaks and valleys in the same weeks – no upward trend) and obtained results that differed little from those when such normalization was not performed.

The aggregate traffic data for each of the four blogs are shown in Figure 2. We can see that each blog is characterized by an increasing number of visits per week. The growth in new visitors

[7] This separation is done according to Google Analytics classification; that is, a visitor is classified as new if there is no cookie associated with the visitor's browser. That approach may create an imperfect measure of visitors' types, as when they have deleted their cookies or browse the Web in "incognito" mode. However, our model indirectly accounts for such leakage through α and κ, thereby taking care of users who are incorrectly classified. See DBS Interactive (2016) for more information.

is due mainly to the presence of more content, since the traffic per post increases only slightly. In contrast, the growth of returning visitors is due to higher traffic per post of the newer posts. This pattern is in line with the assumptions of our model.

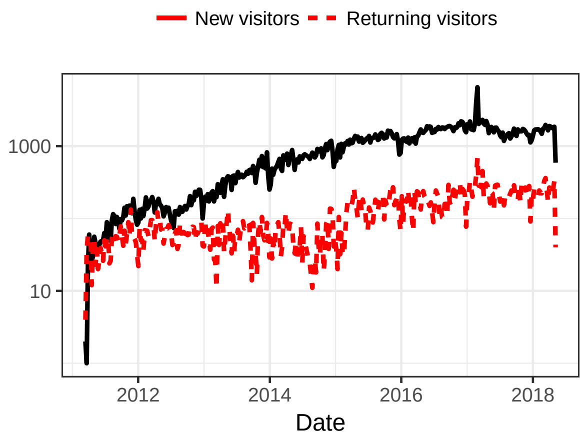

(a) Expatriatus

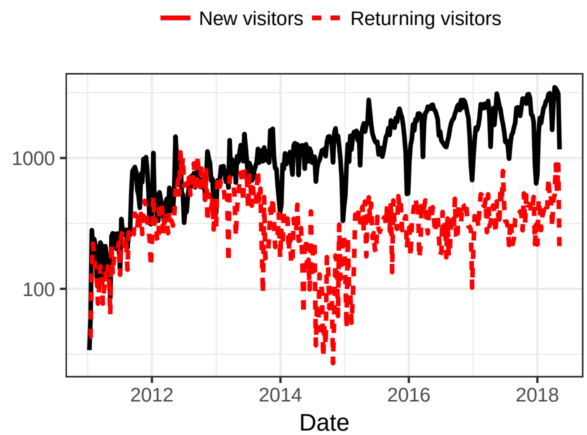

(b) Economics

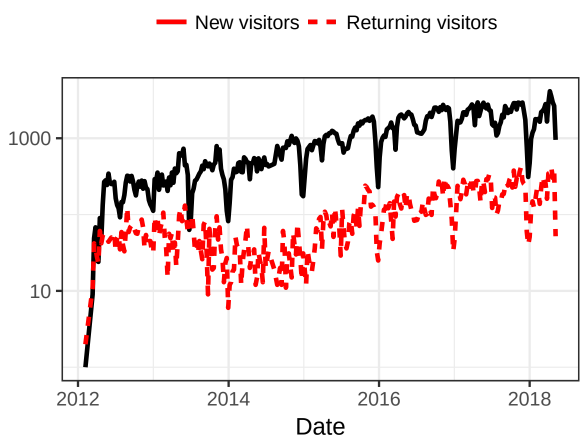

(c) Business Ethics

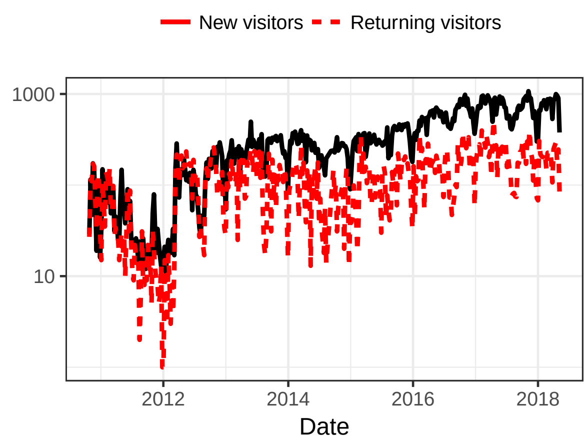

(d) Family Business

Figure 2: Total visits per week for new and returning visitors to the four blogs.

At the disaggregated level – that is, when looking at traffic to posts published in a given week – there is substantial heterogeneity, which can be attributed to the intrinsic uncertainty about quality in the generation of digital content. That is, the audience will find some posts relatively more attractive and will find other posts less so. We apply our model to this disaggregated data, as explained next.

Before proceeding with the estimation, we should mention that our traffic data can be taken as exogenous in the sense that a post's existence and content are not reactions of the blog writer to previous traffic data. There are three main reasons for adopting this view. First, writers of the blog posts did not have access to any traffic statistics and so could hardly have adapted their writing to replicate previously successful posts. Second, the bloggers were business school faculty who were not directly rewarded for blog success; hence they had no incentive to maximize blog traffic. Third, the bloggers usually wrote about their own research and its relation to current events; therefore, the content was usually a "push" of a fixed or slowly changing research agenda. When the content provider is instead reacting to clickstream data, the estimation problem is more nuanced (see Sen and Yildirim 2015).

4.2 Estimation Approach

To capture the heterogeneity among posts, we introduce a fixed effect z_s for a blog posting made at time s. We estimate the equations $r_{st} = B_t z_s \beta_{\text{ret}}^{t-s}$ and $n_{st} = A_t z_s \beta_{\text{new}}^{t-s}$ with respect to the parameters $\zeta_s := \ln(z_s)$, $\phi_{\text{ret}} := \ln(\beta_{\text{ret}})$, $\phi_{\text{new}} := \ln(\beta_{\text{new}})$, α, κ, and A_1. To estimate these parameters, we use the Poisson regression function in the `glm` package of R. Note that the regression cannot be run directly because it uses A_t and B_t, which implicitly depend on α, κ, and A_1 from Equations (1) and (2). However, if we take α and κ as fixed then both A_t and B_t are fixed (with a free scaling factor A_1). In this way, we can estimate ζ_s, ϕ_{ret}, ϕ_{new}, and A_1 for several combinations of α and κ and then determine which yield the best results. We select as the final model the one that minimizes the Akaike information criterion (AIC).

4.3 Results

We estimate the model's parameters for each of the four blogs independently; for this purpose we use the traffic of the first 10 weeks of each post (for all the weeks in which there was a post).[8] The results are summarized in Table 1. To quantify the model's goodness-of-fit, we report its deviance (i.e., twice the difference in log-likelihood between the model and its saturated version, which provides the highest possible likelihood) as compared with the null deviance from model 0 that incorporates only a single variable. We also report the deviance of another benchmark, model $0'$, which features Poisson traffic of rate λ_v (for each visitor type) that is independent of s and t. Finally, we report the AIC value for our model.

The results reported in Table 1 provide a set of metrics that characterize the different blogs and that merit further discussion. First, we find that all blogs have a very high α, which implies that new visitors become followers of the blog and remain followers for a long time (and with a long half-life). Also, the content's amplification effect on the pool of new visitors differs considerably across the blogs. The Expatriatus and Business Ethics content is strongly amplified (values for κ

[8]The visits when a post is older are few and hence extremely erratic, with many zeros. To concentrate on the decline in initial visits, we focus on the first 10 weeks. Similar results are obtained when our estimations are based on the first 5 or 15 weeks.

	Blog			
	Expatriatus	Economics	Business Ethics	Family Business
α	0.9724	0.9966	0.9910	0.9974
κ	0.6892	0.3166	0.8599	0.4713
A_1	1.1	342.3	134.4	28.7
β_{ret}	0.7133	0.3973	0.5804	0.5672
β_{new}	0.7344	0.6145	0.7802	0.6414
Average of ζ_s	−4.63	−4.37	−0.59	−6.28
Number of posts	283	1189	109	282
Number of weeks with new posts	276	383	109	247
Number of weeks	374	383	323	394
Number of observations	112,330	147,075	51,651	91,032
Number of variables	281	388	114	252
AIC	52,395	120,992	18,575	58,336
Deviance	37,110	92,146	12,015	42,593
Deviance model 0	124,849	541,220	41,853	138,776
Deviance model 0′	117,309	541,188	37,735	138,582

Table 1: Estimates and goodness of fit for each of the four blogs.

of 0.69 and 0.86, respectively), which explains the rapid growth in new visitors to these blogs, as plotted in Figure 2. The other two blogs, Economics and Family Business, exhibit less amplification ($\kappa = 0.32$ and 0.47, respectively); this result suggests that, for these blogs, the size of A_t remains relatively more stable over time.

Second, we see that traffic from returning customers tends to decay faster than traffic from new customers (i.e., $\beta_{\text{ret}} < \beta_{\text{new}}$). Here, too, there are differences across blogs. In particular, the content in Economics decays the fastest (visits from returning customers decline by 61%/week and by 39%/week for new visitors); this decline is a reflection of the blog's focus on current economic issues, which are quickly superseded by more recent events. In contrast, content on the other blogs is of a less time-sensitive nature and therefore decays more slowly (for Expatriatus, e.g., visits from returning and new visitors decline by 29%/week and 27%/week, respectively).

Combining these two results allows us to interpret the evolution of blog traffic. For Expatriatus, content is produced in 73% of the weeks in the sample (276/374), and new and returning traffic keeps coming at even growing rates, which implies that traffic per post keeps increasing. On the other hand, for Economics, content is produced every single week in the sample, but traffic is only marginally larger than Expatriatus: this is due to the lesser virality and also the faster decay rates (hence a post does not create the same amount of visits over the life of the content).

Finally, the model's goodness of fit is reasonable: the ratio of its deviance to that of model 0 ranges from 69% to 82%. Given that we are fitting visits over 10 weeks for two streams of customers with few parameters (one per post ζ_s, plus five global parameters $\alpha, \kappa, A_1, \phi^{ret}, \phi^{new}$), the model through its dynamic structure does capture adequately visit variation without requiring post-specific trends.

4.4 Optimizing the Release of Content

In this section we describe a case study of the model undertaken at IESE Business School using data collected through June 2015. The study focuses on Expatriatus and Economics, which are not only the most visited of the four blogs described in Table 1 but also reflect two different blogging styles, as we shall discuss in more detail. The objective of this case study is academic and for purposes of illustration, yet our results reveal the potential of each blog when managed optimally. Such information can be used to detect blogging talent and allocate resources accordingly. The results can also inform bloggers about how best to pace themselves, which is a common challenge when managing online content (O'Neill 2011).

To simplify the calculations, we assume a quadratic cost cz^2 and no virality ($\kappa = 0$). Since the bloggers are business faculty, it is reasonable to assume increasing marginal costs. The parameter estimation under $\kappa = 0$ is strongly similar to that in Table 1, with the main difference being that the follower retention parameter α for Expatriatus is 0.91 instead of 0.97 (for Economics, α remains at 0.99).

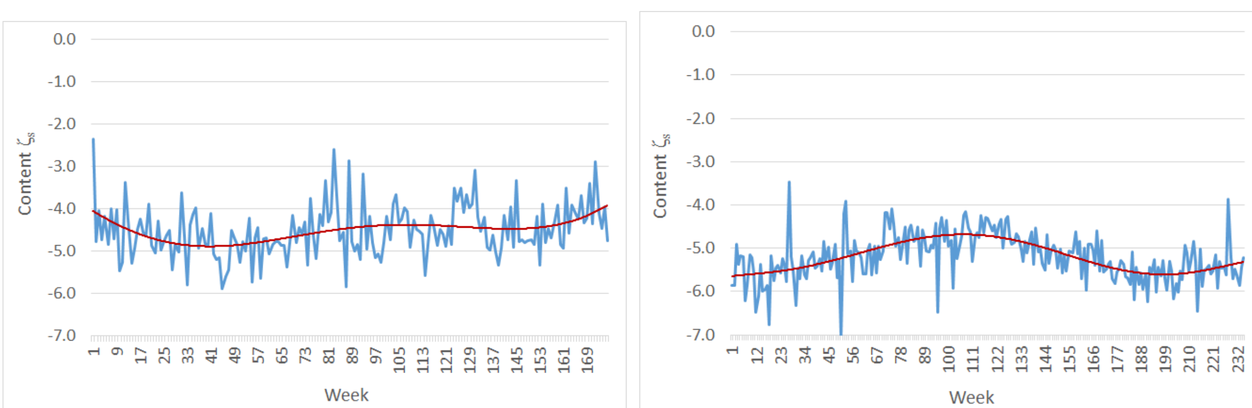

(a) Expatriatus (b) Economics

Figure 3: Variation in ζ_s across posts for two of the blogs; the smooth line is a fitted polynomial of order 6.

Figure 3 plots, for both Expatriatus and Economics, the value of ζ_s for the different posts. Interestingly, the model is able to capture the different behaviors of these two blogs. Namely, Expatriatus largely follows a stable policy but with a slight dip around week 40 and a small surge in the final weeks of the period studied; these estimates accord well with the blog's relatively steady posting frequency. In contrast, the graph for Economics clearly exhibits a dynamic pattern: its activity is below average at first, then above, and finally below again. These estimates are likewise in accordance with the blog's actual activity, since after a ramp-up phase it made many contributions per week and later reduced its pace.

The actual cost of generating the content is not available. We therefore assume a quadratic cost structure and derive the cost of blogging as follows. Let T be the number of weeks in the case study for which data are available: $T = 177$ weeks for Expatriatus and $T = 236$ for Economics. For each blog, let B_t^{cur} denote the expected follower base at time t under their respective *current* content release policies. In particular, we have that $B_T^{\text{cur}} = 525$ for Expatriatus and $B_T^{\text{cur}} = 12{,}632$ for Economics. Now let c^*, which we refer to as the *imputed* blogging cost, be the cost such that the Riccati control introduced in Section 3.2 yields the same (expected) follower base by time T as does the current policy. The assumption that underlies our calculation of c^* is that the blogger was aiming to achieve, by time T, the optimal follower base. In other words, the blogger was thinking optimally in terms of the target follower base B_T^{cur} but adopted a (suboptimal) policy to reaching that goal. Note that we use the infinite-horizon Riccati control because both blogs expected to continue posting content after time T. That said, our results did not change qualitatively when we performed the analysis using the finite-horizon Riccati control (i.e., as if no blogging occurred after time T).

To monetize a blog's traffic, we use an industry average cost-per-thousand (CPM) views of \$6.50 (see Roels and Skowrup 2009). We also assume that new and returning traffic are worth the same, so $p_{\text{new}} = p_{\text{ret}} = 0.0065$. Since the true optimal policy is unknown, instead we use the Riccati control in infinite horizon described in §A.2, which is the optimal policy for the deterministic counterpart W^D. Finally, we set the discount factor δ equal to 0.99.

The upper two graphs in Figure 4 illustrate the content/blogging decisions over time for the actual policy and the Riccati control, where we have plotted the average of 500 replications.[9] The values are normalized by the *constant policy*, under which a fixed amount of content is released in each period in order to reach the target B_T^{cur} by time T. The figure's middle two graphs plot the (expected) follower-base trajectories. By design, all three policies – the actual policy, the Riccati control, and the constant policy – reach the same follower base at time T. The lower two graphs show the cumulative profit generated by these three policies. Those profits have been normalized by the profit attained under the Riccati control at time T.

We find that Expatriatus is blogging less than the optimal amount (except for the very early

[9]We use the smoothed values shown in Figure 3 for the actual content decisions.

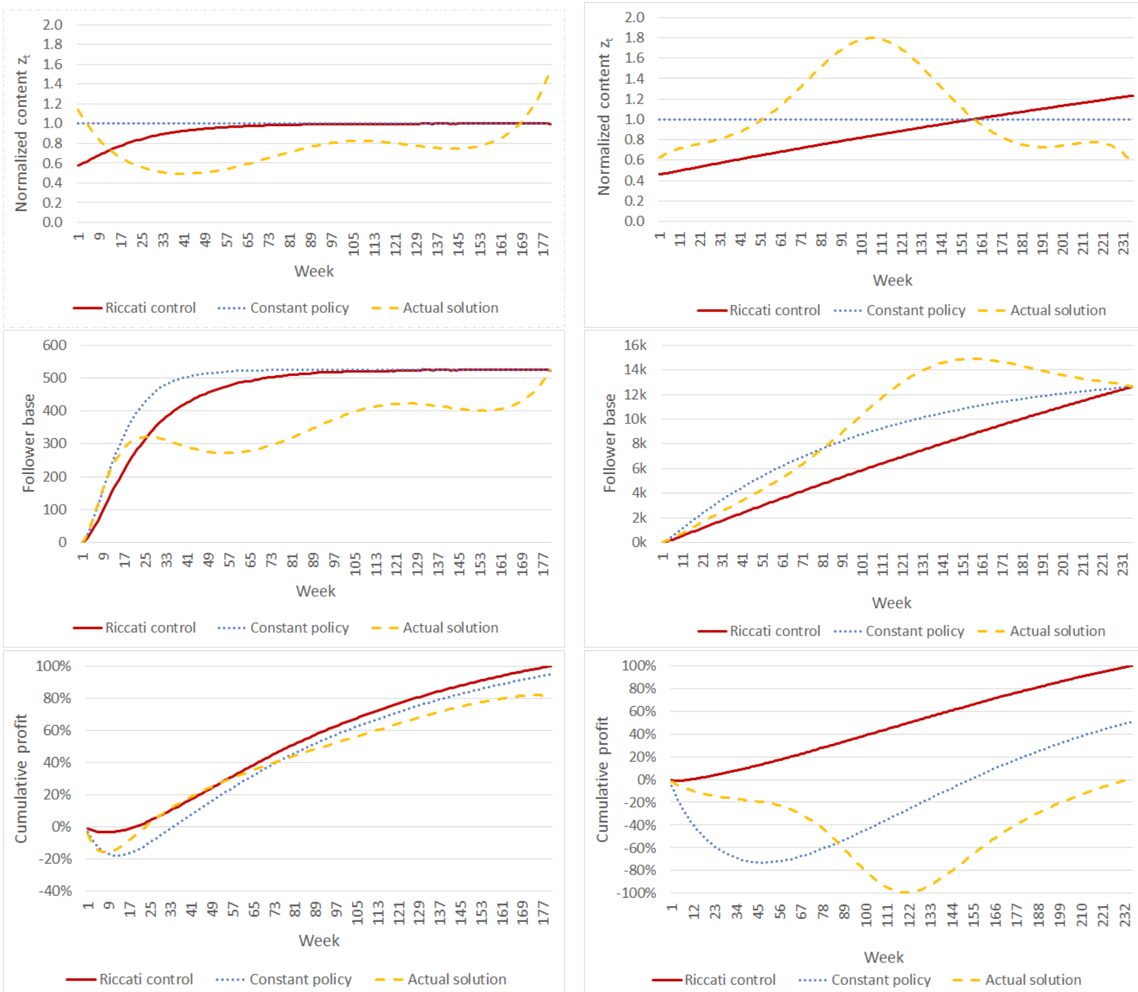

(a) Expatriatus (b) Economics

Figure 4: Normalized content decisions (top figures), follower base trajectories (middle figures), and the corresponding cumulative profits (bottom figures) for the actual policy (dashed line), the Riccati control (solid line), and the constant policy (dotted line).

and late periods) and builds a follower base at a slower pace than it would by adopting the near-optimal Riccati policy. On average, the actual policy released 19% less content than the Riccati control. The impact on profits was significant: the slope at which profits accumulated was smaller under the actual than the Riccati policy because the follower base was also smaller, which reduced the number of returning visitors. By time T, the current policy had underperformed by 20% in terms of profit. Had this blog employed the constant policy, the loss would have been only 5%. As a benchmark, the rule-of-thumb heuristic (c.f. Section 3.3), which is not shown in Figure 4, results in a profit loss of 1% as compared with the Riccati control. It is interesting that, since the follower retention rate of Expatriatus is relatively low ($\alpha = 0.91$), the Riccati policy reaches steady state quickly (at about week 150). We conclude that, for future periods, the Expatriatus blog should adopt the constant policy – which, as expected, corresponds to the steady-state solution z_∞ given in Proposition 2.

In contrast to Expatriatus, Economics systematically blogged more than the optimal amount during the first three years of its existence. Indeed, until week 160 its actual policy released 66% more content (on average) than the Riccati control; in the remaining weeks, the blog released 30% less, on average. Hence, the actual policy did the exact opposite of what is (nearly) optimal: at its inception it blogged too much and for too long; then, after it had amassed a significant amount of followers, it blogged too little. As a result, the blog has barely broken even after more than four years. Here the constant policy would yield better results than the actual policy; however, the constant policy would achieve only about half of the attainable profit and so would not yield as great an improvement as in the case of Expatriatus. Recall that Economics has a high follower retention rate ($\alpha = 0.99$) and therefore reaches steady state far in the future, when profits are much less important owing to the discount factor δ. So following a dynamic policy like the Riccati control (or the simple rule-of-thumb heuristic) pays off handsomely – and more so than when α is low – because most of the profit is accrued while building the follower base (i.e., before reaching steady state). Moreover, the optimal follower-base growth is achieved gradually; in fact, that growth is nearly linear during the first few years.

These numerical results provide some useful guidance for content release. Common wisdom favors releasing abundant content in the early periods to "start with a bang". Yet as suggested by the Riccati policy plotted in Figure 4, it is actually better to start small because exerting additional effort is not worthwhile when there are relatively few followers. As soon as more followers appear, however, the blogger should gradually increase the strength of the content; doing so generates more visits, followers, and profits despite the greater costs. Put simply, every lucky break pushes the content provider to work harder. This dynamic eventually culminates when the size of the follower base reaches the stationary level, after which it is no longer economical to increase the strength of the content releases. It is especially important to follow this gradual growth path when the retention rate (α) is high, as evidenced by the case of Economics. In contrast, if retention is low (i.e., if attrition is high) then a constant policy based on the stationary solution z_∞ performs reasonably well. Finally, in the finite-horizon case, the best approach is to slow down and reduce the content strength as that horizon approaches; the reason is that the value of building up a follower base declines because of the shorter time window over which it can be monetized.

5. Conclusions and Further Research

In this paper we develop a dynamic model for a content provider looking to optimize the release of content over time in the presence of followers. In our model, content has two main functions: it directly generates traffic, and hence profits; and it indirectly attracts new visitors who become followers of the content provider, which in turn leads to future traffic and profits. We formulate the problem as a stochastic DP and characterize the optimal policy's properties. We also develop tractable solution approaches by interpreting the problem as a linear-quadratic Riccati system and by establishing lower and upper bounds. In conjunction with these theoretical results, we study the actual behavior of several blogs and develop an estimation approach for the model parameters. These parameters are then used to optimize blogging activity and provide practical recommendations. Our main insight is that content release should gradually increase over time until an appropriate steady-state level is reached.

There are many directions in which this research can be extended. One immediate possibility would be to interpret our model not in terms of content and a follower base but rather in terms of advertising effort and customer awareness (cf. Pekelman and Sethi 1978). Our results could then be applied in the marketing field and integrated with empirical studies in that literature. Among other applications, the model could be used to support resource allocation decisions.

We must acknowledge that one possible issue with our approach is that it relies on point estimates of the parameters, which serve as input to the optimization model. Those parameters could be highly uncertain a priori, and any estimation based on noisy data could affect identification of the optimal policy. One way to address this concern would be to reformulate the problem using robust optimization (Bertsimas and Thiele 2006). An alternative is to combine the predictive and prescriptive steps into one optimization problem, as in Bertsimas and Kallus (2014) and Elmachtoub and Grigas (2017). That could well be a fruitful avenue of work for scholars to explore.

Finally, another line of research worth mentioning is that, although we used real data to calibrate our model, additional studies are needed so that we can better describe content evolution – especially the decay of its attractiveness. In particular, our model assumes that visits decay exponentially over time. However, it could well be that this decay process (i) is not exactly exponential (we have observed that very old content seems to decay less than does more recent content) and (ii) is driven by controllable factors, such as competition from other content released at the same time. For a similar idea, see Krider and Weinberg (1998) on the movie industry and Caro et al. (2014) on the apparel industry (although neither of those papers considers followers).

Acknowledgments

We are grateful to participants in seminars at Columbia GSB, Cornell University, the University of Florida, Santa Clara University, INSEAD, Harvard Business School, COER (Wharton), UT Dallas, UT Austin, Boston University, and the University of Chile for their helpful feedback. We also thank three anonymous referees and an associate editor for their comments and insights.

References

Afèche, P., M. Araghi, and O. Baron. 2017. Customer acquisition, retention, and service access quality: Optimal advertising, capacity level, and capacity allocation. *Manufacturing & Service Operations Management* 19 (4): 674–691.

Aggarwal, R., R. Gopal, R. Sankaranarayanan, and P. V. Singh. 2012. Blog, blogger, and the firm: Can negative employee posts lead to positive outcomes? *Information Systems Research* 23 (2): 306–322.

Amman, H. M., and H. Neudecker. 1997. Numerical solutions of the algebraic matrix riccati equation. *Journal of Economic Dynamics and Control* 21:363–369.

Bandari, R., S. Asur, and B. A. Huberman. 2012. The pulse of news in social media: Forecasting popularity. In *CoRR*.

Bass, F. M. 1969. A new product growth for model consumer durables. *Management science* 15 (5): 215–227.

Bass, F. M., T. V. Krishnan, and D. C. Jain. 1994. Why the bass model fits without decision variables. *Marketing science* 13 (3): 203–223.

Bernstein, F., and V. Martínez-de-Albéniz. 2017. Dynamic product rotation in the presence of strategic customers. *Management Science* 63 (7): 2092–2107.

Bertsekas, D. P. 1973. Stochastic optimization problems with nondifferentiable cost functionals. *Journal of Optimization Theory and Applications* 12 (2): 218–231.

Bertsekas, D. P. 2000. *Dynamic programming and optimal control*. 2nd ed. Athena Scientific, Belmont MA.

Bertsimas, D., and N. Kallus. 2014. From predictive to prescriptive analytics. *arXiv preprint arXiv:1402.5481*.

Bertsimas, D., and A. Thiele. 2006. A robust optimization approach to inventory theory. *Operations Research* 54 (1): 150–168.

Besbes, O., Y. Gur, and A. Zeevi. 2015. Optimization in online content recommendation services: Beyond click-through rates. *Manufacturing & Service Operations Management* 18 (1): 15–33.

Burns, K. B., H. Enright, J. F. Hayes, K. McLaughlin, and C. Shi. 1997. The art and science of retail renewal. *McKinsey Quarterly* 1:100–113.

Caro, F., V. Martínez-de Albéniz, and P. Rusmevichientong. 2014. The assortment packing problem: Multiperiod assortment planning for short-lived products. *Management Science* 60 (11): 2701–2721.

Chan, T. Y., C. Wu, and Y. Xie. 2011. Measuring the lifetime value of customers acquired from google search advertising. *Marketing Science* 30 (5): 837–850.

Çınar, E., and V. Martínez-de-Albéniz. 2013. A closed-loop approach to dynamic assortment planning. Working paper, IESE Business School.

DBS Interactive 2016, October 4. Google analytics made easy: New visitors vs. returning visitors. https://www.dbswebsite.com/blog/2016/10/04/google-analytics-made-easy-new-visitors-vs-returning-visitors/.

Elmachtoub, A. N., and P. Grigas. 2017. Smart "predict, then optimize". *arXiv preprint arXiv:1710.08005*.

Ghose, A., and S. Yang. 2009. An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Science* 55 (10): 1605–1622.

Goel, S., A. Anderson, J. Hofman, and D. J. Watts. 2015. The structural virality of online diffusion. *Management Science* 62 (1): 180–196.

Goetz, M., J. Leskovec, M. McGlohon, and C. Faloutsos. 2009. Modeling blog dynamics. In *ICWSM*.

Guo, L., E. Tan, S. Chen, X. Zhang, and Y. E. Zhao. 2009. Analyzing patterns of user content generation in online social networks. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, 369–378. ACM.

Huang, W., V. Kulkarni, and J. M. Swaminathan. 2008. Managing the inventory of an item with a replacement warranty. *Management Science* 54 (8): 1441–1452.

Joo, M., K. C. Wilbur, B. Cowgill, and Y. Zhu. 2014. Television advertising and online search. *Management Science* 60 (1): 56–73.

Kendrick, D. 1981. *Stochastic control for economic models*. McGraw-Hill New York.

Krankel, R. M., I. Duenyas, and R. Kapuscinski. 2006. Timing successive product introductions with demand diffusion and stochastic technology improvement. *Manufacturing & Service Operations Management* 8 (2): 119–135.

Krider, R. E., and C. B. Weinberg. 1998. Competitive dynamics and the introduction of new products: The motion picture timing game. *Journal of Marketing Research* 35:1–15.

Leskovec, J., M. McGlohon, C. Faloutsos, N. S. Glance, and M. Hurst. 2007. Patterns of cascading behavior in large blog graphs. In *SDM*, Volume 7, 551–556. SIAM.

Liaukonyte, J., T. Teixeira, and K. C. Wilbur. 2015. Television advertising and online shopping. *Marketing Science* 34 (3): 311–330.

Lobel, I., J. Patel, G. Vulcano, and J. Zhang. 2015. Optimizing product launches in the presence of strategic consumers. *Management Science* 62 (6): 1778–1799.

Martínez-de-Albéniz, V., and A. Valdivia. 2016. The impact of exhibition scheduling on museum attendance. *Manufacturing & Service Operations Management* Forthcoming:NA.

Mediakix 2017, December 15. How much is instagram worth? http://mediakix.com/2017/12/how-much-is-instagram-worth-market-cap/.

Mendelson, H., and K. Moon. 2016. Modeling growth in service operations: Evidence from the app economy. Graduate Business School, Stanford University, working paper.

Moe, W. W., and P. S. Fader. 2004. Capturing evolving visit behavior in clickstream data. *Journal of Interactive Marketing* 18 (1): 5–19.

Nerlove, M., and K. J. Arrow. 1962. Optimal advertising policy under dynamic conditions. *Economica* 29 (114): 129–142.

Olsen, T. L., and R. P. Parker. 2008. Inventory management under market size dynamics. *Management Science* 54 (10): 1805–1821.

O'Neill, M. 2011, January 27. Scheduling your way to YouTube success. http://www.adweek.com/digital/scheduling-youtube/.

Pekelman, D., and S. Sethi. 1978. Advertising budgeting, wearout and copy replacement. *Journal of the Operational Research Society* 29 (7): 651–659.

Roels, G., and T. Skowrup. 2009. Break.com. Technical report, UCLA Anderson School of Management.

Rutz, O. J., and R. E. Bucklin. 2011. From generic to branded: A model of spillover in paid search advertising. *Journal of Marketing Research* 48 (1): 87–102.

Rutz, O. J., and M. Trusov. 2011. Zooming in on paid search ads-a consumer-level model calibrated on aggregated data. *Marketing Science* 30 (5): 789–800.

Satoh, D. 2001. A discrete bass model and its parameter estimation. *Journal of the Operations Research Society of Japan* 44 (1): 1–18.

Sen, A., and P. Yildirim. 2015. Clicks and editorial decisions: How does popularity shape online news converage? The Wharton School, University of Pennsylvania, working paper.

Smith, J. E., and K. F. McCardle. 2002. Structural properties of stochastic dynamic programs. *Operations Research* 50 (5): 796–809.

Sun, M., and F. Zhu. 2013. Ad revenue and content commercialization: Evidence from blogs. *Management Science* 59 (10): 2314–2331.

Topkis, D. M. 1998. *Supermodularity and complementarity.* Princeton, NJ: Princeton University Press.

Wilson, L. O., and J. A. Norton. 1989. Optimal entry timing for a product line extension. *Marketing Science* 8 (1): 1–17.

Appendix

A. Extensions

A.1 General formulation

In the proof of Theorem 1, it is shown that W_t can be written in an expanded way. When:

- the content decision suffers from random shocks and results in strength $\varepsilon_1 z_t + \varepsilon_0$, where $\varepsilon_0, \varepsilon_1 \geq 0$ are random variables;
- the decay rates are $\beta_{\text{ret},s,t} \geq 0$ instead of β_{ret}^{t-s}, and $\beta_{\text{new},s,t} \geq 0$ instead of β_{new}^{t-s};
- the arrival process follows another functional form $\mathcal{Q}(\lambda)$, provided that for all sample paths, $\mathcal{Q}(x \cdot y)$ is increasing and supermodular in (x, y);
- the state variable dynamics are written as:
  $A_{t+1} = \kappa_{0,t} A_t + \kappa_{1,t} \sum_{s=1}^{t} \mathcal{Q}\Big(A_t z_s \beta_{\text{new},s,t}\Big) + \kappa_{2,t} \sum_{s=1}^{t} \mathcal{Q}\Big(B_t z_s \beta_{\text{ret},s,t}\Big)$ and
  $B_{t+1} = \alpha_{0,t} A_t + \alpha_{1,t} \sum_{s=1}^{t} \mathcal{Q}\Big(A_t z_s \beta_{\text{new},s,t}\Big) + \alpha_{2,t} \sum_{s=1}^{t} \mathcal{Q}\Big(B_t z_s \beta_{\text{ret},s,t}\Big)$;

then we can write the value function as

$$V_t(A_t, B_t, z_1, \ldots, z_{t-1}) = \max_{z_t \geq 0} \left\{ \begin{array}{l} p_{\text{ret}} B_t \left(\displaystyle\sum_{s=1}^{t-1} z_s \beta_{\text{ret},s,t} + \mathbb{E}[\varepsilon_1] z_t + \mathbb{E}[\varepsilon_0] \right) \\ + p_{\text{new}} A_t \left(\sum_{s=1}^{t-1} \gamma_s \beta_{\text{news},s,t} + \mathbb{E}[\varepsilon_1] z_t + \mathbb{E}[\varepsilon_0] \right) \\ - c(z_t) + \delta\, \mathbb{E}\left[V_{t+1}\left(A_{t+1}, B_{t+1}, z_1, \ldots, z_{t-1}, \varepsilon_1 z_t + \varepsilon_0 \right) \right] \end{array} \right\} \qquad (12)$$

and $V_{T+1} \equiv 0$, with $T < \infty$. Theorem 1 extends to this formulation as well.

Theorem 2. *For all* $t = 1, \ldots, T$,

(a) V_t *is nondecreasing and supermodular in* $(A_t, B_t, z_1, \ldots, z_{t-1})$ *and convex in* A_t, B_t.

(b) z_t^* *is component-wise nondecreasing in* $(A_t, B_t, z_1, \ldots, z_{t-1})$.

The proof is similar to that of Theorem 1.

A.2 Quadratic Cost and the Riccati Recursion

Consider the deterministic counterpart W_t^D introduced in §3.2 with quadratic cost $c(z_t) = c\,z_t^2$ and $\kappa = 0$. Then, one can write $z_t^*(x_t) = u_t x_t + \eta_t$ and $W_t^D(x_t) = x_t' K_t x_t + q_t x_t + \nu_t$, where $x_t' = (B_t, r_t, n_t)$ is the state, K_t is a 3×3 matrix, u_t and q_t are 1×3 vectors, and η_t and ν_t are scalars. The unknown terms $K_t, u_t, q_t, \eta_t, \nu_t$, can be computed using the following Riccati backward recursion starting in $T+1$ (Kendrick 1981):

$$
\begin{aligned}
K_t &= Q + \delta M' K_{t+1} M + \frac{(R' + 2\delta M' K_{t+1} h)(R + 2\delta h' K_{t+1} M)}{4(c - \delta h' K_{t+1} h)} && (13)\\
q_t &= H + \delta q_{t+1} M + \frac{D + \delta q_{t+1} h}{2(c - \delta h' K_{t+1} h)}(R + 2\delta h' K_{t+1} M) && (14)\\
\nu_t &= \frac{(D + \delta q_{t+1} h)^2}{4(1-\delta)(c - \delta h' K_{t+1} h)} && (15)\\
u_t &= \frac{R + 2\delta h' K_{t+1} M}{2(c - \delta h' K_{t+1} h)} && (16)\\
\eta_t &= \frac{D + \delta q_{t+1} h}{2(c - \delta h' K_{t+1} h)}, && (17)
\end{aligned}
$$

where $M := \begin{pmatrix} \alpha & 0 & A\beta_{\text{new}} \\ 0 & \beta_{\text{ret}} & 0 \\ 0 & 0 & \beta_{\text{new}} \end{pmatrix}$, $h := \begin{pmatrix} A \\ 1 \\ 1 \end{pmatrix}$, $Q := \frac{1}{2}\begin{pmatrix} 0 & p_{\text{ret}}\beta_{\text{ret}} & 0 \\ p_{\text{ret}}\beta_{\text{ret}} & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$, $R := (p_{\text{ret}}, 0, 0)$, $H := (0, 0, p_{\text{new}} A\beta_{\text{new}})$ and $D := p_{\text{new}} A$.

When the horizon is infinite ($t \to \infty$), the optimal profit-to-go is also quadratic $W^D(x) = x'Kx + qx + \nu$, where K and q satisfy the fixed-point equations:

$$
\begin{aligned}
K &= Q + \delta M' K M + \frac{(R' + 2\delta M' K h)(R + 2\delta h' K M)}{4(c - \delta h' K h)} && (18)\\
q &= H + \delta q M + \frac{(D + \delta q h)(R + 2\delta h' K M)}{2(c - \delta h' K h)}. && (19)
\end{aligned}
$$

Since the dynamic program (7) in general is not concave, the finite-horizon recursions (13) and (14) might not converge to the fixed-points K and q (Amman and Neudecker 1997). However, this is not an issue if the indefinite matrix Q is not too dominant, which is likely in practice since β_{ret} is usually lower than β_{new}.[10] Finally, it is worth pointing out that the connection of customer behavior dynamics with the Riccati equation has also been made in other contexts, such as the Bass diffusion model (Satoh 2001).

[10] In §4 we find support for $\beta_{\text{ret}} < \beta_{\text{new}}$, see Table 1.

A.3 An Upper Bound

In this section we obtain an upper bound for the profit-to-go function W_t, for the case when $\kappa = 0$, so that $A_t = A$ (no virality). Applying the envelope theorem on (7) yields[11]

$$\frac{dW_t}{dB_t} = p_{\text{ret}}\left(z_t^* + \beta_{\text{ret}} r_t\right) + \delta\alpha\mathbb{E}\left[\frac{dW_{t+1}}{dB_{t+1}}\Big(\alpha B_t + \mathcal{P}\big(A(z_t^* + \beta_{\text{new}} n_t)\big), z_t^* + \beta_{\text{ret}} r_t, z_t^* + \beta_{\text{new}} n_t\Big)\right]. \quad (20)$$

The recursive structure of Equation (20) allows to derive an upper bound on W_t.

Proposition 3. *Assume that $z_t^* \leq z^{max}$ for all t, B_t, r_t, n_t. Then, for all t:*

(a) $\dfrac{dW_t}{dB_t} \leq N z^{max} + N' r_t$ *with* $N = \dfrac{p_{\text{ret}}}{(1-\delta\alpha)(1-\delta\alpha\beta_{\text{ret}})}$ *and* $N' = \dfrac{p_{\text{ret}}\beta_{\text{ret}}}{1-\delta\alpha\beta_{\text{ret}}}$.

(b) $W_t \leq W_t^{UB}$ *defined as*

$$W_t^{UB}(B_t, r_t, n_t) = \max_{z_t \geq 0}\left\{\begin{array}{l} p_{\text{new}} A\left(1 z_t + \beta_{\text{new}} n_t\right) + p_{\text{ret}} B_t\left(z_t + \beta_{\text{ret}} r_t\right) - c(z_t) \\ +\delta\big(N z^{max} + N'(z_t + \beta_{\text{ret}} r_t)\big)\psi\Big(A(z_t + \beta_{\text{new}} n_t)\Big) \\ +\delta W_{t+1}^{UB}(\alpha B_t + A(z_t + \beta_{\text{new}} n_t), z_t + \beta_{\text{ret}} r_t, z_t + \beta_{\text{new}} n_t) \end{array}\right\} \quad (21)$$

with $W_{T+1}^{UB} \equiv 0$ *and* $\psi(\lambda) = \mathbb{E}\left[\max\left\{0, \mathcal{P}(\lambda) - \lambda\right\}\right]$.

Proposition 3 establishes an upper bound for W_t using a deterministic DP with identical transitions to that of the reduced-state formulation (7), but with a different revenue function. Specifically, revenue is increased by $\delta\big(N z^{max} + N'(z_t + \beta_{\text{ret}} r_t)\big)\psi\Big(A(z_t + \beta_{\text{new}} n_t)\Big)$, which is an upper bound of the additional revenue due to arrival uncertainty (from the Poisson shocks to the number of new followers). Interestingly, one can compute this additional term using that $\psi(\lambda) = \dfrac{\lambda^{k_0+1} e^{-\lambda}}{k_0!}$ with $k_0 = \lfloor\lambda\rfloor$. Finally, note that the existence of z^{max} and the finiteness of the upper bound occurs whenever the cost $c(z)$ is sufficiently convex, e.g., when $\lim_{z \to z^{max}} c(z) = +\infty$.

Since Proposition 3 holds for all t, in the limit we have $W \leq W^{UB}$, where W^{UB} is the infinite horizon problem that follows from (21). We can derive a result similar to Proposition 2 for the dynamic program (21) and W^{UB}.

Proposition 4. *If* (21) *admits a finite stationary solution z^{UB}, it is given by the solution to the equation:*

$$c'(z^{UB}) = k + m z^{UB} + \frac{\delta N'}{1-\delta\beta_{\text{ret}}}\psi + \delta A\left(N z^{max} + N'\frac{z^{UB}}{1-\beta_{\text{ret}}}\right)\left(1 + \beta_{\text{new}}\right)\psi'. \quad (22)$$

where m and k are defined in Proposition 2.

[11] This applies when W_t is finite. In that case, we can exchange the differentiation and expectation operators because W_t is continuously differentiable in B_t for all t. This can be shown by simple induction, because, for any realization of the randomness, both the current profit function $p_{\text{ret}} B_t\left(z_t + \beta_{\text{ret}} r_t\right) + p_{\text{new}} A\left(z_t + \beta_{\text{new}} n_t\right) - c(z_t)$ and the state B_{t+1} are continuously differentiable in B_t, and W_{t+1} is convex in B_t from Theorem 1, see Bertsekas (1973) for more details.

Establishing the overall quality (or tightness) of the upper bound W_t^{UB} is not straightforward. However, for one important special case the gap between the optimal profit-to-go W and the upper bound W^{UB} can be bounded by a closed-form expression. To see this, let $x^{UB} := (B^{UB}, r^{UB}, n^{UB})$ be the (steady) state associated with the stationary solution z^{UB} given by Proposition 4 and let $W^D_{z^{UB}}$ be the profit-to-go of the deterministic counterpart under the stationary policy z^{UB}. Then, we have that:

$$\begin{aligned}\frac{W^{UB}\big(x^{UB}\big) - W\big(x^{UB}\big)}{W^{UB}\big(x^{UB}\big)} &\leq \frac{W^{UB}\big(x^{UB}\big) - W^{D}\big(x^{UB}\big)}{W^{UB}\big(x^{UB}\big)} \\ &\leq \frac{W^{UB}\big(x^{UB}\big) - W^{D}_{z^{UB}}\big(x^{UB}\big)}{W^{UB}\big(x^{UB}\big)} \\ &= \Gamma(z^{UB}) := \frac{\Psi(z^{UB})}{z^{UB}\Delta - c\big(z^{UB}\big) + \Psi(z^{UB})}, \end{aligned} \tag{23}$$

where $\Psi(z^{UB}) := \delta\left(Nz^{max} + N'\frac{z^{UB}}{1-\beta_{\text{ret}}}\right)\psi\Big(\frac{Az^{UB}}{1-\beta_{\text{new}}}\Big)$, $\Delta := A\left(p_{\text{ret}}\frac{E^0_{\text{ret}}}{1-\alpha} + p_{\text{new}}\right)E^0_{\text{new}}$, and $E^0_{\text{ret}}, E^0_{\text{new}}$ are defined in Proposition 2. The first inequality follows from Proposition 1 and the second is due to the optimality of W^D.

A.4 Fixed costs: shutting down might be optimal

In this section, we study the case where there are fixed costs for each period. That is, there is a continuation charge $c_t \geq 0$ in period t if the platform is to continue operating in the next period. This means that the content manager can choose at any point to shut down and save future fixed costs. The charge c_t can represent actual investments needed to set up the platform. For instance, suppose that launching the site requires an investment I. Then, $c_1 = I$ and $c_t = 0$ for $t > 1$. An alternative interpretation of c_t is that it represents the content provider's outside option, i.e., the net present value of the opportunity cost of the blogger's time.

With fixed costs, the formulation of the problem thus changes into:

$$W_t^{stop}(A_t, B_t, r_t, n_t) = \max_{z_t \geq 0}\left\{\begin{array}{l} p_{\text{new}}A_t\left(z_t + \beta_{\text{new}}n_t\right) + p_{\text{ret}}B_t\left(z_t + \beta_{\text{ret}}r_t\right) \\ -c(z_t) + \delta\,\left(\mathbb{E}\big[W_{t+1}^{stop}(A_{t+1}, B_{t+1}, r_{t+1}, n_{t+1})\big] - c_t\right)^+ \end{array}\right\} \tag{24}$$

where $x^+ = \max\{x, 0\}$. Note that one can replace c_t in Equation (24) with $\frac{c_t}{\delta}$ to account for discounting. It turns out that Theorem 1 can be extended to the model with fixed costs.

Theorem 3. *For all t,*

(a) *W_t^{stop} is nondecreasing and supermodular in (A_t, B_t, r_t, n_t) and unidirectionally convex in A_t and B_t.*

(b) z_t^* *is component-wise nondecreasing in* (A_t, B_t, r_t, n_t).

With this result, for every t, there are states in the set $\mathcal{S}_t := \{(A_t, B_t, r_t, n_t) \,|\, W_t^{stop}(A_t, B_t, r_t, n_t) \le c_t\}$ such that it is optimal to shut down. The set $\mathcal{S}_t$ is non-empty when $c_t > 0$, and this means that a low realization of the stochastic processes may lead to a sufficient decrease in A_t and B_t so that it is optimal to close the platform.

B. Proofs

Proof of Theorem 1

Proof. To prove the result, we reformulate the DP into an expanded format, so that the relationship with $z_1, \ldots, z_t$ is apparent. Specifically, we have that

$$W_t(A_t, B_t, r_t, n_t) = V_t(A_t, B_t, z_1, \ldots, z_{t-1})$$

where

$$V_t(A_t, B_t, z_1, \ldots, z_{t-1}) = \max_{z_t \ge 0} \left\{ \begin{array}{l} p_{\text{ret}} B_t \left(\sum_{s=1}^{t} z_s \beta_{\text{ret}}^{t-s} \right) + p_{\text{new}} A_t \left(\sum_{s=1}^{t} z_s \beta_{\text{new}}^{t-s} \right) - c(z_t) \\ +\delta\, \mathbb{E} \left[V_{t+1} \left(\begin{array}{c} A_t + \kappa \sum_{s=1}^{t} \mathcal{P}\Big(A_t z_s \beta_{\text{new}}^{t-s}\Big) + \kappa \sum_{s=1}^{t} \mathcal{P}\Big(B_t z_s \beta_{\text{ret}}^{t-s}\Big), \\ \alpha B_t + \sum_{s=1}^{t} \mathcal{P}\Big(A_t z_s \beta_{\text{new}}^{t-s}\Big), z_1, \ldots, z_t \end{array} \right) \right] \end{array} \right\},$$

The proof builds on Lemma 2.6.4 in Topkis (1998). The central argument is to write $V_{t+1} = g(y_1, y_2, x)$ where $x = (A_t, B_t, z_1, \ldots, z_t)$, $y_1 = f_1(x) = A_t + \kappa \sum_{s=1}^{t} \mathcal{P}\Big(A_t z_s \beta_{\text{new}}^{t-s}\Big) + \kappa \sum_{s=1}^{t} \mathcal{P}\Big(B_t z_s \beta_{\text{ret}}^{t-s}\Big)$, $y_2 = f_2(x) = \alpha B_t + \sum_{s=1}^{t} \mathcal{P}\Big(A_t z_s \beta_{\text{new}}^{t-s}\Big)$. We then establish by induction that f_1 and f_2 are nondecreasing and supermodular in x, while g is supermodular in x and nondecreasing and convex in y_1 (taking y_2 and x fixed) and y_2 (taking y_1 and x fixed).

Formally, we show by induction that V_t is nondecreasing and supermodular in $(A_t, B_t, z_1, \ldots, z_{t-1})$ and convex in A_t and B_t (all other variables fixed).

This is clearly true for the terminal value function $V_{T+1} = 0$.

Assume it is true for $t + 1 \le T + 1$. Consider a sample path: we can write $\varphi_{st,\text{new}} := \mathcal{P}\Big(A_t z_s \beta_{\text{new}}^{t-s}\Big) = \int_0^{A_t z_s \beta_{\text{new}}^{t-s}} d\Lambda_u^{st,\text{new}}$ and $\varphi_{st,\text{ret}} := \mathcal{P}\Big(B_t z_s \beta_{\text{ret}}^{t-s}\Big) = \int_0^{B_t z_s \beta_{\text{ret}}^{t-s}} d\Lambda_u^{st,\text{ret}}$ where $\Lambda_u^{\cdot}$ is a Poisson counting process of rate 1 (i.e., arrivals occur at given times $\ell_1, \ell_2, \ldots$). In that sample path, $\varphi_{st,\text{new}}$ and $\varphi_{st,\text{ret}}$ are nondecreasing and supermodular in $(A_t, B_t, z_1, \ldots, z_t)$. Consider $A_l \le A_h$: $\varphi_{st,\text{new}}(A_h) - \varphi_{st,\text{new}}(A_l) = \int_0^{(A_h - A_l) z_s \beta_{\text{new}}^{t-s}} d\Lambda_u^{st,\text{new}'}$ (for a shifted Poisson process). This is increasing in z_s for all sample paths, hence the chosen one. The argument is identical for $\varphi_{st,\text{ret}}$ and the pair (B_t, z_s).

By the inductive hypothesis, V_{t+1} is nondecreasing and supermodular in $(A_{t+1}, B_{t+1}, z_1, \ldots, z_t)$, and unidirectionally convex in A_{t+1} and B_{t+1}. Applying Lemma 2.6.4 in Topkis (1998) yields that in that sample path,

$$V_{t+1}\left(A_t + \kappa \sum_{s=1}^{t} \varphi_{st,\text{new}} + \kappa \sum_{s=1}^{t} \varphi_{st,\text{ret}}, \alpha B_t + \sum_{s=1}^{t} \varphi_{st,\text{new}}, z_1, \ldots, z_t\right)$$

is supermodular in $(A_t, B_t, z_1, \ldots, z_t)$. It is also nondecreasing. Hence, since supermodularity is preserved by taking sample path averages (Corollary 2.6.2 in Topkis 1998) and addition, we obtain that

$$\begin{aligned} & p_{\text{ret}} B_t \left(\sum_{s=1}^{t} z_s \beta_{\text{ret}}^{t-s}\right) + p_{\text{new}} A_t \left(\sum_{s=1}^{t} z_s \beta_{\text{new}}^{t-s}\right) - c(z_t) \\ & + \delta\, \mathbb{E}\left[V_{t+1} \left(\begin{array}{c} A_t + \kappa \sum_{s=1}^{t} \mathcal{P}\Big(A_t z_s \beta_{\text{new}}^{t-s}\Big) + \kappa \sum_{s=1}^{t} \mathcal{P}\Big(B_t z_s \beta_{\text{ret}}^{t-s}\Big), \\ \alpha B_t + \sum_{s=1}^{t} \mathcal{P}\Big(A_t z_s \beta_{\text{new}}^{t-s}\Big), z_1, \ldots, z_t \end{array} \right) \right] \end{aligned}$$

is nondecreasing and supermodular in $(A_t, B_t, z_1, \ldots, z_t)$. Proposition 4 of Smith and McCardle (2002) establishes that V_t is nondecreasing and supermodular in $(A_t, B_t, z_1, \ldots, z_{t-1})$. As a result, z_t^* is nondecreasing in A_t and B_t.

To complete the induction, recalling that V_t is finite and $\mathcal{C}^2$ (hence differentiation and expectation operators can be exchanged), we use the envelope theorem to establish the convexity of V_t with respect to A_t (the argument for B_t is similar). Namely, letting $\tilde{P} = \mathcal{P}\Big(\sum_{s=1}^{t-1} z_s \beta_{\text{new}}^{t-s} + z_t^*\Big)$, we have that

$$\frac{\partial V_t}{\partial A_t} = p_{\text{new}} \left(\sum_{s=1}^{t-1} z_s \beta_{\text{new}}^{t-s} + z_t^*\right) + \delta\, \mathbb{E}\left[\Big(1 + \kappa \tilde{P}\Big) \frac{\partial V_{t+1}}{\partial A_{t+1}}\right] + \delta\, \mathbb{E}\left[\tilde{P} \frac{\partial V_{t+1}}{\partial B_{t+1}}\right].$$

This is clearly nondecreasing in A_t, because z_t^* increases in A_t, $\tilde{P}$ stochastically increases in A_t and $\frac{\partial V_{t+1}}{\partial A_{t+1}}, \frac{\partial V_{t+1}}{\partial B_{t+1}} \geq 0$, and V_{t+1} is convex in A_{t+1}, B_{t+1} and supermodular in (A_{t+1}, B_{t+1}) (hence all second derivatives are positive).

This completes the induction. Collapsing variables $z_1, \ldots, z_{t-1}$ into r_t, n_t yields the desired result. ■

Proof of Proposition 1

Proof. We can prove the lower bound using Jensen's inequality twice. First, we use the convexity of V_{t+1} (as defined in the proof of Theorem 1) with respect to B_{t+1} from Theorem 1 to show by induction that $V_t \geq V_t^{D-B}$ where V_t^{D-B} is defined by the recursion

$$\begin{aligned} & V_t^{D-B}(A_t, B_t, z_1, \ldots, z_{t-1}) \\ = & \max_{z_t \geq 0} \left\{ \begin{array}{l} p_{\text{ret}} B_t \left(\sum_{s=1}^{t} z_s \beta_{\text{ret}}^{t-s} \right) + p_{\text{new}} A_t \left(\sum_{s=1}^{t} z_s \beta_{\text{new}}^{t-s} \right) - c(z_t) \\ + \delta \left[V_{t+1}^{D-B} \left(A_t + \kappa \sum_{s=1}^{t} \mathcal{P}\Big(A_t z_s \beta_{\text{new}}^{t-s} \Big) + \kappa \sum_{s=1}^{t} \mathcal{P}\Big(B_t z_s \beta_{\text{ret}}^{t-s} \Big), \alpha B_t + A_t \sum_{s=1}^{t} z_s \beta_{\text{new}}^{t-s}, z_1, \ldots, z_t \right) \right] \end{array} \right\}. \end{aligned}$$

Indeed, since

$$\begin{aligned} & V_t(A_t, B_t, z_1, \ldots, z_{t-1}) \\ \geq & \max_{z_t \geq 0} \left\{ \begin{array}{l} p_{\text{ret}} B_t \left(\sum_{s=1}^{t} z_s \beta_{\text{ret}}^{t-s} \right) + p_{\text{new}} A_t \left(\sum_{s=1}^{t} z_s \beta_{\text{new}}^{t-s} \right) - c(z_t) \\ + \delta \left[V_{t+1} \left(A_t + \kappa \sum_{s=1}^{t} \mathcal{P}\Big(A_t z_s \beta_{\text{new}}^{t-s} \Big) + \kappa \sum_{s=1}^{t} \mathcal{P}\Big(B_t z_s \beta_{\text{ret}}^{t-s} \Big), \alpha B_t + A_t \sum_{s=1}^{t} z_s \beta_{\text{new}}^{t-s}, z_1, \ldots, z_t \right) \right] \end{array} \right\}, \end{aligned}$$

and $V_{T+1} = V_{T+1}^{D-B} \equiv 0$, then it follows by induction that $V_t \geq V_t^{D-B}$. Furthermore, it is possible to show (like in Theorem 1) that V_t^{D-B} is convex in A_t. Applying the same argument again we obtain that $V_t^{D-B} \geq V_t^D$. Thus, $V_t \geq V_t^D$, and hence $W_t \geq W_t^D$. ∎

Proof of Proposition 2

Proof. For any constant policy z_∞, according to (3)-(6), the state must converge to $r_\infty = \frac{1}{1-\beta_{\text{ret}}} z_\infty, n_\infty = \frac{1}{1-\beta_{\text{new}}} z_\infty, B_\infty = \frac{A}{(1-\alpha)(1-\beta_{\text{new}})} z_\infty$.

Now suppose an interior stationary solution z_∞ exists. Then it must satisfy the following first-order conditions of the Bellman equation at $(B_\infty, r_\infty, n_\infty)$:

$$c'(z_\infty) = p_{\text{ret}} B_\infty + p_{\text{new}} A + \delta \left(A \frac{\partial W^D}{\partial B} + \frac{\partial W^D}{\partial r} + \frac{\partial W^D}{\partial n} \right). \tag{25}$$

Using the envelope theorem at $(B_\infty, r_\infty, n_\infty)$ we have:

$$\frac{\partial W^D}{\partial B} = p_{\text{ret}} \left(z_\infty + \beta_{\text{ret}} r_\infty \right) + \delta\alpha \frac{\partial W^D}{\partial B}, \quad \frac{\partial W^D}{\partial r} = p_{\text{ret}} B_\infty \beta_{\text{ret}} + \delta \beta_{\text{ret}} \frac{\partial W^D}{\partial r},$$

and

$$\frac{\partial W^D}{\partial n} = p_{\text{new}} A \beta_{\text{new}} + \delta \beta_{\text{new}} \left(A \frac{\partial W^D}{\partial B} + \frac{\partial W^D}{\partial n} \right)$$

Hence, $\frac{\partial W^D}{\partial B} = \frac{p_{\text{ret}}}{(1-\delta\alpha)(1-\beta_{\text{ret}})} z_\infty, \frac{\partial W^D}{\partial r} = \frac{p_{\text{ret}} \beta_{\text{ret}} A}{(1-\beta_{\text{new}})(1-\delta\beta_{\text{ret}})(1-\alpha)} z_\infty$, and $\frac{\partial W^D}{\partial n} = \frac{p_{\text{new}} A \beta_{\text{new}}}{1-\delta\beta_{\text{new}}} + \frac{\delta \beta_{\text{new}} p_{\text{ret}} A}{(1-\beta_{\text{ret}})(1-\delta\alpha)(1-\delta\beta_{\text{new}})} z_\infty$.

Replacing these expressions in (25) and rearranging terms we obtain (10) where

$$m := p_{\text{ret}} A \left[\frac{1}{(1-\alpha)(1-\beta_{\text{new}})(1-\delta\beta_{\text{ret}})} + \frac{\delta}{(1-\delta\alpha)(1-\beta_{\text{ret}})(1-\delta\beta_{\text{new}})} \right]$$

and $k := \dfrac{p_{\text{new}}A}{1-\delta\beta_{\text{new}}}$.

For the convergence starting at $B_1 = r_1 = n_1 = 0$ notice that $B_2 = Az_1 \geq B_1, r_2 = z_1 \geq r_1$ and $n_2 = z_1 \geq n_1$. Hence, from Theorem 1 for the infinite-horizon value function W we have that $z_2^* \geq z_1^*$. Using induction it follows that $B_{t+1} \geq B_t, r_{t+1} \geq r_t, n_{t+1} \geq n_t$, and therefore $z_{t+1}^* \geq z_t^*$, $\forall t \geq 0$. Hence, the state and control paths are monotone nondecreasing sequences. If z_∞ is finite, then z_t^* is bounded and converges to z_∞. Similarly, the state converges to $(B_\infty, r_\infty, n_\infty)$. This completes the proof. ■

Proof of Proposition 3

Proof. We prove the result by induction. If (a) is true at $t+1$ for $N_1, N_2 \geq 0$, then Equation (20) implies that

$$\frac{dW_t}{dB_t} \leq p_{\text{ret}}(z^{max} + \beta_{\text{ret}} r_t) + \delta\alpha(N_1 z^{max} + N_2 z^{max} + N_2 \beta_{\text{ret}} r_t).$$

Hence, the same result is true at t for $N_1 = p_{\text{ret}} + \delta\alpha(N_1 + N_2)$ and $N_2 = p_{\text{ret}}\beta_{\text{ret}} + \delta\alpha\beta_{\text{ret}}N_2$. We can solve for N_1 and N_2 to obtain $N_2 = \dfrac{p_{\text{ret}}\beta_{\text{ret}}}{1-\delta\alpha\beta_{\text{ret}}}$ and $N_1 = \dfrac{p_{\text{ret}} 1 + \delta\alpha N_2}{1-\delta\alpha}$. Given that $dW_{T+1}/dB_{T+1} \equiv 0 \leq N_1 z^{max} + N_2\, r_{T+1}$, result (a) follows for all t.

For part (b), consider a given constant λ and note that for any $\epsilon \leq 0$ we have that $W_{t+1}(\alpha B_t + \lambda + \epsilon, r_{t+1}, n_{t+1}) \leq W_{t+1}(\alpha B_t + \lambda, r_{t+1}, n_{t+1})$, which follows from Theorem 1 because W_{t+1} is nondecreasing in B_{t+1}. On the other hand, part (a) and the fundamental theorem of calculus for a given λ implies that $W_{t+1}(\alpha B_t + \lambda + \epsilon, r_{t+1}, n_{t+1}) \leq W_{t+1}(\alpha B_t + \lambda, r_{t+1}, n_{t+1}) + \epsilon(N_1 z^{max} + N_2\, r_{t+1})$ when $\epsilon \geq 0$. Now let $\lambda = A(z_t + \beta_{\text{new}} n_t)$ and $\epsilon = \mathcal{P}(\lambda) - \lambda$. This yields that $W_t \leq W_t^{UB}$ where

$$W_t^{UB}(B_t, r_t, n_t) = \max_{z_t \geq 0} \left\{ \begin{array}{l} p_{\text{ret}} B_t\,(z_t + \beta_{\text{ret}} r_t) + p_{\text{new}} A\,(z_t + \beta_{\text{new}} n_t) - c(z_t) \\ +\delta W_{t+1}^{UB}(\alpha B_t + A(z_t + \beta_{\text{new}} n_t), z_t + \beta_{\text{ret}} r_t, z_t + \beta_{\text{new}} n_t) \\ +\delta\big(N_1 z^{max} + N_2(z_t + \beta_{\text{ret}} r_t)\big) E\Big[\max\{0, \mathcal{P}(\lambda) - \lambda\}\Big] \end{array} \right\},$$

We define $\psi(\lambda) := E\Big[\max\{0, \mathcal{P}(\lambda) - \lambda\}\Big]$. This generates Equation (21). Result (b) immediately follows by induction. ■

Proof of Proposition 4

Proof. We proceed as in the proof of Proposition 2: for any constant policy z^{UB}, the state converges to $r_\infty = \dfrac{1}{1-\beta_{\text{ret}}} z^{UB}, n_\infty = \dfrac{1}{1-\beta_{\text{new}}} z^{UB}, B_\infty = \dfrac{A}{(1-\alpha)(1-\beta_{\text{new}})} z^{UB}$.

When an interior stationary solution z^{UB} exists, it satisfies first-order conditions:

$$\begin{aligned} c'(z^{UB}) = \;& p_{\text{ret}} B_\infty + p_{\text{new}} A + \delta \left(A \tfrac{\partial W^{UB}}{\partial B} + \tfrac{\partial W^{UB}}{\partial r} + \tfrac{\partial W^{UB}}{\partial n}\right) \\ & +\delta N_2 \psi + \delta\big(N_1 z^{max} + N_2 \tfrac{z^{UB}}{1-\beta_{\text{ret}}}\big) \tfrac{\partial \psi}{\partial z}. \end{aligned} \tag{26}$$

Note that ψ and $\partial\psi/\partial z$ are evaluated at $\left(z^{UB}, \frac{z^{UB}}{1-\beta_{\text{new}}}\right)$. Using the envelope theorem at $(B_\infty, r_\infty, n_\infty)$ we have:

$$\frac{\partial W^{UB}}{\partial B} = p_{\text{ret}}\left(z^{UB} + \beta_{\text{ret}} r_\infty\right) + \delta\alpha\frac{\partial W^{UB}}{\partial B},$$

$$\frac{\partial W^{UB}}{\partial r} = p_{\text{ret}} B_\infty \beta_{\text{ret}} + \delta\beta_{\text{ret}}\frac{\partial W^{UB}}{\partial r} + \delta N_2 \beta_{\text{ret}}\psi,$$

and

$$\frac{\partial W^{UB}}{\partial n} = p_{\text{new}} A \beta_{\text{new}} + \delta\beta_{\text{new}}\left(A\frac{\partial W^{UB}}{\partial B} + \frac{\partial W^{UB}}{\partial n}\right) + \delta\big(N_1 z^{max} + N_2\frac{z^{UB}}{1-\beta_{\text{ret}}}\big)\frac{\partial \psi}{\partial n}.$$

Hence,

$$\frac{\partial W^{UB}}{\partial B} = \frac{p_{\text{ret}}}{(1-\delta\alpha)(1-\beta_{\text{ret}})} z^{UB},$$

$$\frac{\partial W^{UB}}{\partial r} = \frac{1}{(1-\delta\beta_{\text{ret}})}\left[\frac{p_{\text{ret}}\beta_{\text{ret}} A}{(1-\alpha)(1-\beta_{\text{new}})} z^{UB} + \delta N_2 \beta_{\text{ret}}\psi\right],$$

and

$$\frac{\partial W^{UB}}{\partial n} = \frac{1}{(1-\delta\beta_{\text{new}})}\left[\begin{array}{c} p_{\text{new}} A\beta_{\text{new}} + \frac{\delta\beta_{\text{new}} p_{\text{ret}} A}{(1-\delta\alpha)(1-\beta_{\text{ret}})} z^{UB} \\ +\delta\big(N_1 z^{max} + N_2\frac{z^{UB}}{1-\beta_{\text{ret}}}\big)\frac{\partial\psi}{\partial n}\end{array}\right].$$

Replacing these expressions in (26) yields Equation (22), i.e.,

$$c'(z^{UB}) = \quad k + m z^{UB} + \tfrac{\delta N_2}{1-\delta\beta_{\text{ret}}}\psi + \delta\big(N_1 z^{max} + N_2\tfrac{z^{UB}}{1-\beta_{\text{ret}}}\big)\left(\tfrac{\partial\psi}{\partial z} + \tfrac{\delta}{(1-\delta\beta_{\text{new}})}\tfrac{\partial\psi}{\partial n}\right).$$

where k and m are the same as in Proposition 2. ■

Proof of Theorem 3

Proof. This result replicates the proof of Theorem 1, with the additional property that

$$\Big(\mathbb{E}\big[W_{t+1}^{stop}(A_{t+1}, B_{t+1}, r_{t+1}, n_{t+1})\big] - c_t\Big)^+$$

is nondecreasing and supermodular in $(A_{t+1}, B_{t+1}, r_{t+1}, n_{t+1})$, and unidirectionally convex in A_{t+1} and B_{t+1}. Monotonicity and convexity are obviously preserved by the operator $\max\{\cdot, 0\}$. On the other hand, supermodularity is also preserved, as shown next. Assume $f(x,y)$ is supermodular and nondecreasing in both variables. Then, letting $x_1 \le x_2$ and $y_1 \le y_2$, we have $f(x_2,y_1) - f(x_1,y_1) \le f(x_2,y_2) - f(x_1,y_2)$ from supermodularity. Let $f^+ = \max\{f, 0\}$. Because f is nondecreasing, then $f(x_1,y_1) \le f(x_1,y_2) \le f(x_2,y_2)$ and $f(x_1,y_1) \le f(x_2,y_1) \le f(x_2,y_2)$. Assume that $f(x_1,y_2) \le f(x_2,y_1)$ (the other case is similar). Then there are five possible cases:

- If $f(x_1,y_1) \ge 0$ then $f^+(x_2,y_1) - f^+(x_1,y_1) = f(x_2,y_1) - f(x_1,y_1) \le f(x_2,y_2) - f(x_1,y_2) = f^+(x_2,y_2) - f^+(x_1,y_2)$;

- If $f(x_1, y_1) \leq 0 \leq f(x_1, y_2)$ then $f^+(x_2, y_1) - f^+(x_1, y_1) = f(x_2, y_1) \leq f(x_2, y_1) - f(x_1, y_1) \leq f(x_2, y_2) - f(x_1, y_2) = f^+(x_2, y_2) - f^+(x_1, y_2)$;
- If $f(x_1, y_2) \leq 0 \leq f(x_2, y_1)$ then $f^+(x_2, y_1) - f^+(x_1, y_1) = f(x_2, y_1) \leq f(x_2, y_2) = f^+(x_2, y_2) - f^+(x_1, y_2)$;
- If $f(x_2, y_1) \leq 0 \leq f(x_2, y_2)$ then $f^+(x_2, y_1) - f^+(x_1, y_1) = 0 \leq f(x_2, y_2) = f^+(x_2, y_2) - f^+(x_1, y_2)$;
- If $f(x_2, y_2) \leq 0$ then $f^+(x_2, y_1) - f^+(x_1, y_1) = 0 = f^+(x_2, y_2) - f^+(x_1, y_2)$.

This completes the proof. ■