LEARNING TO TAKE A BREAK: SUSTAINABLE OPTIMIZATION OF LONG-TERM USER ENGAGEMENT

Anonymous authors
Paper under double-blind review

ABSTRACT

Optimizing user engagement is a key goal for modern recommendation systems, but blindly pushing users towards increased consumption risks burn-out, churn, or even addictive habits. To promote digital well-being, most platforms now offer a service that periodically prompts users to take a break. These, however, must be set up manually, and so may be suboptimal for both users and the system. In this paper, we propose a framework for optimizing long-term engagement by learning individualized breaking policies. Using Lotka-Volterra dynamics, we model users as acting based on two balancing latent states: *drive*, and *interest*—which must be conserved. We then give an efficient learning algorithm, provide theoretical guarantees, and empirically evaluate its performance on semi-synthetic data.

1 INTRODUCTION

As consumers of content, we have come to rely extensively on algorithmic recommendations. This has made the task of recommending—in a relevant, timely, and personalized manner—key to the success of modern media platforms. Most commercial systems are built with the primary aim of optimizing user *engagement*, a process in which machine learning plays a central role. But alongside their many successes, recommendation systems have also been scrutinized for heedlessly driving users towards excessive and often undesired levels of consumption. This has raised awareness as to the need for redesigning recommendation systems in ways that actively promote digital well-being.

How can media platforms balance business goals with user well-being? One prominent approach, which is now offered by most major platforms, is to periodically prompt users to take breaks (Constine, 2018; Perez, 2018). The idea behind breaks is that occasional disruptions curb the inertial urge for perpetual consumption, and can therefore aid in reducing 'mindless scrolling' (Rauch, 2018), or even addiction (Montag et al., 2018; Ding et al., 2016). As a general means for promoting well-being, breaking is psychologically well-grounded (e.g., Danziger et al., 2011; Sievertsen et al., 2016). But for platforms, breaks serve a utilitarian purpose: their goal is to foster *long-term engagement* by compensating for the myopic nature of conventional recommendation algorithms, which are typically trained to optimize *immediate engagement*. Since breaking schedules are applied heuristically on top of existing recommendation policies—and typically need to be set up manually by users—current solutions unlikely utilize their full potential (Monge Roffarello & De Russis, 2019).

In this paper, we propose a disciplined learning framework for responsible and sustainable optimization of long-term user engagement by controlling breaks. Our point of departure is that *sustained engagement necessitates sustained user well-being*, and here we advocate for breaks as a means to establish both. Focusing on feed-based recommendation, our framework optimizes long-term engagement by learning an optimal *breaking policy* that prescribes individualized breaking schedules.

The challenge in learning to break is that the effects of recommendations on users can slowly accumulate over time, deeming as ineffectual policies that rely on clear signs of over-exposure. To be preemtive, we argue that breaks must be scheduled in a way that anticipates the future trajectory of user behavior, and early on. To achieve this, we introduce a novel class of behavioral models based on Lokta-Volterra (LV) dynamical systems (Lotka, 1910). These depict users as acting based on two balancing forces: *drive to consume* and *intrinsic interest*, with corresponding latent states. Intuitively, high interest increases drive to consume, but prolonged consumption decreases interest;

together, these describe how user behavior varies over time and in response to recommendations. Our model captures the notion that interest can exhaust long before over-consumption is observed. This arms our approach with the persceince needed to prevent burn-out by ensuring that interest is sustainably preserved; thus, whereas current solutions target the symptom—ours aims for the cause.

Our proposed learning algorithm consists of two steps: First, we embed user interaction sequences in 'LV-space'—the set of all possible trajectories that our behavioral model class can express. Then, we optimize individualized breaking policies by solving an optimal control problem over this latent space, in which the control variable is a breaking schedule applied on top of the existing recommendation scheme. Here the challenge is that different breaking policies can lead to different counterfactual trajectories, of which observational data is only partially informative. Since our goal considers long-term outcomes, our solution is to optimize directly for *counterfactual steady-states*. From a behavioral perspective, we view this as aiming to steer towards sustainable habits; from a computational perspective, under our choice of policy class, this enables tractable learning.

As we show, the optimization landscape of LV equilibria admits a compact representation, whose main benefit is that it can be fully described by *predictions* of individualized user engagement rates. Practically, this is advantageous, as it circumvents the need to take arbitrary and costly exploration steps, and enables learning using readily available predictive tools (e.g., Gupta et al., 2006). We make use of a small set of learned predictive models, each trained on a small and minimally-invasive experimental dataset, which allow us to tune our policy to suit different conditions. The final learned policy has an intuitive interpretation: it takes as input a small set of predictions for a user, and via careful interpolation, applies a decision rule that anticipates the effects of breaking on future outcomes (c.f. conventional approaches, which take in predictions and apply the myopic argmax rule).

Our main theoretical result is a bound on the expected long-term engagement of our learned breaking policy, relative to the optimal policy in the class. We show that the gap decomposes into three distinct additive terms: (i) predictive error, (ii) modeling error (i.e., embedding distortion), and (iii) variance around the (theoretical) steady state. These provide an intuitive interpretation of the bound, as well as means to understand the effects of different modeling choices. Our proof technique relies on carefully weaving LV equilibrium analysis within conventional concentration bounds for learning.

Finally, we provide an empirical evaluation of our approach on semi-synthetic data. Using the MoiveLens 1M dataset, we generate discrete time-series data in a way that captures the essence of our behavioral model, but is different from the actual continuous-time dynamics we optimize over. Results show that despite this gap, our approach improves significantly over myopic baselines, and often closely matches an optimal oracle. We also study the role of experimental treatments, and analyze how different users benefit differently from breaks. Taken together, these demonstrate the potential utility of our approach. Code is available at: https://github.com/lvml-iclr-2023/lvml.

Broader perspective. At a high level, our work argues for viewing recommendation as a task of *sustainable resource management*. As other cognitive tasks, engaging with digital content requires the availability of certain cognitive resources—attentional, executive, or emotional. These resources are inherently *limited*, and prolonged engagement depletes them (Kahneman, 1973; Muraven & Baumeister, 2000); this, in turn, can reduce the capacity of key cognitive processes (e.g., perception, attention, memory, self-control, and decision-making), and in the extreme—cause ego depletion (Baumeister et al., 1998) or cognitive fatigue (Mullette-Gillman et al., 2015). As a means to allow resources to replenish, 'mental breaks' have been shown to be highly effective (Bergum & Lehr, 1962; Hennfng et al., 1989; Gilboa et al., 2008; Ross et al., 2014; Helton & Russell, 2017).

Nevertheless, traditional approaches to recommendation remain agnostic to the idea that recommending takes a cognitive toll: they simply recommend at each point in time the item predicted to be most engaging (Robertson, 1977). As an alternative, our approach explicitly models recommendation as a process which draws on these resources, and therefore—must also conserve them. The subclass of 'Predator-Prey' LV dynamics which we draw on are used extensively in ecology for modeling the dynamics of interacting populations, and demonstrate how over-predation can ultimately lead to self-extinction by eliminating the prey population—but also show how enabling resources to naturally replenish ensures sustainable relations. As such, here we advocate for studying recommendation systems as human-centric *ecosystems*, and take one step towards their sustainable design.

1.1 RELATED WORK

User dynamics: latent states and feedback. A recent body of work aims to capture time-varying behavior by modeling users as acting based on dynamic latent states. Broadly, works in this field model the effects of recommendations as either shifting user preferences via positive-only feedback (Jiang et al., 2019; Kalimeris et al., 2021; Sanna Passino et al., 2021; Dean & Morgenstern, 2022), or reducing willingness to consume via negative-only feedback, e.g. via boredom, satiation, or fatigue (Wang & Lin, 2003; Warlop et al., 2018; Kleinberg & Immorlica, 2018; Cao et al., 2020; Leqi et al., 2021). While these restrict behavior to expressing a unidirectional effect, our approach integrates both types of feedfback and models internal states as competing but balancing forces, which we argue is more realistic. This draws connections to recent attempts of injecting psychological modeling into recommendation system design (Kleinberg et al., 2022; Dubey et al., 2022; Curmei et al., 2022). Works in this field often combine theoretical analysis with simulation studies (Schmit & Riquelme, 2018; Chaney et al., 2018; Mansoury et al., 2020; Krauth et al., 2020), and here we follow suit.

Lotka-Volterra dynamics: modeling, learning, and control. The study of ecosystem dynamics and their conservation has a long and rich history, in which LV analysis plays an integral role (see e.g. Hofbauer et al. (1998); Takeuchi (1996)). LV systems our used primarily for modeling biological ecosystems, but are also used in economics (Weibull, 1997; Samuelson, 1998), finance (Farmer, 2002; Scholl et al., 2021), and behavioral modeling (e.g., drug addiction and relapse (Duncan et al., 2019)). In terms of learning, Gorbach et al. (2017) and Ryder et al. (2018) propose variational techniques for dynamical systems, but do not consider control. Our work aims to directly learn optimal policies, for which we draw on recent advances in turnpike optimal control (Trélat & Zuazua, 2015).

2 LEARNING SETTING

We consider a sequential recommendation setting in which users interact with a stream of recommended items over time. New users u are sampled iid from some unknown distribution D, and time for each user is measured relative to their time of joining. Interactions occur at discrete time-points in continuous time, $t \in \mathbb{R}_+$, and upon user request: when a user u queries the system for additional content at time t, the system responds by presenting a recommended item $x(t) \in \mathcal{X}$. We assume recommendations are governed by an existing and fixed *recommendation policy* ψ, which determines $x(t)$ given a request from u at time t. We allow ψ to be stochastic, and make no additional assumptions on its structure or mechanics. User interactions are therefore described by a sequence of pairs $\{(t_i, x_i)\}_i$, where $x_i = x(t_i) \sim \psi(u; t_i) \in \mathcal{X}$ is the item recommended to u at time t_i.

Our key modeling assumption is that subsequent request times, $t_{i+1} = t_i + \Delta t_i$ for $\Delta t_i > 0$, are determined jointly by the user's 'state' at time t_i and the recommended item x_i (note this means Δt_i can depend on t_i). We consider user states as latent and in the abstract, but broadly expect 'good' recommendations to entail frequent requests by inducing small values of Δt_i. Together, user u's choice behavior, coupled with the policy ψ, induce a temporal point process (TPP) which governs the generation of interaction sequences of duration t as $\{(t_i, x_i) \mid t_i \in [0, t]\} \sim \mathrm{TPP}_\psi(u; t)$.

For each user, the system collects data until some (relative) fixed time T_0, and seeks to optimize engagement in the subsequent interval $[T_0, T_0 + T]$, where T is a predetermined time horizon. We denote the corresponding input sequences by $S_u^0 = \{(t_i, x_i) \mid t_i \in [0, T_0)\}$, and target sequences by $S_u = \{(t_i, x_i) \mid t_i \in [T_0, T]\}$. Defining by $\frac{1}{T}|S_u|$ the *engagement rate* of u for the chosen time horizon T, our goal in learning will be to maximize expected *long-term user engagement rate*, namely $\mathbb{E}_{u\sim D}\, \mathbb{E}\left[\frac{1}{T}|S_u|\right]$ for $(S_u^0, S_u) \sim \mathrm{TPP}_\psi(u; T_0 + T)$, in expectation over new users.

Breaking policies. The unique aspect of our learning problem is that our only means for increasing engagement is by suggesting *breaks*. Our task will be to learn a *breaking policy* $\pi \in \Pi$ which can override ψ: when user u queries the system at time t_i, the policy $\pi(u; t_i) \in \{0, 1\}$ determines whether to present the intended item $x(t) \sim \psi(u; t_i)$ (for $\pi = 0$), or suggest a break instead ($\pi = 1$).[1] Denoting the overall composed policy by $\pi \circ \psi$, our aim is to learn a breaking policy π that increases engagement by complementing an existing ψ. Hence, our learning objective is:

$$\arg\max_{\pi\in\Pi} \mathbb{E}_{u\sim D}\, \mathbb{E}_{\mathrm{TPP}}\left[\tfrac{1}{T}|S_u|\right], \qquad (S_u^0, S_u) \sim \mathrm{TPP}_{\pi\circ\psi}(u; T_0 + T) \tag{1}$$

[1] This is similar in spirit to the 'learning to defer' paradigm (Madras et al., 2018), but in a different context.

Broadly, we expect breaks to negatively affect short-term engagement (i.e., entail longer ensuing Δt_i), but have the potential to improve engagement in the long run if scheduled appropriately.

Data and exploration. For learning an optimal breaking policy, we assume the system has access to a dataset of previously logged interactions $\mathcal{S} = \{(S_1^0, S_1), \ldots, (S_m^0, S_m)\}$ for m train-time users $u_1, \ldots, u_m \sim D$, collected under a 'clean' recommendation policy ψ. It will be convenient to assume that the system 'featurizes' user input sequences S_u^0 via a vector mapping ϕ used in learning. For clarity, we overload notation and represent users as $u = \phi(S_u^0) \in \mathbb{R}^d$. This allows us to support additional user feedback (e.g., ratings) or information (e.g., demographics) as input to ϕ.

Since Eq. (1) is a policy problem (note breaks affect outcomes), learning requires some form of exploration or experimentation. Here we aim for experimentation to be simple and minimal. Specifically, we assume access to a small number of N additional datasets, $\mathcal{S}^{(j)} = \{(u_k, S_k)\}_{k=1}^{m_j}$, collected for different users, and under composed policies $\pi_j \circ \psi$ for various predetermined π_j. These datasets can either be given at the onset, or collected as part of the learning process; our only usage of them will be for learning predictors $\hat{y} = f_j(u)$ of engagement rate $y = \frac{1}{T}|S_u|$. Ideally, we would like to make do with only a few $\mathcal{S}^{(j)}$ of small size and that can be gathered concurrently; our results show that even a single additional dataset can be highly informative, and in some cases sufficient.

Learning to break. To effectively optimize Eq. (1), learning must anticipate the effects of different breaking policies $\pi \in \Pi$ on future sequences S_u of unobserved users u. It will therefore be useful to distinguish between the policy π itself, which determines when breaks are applied, and a component for estimating individualized counterfactual engagement rates $\frac{1}{T}|\hat{S}_u(\pi)|$ induced by π, which will aid in choosing a good policy. Our focus will be on learning simple policies coupled with rich and task-appropriate models of responsive user behavior. In particular, we set Π to include all *stationary policies*, $\pi(u) = \pi(p_u)$, which for each user u recommends a break with a time-independent personalized probability $p_u = p(u) \in [0, 1]$. Stationary policies are interpretable, amenable to efficient optimization, and straightforward to implement. As we show, despite their simplicity, such policies can be quite expressive when coupled with an appropriate behavioral model.

3 Engagement via 'Predator-Prey' dynamics

To optimize engagement with breaks, we propose to model user behavior as a dynamical system of Lotka-Volterra (LV) Predator-Prey equations. The model operates over continuous dynamics, which has analytic and optimizational benefits; we later make the connection back to discrete sequences.

Behavioral model. Our model of user behavior views each user as acting based on two types of time-dependent latent variables: internal *drive* for consumption, denoted by $\lambda(t) \in \mathbb{R}_+$; and intrinsic *interest*, denoted by $q(t) \in [0, 1]$. Intuitively, we expect that engaging experiences will reinforce a user's desire to further consume; conversely, excessive exposure to content should slowly 'erode' her interest—but left alone, will allow it to replenish (Thoman et al., 2011). Thus, drive and interest act as balancing forces. Together, we model the time-dependent relations between $\lambda(t)$ and $q(t)$ as:

$$\frac{\mathrm{d}\lambda}{\mathrm{d}t} = -\alpha\lambda + \beta\lambda q, \qquad \frac{\mathrm{d}q}{\mathrm{d}t} = \gamma q(1-q) - \delta\lambda q \tag{2}$$

where $\theta = (\alpha, \beta, \gamma, \delta) \geq 0$ parameterize the dynamics. For λ (drive, or 'predator'), α determines its natural decay rate, and β its interest-dependent self-reinforcement rate; for q (interest, or 'prey'), γ specifies its natural replenishment rate, and δ its rate of depletion from consumption. Note the two equations are coupled: q mediates the reinforcement of λ, and λ mediates the depletion of q.

Dynamical properties. The LV model in Eq. (2) describes consumption as a cycle: when interest $q(t)$ is high, drive to consume $\lambda(t)$ increases, resulting in positive feedback; conversely, when $\lambda(t)$ is high, $q(t)$ decreases, which expresses negative feedback. In general, λ grows until interest is too low to sustain consumption, at which point consumption drops sharply, allowing interest to recover—and the cycle repeats. The cycling behavior exhibits oscillations in λ and q, with one lagging after the other. A typical trajectory is illustrated in Figure 1. Note how the drop in λ occurs only some time after q is depleted; hence, anticipating (and preventing) the collapse of λ requires knowledge (and conservation) of q. Thus, q serves as a resource: necessary for consumption, and of limited supply.

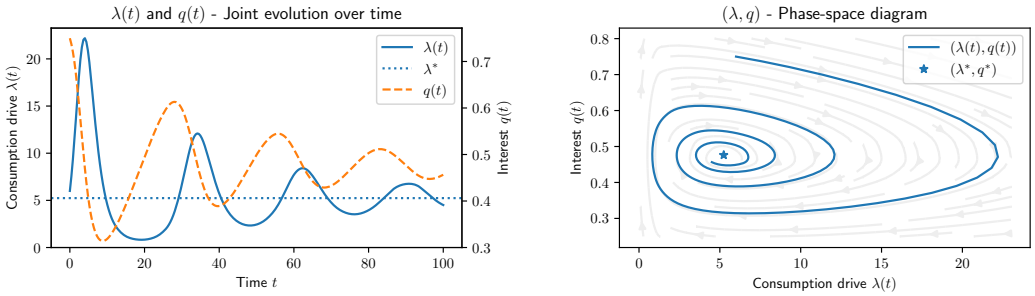

Figure 1: Characteristic LV dynamics (Eq. (2)). **(Left)** Temporal relations between *consumption drive* $\lambda(t)$, *intrinsic interest* $q(t)$, and equilibrium λ^*. Note how λ drops only some time after q has depleted. **(Right)** Evolution of the same system in phase space (λ, q), with its equilibrium (λ^*, q^*).

Over time, and if no interventions are applied, the magnitude of oscillations decreases, and the system naturally approaches a stable equilibrium, denoted (λ^*, q^*), determined by the system parameters $\theta = (\alpha, \beta, \gamma, \delta)$ and which attracts all initial conditions $\lambda(0), q(0)$ (Takeuchi, 1996). Equilibrium plays a key role in our approach, as it captures the notion of habits, which we aim to improve.

Engagement maximization as optimal control. To optimize engagement (Eq. (1)), we propose to optimize λ^* as an alternative proxy. Broadly, our approach will be to associate with each user u a dynamical system parameterized by $\theta_u = (\alpha_u, \beta_u, \gamma_u, \delta_u)$, and recommend breaks which lead to high values of the corresponding $\lambda_u^* = \lambda_{\theta_u}^*$. Intuitively, if we think of $\frac{1}{T}|S_u|$ as an 'empirical' rate (determined by the Δt_i in S_u), then λ_u^* is its continuous theoretical counterpart, and so ideally we would like to find a θ_u for which λ_u^* is the limiting behavior of $\frac{1}{T}|S_u|$ (i.e., when $\Delta t \to 0$ and $T \to \infty$). In practice, we expect λ_u^* to be a useful target when the observed $\frac{1}{T}|S_u|$ exhibits habits that are 'close enough' to the theoretical equilibrium, and in Sec. 4.3 we make this condition precise.

Since our goal is to optimize a breaking policy, we must also be precise about the way breaks affect dynamics in our model. Towards this, we introduce into Eq. (2) a *control variable*, $p(t) \in [0, 1]$, which acts as a 'gate' that controls the mediation strength between $\lambda(t)$ and $q(t)$ via:

$$\frac{\mathrm{d}\lambda}{\mathrm{d}t} = -\alpha\lambda + \beta\lambda q(1-p), \qquad \frac{\mathrm{d}q}{\mathrm{d}t} = \gamma q(1-q) - \delta\lambda q(1-p) \tag{3}$$

When $p > 0$, breaking has a dual effect: it decelerates drive λ, and at the same time, lets q recover.

Our goal will now be to solve the optimal control problem of finding $p(t)$ which maximizes λ^*. Note that not every dynamic controlling schedule guarantees convergence to some λ^*. Hence, we focus on *fixed controls*, $p \in [0, 1]$, which we prove converge to equilibrium, and having *closed form*.

Lemma 1. *Let $\theta = (\alpha, \beta, \gamma, \delta)$ define an LV system as in Eq. (3). Then for any $p \in [0, 1]$, we have:*

$$\lambda_\theta^*(p) = \frac{\gamma}{\delta}\frac{1}{1-p}\left(1 - \frac{\alpha}{\beta}\frac{1}{1-p}\right) \quad \text{if } p \in [0, 1-\alpha/\beta]\,, \text{ and zero otherwise.} \tag{4}$$

Proof in Appendix A.1. For optimization, Lemma 1 is useful since it depicts λ^* as a function of p, parameterized by θ. Thus, given some θ_u for user u, our objective is to solve the control problem:

$$p_u = \arg\max_{p \in [0,1]} \lambda_{\theta_u}^*(p) \tag{5}$$

Optimizing engagement now reduces to solving Eq. (5); given p_u, we make the connection back to discrete time by setting the breaking policy to be $\pi(p_u)$, which recommends breaks at rate p_u.

4 Learning Optimal Breaking Policies

We now turn to presenting our learning algorithm. Our approach to optimizing engagement consists of two steps: (i) associating with each user u a set of LV parameters θ_u, and then (ii) finding p_u which maximizes $\lambda_{\theta_u}^*(p)$, and plugging into $\pi(p_u)$. In practice, we add an intermediate prediction step, which allows us to 'shortcut' directly from observations to optimal policies. We conclude with analysis showing when learned policies $\pi(p_u)$ lead to good expected engagement $\mathbb{E}_{u\sim D}\,\mathbb{E}_{\mathrm{TPP}}\left[\frac{1}{T}|S_u|\right]$.

4.1 EMBEDDING USERS IN LV SPACE

To find θ_u, a seemingly reasonable approach would be to fit an LV trajectory to the initial sequence S_u^0, i.e., by solving $\min_\theta \sum_i |\Delta t_i - \lambda_\theta(t_i)|$ for the observed $t_i \in S_u^0$. This can be interpreted as embedding S_u^0 in 'LV-space' by finding the nearest continuous trajectory, for which θ provides a compact representation. However, a key issue with this approach is that S_u^0 contains *past* observations made under a *single* policy $\pi(p)$ (e.g., the 'clean' policy $\pi(0)$). To optimize *future* engagement, the learned θ_u must correctly account for the affects of *general* p on possbile ensuing sequences S_u.

As an alternative, we propose to find θ_u by fitting the entire *equilibrium curve* of $\lambda_\theta^*(p)$ (Eq. (4)). Formally, for each p, define the *expected empirical engagement rate* $\bar{\lambda}_u(p)$ as:

$$\bar{\lambda}_u(p) = \mathbb{E}_\pi\left[\tfrac{1}{T}|S_u|\right], \qquad S_u \sim \mathrm{TPP}_{\pi(p)\circ\psi}(u;T) \tag{6}$$

i.e., $\bar{\lambda}$ is the rate of *counterfactual* future trajectories for all possible choices of p. Ideally, we would like to find θ for which the learned curve $\lambda_\theta^*(p)$ closely aligns with that of $\bar{\lambda}_u(p)$ across $p \in [0,1]$:

$$\bar{\theta}_u = \arg\min_\theta \|\bar{\lambda}_u - \lambda_\theta^*\| \tag{7}$$

for some function norm $\|\cdot\|$, and for which $\lambda^*_{\bar{\theta}_u}$ and $\bar{\lambda}_u$ have similar maximizing p (since we aim for optimizing λ_θ^* to work well for $\bar{\lambda}_u$). Unfortunately, $\bar{\lambda}_u$ is a theoretical construct, and so $\bar{\theta}_u$ cannot be obtained from observations. Hence, we propose to replace $\bar{\lambda}_u$ with a finite set of *predictions*.

The role of prediction. Recall our input consists of a primary dataset $\mathcal{S}^{(0)} = \mathcal{S} = \{(u_i, S_i)\}_{i=1}^m$, as well as additional 'experimental' datasets $\mathcal{S}^{(j)}$ collected under different $\pi_j = \pi(p_j)$ and of sizes m_j for $j = 1, \ldots, N$. We can use these to learn individualized, policy-specific predictors $f_j(u) = f_{p_j}(u)$, trained to predict for each user u her engagement rate $y = \frac{1}{T}|S_u|$ under π_j. For example, if we train f_j to minimize the squared error $\sum_k (f_j(u_k) - y_k)^2$ on pairs $(u_k, y_k) \in S^{(j)}$, then $f_j(u)$ should be a reasonable estimator of the expected $\bar{\lambda}_u(p_j)$. Hence, for a given u, a finite set of pairs $\{(p_j, f_j(u))\}_{j=1}^N$ gives points to which we can fit θ to λ_θ^*. Our final criterion for choosing $\hat{\theta}_u$ is:

$$\hat{\theta}_u = \arg\min_\theta \|\boldsymbol{f}(u) - \lambda_\theta^*\| = \arg\min_\theta \sum_{j=1}^N \left(f_j(u) - \lambda_\theta^*(p_j)\right)^2 \tag{8}$$

given here with the ℓ_2 vector norm, and where $\boldsymbol{f}(u) = (f_1(u), f_2(u), \ldots, f_N(u)) \in \mathbb{R}_+^N$. As we will see next, optimizing over p requires only the ratios α/β and γ/δ, in which λ^* is quadratic. Hence, Eq. (8) can be efficiently solved using a polynomial Non-Negative Least Squares (NNLS) regression solver (Chen & Plemmons, 2010).

The role of experimentation. Two parameters control the goodness of fit for $\hat{\theta}_u$: the number of experimental datasets, N, and their sizes, m_j for $j \in [N]$. In general, increasing N provides more data points for solving Eq. (8), and increasing each m_j reduces noise for that point (i.e., $f_u(p)$ should be closer to $\bar{\lambda}_u$). But experimentation is costly, and so in reality N and the m_j may be small. As motivation, we next show that under realizability and for accurate predictions, $N = 2$ suffices.

Proposition 1. *Fix $N = 2$, and let $p_0, p_1 \in [0, 1 - \alpha/\beta]$. For a user u, if (i) exists θ_u s.t. $\bar{\lambda}_u = \lambda^*_{\theta_u}$, and (ii) $f_i(u) = \bar{\lambda}_u(p_i)$ for $i = 1, 2$, then $\hat{\theta}_u$ is optimal, i.e., solving Eq. (8) gives $\hat{\theta}_u = \bar{\theta}_u$.*

Proof in Appendix A.1, and relies on Lemma 1. Next, we discuss how to obtain $\pi(u)$ from θ_u.

4.2 FROM PREDICTIONS TO OPTIMAL POLICIES

Recall our goal is to learn an *individualized* breaking policy π that for each user u applies an appropriate, personalized breaking schedule $\pi(u)$. To be able to generalize to unseen u, the conventional approach is to learn a parameterized policy $\pi(u;\eta)$, where η is optimized on training data, and applied to new users at test-time. Our approach, by relying on predictions, circumvents the need to learn parameterized policies: once the predictors $\{f_j(u)\}_j$ have been learned, the policy problem decomposes over users, and optimal policies $\pi(u)$ are determined independently for each user.

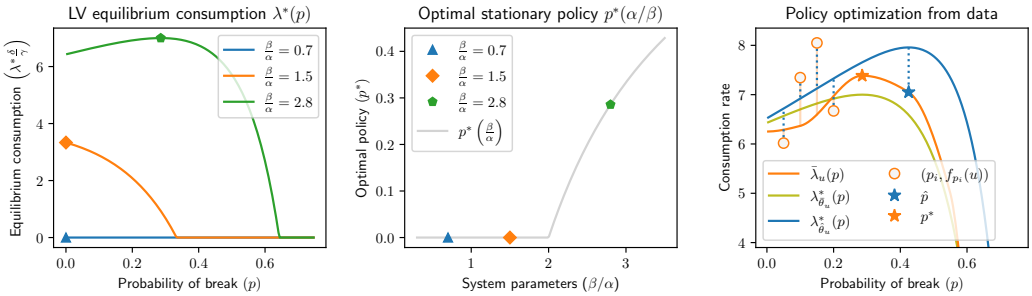

Figure 2: **(Left)** Equilibrium curves $\lambda^*(p)$ and optimal policies p^* (*markers*) for user types (θ_u) that: benefit from breaks (*green*), do not require breaks (*orange*), and will inevitably churn (*blue*). **(Center)** The optimal policy p^* for all β/α, showing a second-order phase change at $\beta/\alpha = 2$. **(Right)** An illustration of the true counterfactural engagement curve (*orange*),

Optimal policies. Given θ_u, our next result establishes a closed form solution for the optimal p_u. Note that Eq. (4) shows $\lambda^*_\theta(p)$ is piecewise-polynomial in $z = 1/(1-p)$. Solving for z, we get:

Lemma 2. *Let $\theta = (\alpha, \beta, \gamma, \delta)$ define an LV system as in Eq. (3). Then the optimal p^* is given by:*

$$p^*(\theta) = \arg\max_{p\in[0,1]} \lambda^*_\theta(p) = 1 - 2\frac{\alpha}{\beta} \quad \text{if} \quad \frac{\alpha}{\beta} \leq \frac{1}{2}, \text{ and zero otherwise.} \tag{9}$$

Proof in Appendix A.1. Hence, once θ_u is found, our learned policy is defined as:

$$\pi(u) = \pi(\hat{p}_u), \qquad \text{where } \hat{p}_u = p^*(\hat{\theta}_u) \tag{10}$$

Altogether, to compute $\hat{p}_u$, the least squares approach in Eq. (8) is applied to obtain the polynomial coefficients $\left(\frac{\gamma_u}{\delta_u}, \frac{\alpha_u\gamma_u}{\beta_u\delta_u}\right)$, then $\frac{\alpha_u}{\beta_u}$ is estimated from their ratio, and then plugged into Eq. (9) to obtain $\hat{p}_u$. For $N = 2$, $\pi(u)$ has a closed-form formulation as a function of predictions (Appendix A.2). Additionally, note that by Eq. (10), $\hat{p}_u$ exhibits a phase transition at $\frac{\alpha_u}{\beta_u} = \frac{1}{2}$, below which $\hat{p}_u = 0$.

Corollary 1. *In LV space, $\pi(u)$ partitions users to those who require breaks, and those who don't.*

Corollary 2. *In the realizable case of Prop. 1, $\pi(\hat{p}_u)$ idempodently improves over the myopic $\pi(0)$.*

Thus, the optimal policy can be interpreted as suggesting breaks only when it deems them necessary. Figure 2 illustrates $\lambda^*(p)$ curves, phase shift, and optimal and learned policies for various user types.

4.3 Theoretical Guarantees

We are now ready to state our main theoretical result, which bounds the expected long-term engagement obtained by our learned policy, $\hat{\pi}$. Our bound shows that the gap between $\hat{\pi}$ and the optimal policy, π^{opt}, is goverened by three additive terms, each relating to a different aspect of our approach: modeling error (ε_{LV}), predictive error ($\varepsilon_{\text{pred}}$), and deviation from expected behavior (ε_{dev}). A description and interpretation of each term follows shortly. For simplicity, we focus on $N = 2$.

Theorem 1 (Informal)**.** *Let $p_0, p_1 \in [0, 1]$, and denote by $\pi^{\text{opt}} \in \Pi$ be the optimal stationary policy. Then for the learned breaking policy $\hat{\pi}$, we have:*

$$\mathbb{E}_{u,\pi^{\text{opt}}}\left[\tfrac{1}{T}|S_u|\right] - \mathbb{E}_{u,\hat{\pi}}\left[\tfrac{1}{T}|S_u|\right] \leq \frac{\eta_{\text{TPP}}}{|p_1 - p_0|}(\varepsilon_{\text{LV}} + \varepsilon_{\text{pred}} + \varepsilon_{\text{dev}})$$

where η_{TPP} is a TPP-specific constant scale factor.

Formal statement, precise definitions, and proof are given in Appendix A.3. The proof consists of three main steps: We begin with a clean LV system at $T = \infty$, and quantify the downstream effect of perturbing the optimal policy. Then, we plug in the learned policy, and bound the gap due to predictive errors and finite T. The final step makes the transition from continuous dynamics to the discrete TPP. We next detail the role of each of the five terms in the bound, and how they can be controlled.

- **Predictive error**: Since targets $y = \frac{1}{T}|S_u|$ are real, $\varepsilon_{\text{pred}}$ is simply the expected regression error over users, measured here in RMSE. As is standard, $\varepsilon_{\text{pred}}$ can be reduced by increasing the number of samples m, or by learning more expressive predictive functions f (e.g., larger neural nets).
- **Modeling error**: LV dynamics permit tractable learning; but as any hypothesis class, this trades off with model capacity. Here, ε_{LV} quantifies the error due limited expressive power. Further reducing ε_{LV} can be achieved by consider richer dynamic models—a challenge left for future work.
- **Deviation from expectation**: The learned $\hat{p}_u$ rely on predicted equilibrium, but are trained on finite-horizon data. In expectation, ε_{dev} captures how finite sequences deviate from their mean. As a rule of thumb, we expect larger T to reduce this form of noise, but this cannot be guaranteed.
- **Sensitivity** : For $N = 2$, the term $|p_1 - p_0|$ quantifies the added value of exploring beyond the default breaking policy of p_0. Intuitively, when the points are farther away, fitting the equilibrium curve is easier. Thus, for the likely case of $p_0 = 0$, the experimental breaking rate p_1 should be chosen to balance between performance gain and overexposure of experimental subjects to breaks.

5 Experiments

We conclude with an empirical evaluation of our approach on semi-synthetic data. We use the MovieLens 1M dataset to generate recommendations and simulate user behavior, this enabling us to evaluate counterfactual outcomes under different policies. See Appendix B for additional details.

5.1 Experimental setup

Data. The MovieLens 1M dataset (Harper & Konstan, 2015) includes 1,000,209 ratings provided by 6,040 users and for 3,540 items, which we use to obtain features, determine the dynamics, and emulate ϕ. We sample and hold out 30% of all ratings r_{ux} via user-stratified sampling, to which we apply Collaborative Filtering (CF) to get user features u and item features x that approximate $u^\top x \approx r_{ux}$ ($d = 8$, RMSE $= 0.917$, $r \in [1, 5]$). This mimics a process where ϕ is based on items recommended by ψ and rated by users. We then take the remaining data points and randomly assign 1,000 users to the test set, on which we evaluate policies. The remaining users are assigned to the train set, which is then further partitioned into the main $\mathcal{S}$ and the different $\mathcal{S}^{(j)}$ per experimental condition.

Recommendation policy and user behavior. Following Kalimeris et al. (2021); Hron et al. (2022), we set ψ to recommend based on learned softmax scores, $\text{softmax}(u^\top x)$, taken over all non held-out x having ratings for u. For user behavior, we generate discrete interaction sequences S_u in a way that relates to our dynamic model, but is nonetheless quite distinct. Specifically, each user u is associated with *discrete-time* latent-state variables, updated via an LV discretization scheme (see Appendix B.5). These, together with recommended items $x_i = x(t_i)$, determine consecutive $\Delta t_i \in S_u$. Importantly, we let Δt_i depend on u's rating for x_i, namely r_{u,x_i}, which we interpret as u's utility from consuming x_i. We parameterize this dependence via 'immediate', item-specific consumption gain parameters β_{ui} that depend on r_{u,x_i}, and for simplicity set $\alpha_u, \gamma_u, \delta_u$ to be fixed.

Note that since β_{ui} depends on the recommended x_i, it varies in time, and hence there is no single β that underlies the dynamics: even in the limit ($\Delta t \to 0, T \to \infty$), user behavior cannot be described by a continuous LV system, which implies $\varepsilon_{\text{LV}} > 0$ (see Appendix Figure 5). Since the baseline RMSE is high, we set $\beta_{ui} \propto \tilde{r}^2_{u,x_i}$, where $\tilde{r}_{u,x_i} = \kappa r_{u,x_i} + (1 - \kappa)u^\top x$, so that κ interpolates between predicted ratings ($\kappa = 0$) and true ratings ($\kappa = 1$). This allows us to (indirectly) control $\varepsilon_{\text{pred}}$. For all experiments we use $T = 100$, and so expect a roughly fixed $\varepsilon_{\text{dev}} > 0$.

Methods. We compare our approach (`LV`) to several baselines: (i) a `default` policy which myopically optimizes for immediate engagement (and so does not break); (ii) a 'safety switch' policy (`safety@`τ) that breaks once consumption surpasses a threshold τ; (iii) a prediction-based policy (`best-of`) that chooses the best observed $p_u = \arg\max_{p_j} f_j(u)$ (rather than optimizing over $p_u \in [0, 1]$); and (iv) an `oracle` benchmark which directly optimizes the (generally unknown) true underlying dynamics. We measure mean long-term engagement rate (LTE) for each approach, and report averages and standard errors computed over 10 random splits. Performance is measured relative to the `default` baseline as it represents no change in policy (typical absolute values are LTE ≈ 10).

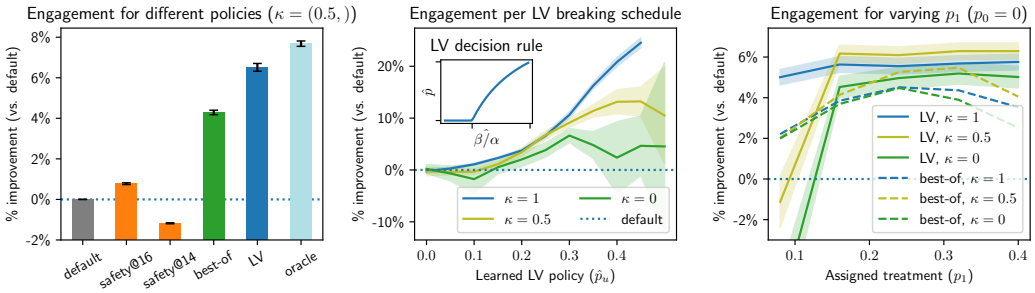

Figure 3: Results on the MovieLens 1M dataset. **(Left)** Performance gain of different approaches (relative to `default` policy). **(Center)** Performance of `LV` by user group, partitioned by learned policies $\hat{p}_u$. **(Right)** Sensitivity to an increasingly aggressive experimental p_1 ($N = 2, p_0 = 0$).

5.2 Results and Analysis

Main results. Figure 3 (left) compares the performance of our method to other policies. Here we set $p_0 = 0$, use $N = 3$ with $p_j \in \{0.05, 0.1, 0.15\}$, and consider $\kappa = 0.5$ (note κ affects all policies). As can be seen, our approach significantly improves over `default` (+6.46%). For `safety@`τ, improvement over the optimal $\tau = 16$ (+5.63%; chosen in hindsight) shows the importance of being preemptive; for the slightly smaller $\tau = 14$, breaks are harmful. The gap from `best-of` (+2.2%) quantifies the gain from the optimization step in Eq. (9), and the close performance to `oracle` (-1.17%) suggests that optimizing the empirical curve (Eq. (8)) works well as as proxy.

User types. Figure 3 (center) shows for our approach how gain in LTE varies across learned breaking policies $\hat{p}_u > 0$. For increasingly-accurate predictions ($\kappa \in \{0, 0.5, 1\}$), the main plot shows performance gains for each group of users, partitioned by their $\hat{p}_u$ values (binned; plot shows average and unit standard deviation per bin). Gains until $\hat{p}_u \leq 0.15$ are mild, but for $\hat{p}_u > 0.15$, the general trend is positive: users who are deemed to require more frequent breaks, benefit more from breaking. Gains until $\hat{p}_u \leq 0.3$ increase for all κ, but for $\hat{p}_u > 0.3$, extrapolation becomes difficult: note the higher variation within each κ, as well as significant differences across κ. This highlights the importance of accurate predictions for inferring optimal $\hat{p}_u$ when the experimental p_i are small. The inlaied plot shows that, in line with our theory, $\hat{p}_u$ exhibits an empirical phase shift in the estimated $\hat{\theta}_u$.

Treatments. Figure 3 (right) shows the effects of experimental treatments on performance. Focusing on $N = 2$, we fix $p_0 = 0$, and consider increasingly aggressive experimentation by varying $p_1 \in (0, 0.4]$. For our approach, increasing p_1 helps, which is anticipated by our theoretical bound. For the `best-of` approach, larger p_1 also helps, but exhibits population-level optimum ($p_1 \approx 0.24$), which is easy to 'overshoot'. Note that when prediction accuracy is low ($\kappa = 0$), experimentation is essential: if p_1 is not sufficiently large, then performance can sharply deteriorate.

6 Discussion

Our paper studies the novel problem of learning optimal breaking policies for recommendation. We posit a tight connection between long-term engagement and user well-being, and argue that optimizing the former requires careful consideration of the latter. Viewing cognitive capacity as a limited but conservable resource, we propose Lotka-Volterra dynamics as a behavioral model that enables effective, responsible, and sustainable optimization of recommendation ecosystems. As any policy task that involves humans, care must be taken regarding potential risks. While our experiments show that optimizing engagement also improves well-being, this need not always be the case; in fact, merely measuring well-being in reality is challenging. Breaks, as interventions, are presumably 'safe', in the sense that at worst they may lead to suboptimal performance (for system) or satisfaction (for users). But breaks can also be used nefariously, e.g., by enabling temporally-varying rewards (Eyal, 2019). As such, breaking policies should be deployed accountably and transparently.

REFERENCES

Roy F Baumeister, Ellen Bratslavsky, Mark Muraven, and Dianne M Tice. Ego depletion: is the active self a limited resource. *Journal of personality and social psychology.*, 74(5), 1998.

Bruce O Bergum and Donald J Lehr. Vigilance performance as a function of interpolated rest. *Journal of Applied Psychology*, 46(6):425, 1962.

Junyu Cao, Wei Sun, Zuo-Jun (Max) Shen, and Markus Ettl. Fatigue-aware bandits for dependent click models. In *AAAI*, pp. 3341–3348, 2020.

Allison JB Chaney, Brandon M Stewart, and Barbara E Engelhardt. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM conference on recommender systems*, pp. 224–232, 2018.

Donghui Chen and Robert J Plemmons. Nonnegativity constraints in numerical analysis. In *The birth of numerical analysis*, pp. 109–139. World Scientific, 2010.

Josh Constine. Instagram says 'you're all caught up' in first time-well-spent feature. techcrunch, 2018. URL https://techcrunch.com/2018/05/21/scroll-responsibly/. [Online; accessed 23-September-2022].

Mihaela Curmei, Andreas A. Haupt, Benjamin Recht, and Dylan Hadfield-Menell. Towards psychologically-grounded dynamic preference models. In *Proceedings of the 16th ACM Conference on Recommender Systems*, RecSys '22, pp. 35–48, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392785. doi: 10.1145/3523227.3546778.

Shai Danziger, Jonathan Levav, and Liora Avnaim-Pesso. Extraneous factors in judicial decisions. *Proceedings of the National Academy of Sciences*, 108(17):6889–6892, 2011. doi: 10.1073/pnas.1018033108.

Sarah Dean and Jamie Morgenstern. Preference dynamics under personalized recommendations. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, EC '22, pp. 795–816, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391504. doi: 10.1145/3490486.3538346.

Xiang Ding, Jing Xu, Guanling Chen, and Chenren Xu. Beyond smartphone overuse: identifying addictive mobile apps. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pp. 2821–2828, 2016.

Rachit Dubey, Thomas L. Griffiths, and Peter Dayan. The pursuit of happiness: A reinforcement learning perspective on habituation and comparisons. *PLoS Comput. Biol.*, 18(8), 2022. doi: 10.1371/journal.pcbi.1010316.

Jacob P Duncan, Teresa Aubele-Futch, and Monica McGrath. A fast-slow dynamical system model of addiction: Predicting relapse frequency. *SIAM Journal on Applied Dynamical Systems*, 18(2): 881–903, 2019.

Nir Eyal. *Hooked : how to build habit-forming products / Nir Eyal with Ryan Hoover.* Portfolio/Penguin, New York, updated edition edition, 2019. ISBN 9781591847786.

J Doyne Farmer. Market force, ecology and evolution. *Industrial and Corporate Change*, 11(5): 895–953, 2002.

Simona Gilboa, Arie Shirom, Yitzhak Fried, and Cary Cooper. A meta-analysis of work demand stressors and job performance: examining main and moderating effects. *Personnel psychology*, 61(2):227–271, 2008.

Nico S Gorbach, Stefan Bauer, and Joachim M Buhmann. Scalable variational inference for dynamical systems. *Advances in neural information processing systems*, 30, 2017.

Sunil Gupta, Dominique Hanssens, Bruce Hardie, Wiliam Kahn, V Kumar, Nathaniel Lin, Nalini Ravishanker, and S Sriram. Modeling customer lifetime value. *Journal of service research*, 9(2): 139–155, 2006.

F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.

William S Helton and Paul N Russell. Rest is still best: The role of the qualitative and quantitative load of interruptions on vigilance. *Human factors*, 59(1):91–100, 2017.

Robert A Hennfng, Steven L Sauter, Gavriel Salvendy, and Edward F Krieg Jr. Microbreak length, performance, and stress in a data entry task. *Ergonomics*, 32(7):855–864, 1989.

Josef Hofbauer, Karl Sigmund, et al. *Evolutionary games and population dynamics*. Cambridge university press, 1998.

Jiri Hron, Karl Krauth, Michael I Jordan, Niki Kilbertus, and Sarah Dean. Modeling content creator incentives on algorithm-curated platforms. *arXiv preprint arXiv:2206.13102*, 2022.

Nicolas Hug. Surprise: A python library for recommender systems. *Journal of Open Source Software*, 5(52):2174, 2020.

Ray Jiang, Silvia Chiappa, Tor Lattimore, András György, and Pushmeet Kohli. Degenerate feedback loops in recommender systems. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 383–390, 2019.

Daniel Kahneman. *Attention and effort*, volume 1063. Citeseer, 1973.

Dimitris Kalimeris, Smriti Bhagat, Shankar Kalyanaraman, and Udi Weinsberg. Preference amplification in recommender systems. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 805–815, 2021.

Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. The challenge of understanding what users want: Inconsistent preferences and engagement optimization. *arXiv preprint arXiv:2202.11776*, 2022.

Robert Kleinberg and Nicole Immorlica. Recharging bandits. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 309–319. IEEE, 2018.

Karl Krauth, Sarah Dean, Alex Zhao, Wenshuo Guo, Mihaela Curmei, Benjamin Recht, and Michael I Jordan. Do offline metrics predict online performance in recommender systems? *arXiv preprint arXiv:2011.07931*, 2020.

Liu Leqi, Fatma Kilinc Karzan, Zachary Lipton, and Alan Montgomery. Rebounding bandits for modeling satiation effects. *Advances in Neural Information Processing Systems*, 34:4003–4014, 2021.

Alfred J. Lotka. Contribution to the theory of periodic reactions. *The Journal of Physical Chemistry*, 14(3):271–274, 1910. doi: 10.1021/j150111a004.

David Madras, Toni Pitassi, and Richard Zemel. Predict responsibly: improving fairness and accuracy by learning to defer. *Advances in Neural Information Processing Systems*, 31, 2018.

Masoud Mansoury, Himan Abdollahpouri, Mykola Pechenizkiy, Bamshad Mobasher, and Robin Burke. Feedback loop and bias amplification in recommender systems. In *Proceedings of the 29th ACM international conference on information & knowledge management*, pp. 2145–2148, 2020.

Alberto Monge Roffarello and Luigi De Russis. The race towards digital wellbeing: Issues and opportunities. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, pp. 1–14, 2019.

Christian Montag, Zhiying Zhao, Cornelia Sindermann, Lei Xu, Meina Fu, Jialin Li, Xiaoxiao Zheng, Keshuang Li, Keith M Kendrick, Jing Dai, et al. Internet communication disorder and the structure of the human brain: Initial insights on wechat addiction. *Scientific reports*, 8(1): 1–10, 2018.

O'Dhaniel A Mullette-Gillman, Ruth LF Leong, and Yoanna A Kurnianingsih. Cognitive fatigue destabilizes economic decision making preferences and strategies. *PloS one*, 10(7):e0132022, 2015.

Mark Muraven and Roy F Baumeister. Self-regulation and depletion of limited resources: Does self-control resemble a muscle? *Psychological bulletin*, 126(2):247, 2000.

Sarah Perez. Apple unveils a new set of 'digital wellness' features for better managing screen time. *TechCrunch, June*, 4, 2018. URL https://techcrunch.com/2018/06/04/apple-unveils-a-new-set-of-digital-wellness-features-for-better-managing-screen-time/. [Online; accessed 23-September-2022].

Jennifer Rauch. *Slow media: Why slow is satisfying, sustainable, and smart*. Oxford University Press, 2018.

Stephen E Robertson. The probability ranking principle in ir. *Journal of documentation*, 1977.

Hayden A Ross, Paul N Russell, and William S Helton. Effects of breaks and goal switches on the vigilance decrement. *Experimental brain research*, 232(6):1729–1737, 2014.

Tom Ryder, Andrew Golightly, A Stephen McGough, and Dennis Prangle. Black-box variational inference for stochastic differential equations. In *International Conference on Machine Learning*, pp. 4423–4432. PMLR, 2018.

Larry Samuelson. *Evolutionary games and equilibrium selection*, volume 1. MIT press, 1998.

Francesco Sanna Passino, Lucas Maystre, Dmitrii Moor, Ashton Anderson, and Mounia Lalmas. Where to next? a dynamic model of user preferences. In *Proceedings of the Web Conference 2021*, WWW '21, pp. 3210–3220, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383127. doi: 10.1145/3442381.3450028.

Sven Schmit and Carlos Riquelme. Human interaction with recommendation systems. In *International Conference on Artificial Intelligence and Statistics*, pp. 862–870. PMLR, 2018.

Maarten P. Scholl, Anisoara Calinescu, and J. Doyne Farmer. How market ecology explains market malfunction. *Proceedings of the National Academy of Sciences*, 118(26):e2015574118, 2021. doi: 10.1073/pnas.2015574118.

Hans Henrik Sievertsen, Francesca Gino, and Marco Piovesan. Cognitive fatigue influences students' performance on standardized tests. *Proceedings of the National Academy of Sciences*, 113 (10):2621–2624, 2016.

Yasuhiro Takeuchi. *Global dynamical properties of Lotka-Volterra systems*. World Scientific, 1996.

Dustin B Thoman, Jessi L Smith, and Paul J Silvia. The resource replenishment function of interest. *Social Psychological and Personality Science*, 2(6):592–599, 2011.

Emmanuel Trélat and Enrique Zuazua. The turnpike property in finite-dimensional nonlinear optimal control. *Journal of Differential Equations*, 258(1):81–114, 2015.

Jyun-Cheng Wang and Juo-Ping Lin. Are personalization systems really personal?-effects of conformity in reducing information. In *36th Annual Hawaii International Conference on System Sciences, 2003. Proceedings of the*, pp. 10–pp. IEEE, 2003.

Romain Warlop, Alessandro Lazaric, and Jérémie Mary. Fighting boredom in recommender systems with linear reinforcement learning. *Advances in Neural Information Processing Systems*, 31, 2018.

Jörgen W Weibull. *Evolutionary game theory*. MIT press, 1997.