

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 preemptive, 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 Lotka-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 prescience 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. 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; Henning 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 feedback 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 are 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 independently 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 mean ‘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 $\{(t_i, x_i) \mid t_i \in [0, t]\} \sim \text{TPP}_\psi(u; t)$.

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 $S_u = \{(t_i, x_i) \mid t_i \in [T_0, T_0 + T]\}$. 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 $\mathbb{E}_{u \sim D} \mathbb{E}[\frac{1}{T}|S_u|]$ for $(S_u^0, S_u) \sim \text{TPP}_\psi(u; T_0 + T)$, 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:

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

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.

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

Data and exploration. For learning an optimal breaking policy, we assume the system has access to a dataset of previously logged interactions $\mathcal{S} = \{(S_1^0, S_1), \dots, (S_m^0, S_m)\}$ for m train-time users $u_1, \dots, u_m \sim D$, 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:

$$\frac{d\lambda}{dt} = -\alpha\lambda + \beta\lambda q, \quad \frac{dq}{dt} = \gamma q(1 - q) - \delta\lambda q \quad (2)$$

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.

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 pa-

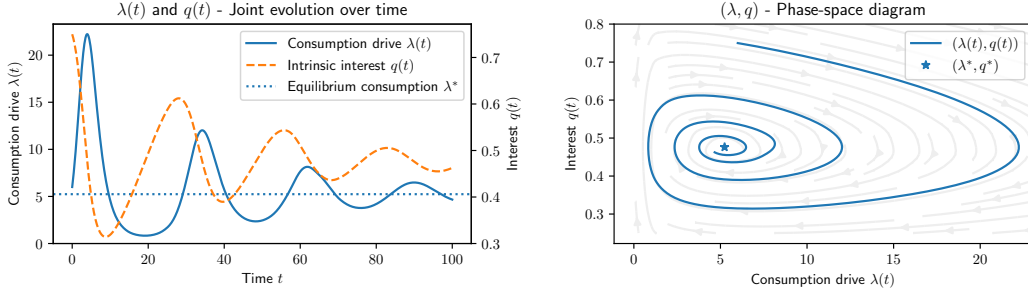


Figure 1: Characteristic LV dynamics (Eq. (2)). **(Left)** Temporal relations between *consumption drive* $\lambda(t)$, *intrinsic interest* $q(t)$, and equilibrium λ^* . Note how λ drops only some time after q has depleted. **(Right)** Evolution of the same system in phase space (λ, q) , with its equilibrium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

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 $\Delta t \rightarrow 0$ and $T \rightarrow \infty$). 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:

$$\frac{d\lambda}{dt} = -\alpha\lambda + \beta\lambda q(1-p), \quad \frac{dq}{dt} = \gamma q(1-q) - \delta\lambda q(1-p) \quad (3)$$

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:

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

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:

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

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 $\mathbb{E}_{u \sim D} \mathbb{E}_{\text{TPP}} \left[\frac{1}{T}|S_u| \right]$.

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 possible 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:

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

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 $p \in [0, 1]$:

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

for some function norm $\|\cdot\|$, and for which λ_{θ}^* 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 $j = 1, \dots, N$. 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 $(u_k, y_k) \in \mathcal{S}^{(j)}$, then $f_j(u)$ should be a reasonable estimator of the expected $\lambda_u(p_j)$. 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:

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

given here with the ℓ_2 vector norm, and where $\mathbf{f}(u) = (f_1(u), f_2(u), \dots, f_N(u)) \in \mathbb{R}_+^N$. 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 counterfactual 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 piece-wise 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:

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

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

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

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, $\pi(\hat{p}_u)$ idempotently 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 governed 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:

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

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, 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 3,706 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 S and the different $S^{(j)}$ 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, $\text{softmax}_x(\hat{r}_u)$, 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 ($\Delta t \rightarrow 0, T \rightarrow \infty$), 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 $\beta_{ui} \propto \tilde{r}_{u,x_i}^2$, 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 $\text{LTE} \approx 10$).




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 (+6.48%). For `safety@ τ` , improvement over the optimal $\tau = 16$ (+5.68%; chosen in hindsight) shows the importance of being preemptive; for the slightly smaller $\tau = 14$, breaks are harmful. The gap from `best-of` (+2.19%) quantifies the gain from the optimization step in Eq. (9), and the close performance to `oracle` (-0.918%) 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 in-laid 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 Henning, 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’Dhanial 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.
- Mengting Wan and Julian McAuley. Item recommendation on monotonic behavior chains. In *Proceedings of the 12th ACM conference on recommender systems*, pp. 86–94, 2018.
- Mengting Wan, Rishabh Misra, Ndapa Nakashole, and Julian McAuley. Fine-grained spoiler detection from large-scale review corpora. *arXiv preprint arXiv:1905.13416*, 2019.
- 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.
- Jürgen W Weibull. *Evolutionary game theory*. MIT press, 1997.

A DEFERRED PROOFS

In this section, we formalize the model presented in [section 3](#), and prove the claims presented in [section 4](#). The section ends with a formal proof of [Theorem 1](#).

A.1 PROPERTIES OF LOTKA-VOLTERRA SYSTEMS

Definition A.1 (Static policy equilibrium). *Let $\lambda(t), q(t)$ denote a Lotka-Volterra model characterized by parameters $\theta = (\alpha, \beta, \gamma, \delta) \in \mathbb{R}_+^4$, as defined in [Equation 2](#). Let $p \in [0, 1]$, and denote by π_p the static policy corresponding to p . For $\lambda(0), q(0) > 0$, the static equilibrium of the system is defined as:*

$$\begin{aligned}\lambda^*(p; \theta) &= \lim_{t \rightarrow \infty} \lambda(t) \\ q^*(p; \theta) &= \lim_{t \rightarrow \infty} q(t)\end{aligned}$$

We denote $\lambda^*(p) = \lambda^*(p; \theta)$ when θ is clear from the context. We denote $\lambda^*(p; u) = \lambda^*(p; \theta_u)$ when a user $u \in \mathcal{U}$ characterized by parameters θ_u is given and clear from the context.

Proposition A.1 (Global stability). *$\lambda^*(p; \theta)$ exists and uniquely defined for all $\theta \in \mathbb{R}_+^4$, $p \in [0, 1]$ and for all initial conditions $\lambda(0), q(0) > 0$.*

Proof. See ([Takeuchi, 1996](#), Section 3.2). □

Lemma A.1 (Equilibrium of LV behavioral model. Formal proof of [Lemma 1](#)). *Assume a Lotka-Volterra model characterized by $\theta = (\alpha, \beta, \gamma, \delta) \in \mathbb{R}_+^4$, and let $p \in [0, 1]$ denote the proportion of interactions in which a forced break is served. The static equilibrium of the model under static policy π_p is given by:*

$$\begin{aligned}\lambda^*(p) &= \begin{cases} \frac{\gamma}{\delta} \frac{1}{1-p} \left(1 - \frac{\alpha}{\beta} \frac{1}{1-p}\right) & p \in \left[0, 1 - \frac{\alpha}{\beta}\right] \\ 0 & \text{otherwise} \end{cases} \\ q^*(p) &= \begin{cases} \frac{\alpha}{\beta} \frac{1}{1-p} & p \in \left[0, 1 - \frac{\alpha}{\beta}\right] \\ 1 & \text{otherwise} \end{cases}\end{aligned}$$

Proof. The LV dynamical system is given by [Equation 3](#):

$$\begin{aligned}\frac{d\lambda}{dt} &= -\alpha\lambda + \beta(1-p)\lambda q \\ \frac{dq}{dt} &= \gamma q(1-q) - \delta(1-p)\lambda q\end{aligned}$$

when $p \in \left[0, 1 - \frac{\alpha}{\beta}\right]$ we equate $\frac{d\lambda}{dt} = 0$, $\frac{dq}{dt} = 0$ and obtain the result. The solution is guaranteed to be valid, as both $\lambda^*(p) > 0$ and $q^*(p) \in [0, 1]$.

Conversely, when $p \notin \left[0, 1 - \frac{\alpha}{\beta}\right]$, there exists $\epsilon > 0$ such that $\frac{d}{dt} \log \lambda < -\epsilon < 0$ for all $\lambda > 0$, $q \in [0, 1]$. From this we obtain that $\log \lambda(t)$ tends towards $-\infty$, and therefore $\lambda(t)$ tends towards 0, and $\lambda^*(p) = 0$ as required. When $\lambda(t)$ is close to zero, the interaction terms vanish in the $\frac{dq}{dt}$ equation, and $q(t)$ grows logistically towards 1. □

Proposition A.2 (Equilibrium bounds). *For a Lotka-Volterra model, the static equilibrium $\lambda^*(p)$ is bounded by:*

$$0 \leq \lambda^*(p) \leq \frac{\beta\gamma}{4\alpha\delta}$$

Proof. Denote $x = \frac{1}{1-p}$. From [Lemma A.1](#), for $x \in \left[1, \frac{\beta}{\alpha}\right]$ the equilibrium consumption $\lambda^*(x)$ is given by:

$$\lambda^*(x) = \frac{\gamma}{\delta} x \left(1 - \frac{\alpha}{\beta} x\right)$$

and is zero otherwise. The equilibrium is a quadratic function of x with roots $x \in \left\{0, \frac{\beta}{\alpha}\right\}$, and therefore attains its maximum at $x = \frac{\beta}{2\alpha}$. Plugging back the maximizing x into λ^* we obtain the upper bound. Lower bound is attained as the equilibrium in [Lemma A.1](#) is clipped by 0 from below. \square

Lemma A.2 (Optimal static policy. Formal proof of [Lemma 2](#)). *The optimal static policy for a Lotka-Volterra system is given by:*

$$p_{\text{opt}} = \begin{cases} 1 - 2\frac{\alpha}{\beta} & \frac{\alpha}{\beta} \leq \frac{1}{2} \\ 0 & \frac{\alpha}{\beta} > \frac{1}{2} \end{cases}$$

And the optimal equilibrium engagement rate is given by:

$$\lambda_{\text{opt}}^* = \begin{cases} \frac{\beta\gamma}{4\alpha\delta} & \frac{\alpha}{\beta} \leq \frac{1}{2} \\ \frac{\gamma}{\delta} \left(1 - \frac{\alpha}{\beta}\right) & \frac{\alpha}{\beta} > \frac{1}{2} \end{cases}$$

Proof. Denote $x = \frac{1}{1-p}$. From [Proposition A.2](#), the global maximum of $\lambda^*(x)$ is attained at $x = \frac{\beta}{2\alpha}$. Consider two cases: When $\frac{\alpha}{\beta} \leq \frac{1}{2}$, we obtain that $x_{\text{opt}} = \frac{\beta}{2\alpha} \geq 1$, and therefore $p_{\text{opt}} = 1 - \frac{1}{x} \in [0, 1]$. From this we obtain that in this case the global maximum is attained on the simplex, and given by the formula from [Proposition A.2](#). Conversely, when $\frac{\alpha}{\beta} > \frac{1}{2}$, we obtain $p = 1 - \frac{1}{x} < 0$, and therefore x_{opt} translates to a negative value of p . As $\lambda^*(p)$ is uni-modal, the optimal policy restricted to the simplex $[0, 1]$ in this case is attained on the closest boundary point $p = 0$.

[Figure 2](#) provides graphical intuition for this proof (left and center subplots). \square

Proposition A.3 (Inference of α/β from two-treatment equilibrium data. Formal proof of [Proposition 1](#)). *Let $\lambda(t), q(t)$ be a Lotka-Volterra model, let $p_1, p_2 \in [0, 1]$. Denote by $\lambda^*(p_1), \lambda^*(p_2)$ the static equilibrium rates corresponding to static policies π_{p_1}, π_{p_2} , and assume $\lambda^*(p_1), \lambda^*(p_2) > 0$. The parameter ratio $\frac{\alpha}{\beta}$ is given by the following formula:*

$$\frac{\alpha}{\beta} = \frac{(1-p_2)\lambda^*(p_2) - (1-p_1)\lambda^*(p_1)}{\frac{1}{1-p_1} - \frac{1}{1-p_2}}$$

Proof. From [Lemma A.1](#), the equilibrium consumption $\lambda^*(p)$ is given by:

$$\begin{aligned} \lambda^*(p) &= \frac{\gamma}{\delta} \frac{1}{1-p} \left(1 - \frac{\alpha}{\beta} \frac{1}{1-p}\right) \\ &= \frac{\gamma}{\delta} \frac{1}{1-p} - \frac{\alpha}{\beta} \frac{\gamma}{\delta} \left(\frac{1}{1-p}\right)^2 \end{aligned}$$

When $\lambda^*(p_i)$ is observed for different policies $p_1, \dots, p_m \in \left[0, 1 - \frac{\alpha}{\beta}\right]$, we obtain a polynomial regression problem for the parameters $\frac{\alpha}{\beta}$ and $\frac{\gamma}{\delta}$, which can be solved e.g using Non-Negative Least Squares.

When $m = 2$, we obtain a system of two linear equations. Apply Cramer's rule to obtain:

$$\frac{\gamma}{\delta} = \frac{\frac{\lambda^*(p_2)}{(1-p_1)^2} - \frac{\lambda^*(p_1)}{(1-p_2)^2}}{\frac{1}{(1-p_1)(1-p_2)^2} - \frac{1}{(1-p_1)^2(1-p_2)}} = \frac{(1-p_2)^2\lambda^*(p_2) - (1-p_1)^2\lambda^*(p_1)}{p_2 - p_1} \quad (11)$$

$$\frac{\alpha}{\beta} \frac{\gamma}{\delta} = \frac{\frac{\lambda^*(p_2)}{(1-p_1)} - \frac{\lambda^*(p_1)}{(1-p_2)}}{\frac{1}{(1-p_1)(1-p_2)^2} - \frac{1}{(1-p_1)^2(1-p_2)}} = (1-p_1)(1-p_2) \frac{(1-p_2)\lambda^*(p_2) - (1-p_1)\lambda^*(p_1)}{p_2 - p_1} \quad (12)$$

And therefore $\frac{\alpha}{\beta}$ is given by:

$$\frac{\alpha}{\beta} = \frac{\frac{\lambda^*(p_2)}{(1-p_1)} - \frac{\lambda^*(p_1)}{(1-p_2)}}{\frac{\lambda^*(p_2)}{(1-p_1)^2} - \frac{\lambda^*(p_1)}{(1-p_2)^2}} = (1-p_1)(1-p_2) \frac{(1-p_2)\lambda^*(p_2) - (1-p_1)\lambda^*(p_1)}{(1-p_2)^2\lambda^*(p_2) - (1-p_1)^2\lambda^*(p_1)}$$

□

A.2 MODEL FITTING FROM ENGAGEMENT PREDICTIONS

Notations. In this section only, we use the common notation $q = 1 - p$ to denote complementary probabilities.

Definition A.2 (Empirical value of α/β). *For single-channel experiments with forced-break probabilities p_1, p_2 , denote $\lambda_i = \lambda^*(p_i)$, $f_i = f_{p_i}(u)$, $q_i = 1 - p_i$. The empirical value of the $\frac{\alpha}{\beta}$ parameter is given by the following formula:*

$$\frac{\hat{\alpha}}{\hat{\beta}} = \frac{q_1 q_2 (q_1 f_1 - q_2 f_2)}{q_1^2 f_1 - q_2^2 f_2}$$

Proposition A.4 (α/β estimation error from prediction errors). *Given a single-channel Lotka-Volterra system with parameter $\frac{\alpha}{\beta} \geq 1$. Let $p_1, p_2 \in [1, \frac{\alpha}{\beta}]$, denote $\lambda_i^* = \lambda^*(p_i) \in \mathbb{R}_+$, and let $f_i = \lambda_i^* + \varepsilon_i$ be the predicted engagement rates corresponding to p_1, p_2 . When $|\varepsilon_1|, |\varepsilon_2| \leq \varepsilon \leq \frac{\gamma}{\delta} \frac{|p_1 - p_2|}{4}$, the estimation error is bounded by:*

$$\left| \frac{\alpha}{\beta} - \frac{\hat{\alpha}}{\hat{\beta}} \right| \leq \frac{\varepsilon}{|p_1 - p_2|} \frac{\beta \delta}{\alpha \gamma}$$

Proof. denote $q_i = 1 - p_i$. The value of $\frac{\alpha}{\beta}$ is given by [Proposition A.3](#):

$$\frac{\alpha}{\beta} = \frac{q_1 q_2 (q_1 \lambda_1^* - q_2 \lambda_2^*)}{q_1^2 \lambda_1^* - q_2^2 \lambda_2^*}$$

And the estimator for $\frac{\alpha}{\beta}$ is obtained by replacing the true value with their predictions:

$$\begin{aligned} \frac{\hat{\alpha}}{\hat{\beta}} &= \frac{q_1 q_2 (q_1 f_1 - q_2 f_2)}{q_1^2 f_1 - q_2^2 f_2} \\ &= \frac{q_1 q_2 (q_1 (\lambda_1^* + \varepsilon_1) - q_2 (\lambda_2^* + \varepsilon_2))}{q_1^2 (\lambda_1^* + \varepsilon_1) - q_2^2 (\lambda_2^* + \varepsilon_2)} \end{aligned}$$

The estimation error is given by:

$$\begin{aligned} \left| \frac{\alpha}{\beta} - \frac{\hat{\alpha}}{\hat{\beta}} \right| &= \left| \frac{q_1^2 q_2^2 (q_1 - q_2) (\varepsilon_2 \lambda_1^* - \varepsilon_1 \lambda_2^*)}{(q_1^2 \lambda_1^* - q_2^2 \lambda_2^*) (q_1^2 \lambda_1^* - q_2^2 \lambda_2^* - (q_1^2 \varepsilon_1 - q_2^2 \varepsilon_2))} \right| \\ &= \underbrace{(q_1 q_2)^2}_{\equiv(i)} \underbrace{\left| \frac{q_1 - q_2}{q_1^2 \lambda_1^* - q_2^2 \lambda_2^*} \right|}_{\equiv(ii)} \underbrace{|\varepsilon_2 \lambda_1^* - \varepsilon_1 \lambda_2^*|}_{\equiv(iii)} \underbrace{\left| \frac{1}{q_1^2 \lambda_1^* - q_2^2 \lambda_2^* - (q_1^2 \varepsilon_1 - q_2^2 \varepsilon_2)} \right|}_{\equiv(iv)} \end{aligned}$$

We now proceed to bound each factor:

- For (i), the term $(q_1 q_2)^2$ is bounded by 1 since $q_1, q_2 \in [0, 1]$.
- For (ii), the term $\left| \frac{q_1 - q_2}{q_1^2 \lambda_1^* - q_2^2 \lambda_2^*} \right|$ is equal to $(\frac{\gamma}{\delta})^{-1}$ by Eq. (11).
- For (iii), from [Proposition A.2](#) we obtain the bound $0 \leq \lambda_i^* \leq \frac{\beta \gamma}{4 \alpha \delta}$, and therefore the term $|\varepsilon_2 \lambda_1^* - \varepsilon_1 \lambda_2^*|$ is bounded by $2 \left(\frac{\beta \gamma}{4 \alpha \delta} \right) \varepsilon = \frac{\beta \gamma}{2 \alpha \delta} \varepsilon$.

- For (iv), the term $\left| \frac{1}{q_1^2 \lambda_1^* - q_2^2 \lambda_2^* - (q_1^2 \varepsilon_1 - q_2^2 \varepsilon_2)} \right|$ is equal to:

$$\begin{aligned}
\text{(iv)} &\equiv \left| \frac{1}{q_1^2 \lambda_1^* - q_2^2 \lambda_2^* - (q_1^2 \varepsilon_1 - q_2^2 \varepsilon_2)} \right| \\
&= \frac{1}{|p_1 - p_2|} \left| \frac{q_1^2 \lambda_1^* - q_2^2 \lambda_2^* - (q_1^2 \varepsilon_1 - q_2^2 \varepsilon_2)}{p_1 - p_2} \right|^{-1} \\
&= \frac{1}{|p_1 - p_2|} \left| \underbrace{\frac{q_1^2 \lambda_1^* - q_2^2 \lambda_2^*}{p_1 - p_2}}_{\text{Eq. (11)}} - \frac{q_1^2 \varepsilon_1 - q_2^2 \varepsilon_2}{p_1 - p_2} \right|^{-1} \\
&= \frac{1}{|p_1 - p_2|} \left| \frac{\gamma}{\delta} - \frac{q_1^2 \varepsilon_1 - q_2^2 \varepsilon_2}{p_1 - p_2} \right|^{-1}
\end{aligned}$$

Note that $\left| \frac{q_1^2 \varepsilon_1 - q_2^2 \varepsilon_2}{p_1 - p_2} \right| \leq \frac{2\varepsilon}{|p_1 - p_2|}$. When ε is small enough, and specifically when the bound $\varepsilon \leq \frac{\gamma}{8} \frac{|p_1 - p_2|}{4}$ holds, we obtain:

$$\left| \frac{\gamma}{\delta} - \frac{q_1^2 \varepsilon_1 - q_2^2 \varepsilon_2}{p_1 - p_2} \right|^{-1} \leq \frac{\delta}{\gamma} \left| 1 - \frac{1}{2} \right|^{-1} \leq 2 \frac{\delta}{\gamma}$$

and therefore:

$$\text{(iv)} \leq \frac{2}{|p_1 - p_2|} \frac{\delta}{\gamma}$$

Aggregating results (i)-(iv) above, we obtain the overall bound:

$$\begin{aligned}
\left| \frac{\alpha}{\beta} - \frac{\hat{\alpha}}{\beta} \right| &= \underbrace{(q_1 q_2)^2}_{\leq 1} \underbrace{\left| \frac{q_1 - q_2}{q_1^2 \lambda_1^* - q_2^2 \lambda_2^*} \right|}_{=\frac{\delta}{\gamma}} \underbrace{|\varepsilon_2 \lambda_1^* - \varepsilon_1 \lambda_2^*|}_{\leq \frac{\beta \gamma}{2\alpha \delta} \varepsilon} \underbrace{\left| \frac{1}{q_1^2 \lambda_1^* - q_2^2 \lambda_2^* - (q_1^2 \varepsilon_1 - q_2^2 \varepsilon_2)} \right|}_{\leq \frac{2}{|p_1 - p_2|} \frac{\delta}{\gamma}} \\
&\leq \frac{\varepsilon}{|p_1 - p_2|} \frac{\beta \delta}{\alpha \gamma}
\end{aligned}$$

□

Proposition A.5 (Cost of α/β estimation error). *Let $\frac{\alpha}{\beta}$ be the engagement ratio parameter of a one-channel Lotka-Volterra system, and let $\left(\frac{\hat{\alpha}}{\beta}\right)$ be an estimate of these parameters. Let λ_{opt}^* be the engagement rate of the optimal static policy, and denote $\lambda^*(x) = \lambda^*(\hat{p}(x))$. When $\left| \frac{\alpha}{\beta} - \left(\frac{\hat{\alpha}}{\beta}\right) \right| \leq \min \left\{ \frac{\alpha}{2\beta}, 1 \right\}$ The price of estimation error is bounded by:*

$$\lambda_{\text{opt}}^* - \lambda^* \left(\left(\frac{\hat{\alpha}}{\beta}\right) \right) \leq \left(\frac{\gamma}{\delta}\right) \min \left\{ \left(2\frac{\alpha}{\beta}\right)^{-2} \left| \frac{\alpha}{\beta} - \left(\frac{\hat{\alpha}}{\beta}\right) \right|, \left(4\frac{\alpha}{\beta}\right)^{-1} \right\}$$

Proof. Denote $r = \frac{\alpha}{\beta}$, $x = \left(\frac{\hat{\alpha}}{\beta}\right)$, and assume without loss of generality that $\frac{\gamma}{\delta} = 1$ and $r \leq 1$. The optimal equilibrium engagement rate is given by:

$$\lambda_{\text{opt}}^* = \begin{cases} \frac{1}{4r} & r \in (0, \frac{1}{2}] \\ 1 - r & r \in (\frac{1}{2}, 1] \end{cases}$$

The chosen policy $\hat{p}(x)$ is given by:

$$\hat{p}(x) = \begin{cases} 1 - 2x & x \in [0, \frac{1}{2}] \\ 0 & \text{otherwise} \end{cases}$$




Figure 4: Graphical illustration of [Proposition A.5](#). Cost of estimation error for different values of $\frac{\alpha}{\beta}$, and their corresponding upper bounds given by the claim.

Assume without loss of generality that $x \in [0, \frac{1}{2}]$, as values of x outside the interval can be clipped to its edges without affecting the result. The equilibrium engagement rate of the selected policy is given by:

$$\lambda^*(x) = \lambda^*(\hat{p}(x)) = \begin{cases} 0 & x \in [0, \frac{r}{2}] \\ \frac{1}{2x} (1 - \frac{r}{2x}) & x \in (\frac{r}{2}, \frac{1}{2}] \end{cases}$$

Denote $\Delta(x) = \lambda_{\text{opt}}^* - \lambda^*(x)$. We obtain:

$$\Delta(x) = \lambda_{\text{opt}}^* - \lambda^*(x) = \begin{cases} \frac{1}{4r} & r \in (0, \frac{1}{2}], x \in [0, \frac{r}{2}] \\ \frac{(x-r)^2}{4x^2r} & r \in (0, \frac{1}{2}], x \in (\frac{r}{2}, \frac{1}{2}] \\ (1-r) & r \in (\frac{1}{2}, 1], x \in [0, \frac{r}{2}] \\ (1-r) - \frac{1}{2x} (1 - \frac{r}{2x}) & r \in (\frac{1}{2}, 1], x \in (\frac{r}{2}, \frac{1}{2}] \end{cases}$$

Observe that $\frac{1}{4r} \geq 1-r$ for all $r \in (0, 1]$, and therefore we obtain for all x, r :

$$\Delta(x) \leq \frac{1}{4r} \quad (13)$$

From the convexity of $\Delta(x)$ in the region around $x = r$ we obtain:

$$\Delta(x) \leq \frac{1}{2r^2} |x - r| \quad (14)$$

Finally, combining the two bounds yields the final result. A geometric interpretation of this claim is illustrated in [Figure 4](#). \square

A.3 OPTIMAL STATIONARY POLICY FROM ENGAGEMENT PREDICTIONS

Definition A.3 (Expected observable rate). *Let $u \in \mathcal{U}$, $p \in [0, 1]$, and $T > 0$. Let $p \in [0, 1]$, denote the corresponding static policy by π_p . The expected observable rate $\bar{\lambda}_u(p; T)$ is defined as:*

$$\bar{\lambda}_u(p; u) = \mathbb{E}_\pi \left[\frac{1}{T} |\text{TPP}_{\pi_p}(u; T)| \right]$$

where expectation is taken over the stochastic decisions of π_p .

Definition A.4 (Lokta-Volterra approximation of TPP). *Let $u \in \mathcal{U}$, and $T > 0$. Denote by p^* the maximizer of expected observable rate:*

$$p^* = \arg \max_{p \in [0, 1]} \bar{\lambda}_u(p; u)$$

The LV approximation of $\text{TPP}(u; T)$ is defined as:

$$\theta_u^* = \arg \min_{\theta} \max_{p \in [0, 1]} |\bar{\lambda}_u(p; u) - \lambda^*(p; \theta)|$$

such that $\arg \max_p \lambda^*(p; \theta) = p^*$. The corresponding approximation error is defined as:

$$\varepsilon_{\text{LV}, u} = \max_{p \in [0, 1]} |\bar{\lambda}_u(p; u) - \lambda^*(p; \theta_u^*)|$$

Notations. When u is clear from the context, we denote $\theta^* = \theta_u^*$, $\varepsilon_{LV} = \varepsilon_{LV,u}$. We use α^*, β^*, \dots to refer to the corresponding parts of the Lokta-Volterra parameters vector θ^* .

We are now ready to state and prove the main theorem for this section:

Theorem A.1 (Regret bound for learned static policy. Formal version of [Theorem 1](#)). *Let $p_1, p_2 \in [0, 1]$ denote two static forced-break policies, and denote by \mathcal{U} the set of users, and assume they remain engaged under the stationary policies $\pi(p_1)$ and $\pi(p_2)$. Assume $S_u(p; T) \sim \text{TPP}_{\pi_p \circ \psi}(u; T)$, and let $\mu = \left(\max_{u \in \mathcal{U}} \frac{\tilde{\gamma}_u}{\delta_u}\right) \cdot \left(\max_{u' \in \mathcal{U}} \frac{\tilde{\delta}_{u'}}{\tilde{\gamma}_{u'}}\right)$, $\nu = \max_{u \in \mathcal{U}} \left(\frac{\tilde{\beta}_u}{\tilde{\alpha}_u}\right)$.*

Let $f_{p_1}, f_{p_2} : \mathcal{U} \rightarrow \mathbb{R}_+$ be functions predicting $\frac{1}{T}|S_u(p_1; T)|, \frac{1}{T}|S_u(p_2; T)|$, respectively. Denote the learned policy by \hat{p} , and the optimal policy by p^ .*

If (i) the expected RMSE of f_{p_1}, f_{p_2} is bounded by $\varepsilon_{\text{pred}}$, (ii) the average absolute deviation of $\frac{1}{T}|\text{TPP}(u; T)|$ is bounded by ε_{dev} , and (iii) the expected LV approximation error of the system is bounded by ε_{LV} , then the learned policy \hat{p} has bounded regret:

$$\mathbb{E}_{u, \pi} \left[\left| \frac{1}{T}|S_u(p^*; T)| - \frac{1}{T}|S_u(\hat{p}; T)| \right| \right] \leq \frac{\eta_{\text{TPP}}}{|p_1 - p_2|} (\varepsilon_{\text{pred}} + \varepsilon_{\text{dev}} + \varepsilon_{LV})$$

where expectation is taken over stochastic choices of policies, and $\eta_{\text{TPP}} = g(\mu, \nu) \in \text{poly}(\mu, \nu)$.

Proof. By assumption (i), the functions f_{p_1}, f_{p_2} have bounded expected RMSE:

$$\mathbb{E}_u \left[\left(f_{p_i}(u) - \frac{1}{T}|S_u(p_i; T)| \right)^2 \right] \leq \varepsilon_{\text{pred}}^2 \quad (15)$$

Applying Jensen's inequality with the convex function $\varphi(x) = x^2$ yields:

$$\left(\mathbb{E}_u \left[\left| f_{p_i}(u) - \frac{1}{T}|S_u(p_i; T)| \right| \right] \right)^2 \leq \mathbb{E}_u \left[\left(f_{p_i}(u) - \frac{1}{T}|S_u(p_i; T)| \right)^2 \right]$$

Combining with Eq. (15) and taking the square root, we obtain an upper bound on the expected absolute error:

$$\mathbb{E}_u \left[\left| f_{p_i}(u) - \frac{1}{T}|S_u(p_i; T)| \right| \right] \leq \varepsilon_{\text{pred}} \quad (16)$$

Let $\Delta_f = |f_{p_i}(u) - \lambda^*(p_i)|$ apply the triangle inequality to obtain:

$$\begin{aligned} \Delta_f &= |f_{p_i}(u) - \lambda^*(p_i)| \\ &\leq \left| f_{p_i}(u) - \frac{1}{T}|S_u(u; T)| \right| + \left| \frac{1}{T}|S_u(u; T)| - \bar{\lambda}(p_i; u) \right| + \left| \bar{\lambda}(p_i; u) - \lambda^*(p_i) \right| \end{aligned}$$

Denote $\varepsilon_f = \varepsilon_{\text{pred}} + \varepsilon_{\text{dev}} + \varepsilon_{LV}$. Applying the triangle inequality and using the bounds in Eq. (16) together with assumptions (ii), (iii), we obtain:

$$\begin{aligned} \mathbb{E}_{u, \pi} [\Delta_f] &\leq \mathbb{E}_u \left[\left| f_{p_i}(u) - \frac{1}{T}|S_u(u; T)| \right| \right] \\ &\quad + \mathbb{E}_{u, \pi} \left[\left| \frac{1}{T}|S_u(u; T)| - \bar{\lambda}(p_i; u) \right| \right] \\ &\quad + \mathbb{E}_u \left[\left| \bar{\lambda}(p_i; u) - \lambda^*(p_i; \theta_u^*) \right| \right] \\ &\leq \varepsilon_{\text{pred}} + \varepsilon_{\text{dev}} + \varepsilon_{LV} = \varepsilon_f \end{aligned} \quad (17)$$

Denote $\theta_u^* = (\alpha, \beta, \gamma, \delta)$. The empirical value $\left(\frac{\hat{\alpha}}{\hat{\beta}}\right)$ of $\left(\frac{\alpha}{\beta}\right)$ is given by [Definition A.2](#). Denote the estimation error by $\Delta_{\frac{\alpha}{\beta}} = \left| \left(\frac{\hat{\alpha}}{\hat{\beta}}\right) - \left(\frac{\alpha}{\beta}\right) \right|$.

By [Proposition A.4](#), the following pointwise upper bound on $\Delta_{\frac{\alpha}{\beta}}$ applies when $\Delta_f \leq \frac{\gamma}{\delta} \frac{|p_1 - p_2|}{4}$:

$$\Delta_{\frac{\alpha}{\beta}} \leq \frac{\Delta_f}{|p_1 - p_2|} \frac{\beta\delta}{\alpha\gamma} \quad (18)$$

Plugging in the bound on the expected value of Δ_f into Eq. (18), we obtain in expectation:

$$\begin{aligned} \mathbb{E}_{u, \pi} \left[\Delta_{\frac{\alpha}{\beta}} \mid \Delta_f \leq \frac{\gamma}{\delta} \frac{|p_1 - p_2|}{4} \right] &\leq \mathbb{E}_{u, \pi} \left[\frac{\Delta_f}{|p_1 - p_2|} \frac{\beta\delta}{\alpha\gamma} \mid \Delta_f \leq \frac{\gamma}{\delta} \frac{|p_1 - p_2|}{4} \right] \\ &\leq \frac{\varepsilon_f}{|p_1 - p_2|} \max_u \frac{\beta\delta}{\alpha\gamma} \end{aligned} \quad (19)$$

Next, we apply [Proposition A.5](#). Denote $\Delta_{\lambda^*} = \lambda^*(p^*) - \lambda^*(\hat{p})$, and define the following probability event:

$$A = \left(\Delta_f \leq \frac{\gamma |p_1 - p_2|}{\delta} \frac{1}{4} \right) \text{ and } \left(\Delta_{\frac{\alpha}{\beta}} \leq \frac{1}{2\nu} \right)$$

Note that the bound in [Proposition A.5](#) is represented as a minimum between two functions, one linear in ε and one constant. To leverage this property, apply the law of total expectation:

$$\mathbb{E}_{u,\pi}[\Delta_{\lambda^*}] = \mathbb{E}_{u,\pi}[\Delta_{\lambda^*} | A] \mathbb{P}[A] + \mathbb{E}_{u,\pi}[\Delta_{\lambda^*} | \bar{A}] \mathbb{P}[\bar{A}] \quad (20)$$

Under A , the first term in Eq. (20) can be bounded by the linear term in [Proposition A.5](#). Taking $\mathbb{P}[A] \leq 1$ and combining with equation Eq. (18):

$$\begin{aligned} \mathbb{E}_{u,\pi}[\Delta_{\lambda^*} | A] \mathbb{P}[A] &\leq \mathbb{E}_{u,\pi}[\Delta_{\lambda^*} | A] \\ &\leq \mathbb{E}_{u,\pi} \left[\frac{\beta^2 \gamma}{2\alpha^2 \delta} \Delta_{\frac{\alpha}{\beta}} | A \right] \\ &\leq \mathbb{E}_{u,\pi} \left[\frac{\beta^2 \gamma}{2\alpha^2 \delta} \frac{\Delta_f}{|p_1 - p_2|} \frac{\beta \delta}{\alpha \gamma} | A \right] \\ &\leq \frac{\nu^3}{2|p_1 - p_2|} \varepsilon_f \end{aligned} \quad (21)$$

The expectation factor in the second term of Eq. (20) can be bounded by the constant term in [Proposition A.5](#):

$$\mathbb{E}_{u,\pi}[\Delta_{\lambda^*} | \bar{A}] \leq \frac{1}{4} \max_u \frac{\beta \gamma}{\alpha \delta} \leq \frac{\nu}{4} \max_u \frac{\gamma}{\delta} \quad (22)$$

Decompose the probability factor $\mathbb{P}[\bar{A}]$ using the law of total probability:

$$\begin{aligned} \mathbb{P}[\bar{A}] &= \mathbb{P} \left[\Delta_f > \frac{\gamma |p_1 - p_2|}{\delta} \frac{1}{4} \right] + \mathbb{P} \left[\Delta_{\frac{\alpha}{\beta}} > \frac{1}{2\nu} \mid \Delta_f \leq \frac{\gamma |p_1 - p_2|}{\delta} \frac{1}{4} \right] \mathbb{P} \left[\Delta_f \leq \frac{\gamma |p_1 - p_2|}{\delta} \frac{1}{4} \right] \\ &\leq \mathbb{P} \left[\Delta_f > \frac{\gamma |p_1 - p_2|}{\delta} \frac{1}{4} \right] + \mathbb{P} \left[\Delta_{\frac{\alpha}{\beta}} > \frac{1}{2\nu} \mid \Delta_f \leq \frac{\gamma |p_1 - p_2|}{\delta} \frac{1}{4} \right] \end{aligned}$$

Apply Markov's inequality $\mathbb{P}[|X| \geq a] \leq \frac{\mathbb{E}[|X|]}{a}$ on the probabilities to obtain:

$$\begin{aligned} \mathbb{P} \left[\Delta_f > \frac{\gamma |p_1 - p_2|}{\delta} \frac{1}{4} \right] &\leq \mathbb{E}_{u,\pi}[\Delta_f] \left(\frac{\gamma |p_1 - p_2|}{\delta} \frac{1}{4} \right)^{-1} \\ &\stackrel{\text{by Eq. (17)}}{\leq} \frac{\varepsilon_f}{|p_1 - p_2|} \max_u \frac{\delta}{\gamma} \end{aligned} \quad (23)$$

$$\begin{aligned} \mathbb{P} \left[\Delta_{\frac{\alpha}{\beta}} > \frac{1}{2\nu} \mid \Delta_f \leq \frac{\gamma |p_1 - p_2|}{\delta} \frac{1}{4} \right] &\leq \mathbb{E}_{u,\pi} \left[\Delta_{\frac{\alpha}{\beta}} \mid \Delta_f \leq \frac{\gamma |p_1 - p_2|}{\delta} \frac{1}{4} \right] \\ &\stackrel{\text{by Eq. (19)}}{\leq} \frac{\varepsilon_f}{|p_1 - p_2|} \max_u \frac{\beta \delta}{\alpha \gamma} \\ &\leq \frac{\varepsilon_f}{|p_1 - p_2|} \nu \max_u \frac{\delta}{\gamma} \end{aligned} \quad (24)$$

Plugging back equations Eq. (21), Eq. (22), Eq. (23), Eq. (24) into equation Eq. (20), we obtain bounds for each term:

$$\mathbb{E}_{u,\pi}[\Delta_{\lambda^*}] = \underbrace{\mathbb{E}_{u,\pi}[\Delta_{\lambda^*} | A] \mathbb{P}[A]}_{\text{by Eq. (21)}} + \underbrace{\mathbb{E}_{u,\pi}[\Delta_{\lambda^*} | \bar{A}]}_{\text{by Eq. (22)}} \underbrace{\mathbb{P}[\bar{A}]}_{\text{by Eq. (23), Eq. (24)}} \quad (25)$$

we obtain:

$$\mathbb{E}_{u,\pi}[\Delta_{\lambda^*}] \leq \frac{\varepsilon_f}{|p_1 - p_2|} \left(\frac{\nu^3}{2} + \left(\nu + \frac{\nu^2}{4} \right) \mu \right) = \varepsilon_{\lambda^*}$$

To obtain the regret bound on the empirical rates, we apply assumptions (ii), (iii) once again to bound the expected difference between $\lambda^*(p)$ and $\frac{1}{T}|S_u(p; T)|$, and apply the triangle inequality:

$$\mathbb{E}_{u,\pi} \left[\left| \frac{1}{T}|S_u(p^*; T)| - \frac{1}{T}|S_u(\hat{p}; T)| \right| \right] \leq \varepsilon_{\lambda^*} + 2(\varepsilon_{\text{dev}} + \varepsilon_{\text{LV}})$$

Note that $\frac{\nu}{|p_1 - p_2|} > 1$, as $\frac{\beta}{\alpha} \geq 1$ since all the users are assumed to remain engaged in the long term, and $|p_1 - p_2| \leq 1$ as $p_1, p_2 \in [0, 1]$. Therefore, the function $\eta_{\text{TPP}} = g(\mu, \nu) = \left(\frac{\nu^3}{2} + \left(\nu + \frac{\nu^2}{4} \right) \eta + 2\nu \right)$ satisfies:

$$\mathbb{E}_{u,\pi} \left[\left| \frac{1}{T}|S_u(p^*; T)| - \frac{1}{T}|S_u(\hat{p}; T)| \right| \right] \leq \frac{\eta_{\text{TPP}}}{|p_1 - p_2|} (\varepsilon_{\text{pred}} + \varepsilon_{\text{dev}} + \varepsilon_{\text{LV}})$$

□

B EXPERIMENTAL DETAILS

B.1 DATA

MovieLens-1M We base our main experimental environment on the MovieLens-1M dataset, which is a standard benchmark dataset used widely in recommendation system research (Harper & Konstan, 2015). The dataset includes 1,000,209 ratings provided by 6,040 users and for 3,706 movies. Ratings are in the range $\{1, \dots, 5\}$, and all users in the dataset have at least 20 reported ratings. The dataset is publicly available at: <https://grouplens.org/datasets/movielens/1m/>.

Data partitioning. To learn latent user and item features, 30% of all ratings were drawn at random. Stratified sampling was applied to ensure that all users and items were covered, and so that each users have roughly the same proportion of ratings used for this step. These ratings were only used only for learning a CF model, and were discarded afterwards. The remaining 70% data points were used for training and testing. For these, we first randomly sampled 1,000 users to form the test set. Then, the remaining users were partitioned into the main train set \mathcal{S} , which included 70% ($\approx 3,528$) of these users, and the experimental treatment sets $\mathcal{S}^{(j)}$, each including 10% (≈ 504) users for $N = 3$. This procedure was repeated 10 times with different random seeds. We report average results, together with 95% t-distribution confidence intervals representing variation between runs.

B.2 IMPLEMENTATION DETAILS

- **Hardware:** All experiments were run on a single laptop, with 16GB of RAM, M1 Pro processor, and with no GPU support.
- **Runtime:** A single run consisting the entire pipeline (data loading and partitioning, collaborative filtering, training classifiers, simulating dynamics, learning policies, measuring and comparing performance) takes roughly 23 minutes. The main bottleneck is the discrete LV simulation, taking roughly 70% of runtime to compute, mostly due to bookkeeping necessary for the non-stationary baselines. Simulation code was optimized using the NUMBA jit compiler, which improves runtime.
- **Optimization packages:**
 - **Collaborative filtering (CF):** We use the SURPRISE package (Hug, 2020), which includes an implementation of the SVD algorithm for CF. All parameters were set to default values.
 - **Regression:** We use the SCIKIT-LEARN implementation of linear regression for predicting long-term engagement from user features (i.e the prediction models $f_j(u)$ in Eq. (8)). All parameters were set to default values.
 - **Non-Negative Least Squares (NNLS):** We use the SCIPY.OPTIMIZE implementation of NNLS. The algorithm was used with its default parameters.
- **Code:** Code for reproducing all of our figures and experiments is available in the following repository: <https://github.com/lvml-iclr-2023/lvml>.

B.3 OTHER BASELINES

- **Safety:** In each step of the TPP simulation, look k step back, and calculate the empirical rate $\tilde{\lambda}_i = \frac{k}{t_i - t_{i-k}}$. If this rate exceeds the threshold $\tilde{\lambda}_i > \tau$, the policy enters a ‘cool-down’ policy state, serving only forced breaks until the next time period. In our experiments, we used thresholds $\tau \in \{14, 16\}$, $k = 10$ look-behind steps, and defined the cool-down period as 0.5 time units.
- **Oracle:** To estimate the effect of perfect predictions, we implement an oracle predictor $f_p^{\text{oracle}}(u)$ which has access to the latent user parameters. For a given u and for each p , the predictor outputs the infinite-horizon LV equilibrium for u , namely $f_p^{\text{oracle}}(u) = \lambda^*(p; \tilde{\theta}_u)$. We define $\tilde{\theta}_u = (\alpha_u, \tilde{\beta}_u, \gamma_u, \delta_u)$, where $\alpha_u, \gamma_u, \delta_u$ are the unobserved parameters for the given user, and $\tilde{\beta}_u$ is the expected value of β_{ux} induced by the distribution over recommended items x induced by the recommendation policy ψ . We view $\tilde{\theta}_u$ as a useful proxy for the otherwise unattainable $\bar{\theta}_u$.

B.4 HYPERPARAMETERS

- **Collaborative filtering:** We used $d = 8$ latent factors and enabled bias terms, which ensured performance is close to the benchmark of RMSE = 0.873 reported in the SURPRISE documentation. We used the vanilla SVD solver, with all hyperparameters set to their default values.
- **Recommendation policy:** Softmax temperature was set to 0.5.
- **Prediction:** We trained regressors $f(u)$ on input feature vectors consisting of three components:

$$u = (\tilde{v}_u, b_u, \hat{\rho}_u) \in \mathbb{R}^{d+2} \quad (26)$$

The three components are: (i) SVD latent user factors $\tilde{v}_u \in \mathbb{R}^d$, (ii) SVD user bias term $b_u \in \mathbb{R}$, (iii) an additional feature consisting of the average predicted ratings for unseen items $\hat{\rho}_u$ weighted by recommendation probability, which we found to slightly improve predictive performance:

$$\hat{\rho}_u = \sum_{x \in \text{holdout}(u)} \hat{r}_{ux} \cdot \text{softmax}_x(\hat{r}_u) \quad (27)$$

Where $\text{holdout}(u)$ is the set of unseen items corresponding to user u , $\hat{r}_{ux} \in [1, 5]$ are the predicted ratings, and $\hat{r}_u \in [1, 5]^{|\text{holdout}(u)|}$ is the vector of all predicted ratings used for softmax recommendation as described in section 5. We chose to focus on linear models since the treatment datasets are relatively small (each $|\mathcal{S}^{(j)}| \approx 500$), and since other model classes (including boosted trees and MLPs) did not perform significantly better.

- **Discrete TPP:** Interaction sequences for each user were generated according to an LV discretization scheme $\text{TPP}^{\text{LV}}(u; T)$, described in detail in the next section. Latent states were initialized randomly with relative uniform noise around the theoretical LV equilibrium point $(\lambda_0, q_0) = ((1 + \xi_\lambda)\lambda^*, (1 + \xi_q)q^*)$, where $\xi_\lambda, \xi_q \sim \text{Uniform}(-0.1, 0.1)$. Latent states were updated each $B = 10$ recommendations to stabilize noise (see Figure 5). When x is recommended to u at time t , latent states and Δt are set according to $\beta_u(t)$, which depends on ratings r_{ux} (true or mixed with predictions $u^\top x$ via κ). Specifically, we use $\beta_u(t) = r_{ux}^2/5 \in \{0.2, 0.8, 1.8, 3.2, 5\}$, which is convex, to accentuate the role of low ratings since they are underrepresented in the data. For $B \geq 1$, we take the effective $\beta_u(t)$ to be the average over the B items recommended in that step. We set $\alpha = 1.3$, and chose $\gamma = 0.2, \delta = 0.01$ (which together determine scale) so that typical values for engagement rate $\frac{1}{T}|S_u|$ are on the order of ≈ 10 for the chosen $T = 100$.

B.5 DISCRETE TPP FOR LOKTA-VOLTERRA SIMULATION

The Temporal Point Process (TPP) we use for simulating user interaction sequences S_u is based on a discretization of the LV system described in Eq. (2), using the forward Euler method with variable step sizes. We denote this process by $\text{TPP}^{\text{LV}}(u; T)$, and present the sampling procedure in Algorithm 1.

Each user is associated with discrete latent states λ_i, q_i , and parameters $\alpha_u, \gamma_u, \delta_u$. Initial states λ_0, q_0 are set randomly. At each step, and in time t_i , the system recommends $x_i = x(t_i)$, which triggers updates in latent states, and determines the next time of interaction t_{i+1} . As noted, these update depend on item-specific parameters β_{u, x_i} .




Figure 5: Example discrete sequence $S_u \sim \text{TPP}^{\text{LV}}(u; T)$, compared to continuous LV dynamics. Our TPP produces discrete sequences that are qualitatively different from their continuous-time analogs (blue lines). Nonetheless, it captures the general properties of our proposed behavioral model: note how cumulative averaging behavior (orange dashes) exhibits ‘habit formation’, which our equilibrium approach targets (blue dots). For the same initial conditions $\lambda(0), q(0)$, the figure shows how varying the number of recommended items per step (B) ‘smooths’ the discrete behavior (left: $B = 1$, center: $B = 10$). For fixed $\beta_u(t) = \beta_u$, when $B \rightarrow \infty$, and when $\Delta t \rightarrow 0$, TPP sequences approach a continuous LV trajectory; in general, and particularly when $\beta_u(t)$ varies by step and per recommended items—this is not the case.

Under stationary policy $\pi(p)$, the system recommends an item with probability $(1-p)$, and suggests a break with probability p . The simulator considers B recommendation opportunities at each step. For each $k \in \{1, \dots, B\}$, denote by $I_k \in \{0, 1\}$ the break indicator, equal to 0 when a break is recommended at the k -th slot in the batch. Denote by $x \sim \psi$ the item recommended by the underlying policy ψ , and by $\beta(x) = r_{ux}^2/5$ the corresponding LV hyperparameter as defined above.

Algorithm 1: Sample from $\text{TPP}^{\text{LV}}(u; T)$

Input: Break probability $p \in [0, 1]$
 Stationary content recommendation probability ψ
 Lotka-Volterra parameters $\theta_u = (\alpha, \beta, \gamma, \delta)$
 Time horizon $T > 0$

Output: Interaction sequence $S_u \sim \text{TPP}_{\pi(p) \circ \psi}^{\text{LV}}(u; T)$

Initialize $i = 0, t_0 = 0, S_u = \{\}$;

while $t_i < T$ **do**

foreach $k \in \{1, \dots, B\}$ **do**

$I_k \sim \text{Bernoulli}(1-p)$;

$x_k \sim \psi$;

$\beta_k \leftarrow \beta(r_{ux_k})$;

end

$\Delta t_i \leftarrow \lambda_i^{-1}$;

$\lambda_{i+1} \leftarrow \lambda_i + \left(-\alpha + \frac{\sum_{k=1}^B I_k \beta_k}{B} q_i\right) \lambda_i \Delta t_i$;

$q_{i+1} \leftarrow q_i + \left(\gamma(1-q_i) - \frac{\sum_{k=1}^B I_k \delta}{B} \lambda_i\right) q_i \Delta t_i$;

$t_{i+1} \leftarrow t_i + \lambda_{i+1}^{-1}$;

$S_u \leftarrow S_u \cup \{(t_i, (x_1, \dots, x_B), (I_1, \dots, I_B))\}$;

$i \leftarrow i + 1$;

end

C ADDITIONAL EMPIRICAL EVALUATION

In this section, we provide additional empirical results as further indicators of potential applicability:

- In [subsection C.1](#), we repeat the experiments presented in [section 5](#) using a different dataset, showing that qualitative results remain consistent across recommendation modalities.
- In [subsection C.2](#), we demonstrate how to leverage additional problem structure commonly available in practice, and develop adaptive breaking policies which improve over time.
- Finally, in [subsection C.3](#) we present a stateless consumption model which is unrelated to Lotka-Volterra and does not promote breaks. Evaluation of the LV policy optimization method under this model shows that the optimization procedure does not degrade performance when breaks are not beneficial, thus extending [Corollary 2](#) beyond the realizable case.

C.1 ADDITIONAL DATASET

The Goodreads dataset. The Goodreads book recommendations data set is a common benchmark dataset used in recommendation systems research ([Wan & McAuley, 2018](#); [Wan et al., 2019](#)). Ratings are in the range $\{1, \dots, 5\}$, and the dataset is filtered to only include users with at least 20 reported ratings. We use the official comic-books genre subset of the dataset, which includes 3,679,076 ratings provided by 41,932 users and for 87,565 books after pre-processing. The dataset is publicly available at: <https://sites.google.com/eng.ucsd.edu/ucsdbookgraph/>. Pre-processing code is available in our code repository: <https://github.com/lvml-iclr-2023/lvml>.

Evaluation. For evaluation, we follow the same procedure described in [Appendix B](#), maintaining identical data processing procedures, constants, and hyperparameters. This procedure was repeated 10 times with different random seeds. We report average results, together with 95% t-distribution confidence intervals.

Results. [Figure 6](#) shows results for the Goodreads experiment, in the same format as [Figure 3](#). Results exhibit performance and trends that are qualitatively similar to the MovieLens experiment in [section 5](#). The LV policy optimization method achieved better performance on this dataset (right pane): The performance of the LV is closer to oracle (-0.526% in Goodreads compared to -0.918% in ML1M), and the gap from the best-of method is slightly larger (+5.77% in Goodreads compared to +2.19% in ML1M). Varying user types (center pane) shows less variation across values of κ , indicating that predictors achieve satisfactory performance even when the κ is low and prediction is harder (see [subsection 5.1](#)). Varying treatments (right pane) also coincides with the observations: LV policy optimization on the Goodreads dataset exhibits less performance degradation as $p_1 \rightarrow 0$, suggesting that less data may be sufficient for optimization. We attribute these results to the larger size of the dataset (approximately 3.6M interactions, compared to 1M interactions in ML1M), and to the possibility that a stronger structure may exist in this recommendation modality compared to general movie recommendation.




Figure 6: Results on the Goodreads comic-books dataset (compare to [Figure 3](#)). **(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 ($N = 2, p_0 = 0$).

C.2 LEVERAGING ADDITIONAL STRUCTURE USING ADAPTIVE POLICIES

One benefit of our approach in [section 4](#) is that it requires only predictions of long-term engagement. In reality, however, other sources of information may also be available to the learner, which may be useful for improving decisions regarding when to recommend breaks. Here we consider additional information in the form of explicit user feedback, collected by the system over the course of interaction. We model users as providing to the system feedback regarding item quality, namely true ratings r_{ux} , for some of the recommended items x they consume. For simplicity, we assume users report ratings with probability $p_r \in [0, 1]$, independently for each item, where we vary p_r across experimental conditions. Intuitively, we would expect that, over time, user-reported ratings help improve our learned breaking policy; however, over-relying on few data and at an early stage may hinder performance. Here we examine when and how such feedback should be utilized.

Adaptive policy optimization. To extend our approach to incorporate ratings-based feedback, we propose an adaptive explore-then-exploit approach which can leverage sparse explicit feedback. Fixing an *adaptation time* $T_0 \in [0, T]$ (a hyperparameter), the method applies the standard learned breaking policy $\hat{\pi}$ (Equation 9) in the time frame $t \leq T_0$, during which it also collects and stores user-reported ratings. Then, at time $t = T_0$, the method updates the learning policy on the basis of the new ratings data, and applies this policy until the eventual time T . In particular, ratings are used to improve the estimated component $\hat{\rho}_u$ in the user feature vector u' (described in Equation 27). In this way, the method uses the existing engagement predictors $f_p(u)$ without requiring any retraining.

A formal description of this method is provided in [Algorithm 2](#). Here we denote the set of ratings collected until time t by $r_{[0,t]} \in \{1, \dots, 5\}^*$; the average rating at time t by $\bar{\rho}_{(0,t)} = \frac{1}{|r_{[0,t]}|} \sum_{r \in r_{[0,t]}} r$; and the set of long-term engagement predictions for a feature vector u by $\mathbf{f}(u) = \{(p_i, f_{p_i}(u))\}$ (as in [subsection 4.1](#)).

Algorithm 2: Adaptive policy optimization using sparse rating signals

Input: Initial break probability $p \in [0, 1]$

Time horizon $T > 0$

Adaptation time $T_0 \in [0, T]$

User feature vector $u = (\tilde{v}, b, \hat{\rho})$

1. Collect ratings data $r_{[0,T_0]}$ with breaking policy $\pi(p)$ until time T_0 ;
 2. Construct updated feature vector $u' = (\tilde{v}, b, \bar{\rho}_{[0,T_0]})$ using the average rating $\bar{\rho}_{[0,T_0]}$;
 3. Compute updated long-term engagement predictions $\mathbf{f}(u')$;
 4. Use the LV policy optimization method to obtain updated policy $p' = p^*(\mathbf{f}(u'))$;
 5. Use breaking policy $\pi(p')$ for the remaining time $t \in [T_0, T]$;
-

Evaluation. We evaluate the adaptive policies using the MovieLens-1M experimental setup, as described in [section 5](#) and [appendix B](#). We make use of the same linear regression predictors $f_p(u)$ from the main part of the experiment, and maintain the same time horizon $T = 100$. We vary the ratings density parameter p_r in the range $[0, 1]$, and vary the adaptation time T_0 across three distinct values $T_0 \in \{0.5, 5, 50\}$.

Results. Results are presented in [Figure 7](#). In the results, we observe that LV-adaptive with $T_0 = 5$ represents an optimal point within the hyperparameter space. For this choice, results show significant gains over the non-adaptive LV method even under moderate rating densities, and quick convergence towards the LV-true-ratings upper bound. In contrast, LV-adaptive with $T_0 = 0.5$ represents a “premature optimization” scenario, attempting to adapt the policy without allowing the estimated average $\bar{\rho}_{[0,T_0]}$ enough time to converge to its expected value. Similarly, $T_0 = 50$ represents the result of “late optimization”, which has accurate estimation of $\bar{\rho}_{[0,T_0]}$ but not enough time to benefit from the updated policy. Finally

C.3 DISTINCT BEHAVIORAL MODEL

In this work, we propose the LV model as a behavioral hypothesis class for counterfactual prediction of long-term engagement. As such, using the LV model within the learning-to-break optimization




Figure 7: Adaptive policy evaluation for varying rating density $p_r \in [0, 1]$. Solid blue line represents the adaptive policy described in Algorithm 2, with adaptation time $T_0 = 5$. Dashed lines represent the performance of adaptive policies under different choices of T_0 . Cyan horizontal line represents LV-true-ratings, a non-adaptive method with oracle access to the true average rating. The remaining horizontal lines represent the performance of selected non-adaptive policies, as presented in Figure 3 (left pane).

framework is a design choice, to be made at the discretion of the learner; our LV model would be a good choice if it fits the data better than alternative model classes, given the amount of available data. This relation is made precise in our error bound (Theorem 1), which bounds the error when using the LV model for *any* underlying TPP (i.e., we make no assumptions about the true underlying data generation process).

Our main experiments evaluate our approach on data that is not LV dynamics, but nonetheless, bear some resemblance. To complement these results, here we run an additional experiment in which we empirically evaluate our LV-based approach on data that is generated by a behavioral model that is entirely distinct. In particular, we consider a user model in which consumption decisions only depend on the quality of recommended items, without any dependence on internal states—i.e., it is *stateless*. This conforms to behavioral models which are implicitly assumed in conventional recommendation methods such as collaborative filtering.

Stateless content consumption. The data generation process is formalized in Algorithm 3. As a means of capturing stateless behavior, we define a “close-range” temporal point process $\text{TPP}^{\text{CR}}(u; T)$ which generates sequences of user interactions based solely on the average rating of recommended items in a batch. Here we relate rating with utility, and assume that items having higher utility (and therefore higher ratings) induce more frequent interactions; in the same way, we consider break prompts as items having zero utility, and hence zero rating. At time t_i , and given a batch of size B with $k \leq B$ recommended items and $B - k \geq 0$ breaks prompts, users acting according to the stateless behavioral will consume the next batch of content at time:

$$t_{i+1} = t_i + \frac{1}{\tau} \left(\frac{1}{B} \sum_{j=1}^k r_{ux_j} \right)^{-1} \quad (28)$$

where $(r_{ux_1}, \dots, r_{ux_k}) \in \{1, \dots, 5\}^k$ are the ratings of the items recommended at time k , and $\tau > 0$ is a constant latent parameter to be learned from data. The breaking policy π decides on the number of items k to recommend on each step. Since $r_{ux} \geq 1$ for all user-item pairs, the time difference $\Delta t_i = t_{i+1} - t_i$ in Equation 28 is minimized by taking $k = B$. This shows that the optimal breaking policy under this behavioral model is the default one, which does not prompt the user to break.

Evaluation. We evaluate the stateless behavioral model using the MovieLens-1M experimental setup, as described in section 5. We set $\tau = 4$ for all users, utilize linear regression for engagement prediction, and maintain all hyperparameters without change. Data processing steps are performed as described in appendix B, and the chosen breaking policies are evaluated.

Algorithm 3: Sample from $\text{TPP}^{\text{CR}}(u; T)$

Input: Break probability $p \in [0, 1]$
 Stationary content recommendation probability ψ
 Scalar parameter $\theta_u = \tau > 0$
 Time horizon $T > 0$

Output: Interaction sequence $S_u \sim \text{TPP}_{\pi(p) \circ \psi}^{\text{CR}}(u; T)$

Initialize $i = 0, t_0 = 0, S_u = \{\}$;

while $t_i < T$ **do**

foreach $k \in \{1, \dots, B\}$ **do**

$I_k \sim \text{Bernoulli}(1 - p)$;

$x_k \sim \psi$;

end

$t_{i+1} \leftarrow t_i + \frac{1}{\tau} \left(\frac{1}{B} \sum_{k=1}^B I_k r_{ux_k} \right)^{-1}$;

$S_u \leftarrow S_u \cup \{(t_i, (x_1, \dots, x_B), (I_1, \dots, I_B))\}$;

$i \leftarrow i + 1$;

end




Figure 8: Fitting an LV equilibrium curve on the empirical consumption rates of $\text{TPP}^{\text{CR}}(u; T)$. The plots are a realization of the schematic diagram in Figure 2 using data from the MovieLens-1M experiment. **(Left)** Prediction with unbiased engagement predictions. **(Right)** User with extremely low engagement, for which the system tends to over-predict. In both cases, the LV policy optimization method selects the optimal policy $p = 0$.

Results. Despite the difference between the true TPP and our choice of model class, the LV policy optimization method successfully learned the optimal no-breaks—which in this case, is the policy $p = 0$ for all users. Further detail is provided by Figure 8, which illustrates the policy optimization steps under $\text{TPP}^{\text{CR}}(u; T)$ for typical users with unbiased and biased engagement predictions. In both examples, the optimal points of both curves coincide, and the optimal policy is selected despite poor point-wise fit. Combined, these results show that our approach is “safe”, in the sense that when breaking is sub-optimal, the learned breaking policy does not override the default policy.