Information Provision with Information Overload

N. Bora Keskin[1] Chenghuai Li[1] Meng Li[2]

[1]Duke University, Fuqua School of Business, e-mail: {bora.keskin, chenghuai.li}@duke.edu

[2]University of Houston, C.T. Bauer College of Business, e-mail: mli@bauer.uh.edu

Abstract

Motivated by practical applications in e-commerce, we consider a product promotion problem where a product showcase webpage is personalized for each arriving consumer during a campaign. The product has various characteristics that may affect the consumers' willingness to purchase, and each consumer may weigh these characteristics differently. However, showing all the product characteristics to consumers can sometimes lead to *information overload*, which means consumers can be less interested in the product when facing too much information on the product. Based on the evidence from a field experiment, we build a feature-based logit model to capture the information overload effect in the consumer choice behavior such as clicking. Faced with this probabilistic choice model with unknown parameters, we design a statistical learning algorithm that carefully perturbs the showing characteristics of the product and then exploits the estimated model to optimize the product display page for each arriving consumer. By dedicating $\sqrt{T}$ periods to exploration, the algorithm leveraging the maximum likelihood estimation achieves the expected cumulative regret of order $\sqrt{T}\log T$. Finally, we conduct a B2B field experiment in collaboration with a medical agent who sells tissue simulation equipment for medical imaging and radiation therapy to institutions in China. In this field experiment, we randomly display different versions of a product webpage to arriving customers, where these versions differ in the amount of displayed product information. Observing the information overload phenomenon from the experimental data, we empirically test the performance of our proposed algorithm.

Keywords: information overload, binary logit, dynamic learning, regret, field experiment.

1 Introduction

Information disclosure refers to the practice of sharing or releasing data, often personal or sensitive, by individuals or organizations. This can range from voluntarily sharing personal details on social media to companies disclosing product information to consumers. However, information overload occurs when the volume and velocity of information available surpasses an individual's or organization's capacity to process it effectively (Eppler and Mengis 2008, Bawden and Robinson 2020), leading to confusion, anxiety, and stress, which impairs decision-making capabilities. This phenomenon is exacerbated in the digital era, characterized by the abundant availability and rapid dissemination of information via social media and other online platforms. This overload presents significant challenges across various domains: on social media platforms, users are bombarded with a continuous stream of updates, photos, videos, and tweets; the proliferation of internet news sources inundates users with information on a myriad of topics, from politics to entertainment; in online learning, students face an excess of educational material; in health and wellness, the vastness of online information makes discerning reliable sources difficult; and in e-commerce, consumers are overwhelmed by extensive product details, reviews, and pricing options.

The information overload exists widely in various fields. In psychology, Milgram (1970) observes that signal overload in urban environments contributes to antisocial behavior. In the medical sector, Laker et al. (2018) identify that the overuse of electronic health records and other clinical data impedes the timeliness of doctors' clinical decision-making. Similarly, empirical evidence in business, such as A/B testing results from B2C (business-to-consumer) and B2B (business-to-business) markets, substantiates the presence of information overload.

1.1 Motivating Example

We procure the data from Air China, who conducted an A/B test about the display mode of the passenger's fare page. In this A/B test, they randomly display two versions of the fare page to the website visitors. One version is the original version Air China has used before, called version A; the other version is the *lite* version, called version B (see the screenshots, i.e., Figures 1a and 1b). In short, the lite version only displays the most basic information explicitly but hides the other more advanced detailed information in web hyperlinks or pop-ups, which only appear when the passenger clicks on the corresponding buttons, compared to the original version.

(a) Original version (version A). **(b) Lite version (version B).**

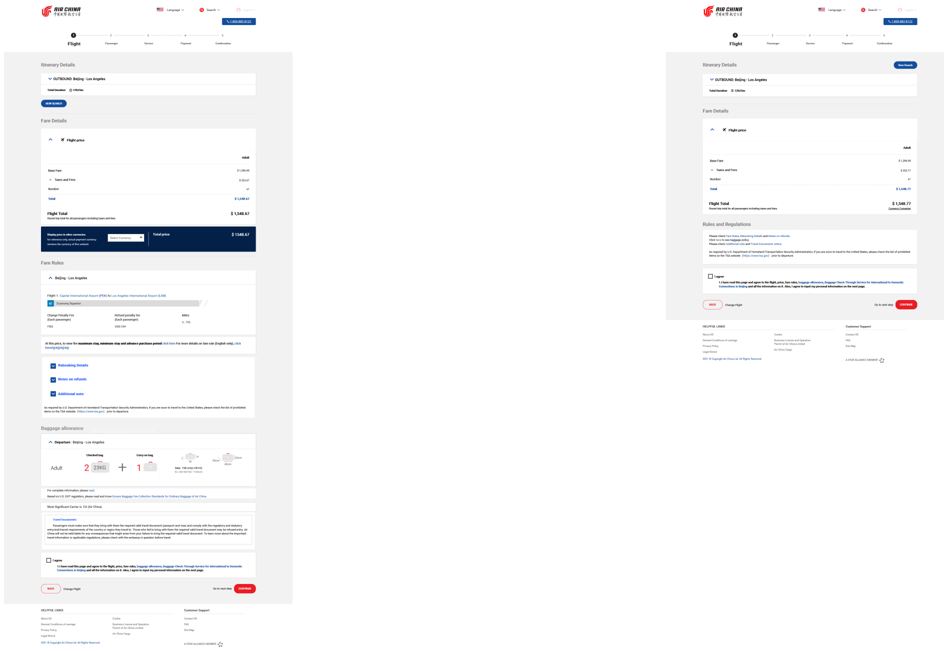

Figure 1: Air China fare page.

The commonalities are that both versions display the "Fare Details" module that entails the flight price (including the base fare price, taxes and fees, the total price, etc.) and the signature box for passengers to agree with the terms. However, in addition to the flight price, version A also displays other modules explicitly on the fare page, such as "Fare Rules", "Baggage Allowance", and "Travel Documents". Specifically, "Fare Rules" includes information about the seat classes for each flight, the change penalty fee, the refund penalty fee, rebooking details, notes on refunds, and additional notes. "Baggage Allowance" includes information about the maximum weight and size constraints on the checked bag and carry-on bag. "Travel Documents" contains a statement that reminds passengers to bring with them the required valid travel document (passport and visa) and comply with the regulatory and statutory entry/exit/transit requirements of the country or region they travel to. In comparison, version B does *not* show these modules explicitly on the fare page but only provides hyperlinks to the pop-ups in case the passengers are interested. Because of this difference, the lite version (version B) has a shorter fare page displaying less information to passengers.

Air China tests whether the lite version of the fare page can increase the conversion rate (the number of order checkouts divided by the number of visits to the webpage, multiplied by 100 to get a percentage), or decrease the bounce rate (the percentage of passengers who land on the webpage but just leave without triggering another request such as ordering a plane ticket). In the experiment period (from October 25, 2021, to November 24, 2021), 81078 visitors were assigned to the group of the website with original fare pages, and they placed 1392 orders; at the same time, 87244 visitors were assigned to the group of the website with lite fare pages, and 1918 orders are placed. Table 1 demonstrates that more information would lead to less order and a smaller conversion rate.

Table 1: Statistical Evidence for Impact of Information on Consumer Behavior

	Unique visitors to website	Orders	Rate
Original version (A)	81078	1392	1.72%
Lite version (B)	87244	1918	2.20%
Comparison	N/A	N/A	$\uparrow 27.91\%$
p-value	N/A	N/A	1.44×10^{-12}

1.2 Our Setting

We study how to dynamically adjust the product elements displayed on the recommendation page to maximize customers' willingness to click (purchase). To this end, we construct a recommendation system framework with one seller and T customers. The product display webpage is personalized for each arriving customer. The product has different features, and consumers have different weights for these features, which affects the customers' click behavior. Showing all the features of a product to consumers can sometimes cause decision fatigue and overwhelm customers, i.e., the phenomenon of information overload. To avoid losses caused by information overload, the seller can filter and display effective product features while optimizing the current push mode based on historical data. In detail, customer clicks can be characterized as binary click behavior. We establish a new feature-based logit model that uses click probability to measure the degree of match between the recommendation list and consumer interests.

Our feature-based logit model incorporates the information overload phenomenon by explicitly characterizing the cognitive burden of displaying additional product features on customers' willingness to click. We consider an online setting, where the seller can learn from past observations of customer responses to different product display pages. The customers' binary click behavior is viewed as bandit feedback; the seller only obtains feedback for her chosen display of product features and does not observe (counterfactual) responses to alternative product displays for each customer. Therefore, the seller needs to carefully experiment with different product display pages to obtain accurate estimates of customers' matching weights for product features and the magnitude of the information overload effect (exploration) so that she can optimize the product display page and maximize customers' willingness to click (exploitation).

In this paper, we propose a new statistical learning algorithm in the bandit setting leveraging maximum likelihood estimation. In the presence of the information overload effect, our setting is different from the contextual generalized linear bandit problem. This feature requires us to introduce new conditions for selecting experimenting product displays in the exploration stage, so that we can effectively learn the customers' weights for product features and estimate the information overload parameter relating to the total number of chosen features. We measure the algorithm's performance using the standard notion of expected cumulative regret, which is the total expected loss in reward compared to a clairvoyant' oracle that knows all the model parameters. Our main result establishes that our algorithm achieves the expected cumulative regret that scales in the rate of $\sqrt{T}\log T$ over the total number of customers T.

Finally, to demonstrate the information overload phenomenon in practical applications, we conducted a B2B field experiment in collaboration with a medical agent who sells tissue simulation equipment for medical imaging and radiation therapy to institutions in Mainland China. We select a popular product and randomly display different versions of the product page to arriving customers. In addition to the basic information about the product, these versions differ in whether displaying additional graphical or text information to customers. By collecting the clicking data of these customers, we find that while the sole graphical or text information boosts the customers' clicking rate, adding both graphical and text information can cause the customers' clicking rate to drop. Therefore, we calibrate our model with the experimental data and empirically test our proposed algorithm's performance.

1.3 Related Literature

The concept of information overload has been a focal point in economics and marketing research. Early studies by Jacoby et al. (1974) and Malhotra (1982) have highlighted the potential dysfunctions arising from providing consumers with excessive information. This body of literature suggests that an overload of information can negatively impact consumer behavior and decision-making processes. Central to this discussion is the classical information-processing theory on consumer choice (Bettman 1979), which posits that consumers have limited capacities in terms of working memory and computational capabilities. Jacoby (1984) assert that information overload can occur in various settings, affecting consumer behavior. Iyengar and Lepper (2000) and Kuksov and Villas-Boas (2010) demonstrate through experiments and analytical models that offering fewer product choices could actually enhance consumer purchasing likelihood, as excessive options can increase search and evaluation costs, deterring consumers from making choices. Overall, there is a consensus in this literature that the burden of information load can adversely affect individual performance.

There is an increasing number of models that incorporate information overload for consumers. Schmit and Johari (2018) introduce the consumer abandonment behavior caused by too much information into a learning model. Cao and Sun (2019) observe that overexposure to marketing activities could lead to customer dissatisfaction and integrate this into a dynamic model where marketing fatigue results in platform abandonment. Goutam et al. (2019) develop a Markov Chain Model to capture the choice overload of consumers. Research also focuses on recommendation mechanisms, where sellers devise optimal information strategies to promote product purchases. Gao et al. (2021) explore assortment optimization for impatient customers, while Feldman and Segev (2022) employ polynomial-time approximation schemes to enhance algorithmic accuracy in this context. However, our concern lies more with how the extent of product information disclosure affects consumer purchasing willingness. Bastani et al. (2022) delve into learning personalized product recommendations considering customer disengagement, offering a nuanced perspective on balancing information provision with customer engagement. Küçükgül et al. (2022) use a dynamic Bayesian persuasion framework to study optimal dynamic information provision for maximizing revenue.

2 Problem Formulation

Consider a firm, hereafter referred to as the *seller*, that promotes a product online and displays a *varying* recommendation page to T customers who arrive sequentially. Regarding each recommendation opportunity as a "period," the seller has a discrete time horizon of T periods and dynamically adjusts the product elements shown on the recommendation page over the time horizon. For expositional convenience, we also call the customer arriving in period t as customer t.

At the beginning of period t, the seller observes a n-dimensional vector of features pertaining to the arriving customer t. Let $X_t = (X_{t,1}, X_{t,2}, \ldots, X_{t,n})^T$ denote the customer's feature vector, where each component denotes one customer feature, such as the customer's self-identified gender, age range, location, and data pertaining to the customer's past purchases. We assume that $\{X_t, t \in [T]\}$ are independent and identically distributed with a compact support $\mathscr{X} \subset \mathbf{B}_0(x_{max}) \subset \mathbb{R}^n$, where $\mathbf{B}_0(x_{max})$ is a n-dimensional ball of radius $x_{max} > 0$. Without loss of generality, we normalize $\mathbb{E}[X_t]$ to zero. We also denote the covariance matrix of X_t by $\Sigma_X := \mathbb{E}[X_t X_t^T]$ and assume that Σ_X is symmetric and positive-definite. For convenience, we introduce an augmented customer feature vector $\bar{X}_t := [1\ X_t^T]^T$. Since customers' preferences over the product are correlated to their own features, we use $\bar{Z}_t := \bar{\Theta}\bar{X}_t$ to denote a customer's preference vector over various product features, where $\bar{\Theta}$ is a $m \times (n+1)$ fixed matrix with unknown parameters. Note that $\bar{Z}_{t,i}$ denotes the customer's preference score for the i^{th} feature of the product that may or may not be shown by the seller on the recommendation page.

For each customer, upon observing $X_t = x_t$, the seller can choose to display some of the m product features, such as a photo of the product or a piece of brand description, on the recommendation page and hide the other under an "expand more" button. Precisely, the seller's decision vector in period t is an m-dimensional binary vector $a_t \in \{0,1\}^m$, where $a_{t,i} = 1$ if and only if the seller displays the i^{th} product feature to customer t. Given the seller's decision a_t, customer t's preference for the product is

$$\bar{U}_t := \bar{\alpha} \cdot \bar{x}_t + a_t \cdot \bar{z}_t,$$

where $\bar{\alpha} \in \mathbb{R}^{n+1}$ is an unknown vector related to the intercept of customers' preference function, $\bar{x}_t = [1\ x_t^T]^T$, $\bar{z}_t = \bar{\Theta}\bar{x}_t$, and $u \cdot v$ denotes the inner product of two vectors with the same dimension. Note that $\bar{U}_t$ captures feature-dependent customer's general interest for the product (through $\bar{\alpha} \cdot \bar{x}_t$) and feature-dependent preference scores for each display option of product features (through $\bar{z}_{t,i}$, $i \in [m]$).

After viewing the product recommendation page, the customer can click on the button that brings her to the next step (e.g., shopping cart) or just leave. Hereafter, we will refer to the customer's behavior as a click. The probability of customer t's click is affected by her preference for the product $\bar{U}_t$, which depends on the customer's features x_t and the seller's display of product features a_t. Moreover, motivated by the field experiment evidence in Section 1, information overload is an important factor that can influence the customer's click behavior. When a customer is recommended a product or service, too much information (e.g., overly detailed descriptions) can make the recommendation less attractive. Therefore, to incorporate the information overload effect, we use a *new* binary logit model to characterize the customer's click behavior. Specifically, given a customer's feature vector $x \in \mathscr{X}$, the click probability in response to the seller's decision $a = (a_1, a_2, \ldots, a_m)^T \in \{0,1\}^m$ is

$$\bar{P}(a; x, \bar{\Theta}, \bar{\alpha}, \beta) := \frac{e^{\bar{U}}}{e^{\beta g(\sum_{i=1}^{m} \mathbb{1}\{a_i=1\})} + e^{\bar{U}}}, \tag{1}$$

Where: $\bar{U} = \bar{\alpha} \cdot \bar{x} + a \cdot \bar{z}$ (note that $\bar{z} = \bar{\Theta}\bar{x}$ and $\bar{x} = [1\ x^T]^T$), $\beta > 0$ is an unknown parameter characterizing the magnitude of information overload effect, and $g(\cdot)$ is an increasing function with $g(0) = 0$. Common choices of $g(\cdot)$ include simple convex functions, such as $g(x) = x^2$ and $g(x) = e^x$. The nonlinearity of $g(\cdot)$ offers us more flexibility in modeling the information overload effect. When more product features are displayed by the seller, the term $\exp(\beta g(\sum_{i=1}^{m} \mathbb{1}\{a_i = 1\}))$ in the denominator will be larger and hence decrease the click probability. We define $c := \sum_{i=1}^{m} \mathbb{1}\{a_i = 1\}$ as the total number of selected product features ("c" represents "count"), and the term $\exp(\beta g(c))$ reflects the information overload effect in the probabilistic click model (1).

For convenience, we use $\bar{\Theta}_i$ to denote the i^{th} row of $\bar{\Theta}$. Then, let $\bar{\theta} := (\bar{\Theta}_1, \ldots, \bar{\Theta}_m, \bar{\alpha}^T, \beta)^T$ and $\bar{\theta}_{box}$ be a compact rectangle in $\mathbb{R}^{(n+1)(m+1)+1}$ from which the value of $\bar{\theta}$ is chosen, where $\bar{\theta}_{box} \supset \{\tilde{\bar{\theta}} \in \mathbb{R}^{[(n+1)(m+1)+1]} \mid \|\tilde{\bar{\theta}} - \bar{\theta}\|_2 \leq 1\}$. Now, we impose a margin condition on the customer feature distribution adapted from the bandit literature. This assumption ensures that the density of the customer feature distribution should be bounded near a decision boundary, i.e., the intersection of a hypersurface $\{x \in \mathbb{R}^n : \bar{P}(a_1; x, \bar{\theta}) - \bar{P}(a_2; x, \bar{\theta}) = 0\}$ and $\mathscr{X}$ for any different product displays $a_1 \neq a_2 \in \{0,1\}^m$. (We allow the customer feature distribution to have point masses on decision boundaries.)

Assumption 1. (MARGIN CONDITION) *There exists a constant $C_* \in \mathbb{R}^+$ such that for all product displays a_1, a_2 in*

$\{0,1\}^m$ *where* $a_1 \neq a_2$, $\mathbb{P}_X\left[0 < \left|\bar{P}(a_1; X, \bar{\theta}) - \bar{P}(a_2; X, \bar{\theta})\right| \leq \kappa\right] \leq C_* \kappa$ *for all* $\kappa \in \mathbb{R}^+$.

This assumption requires that the probability of the customer feature vector X_t drawn near a decision boundary is proportional to how close the rewards resulting from two product displays are. When the realization of the customer feature vector X_t is close to (but not on) a decision boundary, i.e., $X_t \in \{x \in \mathscr{X} : 0 < \left|\bar{P}(a_1; x, \bar{\theta}) - \bar{P}(a_2; x, \bar{\theta})\right| \leq \kappa\}$, even small errors in parameter estimation can cause the seller to choose the suboptimal product display (between the two choices a_1 and a_2) because the rewards for both displays are nearly equal. Therefore, if a disproportionate fraction of observed customer feature vectors are drawn near these decision boundaries (hypersurfaces), the seller can perform poorly. This assumption was introduced into the classification literature by Tsybakov (2004) and highlighted in bandit settings by Goldenshluger and Zeevi (2013), Bastani and Bayati (2020). In particular, if we restrict our attention to product displays with only a single feature, i.e., $a_1, a_2 \in \{\mathbf{e}_i\}_{i=1}^m$ and the linear information overload effect, i.e., $g(x) = x$, our problem is reduced to a contextual generalized linear bandit problem and Assumption 1 degenerates to the margin condition assumption in Bastani and Bayati (2020). In practical applications, most customer feature distributions can satisfy Assumption 1, including continuous customer features with bounded density in the support and categorical customer features with discrete distribution.

In period t, the seller observes customer t's binary response to the product recommendation page, either click or not. Let $F_t \in \{0,1\}$ denote customer t's bandit feedback. Then, F_t follows a Bernoulli distribution with mean $\bar{P}(a_t; x_t, \bar{\theta})$. The seller's objective is to maximize the total click-through rate while facing uncertainty in the market. For illustration, suppose that the seller earns a unit revenue from every click. If the seller knew $\bar{\theta}$ in advance, referred to as the clairvoyant, then she can maximize the expected single-period revenue by solving $\max_{a \in \{0,1\}^m} \bar{P}(a; x, \bar{\theta})$ upon observing the customer's feature $x \in \mathscr{X}$. An optimal solution $a^*(x, \bar{\theta})$ must exist since the feasible set is a finite set, and we define $\bar{P}^*(x, \bar{\theta}) := \bar{P}(a^*(x, \bar{\theta}); x, \bar{\theta})$. To solve $a^*(x, \bar{\theta})$, the clairvoyant must weigh the tradeoff between increasing the product's overall attraction by including attractive product features (with positive $\bar{z}_i$) and the negative effect of overloading (β). Although for the general function $g(\cdot)$, there does *not* exist a closed-form solution for $a^*(x, \bar{\theta})$, we have the following proposition.

Proposition 1. (Clairvoyant's Optimal Display) *Suppose that the seller observes the customer's feature vector* $x \in \mathscr{X}$. *Let* $\bar{z}_{i_1} \geq \bar{z}_{i_2} \geq \ldots \geq \bar{z}_{i_n}$ *be the sorted sequence of* $\bar{z}_i := (\bar{\Theta}\bar{x})_i$, $i = 1, \ldots, m$. *Then,* $a^*(x, \bar{\theta}) :=$ *argmax*$_{a \in \{0,1\}^m} \bar{P}(a; x, \bar{\theta})$ *must be of the following form*

$$a^*_{i_l}(x, \bar{\theta}) = \begin{cases} 1, & \text{if } l \leq k \\ 0, & \text{if } l > k \end{cases}, \quad \text{for some } k \in \{0, \ldots, m\}.$$

Now, we specify the seller's admissible display policies. First, we construct a sequence of historical data vectors as $H_0 = (X_1)$ and $H_t = (a_1, \ldots, a_t, F_1, \ldots, F_t, X_1, \ldots, X_{t+1})$ for $t \in [T]$. Then, we define an *admissible policy* as a sequence of functions $\pi = (\pi_1, \pi_2, \ldots)$, where $\pi_t : \mathbb{R}^{(t-1)(m+n+1)+n} \to \{0,1\}^m$ is a measurable function that maps H_{t-1} to the seller's display for customer t, a_t^π. Thus, under policy π, $a_t^\pi = \pi_t(H_{t-1})$, for $t \in [T]$. We denote by Π the set of all admissible policies.

Given policy $\pi \in \Pi$ and $\bar{\theta} \in \bar{\theta}_{box}$, define a probability measure $\mathbb{P}^\pi_{\bar{\theta}}(\cdot)$ on the sample space of customers' binary feedback sequences $F = (F_1, F_2, \ldots)$ such that

$$\mathbb{P}^\pi_{\bar{\theta}}(F_1 = f_1, \ldots, F_T = f_T) = \prod_{t=1}^T \left(\frac{e^{\bar{U}_t}}{e^{\beta g(c_t)} + e^{\bar{U}_t}}\right)^{f_t} \left(1 - \frac{e^{\bar{U}_t}}{e^{\beta g(c_t)} + e^{\bar{U}_t}}\right)^{1-f_t},$$

for $f_1, f_2, \ldots, f_T \in \{0,1\}$. The seller's conditional expected revenue loss in T periods relative to a clairvoyant who knows the underlying parameter vector $\bar{\theta}$ is defined as

$$\Delta^\pi_{\bar{\theta}}(T; \mathbf{X}_T) = \mathbb{E}^\pi_{\bar{\theta}}\left(\sum_{t=1}^T [\bar{P}^*(X_t, \bar{\theta}) - \bar{P}(a_t^\pi; X_t, \bar{\theta})] \middle| \mathbf{X}_T\right), \tag{2}$$

for $\bar{\theta} \in \bar{\theta}_{box}$, $\pi \in \Pi$, and $\mathbf{X}_T = (X_1, \ldots, X_T) \in \mathscr{X}^T$, where: $\mathbb{E}^\pi_{\bar{\theta}}(\cdot)$ is the expectation operator associated with $\mathbb{P}^\pi_{\bar{\theta}}(\cdot)$. We call this performance metric (2) as the seller's T-period *conditional regret*, which is a random variable depending on the realization of $\mathbf{X}_T = (X_1, \ldots, X_T)$. The seller's objective is to minimize its T-period *expected regret*

$$\Delta^\pi_{\bar{\theta}}(T) = \mathbb{E}_X[\Delta^\pi_{\bar{\theta}}(T; \mathbf{X}_T)],$$

for $\bar{\theta} \in \bar{\theta}_{box}$ and $\pi \in \Pi$, where: $\mathbb{E}_X[\cdot]$ is the expectation operator associated with the probability measure governing $\{X_t, t = 1, 2, \ldots\}$. Throughout the sequel, we will use the expectation notation $\mathbb{E}^\pi_{X,\bar{\theta}}(\cdot) := \mathbb{E}_X\left(\mathbb{E}^\pi_{\bar{\theta}}(\cdot)\right)$, and we let $\mathbb{P}^\pi_{X,\bar{\theta}}(\cdot)$ be the probability measure associated with $\mathbb{E}^\pi_{X,\bar{\theta}}(\cdot)$.

3 Designing Explore-Then-Commit Policies

To attract more customer clicks while addressing the potential problem of information overload, we design a type of explore-then-commit policy. We first conduct product feature display experiments and estimate the unknown parameters $\bar{\theta}$ via likelihood maximization, and then we optimize product information provisions to the later visitors based on the estimated model. In addition to learning about the unknown weights $\bar{\Theta}$ and $\bar{\alpha}$, the algorithm provides a certain degree of variation in the number of product features in the exploration stage so that it can better estimate the impact of overloading information on customers' click behavior.

We first describe the maximum likelihood estimation procedure. Given the observed customer features $(x_1, \ldots, x_t)$, the seller's past displays $(a_1, \ldots, a_t)$, and the customers' click responses $(F_1, \ldots, F_t) = (f_1, \ldots, f_t)$, the log-likelihood mapping, $\tilde{\bar{\theta}} = (\tilde{\bar{\Theta}}_1, \ldots, \tilde{\bar{\Theta}}_m, \tilde{\bar{\alpha}}^T, \tilde{\beta})^T \mapsto \bar{\ell}_t(\tilde{\bar{\theta}})$, is

$$\bar{\ell}_t(\tilde{\bar{\theta}}) := \sum_{s=1}^{t} (1 - f_s)\tilde{\beta} g(c_s) + f_s \tilde{\bar{U}}_s - \log\left(e^{\tilde{\beta} g(c_s)} + e^{\tilde{\bar{U}}_s}\right) \quad \text{for} \quad \tilde{\bar{\theta}} \in \mathbb{R}^{(n+1)(m+1)+1},$$

where $\tilde{\bar{U}}_s := \tilde{\bar{\alpha}} \cdot \bar{x}_s + \sum_{i=1}^m a_{s,i} \left(\tilde{\bar{\Theta}} \bar{x}_s\right)_i$ for $s \in [t]$. We show that $\bar{\ell}_t(\tilde{\bar{\theta}})$ is strictly concave, and we denote the maximum likelihood estimator (MLE) as $\hat{\bar{\theta}}_{t+1} = \operatorname{argmax}_{\tilde{\bar{\theta}} \in \mathbb{R}^{(n+1)(m+1)+1}} \left\{\bar{\ell}_t(\tilde{\bar{\theta}})\right\}$. Moreover, let $\hat{\bar{\vartheta}}_{t+1} = \mathscr{P}_{\bar{\theta}_{box}} \left\{\hat{\bar{\theta}}_{t+1}\right\}$ be the truncated estimator, where $\mathscr{P}_{\bar{\theta}_{box}} : \mathbb{R}^{(n+1)(m+1)+1} \to \bar{\theta}_{box}$ is the projection mapping from $\mathbb{R}^{(n+1)(m+1)+1}$ onto $\bar{\theta}_{box}$. After deriving the truncated MLE $\hat{\bar{\vartheta}}_{t+1}$, policy π can exploit the estimated model to optimize product feature displays for the remaining periods.

Due to the nature of explore-then-commit algorithms, we need to carefully design product displays in the exploration stage so that we can obtain a sufficiently accurate estimate for the model. To do so, we connect the estimation accuracy to the empirical Fisher information matrix. In particular, we show that with high probability, the minimum eigenvalue of the Fisher information, $\mu_{\min}\left(\bar{\mathscr{I}}_t\right)$, is lower bounded by

$$\min\left\{\mu_{\min}(\Sigma_X) \cdot \mu_{\min}\left(\sum_{s=1}^{t} \begin{pmatrix} a_s a_s^T & a_s \\ a_s^T & 1 \end{pmatrix}\right), \mu_{\min}\left(\sum_{s=1}^{t} \begin{pmatrix} g(c_s)^2 & -g(c_s) a_s^T & -g(c_s) \\ -g(c_s) a_s & a_s a_s^T & a_s \\ -g(c_s) & a_s^T & 1 \end{pmatrix}\right)\right\}.$$

Based on the above relationship, we hope to display the product features in the first t periods to ensure that both $\mu_{\min}\left(\sum_{s=1}^t (a_s^T,\ 1)^T (a_s^T,\ 1)\right)$ and $\mu_{\min}\left(\sum_{s=1}^t (-g(c_s),\ a_s^T,\ 1)^T (-g(c_s),\ a_s^T,\ 1)\right)$ increase in the order of t. To achieve it, we propose the following conditions: (i) we select $m+1$ experimental display vectors, say $\hat{a}_1, \ldots, \hat{a}_{m+1} \in \{0,1\}^m$, such that $(\hat{a}_1^T, 1)^T, \ldots, (\hat{a}_{m+1}^T, 1)^T$ form a basis of $\mathbb{R}^{m+1}$; (ii) we select $m+2$ experimental display vectors (which can overlap with the previous $m+1$ display vectors), say $\hat{a}'_1, \ldots, \hat{a}'_{m+2} \in \{0,1\}^m$, such that $(-g(\hat{c}'_1), (\hat{a}'_1)^T, 1)^T, \ldots, (-g(\hat{c}'_{m+2}), (\hat{a}'_{m+2})^T, 1)^T$ form a basis of $\mathbb{R}^{m+2}$, where $\hat{c}'_i = \sum_{j=1}^m \hat{a}'_{i,j}$; (iii) we display each of these experimental vectors sufficiently many times in the exploration stage. By following the above guidelines, we can make $\mu_{\min}\left(\bar{\mathscr{I}}_t\right) \gtrsim t$ with high probability. Consequently, if the explore-then-commit policy π satisfies the above conditions in the exploration stage, the estimation error of the truncated MLE will decay at the rate $\sqrt{\frac{\log t}{t}}$ with respect to the number of experimentation periods.

Proposition 2. (ESTIMATION ERROR) *Suppose that a policy $\pi \in \Pi$ devotes the first t periods to experimentation. If the product displays are chosen deterministically in the experimental stage and:*

1. *There exist $\underline{l}_1, C_1 > 0$ and $m+1$ different product display vectors $\hat{a}_1, \hat{a}_2, \ldots, \hat{a}_{m+1} \in \{0,1\}^m$, where concatenated vectors $(\hat{a}_1^T, 1)^T, \ldots, (\hat{a}_{m+1}^T, 1)^T$ form a basis of $\mathbb{R}^{m+1}$, such that for any $t \geq C_1$, the display vector $\hat{a}_i$ has been selected at least $\underline{l}_1\, t$ times for all $i \in [m+1]$;*

2. *There exist $\underline{l}_2, C_2 > 0$ and $m+2$ different product display vectors $\hat{a}'_1, \hat{a}'_2, \ldots, \hat{a}'_{m+2} \in \{0,1\}^m$, where concatenated vectors $(-g(\hat{c}'_1), (\hat{a}'_1)^T, 1)^T, \ldots, (-g(\hat{c}'_{m+2}), (\hat{a}'_{m+2})^T, 1)^T$ form a basis of $\mathbb{R}^{m+2}$, such that for any $t \geq C_2$, the display vector $\hat{a}'_i$ has been selected at least $\underline{l}_2\, t$ times for all $i \in [m+2]$,*

then there exist finite and positive constants κ_2, ρ_2, and t_1 such that

$$\mathbb{P}^{\pi}_{X,\bar{\theta}}\left(\|\hat{\bar{\theta}}_{t+1} - \bar{\theta}\|^2 \leq \frac{\rho_2[(n+1)(m+1)+1]\log t}{t}\right) \geq 1 - \kappa_2[(n+1)(m+1)+1]\frac{\log t}{t},$$

for all $\bar{\theta} \in \bar{\theta}_{box}$ and $t \geq t_1$, where $\|\cdot\|$ denotes the Euclidean norm.

Now, we design a specific policy that satisfies the conditions specified in Proposition 2. The policy LSAoN, i.e., Algorithm 1, uses the first $\lceil T^{\lambda} \rceil$ periods to experiment with various product features: the policy either displays one product feature at a time (i.e., $a_t = \mathbf{e}_i$ for $i \in [m]$), displays all the product features simultaneously (i.e., $a_t = \mathbf{1}$), or displays none of the additional product features (i.e., $a_t = \mathbf{0}$). We can verify that the concatenated vectors

$$(\mathbf{e}_1^T, 1)^T, \ldots, (\mathbf{e}_m^T, 1)^T, (\mathbf{1}^T, 1)^T$$

form a basis of $\mathbb{R}^{m+1}$ as long as $m > 1$. Furthermore, we notice that the set of vectors $\{\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_m, \mathbf{1}, \mathbf{0}\} \subset \{0,1\}^m$ make sure that the concatenated vectors

$$(-g(1), \mathbf{e}_1^T, 1)^T, \ldots, (-g(1), \mathbf{e}_m^T, 1)^T, (-g(m), \mathbf{1}^T, 1)^T, (-g(0), \mathbf{0}^T, 1)^T$$

form a basis of $\mathbb{R}^{m+2}$ as long as $mg(1) \neq g(m)$. Note that if $g(\cdot)$ is strictly convex, then we always have $g(m) > mg(1)$ for $m > 1$. Due to the cyclic nature of the exploration stage, the number of appearing times for each display vector in $\{\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_m, \mathbf{1}, \mathbf{0}\}$ is proportional to the length of the exploration stage. Therefore, the LSAoN policy generates a sufficiently accurate estimate after the exploration stage based on Lemma 2. Then, we can prove the following expected regret guarantee for the LSAoN policy.

Theorem 1. (REGRET) *Let π be the LSAoN policy with the first $\lceil \sqrt{T} \rceil$ periods devoted to experimentation. Then, there exist finite and positive constants $\bar{C}$ and $\bar{t}_2$ such that*

$$\Delta_{\bar{\theta}}^{\pi}(T) \leq \bar{C} 2^m [(n+1)(m+1)+1] \sqrt{T} \log T$$

for all $\bar{\theta} \in \bar{\theta}_{box}$ and $T \geq \bar{t}_2$.

Algorithm 1 (LSAoN policy: Learning with Single, All, or Null Product Features)

INPUT: (i) $\mathbf{m}$: the total number of product features. (ii) $\mathbf{T}$: time horizon. (iii) $\lambda \in (0,1)$: a parameter adjusting the length of exploration stage; the first $\lceil T^\lambda \rceil$ periods are for exploration.
(Learning)
for $t = 1, \ldots, \lceil T^\lambda \rceil$ **do**
  Compute the residual of t module $(m+2)$, i.e., set $r \equiv t \bmod (m+2)$.
  if $r = 0$ **then**
    Set $a_t = \mathbf{0}$.
  else if $r \in \{1, 2, \ldots, m\}$ **then**
    Set $a_t = \mathbf{e}_r$
  else
    Set $a_t = \mathbf{1}$.
  end if
end for
(Estimation) Compute the log likelihood $\bar{\ell}_{\lceil T^\lambda \rceil}\left(\tilde{\bar{\theta}}\right)$ for $\tilde{\bar{\theta}} \in \mathbb{R}^{[(n+1)(m+1)+1]}$ and the maximum likelihood estimate $\hat{\bar{\theta}}_{\lceil T^\lambda \rceil + 1} := \operatorname{argmax}_{\tilde{\bar{\theta}}} \bar{\ell}_{\lceil T^\lambda \rceil}\left(\tilde{\bar{\theta}}\right)$.
Set the projected estimate $\hat{\bar{\vartheta}}_{\lceil T^\lambda \rceil + 1} = \mathscr{P}_{\bar{\theta}_{box}}\left\{\hat{\bar{\theta}}_{\lceil T^\lambda \rceil + 1}\right\}$.
(Earning)
for $t = \lceil T^\lambda \rceil + 1, \ldots, T$ **do**
  Upon observing the customer's feature vector $X_t = x_t$, select the optimal display of product features based on the estimated model, i.e., set $a_t = \operatorname{argmax}_{a \in \{0,1\}^m} \bar{P}\left(a; x_t, \hat{\bar{\vartheta}}_{\lceil T^\lambda \rceil + 1}\right)$.
end for

---

4 Numerical Study

4.1 B2B Field Experiment and Model Calibration

In this section, we introduce a B2B field experiment we conducted in collaboration with a medical equipment agent. This agent has the authority to sell medical equipment in Mainland China on behalf of a worldwide leader in tissue simulation for medical imaging and radiation therapy, and its targeted clients are mainly hospitals and other medical institutions. The scope of products includes various phantoms for every imaging modality. For our purpose, we select a popular product, a multi-modality abdominal biopsy phantom.

On the agent's website, the product page displays relevant information about the product, and the customers can click on the product graph and the buttons to consult online and contact the agent. The agent's question is how much product information it should display on the product page to attract customers. In addition to the basic information, such as the product's name, brief illustrations of its use, category, product's manufacturer (brand), product origin, etc., the agent can also display additional graphical information (various graphs of medical imaging) and/or additional text information (detailed itemized characteristics about the product). In our field experiment, we call the additional text information our first treatment and call the additional graphical information our second treatment. Then, based on whether we apply each treatment, we have four versions of the product page ($4 = 2 \times 2$). For simplicity, we can denote these versions as $\{a_1, a_2\}$, where $a_j = 1$ means that we apply the j^{th} treatment and $a_j = 0$ otherwise.

When each customer arrives at the product page, the system will randomly display one of the four versions to the customer with equal probability. We monitor whether the customer clicks on any place in the prespecified area of the product page, and the customer's clicking is a binary indicator of her/his interest in this product. Therefore, we can observe the customers' clicking behavior in response to different versions of the product page. We ran our experiment from September 7th to October 28th, 2023, and recorded the number of customer visits and clicks as follows.

Versions of product page	Number of visits	Number of clicks	Clicking rate
$\{0, 0\}$	56	12	21.43%
$\{0, 1\}$	51	12	23.53%
$\{1, 0\}$	58	20	34.48%
$\{1, 1\}$	54	11	20.37%

Table 2: Data of field experiment. In version $\{0, 0\}$, there is no additional information. In version $\{0, 1\}$, there is *only* additional graphical information. In version $\{1, 0\}$, there is *only* additional text information. In version $\{1, 1\}$, there are *both* additional text and graphical information.

Based on the data in Table 2, we notice that additional text or graphical information increases the clicking rate compared to the plain version without any additional information. In particular, the text information seems to have a stronger promotion effect than the graphical information. However, if we add both text and graphical information to the product page, the clicking rate decreases relative to the plain version. This observation demonstrates a kind of information overload phenomenon in our field experiment: when either treatment in isolation increases the clicking rate of customers, they together cause a negative effect in attracting customers.

We now fit our model using the above data in Table 2. Since our field experiment has no customer contextual information, we can utilize a special case of clicking model (1). Given the agent's display of the product features for customer t, $a_t = (a_{t,1}, a_{t,2}) \in \{0, 1\}^2$, customer t's preference for this product is denoted as $\bar{U}_t = \alpha + a_{t,1}\theta_1 + a_{t,2}\theta_2$, where $\bar{\alpha} \in \mathbb{R}$ is a constant, and $\theta_1, \theta_2 \in \mathbb{R}$ denote the weights of these two additional product features (text and graphical information) to customers. Then, the clicking probability of customer t is expressed as

$$\bar{P}\left(a_t; \alpha, \theta_1, \theta_2, \beta\right) = \frac{\exp\left(\bar{U}_t\right)}{\exp\left(\beta g\left(\sum_{i=1}^{2} a_{t,i}\right)\right) + \exp\left(\bar{U}_t\right)},$$

where $\beta > 0$ characterizes the degree of information overload. For calibration, we have a certain flexibility in choosing function $g(\cdot)$, and we use $g(t) = t^2$ for simplicity. By maximum likelihood estimation as discussed in Section 3, we obtain the following estimates for model parameters: $\hat{\alpha} = -1.2993,\ \hat{\theta}_1 = 1.0785,\ \hat{\theta}_2 = 0.5417,$ and $\hat{\beta} = 0.4210$.

4.2 Performance of the LSAoN Policy

Next, we illustrate the performance of the LSAoN policy (see Algorithm 1) in our calibrated model. We know that in the calibrated model, the best display option is only to show the additional text information, i.e., $a^* = \{1, 0\}$, and the optimal expected click probability (equivalently, optimal single-period revenue given unit revenue for each click) is $\bar{P}^*(\bar{\theta}) := \bar{P}(a^*; \alpha, \theta_1, \theta_2, \beta) = 0.3448$. In our simulation study, we let $\lambda = 0.5$, and the algorithm uses the first $\lceil\sqrt{T}\rceil$ periods to experiment with different displays for arriving customers.

Figure 2 depicts the average performance of the LSAoN policy for the calibrated instance. Note that for each T, simulation results are conducted over 1000 replications to compute the average regret. We consider selling horizons ranging from $T = 2,500$ to $T = 50,000$. The regret in Figure 2 seems to be of order $\sqrt{T}$, as predicted by Theorem 1.

5 Concluding Remark

Motivated by a field experiment conducted in the airline industry, we study the prevalent information overload problem in the context of personalized digital advertising and product recommendation. In particular, when the seller displays information for the promoted product, showcasing all the product characteristics can sometimes impair customers' willingness to purchase, thereby decreasing the click-through rate. We first establish a new feature-based logit model to

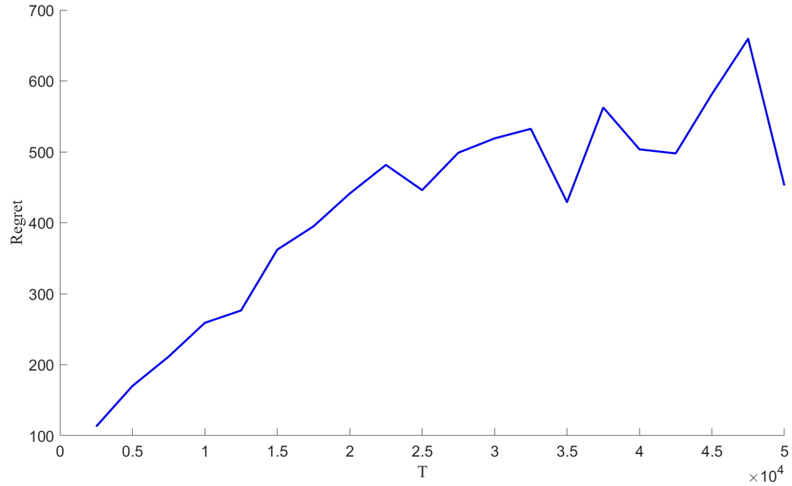

Figure 2: Performance of the LSAoN policy. The graph illustrates the dependence of the regret on T for the calibrated instance.

explicitly capture the information overload behavioral aspect of customers' clicks. Then, we consider an online setting where the seller needs to learn from data about customers' interests in various product features and the magnitude of information overload while maximizing the cumulative click-through rate. To address the exploration-exploitation tradeoff, we propose explore-then-commit type algorithms by carefully designing product display experiments in the exploration stage. Using maximum likelihood estimation, our algorithms can accurately estimate the unknown problem parameters and achieve the expected cumulative regret of order $\sqrt{T}\log T$.

To test the value of study in practical settings, we conducted a B2B field experiment in collaboration with a medical agent who sells tissue simulation equipment for medical imaging and radiation therapy to institutions in Mainland China. By randomly displaying different webpage versions of a popular product to arriving customers, we observe the information overload phenomenon from our collected click data. In particular, although a sole graphical or text information boosts the customers' click rate, displaying them simultaneously can backfire. We calibrate our model with the experimental data and find our proposed algorithm performs well empirically.

References

Bastani, H. and Bayati, M. (2020). Online decision making with high-dimensional covariates. *Operations Research*, 68(1):276–294.

Bastani, H., Harsha, P., Perakis, G., and Singhvi, D. (2022). Learning personalized product recommendations with customer disengagement. *Manufacturing & Service Operations Management*, 24(4):2010–2028.

Bawden, D. and Robinson, L. (2020). *Information overload: An overview*. Oxford University Press.

Bettman, J. R. (1979). *Information processing theory of consumer choice*. Addison-Wesley Pub. Co.

Cao, J. and Sun, W. (2019). Dynamic learning of sequential choice bandit problem under marketing fatigue. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3264–3271.

Eppler, M. J. and Mengis, J. (2008). The concept of information overload-a review of literature from organization science, accounting, marketing, mis, and related disciplines (2004) the information society: An international journal, 20 (5), 2004, pp. 1–20. *Kommunikationsmanagement im Wandel: Beiträge aus 10 Jahren= mcminstitute*, pages 271–305.

Feldman, J. and Segev, D. (2022). The multinomial logit model with sequential offerings: Algorithmic frameworks for product recommendation displays. *Operations Research*, 70(4):2162–2184.

Gao, P., Ma, Y., Chen, N., Gallego, G., Li, A., Rusmevichientong, P., and Topaloglu, H. (2021). Assortment optimization and pricing under the multinomial logit model with impatient customers: Sequential recommendation and selection. *Operations research*, 69(5):1509–1532.

Goldenshluger, A. and Zeevi, A. (2013). A linear response bandit problem. *Stochastic Systems*, 3(1):230–261.

Goutam, K., Goyal, V., and Soret, A. (2019). A generalized markov chain model to capture dynamic preferences and choice overload. *arXiv preprint arXiv:1911.06716*.

Iyengar, S. S. and Lepper, M. R. (2000). When choice is demotivating: Can one desire too much of a good thing? *Journal of personality and social psychology*, 79(6):995.

Jacoby, J. (1984). Perspectives on information overload. *Journal of consumer research*, 10(4):432–435.

Jacoby, J., Speller, D. E., and Kohn, C. A. (1974). Brand choice behavior as a function of information load. *Journal of marketing research*, 11(1):63–69.

Küçükgül, C., Özer, Ö., and Wang, S. (2022). Engineering social learning: Information design of time-locked sales campaigns for online platforms. *Management Science*, 68(7):4899–4918.

Kuksov, D. and Villas-Boas, J. M. (2010). When more alternatives lead to less choice. *Marketing Science*, 29(3):507–524.

Laker, L. F., Froehle, C. M., Windeler, J. B., and Lindsell, C. J. (2018). Quality and efficiency of the clinical decision-making process: Information overload and emphasis framing. *Production and Operations Management*, 27(12):2213–2225.

Malhotra, N. K. (1982). Information load and consumer decision making. *Journal of consumer research*, 8(4):419–430.

Milgram, S. (1970). The experience of living in cities: A psychological analysis. In *Annual Meeting of the American Psychological Association., Sep, 1969, Washington, DC, US; This paper is based on an Invited Address presented to the Division of General Psychology at the aforementioned meeting.* American Psychological Association.

Schmit, S. and Johari, R. (2018). Learning with abandonment. In *International Conference on Machine Learning*, pages 4509–4517. PMLR.

Tsybakov, A. B. (2004). Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166.